

SWOT Science Data Products

User Handbook

JPL D-109532

Revision A

March 4, 2025



Jet Propulsion Laboratory
California Institute of Technology



See Sect. [1.3](#) for how to cite this document.

©2025 California Institute of Technology. Government Sponsorship Acknowledged.

Document Custodian:


Electronic Signature on File

Curtis Chen
JPL Algorithm System Engineer

Approved by:

Electronic Signature on File

Shailen Desai
JPL Measurement System Engineer

 E-SIGNED by Nicolas Picot
on 2025-04-03 06:13:36 PDT

Nicolas Picot
CNES Measurement System Engineer

Paper copies of this document may not be current and should not be relied on for official purposes. The current version is in the JPL Engineering Product Data Management system (EPDM: <https://epdm.jpl.nasa.gov>) and the CNES Products Data Management System.

EPDM Electronic Signatures

Status ^	Performer ◆	A... ◆	Due Date ◆	End Date ◆	Comments ◆
Approve	Chen, Curtis W (curtis)	Docu...	25-Mar-2025 10:00	17-Mar-2025	Approve
Approve	Desai, Shailen D (sdesai)	Docu...	25-Mar-2025 10:00	17-Mar-2025	Approved

Change Log

Version	Date	Sections Changed	Reason for Change
Initial Release	2024-05-02	All	Initial Release Approved for public release (URS325006/CL#24-2261)
Revision A	2025-03-04	All	Minor clarifications Approved for public release (URS332012/CL#25-1111)

Contents

Change Log	3
1 Introduction	9
1.1 Document Scope and Intent	9
1.2 Document Organization	10
1.3 Citing This Document	11
1.4 Document Conventions	11
2 Overview of Science Sensors and Data Products	12
2.1 Science Sensors	12
2.2 Science Data Products	14
3 Definitions and Conventions	18
3.1 Terms	18
3.1.1 Ground Track and Nadir Track	18
3.1.2 Orbit	18
3.1.3 Revolution	18
3.1.4 Pass (Ascending and Descending)	18
3.1.5 Orbit Phase (Calibration Orbit and Nominal or Science Orbit)	19
3.1.6 Cycle	19
3.1.7 Beta Angle	19
3.1.8 Yaw Flip	20
3.1.9 Forward and Backward	20
3.1.10 Along Track	20
3.1.11 Cross Track	21
3.1.12 Heading	21
3.1.13 Swath and Half Swath	21
3.1.14 Left and Right	21
3.1.15 Horizontal (H) and Vertical (V) Polarization	23
3.1.16 Point Target Response (PTR)	23
3.1.17 Resolution	23
3.1.18 Sampling	25
3.1.19 Pixel and Line	25
3.1.20 Altitude	25

3.1.21	Range	25
3.1.22	Corrected Range	25
3.1.23	Height	25
3.1.24	Sea Surface Height Anomaly	26
3.1.25	Water Surface Elevation	26
3.1.26	Good, Suspect, Degraded, and Bad Quality	27
3.2	Data Products, Granules, and Files	27
3.3	Coordinate Systems and Representations	28
3.3.1	GCRF	28
3.3.2	ITRF	28
3.3.3	Reference Ellipsoid	29
3.3.4	LLH	29
3.3.5	TCN and HCL	30
3.3.6	KMSF	30
3.3.7	UTC and TAI Time Conventions	30
3.3.8	Correction Convention	32
3.4	Reference Orbit	33
4	Mission Description	34
4.1	Mission Phases and Cycle Numbering	34
4.2	Reference Orbits and Pass Numbering	35
4.3	Orbit Characteristics	36
4.4	Satellite Description	36
4.4.1	Ka-Band Radar Interferometer (KaRIn)	40
4.4.2	Poseidon-3C Nadir Altimeter (NAlt)	40
4.4.3	Advanced Microwave Radiometer	40
4.4.4	DORIS	43
4.4.5	Global Positioning System Payload (GPSP)	43
4.4.6	Laser Retroreflector Array (LRA)	44
4.5	Spacecraft Events	45
4.5.1	Propulsive Maneuvers	45
4.5.2	Gyro Calibration Maneuvers	46
4.5.3	Yaw Flips	46
4.5.4	Solar Array Rotations	46
4.5.5	Eclipse Transitions	47
5	Product Distribution and Fidelity	49
5.1	Data Processing Centers	49
5.2	Data Distribution	49
5.2.1	NASA PO.DAAC	49
5.2.2	CNES AVISO and Hydroweb.next	50
5.3	Fidelity and Latency of Products	51
5.3.1	SAT_COM Product	51
5.3.2	L1_DORIS_RINEX and L1_GPSP_RINEX Products	51

5.3.3	ATTD_RECONST Product	52
5.3.4	MOE and POE Products	52
5.3.5	L2_RAD Products	52
5.3.6	L2_NALT Products	52
5.3.7	KaRIn LR and HR Products	53
5.4	File Naming Conventions	55
5.4.1	KaRIn, ATTD_RECONST, L1_GPSP_RINEX and L2_RAD Products	55
5.4.2	NAlt Products	57
5.4.3	POE, MOE, and SAT_COM Products	58
5.4.4	L1_DORIS_RINEX Product	58
5.5	Known Issues Affecting Data Availability	58
6	Standard Science Data Products	60
6.1	L1_DORIS_RINEX	60
6.2	L1_GPSP_RINEX	60
6.3	SAT_COM	61
6.4	ATTD_RECONST	61
6.5	POE and MOE	63
6.6	L2_RAD_GDR, L2_RAD_IGDR and L2_RAD_OGDR	63
6.7	L2_NALT_GDR, L2_NALT_IGDR and L2_NALT_OGDR	64
6.8	L1B_LR_INTF	67
6.9	L2_LR_SSH	70
6.10	L1B_HR_SLC	71
6.11	L2_HR_PIXC	75
6.12	L2_HR_RiverSP	79
6.13	L2_HR_LakeSP	82
6.14	L2_HR_PIXCVec	83
6.15	L2_HR_Raster	87
6.16	L2_HR_RiverAvg	87
6.17	L2_HR_LakeAvg	89
6.18	L2_HR_FPDEM	89
7	Operational Data Processing	91
7.1	Algorithm Flow	91
7.2	Auxiliary Inputs	94
7.3	Intermediate Data Products	95
8	KaRIn Measurement Principles	98
8.1	Radar Fundamentals	98
8.1.1	Basic Radar Ranging	98
8.1.2	Echo Phase	99
8.1.3	Range Compression and Resolution	101
8.1.4	Echo Power	106
8.2	SAR Imaging	107
8.2.1	Side-Looking Radar Imaging	107

8.2.2	Azimuth Compression	112
8.3	Interferometric SAR Topographic Mapping	123
8.3.1	Cross-Track Interferometry	124
8.3.2	Interferometric Coherence	128
8.3.3	Coregistration	128
8.3.4	Flattening and Reference Locations	129
8.3.5	Phase Unwrapping and Ambiguity Resolution	131
8.3.6	Height Reconstruction and Geolocation	133
8.4	Interferometric Decorrelation and Looks	137
8.4.1	SNR Decorrelation	137
8.4.2	Geometric Decorrelation	138
8.4.3	Volumetric Decorrelation	139
8.4.4	Phase Noise	140
9	KaRIn Design Overview	141
9.1	KaRIn Hardware Design	141
9.1.1	KaRIn Antennas	141
9.1.2	KaRIn Radar Electronics	146
9.1.3	Other KaRIn Hardware	147
9.2	On-Board Processing	147
9.2.1	LR OBP	147
9.2.2	HR OBP	148
9.2.3	OBP Doppler-Centroid Estimation, Tables, and Processing Intervals	149
9.2.4	FPGA Reconfigurations	151
9.2.5	Radiation Hits	151
9.3	HR Mask	153
9.4	KaRIn Ground Processing Design	153
9.4.1	KaRIn LR Processing	153
9.4.2	Crossover Calibration	154
9.4.3	KaRIn HR Processing	156
10	Phenomenology	159
10.1	Speckle	159
10.2	Layover	161
10.3	Surfboard Effect	163
10.4	Sea State Bias	165
10.5	Rain Attenuation	165
10.6	Dark Water	167
10.7	Specular Ringing	170
10.8	Dark Water Misclassification Due to Reference DEM Errors	171
10.9	Bright Land	174
10.10	Phase Unwrapping Artifacts	175
10.11	Target Motion	177
10.12	LR Errors Over Land	179

11 Geophysical Effects	182
11.1 Media Delays	182
11.1.1 Dry Troposphere	183
11.1.2 Wet Troposphere	183
11.1.3 Ionosphere	184
11.2 Sea State Bias Correction	185
11.3 Geophysical Models	185
11.3.1 Geoid	185
11.3.2 Mean Sea Surface	187
11.3.3 Mean Dynamic Topography	189
11.3.4 Tides	189
11.3.5 Ocean Response to Atmospheric Forcing	193
Bibliography	195
Acknowledgements	202
A Acronyms	203

Chapter 1

Introduction

The Surface Water and Ocean Topography (SWOT) mission is a partnership between the physical oceanography and hydrology communities[1]. The primary objective of SWOT is to provide a global survey of water on the surface of the Earth. This includes observing spatial and temporal changes in inland freshwater bodies as well the fine structure of the ocean's surface topography. SWOT was jointly developed by the National Aeronautics and Space Administration (NASA) and the French Space Agency, Centre National d'Études Spatiales (CNES), with contributions from the Canadian Space Agency and United Kingdom Space Agency.

SWOT aims to contribute to the understanding of the global water cycle by providing measurements of water surface topography over both the oceans and inland water. As described in [2], the primary oceanographic objectives are to characterize the mesoscale and sub-mesoscale ocean circulation determined from the ocean surface topography at spatial resolutions of 15 km. The primary hydrology measurement objectives are to:

- Provide a global inventory of all terrestrial surfaces whose surface area exceeds $250 \times 250 \text{ m}^2$ and rivers whose width exceeds 100 m.
- Measure the global storage change in terrestrial surface water bodies at sub-monthly, seasonal, and annual time scales.
- Estimate the global change in river discharge at sub-monthly, seasonal, and annual time scales.

The SWOT mission objectives call for an unprecedented combination of high vertical accuracy, fine horizontal resolution, and global coverage. In order to meet these demands, SWOT builds upon traditional nadir ocean altimeter missions by including a first-of-its-kind measurement architecture based on spaceborne synthetic aperture radar (SAR) interferometry from near-nadir viewing geometries. This measurement is implemented by the Ka-band Radar Interferometer (KaRIn) instrument on SWOT.

1.1 Document Scope and Intent

This document is written for users of SWOT science data products. It is intended to provide a general, high-level introduction to SWOT data and to serve as a reference for information that

facilitates the use of the data. Material in this document is therefore often neither comprehensive nor detailed. Rather, this document is meant to give a relatively condensed, self-contained overview to help readers get started with SWOT data. Readers should seek further depth from other sources, which are often referenced from this document, as needed.

In particular, a number of SWOT documents describe the details of the various data products that are available to users (see Ch. 6). These product description documents (PDDs) define key aspects of each specific data product, such as the product contents, representation, and file format, that will likely be of keen interest to users. Similarly, a number of algorithm theoretical basis documents (ATBDs) describe aspects of the processing approaches used to generate the data products from raw SWOT data (see Ch. 7). This document can be considered a parent of the PDDs and ATBDs. The emphasis of this handbook is on information that is common across multiple PDDs and/or ATBDs.

1.2 Document Organization

The remainder of this document is organized as follows. Chapters 2–7 give an overview of SWOT data products and the processing through which the products are generated. These chapters are intended to orient users with respect to what SWOT data products are available and what distinguishes them. Chapters 8–11 then give an introduction to the principles and characteristics of the novel wide-swath KaRIn measurements. These chapters are intended help users understand how the measurements are made, thereby giving insight into how the data can be interpreted. Specifically,

- Chapter 2 gives a very brief overview of the SWOT instruments and available data products.
- Chapter 3 defines terms, conventions, and coordinate systems.
- Chapter 4 provides an introduction to the design of the SWOT mission.
- Chapter 5 discusses data distribution and data availability.
- Chapter 6 describes the set of SWOT standard data products (SDPs).
- Chapter 7 introduces the operational data processing from which data products are generated.
- Chapter 8 provides a brief tutorial of the measurement principles underlying the KaRIn measurement.
- Chapter 9 gives an overview of the end-to-end design of the KaRIn measurement, including both the instrument hardware and the ground algorithms.
- Chapter 10 describes some aspects of phenomenology with which users may need to be familiar in order to interpret SWOT data properly.
- Chapter 11 discusses geophysical effects that may need to be considered when processing or interpreting the SWOT data.

1.3 Citing This Document

Please cite this document as follows:

JPL D-109532, Revision A, “SWOT Science Data Products User Handbook,” Jet Propulsion Laboratory Internal Document, Pasadena, CA, 2025.

1.4 Document Conventions

Hyperlinks referring to sections, figures, tables, equations, and external URLs in this document are indicated by [blue](#) text. Links to references in the bibliography of this document are indicated by [green](#) text.

Chapter 2

Overview of Science Sensors and Data Products

This chapter gives a brief, high-level overview of the SWOT measurement system and science data products in order to provide context for the material in the following chapters.

2.1 Science Sensors

The SWOT measurement system consists of the following instruments:

- Ka-band Radar Interferometer (KaRIn)
- Dual-frequency (Ku- and C-band) pulse-limited Nadir Altimeter (NAlt)
- Three-frequency Advanced Microwave Radiometer (AMR)
- Doppler Orbitography and Radiopositioning Integrated by Satellite (DORIS) receiver
- Global Positioning System Payload (GPSP) receiver
- Laser Retroreflector Array (LRA)

Very simply, the SWOT measurement concept, as shown in Figure 2.1, is to determine the geocentric height of the water surface from the difference between the geocentric position of the satellite and the range between the satellite and water surface. Measurements of the range from the satellite to the water surface are provided by two independent radar instruments, KaRIn and NAlt, with complementary spatial coverage along the satellite ground track. The KaRIn and NAlt instruments both take advantage of the higher average reflectivity of water as compared to land to generate measurements of water height.

The novel KaRIn synthetic aperture radar interferometer serves as the primary motivation for the SWOT mission. It effectively provides range measurements across two 50 km swaths from 10 to 60 km on each side of the nadir ground track, as shown in Figure 2.1. KaRIn transmits Ka-band pulses from one of two antennas separated by a baseline length of approximately 10 m. The relative delay, or phase difference, between the reflected signals received by each of




Figure 2.1: Overview of SWOT Measurement System.

the two antennas with their known baseline length provides measurements of the water surface height, in effect through triangulation. KaRIn includes both a low-rate (LR) data stream that is primarily intended for oceanography and a high-rate (HR) data stream that is primarily intended for terrestrial hydrology.

The NAlt instrument complements KaRIn by providing range measurements along the nadir ground track, as shown in Figure 2.1. It has significant heritage from the Jason-3 mission[3] and its predecessors. The NAlt is primarily aimed towards measuring sea surface height along the nadir ground track over the ocean, conveniently bridging the gap between the KaRIn swath measurements on each side of nadir. Nevertheless, nadir altimeters have also been increasingly used to successfully recover measurements of water height over inland water bodies along the nadir ground track (e.g., [4]), though with significantly lower spatial resolution than is available from KaRIn and only at nadir.

High accuracy estimates of the satellite position are typically determined using measurements from the onboard DORIS and GPSP tracking systems, with the LRA measurements typically reserved for independent validation. Each of the three tracking systems have strong heritage from similar instruments onboard the Sentinel-6 satellite altimeter mission[5].

The range measurements from both KaRIn and NAlt are subject to delays caused by the troposphere and ionosphere. The primary purpose of the AMR is to provide measurements of the range delays and two-way atmospheric attenuation due to the wet troposphere at the Ku-, C-, and Ka-band radar frequencies used by NAlt and KaRIn. The SWOT AMR provides two independent, off-nadir troposphere measurements, one at the middle of each of the two KaRIn swaths. The AMR measurements are subject to land contamination and therefore primarily

aimed towards recovery of wet-troposphere effects over the open oceans. (They may also be valid over large inland water bodies such as large lakes.) Over inland water, however, the wet troposphere delays and atmospheric attenuation are generally determined from meteorological models. Range delays due to the dry troposphere, for both the NAlt and KaRIn measurements, are also determined from meteorological models over all surfaces. A by-product of the two NAlt frequencies is a direct measurement of the ionosphere delay at nadir, which is then used to correct the NAlt range measurements over oceans. The ionosphere range delay correction for the global KaRIn measurements is computed using global maps of total electron content.

2.2 Science Data Products

A list of the SWOT science data products is provided in Tables 2.1, 2.2, and 2.3. They include four products aimed towards oceanography users, nine products aimed towards hydrology users, and five products with satellite orbit position and velocity time series, tracking data, and attitude data. Each of these products is often referred to by their respective short names, which are also provided in these tables. The most recent versions of this document, the PDDs, and ATBDs are available at the PO.DAAC [SWOT Dataset Information](#) website.

Short Name	Description
L2_LR_SSH [6; 7]	KaRIn sea surface heights from the swath spanning 60 km on each side of nadir with a nadir gap. Also provides sea surface height anomaly, wind speed, and significant wave height, on a geographically fixed, swath-aligned 2 km×2 km grid, as well as sea surface height on a 250 m×250 m native grid.
L1B_LR_INTF [8; 9]	KaRIn low rate data processed to interferograms for each of the nine Doppler beams formed and spatially averaged by the On Board Processor, corrected on the ground for phase biases that are inherent to the processing undertaken on board. The geometry of the measurements is also reported for use in subsequent processing.
L2_NALT_GDR, L2_NALT_IGDR, L2_NALT_OGDR [10]	NAlt Geophysical Data Record (GDR) products similar to those from ongoing nadir altimeter missions such as Jason-3. Provides sea surface height anomaly, significant wave height and wind speed measurements from the nadir altimeter. The GDR uses restituted auxiliary data and the Precise Orbit Ephemeris (POE) and is available with a latency of less than 90 days. The IGDR uses the Medium-accuracy (preliminary) Orbit Ephemeris (MOE) and is available with a latency of less than 1.5 days. The OGDR uses the onboard DORIS orbit ephemeris and predicted values for some auxiliary data, does not have ionosphere model values, and is available with a latency of less than 7 hours.
L2_RAD_GDR, L2_RAD_IGDR, L2_RAD_OGDR [11; 12]	AMR Geophysical Data Record (GDR) products similar to those from the Sentinel-6 mission. The AMR products provide radiometer measurements of wet troposphere content, atmospheric attenuation corrections to backscatter, cloud liquid water, water vapor content, and wind speed. The GDR provides values based upon analyzed instrument calibrations and the Precise Orbit Ephemeris (POE) and is available with a latency of less than 90 days. The IGDR uses preliminary instrument calibrations and the Medium-accuracy (preliminary) Orbit Ephemeris (MOE) and is available with a latency of less than 1.5 days. The OGDR uses preliminary calibrations and the onboard DORIS orbit ephemeris and is available with a latency of less than 7 hours.

Table 2.1: Description of SWOT oceanography science data products.

Short Name	Description
L2_HR_PIXC [13; 14]	KaRIn high rate point cloud of water mask pixels, or pixel cloud (PIXC), with geolocated heights, backscatter, geophysical fields, and flags.
L2_HR_RiverSP [15; 16]	KaRIn shapefiles of river reaches (approximately 10 km long) and nodes (approximately 200 m spacing) that are identified in a prior river database. Reach attributes include water surface elevation, slope, width, and derived discharge.
L2_HR_RiverAvg [17]	KaRIn cycle average and aggregation of pass-level river reach data within predefined hydrological basins.
L2_HR_LakeSP [18; 19]	KaRIn shapefiles of lakes that are identified in a prior lake database and detected features that are not in the prior river or lake databases. Lake attributes include water surface elevation, area, derived storage change.
L2_HR_LakeAvg [20]	KaRIn cycle average and aggregation of pass-level lake data within predefined hydrological basins.
L2_HR_PIXCVec [21]	KaRIn auxiliary information for the pixel cloud product indicating to which water bodies the pixels are assigned in the river and lake products. Also includes height-constrained pixel geolocation after reach- or lake-scale averaging.
L2_HR_Raster [22; 23]	KaRIn rasterized water surface elevation and inundation extent in geographically fixed tiles at resolutions of 100 m and 250 m on Universal Transverse Mercator (UTM) projection grids. Provides rasters with water surface elevation, area, water fraction, backscatter, and geophysical information. On-demand processing is also available to users for different resolutions, sampling grids, scene sizes, and file formats.
L2_HR_FPDEM [24]	KaRIn floodplain digital elevation model (FPDEM), derived from multiple cycles of SWOT acquisitions. The expected resolution is approximately 50–100 m. A large fraction of the DEM pixels will be void. Provides height and quality flag for each pixel.
L1B_HR_SLC [25; 26]	KaRIn high rate data processed to single-look complex (SLC) synthetic aperture radar (SAR) images for each antenna.

Table 2.2: Description of SWOT hydrology science data products.

Short Name	Description
POE and MOE [27]	Precise Orbit Ephemeris (POE) providing time series of position and velocity vectors of the satellite center of mass. Available with a latency of less than 35 days. Medium-accuracy Orbit Ephemeris (MOE) providing time series of position and velocity vectors of the satellite center of mass. Available with a latency of less than 1.5 days.
ATTD_RECONST [28]	Satellite attitude reconstructed from a combination of onboard gyro and star tracker data. Available with a latency of less than 1.5 days.
SAT_COM [29]	Satellite center of mass position with respect to a satellite reference frame. Historical file updated daily.
L1_DORIS_RINEX [30]	DORIS tracking data. Available with a latency of less than 2 days.
L1_GPSP_RINEX [31]	GPSP tracking data. Available with a latency of less than 2 days.

Table 2.3: Description of SWOT orbit position, attitude, and tracking data science data products.

Chapter 3

Definitions and Conventions

This chapter gives definitions and conventions used throughout this document as well as in other SWOT documents such as the PDDs and ATBDs.

3.1 Terms

This section provides definitions of various terms as they are often used in the context of SWOT.

3.1.1 Ground Track and Nadir Track

The terms *ground track* or *nadir track* refer to the locus of points on the Earth surface directly under the SWOT spacecraft as it orbits the Earth. The ground track may be referenced to different models of the Earth surface in different contexts, but it is often defined in terms of latitude and longitude with respect to the reference ellipsoid (see Sect. 3.3.3).

3.1.2 Orbit

The term *orbit* can refer to the generic characteristics of the SWOT orbital trajectory around the Earth, for example as shown in Table 4.2. The term *orbit* is also used in some contexts to refer to one specific revolution (see Sect. 3.1.3) or to refer to one of the SWOT orbit phases (see Sect. 3.1.5). The meaning should generally be clear in context.

3.1.3 Revolution

The term *revolution* typically refers to a particular orbital period (approximately 100 min) that encompasses one trip of the SWOT spacecraft around the Earth.

3.1.4 Pass (Ascending and Descending)

The term *pass* refers to a particular half of an orbital revolution from either the southernmost to northernmost or the northernmost to southernmost latitudes covered by the orbit. Passes during which the spacecraft travels from southern to northern latitudes are called *ascending passes*, while passes during which the spacecraft travels from northern to southern latitudes are called




Figure 3.1: Illustration showing the definition of the beta (β) angle.

descending passes. Note that the terms “ascending” and “descending” here do not refer to the orbital altitude.

3.1.5 Orbit Phase (Calibration Orbit and Nominal or Science Orbit)

The term *orbit phase* refers to either of two spans of time during which the SWOT orbit is controlled to follow a predefined ground track that repeats (to within ± 1 km) periodically. The orbital altitude is similar but slightly different between these two orbit phases (see Table 4.2), so the ground track and repeat period differ between them. Shortly after launch, the orbit is designed to have a repeat period of approximately 1 day. This is called the *calibration orbit*. After several months in the calibration orbit, a series of orbital maneuvers is undertaken to transition to the *nominal orbit* (also called the *science orbit*), which has a repeat period of approximately 21 days. The calibration orbit gives only sparse spatial coverage of the Earth, but its short revisit times are advantageous for calibration and validation activities. These short revisit times may also be useful for science objectives that relate to phenomena with temporal variations at scales of days, provided that these phenomena can be studied over the limited spatial coverage of the calibration orbit. The nominal orbit provides near global spatial coverage at the expense of temporal coverage (i.e., lengthier revisit times). Additional details on the orbit design are given in Sect. 4.2.

The calibration and nominal orbit phases are sometimes called the “1-day” and “21-day” orbits. The calibration orbit is also sometimes called the “fast-sampling” orbit.

3.1.6 Cycle

The term *cycle* refers to the set of passes or revolutions in one repeat cycle of the orbit. The repeat cycle depends on the orbit phase (see Sect. 3.1.5). A cycle of the calibration orbit is approximately 1 day long and includes 14 revolutions or 28 passes. A cycle of the nominal orbit is approximately 21 days long and includes 292 revolutions or 584 passes. See Ch. 4 for details on the cycle numbering scheme used by SWOT and for more precise details of the orbit design.

3.1.7 Beta Angle

The term *beta angle* or (*β angle*) refers to the angle between the vector from the Earth to the sun and the plane containing the spacecraft orbit around the Earth (see Fig. 3.1). The beta angle is given by

$$\beta = \sin^{-1}(\hat{r}_{ES} \cdot \hat{n}_o), \quad (3.1)$$

where \hat{r}_{ES} is a unit vector from the Earth to the Sun, and \hat{n}_o is a unit vector that is normal to the spacecraft orbit plane and pointed in the right-handed direction relative to the spacecraft orbital motion. The dot operator (“ \cdot ”) represents the vector inner product or scalar product. The beta angle is zero when the vector to the sun lies within the plane of the orbit. The beta angle evolves gradually as the orbit precesses and the Earth revolves around the Sun.

The beta angle affects the solar illumination of the spacecraft and therefore influences the thermal environment of the SWOT spacecraft and its instruments. The beta angle may therefore be correlated with some systematic errors in the instruments. The beta angle is also related to the pattern of spacecraft eclipse transitions during which the spacecraft goes into or out of the Earth’s shadow. Eclipse events can cause thermal transients that may influence instrument performance. For example, KaRIn performance is expected to be impacted by eclipse events, so its measurements are flagged during and shortly after eclipse transitions (see Sect. 4.5.5). Additionally, spacecraft yaw flips (see Sect. 3.1.8) and solar array rotations (see Sect. 4.5.4) are driven by the variations in the beta angle.

3.1.8 Yaw Flip

The term *yaw flip* refers to a change in the nominal attitude of the SWOT spacecraft by 180° about an axis that is parallel to the nadir vector. The design of the SWOT spacecraft is such that, for thermal reasons, one side of the spacecraft is intended to face away from the sun preferentially. As the beta angle changes, and as the vector to the sun moves from one side of the orbit plane to the other, it is necessary for the spacecraft to execute yaw flips near $\beta = 0$ to keep the anti-sun side of the spacecraft indeed facing away from the sun. Yaw flips occur around every 2.5 months (approximately 78 days). From one yaw flip to the next, the side-specific KaRIn antennas and instrument hardware that image the left and right sides of the swath (see Sect. 3.1.14) alternate.

3.1.9 Forward and Backward

The terms *forward* (or “head first”) and *backward* (or “tail first”) are sometimes used to refer to the spacecraft yaw state or orientation with respect to the spacecraft flight direction. As a consequence of yaw flips, on average the spacecraft is flying forward about half the time, and backward the other half. The forward direction corresponds to yaw angles near zero, while the backward direction corresponds to yaw angles near 180° (see Sect. 3.3.6). SWOT is designed so that measurements in either yaw state should be equally accurate.

3.1.10 Along Track

The term *along track* refers to the direction that is parallel to the SWOT nadir track on the ground. The along-track direction is typically positive in the SWOT flight direction. The along-track distance between two targets on the Earth surface is usually taken to be the distance along the curved surface of the planet, although this distance is usually approximate in any case given the cross-track curvature of the nadir track as well as the eccentricity of the reference ellipsoid (see Sect. 3.3.3).

3.1.11 Cross Track

The term *cross track* (sometimes also *across track*) refers to the direction along the Earth surface that is perpendicular to the SWOT nadir track on the ground. In some contexts, the cross-track direction is positive to the right (see Sect. 3.1.14) whether on the left or right sides of the swath, while in other contexts, the cross-track direction is positive away from nadir in either direction. The convention should be clear from the context where the distinction is important. For targets on the Earth surface, the cross-track distance from nadir is usually defined as the distance over a locally spherical approximation to the Earth surface rather than a Cartesian distance, though the additional length over the curved surface is relatively small. The Cartesian distance is assumed in reference to the TCN frame described in Sect. 3.3.5. The cross-track distance provided in SWOT LR and HR products is positive to the right.

3.1.12 Heading

The term *heading* is defined as the angle of the SWOT nadir track with respect to true north at a given nadir location. Headings of 0°, 90°, 180°, and 270° (or −90°, equivalently) correspond to northward, eastward, southward, and westward travel over the surface, respectively. The SWOT orbit is such that the spacecraft heading varies between approximately 8° for ascending passes near the equator and 172° for descending passes near the equator (see Sect. 4.2). Note that the heading is independent of the spacecraft attitude (particularly the yaw state). The heading depends only on the spacecraft velocity vector relative to the local Earth frame.

3.1.13 Swath and Half Swath

The term *swath* refers to the span of KaRIn measurements in the cross-track direction. The KaRIn requirements specify a swath width of ± 60 km from the nadir track. The term *half swath* refers to the portion of the KaRIn swath to only one side or the other of the nadir track. In some contexts, the term *swath* is also used to refer to measurements to only one side of nadir (as in “left swath” or “right swath”), and the term *full swath* is used to refer to measurements on both sides. The meaning should be clear in context.

3.1.14 Left and Right

The terms *left* and *right* refer to the directions to the left and right as if standing on the Earth surface on the SWOT nadir track and facing in the direction of the spacecraft flight direction (see Fig. 3.2). Which channels in the KaRIn and AMR hardware collect data for the left and right sides depends on the yaw state, or orientation, of the spacecraft (see Sect. 3.1.8). Namely, which of the KaRIn or AMR antennas observe the left side and which observe the right side depend on whether the spacecraft is flying forward or backward (see Sect. 3.1.9).

KaRIn LR and HR data products are defined in terms of the left and right sides of the swath regardless of the spacecraft yaw state. On the other hand, AMR data products are defined in terms of the hardware channels collecting the data.

Note that when viewed on a typical map projection with north pointed toward the top of the screen or the page, the left half swath would be to the east of the nadir track for descending




Figure 3.2: Illustration of how left and right are defined for different pass directions and yaw states: (a) nominal yaw in an ascending pass; (b) nominal yaw in a descending pass; (c) flipped yaw in an ascending pass; (d) flipped yaw in a descending pass. The KMSF $+x$ and $+y$ axes (see Sect. 3.3.6) and the KaRIn half-swath polarizations (H and V for horizontal and vertical), which depend on the yaw state, are shown in green. North is assumed to be toward the top of the page for all four cases.

passes (SWOT moving toward the south) and would therefore appear to the right on the screen or the page.

3.1.15 Horizontal (H) and Vertical (V) Polarization

The terms *horizontal* (H) and *vertical* (V) in reference to a KaRIn half-swath polarization indicate the orientation of the electric field (i.e., the polarization) of the radar signal used by KaRIn for that side. The KaRIn signals for each side use opposite, linear electromagnetic polarizations in order to better isolate the interleaved measurements for the two sides from each other.

The nomenclature of calling these polarizations horizontal and vertical follow standard representation conventions in SAR imaging. The horizontally polarized signal has its electric field vector nominally parallel to the along-track direction and therefore perpendicular to the local vertical direction. The vertically polarized signal has its electric field vector nominally perpendicular to the along-track direction and therefore partially aligned with the local vertical direction. (The electric field vectors of both are perpendicular to the direction of signal propagation as is always the case for free-space propagation of electromagnetic waves.) However, at the steep incidence angles of the KaRIn observations, even the vertically polarized signal has its electric field vector nearly parallel to the horizontal (cross-track) direction.

3.1.16 Point Target Response (PTR)

The term *point target response* (PTR) refers to the response of a radar measurement system to a hypothetical, idealized, infinitely narrow, point-like target. The PTR typically has a main lobe as well as sidelobes, whose characteristics are determined by the system design. The overall observation is typically modeled as the linear combination (superposition) of the responses to many individual point-like targets with additive noise. The shape of the PTR determines the ability of the measurement system to resolve small features on the surface and to distinguish contrasting parcels of the surface with different properties. See Ch. 8 for additional discussion. The PTR is similar to the point spread function (PSF) in optical systems. The PTR is also known as the impulse response or impulse response function (IRF).

3.1.17 Resolution

The term *resolution* describes the ability of the KaRIn measurement to distinguish spatial variations on the Earth surface. The precise mathematical definition of the resolution can vary with context, but, within this document and the SWOT PDDs and ATBDs, it is typically the full-width at half maximum (FWHM) or “3 dB width” of the power of the PTR in a given direction (see Fig. 3.3). Note that the half-power resolution is therefore the width at $1/\sqrt{2}$ of the peak magnitude of the field, voltage, or complex envelope of the PTR. The resolution defined in this manner does not depend on the level of additive noise. The resolution may be different in different directions (e.g., the along-track resolution of KaRIn measurements may be finer than the cross-track resolution). The resolution is not necessarily the same as the sampling (see Sect. 3.1.18).

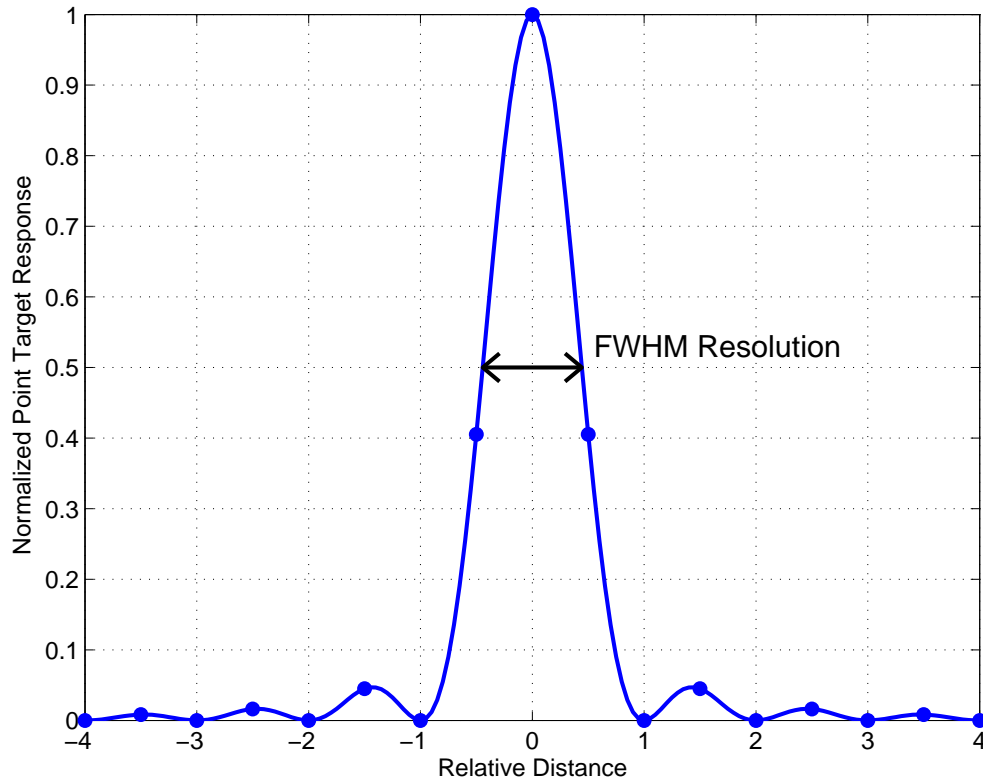


Figure 3.3: Illustration of the definition of the resolution. The continuous blue curve is an ideal $\text{sinc}^2(x)$ curve that is meant to represent the power of the PTR of the measurement system in the absence of noise, where $\text{sinc}(x) = \sin(\pi x)/(\pi x)$. The FWHM resolution is the width of the main lobe of the PTR at half of its peak power. Blue dots represent a notional sampling of the signal. In this case, the sampling rate is exactly twice the Nyquist limit for the complex response before power detection such that the sample spacing is approximately half the resolution.

3.1.18 Sampling

The term *sampling* describes the spatial or temporal discretization of the measurements. The sampling may be finer than the resolution (see Sect. 3.1.17). In such cases, the data would be said to be oversampled, and neighboring samples would be highly correlated with each other. This is illustrated in Fig. 3.3.

3.1.19 Pixel and Line

In some SWOT documents, the terms *pixel* and *line* refer to the cross-track and along-track sample indices of a 2-D array of KaRIn data, respectively. This usage should not be confused with the colloquial meaning of the term “pixel” that refers to any single measurement from a 2-D image. The meanings of the terms should be clear in context.

3.1.20 Altitude

The term *altitude* refers to the distance of a point (usually the satellite center of mass) above the reference ellipsoid (see Sect. 3.3.3). In some contexts, the altitude may be referenced to something other than the reference ellipsoid, but this would be noted explicitly.

3.1.21 Range

The term *range* refers to the distance between a radar antenna (NAIt or KaRIn) and a target, which, for SWOT, is typically some location on the surface of the Earth. The term may refer to the range as measured by the radar (e.g., the NAIt measurement is referred to as the “range” or “altimeter range”), or it may refer to the true or assumed range. The precise definition is noted explicitly in context where the distinction is important.

3.1.22 Corrected Range

The term *corrected range* has traditionally referred to the range measured by the radar after applying corrections for geophysical effects that impact the radar signals (see Sects. 11.1 and 11.2). These include the dry troposphere (τ_{dry}), wet troposphere (τ_{wet}), ionosphere (τ_{iono}), and sea state bias (τ_{ssb}) corrections. These corrections follow the sign conventions described in Sect. 3.3.8:

$$\text{Corrected Range} = \text{Measured Range} + \tau_{\text{dry}} + \tau_{\text{wet}} + \tau_{\text{iono}} + \tau_{\text{ssb}}. \quad (3.2)$$

For example, the NAIt products report the measured range and the above-mentioned corrections, so the user may construct corrected range if desired. The KaRIn products do not report range. Note also that the KaRIn LR products do apply a sea-state-bias correction, while the HR products do not.

3.1.23 Height

The term *height* refers to the distance of the water surface above the reference ellipsoid (see Sect. 3.3.3). This includes the “sea surface height” (SSH) over the ocean and the “water surface height” over inland water. The height of any given location on the Earth surface (including

ocean, inland water, and land) would be expected to change over time due to tide effects (see Sect. 11.3.4).

In the simplest case of a nadir range measurement (e.g., from the NAlt), the height H (SSH or water surface height) is as follows:

$$H = \text{Altitude} - \text{Corrected Range.} \quad (3.3)$$

3.1.24 Sea Surface Height Anomaly

In this document and in the NAlt and KaRIn LR products and documentation, the term *sea surface height anomaly* (SSHA) refers to the elevation of the ocean surface relative to the mean sea surface (see Sect. 11.3.2) after subtracting tide effects (see Sect. 11.3.4) and the so-called dynamic atmosphere correction (see Sect. 11.3.5.2), as shown in the equation below. In this equation, SSH refers to the sea surface height (see Sect. 3.1.23), and the other terms are self-explanatory. In the case of SSHA, the pole tide height refers to the sum of the solid Earth pole tide height, the ocean pole tide height, and the load pole tide height. The internal tide refers to the coherent internal tide:

$$\begin{aligned} \text{SSHA} = & (\text{SSH} \\ & - \text{mean_sea_surface_height} \\ & - \text{solid_earth_tide_height} \\ & - \text{ocean_tide_height} \\ & - \text{load_tide_height} \\ & - \text{internal_tide_height} \\ & - \text{pole_tide_height} \\ & - \text{dynamic_atmosphere_correction}). \end{aligned} \quad (3.4)$$

3.1.25 Water Surface Elevation

In this document and in KaRIn HR products and documentation, the term *water surface elevation* (WSE) refers to the elevation of an inland water surface relative to the geoid (see Sect. 11.3.1) after subtracting tide effects (see Sect. 11.3.4), as shown in the equation below. In this equation, H refers to the water surface height (see Sect. 3.1.23), and the other terms are self-explanatory. In the case of WSE, the pole tide height refers to the sum of the solid Earth pole tide height and the load pole tide height:

$$\begin{aligned} \text{WSE} = & (H \\ & - \text{geoid_height} \\ & - \text{solid_earth_tide_height} \\ & - \text{load_tide_height} \\ & - \text{pole_tide_height}). \end{aligned} \quad (3.5)$$

3.1.26 Good, Suspect, Degraded, and Bad Quality

In the context of KaRIn LR and HR products, quality flags often indicate *good*, *suspect*, *degraded*, and *bad* measurements. These terms have specific meanings in this context, as given below. Note, however, that non-KaRIn products may not make similar distinctions between “suspect” and “degraded” data.

- A *good* measurement is typically not associated with any off-nominal conditions such as failed quality or self-consistency checks, missing or flagged input data, or atypical logic during data processing. *Good* measurements are thus sometimes also called “nominal” measurements. Users would normally use *good* data. By convention, *good* values are indicated by quality-flag values of zero in most SWOT products (see the PDD of a specific product for details on the definitions of quality flags).
- A *suspect* measurement is typically associated with a condition for which the quality of the measurement is cast in doubt. Such measurements may have errors that are larger than normal, but they may also have errors that are in family with those of *good* measurements. Thus, *suspect* measurements are not necessarily known to be worse than *good* measurements. As an example, a measurement may be flagged as *suspect* because it failed a quality check that is not always a reliable indicator of the true quality of the data, and downstream data may be flagged as *suspect* because they used *suspect* upstream data as input. Users may choose to use *suspect* data with caution.
- A *degraded* measurement is typically associated with a problem that is known to give larger than typical error. For example, if a geophysical correction is not available, a measurement might be produced anyway but be flagged as *degraded*. *Degraded* data should typically not be used unless there are insufficient *good* or *suspect* data available, and *degraded* data should be used only with considerable caution.
- A *bad* measurement is typically associated with a major problem that often makes the measurement completely unreliable. *Bad* measurements are often null filled. When they are not null filled, they may be associated with gross errors. Measurements that are flagged as *bad* should generally not be used.

In this document, the terms *good*, *suspect*, *degraded*, and *bad* are often italicized when they are used to indicate specific meanings in terms of quality states (as defined as in this section) rather than used colloquially.

Additional information on what conditions led to particular quality flag values is often available in KaRIn LR and HR products. Users should consult the PDDs of specific products for details.

3.2 Data Products, Granules, and Files

The SWOT mission is intended to meet the needs of different user communities. Consequently, SWOT data are available in many forms, each involving different levels and/or types of processing, which may emphasize different aspects of the SWOT observations. In this context, the term

data product (or often simply *product*) refers to a particular form of data after a given set of standard processing. The data product definition includes the format(s) of the file(s) that are present in the product, the names of data fields present in the file(s), the representation conventions for the data fields, metadata, and so forth. An overview of the set of available SWOT science data products is given in Tables 2.1, 2.2, and 2.3, while a more extensive introduction to the products is given in Ch. 6. The detailed definition of each product is specified in the PDD for that product (see the PDD references provided in these three tables).

While a data product represents the form of the data, users work with specific granules of the products, where the term *granule* in this context represents the downloadable unit of the data product. That is, a granule of data from a given product covers a predetermined temporal or spatial extent of SWOT measurement data. Different data products have different granule sizes. For example, the L2_LR_SSH product is available in pass-sized granules, while the L2_HR_PIXC product is available in 64 km \times 64 km tile-sized granules. The granule definitions for each product are given in the PDD, with additional details available in [32].

A granule of a particular data product may consist of one or more *files*. For products with more than one file per granule, the files in a granule typically cover the same (or very similar) temporal and/or spatial extent but provide complementary information to one another. For example, the L2_LR_SSH product consists of four files per granule (Basic, WindWave, Expert, and Unsmoothed) that are tailored to different use cases; these files may be used individually or together. Similarly, the L2_HR_RiverSP product, which is provided in a shapefile format, consists of a reach-level shapefile and a node-level shapefile, with each shapefile itself comprising multiple files (having extensions .shp, .shx, .dbf, .prj, and possibly .shp.xml).

3.3 Coordinate Systems and Representations

This section provides a brief overview of several coordinate systems that are used in SWOT data products.

3.3.1 GCRF

The Geocentric Celestial Reference Frame (GCRF) is an Earth-centered, inertial coordinate frame whose orientation is fixed to the stars. Therefore, the GCRF coordinates of an arbitrary point on the Earth surface vary as the planet rotates. The $+z$ axis points toward the north pole. The $+x$ axis points toward the mean equinox of Earth at 2000-01-01 12:00 Terrestrial Time (TT), which is equivalent to 2000-01-01 11:58:55.816 UTC or 2000-01-01 11:59:27.816 TAI.

3.3.2 ITRF

The International Terrestrial Reference Frame (ITRF) is an Earth-centered, Earth-fixed (ECEF), Cartesian coordinate frame. This frame is also sometimes denoted XYZ. The frame has its origin at the center of mass of the planet. The $+z$ axis points toward the north pole, the $+x$ axis points along the Greenwich (prime) meridian, and the $+y$ axis completes the right-handed system. The frame rotates with the Earth such that a point on the Earth surface has constant coordinates

Ellipsoid Parameter	Value
Semi-major axis	6378137 m
Semi-minor axis	6356752.314245 m
Flattening	0.00335281066474748
Inverse Flattening	298.257223563

Table 3.1: Parameters for the WGS 84 reference ellipsoid used by SWOT products as of the beginning of the mission.

in this frame, notwithstanding the motions of the surface itself (for example, due to tides; see Sect. 11.3.4) relative to the center of the Earth.

As of the beginning of the mission, SWOT products adopted the 2014 version of the ITRF (ITRF14) [33] but may transition to the 2020 version of the ITRF (ITRF20) [34]. ITRF14 is defined as a linear function of time, while ITRF2020 includes a seasonal (annual and semiannual) geo-center model along with the linear function of time. The frame of measurements provided in SWOT products is as realized by the adopted ITRF at the measurement epoch, consistent with other nadir altimeter products.

3.3.3 Reference Ellipsoid

SWOT products generally contain the parameters of the reference ellipsoid that defines latitude, longitude, and ellipsoid-relative height in the products. As of the beginning of the mission, the SWOT products adopted the WGS 84 (World Geodetic System 1984) ellipsoid [35] as the reference ellipsoid. The parameters of the reference ellipsoid are given in Table 3.1. However, the parameters for the reference ellipsoid that are given in the metadata of the product file itself should be considered definitive and may supersede this document.

3.3.4 LLH

Positions in SWOT products are often represented in latitude-longitude-height (LLH) coordinates. The geodetic latitude and the ellipsoid-relative height in these coordinates are defined with respect to the reference ellipsoid described in Sect. 3.3.3. See [36] for a method to convert between ECEF and LLH coordinates.

Latitudes are given as signed values that increase to the north of the equator. Longitude values are usually represented as angles between -180 and $+180^\circ$ in KaRIn hydrology (HR) products, where it is desirable to place the longitude wrap in the middle of the Pacific Ocean, away from land. On the other hand, longitude values are usually represented as angles between 0 and 360° in KaRIn oceanography (LR), NAlt, and AMR products, where it is undesirable to place the longitude wrap in the middle of the Pacific Ocean. Longitude values increase to the east of the Greenwich (prime) meridian.

Because of the SWOT orbit inclination (see Sect. 3.4), latitudes beyond approximately $\pm 78^\circ$ are not observed.

3.3.5 TCN and HCL

Some SWOT processing algorithms use a local, Cartesian frame whose axes are the track, cross-track, and nadir (TCN) directions. This frame is oriented so that the $+z$ axis (N) is pointed toward nadir from the spacecraft, with the downward direction defined by the reference ellipsoid (see Sect. 3.3.3). The $+x$ axis (T) is parallel to the horizontal component of the spacecraft ECEF velocity vector. The $+y$ axis (C) completes the right-handed system such that it points to the right (see Sect. 3.1.14) of the nadir track.

The origin is usually taken to be at or near the spacecraft center of mass because this frame is used for orienting more so than for positioning. The exact location of the origin does not matter much as long as the origin does not deviate very far horizontally from the spacecraft or the nadir point.

The TCN coordinate system is analogous to an HCL coordinate system nomenclature that is sometimes used in other applications. The HCL coordinate system refers to the radial, cross-track and along-track axes, respectively, where $H = -N$, and $L = T$.

3.3.6 KMSF

The KaRIn metering structure frame (KMSF) is a KaRIn body-fixed frame that is often used to define coordinates with respect to the instrument (see Fig. 3.4). The KMSF origin is approximately 41 cm above the midpoint of the KaRIn interferometric baseline. The KMSF $+z$ direction points approximately downward when KaRIn is in its nominal science orientation. The frame is oriented such that the KaRIn interferometric baseline is approximately parallel to the KMSF y axis. The KaRIn antennas at each end of the baseline are often called the $+y$ and $-y$ antennas based on their locations in KMSF coordinates. The KMSF $+x$ axis is either approximately parallel (when flying forward) or approximately antiparallel (when flying backward) to the spacecraft Earth-relative velocity vector, depending on the yaw state, or orientation, of the spacecraft (see Sect. 3.1.8). The KMSF principal axes are approximately aligned with the axes of the TCN frame when the yaw is near 0. The sun nominally shines on the $-y$ side of the spacecraft.

The KaRIn radar signal is horizontally polarized (H) for the half swath on the $-y$ side of the spacecraft and vertically polarized (V) for the half swath on the $+y$ side of the spacecraft. However, the H and V polarizations switch between the left and right sides (see Sect. 3.1.14) depending on whether the spacecraft is yaw flipped.

The AMR has two active strings which together provide measurements on the left and right sides of nadir, approximately half-way across each of the two KaRIn half swaths. The two AMR strings are named $+y$ and $-y$, in reference to the half-swath measurement locations on the respective sides of the KMSF body-fixed frame.

3.3.7 UTC and TAI Time Conventions

Many SWOT products give times in both UTC (coordinated universal time) and TAI (international atomic time) representations. That is, the UTC time of a given SWOT measurement in a particular product may be provided in a variable called *time*, while the TAI time of the same measurement may be provided in that product in a variable called *time_tai*. (The variable names




Figure 3.4: Illustration of the KaRIn Metering Structure Frame (KMSF). The spacecraft velocity direction is not shown because it could be approximately parallel or approximately antiparallel to the KMSF $+x$ direction depending on the yaw state of the spacecraft.

UTC Date	TAI Date	<i>time</i>	<i>time_tai</i>	<i>time_tai - time</i>
2000-01-01 00:00:00	2000-01-01 00:00:32	0.0	32.0	32
2016-12-31 23:59:59	2017-01-01 00:00:35	536543999.0	536544035.0	36
2016-12-31 23:59:59.5	2017-01-01 00:00:35.5	536543999.5	536544035.5	36
2016-12-31 23:59:60	2017-01-01 00:00:36	536543999.0	536544036.0	37
2017-01-01 00:00:00	2017-01-01 00:00:37	536544000.0	536544037.0	37
2017-01-01 12:00:00	2017-01-01 12:00:37	536587200.0	536587237.0	37

Table 3.2: Examples of times represented in both UTC and TAI following SWOT product conventions. A leap second was inserted on December 31, 2016.

here are examples; the actual names of the variables containing time information are specific to individual products, and readers should consult the PDDs for those products as needed.)

UTC times in SWOT products are represented as the non-continuous number of UTC seconds since January 1, 2000 00:00:00 UTC, which is equivalent to January 1, 2000 00:00:32 TAI. The UTC time exhibits a jump when a leap second occurs. Midnight UTC will always be an integer multiple of 86400 s, which is the number of seconds in one day. Information about the number of leap seconds over the time coverage of the granule is given in the product metadata for each granule.

TAI times in SWOT products are represented as the continuous number of TAI seconds since January 1, 2000 00:00:00 TAI, which is equivalent to December 31, 1999 23:59:28 UTC. The TAI time does not exhibit jumps when leap seconds are added to (or removed from) UTC times.

Note that the times in SWOT products are referenced to midnight—not noon—on January 1, 2000.

Table 3.2 provides some examples for the corresponding values of the UTC time (variable *time*) and the TAI time (variable *time_tai*) as well as the difference between UTC and TAI time. Some products provide this difference as an attribute name *tai_utc_difference*. Note that *time* has a 1 s regression during a leap second transition, while *time_tai* is continuous. That is, when a positive leap second is inserted, two different instances will have the same value for the variable *time*, making *time* non-unique by itself; the information on leap seconds given in the product metadata can be used to resolve this.

3.3.8 Correction Convention

All environmental and instrument corrections are computed, and provided in data products, such that they should be added to the quantity that they correct. That is, a correction is applied to a measured value as follows:

$$\text{Corrected Value} = \text{Measured Value} + \text{Correction.} \quad (3.6)$$

For example, a correction to radar range for an effect that lengthens the apparent signal path (e.g., ionosphere or troposphere delay) is computed as a negative value. Adding this negative value to the uncorrected (measured) range shortens the range to the corrected value. This same convention applies to the KaRIn crossover calibration correction, whereby this correction is added to uncorrected sea surface and inland water heights.

Note, however, that some geophysical effects (see Sect. 11.3) such as tides and the so-called dynamic atmosphere correction are sometimes referred to as “corrections,” but these are effects that are subtracted from the observed physical surface height as indicated by Eqs. (3.4) and (3.5), so they do not necessarily follow the “correction” convention of Eq. (3.6).

3.4 Reference Orbit

During the calibration and science phases of the mission (see Sect. 3.1.5), the SWOT spacecraft is intended to follow predefined, repeating reference-orbit ground tracks for each phase. The SWOT spacecraft is controlled via propulsive station-keeping maneuvers to keep deviations of the actual SWOT ground track from the reference orbit under ± 1 km most of the time. The control window can be extended to ± 2.5 km for 10% of the time to meet power-limitation constraints as necessary, but this is not expected to occur nominally. Details on the orbit design are given in Sect. 4.2.

Chapter 4

Mission Description

This chapter describes aspects of the SWOT mission design that may be relevant to users of SWOT data. This chapter includes discussion of the SWOT orbit design and representation conventions, the payload design, and operational spacecraft events that have bearing on the SWOT data.

4.1 Mission Phases and Cycle Numbering

The primary SWOT mission phases are listed in Table 4.1. The mission began with the Launch and Early Operations Phase (LEOP). This phase is assigned a cycle number of 990. During this phase, various functional checks were performed on the satellite platform, and a series of maneuvers were executed to insert the satellite into the Calibration (1-day) Orbit.

The Calibration Orbit Phase consists of the Commissioning and Calibration Phases. These phases are assigned cycle numbers 400–578. The science instruments were turned on over the course of the Commissioning Phase. During this phase, various changes were made to the configuration of the SWOT satellite, especially within the KaRIn instrument. The final changes to the in-flight configuration of the KaRIn instrument that affect the ground calibration of the KaRIn data were made on March 30, 2023, (cycle 475) thereby marking the start of the Calibration Phase. KaRIn data collected during the Calibration Phase were used to calibrate and to validate the measurement. That is, data collected between cycles 475–578 were used to estimate calibration parameters that are applied in KaRIn ground processing. These calibration parameters are applicable to data from cycle 475 onward but not to data before cycle 475 because of the on-board KaRIn configuration changes prior to cycle 475. The calibration parameters are retroactively applied to the KaRIn data from which they were estimated when reprocessed SWOT products are generated (see Sect. 5.3).

From July 11–21, 2023 a series of maneuvers was performed to insert the satellite into the Nominal (21-day) Orbit. The transition period between the calibration and nominal orbit phases is assigned a cycle number of 999. Thereafter, cycles from the nominal orbit are numbered sequentially from 001. Cycle numbers up to 399 are reserved for the nominal orbit (SWOT would not reach cycle 399 until the year 2046).

No science data products will be released for cycles 900 and 999, due to the drifting satellite

Phase	Start Time (UTC)	Repeat Cycles
Launch and Early Operations (LEOP)	2022-12-16 11:46:47	990
Calibration (1-day) Orbit	2023-01-15 06:01:50	400-578
Commissioning Phase	2023-01-15 06:01:50	400-475
Calibration Phase	2023-03-30 19:05:00	475-578
Orbit Maneuvers and Drift	2023-07-11 02:13:59	999
Nominal (Science or 21-day) Orbit	2023-07-21 05:33:46	001-[399]

Table 4.1: Dates and cycle numbers of SWOT mission phases.

orbit, numerous maneuvers, and changes to the satellite platform configuration. KaRIn products from the Commissioning Phase (cycles 400–475) may be generated on a best-efforts basis. Should data from Commissioning Phase become available, users should treat them with extreme caution, as the ground processing is not likely to capture the various changes to the KaRIn instrument configuration during this period. As such, nominal KaRIn science data products are expected only for cycles 475–578 of the Calibration Orbit and cycles 001 onward for the Nominal Orbit. All of the other (non-KaRIn) science data products are available for the Calibration and Nominal Orbit Phases (cycles 400-578, and 001 onward) whenever data were collected.

4.2 Reference Orbits and Pass Numbering

As defined in Sect. 3.1.5, the SWOT mission includes two reference orbits, a calibration orbit and a nominal orbit. The calibration orbit is designed to provide short repeat times (approximately 1 day) at the expense of limited geographic coverage in order to facilitate checkout and calibration activities. On the other hand, the nominal orbit is designed to achieve near-global coverage, albeit with a longer repeat period (approximately 21 days). Parameters for the reference orbits from these two phases are given in Table 4.2. In this table, the reference altitude is defined as the mean semimajor axis minus the equatorial radius of the reference ellipsoid. However, periapsis occurs at high latitudes, not over the equator, so the actual altitude above the ellipsoid is always larger than the reference altitude.

Each cycle of the calibration orbit has 28 passes that are numbered from 001 to 028. The repeat period of the calibration orbit is 0.99349 days. Consequently, the time of day at which a given location on the Earth is imaged becomes approximately 9 minutes earlier with each successive day.

Each cycle of the nominal orbit has 584 passes that are numbered from 001 to 584. The repeat period of the nominal orbit is 20.86455 days. Consequently, the time of day at which a given location on the Earth is imaged from a given pass becomes approximately 3 hours earlier with each successive cycle.

For both the calibration orbit and the nominal orbit, passes with the same pass number but from different cycles will have nearly identical ground tracks. Passes with odd numbers (001, 003, 005, ...) are ascending passes, while passes with even numbers (002, 004, 006, ...) are descending passes (see Sect. 3.1.4).

The nominal orbit is designed such that pass 332 is aligned with pass 028 of the calibration

Parameter	Calibration Orbit	Nominal Orbit
Mean semimajor axis (km)	7235.3798	7268.7189
Earth equatorial radius (km)	6378.137	6378.137
Reference altitude (km)	857.244	890.582
Maximum altitude above ellipsoid (km)	887.745	921.070
Mean altitude above ellipsoid (km)	871.396	904.713
Minimum altitude above ellipsoid (km)	861.775	895.091
Mean ECEF velocity magnitude (km/s)	7.314	7.297
Mean eccentricity	0.00105	0.00105
Inclination (deg)	77.6	77.6
Maximum spacecraft geodetic latitude (deg)	± 77.66	± 77.66
Nodal period (s)	6131.25	6173.62
Number of revolutions per cycle	14	292
Number of passes per cycle	28	584
Exact repeat cycle duration (days)	0.99349	20.86455
Mean longitude step between successive ground tracks (deg)	-25.714285714	-25.890410959
Mean longitude step between adjacent ground tracks (deg)	25.714285714	1.232876712
Pass length along ellipsoid (km)	19722.934	19721.166
Mean speed of nadir point along ground track (km/s)	6.434	6.389

Table 4.2: Parameters for the reference orbits of the calibration and nominal phases. The WGS 84 ellipsoid is assumed for these parameters.

orbit. Both of these (descending) passes have a reference equator crossing at -144.674° longitude.

4.3 Orbit Characteristics

Plots of the spacecraft altitude above the ellipsoid, the magnitude of the ECEF velocity vector, the nadir point speed along the ground track, the vertical component of the velocity vector, and the heading (see Sect. 3.1.12) are shown in Figs. 4.1, 4.2, 4.3, 4.4, and 4.5. The altitude and velocity magnitude are shown for both the calibration and nominal orbits. The vertical velocity is shown only for the nominal orbit, as the vertical velocity is nearly the same for the calibration orbit. The vertical velocity is defined so that an upward velocity is positive. The terms “ascending” and “descending” in Fig. 4.2 refer to the sign of the change in latitude, not altitude, over a pass (see Sect. 3.1.4).

An example plot of the beta angle (see Sect. 3.1.7) is shown in Fig. 4.6. Note that the amplitude of the variations in beta changes from cycle to cycle. Yaw flips occur when the beta angle crosses zero.

4.4 Satellite Description

The SWOT satellite or flight system consists of the spacecraft bus and the payload module, as shown in Fig. 4.7. The payload module consists of the KaRIn module and the Nadir module. The




Figure 4.1: Reference orbit altitude as a function of latitude for the nominal and calibration orbit phases of the mission. The altitude is defined with respect to the WGS 84 ellipsoid.




Figure 4.2: Reference orbit velocity magnitude as a function of latitude for the nominal and calibration orbit phases of the mission. The velocity is defined in an Earth-fixed frame.




Figure 4.3: Reference orbit speed of the nadir point along the ground track as a function of latitude for the nominal and calibration orbit phases of the mission. The ground track is defined on the WGS 84 ellipsoid.




Figure 4.4: Reference orbit vertical velocity as a function of latitude for ascending and descending passes. The vertical velocity for the nominal phase of the mission is shown; the vertical velocity for the calibration phase is nearly indistinguishable in such a plot.




Figure 4.5: Reference orbit heading. The heading for the nominal phase of the mission is shown; the heading for the calibration phase is nearly indistinguishable in such a plot.

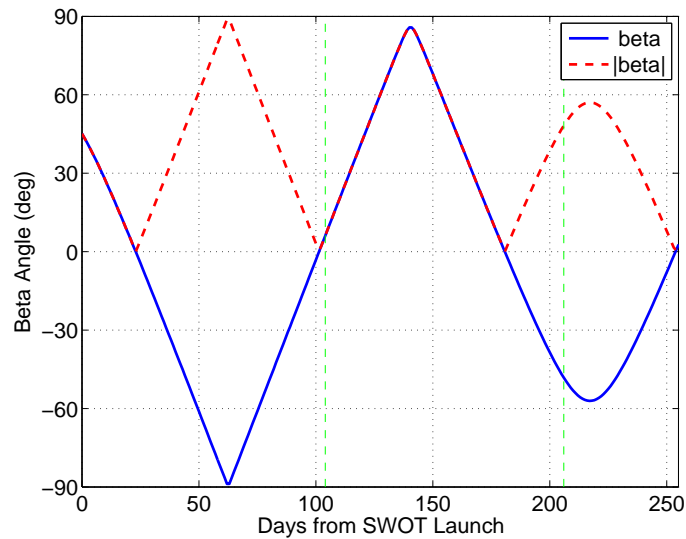


Figure 4.6: Example plot of the beta angle and its absolute value over time. The horizontal axis is the time in days relative to the SWOT launch on 2022-12-16. Dashed, green vertical lines indicate the start of the calibration phase on Day 104 (2023-03-30) and the start of the transition to the nominal orbit on Day 206 (2023-07-10).

Nadir module contains the NAlt, AMR, DORIS, and LRA instruments. The GPSP is inside the spacecraft bus since the payload module is typically facing away from the GPS constellation. In its deployed configuration, the satellite extends approximately 16 m along the KMSF x axis from one end to the other of the two solar panels, approximately 10 m along the y axis from one end to the other of the KaRIn baseline, and approximately 5 m along the z axis. The initial satellite mass at launch was approximately 2100 kg. The SAT_COM product provides a history of the satellite mass since launch.

4.4.1 Ka-Band Radar Interferometer (KaRIn)

KaRIn is a dual-sided, Ka-band, cross-track radar interferometer that was designed expressly to be the centerpiece of the SWOT mission. As is evident from Fig. 4.7, KaRIn is physically quite large. The KaRIn electronics are located within the KaRIn module, which, in the nominal flight orientation, is above the nadir module and below the spacecraft bus. Masts extend to either side of the central module in the KMSF $\pm y$ directions to form the 10 m interferometric baseline.

Refer to Sect. 9.1 for an overview of the KaRIn hardware design as it pertains to scientific interpretation of the data.

4.4.2 Poseidon-3C Nadir Altimeter (NAlt)

The Poseidon-3C nadir altimeter is derived from the Poseidon-3B altimeter onboard the Jason-3 mission, with an improved dynamic echo coding from 8 to 16 bits. It operates at two frequencies, 13.575 GHz (Ku band) and 5.3 GHz (C band), to obtain precise estimates of ionospheric delays and hence corrections. The Poseidon-3C electronics are divided into two (non-redundant) units: the processing unit (PCU) and the radio frequency unit (RFU), as shown in Figure 4.8. Together with the antenna, the three main parts of the instrument are thus:

- The processing unit, which includes the digital chirp generator, base-band demodulator, spectrum analyzer, instrument control unit, and interfaces.
- The RF unit, which includes the up-conversion and down-conversion to/from the Ku- and C-band carrier frequencies, high-power amplifier (solid state), low noise amplification of the received echoes, mixing with a reference chirp (deramp), and gain control of the receive chain.
- The Poseidon-3C dual frequency antenna, which is a center-fed reflector with a 1.2 m diameter, located on the nadir face of the satellite.

4.4.3 Advanced Microwave Radiometer

The SWOT AMR consists of two simultaneously active radiometers. The two radiometers are named the “+y” and “-y” strings and operate independently from each other. Each radiometer string has a single feed that illuminates a common, 1 m primary reflector to produce a beam to one side or the other of nadir. The name of each string references which side of nadir (+y or -y) the beam footprint is on relative to the y axis of the spacecraft body-fixed KMSF coordinate system. The two strings facilitate measurements on the left and right sides of the satellite




Figure 4.7: The SWOT Flight System in its deployed (flight) configuration (background image credit: NASA/JPL-Caltech).



Figure 4.8: Poseidon-3C processing unit (PCU) and radio frequency unit (RFU). Copyright CNES - Thales Alenia Space.

nadir point, approximately halfway across each of the two KaRIn half swaths (see Figs. 1 and 2 from [11]). When flying forward, measurements from the $-y$ and $+y$ strings are approximately 4.5° fore and aft of the satellite nadir point, respectively. As such, the left and right AMR measurements are approximately 4.5° (approximately 70 km) fore and aft of the KaRIn swath, respectively, whether the spacecraft is yaw flipped or not. Consequently, there is a time offset of approximately 11 s between corresponding AMR and KaRIn measurements of the same location on the ground (this offset is accounted for during ground processing). Each string has a linear polarization vector fixed in the instrument coordinate frame, with the two sides orthogonal to each other. The vectors are rotated approximately 45° relative to the vertical and horizontal polarization components at the Earth intercept.

Each of the two strings operate at three frequencies, 18.7, 23.8, and 34.0 GHz. These three frequencies are primarily aimed towards recovery of two-way atmospheric attenuation, water vapor, and cloud liquid burden, respectively. These measurements are geared mainly towards the over-ocean data from the NAlt and KaRIn instruments, as the radiometer measurements are subject to land contamination. They may be valid over large inland water bodies such as lakes, depending on their size, where land contamination effects may not be significant, however. The estimates of atmospheric attenuation are computed at the Ku- and C-band frequencies of the NAlt as well as at the Ka-band frequency of KaRIn.

Data from each string are processed independently, and the L2_RAD_GDR, L2_RAD_IGDR, and L2_RAD_OGDR products provide data from each string in each of two separate groups of data structures within each product granule. The AMR data also include the antenna and brightness temperature measurements from each string. The processing algorithms apply independent calibration coefficients for each string to account for the polarization and emission angle differences between the two beams. Data from these products are then interpolated in time and in space to the locations of the NAlt and KaRIn measurements when they are applied to the L2_NALT_GDR, L2_NALT_IGDR, L2_NALT_OGDR, and L2_LR_SSH products.

The SWOT radiometers are not intended to provide useful measurements over land and



Figure 4.9: DORIS receiver (left panel, Credit T-DMS) and antenna (right panel, Credit CNES).

therefore are not used for any of the KaRIn HR products. Furthermore, the radiometers are not intended to support climatology. However, as described in Sect. 4.5.1, the AMR calibrations take advantage of cold-sky observations of opportunity when there are satellite maneuvers and the spacecraft pitches 90° .

4.4.4 DORIS

The complete DORIS system includes the DORIS on-board package, a network of approximately 60 beacons equally distributed around the world, and a ground system. The on-board package includes an omnidirectional antenna located on the nadir face of the satellite and an electronics box (see Fig. 4.9) that includes the receiver itself and an ultrastable oscillator. The electronics box includes a cold-spare receiver/oscillator string for redundancy. The instrument is equipped with seven dual-frequency channels to allow simultaneous tracking of up to seven beacons. A real-time on-board orbit determination capability, called DIODE navigator, is also embedded into the DORIS instrument and is able to provide orbit solutions with RMS radial accuracy of a few centimeters. DIODE also provides essential navigation data to the KaRIn and Poseidon-3C instruments to ensure their proper operation. The DORIS tracking data are reported in the L1_DORIS_RINEX data products. The DORIS tracking data are used to support the precise orbit determination of the SWOT satellite.

4.4.5 Global Positioning System Payload (GPSP)

The GPSP receiver is a single-string, twelve-channel Global Positioning System TriG receiver [37]. It is connected to a RUAG (now Beyond Gravity) precise orbit determination chokering antenna (see Fig. 4.10). The GPSP antenna is positioned on the zenith (KMSF $-z$) side of the spacecraft at the center of the launcher interface ring.

The GPSP tracks the L1 (1575.42 MHz) and L2 (1227.6 MHz) signals from the GPS constellation of satellites. It can simultaneously track signals from up to twelve GPS satellites. The receiver




Figure 4.10: Global Positioning System Payload receiver and chokering antenna.




Figure 4.11: SWOT Laser Retroreflector Array (LRA).

is configured to track the L1C (C1), L1W (P1), and L2W (P2) GPS signals only. It nominally provides tracking data at a rate of once every 10 s. These data are reported in the L1_GPSP_RINEX data product. The GPSP tracking data are used to support the precise orbit determination of the SWOT satellite.

4.4.6 Laser Retroreflector Array (LRA)

The LRA onboard SWOT is a copy of those flown on the Jason-2, Jason-3, and Sentinel-6 missions. The LRA is located on the nadir face of the satellite. It has nine cubes on a spherical surface, as shown in Fig. 4.11. Each cube is manufactured from SUPRASIL-1 (quartz) and has a diameter of approximately 32 mm. Each has a refractive index of 1.46071 and a group refractive index of 1.4853 for a wavelength of $0.532 \mu\text{m}$. Nominally, no obscuration of the cubes can occur, but they can be obscured by multilayer insulation (MLI or thermal blanketing) at very high off-nadir angles. The back-face coating material is MgF₂.

Satellite laser ranging measurements from the SWOT mission are used to support precise orbit determination of the satellite and/or to validate the precise orbit determination obtained using other tracking systems (e.g., GNSS and DORIS). The [International Laser Ranging Service](#) (ILRS), a service of the International Association of Geodesy, coordinates activities for the international network of laser ranging field stations, with the network representing a global consortium of permanent and mobile field stations. The ILRS collects and distributes laser ranging measurements of the SWOT LRA made by the network.

4.5 Spacecraft Events

Actions taken by the spacecraft and/or conditions affecting the spacecraft may compromise the quality of the SWOT measurements. The times of spacecraft events are known, and instrument data are typically flagged when these events are expected to impact data quality. KaRIn is particularly sensitive to perturbations of its thermal and mechanical stability during these events and for some time period after the event during which the hardware may still be settling back to equilibrium. SWOT spacecraft events can be found at the PO.DAAC [Mission Events](#) website.

Note that it is useful to distinguish the cases that (1) SWOT granules are not available due to extended periods of missing data [see Sect. 5.5]; (2) SWOT granules are available, but some data in those granules are missing or are flagged; (3) for the L2 HR products, SWOT granules are available, but measurements are not reported in those granules because water was not detected. The L2_HR_PIXC product provides flags for distinguishing the latter two cases. The classes of spacecraft events below typically result in the second category above.

4.5.1 Propulsive Maneuvers

Satellite thruster firings (maneuvers) are necessary to maintain the spacecraft orbit within the predefined tolerance of the reference orbit (see Sect. 3.4). These maneuvers are known as station-keeping maneuvers or orbital-control maneuvers (OCMs). They typically occur once every several weeks. Collision-avoidance maneuvers are also planned and executed as needed in order to reduce the risk that the SWOT spacecraft collides with other known objects in space (usually space debris). These are typically less frequent than station keeping maneuvers. A station-keeping maneuver might typically take around 1 hour for the maneuver itself. Collision avoidance maneuvers may take several hours longer because they may involve moving the spacecraft away from the reference orbit and back.

The spacecraft thrusters are located on the satellite $-z$ face. During maneuvers, the thrusters usually need to be oriented along the velocity or anti-velocity directions, so the spacecraft is pitched by $\pm 90^\circ$. These large excursions of the spacecraft attitude from nominal impact data availability from all instruments. In addition, KaRIn data may be degraded for a few additional hours after the end of the maneuver. A benefit of these maneuvers is that the AMR points towards cold sky. These cold-sky measurements are used to assist with the routine calibration of the AMR.

Solar Array Orientation	β
0°	$ \beta < 6^\circ$
12°	$6^\circ < \beta < 25^\circ$
30°	$25^\circ < \beta $

Table 4.3: Approximate relationship between nominal solar array orientations and beta angle. A solar-array orientation of 0° corresponds to having the solar cells face the KMSF $-z$ direction. Positive rotation angles are right handed about the KMSF $-x$ axis, so the solar arrays would face between the $-z$ and $-y$ directions.

4.5.2 Gyro Calibration Maneuvers

Occasionally, the spacecraft is commanded to execute large, rapid attitude variations for the purposes of calibrating the gyros which provide attitude knowledge. These attitude variations point the spacecraft away from nominal and require some settling time afterwards once the nominal pointing is reestablished. Gyro calibration maneuvers may last around 10 hours. KaRIn data are possibly affected for a few additional hours after the end of the maneuvers.

4.5.3 Yaw Flips

As described in Sect. 3.1.8, the satellite attitude is rotated by 180° when the sun crosses the orbit plane ($\beta = 0^\circ$) so that the $-y$ side of the spacecraft generally faces away from the sun. These yaw flips are scheduled to occur when $|\beta| < +0.5^\circ$. They occur approximately once every 78 days, or about five times per year. The typical duration of each yaw flip is approximately 25 minutes. These large, rapid variations in the spacecraft attitude may require significant settling time for KaRIn. Additionally, they include commanded changes to the state of the KaRIn on-board processing so that the on-board processing handles the swath sides correctly. KaRIn data may be affected for a few additional hours after the yaw flips. No impact on other science data is expected.

4.5.4 Solar Array Rotations

Through the life of the mission, the solar arrays are rotated periodically based upon the beta angle to maximize their illumination by the sun. These rotations are performed so that the solar arrays are oriented at one of three discrete steps, as indicated in Table 4.3, with a $\pm 2^\circ$ tolerance on beta angle. The tolerance corresponds to approximately 2 days, which allows for flexibility in scheduling the rotation commands.

There are hence four rotations in each beta cycle of approximately 78 days (approximately 18 rotations per year). Typically, two rotations (30° to 12° followed by 12° to 0°) occur in the two weeks before each yaw flip, and two more rotations (0° to 12° followed by 12° to 30°) occur in the two weeks after the yaw flip. This strategy minimizes the perturbations to the spacecraft state and helps to maintain the stability of the instruments.

The duration of solar array rotations is typically short. KaRIn data may be affected for approximately 11 minutes after each rotation as the platform stabilizes. No impact on other science

data is expected.

The solar array rotations occur along an axis that is approximately parallel to the KMSF x axis (see Sect. 3.3.6). Because the solar array rotations are not intended to change the orientation of the entire spacecraft (but may perturb the spacecraft attitude control slightly), these have less of an impact on the instrument data than maneuvers or yaw flips, and less time is needed for settling after a solar array rotation.

The center of mass of the solar arrays is not aligned with their rotation axis, however. As a result, the orientation of the solar arrays affects the position of the satellite center of mass relative to the spacecraft body frame (e.g., KMSF). Pre-launch models are used to account for these changes to the satellite center of mass, and they are reflected in the SAT_COM product. The impact of the solar array angle on the satellite center of mass is illustrated in Fig. 6.1. These changes to the satellite center of mass are taken into account when computing the precise position of the satellite, as well as during the ground processing of the science data.

4.5.5 Eclipse Transitions

When the spacecraft is in the shadow of the Earth, it is said to be eclipsed by the Earth. Eclipse transitions occur when the spacecraft enters or exists the shadow of the Earth. These transitions cause sudden changes in the temperature of instrument hardware (thermal snaps) which can affect instrument stability. These changes are relatively small, however. Eclipse transitions may occur twice per orbit revolution or not at all depending on the beta angle. Each transition is nearly instantaneous; KaRIn data may be affected for around 2 minutes after each eclipse transition. The spacecraft may occasionally also go through the shadow of the Moon, but this occurs much less frequently.

When eclipse transitions occur, they may affect nearly the same band of latitudes on many successive passes. Therefore, KaRIn data affected by eclipse transitions may appear to have a distinctive pattern when viewed on a map (see Fig. 4.12). The actual impact of eclipse transitions on the LR and HR products may not be significant with respect to user needs, however.

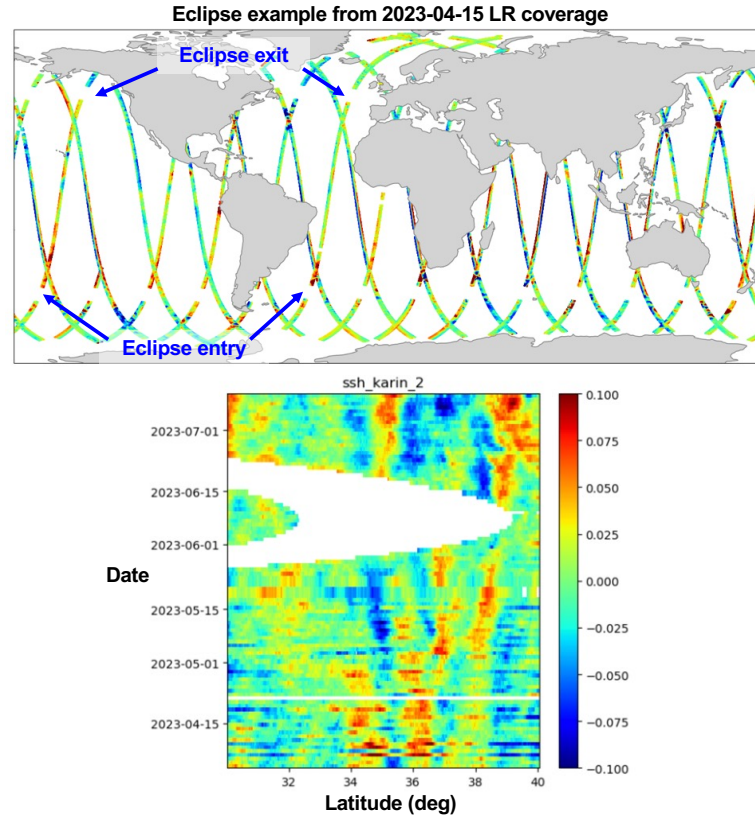


Figure 4.12: Examples of data affected by eclipse transitions. The upper panel shows a map of LR ocean coverage for one particular day (2023-04-15), where data affected by eclipse transitions have been removed. This results in the series of white gaps in the colored swaths (indicated by the blue arrows). The gaps occur at nearly identical latitudes for every pass of the day. The lower panel shows the evolution of the latitudes of the gaps (indicated in white) for one particular pass number over several consecutive cycles within the 1-day orbit. The latitudes affected by the eclipse vary with a time scale on the order of 1 month as the plane of the spacecraft orbit drifts. The colors represent KaRIn SSH with additional corrections applied. Other versions of processing may handle spacecraft events differently, as the measurements may still be quite usable despite the eclipse transitions; this figure is intended only for illustration of the spatial and temporal patterns that can be exhibited by eclipse transitions.

Chapter 5

Product Distribution and Fidelity

This chapter describes the operational generation and distribution of SWOT data products. It includes details on how products can be accessed by the public, the relationships between product latency and product fidelity, file naming conventions, and data availability.

5.1 Data Processing Centers

The processing of SWOT science data products is performed by a joint JPL-CNES science data system (SDS). The science data processing algorithms and auxiliary data are developed by a joint JPL-CNES algorithm development team. The same versions of the science data processing software are used regardless of whether a given product output is generated by a JPL or a CNES SDS processing center. JPL and CNES divide processing responsibilities such that any single product output is generated by only one processing center. This product output is then made available at both the NASA and CNES data distribution centers. As such, identical science data products are distributed by the NASA and CNES data distribution centers.

5.2 Data Distribution

Identical science data products are archived and distributed to users at both the NASA Physical Oceanography Distributed Active Archive Center (PO.DAAC) and the CNES AVISO and Hydroweb.next distribution centers.

5.2.1 NASA PO.DAAC

All of the SWOT datasets described in Tables 2.1, 2.2, and 2.3 are distributed by PO.DAAC and are available through the [NASA Earthdata Search](#) client.

PO.DAAC also hosts a variety of additional documentation (e.g., PDDs and ATBDs) and user resources. Please refer to the following:

- [SWOT Dataset Information](#) for information on data products, along with this user handbook, product description documents (PDDs), and algorithm theoretical basis documents (ATBDs).

- [Mission Events](#) for relevant information about spacecraft events that impact data quality and availability. This page also includes a history of changes to the KaRIn HR downlink mask (see Sect. 9.3).
- [User Resources](#) for resources available to users of SWOT datasets distributed by the PO.DAAC.
- [Data Announcements](#) for PO.DAAC SWOT data news and announcements.

PO.DAAC also provides the following additional resources for users of SWOT datasets.

- [PO.DAAC Cookbook - SWOT chapter](#)
- [PO.DAAC Data Subscriber/Downloader](#), and a [Video tutorial on using the podaac-data-subscriber](#)
- [Earthdata Search \(GUI\)](#)
- [Earthdata Search tutorial](#)
- [Earthaccess python library \(using CMR API on backend\)](#)

General information about Earthdata (AWS) cloud can be found at the following:

- [Obtain Earthdata login account](#)
- [Earthdata Cloud Primer documents](#)
- [Earthdata Common Metadata Repository \(CMR\) API](#)

5.2.2 CNES AVISO and Hydroweb.next

5.2.2.1 CNES AVISO

The SWOT ocean products from Table 2.1 are available at the CNES AVISO distribution center. The L1B_LR_INTF product is available by specific request only. Additional information on SWOT data access can be found via the following AVISO+ webpages:

- [Access to SWOT Ocean Products.](#)
- [SWOT KaRIn LR Ocean Products.](#)

The data can be accessed via the following options with AVISO+ credentials:

- FTP access: ftp://ftp-access.aviso.altimetry.fr:21/swot_products/
- SFTP access: sftp://ftp-access.aviso.altimetry.fr:22/swot_products/
- THREDDS Data Server (TDS) access: <https://tds.aviso.altimetry.fr/thredds/L2/L2-SWOT-DATA/L2-SWOT.html>

AVISO also distributes the L1_GPSP_RINEX, MOE, and POE products. The L1_DORIS_RINEX and SAT_COM products are available from the [International DORIS Service](#).

5.2.2.2 CNES Hydroweb.next

The SWOT KaRIn HR products from Table 2.2 can be accessed via the [CNES hydroweb.next portal](#). Hydroweb.next allows users to visualize and download datasets. It also provides centralized access to a collection of hydrology data products to complement the SWOT HR products.

The following instructions will guide you through your first visit. They are not an exhaustive list of possibilities. They can also be found in the help section.

Once on the [CNES hydroweb.next portal](#), simply type SWOT in the search bar. This will (i) trigger an autocompletion feature proposing some SWOT dataset or (ii) if you press enter, add a SWOT keyword filter and update the results pane accordingly. Control the filters and remove unwanted ones, if any. You may also add spatial or temporal filters.

The visualization of datasets is possible, for some datasets, once imported into your project. Some of the products are available for visualization on the map or directly as timeseries.

Downloading datasets is only possible once you are logged in. You must first create an account, or use your Theia account if you already have one. Once logged in, the download icons will be unlocked. You may either download your datasets individually from the Result pane, or download your entire Project. In either case, this will open the download pane and propose several options:

- download an archive, or a sample archive, in your browser (not recommended if the archive exceeds 10 GB)
- download a python script. All necessary instructions are in the script. You can also modify it later (dataset, region of interest, time restrictions). It requires you to provide an API-Key that you can create from the settings. Your API-Key is private and has no time limitation.

5.3 Fidelity and Latency of Products

Each SWOT product has an associated fidelity and latency. In general, longer-latency products have higher fidelity. Differences in fidelity arise from differences in the input data that are used to generate the product. These are summarized in the sections below. Sect. 5.4 describes how the fidelity of a product is indicated in the product file name.

5.3.1 SAT_COM Product

The SAT_COM data product is a historical file that is updated daily with new entries. Each daily update includes the full history of the satellite mass and center of mass since launch, as well as predicted values based upon planned upcoming spacecraft events. The file with the most recent `<CreationTime>` in the file name should be used.

5.3.2 L1_DORIS_RINEX and L1_GPSP_RINEX Products

The L1_DORIS_RINEX and L1_GPSP_RINEX products are available with a latency of less than 2 days. These products represent the tracking data measurements as provided in the telemetry with minimal, if any, post-processing beyond formatting those data into internationally rec-

ognized Receiver INdependent EXchange (RINEX) formats. Therefore, these products are not expected to be reprocessed and will typically only have one fidelity.

5.3.3 ATTD_RECONST Product

The ATTD_RECONST product provides spacecraft reconstructed attitude. A forward processed version is available with a latency of less than 1.5 days, and a reprocessed version with a latency of less than 15 days. The forward processed version is intended to support generation of the shorter latency L2_RAD_IGDR and KaRIn forward processed products. The reprocessed version is aimed to accommodate late arriving telemetry. When there are no telemetry delays the two versions are expected to be identical. In addition, reprocessing of the ATTD_RECONST products may occur when the processing algorithms are modified. The file with the highest fidelity, version, and product counter should be used.

5.3.4 MOE and POE Products

The MOE product is available with a latency of less than 1.5 days. Nominally, new versions of this product are not generated. The file with the most recent creation date in the file name should be used in the rare cases where the MOE product is reprocessed.

The POE product is available with a latency of less than 35 days. It provides a more accurate representation of the satellite position and velocity than the MOE product. Reprocessed versions of the POE product may be generated through the life of the mission using evolving models of the satellite forces and tracking data measurements, and possibly the reference frame. In particular, reprocessed versions of the POE product may be generated to support reprocessing of the KaRIn, NAlt, and AMR data products. The file with the most recent creation date in the file name should be used.

5.3.5 L2_RAD Products

The L2_RAD products provide AMR measurements and are generated with three different latencies, each with increasingly higher fidelity, as summarized in Table 5.1. Note, in particular, that the L2_RAD_OGDR products geolocate the AMR measurements to nadir as the ATTD_RECONST products are not available at the associated short latencies. Products with the longest latency (highest fidelity) should always take precedence over the shorter latency products. The L2_RAD_GDR products may be reprocessed through the life of the mission, while the L2_RAD_OGDR and L2_RAD_IGDR products will not be reprocessed.

5.3.6 L2_NALT Products

The NALT products provide nadir altimeter measurements and are generated with the three different latencies, each with increasingly higher fidelity, as summarized in Table 5.2. Products with the longest latency (highest fidelity) should always take precedence over the shorter latency products. The L2_NALT_GDR products may be reprocessed through the life of the mission, while the L2_NALT_OGDR and L2_NALT_IGDR products will not be reprocessed.

Auxiliary Data	Impacted Parameter(s)	L2_RAD Product (Latency)		
		OGDR (< 7 hours)	IGDR (< 1.5 days)	GDR (< 90 days)
Orbit Ephemeris	Latitude, Longitude	DORIS Navigator	MOE	POE
Attitude	Latitude, Longitude	Nadir	ATTD_RECONST	
Radiometer Calibration	Antenna Temperatures, Brightness Temperatures, Geophysical Parameters	Preliminary		Calibrated

Table 5.1: Differences between L2_RAD_OGDR, L2_RAD_IGDR, and L2_RAD_GDR products.

The L2_NALT_OGDR, L2_NALT_IGDR, and L2_NALT_GDR products are respectively generated using the L2_RAD_OGDR, L2_RAD_IGDR, and L2_RAD_GDR products as inputs, and are therefore similarly impacted by the adopted radiometer calibration.

In addition to the differences noted in Table 5.2, note also that the L2_NALT_OGDR and L2_NALT_IGDR products do not provide parameters from the Adaptive Retracker, while parameters from all other retrackers are available in all three NAlt products.

5.3.7 KaRIn LR and HR Products

Due to their high data volume and significant computational burden, the KaRIn LR and HR products have two types of fidelity.

- **Forward processing:** KaRIn LR and HR products are routinely generated from the data as they are downlinked with a typical latency of less than 5 days. These products are generated with the currently available processing software and auxiliary data. Due to evolution of the processing algorithms and auxiliary data, the latest forward-processed data may not be completely consistent with earlier forward-processed data. As such, forward-processed products may not be suitable for trending and other analyses that rely on consistency of processing.
- **Reprocessing:** The KaRIn processing algorithms are expected to evolve during the SWOT mission lifetime. As such, periodic reprocessing campaigns of the KaRIn data are planned. During these campaigns, the best available processing algorithms and instrument calibrations will be used to reprocess all data since the beginning of the SWOT calibration phase. These reprocessing campaigns produce products that are generated in a consistent manner. Best efforts will be made to transition the forward processing to the same major release version as the reprocessed products (e.g., they would have the same third character in the CRID field of the product file names as described in Sect. 5.4).

A given granule of SWOT downlink data may therefore be associated with multiple reprocessed versions as well as a forward-processed version. Version identifiers (see Sect. 5.4) are used to distinguish the outputs from the different rounds of processing. For most purposes, the product version that was most recently generated takes precedence over prior versions.

		NAlt Product (Latency)		
Auxiliary Data	Impacted Parameter(s)	OGDR (< 7 hours)	IGDR (< 1.5 days)	GDR (< 90 days)
Orbit Ephemeris	Latitude, Longitude, Altitude, Doppler Correction	DORIS Navigator	MOE	POE
Radiometer Calibration	Wet Troposphere Correction, Atmospheric Attenuation	Preliminary		Calibrated
Meteorological Model Fields	Dry and Wet Troposphere Corrections, Wind Speed, Surface Pressure, Inverse Barometer Correction	Predicted	Analyzed	
	MFWAM Waves	Not Available	Available	
	Sea Ice Concentration	Predicted		Analyzed
Pole Location	Pole Tide Height	Predicted		Analyzed
Dynamic Atmosphere Correction (DAC)	High Frequency Dealiasing Correction	Predicted	Preliminary	Analyzed
Global Ionosphere Maps (GIM)	Model Ionosphere Correction	Not Available	Available	

Table 5.2: Differences between L2_NALT_OGDR, L2_NALT_IGDR, and L2_NALT_GDR products. All changes impact the reported sea surface height.

Auxiliary Data	Impacted Parameter	KaRIn Product	
		Forward Processing	Reprocessing
Orbit Ephemeris	Latitude, Longitude, Altitude	MOE	POE
Radiometer Calibration	Wet Troposphere Correction, Atmospheric Attenuation	Preliminary	Calibrated
Pole Location	Pole Tide Height	Predicted	Analyzed
Dynamic Atmosphere Correction (DAC)	High Frequency Dealiasing Correction	Preliminary	Analyzed
Crossover Calibration	Crossover Calibration Correction	Computed From Earlier Data	Computed from Earlier and Later Data

Table 5.3: Differences between forward processed and reprocessed KaRIn LR and HR products. All changes impact sea surface height anomaly in LR products. The orbit ephemeris, pole location, and crossover calibration changes impact water surface elevations in HR products.

The primary differences between forward and reprocessed products are shown in Table 5.3. In forward processing, the crossover calibration is computed using recent earlier data only. Reprocessing computes the crossover calibration using a window of surrounding (earlier and later) data.

5.4 File Naming Conventions

The SWOT science data products summarized in Sect. 2.2 use one of the following four file naming conventions. See the PDD for a given product for more specific information on the file naming conventions of that product.

5.4.1 KaRIn, ATTD_RECONST, L1_GPSP_RINEX and L2_RAD Products

Most of the SWOT products follow the following file naming convention, with exceptions noted in Sect. 5.4.2, 5.4.3, and 5.4.4:

SWOT_⟨ShortName⟩[_⟨FileID⟩][_⟨GranuleID⟩_⟨BeginTime⟩_⟨EndTime⟩_⟨CRID⟩_⟨Counter⟩].⟨ext⟩

Above, the terms in angle brackets (“⟨⟩”) are placeholders for the fields below:

- ⟨ShortName⟩ is the identifier of the data product, including all of the KaRIn LR and HR, ATTD_RECONST, and L1_GPSP_RINEX data products. The L2_RAD_GDR, L2_RAD_IGDR, L2_RAD_OGDR products use a similar naming convention, but use “GPRAD”, “IPRAD”, and “OPRAD” for the ⟨ShortName⟩, respectively.

- $\langle \text{FileID} \rangle$ is the identifier of the file within the product for products that contain multiple files per granule. Examples include Basic, WindWave, Expert, and Unsmoothed for the L2_LR_SSH product, Reach and Node for the L2_HR_RiverSP product, and Obs, Prior, and Unassigned for the L2_HR_LakeSP product. This field is not related to the different files within a shapefile for products that are provided in shapefile format. As indicated by the square brackets (“[.]”) above, this field and the associated underscore character that separates it from the preceding field are omitted for products that contain only a single file. The L2_HR_Raster product includes several unique specifiers in the FileID field that capture product variants allowed in on-demand processing.
- $\langle \text{GranuleID} \rangle$ is an identifier of the granule that is represented by the file. This identifier usually includes some combination of the cycle number, the pass number, the tile number, the scene number, the continent identifier, and the basin identifier.
- $\langle \text{BeginTime} \rangle$ and $\langle \text{EndTime} \rangle$ are the UTC dates and times that the SWOT observations in the granule start and stop. These values are precise only to 1 s, but more precise values are available in the metadata of the product granule itself.
- $\langle \text{CRID} \rangle$ is the composite release identifier, which provides version information for what software, auxiliary data, and upstream products were used to produce the granule. For publicly released products, the CRID is usually four alphanumeric characters (e.g., “PIC0”, “PGC0”):
 - The first character for products from routine processing is always ‘P’ to denote that the file was generated in the operational production environment. A value of ‘D’ is used for the on-demand L2_HR_Raster product. A value other than ‘P’ and ‘D’ may be used for files generated during testing and development, but such files are generally not released publicly.
 - The second character is usually ‘G’, ‘I’ or ‘O’ to indicate the product fidelity (see Sect. 5.3). Specifically, ‘G’ indicates the highest fidelity and longest latency, ‘I’ indicates medium fidelity and shorter latency, and ‘O’ refers to the lowest fidelity and shortest latency. Currently, KaRIn LR and HR products use ‘I’ for forward-processed products and ‘G’ for reprocessed products. The L1_GPSP_RINEX product is currently only provided with ‘I’ fidelity. The L2_RAD_GDR, L2_RAD_IGDR, and L2_RAD_OGDR products are respectively the ‘G’, ‘I’, and ‘O’ fidelity of AMR products.
 - The third character indicates the product major release version. The value is an uppercase letter that increments with each major product release. Major product releases occur when there are changes to the processing software and/or auxiliary data that have significant impact on the content of data (e.g., important changes in algorithms, calibration parameters, etc.). The first routine public release of KaRIn data is major version C. A prior beta product release was major version B.
 - The fourth character indicates the product minor release version. The value is alphanumeric. The minor version starts from 0 (zero), increments through 9, then moves to A, B, C, and so on. Minor product releases occur when there are changes to the processing software and/or auxiliary data that have only minor impact on

the content of the data (e.g., bug fixes that allow edge-case granules to be produced where production would have failed with a previous version of software). The minor version resets to 0 when the major version is incremented.

- $\langle \text{Counter} \rangle$ is the product counter, which distinguishes multiple versions of the same granule produced within the same processing environment. The value is a two-digit numeric value that is zero padded on the left. The product counter begins with 01 and increments with each new instance of a granule that is generated with the same CRID. The product counter is reset to 01 when the CRID changes. For example, the product counter for a particular product granule increments when a processing problem is encountered, and the instance needs to be rerun. In this case the first attempt will have its product counter set to 01, and the second attempt will have its product counter set to 02. Product granules with the highest product counter should always be used.
- $\langle \text{ext} \rangle$ is the file name extension (e.g., .nc, .shp).

Note that the fields above may include underscore characters within them.

5.4.2 NAlt Products

The SWOT L2_NALT_GDR, L2_NALT_IGDR, L2_NALT_OGDR products adopt a file naming convention that aligns with the heritage of similar products generated for other altimetry missions:

SWOT_ $\langle \text{O/I/G} \rangle$ P $\langle \text{N/R/S} \rangle$ _2P $\langle \text{v} \rangle$ $\langle \text{S/P} \rangle$ $\langle \text{ccc} \rangle$ _ $\langle \text{ppp} \rangle$ _ $\langle \text{BeginTime} \rangle$ _ $\langle \text{EndTime} \rangle$.nc

Above, the terms in angle brackets (“ $\langle \cdot \rangle$ ”) are placeholders for the fields below:

- $\langle \text{O/I/G} \rangle$ is a single character representing the product family with ‘O’, ‘I’, and ‘G’ for the L2_NALT_OGDR, L2_NALT_IGDR, and L2_NALT_GDR, respectively.
- $\langle \text{N/R/S} \rangle$ is a single character representing the product type with ‘N’, ‘R’, and ‘S’ for the “native”, “reduced”, and “sensor” product types.
- $\langle \text{v} \rangle$ is a single character representing the product version. The first release of these products were version f.
- $\langle \text{S/P} \rangle$ is a single character representing the granule duration, where ‘S’ represents “segment” for L2_NALT_OGDR products, and ‘P’ represents “pass” for L2_NALT_IGDR and L2_NALT_GDR products. Here, “segment” means that the granules of the L2_NALT_OGDR products spans the full duration of a granule of telemetry data.
- $\langle \text{ccc} \rangle$ is a three-digit numeric value that is zero padded to the left and represents the cycle number.
- $\langle \text{ppp} \rangle$ is a three-digit numeric value that is zero padded to the left and represents the pass number.
- $\langle \text{BeginTime} \rangle$ and $\langle \text{EndTime} \rangle$ are the UTC dates and times at which the SWOT observations in the granule start and stop. These values are precise only to 1 s, but more precise values are available in the metadata of the product granule itself.

5.4.3 POE, MOE, and SAT_COM Products

The SWOT POE, MOE, and SAT_COM products follow the file naming convention below, in part due to the heritage from prior altimetry missions:

SWOT_⟨ShortName⟩_⟨CreationTime⟩_⟨BeginTime⟩_⟨EndTime⟩.⟨ext⟩

⟨CreationTime⟩ is the UTC date and time that the specific file was generated. ⟨BeginTime⟩ and ⟨EndTime⟩ are the UTC dates and times that the SWOT observations in the granule start and stop. These values are precise only to 1 s, but more precise values are available in the metadata of the product granule itself. The file with the latest ⟨CreationTime⟩ should be used when multiple files have the same ⟨BeginTime⟩ and ⟨EndTime⟩.

⟨ext⟩ is the file name extension, which is currently “.nc” for this set of products.

5.4.4 L1_DORIS_RINEX Product

The L1_DORIS_RINEX products use the following file naming convention:

SWOT_MEP_1P⟨v⟩D⟨SensingDate⟩_000000

⟨v⟩ is the product version format, ranging from ‘a’ to ‘z’. ⟨SensingDate⟩ is the UTC date of the DORIS data in the file.

5.5 Known Issues Affecting Data Availability

Experience with in-flight SWOT operations has revealed the following known issues that can impact data availability over significant time spans. These may cause gaps in the availability of product granules. That is, rather than seeing limited durations over which measurements are missing or are flagged within a product granule, users may find that entire files appear to be missing over the time spans affected by the items below.

- Onboard data storage: Single event upsets (SEUs) with the onboard solid state recorder (SSR) have been experienced since launch. These issues can affect data availability early in the mission (before June 2023). However, mitigation steps have been applied since then and the data outages are now expected to have short durations (less than a few hours). The SSR SEUs impact data from all instruments. Long data outages due to SSR SEUs are noted at the PO.DAAC [Mission Events](#) website.
- Downlink to ground stations: Issues with downlink of data have also been experienced since launch. In most cases, data are recovered during the following pass for some instruments, but the data may arrive late to the ground processing centers. When data are recovered, they may be missing from the forward-processed data, but they will be available in reprocessed data sets. The shorter latency of the forward-processed data makes them less tolerant to delays in receiving data.
- Forward processing: Occasionally, there may be gaps or increased latency in forward-processed KaRIn data due to problems with the processing center itself. In such cases,

reprocessed data should include granules that may have been missing from forward processing.

- **Wildfires at Inuvik Ground Station:** A particularly noteworthy event is the wildfires near the Inuvik (IVK) downlink station in Canada that prevented recovery of data from this station from August 11 to September 15, 2023. The KaRIn HR mask (see Sect. 9.3) was reduced to around half of its nominal coverage during much of this period in order to fit into the downlink capacity of the other ground stations (the reduced mask version was v6c for most of the outage). Around 10 days near the beginning of the outage had little HR coverage at all. As a result, there is significantly lower than normal coverage of KaRIn HR products during this period.
- **Damaged undersea cables affecting Hartebeesthoek Ground Station:** Damage to undersea or submarine data cables prevented the recovery of data from the Hartebeesthoek (HBX or HBK) Ground Station in South Africa from March 15 to April 30, 2024. In response, the KaRIn HR mask (see Sect. 9.3) was updated to remove approximately 25% of its nominal coverage. HR coverage during the first few days of the outage was even less. Two versions of the reduced HR mask (with around 75% of nominal coverage) were used, with the first (v7) primarily removing northern-latitude areas that would likely still be frozen through mid April, and the second (v7b) modifying the coverage to capture the thaw for those areas.
- **Damaged antenna at the Hartebeesthoek Ground Station:** Damage to the antenna prevented the recovery of data from the Hartebeesthoek (HBX or HBK) Ground Station in South Africa from August 20 to September 9, 2024. In response, the KaRIn HR mask (see Sect. 9.3) was updated to remove approximately 25% of its nominal coverage using the v7b mask from August 20 to September 4, 2024. Between September 4 and September 9, 2024, the HR mask was adapted to removed only about 10% of the nominal coverage using a new v08 mask. Data collection resumed with the nominal (100%) downlink mask on September 9, 2024.

A complete history of changes to the HR downlink mask is available at the [Mission Events](#) website. See Sect. 4.5 for a description of spacecraft events that may cause quality flags to be set in the data products.

Chapter 6

Standard Science Data Products

This chapter provides information on SWOT standard science data products that are available to users. The information here is meant to familiarize users with the types of products that are available so that users can decide which SWOT products are most suitable to their needs and seek additional details in the PDDs for those products as necessary.

6.1 L1_DORIS_RINEX

The L1_DORIS_RINEX product [30] provides tracking data from the DORIS instrument after minimal ground processing to unpack its telemetry. See Sect. 4.4.4 for a description of the DORIS instrument.

The L1_DORIS_RINEX product provides tracking data at the rate they are available in the telemetry. The DORIS receiver is nominally configured to provide tracking data at a rate of 0.1 Hz (i.e., every 10 s). Each granule of the product spans one day and is generated after merging DORIS data reported in all telemetry files from that day and removing any duplicate data. The typical size of each granule is approximately 5 MB.

6.2 L1_GPSP_RINEX

The L1_GPSP_RINEX product [31] provides tracking data from the GPSP instrument after minimal ground processing to unpack its telemetry. See Sect. 4.4.5 for a description of the GPSP instrument. The product provides time-tagged pseudorange and carrier-phase tracking data from the L1C, L1W, and L2W GPS signals, as well as the SNRs associated with these data. More specifically, the pseudorange data include the C1C, C1W, and C2W (historically referred to as the C1, P1, and P2) measurements, and the carrier-phase data include the L1C, L1W, and L2W measurements.

The L1_GPSP_RINEX product provides tracking data at the rate they are available in the telemetry. The GPSP receiver is nominally configured to provide tracking data at a rate of 0.1 Hz (i.e., every 10 s). Each granule of the product contains data spanning the coverage of data in each GPSP telemetry file. The typical size of each granule is approximately 10 MB, assuming each file spans approximately 2 hours.

6.3 SAT_COM

The SAT_COM product [29] provides a complete history of the spacecraft mass and center of mass since just before the start of the commissioning phase. The coordinates of the center of mass are in the body-fixed KMSF (see Sect. 3.3.6) coordinate system. These values will evolve due to spacecraft maneuvers and changes in the orientation of the solar panels (see Sect. 4.5.4). The product includes a flag to indicate the type of spacecraft event that is associated with each entry and whether the entry is based upon an upcoming planned event.

The product is updated daily with a new entry added to the product only when there is a new spacecraft event that impacts the mass or center of mass. Values in the product with dates earlier than the most recent update are based upon knowledge of when maneuvers and solar array rotations were performed and of how much propellant was consumed by each maneuver. Values in the product with dates later than the most recent update are based upon planned upcoming maneuvers and solar array rotations. This means that predicted values can be replaced after the spacecraft events have occurred. Users should always use the most recent update to the product (file with most recent creation time) and exercise caution when using predicted values.

Pre-launch models are used to compute the effects on the spacecraft mass and center of mass from satellite maneuvers and solar array rotations. As shown in Fig. 6.1, changes in the solar array orientation primarily impact the y and z components of the satellite center of mass. The satellite center of mass moves along the y and z axes by approximately -12.7 mm and +1.3 mm, respectively, as the array orientation changes from 0° to 12° , and it moves an additional -17.8 mm and +6.8 mm as the array orientation changes from 12° to 30° . In contrast, the impact of routine maneuvers on the mass and center of mass is relatively small. The exception is the series of maneuvers that were required to transition from the 1-day to 21-day repeat orbit in July 2023, during which approximately 16 kg of propellant was used, and the z component of the center of mass shifted by +8 mm.

6.4 ATTD_RECONST

The ATTD_RECONST product [28] provides time series of the spacecraft attitude, defined as the orientation of the body-fixed frame (KMSF; see Sect. 3.3.6) with respect to the inertial frame (GCRF; see Sect. 3.3.1). The product represents the spacecraft attitude using quaternions and provides a quality flag for the quaternions. The ATTD_RECONST PDD [28] provides details on the use of the reported quaternions to represent the rotation matrix between the KMSF and GCRF.

The product provides reconstructed attitude data that are computed by blending data from the onboard gyro and star trackers. The use of gyro data together with star tracker data enables significantly better knowledge of high frequency attitude variations than are available from star tracker data alone. For example, the reconstructed attitude is available every 15.625 ms, while the star tracker data are available approximately every 125 ms. The improved attitude knowledge is required for the most accurate processing of the KaRIn measurements and is necessary to meet the KaRIn measurement requirements. The ATTD_RECONST data are also used to process the radiometer measurements as they are off-nadir, although the star tracker data alone would provide sufficient accuracy for radiometer processing.

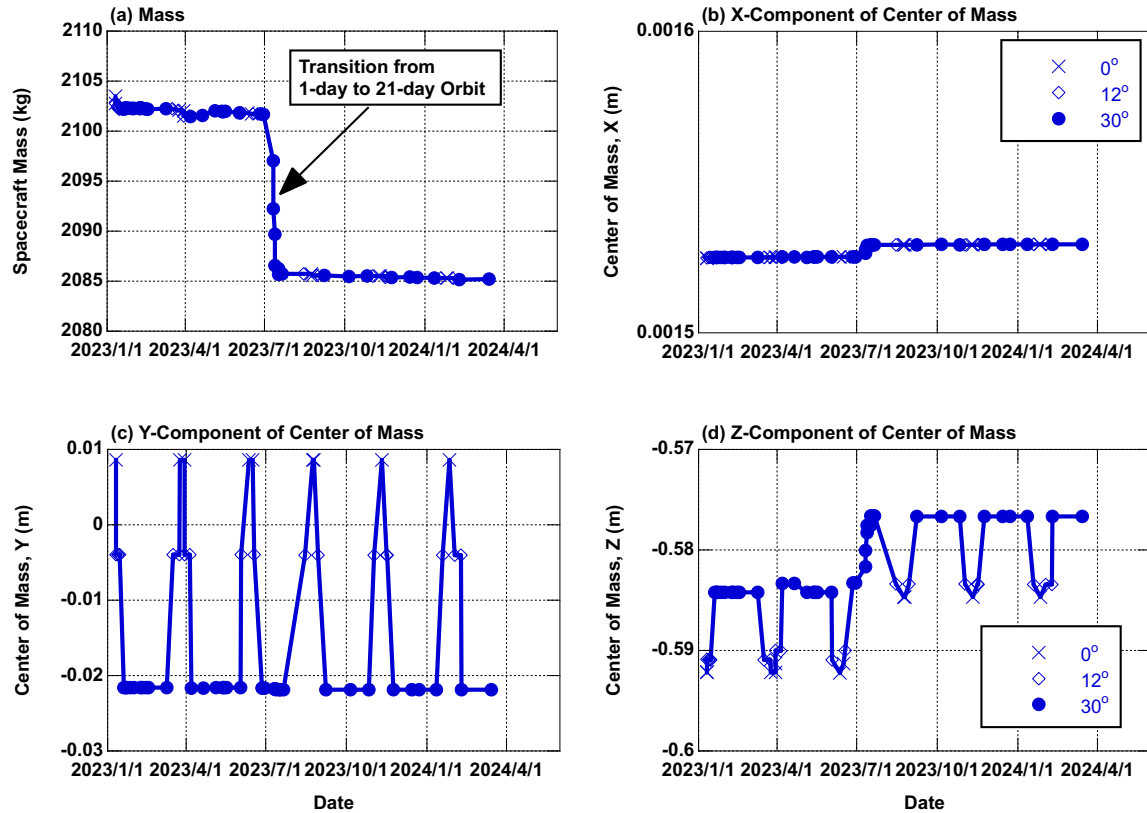


Figure 6.1: SAT_COM data from the first 15 months of the SWOT mission. The product reports the spacecraft mass and the x , y , and z coordinates of the spacecraft center of mass, as shown in panels (a), (b), (c), and (d), respectively. The coordinates are in the body-fixed KaRIn Metering Structure Frame (KMSF). The impact of the solar panel rotations is highlighted with the use of different symbols for each solar panel orientation angle. Solar panel orientation angles of 0° , 12° , and 30° are depicted with crosses, diamonds, and filled circles, respectively.

Each granule of the ATTD_RECONST product is a daily file spanning 26 hours centered at 12:00:00 (TAI). It has a typical volume of 0.3 GB.

6.5 POE and MOE

The POE and MOE products [27] provide time series of high-accuracy estimates of the position and velocity vectors of the SWOT spacecraft center of mass, namely the orbit ephemeris, in the ECEF ITRF (see Sect. 3.3.2). The products also provide a quality flag to indicate off-nominal cases for which the quality of the orbit ephemeris might be degraded due to maneuvers or missing tracking data. The ITRF version used to generate the orbit ephemeris is indicated in the product metadata. The POE and MOE products effectively define the ITRF version for the SWOT surface height measurements. The orbit ephemeris is provided at relatively coarse intervals (e.g., 10 s) as compared to the NAlt and KaRIn measurements. High-order (e.g., > 7) polynomial interpolation of the provided orbit ephemeris is necessary to achieve the best accuracy.

The POE and MOE products differ primarily in the latency at which they are available and consequently in the accuracy of the orbit ephemeris estimates. The MOE product is generated with a latency of less than 1.5 days, while the POE product is generated with a latency of less than 28 days. Reprocessed versions of the POE product may be generated through the life of the SWOT mission, while the MOE product is typically not reprocessed. As noted in Tables 5.1, 5.2 and 5.3, the MOE product is typically used to generate the shorter latency L2_RAD_IGDR, L2_NALT_IGDR, and KaRIn forward-processed products, while the POE product is used to generate the L2_RAD_GDR, L2_NALT_GDR, and KaRIn reprocessed products.

The POE product is expected to have better accuracy than the MOE product. The longer latency of the POE product allows for more accurate ancillary data to be used in the precise orbit determination computations as well as for more robust orbit determination approaches. Detailed descriptions of the orbit determination processing standards are available from the International Earth Rotation Service [38] and the International DORIS Service (IDS) [39]. These include the satellite force models, measurement models, and spacecraft attitude. Detailed descriptions of orbit determination algorithms are extensively available in the published literature (e.g., [40]).

Each POE and MOE granule is a daily file spanning 26 hours centered at 12:00:00 (TAI). It has a typical volume of 0.6 MB.

6.6 L2_RAD_GDR, L2_RAD_IGDR and L2_RAD_OGDR

The L2_RAD_GDR, L2_RAD_IGDR and L2_RAD_OGDR radiometer products [11] provide estimates, from each of the two AMR strings (see Sect. 4.4.3), of the wet troposphere correction, water vapor content, cloud liquid water content, wind speed, and atmospheric attenuation correction to backscatter (σ°) at the KaRIn Ka-band and NAlt Ku- and C-band frequencies. The algorithms for the geophysical parameter estimates (see [12]) are aimed towards open-ocean measurements and are subject to land, rain, and ice contamination. Various quality flags as well as surface type, rain, and ice flags are provided to indicate the quality of the geophysical

estimates. Also provided are the primary radiometer antenna temperature and brightness temperature measurements at each of the three AMR frequencies, 18.7, 23.8, and 34.0 GHz.

The L2_RAD_OGDR, L2_RAD_IGDR, and L2_RAD_GDR products respectively represent products with increasing latency and improved accuracy. The primary differences between these products are described in Sect. 5.3.5. The most important difference is that the L2_RAD_GDR product provides calibrated AMR measurements, while the L2_RAD_OGDR and L2_RAD_IGDR products use a historical calibration from within the last 90 days. In addition, the L2_RAD_OGDR measurements do not account for spacecraft attitude and are geolocated to nadir, while the L2_RAD_IGDR and L2_RAD_GDR products do account for attitude. The L2_RAD_OGDR product is available in telemetry-sized granules. A granule is generated for every telemetry segment as it becomes available. This typically spans 90 minutes of data but can vary depending on the duration of telemetry downlinked from the instrument. Meanwhile, the L2_RAD_IGDR and L2_RAD_GDR products are available in pass-sized granules, spanning approximately 50 minutes of data.

As noted in Sect. 4.4.3, the two AMR strings are pointed off nadir and therefore measure along non-vertical paths through the atmosphere. However, the geophysical parameters in the product correspond to the equivalent vertical atmospheric column directly above the reported geolocation. As such, they should be scaled by an appropriate obliquity factor for non-vertical signal paths [e.g., Eq. (11.1)]. While the algorithms are aimed towards open-ocean measurements, the products provide measurements globally. Useful geophysical estimates may be available over large inland water bodies, depending on their size.

Figure 6.2 shows an example of the 18.7 GHz brightness temperatures from the two AMR strings for a portion of the pass shown in Fig. 6.3. Panel (a) illustrates that the left and right AMR measurements are west and east of the nadir measurement for ascending passes. Panel (b) shows that the left and right AMR measurements are fore and aft of nadir in the along-track direction at any given time. Brightness temperatures over land are significantly higher than those over the ocean. Panels (c) and (d) then show that the left AMR footprint partially crosses the tip of the California coast and goes briefly back over the ocean before coming completely inland soon thereafter. Meanwhile, the right AMR measurements cross completely inland earlier and further south than those on the left. These figures also illustrate that the left and right AMR measurements are interpolated in space when they are applied to the NAlt measurements. For example, the interpolated nadir value lies between the left and right values when shown as a function of latitude [panel (c)], but not when shown as a function of time [panel (d)]. Spatial interpolation is also performed when the AMR measurements are applied to the KaRIn measurements. Spatial variability in the troposphere is expected to be significantly higher than any temporal variations during the approximately 22 s between the left and right AMR measurements in the along-track direction.

6.7 L2_NALT_GDR, L2_NALT_IGDR and L2_NALT_OGDR

The NAlt products [10] provide measurements from the NAlt in a form that is intended for oceanography users. The products provide NAlt measurements of sea surface height anomaly (SSHA) and NAlt estimates of range, backscatter (σ°), significant wave height (SWH), and wind

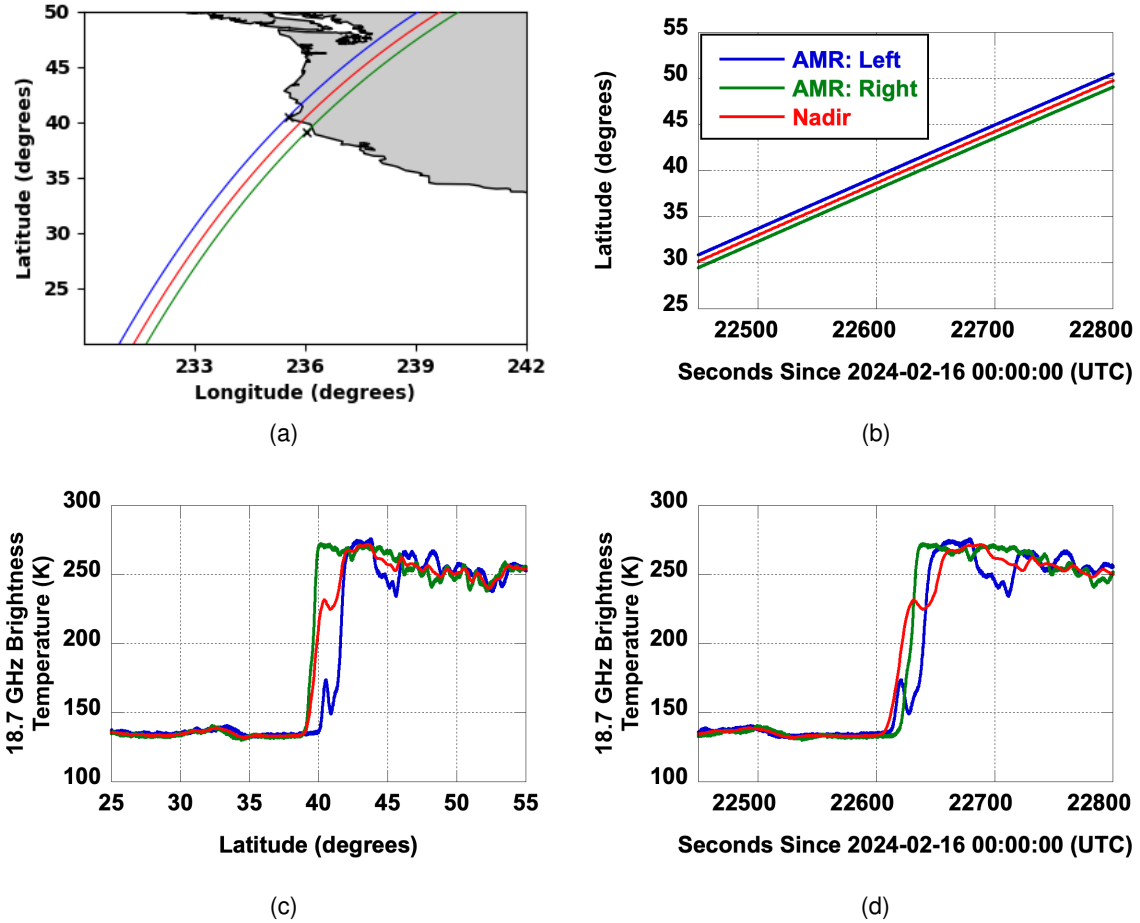


Figure 6.2: Example of AMR measurements for cycle 011, pass 039 (as shown in Fig. 6.3) on approach toward and over the California coast. Blue, green, and red lines represent measurements from the left AMR side, the right AMR side, and after spatial interpolation to nadir. Panels (a) and (b) show the location and time of the AMR measurements relative to nadir. Panels (c) and (d) show the 18.7 GHz brightness temperatures as functions of time and latitude. The two black crosses near the coast in (a) provide a time reference in that they represent left and right measurements occurring at the same time.

speed. The NAlt measurement quantities include quality flags. The products also include geophysical parameters from external sources that are used to compute corrections for the NAlt estimates or that may be useful for interpretation of the NAlt measurements. Such geophysical parameters include information on media delays, tides, and surface gravity waves.

The L2_NALT_OGDR, L2_NALT_IGDR, and L2_NALT_GDR products respectively represent products with increasing latency and improved accuracy. The primary differences between these products are described in Sect. 5.3.6. The most important differences are that the estimates of spacecraft altitude improve with longer latency and that the L2_NALT_GDR product uses calibrated radiometer measurements. In addition, the L2_NALT_OGDR product uses forecast meteorological fields, while the L2_NALT_IGDR and L2_NALT_GDR products use analyzed fields. The L2_NALT_OGDR product is available in telemetry-sized granules. A granule is generated for every telemetry segment as it becomes available. This typically spans 90 minutes of data but can vary depending on the duration of telemetry downlinked from the instrument. Meanwhile, the L2_NALT_IGDR and L2_NALT_GDR products are available in pass-sized granules, spanning approximately 50 minutes of data.

There are three separate files in each of the L2_NALT_IGDR and L2_NALT_GDR products, each with increasing size as additional information is provided. Users may use any file individually as fits their needs. The three files are:

- **Reduced:** The Reduced file provides the most pertinent information for users who are interested primarily in 1 Hz measurements of SSHA, σ° , SWH, and wind speed and who plan to take the measurements at face value. Each Reduced file is very small, at around 0.35 MB per granule.
- **Native:** The Native file provides all of the 1 Hz data in the Reduced file as well as 20 Hz NAlt measurements and more detailed information on instrument and environmental corrections. The additional detail is intended to enable advanced analysis of the data. Each Native file is therefore larger than a corresponding Reduced file at around 7 MB per granule.
- **Sensor:** The Sensor file provides all of the data in the Native file along with the full radar-echo waveforms. The availability of waveforms is intended to allow users to apply their own retracking algorithms to the altimeter waveforms. Each Sensor file is therefore larger than a corresponding Native file at around 43 MB per granule.

The L2_NALT_OGDR product is only available in the Reduced and Native forms.

While the NAlt products are designed primarily for oceanography, they are available globally, including over land. Each 1 Hz frame of the Ku- and C-band range measurements is derived from the linear regression of the respective valid 20 Hz range measurements. An iterative outlier detection scheme is adopted in this linear regression, and the 20 Hz measurements used in the regression are identified in the product. The number of 20 Hz measurements that are used to derive each of the 1 Hz measurements is also provided in the products, as is the root-mean-square of the differences between those 20 Hz measurements and the derived 1 Hz measurement. For specialized applications, such as over land, ice, lakes, and/or rivers, users may choose to perform their own compression algorithm on the 20 Hz measurements.

6.8 L1B_LR_INTF

As described in Sect. 9.2, data from the KaRIn instrument are split into separate low-rate and high-rate data streams within the instrument before downlink from the spacecraft. The low-rate stream is geared primarily toward oceanography. See Sect. 9.2.1 for details on the KaRIn LR on-board processing. The L1B_LR_INTF product [8] provides all relevant information from the LR data stream after ground processing to unpack the KaRIn telemetry, apply initial calibration parameters and geophysical corrections (e.g., for media delays and attenuation), and compensate for interferometric phase biases that arise due to the assumptions of and the limited real-time orbit and attitude knowledge available to the on-board processing.

Most notably, the L1B_LR_INTF data product provides the bias-corrected, complex interferogram measurements for each of the nine Doppler beams formed on-board for each of the left and right sides of the swath. The L1B_LR_INTF product also contains the echo power, noise power, and backscatter (σ°) for each of the nine beams on each side. The product includes an estimate of the volumetric correlation (see Sect. 8.4.3) for each beam on each side as well.

The L1B_LR_INTF product also contains additional information that complements the KaRIn LR data: spacecraft ephemeris and attitude; geophysical corrections such as media delays that were used during L1B processing; reference locations that were assumed for L1B processing; and flags that provide information on data quality.

The intent of the L1B_LR_INTF product is to provide LR information that abstracts some of the engineering details of the implementation of the KaRIn instrument but that remains close to the fundamental measurement quantities of the instrument—namely echo power and interferometric phase. The information in the L1B_LR_INTF product can therefore be processed by expert users who wish to apply their own specialized algorithms for height reconstruction, beam combining, spatial smoothing, and/or resampling in order to meet their individual objectives.

The L1B_LR_INTF product is available in pass-sized granules that cover both swath sides [32] at the native sampling and resolution of the KaRIn instrument (approximately 250 m and 500 m, respectively, for each beam). Because the sampling of the product follows the native sampling of the KaRIn instrument, it is not sampled completely regularly in either space or time. A 3-D reference location is provided for each sample, but this location is based on a reference surface (i.e., a reference DEM; see Sect. 8.3.4), not the 3-D measurement from KaRIn itself. Therefore, the L1B_LR_INTF product does *not* contain the measured height of the Earth surface. Rather, the L1B_LR_INTF product contains all the relevant information so that users can compute their own estimates of the height of the Earth surface from the underlying KaRIn interferograms. With the full-resolution data included for each of the nine beams, granules of the L1B_LR_INTF product are typically around 20 GB each. This is considerably larger than corresponding granules of the downstream L2_LR_SSH product.

Several examples of KaRIn data are shown in this and the following sections. The examples are taken from cycle 011, pass 039 (i.e., 011_039), acquired on 2024-02-16. The data were processed through the operational forward processing system (see Ch. 7) with CRID PIC0.

Figure 6.3 depicts the spatial coverage of the L1B_LR_INTF and L2_LR_SSH products for this example pass. As described in Ch. 4, 011_039 corresponds to an ascending pass (the pass number 039 is odd) from the science (nominal or 21-day) orbit (the cycle number is less than 400). Being an ascending pass, the spacecraft crosses the equator going from south to north. The left and

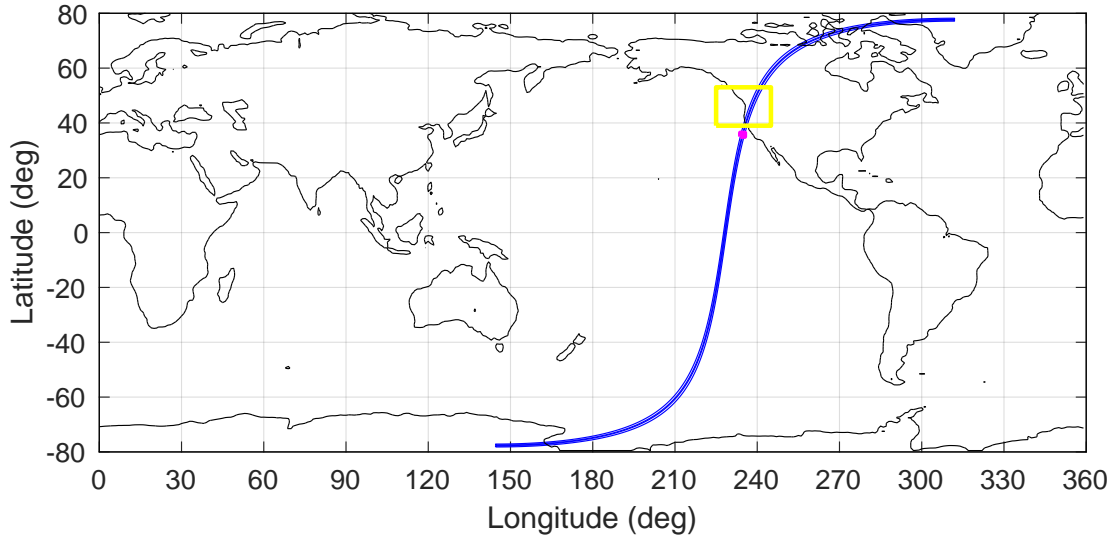


Figure 6.3: Pass coverage of the example KaRIn LR data shown in this chapter. The blue lines indicate the left and right sides of the KaRIn swath for the entire pass of the example LR data. The two sides of the KaRIn swath are shown individually in the figure, but they may be difficult to distinguish visually at the global scale of the map. The small magenta box in the Pacific Ocean off the California coast indicates the region where example L1B_LR_INTF and L2_LR_SSH data are shown in Figs. 6.4 and 6.5. The yellow box over the western edge of North America indicates the region where example L2_HR_RiverSP and L2_HR_LakeSP data are shown in Figs. 6.8(a) and 6.9(a).

right sides of the swath are hence to the west and to the east, respectively, of the nadir track at the equator.

Figure 6.4 shows example data from the L1B_LR_INTF product. The images in this figure cover only a small spatial extent of the pass-sized 011_039 granule. The area covered is indicated by the magenta box in Fig. 6.3. Data for only the central KaRIn OBP Doppler beam on either swath side are shown, although the product contains data for all beams. The phase of the bias-corrected interferogram provides the measurement of the surface height above the reference surface used for interferogram flattening during ground processing (see Sect. 8.3.4). The phase wraps modulo 2π rad. KaRIn measures the interferometric phase very precisely, so subtle variations in phase that may not be discernible visually in the graphic allow small changes in the surface height to be observed in the L2_LR_SSH product. Note that the geolocation shown is based on the reference surface used for processing, not the measured KaRIn heights. This figure shows only data that are flagged as *good*. No data are reported in the product in the white background areas.




Figure 6.4: Example LIB_LR_INTF data from the central Doppler beam: (a) radar backscatter or σ^0 [see Sect. 8.1.4]; (b) normalized complex interferogram after correcting for systematic biases during ground processing. In (b), the grayscale indicates the magnitude of the measured total coherence γ_{tot} (black and white represent 0 and 1, respectively; see Sect. 8.3.2), and the overlaid color hue indicates the wrapped interferometric phase. The graphics do not necessarily capture the full resolution and dynamic range of the data in the product.

6.9 L2_LR_SSH

The L2_LR_SSH product [6] provides KaRIn measurements from the LR data stream in a form that is intended to be as accessible as possible to oceanography users. The product contains KaRIn measurements of sea surface height (SSH) and sea surface height anomaly (SSHA) as well as KaRIn estimates of backscatter (σ^0), significant wave height (SWH), and wind speed. The KaRIn measurement quantities include quality flags and uncertainty estimates. The product also includes geophysical parameters from external sources that are used to compute corrections for the KaRIn estimates or that may be useful for interpretation of the KaRIn measurements. Such geophysical parameters include information on tides, media delays, and surface gravity waves.

The reported SSH and SSHA typically contain large, systematic errors that manifest as nearly linear cross-track trends with very little along-track variation. These systematic errors often overwhelm the spatial variations of the SSHA that are of interest. Because of this, the L2_LR_SSH product includes a crossover correction (see Sect. 9.4.2) that can be used to remove most of the cross-track systematic error. It is important to note, however, that the crossover correction is merely reported separately in the product—it is *not* applied to the SSH or the SSHA that are reported in the product. Therefore, users who read the SSH or SSHA directly from the product can expect large (meter-scale) trends in the cross-track direction that vary slowly in the along track direction. Users may apply the crossover corrections that are reported in the product files to the SSH and/or SSHA in order to remove these trends, but the addition of the correction value to the uncorrected SSH or SSHA to obtain the corrected SSH or SSHA must be done explicitly by the user:

$$\text{Corrected SSH or SSHA} = \text{Reported SSH or SSHA} + \text{Crossover Correction.} \quad (6.1)$$

The L2_LR_SSH product is available in pass-sized granules that cover both swath sides [32]. The data are sampled on a geographically fixed, swath-aligned 2 km grid. Data on this grid have been smoothed from the native KaRIn measurements, so they are less noisy but also at a coarser resolution. KaRIn SSH measurements are also available in a separate file of this product at the native resolution (500 m) of the KaRIn downlink data, but they are also consequently sampled on the native grid of the center Doppler beam of the downlink data. This native grid is neither geographically fixed nor completely regular.

There are four separate files in the L2_LR_SSH product for each granule. Users may use the files individually or in combination as desired for their needs. The four files are:

- **Basic SSH ['Basic']:** The Basic file provides the most pertinent information for users who are interested primarily in the KaRIn SSH and SSHA measurements and who plan to take the measurements at face value. Measurements in the file are given on a 2 km, swath-aligned, geographically fixed grid. The Basic file is relatively small, at around 10 MB per granule.
- **Wind and Wave ['WindWave']:** The WindWave file provides information that is focused on the KaRIn measurements of backscatter, SWH, and wind speed. The file includes model information of related quantities. Measurements in the WindWave file are given on the same 2 km, swath-aligned, geographically fixed grid as used for the Basic file. The WindWave file is relatively small, at around 13 MB per granule.

- **Expert SSH with Wind and Wave ['Expert']:** The Expert file provides a superset of the information in the Basic and WindWave files as well as more detailed information on spacecraft ephemeris and attitude, instrument and environmental corrections, SWOT radiometer measurements, and pertinent geophysical fields. All of the fields in the Basic and WindWave files are reproduced in the Expert file using identical variable names. The additional detail in the Expert file is intended to enable advanced analyses of the data, including the replacement of correction values and the custom recalculation of the measurement quantities. Measurements in the Expert file are given on the same 2 km, swath-aligned, geographically fixed grid as used for the Basic file. The Expert file is therefore larger than the Basic and WindWave files, at around 33 MB per granule.
- **Unsmoothed SSH ['Unsmoothed']:** The Unsmoothed file provides SSH and backscatter information at the native sampling and resolution of the KaRIn downlink data (approximately 250 m sampling and 500 m resolution). The sampling in the Unsmoothed file follows the sampling of the central Doppler beam of the on-board-processed KaRIn data (see Sect. 9.2.1), but the measurements in the Unsmoothed file represent an average of all nine Doppler beams after beam combining (see Sect. 9.4.1). The sampling grid is therefore approximately swath aligned but neither geographically fixed nor perfectly regular. Given its finer resolution, the Unsmoothed file is considerably larger than the other three files in the L2_LR_SSH product, at about 660 MB per granule.

While the L2_LR_SSH product is designed primarily for oceanography, it is available globally, including over land. The assumptions that are built into the processing of the LR data (including both on-board and ground processing) often make the LR data unsuitable for inland water bodies smaller than around 10 km, however. In particular, it should be noted that even water features over land that are large enough to be resolved (i.e., larger than the product resolution) may not have reliable height measurements in the L2_LR_SSH product (see Sect 10.12 for an example). For inland water features with sizes in the range from 1 to 20 km, the accuracy of the L2_LR_SSH product may vary greatly with the viewing geometry (i.e., with the pass number). See [6] for details on the use of LR data over land. Users are advised against thinking of the LR products as coarser-resolution but more precise versions of HR products.

Figure 6.5 shows example data from the L2_LR_SSH product for the same area and the same SWOT pass (011_039) as is shown in Fig. 6.4 (see Sect. 6.8). Figure 6.5 shows only data that are flagged as *good*. The SSHA data shown are those posted on the 2 km geographically fixed grid and available in the Basic or Expert files. No data are reported in the product in the white background areas.

6.10 L1B_HR_SLC

As described in Sect. 9.2, data from the KaRIn instrument are split into separate low-rate and high-rate data streams within the instrument before downlink from the spacecraft. The high-rate stream is geared primarily toward hydrology. See Sect. 9.2.2 for details on the KaRIn HR on-board processing. The L1B_HR_SLC product [25] provides all relevant information from the HR data stream after ground processing to unpack the KaRIn telemetry, apply initial calibration




Figure 6.5: Example L2_LR_SSH data: (a) SSHA, directly as reported in the product, without the crossover correction applied; (b) SSHA after applying the crossover correction. The crossover correction or some other equivalent correction must be applied by the user to the heights from the L2_LR_SSH product in order to avoid large cross-track tilts. The graphics do not necessarily capture the full resolution and dynamic range of the data in the product.

parameters, and form single-look, complex (SLC) images for each of the $+y$ and $-y$ KaRIn receiver channels.

In addition to the individual SLC images, the L1B_HR_SLC product contains information that is necessary for later processing to obtain height and calibrated power measurements: spacecraft ephemeris and attitude; flags that provide information on data quality; radiometric calibration information; noise estimates for each KaRIn receiver channel; and the low-resolution, swath-aligned reference DEM used for image formation.

The SLC images contained in the L1B_HR_SLC product are coregistered with each other and focused relative to the reference DEM, so a high-quality interferogram can be formed by simply multiplying the SLC for one channel by the complex conjugate of the SLC for the other channel. The resulting interferogram would be flattened with respect to the reference DEM (see Sect. 8.3.1). The interferogram is not reported in the L1B_HR_SLC product, however. The SLC images are given in the slant plane (see Sect. 8.2.1) and are therefore *not* geolocated.

The SLC images have a resolution of approximately 0.75 m in slant range, corresponding to a ground-projected cross-track resolution of approximately 70 m at near range and approximately 10 m at far range. The along-track sample spacing and resolution are approximately 3 m and 5 m, respectively (the data are oversampled in along-track).

The intent of the L1B_HR_SLC product is to provide HR information that abstracts some of the engineering details of the implementation of the KaRIn instrument but that remains close to the fundamental measurement quantities of the instrument. The information in the L1B_HR_SLC product can therefore be processed by expert users who wish to apply their own specialized algorithms for interferogram formation, water detection, phase unwrapping, height reconstruction, spatial averaging, and/or resampling in order to meet their individual objectives.

The L1B_HR_SLC product is available in tile-sized granules that cover a single swath side (left or right) over an along-track extent of approximately 68 km [32]. Adjacent tiles overlap in the along-track direction. The tile boundaries are approximately fixed geographically, but the availability of data depends on the reconfigurable HR downlink mask (see Sect. 9.3). Full tiles of the L1B_HR_SLC product (i.e., tiles in which the HR downlink mask does not start or stop within the tile) are typically around 2 GB.

Figure 6.6 shows example data from the L1B_HR_SLC product for the same SWOT pass (011_039) as is shown for other examples of KaRIn data in this chapter (see Sect. 6.8). The phase of the interferogram provides the measurement of the surface height above the reference DEM used for SAR processing (see Sect. 8.3.4). The phase wraps modulo 2π rad and must be unwrapped (see Sect. 8.3.5), however. Note that the L1B_HR_SLC product does not actually contain the interferogram shown in panel (b); this interferogram is computed from the multilooked, normalized conjugate product of the two SLC images in the product [see Eq. (8.116)]. All of the SLC samples within this granule are flagged as *good*.

Because the 2-D arrays in L1B_HR_SLC product are not geolocated, Fig. 6.6 is shown with relative time of data collection on the vertical axis and slant range on the horizontal axis. These are equivalent to “slow time” and “fast time” as described in Sect. 8.2.1.1. The vertical axis is also equivalent to the along-track direction, and because the spacecraft nadir track generally goes from south to north for ascending passes, north is approximately toward the bottom (and slightly to the right) of the images. With slant-range as the horizontal axis, topographic variations in the land surface at far range (toward the right of the image) appear laid over toward the left

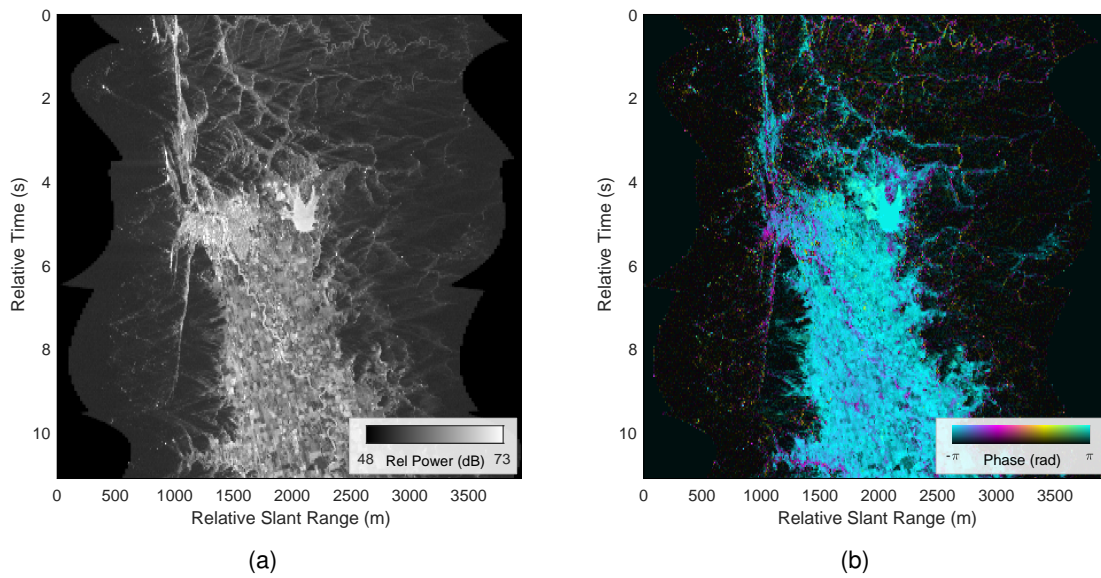


Figure 6.6: Example L1B_HR_SLC data: (a) Uncalibrated echo power from the $-y$ channel; (b) normalized complex interferogram. In (b), the grayscale indicates the magnitude of the total coherence γ_{tot} (black and white represent 0 and 1, respectively; see Sect. 8.3.2), and the overlaid color hue indicates the wrapped interferometric phase. The images are in radar slant-plane coordinates. The graphics do not necessarily capture the full resolution and dynamic range of the data in the product.

(see Sect. 8.2.1.3). Additionally, distances along the ground in cross track are not preserved in the projection to slant range.

Figure 6.6 shows a single L1B_HR_SLC tile granule (011_039_232L), which covers an area slightly larger than that shown by the green boxes in Figs. 6.8(a) and 6.9(a). The tile covers a portion of the Willamette River and Fern Ridge Lake (Reservoir) near the city of Eugene, Oregon. Fern Ridge Lake is the bright feature just above the center of the image. Eugene is visible to the left (east) of Fern Ridge Lake as an area that is somewhat brighter than its surroundings, but not as bright as the water surface of the lake. The Willamette River can be seen flowing downward and somewhat to the right (approximately north) from Eugene.

The bright, linear feature running up and down through the image to the left (toward near range) of Eugene is comprised of a series of echoes from bright, specular features near nadir. Note that echoes may appear even further to the left (at nearer range) than nadir due to topographic variations of the surface. That is, the point on the surface that is nearest to the radar (at the shortest range from the radar) for any given radar position over the pass may not actually be directly underneath the radar given the height variations of the surface.

Close examination of the far left and the far right of the images in Fig. 6.6 shows bands of black pixels in the images with no noise. These bands are zero filled and arise because of variations in the KaRIn starting range. That is, KaRIn records data for each pulse echo over a fixed-width span of time (or of slant range) called the data window (or the receive window). The starting range of the data window varies with along track in order to keep the echoes from the surface within the data window as the overall range to the surface varies. The large-scale range to the surface varies because of variations in both the spacecraft altitude and the surface topography. Consequently, the zero-filled pad on each side of the slant-plane images varies in order to compensate for changes in the starting range and to keep the slant range constant for any given slant-range array index in the SLC images of the L1B_HR_SLC product. Because the width of the data window is constant, the pad at near range becomes narrower when the pad at far range becomes wider and vice versa.

6.11 L2_HR_PIXC

The L2_HR_PIXC product [13] provides geolocated KaRIn height measurements from the HR data stream after pixel-level processing of the 2-D KaRIn SLC images. This processing is optimized for terrestrial hydrology.

In addition to pixel 3-D geolocation information, the L2_HR_PIXC product includes surface classification information that is derived from a combination of both KaRIn and prior (external) information. The classification algorithm includes automated water detection based on the surface reflectivity observed by KaRIn—water is assumed to be more reflective of radar signals than land at the wavelength and incidence angles of the KaRIn measurement. Additionally, the classification algorithm includes the flagging of dark water based on the use of prior information in conjunction with the KaRIn height measurements. The term “dark water” refers to water pixels that are believed to be water but that are too unreflective (“dark”) to be classified as water by the reflectivity-based detection algorithm (see Sect. 10.6 for additional details on dark water). Pixels that are detected as water but that have low interferometric coherence (see Sect. 8.3.2) are

Class	Definition
1	Land
2	Land near water
3	Water near land
4	Open water
5	Dark water
6	Low-coherence water near land
7	Open low-coherence water

Table 6.1: Definitions of numeric surface classification values for the L2_HR_PIXC product.

also classified separately given that water pixels are generally expected to have high coherence. Table 6.1 gives the surface classes defined in the L2_HR_PIXC product.

The L2_HR_PIXC product is derived from the SLC images of the L1B_HR_SLC product, but the data in the L2_HR_PIXC product are not given on a regular grid, nor are they even given directly in 2-D arrays. Rather, based on the computed pixel classification information, only a fraction of the pixels from the 2-D slant-plane arrays (see Sect. 8.2.1) are kept to create the L2_HR_PIXC product. That is, because the L2_HR_PIXC product is intended to provide observations of water, and because water typically covers a small fraction of the land surface, most pixels that are not believed to be water are discarded in order to reduce the volume of the product. Land pixels near water and land pixels covered by a predefined inclusion mask are also kept in the product. (The inclusion mask simply specifies limited areas over which pixels should always be kept.) Pixels that are kept are reported in the product as 1-D lists, but the slant-plane image indices of these pixels are provided in the product, so the pixels can be rearranged into sparse, 2-D, slant-plane images by the user if desired. Alternatively, the 3-D geolocation information of each pixel can be used to represent the data analogous to a point cloud, as might be obtained from a lidar instrument. The 3-D information for each pixel is given with respect to the reference ellipsoid (see Sect. 3.3.3).

Note that false detection of water can cause more pixels than intended to be kept. Conversely, missed detections may cause pixels to be discarded where they might otherwise have been desired to be kept. Missed detections would typically be associated with low reflectivity, however, and the height measurements for such dark pixels are typically much worse than nominal in any case.

In addition to the pixel 3-D geolocation and classification information, the L2_HR_PIXC product contains backscatter (σ°) information, quality flags, and uncertainty estimates. The L2_HR_PIXC product also contains a variety of fields giving geophysical parameters that allow the reported heights to be corrected for geoid variations and solid-Earth, pole, and load tide effects (see Sect. 11.3). Instrument corrections, sensitivities, and ancillary data such as spacecraft ephemeris and attitude that are used during processing are also reported. Along with intermediate variables such as complex interferogram values after only modest spatial averaging, these variables allow users to apply various kinds of customized processing to the SWOT data using only information that is available from the L2_HR_PIXC product.

Processing to compute the L2_HR_PIXC product involves adaptive spatial averaging that depends on the computed classification. For example, spatial averaging may be performed over a

given 2-D window in the slant plane but only over pixels in the window that belong to a predefined set of classes that depends on the class of the pixel at the center of the window. Consequently, the resolution of the L2_HR_PIXC product is variable. Before any pixels are discarded, the sample spacing is approximately 20 m in the along-track direction and between approximately 10 and 70 m in the cross-track direction, with finer sampling farther away from nadir. The resolution may be considerably coarser than the sampling, however.

The intent of the L2_HR_PIXC product is to provide detailed, low-level information on the KaRIn HR measurements of inland water features so that expert users can develop and apply their own specialized processing of SWOT data without necessarily having to deal with the complexities of the detailed KaRIn implementation or the data volumes of the L1B_HR_SLC product.

The L2_HR_PIXC product is available in tile-sized granules that cover a single swath side (left or right) over an along-track extent of 64 km [32]. The L2_HR_PIXC tile boundaries are approximately fixed geographically and coincide with the tiles of the L1B_HR_SLC product, although, unlike L1B_HR_SLC tiles, adjacent L2_HR_PIXC tiles do not overlap in the along-track direction. The availability of L2_HR_PIXC data depends on the reconfigurable HR downlink mask (see Sect. 9.3). Tiles of the L2_HR_PIXC product are typically less than 1 GB, but the size of a single tile can vary greatly depending on the amount of water, and therefore the number of pixels kept, in the observation.

While the L2_HR_PIXC product is designed primarily for terrestrial hydrology, it is available wherever HR data are collected, including over ocean surfaces that are covered by the HR downlink mask. The L2_HR_PIXC product does not contain oceanographically relevant information such as the height of the mean sea surface, estimates of the sea state bias, or corrections for ocean tides, however. Moreover, the assumptions that are built into the processing of the HR data are not necessarily suitable for oceanography. In particular, waves on the ocean surface may introduce artifacts in the HR data due to the coupling of their motion with the HR processing algorithms (see Sect. 10.11). It should also be noted that even if HR data are spatially averaged to the resolution of the LR products, the HR data would be less precise than the LR data due to the KaRIn on-board presumming of the HR data stream (see Sect. 9.2.2). Consequently, users are advised against thinking of the HR products as finer-resolution versions of LR products.

Figure 6.7 shows example data from the L2_HR_PIXC product for the same SWOT pass (011_039) as is shown for other examples of KaRIn data in this chapter (see Sect. 6.8). The tile shown here (011_039_232L) corresponds to the L1B_HR_SLC tile shown in Fig. 6.6. The numeric classification values defined by the L2_HR_PIXC product and shown in panel (b) are given in Table 6.1. For each panel, only samples that are not flagged as *bad* are shown.

In comparing the images of Fig. 6.7 to those of Fig. 6.6, note that the L2_HR_PIXC data are shown in geographic rather than slant-plane coordinates, so the images in Fig. 6.6 appear rotated and distorted with respect to the images of Fig. 6.7. Moreover, as not all slant-plane samples are kept in the L2_HR_PIXC product, many areas in Fig. 6.7 appear white, indicating that no data exist in the product for those locations on the ground. Nevertheless, features such as Fern Ridge Lake and the Willamette River are visible in Fig. 6.7 as they were in Fig. 6.6.

Figure 6.7(b) includes many noteworthy features that are typical of SWOT HR classification results. Many agricultural fields on either side of the Willamette River are incorrectly classified as open water (class 4) due to their high backscatter. Additionally, some areas in the city of Eugene, east of Fern Ridge Lake and at the lower right of panels (b) and (c), are incorrectly detected as

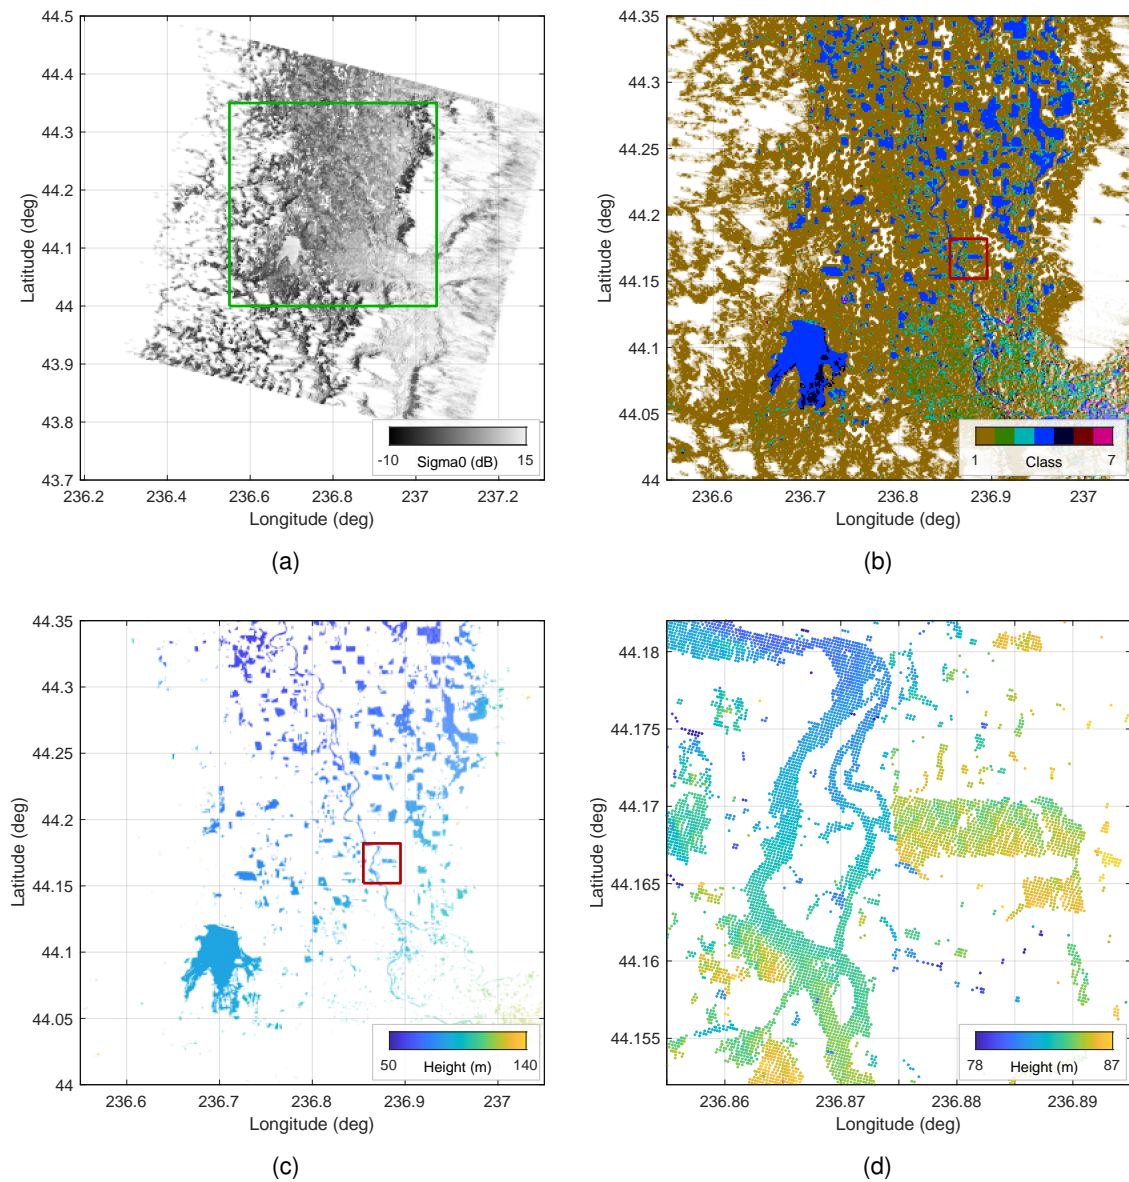


Figure 6.7: Example L2_HR_PIXC data: (a) radar backscatter (σ^0 ; see Sect. 8.1.4) over the tile; (b) surface classification value over the area of the green box in (a); (c) measured height with respect to the reference ellipsoid for pixels in class 4 (open water) for the same area as in (b); (d) measured height with respect to the reference ellipsoid for pixels in classes 3 (water near land) and 4 (open water) for the area of the red boxes in (b) and (c). Fern Ridge Lake appears at the lower left in (b) and (c). The enlargement in (d) includes a portion of the Willamette River. The graphics do not necessarily capture the full resolution and dynamic range of the data in the product.

open water as well. However, the height variations of areas within the city cause many of the false detections there to be decorrelated (see Sect. 8.4), so they are classified as low-coherence water (classes 6 and 7). Conversely, some areas in the southern portion of Fern Ridge Lake have low backscatter in panel (a) and are consequently classified as dark water (class 5) in panel (b). Finally, it is evident that a great deal of land area (class 1) is included in this L2_HR_PIXC granule. This is because land areas near detected water features are generally not discarded, so the false detection of water causes an excessive number of land samples to be stored as well.

The enlargement in Figure 6.7(d) shows pixel-level height values and demonstrates the level of fine-scale detail that is available in the L2_HR_PIXC product. It also illustrates the complexity of separating nearby water features, especially in the presence of false detections of water.

Note that no L2_HR_PIXC pixels in the tile are flagged as *good* (all are at least *suspect*) because the crossover correction (see Sect. 9.4.2) happens to be flagged as suspect for entire time span covering the tile in this example. Some pixels are flagged as *suspect* for other reasons as well. Specular ringing (see Sect. 10.7) also causes some pixels to be flagged as *degraded*.

6.12 L2_HR_RiverSP

The L2_HR_RiverSP product [15] provides KaRIn HR measurements of water surface elevation (WSE), slope, and width for predefined river reaches from a static prior river database (PRD). These measurements are computed with use of an algorithm that determines which pixels from the L2_HR_PIXC product should be attributed to each PRD reach then aggregates the pixel measurements accordingly. This product is intended for terrestrial hydrologists who are interested in rivers.

The locations of reaches in the PRD are defined by centerlines. Reaches are typically about 10 km long, with boundaries chosen in consideration of features such as confluences and dams (although information from the PRD is not necessarily optimal and may sometimes be erroneous). The reach boundary definitions also account for the nominal KaRIn swath coverage of the reference orbit (see Sect. 3.4). The PRD is the same between the calibration and nominal orbits. The PRD is static, so the reach definitions do not change over time, but they may change between different versions of the PRD and hence different versions of the L2_HR_RiverSP product as indicated by the CRID (see Sect. 5.4.1).

Each PRD reach comprises a series of static nodes, with nodes spaced about 200 m apart along the centerline. Therefore, there are usually around 50 nodes per reach. The L2_HR_RiverSP product provides information at both the reach and the node levels.

In addition to WSE, slope, and width, the L2_HR_RiverSP product provides multiple estimates of river discharge (the discharge estimates in early versions of the L2_HR_RiverSP product may be null filled, however). The product also contains quality flags, uncertainty estimates, and details of the processing parameters that were used to compute the primary product measurements. Additional information from the PRD that may be helpful in interpreting the data are provided as well.

The L2_HR_RiverSP product is available in continent-pass granules that cover the intersection of a SWOT pass (including both KaRIn swath sides) with a continent [32]. The continent definitions are hydrological and come from the PRD. The continent definitions are common between the L2_HR_RiverSP and the L2_HR_LakeSP products.

There are separate shapefiles for reaches and for nodes in the L2_HR_RiverSP product for each granule:

- **Reach:** The Reach shapefile provides WSE, slope, width (or area), discharge, and other information for each PRD reach in the granule.
- **Node:** The Node shapefile provides WSE and width (or area) for each PRD node in the granule.

The size of a granule of L2_HR_RiverSP data (including both the Reach and the Node shapefiles) is highly variable, although around 20 MB is typical.

Each reach or node in the PRD, and hence each reach or node in the L2_HR_RiverSP product, has a unique identifier. The identifiers can be used to determine which nodes belong to a given reach.

Note that reaches and nodes may appear in the L2_HR_RiverSP product even if no KaRIn measurement was made for them. This is because the absence of an observation (for example, because water was not detected) may convey useful information in itself. The reach or node vector object in the product would typically be null filled and flagged in such a case.

The PRD in use by SWOT is the SWOT River Database (SWORD) [41]. Therefore, the terms “PRD” and “SWORD” can usually be used interchangeably. The distinction would only be important if an alternative to SWORD were to be created and adopted as the PRD, but this is not envisioned for operational processing.

Figure 6.8 shows example data from the L2_HR_RiverSP product for the same SWOT pass (011_039) as is shown for other examples of KaRIn data in this chapter (see Sect. 6.8). The granule of L2_HR_RiverSP data shown here covers the portion of the pass over North America. For context, the area covered by panel (a) is indicated by the yellow box in Fig. 6.3. Note that not all reaches shown in Fig. 6.8(a) and (b) necessarily have measurements associated with them (the measurements for some reaches may be null filled). Fig. 6.8(b) includes a portion of the Willamette River north (downstream) of Eugene. The identifier of the reach in the red box is 78220000291 as of SWORD version 16, which was used for processing. All measurements in Fig. 6.8(c) and (d) are flagged as *suspect* because all pixels in the L2_HR_PIXC product are flagged as *suspect* (see Sect. 6.11).

Figure 6.8(c) and (d) exhibit some large, aphysical variations in the WSE and width from node to node that are larger than the random fluctuations that can be attributed to noise from the KaRIn instrument. These large deviations are caused by the presence of detected water features that are not actually part of the river but are near enough to the river that they are incorrectly used to estimate the WSE and area of some of the nodes of the river during ground processing. The features may not actually even be water; they could be falsely detected as water due to their high backscatter.

For example, considering the L2_HR_PIXC data from which the L2_HR_RiverSP measurements are computed, the feature at 44.16° latitude, 236.863° longitude just to the west of the river in Fig. 6.7(d) is likely a field that is adjacent to the river (note the difference in height between the field and the river), but it is incorrectly assumed to be part of the river during the estimation of river width. This causes the spike in width around 462.1 km on the horizontal axis of Fig. 6.8(d). Another similar field at 44.181° latitude, 236.857° longitude causes the spikes in

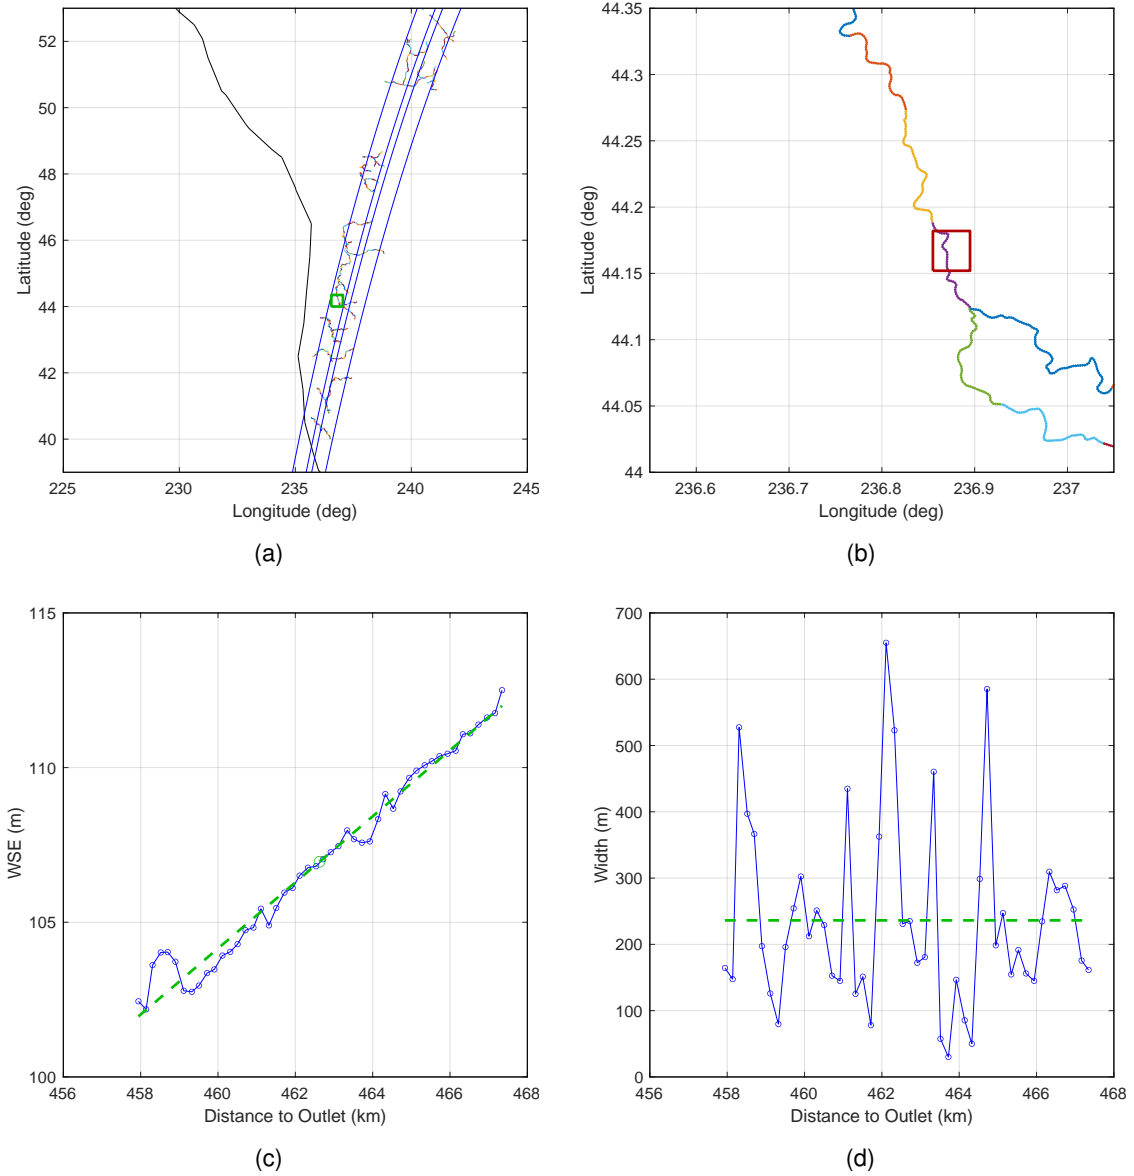


Figure 6.8: Example L2_HR_RiverSP data: (a) Reaches within the continent-pass granule; (b) enlargement of the area within the green box in (a); (c) WSE of the reach that is partially covered by the red box in (b); (d) width of the reach in (c). The black line in (a) represents the west coast of North America, and the blue lines represent the left and right sides of the KaRIn swath. In (a) and (b), colors are assigned arbitrarily to different reaches in order to distinguish them. Dots in (b) represent prior node locations from the PRD. In (c) and (d), each blue marker represents a node measurement, and the green dashed lines indicate the reach-level measurements [WSE and slope in (c) and the width in (d)]. The green box in (a) corresponds to the same area as in Figs. 6.7(b) and 6.7(c). The red box in (b) represents the same area as in Fig. 6.7(d). The graphics do not necessarily capture the full resolution of the data in the product.

WSE and in width at 458.5 km on the horizontal axes of Fig. 6.8(c) and (d). The L2_HR_PIXCVec product (see Sect. 6.14) provides information for determining which pixels are associated with individual node and/or lake objects. The river processing algorithms are likely to evolve over time, so future mitigations to these types of issues may be possible. One should use this example and the specific features in it for illustrative purposes only.

Note that the spatial extent of the continent-pass granule as shown in Fig. 6.8(a) is defined by the continent definitions used by SWORD [41]. Rivers that are further north of those shown in Fig. 6.8(a) are included in a separate granule for continent 8 (North American Arctic) rather than for continent 7 (North America and Caribbean, shown here).

6.13 L2_HR_LakeSP

The L2_HR_LakeSP product [18] provides KaRIn HR measurements of WSE and extent for lakes and lake-like features. Lakes objects are reported as polygons, each of which is associated with a measured WSE and area. These measurements are computed with use of an algorithm that determines which pixels from the L2_HR_PIXC product should be collectively considered as a lake object then aggregates the pixel measurements accordingly. This algorithm uses knowledge of which pixels are already assigned to rivers in the L2_HR_RiverSP product, although some features (e.g., reservoirs) can appear in both the L2_HR_RiverSP and the L2_HR_LakeSP products. The L2_HR_LakeSP product is intended for terrestrial hydrologists who are interesting in lakes.

In addition to WSE and area, the L2_HR_LakeSP product provides estimates of storage change for lakes with sufficient prior information from a static prior lake database (PLD) [42]. The product also contains quality flags, uncertainty estimates, and details of the processing parameters that were used to compute the primary product measurements. Additional information from the PLD that may be helpful in interpreting the data is provided as well.

The L2_HR_LakeSP product is available in continent-pass granules that cover the intersection of a SWOT pass (including both KaRIn swath sides) with a continent [32]. The continent definitions are hydrological and come from the PLD. The continent definitions are common between the L2_HR_RiverSP and the L2_HR_LakeSP products.

There are three separate shapefiles in the L2_HR_LakeSP product for each granule. These files contain overlapping information but are organized differently from one another because there is not always a one-to-one mapping between lake objects in the PLD and lake objects that are observed in the data. For example, one lake in the PLD may appear as two separate lakes in the KaRIn observation or vice versa. The three shapefiles are:

- **Observation-Oriented Shapefile ['Obs']:** The Obs shapefile contains one lake object for each observed lake feature that corresponds (fully or partially) to a lake in the PLD. Multiple observed features may correspond to the same PLD lake, or multiple PLD lakes may correspond to the same observed feature.
- **PLD-Oriented Shapefile ['Prior']:** The Prior shapefile contains one lake object per PLD lake in the granule. If a PLD lake is not observed, it is still present in the Prior shapefile, albeit null filled and without a polygon. Storage change is only reported in the Prior shapefile.

- **Shapefile of Unassigned Features ['Unassigned']:** The Unassigned shapefile contains one lake object for each observed lake feature that does not correspond to a lake in the PLD. The Unassigned shapefile complements the Obs shapefile such that the two shapefiles together represent the complete set of observed features for the granule. The Unassigned shapefile may contain many small features that correspond to false detections of water, however.

The size of a granule of L2_HR_LakeSP data (including all three shapefiles) is highly variable, although 50 MB is typical.

The three shapefiles above have many attributes in common and share a similar structure. Identifiers for each reported object allow features to be related across the different shapefiles. Users may use the different shapefiles individually or in combination as desired for their needs.

Figure 6.9 shows example data from the L2_HR_LakeSP product for the same SWOT pass (011_039) as is shown for other examples of KaRIn data in this chapter (see Sect. 6.8). The granule of L2_HR_LakeSP data shown here covers the portion of the pass over North America. For context, the area covered by panel (a) is indicated by the yellow box in Fig. 6.3. Fern Ridge Lake is at the lower left of Fig. 6.9(b), (c), and (d). Most measurements shown in in Fig. 6.9(c) and (d) are flagged as *suspect*. A few are flagged as *good*. The false detection of water in this tile [see Fig. 6.7(b) and the discussion in Sect. 6.11] causes a large number of “unknown” lakes to be reported in the Unassigned shapefile, as shown in Fig. 6.9(b).

Note that the spatial extent of the continent-pass granule as shown in Fig. 6.9(a) is defined by the continent definitions used by the PLD. Lakes that are further north of those shown in Fig. 6.9(a) are included in a separate granule for continent 8 (North American Arctic) rather than for continent 7 (North America and Caribbean, shown here).

6.14 L2_HR_PIXCVec

The L2_HR_PIXCVec product [21] is a complementary product to the L2_HR_PIXC product. The L2_HR_PIXCVec product is organized like the L2_HR_PIXC product and provides pixel-level information that is computed during the production of the L2_HR_RiverSP and L2_HR_LakeSP products. Notably, the L2_HR_PIXCVec product provides the mapping between L2_HR_PIXC pixels and the L2_HR_RiverSP node and/or L2_HR_LakeSP lake features to which those pixels contribute. The product also gives refined pixel-level geolocation information after feature-scale averaging and height-constrained geolocation. Similarly, the product gives pixel-level ice flags that are derived from feature-scale information.

The L2_HR_PIXCVec product is available in the same tile-sized granules as the L2_HR_PIXC product. The pixels in a given tile of the L2_HR_PIXCVec product map uniquely to the pixels in the corresponding tile of the L2_HR_PIXC product (see Sect. 6.11). The size of a tile of the L2_HR_PIXCVec product varies, though it is usually less than 1 GB.

Figure 6.10 shows example data from the L2_HR_PIXCVec product for the same SWOT pass (011_039) as is shown for other examples of KaRIn data in this chapter (see Sect. 6.8). For comparison, the images of pixel heights reported in the L2_HR_PIXC and L2_HR_PIXCVec products from Fig. 6.7(d) and Fig. 6.10(b) are shown next to each other in Fig. 6.11.

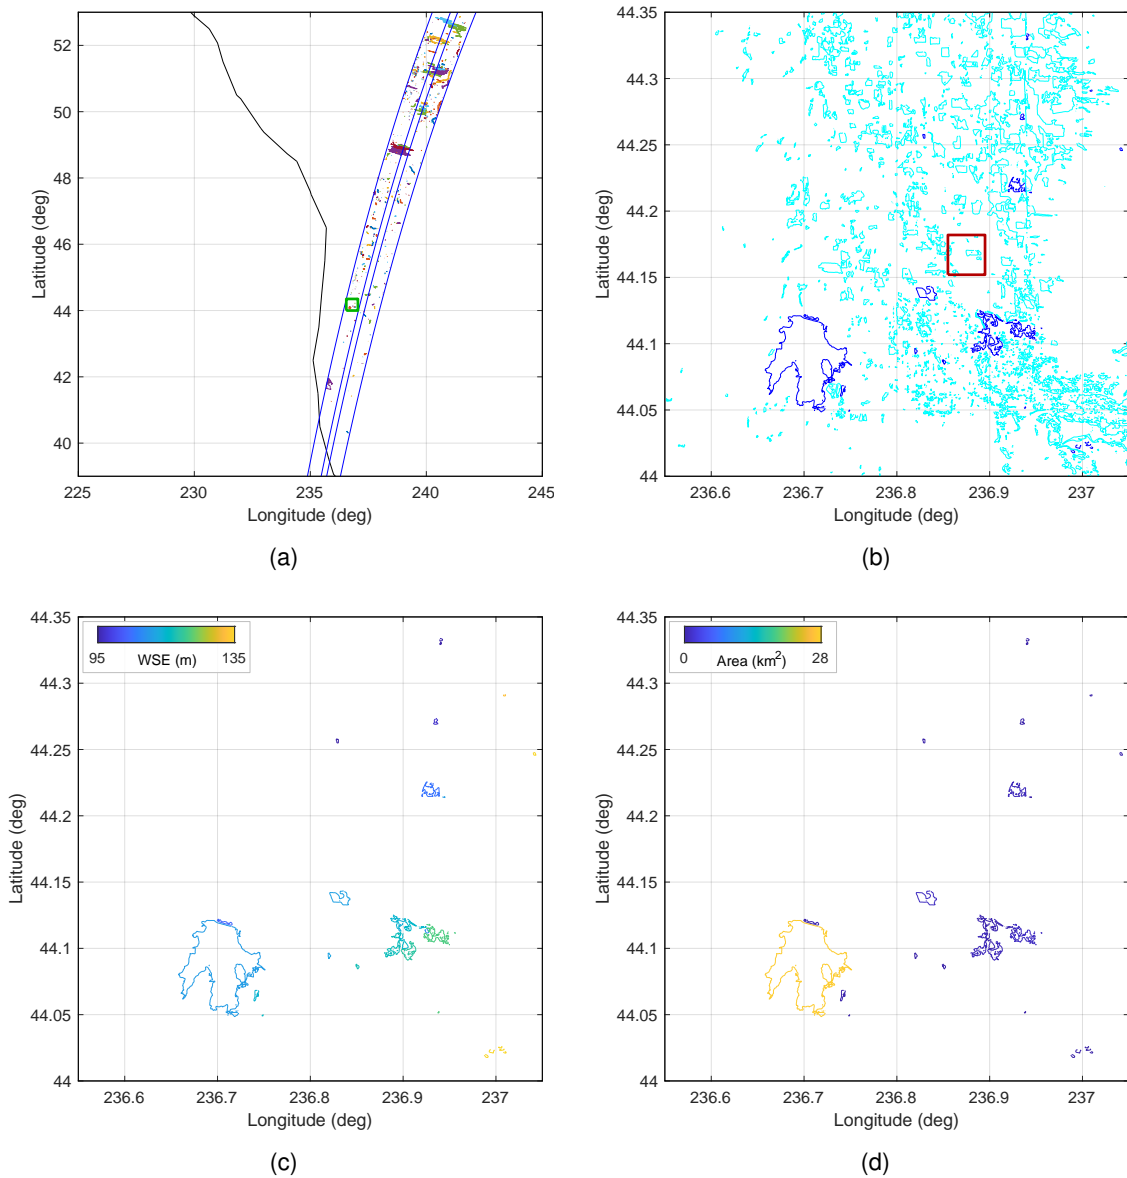


Figure 6.9: Example L2_HR_LakeSP data: (a) Lake objects within the continent-pass granule; (b) enlargement of the area within the green box in (a); (c) WSE of the PLD lake objects in (b); (d) area of the PLD lake objects in (b). The black line in (a) represents the west coast of North America, and the blue lines represent the left and right sides of the KaRIn swath. In (a), colors are assigned arbitrarily to different lake objects in order to distinguish them. In (b), blue and cyan polygons represent lakes from the Obs and Unassigned shapefiles. The latter contains objects that do not correspond to lakes in the PLD. In (c) and (d), the polygon colors represent the WSE or area of the lake objects from the Prior shapefile as indicated by the color bars. The green box in (a) corresponds to the same area as in Figs. 6.7(b) and 6.7(c). The red box in (b) represents the same area as in Fig. 6.7(d). The graphics do not necessarily capture the full resolution and dynamic range of the data in the product.




Figure 6.10: Example L2_HR_PIXCVec data: (a) pixel assignment to nodes and lakes; (b) heights after feature-scale averaging. Dot markers correspond to SWOT observation pixels. In (a), neighboring dots that are the same color were assigned to the same node or lake object during processing. The dot colors in (a) are arbitrary and serve only to distinguish the node or lake features to which pixels were assigned. The river centerline and prior node locations from the PRD are shown by the red line and the red open circles. Polygons from the L2_HR_LakeSP product are indicated by the cyan lines; these polygons are from the Unassigned shapefile and do not correspond to PLD lakes. In (b), the dot color represents the pixel height above the reference ellipsoid after extensive averaging. The area represented is the same as in Fig. 6.7(d), as indicated by the red boxes in Figs. 6.7(b) and (c), 6.8(b), and 6.9(b). The graphics do not necessarily capture the full resolution and dynamic range of the data in the product.

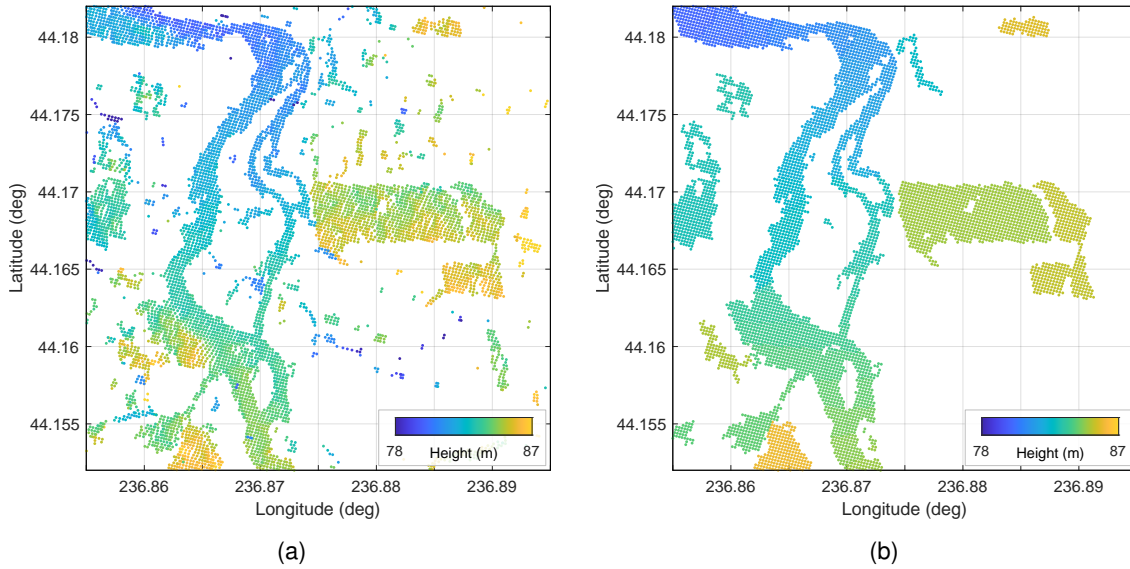


Figure 6.11: Comparison of L2_HR_PIXC (a) and L2_HR_PIXCVec (b) heights and horizontal geolocations. The images are the same as shown in Fig. 6.7(d) and Fig. 6.10(b). They are arranged here side by side for easier comparison. The heights are relative to the reference ellipsoid.

Like Fig. 6.11(a), Fig. 6.11(b) includes only pixels that are not flagged as *bad*. No data are reported in the product in the white background areas. However, many of the spurious L2_HR_PIXC pixels in Fig. 6.11(a) that are not likely to correspond to real water features on the ground have null-filled latitude and longitude values in the L2_HR_PIXCVec product, so these pixels do not appear in Fig. 6.11(b). Such pixels have null-filled latitude and longitude values because the coordinates in the L2_HR_PIXCVec product are derived after height-constrained geolocation that relies on averaging pixel heights over the water features identified in the L2_HR_RiverSP and L2_HR_LakeSP products. That is, only pixels that correspond to known features in the river and/or lake databases or pixels that occur in clusters that are large enough to be included in the L2_HR_LakeSP Unassigned shapefile have valid height-constrained geolocations in the L2_HR_PIXCVec product. Some of these may nonetheless correspond to false detections of water, however.

A close comparison of Fig. 6.11(a) and (b) shows that the pixel geolocations are less noisy in the latter even when considering only the pixels that appear in both. The lower noise in the L2_HR_PIXCVec horizontal geolocation can be observed in its more regular, grid-like arrangement of the pixel coordinates within a given water feature. Relatedly, one can observe that there is less spatial variation in the heights of the pixels within a given water feature in the L2_HR_PIXCVec product as well. This reduction of geolocation noise is the purpose of the height-constrained geolocation. Of course, the L2_HR_PIXCVec product may therefore also be unsuitable for attempting to infer information about fine-scale spatial variations in height internal to a given water feature. Moreover, the smoothing of the data in the L2_HR_PIXCVec product may make it more difficult to determine when pixels were incorrectly assigned to water features,

as in the cases of the agricultural fields attributed to the river in this example (see Sect. 6.11). The height information in the L2_HR_PIXCVec product should therefore only be used with caution.

6.15 L2_HR_Raster

The L2_HR_Raster product [22] provides height, inundation extent, and backscatter (σ°) of water surfaces on a uniformly sampled, 2-D geographic grid. The product is derived from the L2_HR_PIXC product and is intended to be simpler and easier to use than the L2_HR_PIXC product but general enough to allow for the investigation of complex cases such as flood events, during which the assumptions used to generate the L2_HR_RiverSP and L2_HR_LakeSP products may not be appropriate. The L2_HR_Raster product generally has coarser resolution than the L2_HR_PIXC product, however. The L2_HR_Raster product is sparse in the sense that many samples (most land areas) are null filled.

In addition to water height, area, and backscatter, the L2_HR_Raster product provides quality flags, uncertainty estimates, and details of the processing parameters that were used to compute the primary product measurements.

The L2_HR_Raster product from routine, operational processing is available in scene-sized granules that cover both sides of the KaRIn swath (left and right) over an along-track extent of 128 km [32]. Each of these L2_HR_Raster scenes covers a 2×2 set of L2_HR_PIXC tiles. Note that while the scene boundaries are aligned with the SWOT swath, the sampling grid is aligned geographically and is geographically fixed for each scene. The L2_HR_Raster product is produced operationally at each of two different resolutions (100 m and 250 m), which are provided in separate files with otherwise identical format. The size of a L2_HR_Raster file varies, but 50 MB and 15 MB are typical sizes for the files with 100 m and 250 m resolution.

As with the L2_HR_PIXC product, L2_HR_Raster data are available wherever HR data are collected, including over ocean surfaces that are covered by the HR downlink mask. The caveats in Sect. 6.11 about using HR data for oceanography apply to the L2_HR_Raster product, however.

Figure 6.12 shows example data from the L2_HR_Raster product for the same SWOT pass (011_039) as is shown for other examples of KaRIn data in this chapter (see Sect. 6.8). All data that are not flagged as *bad* are shown.

While much of the detected water around the Willamette River in this example is unlikely to be real water, the example shows the potential utility of the L2_HR_Raster product for use in cases such as major flood events, in which the actual distribution of water on the surface does not match the assumptions of the PRD and the PLD.

Note that if comparing the L2_HR_Raster WSE in Fig. 6.12(b) to the L2_HR_PIXC height in Fig. 6.7(c), one should bear in mind that the L2_HR_PIXC height is relative to the ellipsoid and does not have tide effects removed, whereas the L2_HR_Raster WSE is relative to the geoid and does have tide effects removed.

6.16 L2_HR_RiverAvg

While the L2_HR_RiverSP product provides WSE, slope, width, and discharge for PRD reaches on a per-pass basis, the L2_HR_RiverAvg product [17] provides such information on a per-cycle

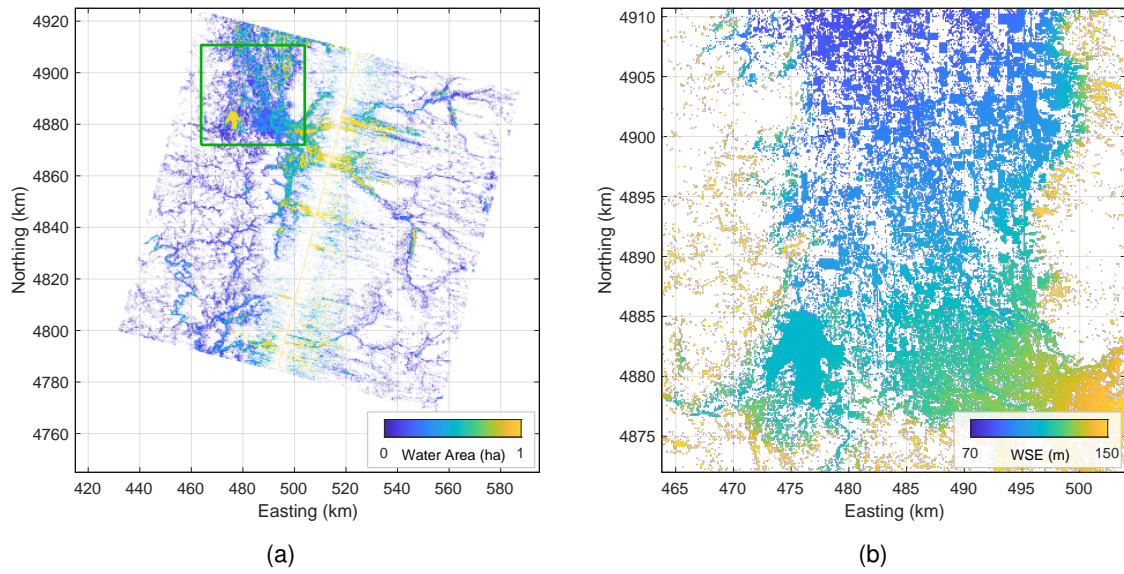


Figure 6.12: Example 100 m L2_HR_Raster data: (a) water area per Raster pixel over the scene granule; (b) WSE over the area corresponding to the green box in (a). The area covered in (b) is the same as in Figs. 6.7(b) and (c), 6.8(b), and 6.9(b)–(d). No data are reported in the product in the white background areas. The graphics do not necessarily capture the full resolution and dynamic range of the data in the product.

basis. A given reach may be observed many different times (from passes with different pass numbers) during a 21-day repeat cycle. The L2_HR_RiverAvg product provides the mean values of the WSE, slope, width, and discharge over those different observations within the cycle. The product also provides the minimum, maximum, and median WSE over the cycle as well as the slope, width, and discharge that coincide with the minimum, maximum, and median WSE. The product also provides quality flags and uncertainty estimates.

The intent of the L2_HR_RiverAvg product is to provide summary information of KaRIn river observations over wide spatial scales and relatively coarse, regular temporal sampling in a manner that is easily digestible for large-scale analyses (e.g., assimilation models).

The L2_HR_RiverAvg product is available in basin-cycle granules [32]. Basins in this context are defined hydrologically by the PRD and are common between the L2_HR_RiverAvg and the L2_HR_LakeAvg products. Each hydrological continent is divided into several basins. A granule of the L2_HR_RiverAvg product gives measurements, summarized over the 21-day repeat cycle as described above, for all reaches within a particular basin for a particular cycle. Node-level measurements are not available in the L2_HR_RiverAvg product.

The L2_HR_RiverAvg product is available only during the science phase (not the calibration phase) of the mission.

6.17 L2_HR_LakeAvg

Similar to the L2_HR_RiverAvg product, the L2_HR_LakeAvg product [20] provides information for PLD lakes on a per-cycle basis. The L2_HR_LakeAvg product provides the mean values of the WSE, area, and storage change over the different observations of a given lake within the cycle. The product also provides the minimum, maximum, and median WSE over the cycle as well as the area and storage change that coincide with the minimum, maximum, and median WSE. The product also provides quality flags and uncertainty estimates.

The intent of the L2_HR_LakeAvg product is to provide summary information of KaRIn lake observations over wide spatial scales and relatively coarse, regular temporal sampling in a manner that is easily digestible for large-scale analyses (e.g., assimilation models).

The L2_HR_LakeAvg product is available in basin-cycle granules similar to the granules of the L2_HR_RiverAvg product (see Sect. 6.16). The L2_HR_LakeAvg product is organized only by PLD lake, equivalent to the Prior shapefile of the L2_HR_LakeSP product; there is no equivalent of the Obs or Unassigned shapefiles of the L2_HR_LakeSP product in the L2_HR_LakeAvg product.

The L2_HR_LakeAvg product is available only during the science phase (not the calibration phase) of the mission.

6.18 L2_HR_FPDEM

The L2_HR_FPDEM product is unlike other KaRIn products in that it is produced only after SWOT has been in operations for a lengthy amount of time. The intent of the product is to provide a DEM of the floodplain near water features that have been observed over the course of the mission to date. The computation of the product assumes that many observations at different water levels from the beginning of the mission to the time of the DEM computation are available

such that the relationship between water height and extent can be used to infer the DEM heights using the so-called “bathtub-ring” approach. It may also be possible to extract unflooded elevations directly in some cases. The product is null filled where no water is observed.

The L2_HR_FPDEM product is provided on $1^\circ \times 1^\circ$ latitude-longitude granules. The product is computed in both gridded and ungridded forms. The gridded version uses a uniform latitude-longitude grid with approximately 30 m posting. The ungridded version gives relevant pixels in point-cloud form.

Chapter 7

Operational Data Processing

This chapter introduces the ground processing system that generates the SWOT science data products described in Ch. 6.

7.1 Algorithm Flow

Figure 7.1 shows the top-level algorithm flow of the ground processing used to create the SWOT science products described in Ch. 6. The KaRIn LR and HR processing chains are encapsulated in the boxes at the lower right.

In these diagrams, rectangular boxes with sharp corners indicate product generation executables (PGEs) or “processors” that create granules of the standard or intermediate data products indicated. Typically, the name of a given processor is related to the names of the main product that it creates (e.g., the PGE_L2_RAD_GDR processor creates the L2_RAD_GDR product). Square brackets around the letter ‘I’ indicate that the algorithm flow is applicable to both a near-real-time or IGDR product as well as the correspond delayed-time or GDR product. For example, “L2_RAD_[I]GDR” represents both L2_RAD_GDR and L2_RAD_IGDR. The colors of the processing boxes indicate which organization leads the development of the processor software, not which organization runs the software operationally to produce data products. Boxes with rounded corners indicate collections of processors that are abstracted for the purposes of illustration. See Sect. 7.2 for a discussion of the auxiliary (“aux”) data indicated, and see Sect. 7.3 for a discussion of the intermediate products. SSALTO (Segment Sol multimissions d’ALTimétrie, d’Orbitographie et de localisation précise) is the CNES multimission ground segment for altimetry. The Level 0 (LOA and LOB) data coming into the flow generally represent instrument telemetry that has only been reformatted or parsed but not processed any further.

As indicated in Fig. 7.1, the KaRIn LR and HR processing chains depend on the orbit (POE or MOE), attitude (ATTD_RECONST), radiometer (L2_RAD_GDR or L2_RAD_IGDR), and nadir-altimeter (L2_NALT_GDR or L2_NALT_IGDR) products. The KaRIn HR processing chain depends on crossover calibration results (INT_LR_XOverCal) from the LR processing chain. Both the LR and HR processing chains depend on KaRIn dynamic-calibration (INT_KCAL_Dyn) processing. Note that the KaRIn L1B products are not precisely geolocated and are intended only for expert users.

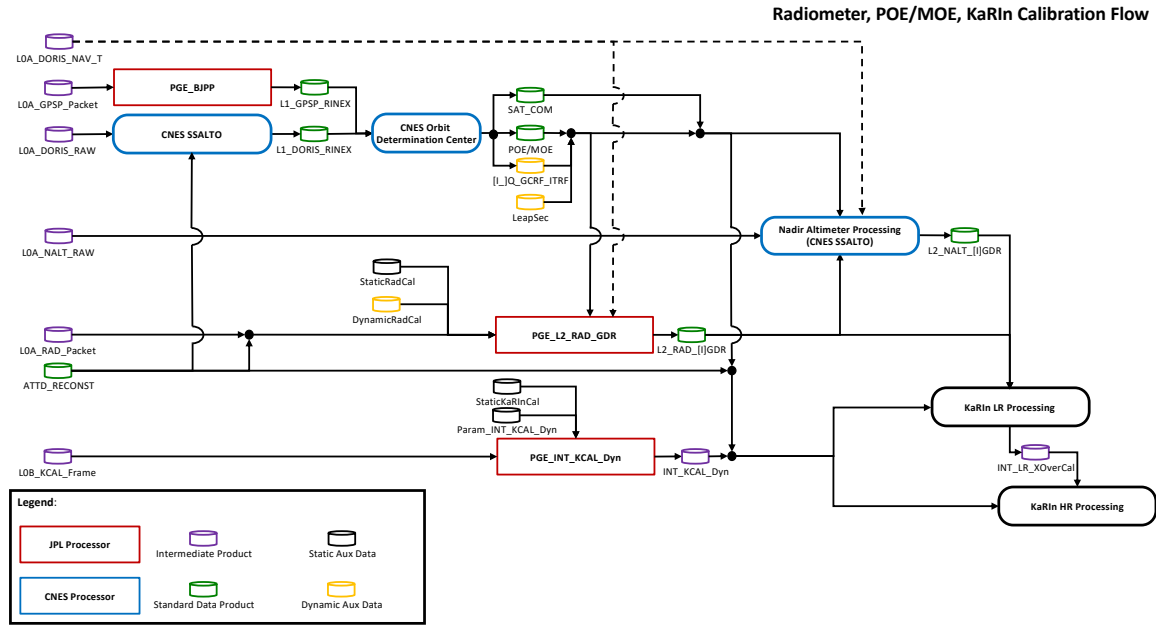


Figure 7.1: Top-level algorithm flow for ground-based science processing.

Figures 7.2 and 7.3 show the algorithm flows for the KaRIn LR and HR processing indicated at the lower right of Fig. 7.1. These diagrams follow the same representation conventions as Fig. 7.1.

KaRIn LR ground processing begins from the spatially averaged, nine-beam interferograms that were formed on board the spacecraft (see Sect. 9.2). The LR L1B processing produces bias-corrected nine-beam interferograms (without measured heights). Height reconstruction, beam combining, and the computation of geophysical corrections occur in LR L2 processing, with the results stored in an intermediate product called *L2_LR_PreCalSSH*. Crossover calibration results are computed from the *L2_LR_PreCalSSH* product, and the results are incorporated into the final *L2_LR_SSH* product.

KaRIn HR processing begins from presumed KaRIn echo data (see Sect. 9.2). The HR L1B processing produces coregistered single-look-complex SAR images for each channel. These are stored in the *L1B_HR_SLC* product. The *L2_HR_PIXC* processor then performs interferogram formation, water detection and classification, adaptive averaging, phase unwrapping and ambiguity resolution, geolocation, computation of geophysical corrections, and pruning to produce the sparse *L2_HR_PIXC* or “pixel-cloud” product. River, lake, raster, and floodplain-DEM processing each then start with the pixel-cloud data. The river processing chain assigns pixels to known river reaches and estimates reach-level measurement quantities. The lake processing chain similarly handles lakes as well as unrecognized features that have been classified as water. The *L2_HR_Raster* processor creates products for use in areas or at times where water features of interest may not be well captured by the river and lake products. Finally, the floodplain-DEM processor creates height maps of areas surrounding water features based on the history of SWOT data collected to date.

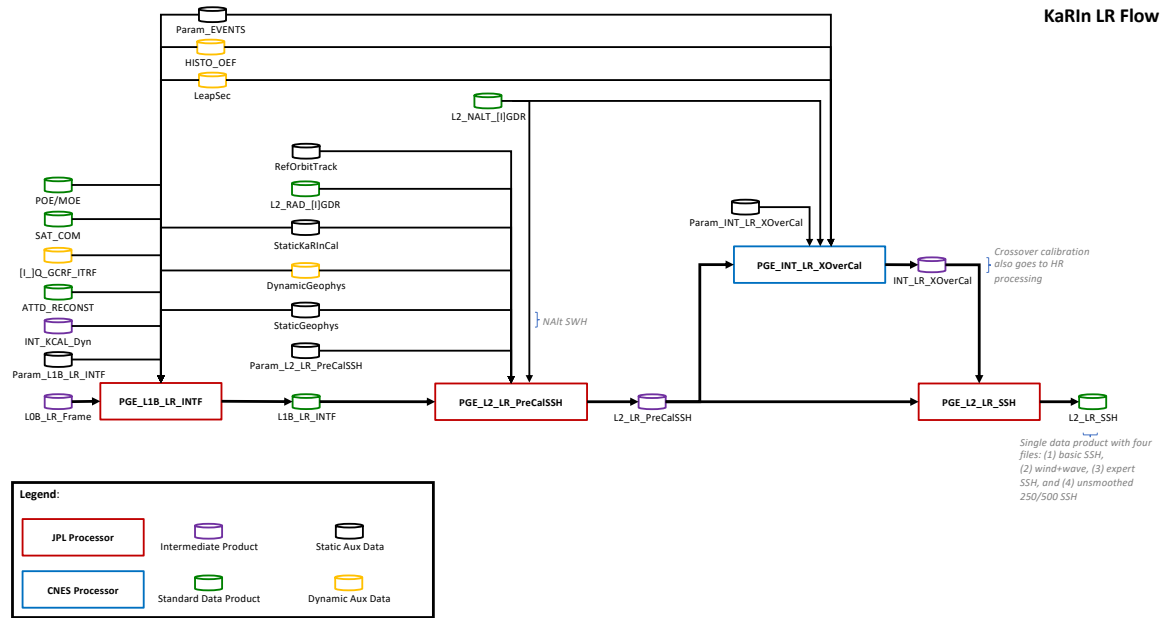


Figure 7.2: KaRIn LR algorithm flow for ground-based science processing.

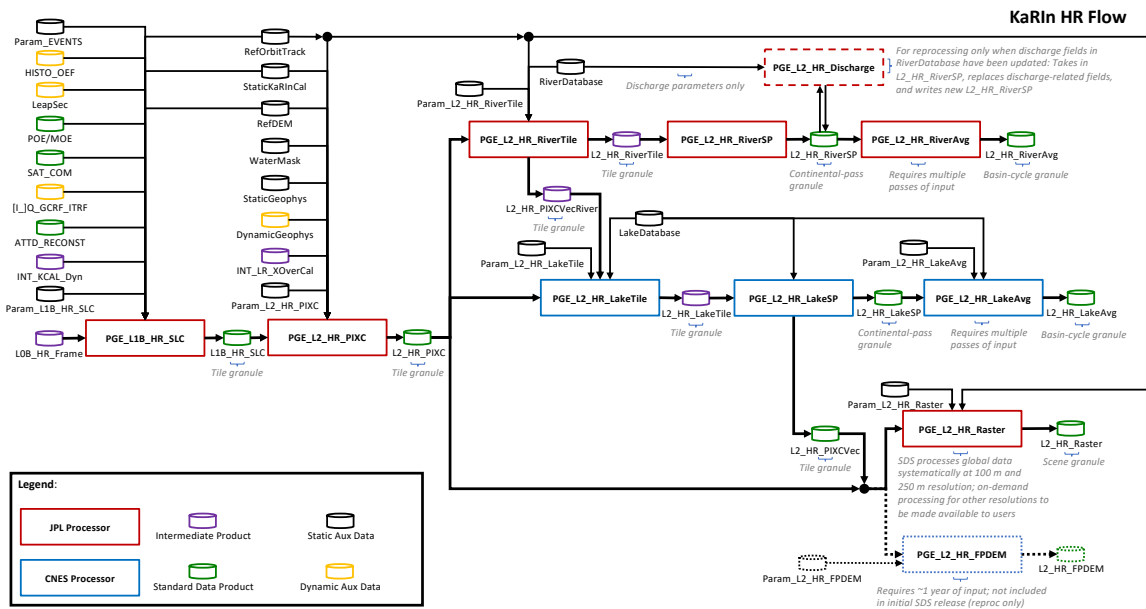


Figure 7.3: KaRIn HR algorithm flow for ground-based science processing.

7.2 Auxiliary Inputs

The processing illustrated in Figs. 7.1, 7.2, and 7.3 relies on a great deal of auxiliary information, which is summarized in this section. Generally, auxiliary data are data that originate from outside the SWOT ground system. Dynamic auxiliary data (DAD), such as weather data, are typically generated operationally by other organizations and ingested into the SWOT processing system. Such data typically change with time. In contrast, static auxiliary data (SAD), such as files of calibration constants, typically do not change with time for a given processing configuration. That is, static auxiliary data are typically created offline (not operationally) by the algorithm team and delivered to the SWOT processing system. The team may update the static auxiliary data occasionally, but the automated operational processing system does not select particular static files based on when the SWOT data were collected except where different static files are needed for the different orbit phases.

Note that the icons for auxiliary data in Figs. 7.1, 7.2, and 7.3 often represent multiple files. When one of these sets of auxiliary data is indicated as an input to a particular processor in the diagrams, not all of the files in the set are necessarily used by the processor.

The sets of static auxiliary data are summarized below:

- LakeDatabase contains the prior lake database (PLD).
- Param_EVENTS contains thresholds that specify how data that are close in time to spacecraft events such as eclipse transitions or maneuvers should be flagged.
- Param_INT_KCAL_Dyn contains algorithm parameters for the INT_KCAL_Dyn processor.
- Param_INT_LR_XOverCal contains algorithm parameters for the INT_LR_XOverCal processor.
- Param_L1B_HR_SLC contains algorithm parameters for the L1B_HR_SLC processor.
- Param_L1B_LR_INTF contains algorithm parameters for the L1B_LR_INTF processor. The data includes copies of some of the OBP tables (see Sect. 9.2) on board the spacecraft as well as the ocean reference surface used for LR ground processing.
- Param_L2_HR_FPDEM contains algorithm parameters for the L2_HR_FPDEM processor.
- Param_L2_HR_LakeAvg contains algorithm parameters for the L2_HR_LakeAvg processor.
- Param_L2_HR_LakeTile contains algorithm parameters for the L2_HR_LakeTile processor.
- Param_L2_HR_PIXC contains algorithm parameters for the L2_HR_PIXC processor.
- Param_L2_HR_Raster contains algorithm parameters for the L2_HR_Raster processor.
- Param_L2_HR_RiverTile contains algorithm parameters for the L2_HR_RiverTile processor.
- Param_L2_LR_PreCalSSH contains algorithm parameters for the L2_LR_PreCalSSH processor.

- RefDEM contains the reference DEM used for HR processing.
- RefOrbitTrack contains files that characterize the nadir track of the reference orbit. This information defines geographically fixed granule boundaries and sampling grids.
- RiverDatabase contains the prior river database (PRD).
- StaticGeophys contains static parameters for a collection of geophysical models. Examples include tide coefficients and gridded geoid heights.
- StaticKaRInCal contains static calibration parameters for KaRIn processing. This includes files of KaRIn antenna patterns and reference chirps for HR processing.
- StaticRadCal contains static calibration parameters for radiometer processing.
- WaterMask contains the prior water probability map, the bright-land mask, and the inclusion mask used for HR classification.

The sets of static dynamic auxiliary data are summarized below:

- DynamicGeophys contains time-varying parameters for geophysical models that are used to provide corrections to and additional information regarding SWOT products. For example, this set of dynamic auxiliary data contains files on Earth surface pressure and temperature that are updated multiple times per day based on weather models; these data are used to account for the propagation delay of the radar signals through the atmosphere at the time of the SWOT observation.
- DynamicRadCal contains radiometer calibration coefficients that evolve over time and are estimated offline.
- Histo_OEF contains information on spacecraft events such as station-keeping maneuvers, yaw flips, and eclipse transitions. These events disturb the spacecraft state and may affect the quality of the SWOT measurements.
- LeapSec contains information on the number of leap seconds in the UTC time system.
- Q_GCRF_ITRF and I_Q_GCRF_ITRF contain the transformation (rotation) between the GCRF and ITRF frames (see sects. 3.3.1 and 3.3.2). This information is needed because the KaRIn attitude information in the ATTD_RECONST product is given with respect to the inertial frame. I_Q_GCRF_ITRF has a shorter latency than Q_GCRF_ITRF, but its estimates of the small variations in the Earth rotation vector are less precise.

7.3 Intermediate Data Products

Intermediate data products are generated by PGEs in the production systems but are not intended for distribution or archive. The information in these products that is relevant to science users is captured by downstream standard science data products, which are archived and

distributed to users. Intermediate products are usually generated to facilitate the design architecture of the data production system. For example, the product definition and science content of the L2_HR_RiverTile intermediate product is identical to that of the L2_HR_RiverSP standard product; the two simply have different granule spatial extents, with the tile-sized L2_HR_RiverTile granules facilitating parallelized production processing. Similarly, the product definition and science content of the L2_LR_PreCalSSH intermediate product is identical to that of the L2_LR_SSH standard product except that the variables related to crossover corrections are null filled in the former but populated in the latter. The existence of intermediate products should generally be transparent to science users of standard products (see Ch. 6). They are described here only for completeness.

The intermediate products are summarized as follows:

- INT_KCAL_Dyn contains dynamic KaRIn calibration information such as loopback data (see Sects. 9.1 and 9.2) and KaRIn instrument parameters that are applicable to both LR and HR processing. Relevant information from this product is reported in the LR and HR standard products.
- INT_LR_XOverCal contains estimates of KaRIn errors from crossover calibration processing. The corrections used are reported in the L2 LR and HR standard products.
- L0A_DORIS_NAV_T contains raw DORIS navigator telemetry packets, which include the real-time orbit ephemeris estimated on board the spacecraft (see Sect. 4.4.4). This orbit ephemeris is used to generate the L2_NALT_OGDR and L2_RAD_OGDR products. The best available ephemeris information is given in the POE standard product.
- L0A_DORIS_Raw contains raw DORIS telemetry packets (see Sect. 4.4.4). The raw tracking data from the telemetry is given in the L1_DORIS_RINEX product.
- L0A_GPSP_Packet contains raw GPSP telemetry packets (see Sect. 4.4.5). The raw tracking data from the telemetry is given in the L1_GPSP_RINEX product.
- L0A_NALT_Raw contains raw telemetry from the nadir altimeter (see Sect. 4.4.2). The nadir altimeter measurements are given in the L2_NALT_GDR standard product.
- L0A_RAD_Packet contains telemetry packets from the radiometer (see Sect. 4.4.3). The radiometer measurements are given in the L2_RAD_GDR standard product.
- L0B_HR_Frame contains unprocessed KaRIn HR or “land telemetry frames. The HR measurements are given in the multiple L2 HR standard products.
- L0B_KCAL_Frame contains unprocessed calibration frames from the KaRIn telemetry (see the INT_KCAL_Dyn intermediate product).
- L0B_LR_Frame contains unprocessed KaRIn LR or “ocean telemetry frames. The LR measurements are given in the L2_LR_SSH standard product.
- L2_HR_LakeTile contains lake (prior, observed, and unidentified-feature) information in tile-sized granule. All relevant information in the L2_HR_LakeTile product is reported in the continent-pass granules of the L2_HR_LakeSP product.

- L2_HR_PIXCVecRiver contains a subset of the information (i.e., that from river processing) contained in the L2_HR_PIXCVec standard data product. The latter also contains information from lake processing.
- L2_HR_RiverTile contains river reach and node observations in tile-sized granules. The same information is reported in the L2_HR_RiverSP product, which is effectively the concatenation of many tile-sized L2_HR_RiverTile granules to form the continent-pass granule of the L2_HR_RiverSP standard product.
- L2_LR_PreCalSSH contains the same information as the L2_LR_SSH product except that the former does not include the estimated crossover corrections (those fields are null filled in the L2_LR_PreCalSSH product).

Chapter 8

KaRIn Measurement Principles

This chapter provides a summary of the fundamental physical principles underlying the measurements that are reported in the KaRIn data products of Ch. 6. Specifics of the KaRIn measurement design are discussed in Ch. 9.

8.1 Radar Fundamentals

This section covers some basic principles of radar measurements.

8.1.1 Basic Radar Ranging

Fundamentally, a radar is a sensor that transmits a radio-wave signal and measures the properties of some target of interest based on the echoes received after the transmitted signal is reflected by the target. The radar signal is usually an electromagnetic transmission with some modulation on a microwave carrier. The carrier frequency f_c has a wavelength λ given by

$$\lambda = \frac{c}{f_c}, \quad (8.1)$$

where c is the speed of light in a vacuum (2.997922458×10^8 m/s). Different radars operate in different frequency bands and therefore with different wavelengths.

Suppose that the radar transmits a signal whose voltage or electric field as a function of time t is given by $s_T(t)$. The received echo $s_R(t)$ from a discrete, point-like target at a distance or range ρ from the radar is generally a delayed and scaled copy of the transmitted signal, so

$$s_R(t) = C s_T(t - \tau), \quad (8.2)$$

where τ is the delay of the receive (Rx) echo relative to the transmit (Tx) signal, and C is a scaling term. Note that the variable C is used throughout this chapter to indicate scaling terms that are unimportant in the corresponding context; C does not necessarily have the same value across different instances in this document. The delay τ is related to the range ρ via the speed c of the electromagnetic signal:

$$\rho = \frac{c\tau}{2}. \quad (8.3)$$




Figure 8.1: Illustration of radar ranging for a single point target.

The factor of 2 in this equation arises because of the round-trip travel of the radar signal. (The signal propagation speed is slightly but noticeably slower in the Earth's troposphere and ionosphere.)

An illustration of radar ranging is shown in Fig. 8.1. In this illustration, the red curve indicates the electric field of a notional transmit signal from the radar as a function of time. The transmit signal in this example is a single cycle of a sinusoid. The blue curve indicates the electric field of the received echo at the radar from a discrete, point-like target at some distance ρ . The received echo is the sum of a delayed, scaled version of the transmitted signal and additive receiver noise. The echo delay τ is related to the range ρ through Eq. (8.3).

8.1.2 Echo Phase

Radars, like radios, are typically heterodyne systems that transmit a signal with some form of modulation on a microwave-frequency carrier. The radar then removes the carrier from the echo after the echo is received. Conceptually, the radar operates as illustrated by the block diagram in Fig. 8.2. The radar derives a pure sinusoidal voltage at the carrier frequency f_c from a stable local oscillator (STALO or “clock”). To generate the transmit signal $s_T(t)$, a baseband waveform $s_C(t)$ is multiplied by the oscillator output then amplified and sent to the transmit antenna. The echo

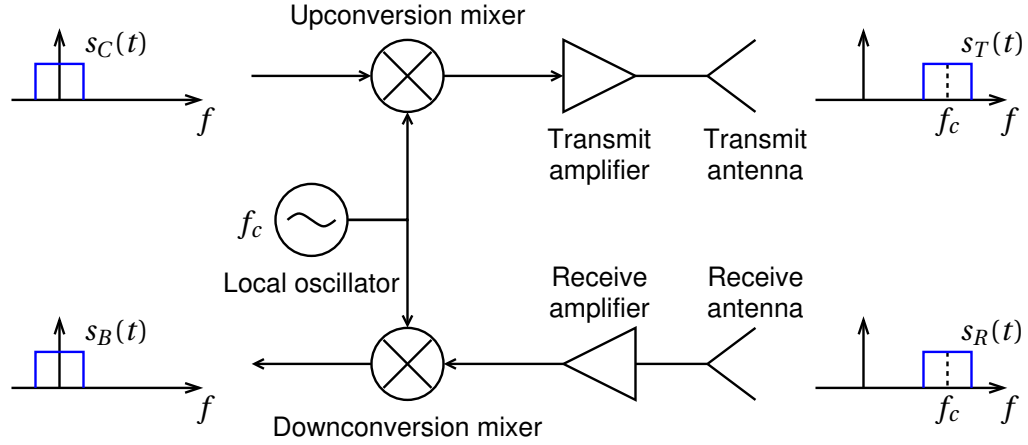


Figure 8.2: Conceptual block diagram of a radar that shows the principle of upconversion and downconversion. The diagram is not intended represent the actual design of any SWOT radar hardware.

$s_R(t)$ at the receive antenna is amplified and multiplied by the output of the same oscillator (but at the time of the delayed echo) in order to remove the carrier frequency and obtain a baseband echo $s_B(t)$. The “mixing” operations of multiplying by the oscillator output thus shift the signals up and down in frequency. Filtering (not shown in the block diagram) suppresses undesired “image” frequencies. These steps are called upconversion and downconversion, respectively.

Note that the block diagram in Fig. 8.2 is highly simplified and for illustration only. An actual radar may have many stages of amplification, upconversion, downconversion, and/or filtering in both the transmit and receive chains. Additionally, the diagram shows separate transmit and receive antennas, but the same antenna may be used for both, and multiple antennas and multiple receive channels may be employed.

The baseband transmit waveform $s_C(t)$ is usually represented by its complex envelope. The complex envelope is a representation of the modulation of a signal, independent of its carrier. The transmit signal, including the carrier, would be the real part of the product of the complex envelope and a complex exponential representing the carrier:

$$s_T(t) = \Re \{ s_C(t) \exp(+j2\pi f_c t) \}. \quad (8.4)$$

Here, $j = \sqrt{-1}$, and $\Re \{ \cdot \}$ takes the real part of its complex argument. Sometimes, the \Re operator is dropped for brevity, as actual voltages and fields are necessarily real. A constant scaling factor associated with amplification has been ignored for simplicity. Eq. (8.4) represents the mixing or upconversion operation in the transmit chain of Fig. 8.2.

From Eq. (8.2), the received echo from a discrete, point-like target is a scaled and delayed copy of the transmitted signal. Therefore,

$$s_R(t) = C \Re \{ s_C(t - \tau) \exp(+j2\pi f_c(t - \tau)) \}. \quad (8.5)$$

The radar receiver removes the carrier from the echo $s_R(t)$ by again multiplying with the carrier to obtain a baseband signal $s_B(t)$ before further processing, as illustrated in the receive chain of

Fig. 8.2. Again ignoring constant scaling factors related to amplification, this downconversion operation is expressed as

$$s_B(t) = \Re \{s_R(t) \exp(-j2\pi f_c t)\}. \quad (8.6)$$

Substituting Eq. (8.5) into Eq. (8.6), the baseband signal $s_B(t)$ for the point target is thus

$$s_B(t) = C \Re \{s_C(t - \tau) \exp(+j2\pi f_c(t - \tau)) \exp(-j2\pi f_c t)\} \quad (8.7)$$

$$= C \Re \{s_C(t - \tau) \exp(-j2\pi f_c \tau)\}. \quad (8.8)$$

In other words, the baseband received echo from a discrete target is a delayed copy of the complex envelope of the transmitted signal multiplied by a constant phase term that is proportional to the delay or range of the target. The phase term is related to the number of sinusoidal radar clock cycles that fit into the target delay. Equivalently, the phase is related to the number of wavelengths of the radar carrier that fit into the round-trip travel distance of the signal. Let ϕ be the target phase, which is the phase of the echo from the point-like target. The target phase can be expressed in terms of the target delay τ or range ρ as

$$\phi = -2\pi f_c \tau \quad (8.9)$$

$$= -\frac{4\pi}{\lambda} \rho, \quad (8.10)$$

where the latter is obtained through substitution with Eqs. (8.1) and (8.3).

Unless noted otherwise (see Sect. 8.3.5), phase is assumed to be a wrapped quantity that is defined only over an interval of 2π rad. That is, a phase value ϕ is equivalent to and interchangeable with any other phase value $\phi + 2M\pi$ where M is an integer. For example, a phase value of $-\pi/2$ is the same as phase values of $3\pi/2$ and $103\pi/2$. The phase may be expressed in degrees or cycles as well as radians, but radians are assumed in mathematical expressions throughout this document unless explicitly stated otherwise.

8.1.3 Range Compression and Resolution

The example of Fig. 8.1 assumes a single, discrete, point-like target. That is, the depth of the target in range is assumed to be much smaller than the duration of the transmit signal. This model of a target is useful as a canonical case, but it is seldom very realistic by itself. In Earth remote sensing, for example, the desired target is typically the Earth surface, or at least a portion of it, and this target is obviously not point like. However, distributed targets such as the Earth surface can often be quite well modeled as collections of random point-like targets that are distributed in space. The echo received by the radar is then well approximated by the sum of the contributions from each elemental point target. The individual targets are generally distributed in range (distance from the radar) due to their distribution in 3-D space. Different targets may also have different reflectivities that affect the scaling of their echoes. Let the scattering distribution $s_S(t)$ give the amplitude scaling integrated over the collection of targets as a function of the target delay. The received echo $s_R(t)$ is then the convolution of the transmit signal $s_T(t)$ with the scattering distribution $s_S(t)$ of the target:

$$s_R(t) = C \int s_S(t - \tau) s_T(\tau) d\tau \quad (8.11)$$

$$= C [s_S(t) * s_T(t)]. \quad (8.12)$$




Figure 8.3: Illustration of radar ranging for multiple point targets.

Here, the “*” operator indicates convolution in shorthand notation.

An illustration of the echo from multiple discrete targets is shown in Fig. 8.3. In this example, the target at a delay of $\tau = 6$ is the same as in Fig. 8.1, but there are additional contributions from targets at delays of $\tau = 3$, 3.6, and 4.6. Moreover, the echo contribution from the target at $\tau = 3$ has twice the magnitude of the three other target echoes. The individual contributions from the targets are shown as thin colored lines. Because the duration of the transmit pulse is 1 in the relative units of the plot, the targets at delays of $\tau = 3$ and 3.6 overlap one another. The trough of the echo for the target at $\tau = 3$ and the crest of echo for the target at $\tau = 3.6$ therefore interfere with each other and partially cancel.

As is evident from Fig. 8.3, the ability to distinguish features in the composite echo depends on the width of the point target response (PTR), which is the response to or observation by the radar of a given point target. For the received echo, the PTR is simply the transmit waveform per Eq. (8.2). It is difficult to resolve variations in range that are smaller than width of the PTR.

In order to enhance the ability of the radar to resolve features at finer scales in range, one might therefore like to reduce the duration of the transmit signal and make the transmit waveform as close to an infinitely narrow impulse as possible. However, such a reduction of the transmit pulse width reduces the radiated energy and increases the bandwidth of the signal. The increase in the bandwidth of the signal increases the bandwidth of the noise given that the filter-

ing of the radar signal is generally matched to the signal bandwidth. (The effects of filtering are ignored for simplicity in Figs. 8.1 and 8.3). Consequently, a reduction in the transmit pulse width gives an unattractive reduction in the signal to noise ratio (SNR).

An alternative that enhances the resolution of the radar without reducing the radiated energy is to transmit a chirp waveform. Suppose that the transmit waveform has a complex envelope $s_C(t)$:

$$s_C(t) = \begin{cases} \exp(j\pi a_s t^2) & \text{if } -\tau_p/2 < t < \tau_p/2 \\ 0 & \text{otherwise} \end{cases}. \quad (8.13)$$

where τ_p is the pulse duration or pulse width, and a_s is a parameter called the chirp slope. The phase of this waveform varies quadratically with time t . Equivalently, this waveform has a sinusoidal variation whose frequency increases linearly with time. This type of waveform is therefore sometimes also called a (gated) linear frequency modulated (LFM) waveform. Note that the instantaneous frequency f of the chirp is defined with respect to the time derivative of the phase as

$$f = \frac{1}{2\pi} \frac{d}{dt} \pi a_s t^2 \quad (8.14)$$

$$= a_s t. \quad (8.15)$$

The bandwidth B_p of the chirp is thus approximately equal to $a_s \tau_p$ (this assumes that $a_s \tau_p^2 \gg 1$ so that edge effects are negligible).

The PTR of a chirp signal by itself in the received echo is not necessarily narrow, so the chirp by itself does not offer particularly fine resolution. That is, when the reflectivity distribution of the target is convolved with the chirp waveform via Eq. (8.12), features of the reflectivity distribution will be smeared out over the extent τ_p of the chirp.

However, the autocorrelation of a chirp has a sharp peak, so after matched filtering, the PTR becomes highly peaked, and fine resolution is obtained. That is, matched filtering, which is known to optimize the SNR, is equivalent to cross correlating with the transmit waveform. Convolution and cross correlation are linear operators, so when the received echo (spread by the chirp) is cross correlated with a copy of the chirp, one obtains the convolution of the reflectivity distribution with the autocorrelation of the chirp. Let $s_F(t)$ be the received echo after downconversion and matched filtering. Ignoring upconversion and downconversion with the carrier for simplicity [i.e., assuming $f_c = 0$ so $s_B(t) = s_R(t)$ and $s_T(t) = s_C(t)$], the above is expressed mathematically as

$$s_F(t) = s_B(t) \star s_C(t) \quad (8.16)$$

$$= s_R(t) \star s_C(t) \quad (8.17)$$

$$= [s_S(t) * s_T(t)] \star s_C(t) \quad (8.18)$$

$$= [s_S(t) * s_C(t)] \star s_C(t) \quad (8.19)$$

$$= s_S(t) * [s_C(t) \star s_C(t)], \quad (8.20)$$

where the \star operator denotes cross correlation (also often written as simply “correlation”). The PTR $h(t)$ of the radar measurement system after matched filtering (which may occur during

ground processing) is the response function that satisfies

$$s_F(t) = s_S(t) * h(t), \quad (8.21)$$

so, from Eq. (8.20), the PTR $h(t)$ is the autocorrelation of the chirp:

$$h(t) = s_C(t) \star s_C(t) \quad (8.22)$$

$$= \int s_C^*(\tau) s_C(t + \tau) d\tau. \quad (8.23)$$

Here, the superscript $*$ indicates complex conjugation.

When the radar carrier is not ignored, the target phase from Eq. (8.10) enters into Eq. (8.21) as

$$s_F(t) = [s_S(t) \exp(-j2\pi f_c t)] * h(t) \quad (8.24)$$

$$= \left[s_S(t) \exp\left(-j\frac{4\pi}{\lambda}\rho\right) \right] * h(t). \quad (8.25)$$

The PTR in general may be complex. Its power is $|h(t)|^2$.

An illustration of a chirp and its autocorrelation are shown in Fig. 8.4. Because the PTR becomes much sharper after the receive echo is matched filtered when a chirp waveform is used, the matched filtering operation in range is often called “range compression” (this term is not related to data compression; range compression does not necessarily reduce the volume of the echo data). The chirp in Fig. 8.4 is intended for illustration purposes; typically, the chirp slope a_s would be larger, giving a higher bandwidth B_p and a commensurately narrower PTR peak relative to the pulse width τ_p .

The spectrum of a chirp waveform has a magnitude that is approximately rectangular in the frequency domain with a bandwidth $B_p = a_s \tau_p$. The Fourier transform of a rectangular function is a sinc function, so the autocorrelation is well approximated by

$$h(t) = \text{sinc}(tB) \quad (8.26)$$

$$= \frac{\sin(\pi tB)}{\pi tB}. \quad (8.27)$$

The PTR $h(t)$ therefore has a mainlobe whose full width at half power in time is approximately $1/B$ (more precisely, $0.8859/B$ for an ideal sinc function). The first sidelobes of the PTR are approximately 13 dB below the peak in power, where the power is proportional to $|h(t)|^2$. If the range resolution δ_ρ is defined approximately as the half-power width of the PTR in range, then from Eq. (8.3),

$$\delta_\rho = \frac{c}{2B}. \quad (8.28)$$

Thus, after range compression, the range resolution δ_ρ depends only on the pulse bandwidth B , not the pulse width τ_p or the range ρ .




Figure 8.4: Illustration of a notional chirp (a) and its autocorrelation (b). The real and imaginary components and the magnitude of the chirp complex envelope are shown in (a).

8.1.4 Echo Power

The received power from a discrete target is given by the so-called radar equation:

$$P_R = P_T \frac{G_T}{4\pi\rho^2} \sigma_{\text{RCS}} \frac{A_R}{4\pi\rho^2} \eta_{\text{loss}}. \quad (8.29)$$

where P_R is the received power (implicitly a function of time), P_T is the radiated transmit power, G_T is the dimensionless power gain of the radar transmit antenna in the direction of the target, σ_{RCS} is the target radar cross section (RCS), A_R is the effective capture area of the radar receive antenna in the direction of the target, and η_{loss} is a dimensionless loss term. The term $G_T/(4\pi\rho^2)$ represents the dissipation of the transmit signal power through spherical spreading as it propagates from the radar to the target. Similarly, the term $A_R/(4\pi\rho^2)$ represents the dissipation of the echo power as it propagates from the target back to the radar. The term η_{loss} captures losses such as within the radar hardware or due to atmospheric attenuation. Note that there are many alternate forms of the radar equation.

The effective capture area of the receive antenna is related to the receive antenna gain G_R (in a given direction) by

$$G_R = \frac{4\pi}{\lambda^2} A_R. \quad (8.30)$$

The RCS σ_{RCS} of a discrete target is defined as the projected area of a hypothetical, perfectly conducting (i.e., completely reflective of radio waves) sphere that scatters the same power as the target does in the direction of the radar. That is, the RCS is a measure of the target reflectivity. The power of the echo from the target is directly proportional to the RCS.

For a distributed target such as the Earth surface, the scattering is characterized by the normalized radar cross section σ° . This quantity is sometimes abbreviated as NRCS or written as “sigma0”. It is usually pronounced as either “sigma nought” or “sigma zero.” This dimensionless quantity is defined as the RCS per unit area on the surface:

$$\sigma^\circ = \frac{\sigma_{\text{RCS}}}{A_S}. \quad (8.31)$$

Here, A_S is the area on the surface over which the scattered energy gives an RCS σ_{RCS} .

From Eq. (8.29), the NRCS can be related to the received echo power by evaluating the following:

$$\sigma^\circ = \frac{(4\pi\rho^2)^2}{P_T G_T A_R \eta_{\text{loss}} A_S} P_R. \quad (8.32)$$

The terms on the right-hand side of this equation other than P_R are often collected into a parameter called the X factor with

$$X = \frac{P_T G_T A_R \eta_{\text{loss}} A_S}{(4\pi\rho^2)^2} \quad (8.33)$$

such that

$$\sigma^\circ = \frac{P_R}{X}. \quad (8.34)$$

Note that an estimate of the power of any additive noise (mostly from the radar receiver) may need to be subtracted from the data in order to obtain an estimate of the power P_R due to the echo, exclusive of noise. The noise estimate would be subtracted before scaling by the X factor to obtain the estimate of σ° .

8.2 SAR Imaging

A traditional altimeter looks downward and measures the distance to the surface as a 1-D profile over the nadir track. By contrast, a traditional synthetic aperture radar (SAR) looks to the side, and, instead of measuring the range to the surface, the SAR measures the echo power as a function of range to form a 2-D image over a swath to the left or to the right of nadir. This section describes the basics of SAR image formation. Additional details can be found in [43].

8.2.1 Side-Looking Radar Imaging

Suppose that a radar flying above the ground has its antenna pointed to the left or to the right so that echoes from only one side of the nadir track are recorded, as illustrated in Fig. 8.5. In this figure, illustration (a) shows the swath that is imaged by the antenna footprint as the radar flies over a set of arbitrary ground targets, which are shown as round blue dots. Illustration (b) shows the three discrete targets that are within the antenna footprint in (a) as viewed with the radar velocity vector pointed into the page. Illustration (b) also shows the received echo power after range compression when those three targets are in the antenna footprint. Noise is neglected for simplicity. There are peaks in the echo profile at delays corresponding to the ranges of the three targets. The magnitude of each peak is proportional to the reflectivity of the corresponding target. Illustration (c) shows the 2-D radar power image obtained by stacking a series 1-D echo profiles together. The two dimensions of the image are slant range and along track; the brightness of the image indicates the backscattered power over the swath.

The term “slant range” refers to the fact that the coordinates of 2-D, side-looking radar images are not natural ground coordinates. Rather, points on the ground are mapped into the image based on their distance from the radar. The resulting image is thus known as a “slant-plane” image. Note that the mapping between ground coordinates and slant-plane coordinates is not linear, so points that are evenly spaced on the ground in cross track are not uniformly spaced in range in the slant-plane image. If the surface is flat and the radar altitude above the ground surface is known, however, then the slant-plane image can be easily mapped back to a horizontal projection in order to obtain a “ground-plane” image that may be easier to interpret.

8.2.1.1 Side-Looking Radar Terminology

A number of special terms are often used in association with side-looking radars. These are summarized here:

- The *look vector*, denoted \vec{l} is the vector from the radar antenna phase center to a given target. The length of the look vector is the range ($|\vec{l}| = \rho$).
- The *azimuth* direction is another term for the along-track direction. This term is often used in conjunction with the range direction in slant-plane coordinates.
- The *look angle* (or sometimes the *elevation* angle), denoted θ_{look} , is the angle of the look vector with respect to the nadir direction at the radar.
- The *incidence angle*, denoted θ_{inc} is the angle between the opposite of the look vector and the local up direction at the target (the angle is typically less than 90°). The incidence

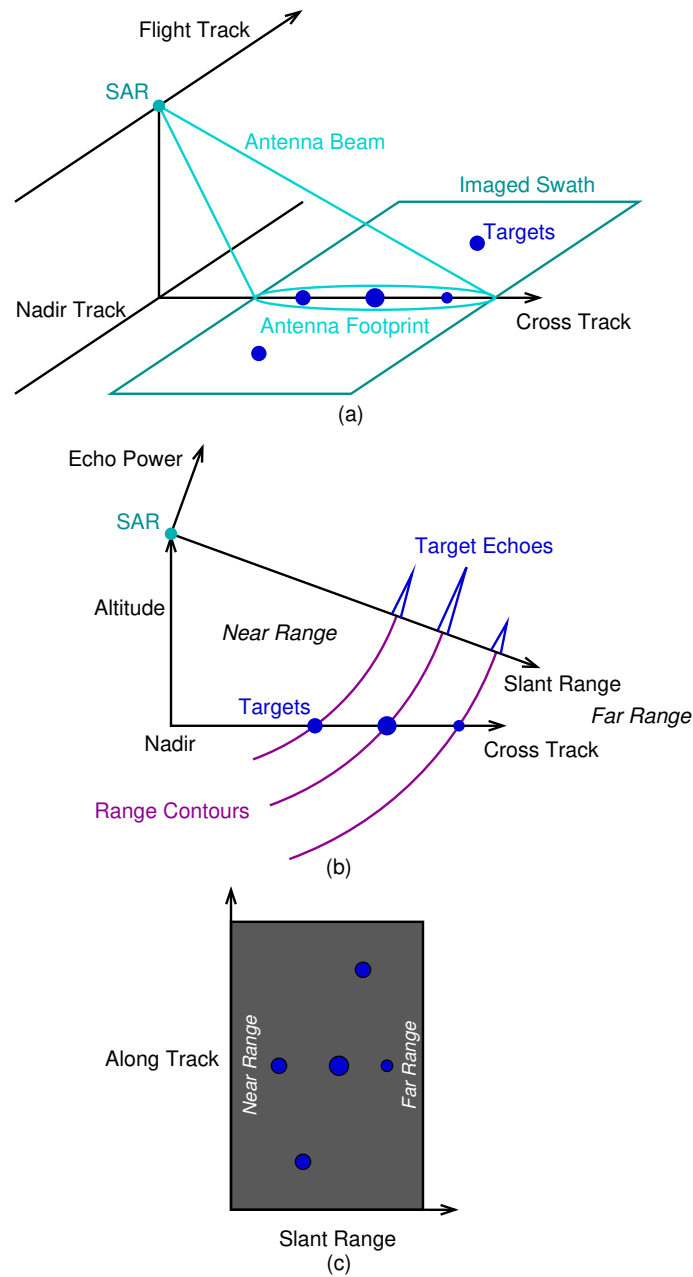


Figure 8.5: Side-looking radar imaging geometry: (a) illustration of the swath imaged as the antenna footprint passes over the ground; (b) a single example profile of the received echo power as a function of slant range; (c) power image formed from many side-looking echoes, with blue indicating higher power than gray.

angle and the look angle are equivalent when considering a flat (planar) Earth surface, but the angles may differ for spaceborne viewing geometries where the curvature of the planet cannot be neglected. In some contexts, the incidence angle refers to the acute angle between the look vector and the local surface normal rather than the local up direction (these differ when the surface slope is nonzero); this is typically clarified when the term used in such a manner.

- A *broadside* viewing geometry is one in which the radar antenna is pointed directly to the side such that targets are observed with look vectors that are exactly perpendicular to the nadir track.
- A *squinted* viewing geometry is one in which the radar antenna is pointed somewhat forward or backward relative to broadside. The squint angle is the angle of the look vector away from a pure broadside viewing geometry.
- *Near range* is the side of an image that is closest to the radar (with the shortest range). This is typically also closest to the nadir track.
- *Far range* is the side of an image that is farthest from the radar (with the longest range). This is typically also farthest from the nadir track.
- *Fast time* refers to the time delay within a given pulse echo such that fast time is associated with the range dimension of a slant-plane image.
- *Slow time* refers to the time over which the radar travels along the flight track such that slow time is associated with the azimuth or along-track dimension of a slant-plane image.

8.2.1.2 Planet Curvature

From spaceborne altitudes, the curvature of the planet is usually not negligible for a side-looking radar. Fig. 8.6 illustrates the definitions of several quantities with respect to a spherical-Earth model.

The planet radius R_E , look angle θ_{look} , incidence angle θ_{inc} , radar altitude h_s , range ρ , and Earth-central angle α are related to one another through the law of sines:

$$\frac{\sin \theta_{\text{look}}}{R_E} = \frac{\sin \theta_{\text{inc}}}{R_E + h_s} = \frac{\sin \alpha}{\rho}. \quad (8.35)$$

Similarly, from the law of cosines,

$$\rho^2 = R_E^2 + (R_E + h_s)^2 - 2R_E(R_E + h_s)\cos(\alpha). \quad (8.36)$$

Finally, it may be noted that

$$\theta_{\text{inc}} = \theta_{\text{look}} + \alpha. \quad (8.37)$$

The expressions above are most applicable if R_E is taken to be the local radius of curvature of the planet. Simply using the equatorial Earth radius (see Table 3.1) everywhere may be sufficient for many purposes, however.




Figure 8.6: Imaging geometry of a side-looking radar assuming a spherical Earth model. The radar velocity is into the page.

These expressions are often helpful in relating geometric quantities of interest in spaceborne SAR analysis. However, much of the description in the remainder of this chapter is tutorial in nature and assumes a flat-Earth geometry for simplicity. The relationships above are provided mainly for reference.

8.2.1.3 Layover

If the imaged surface is not flat, the mapping between the ground plane and the slant plane may not be unique. This gives rise to a phenomenon called layover, which is illustrated in Fig. 8.7. The panels in Fig. 8.7 are similar to those of Fig. 8.5, but Fig. 8.7 depicts a mountain being imaged. Because the peak of the mountain is higher in elevation than the base, it is closer to the radar and therefore appears at a nearer range than the base. Consequently, in the slant-plane radar image, the mountain appears “laid over” toward the near-range side of the image and is therefore said to be in layover. Layover occurs when the cross-track component of the local surface slope exceeds the tangent of the incidence angle θ_{inc} with respect to a flat surface. Layover is therefore typically more problematic for radars with steeper viewing geometries.

When layover occurs, multiple points on the surface have the same range, and there is no longer a one-to-one mapping between cross track and range. That is, echoes from multiple locations on the surface, such as on opposite sides of the mountain in Fig. 8.7(b), arrive at the radar with the same delay. They therefore appear in the same range bin. Without additional information, it is not possible to distinguish which echo contributions arose from which parts of the surface. Layover therefore makes interpretation of the radar images more difficult (see Sect 10.2




Figure 8.7: Layover: (a) illustration of a mountain imaged by a side-looking radar; (b) profile of the mountain as viewed with the radar velocity into the page; (c) slant-plane power image of the mountain with the peak of the mountain appearing nearer in range than the base of the mountain such that the mountain appears “laid over.”

for a discussion of layover that is specific to SWOT).

Foreshortening is a related phenomenon in which topographic variations cause distortions in the slant-plane images relative to the horizontal distribution of targets in the ground plane. Even when the surface cross-track slope is not so steep as to cause layover (no many-to-one mapping from the ground plane to the slant plane), slopes that face toward the radar appear “shortened” in a slant-plane image in that they occupy fewer range bins than would be expected based on their horizontal extent in the ground plane. Conversely, slopes that face away from the radar appear elongated in the slant plane. Foreshortening and layover are direct manifestations of the manner in which a side-looking radar maps surface targets in 3-D space into 2-D range-azimuth images in the slant plane.

8.2.1.4 Resolution of Real Aperture Radar

For a broadside viewing geometry, the ground-projected, cross-track resolution δ_y of the image is related to the range resolution δ_ρ through

$$\delta_y = \frac{\delta_\rho}{\sin \theta_{\text{inc}}}, \quad (8.38)$$

where θ_{inc} is the incidence angle of the signal with respect to the local up direction.

The along-track resolution δ_{xr} of the simple side-looking radar described in this section would be the width of the antenna footprint on the ground. This type of radar is known as a real-aperture radar (RAR). The azimuth beamwidth θ_{Baz} of an aperture antenna whose along-track length is L_x is approximately

$$\theta_{\text{Baz}} = \frac{\lambda}{L_x}, \quad (8.39)$$

so the real-aperture resolution δ_{xr} and the azimuth footprint width D_{az} are

$$\delta_{xr} = D_{\text{az}} \quad (8.40)$$

$$= \theta_{\text{Baz}} \rho \quad (8.41)$$

$$= \frac{\lambda}{L_x} \rho, \quad (8.42)$$

where ρ is the slant range. The along-track resolution of such a radar is therefore rather coarse when the range ρ is large, as it is from space. This motivates synthetic aperture techniques to obtain finer resolution, as described in the next section.

8.2.2 Azimuth Compression

Radar instruments are coherent (they measure both amplitude and phase), so aperture synthesis techniques can be used to obtain much finer azimuth resolution than the real-aperture resolution described in the previous section. Akin to range compression, azimuth compression comprises the operations described in this section to achieve fine resolution in azimuth.

8.2.2.1 Doppler Centroid, Doppler Bandwidth, and Phase History

When the relative velocity between the radar and a given target is not zero, the echo of the transmit signal from the target experiences a Doppler shift. That is, the echo is shifted in frequency relative to the signal carrier frequency f_c by a Doppler frequency f_D given by

$$f_D = \frac{2\vec{v}_r \cdot \hat{l}}{\lambda}, \quad (8.43)$$

where \vec{v}_r is the relative velocity vector between the radar and the target and \hat{l} is the unit look vector from the radar to the target. That is,

$$\vec{v}_r = \vec{v}_s - \vec{v}_t, \quad (8.44)$$

and

$$\hat{l} = \frac{\vec{l}}{\rho}, \quad (8.45)$$

where \vec{v}_s and \vec{v}_t are the velocity vectors of the radar and the target, and the look vector \vec{l} is defined in Sect. 8.2.1.1. The dot operator (“ \cdot ”) in Eq. (8.43) represents the vector inner or scalar product, so $\vec{v}_r \cdot \hat{l}$ is the range rate, which is the rate of change of the range to the target over time. The factor of 2 in Eq. (8.43) arises because of the round-trip travel of the signal from the radar to the target and back to the radar again. As targets on the Earth surface are usually assumed to be stationary in an Earth-fixed frame, \vec{v}_r is usually just the radar velocity \vec{v}_s in the Earth-fixed frame.

For an ideal, broadside SAR geometry, where the radar flight path is level and the antenna is pointed directly to the side, the look vector at the antenna boresight is perpendicular to the velocity vector, so the Doppler frequency f_D is zero. The antenna has a finite azimuth beamwidth [see Eq. (8.39)], however, so it observes a span of angles $\pm\theta_{Baz}/2$ around its boresight. The look vectors over this beamwidth are therefore associated with Doppler frequencies from Eq. (8.43) that satisfy

$$-\frac{2v_s \sin(\theta_{Baz}/2)}{\lambda} \leq f_D \leq \frac{2v_s \sin(\theta_{Baz}/2)}{\lambda}, \quad (8.46)$$

where v_s is the magnitude of the velocity vector \vec{v}_s . The surface is assumed to be stationary ($\vec{v}_r = \vec{v}_s$). The Doppler bandwidth B_{Dop} is thus given by

$$B_{Dop} = \frac{4v_s \sin(\theta_{Baz}/2)}{\lambda} \quad (8.47)$$

$$\approx \frac{2v_s \theta_{Baz}}{\lambda}. \quad (8.48)$$

If the antenna is not pointed perfectly to broadside and/or if the radar velocity vector has a vertical component, then there will be a mean Doppler shift called the Doppler centroid f_{DC} , given by

$$f_{DC} = \frac{2\vec{v}_r \cdot \hat{l}_0}{\lambda}, \quad (8.49)$$

where \hat{l}_0 is the unit look vector at the antenna boresight.




Figure 8.8: Illustration of the range $\rho(x)$ to a target as a function of the radar position x along the flight track as the antenna footprint sweeps over the ground. Each target on the ground contributes to many pulse echoes along the flight track. Conversely, each pulse echo along the flight track contains contributions from many targets on the ground.

For SAR systems, the magnitude of the Doppler shift f_D is typically very small compared to the pulse bandwidth B_p , which is in turn small compared to the radar carrier frequency f_c . The effect of the Doppler shift when considering a single pulse is therefore negligible for a high-level discussion such as given in this chapter.

However, a given target on the ground may contribute to many successive pulse echoes as the radar flies over and as the antenna footprint sweeps over the surface. As illustrated in Fig. 8.8, let $\rho(x)$ be the range from the radar to a target on the ground as a function of the radar position x along the flight track. For simplicity, let x be zero when the target is at broadside, and assume that the flight track is straight and level at an altitude h_s over a flat Earth surface. The cross-track distance to the target is y_t . The range ρ_0 to the target at broadside is thus

$$\rho_0 = \sqrt{h_s^2 + y_t^2}, \quad (8.50)$$

and the range $\rho(x)$ as the radar moves along the flight track is

$$\rho(x) = \sqrt{h_s^2 + y_t^2 + x^2} \quad (8.51)$$

$$= \sqrt{\rho_0^2 + x^2}. \quad (8.52)$$

This relationship between the range to the target and the position along the flight track is called the range history of the target. The range history can also be expressed in terms of slow time (see Sect. 8.2.1.1) by expressing the position x as a function of time t as

$$\rho(t) = \sqrt{\rho_0^2 + (v_s t)^2}, \quad (8.53)$$

where v_s is the radar velocity.

The echo phase is related to the range through Eq. (8.10), so the phase history of the target (i.e., the phase as a function of radar position along the flight track) is

$$\phi(x) = -\frac{4\pi}{\lambda} \sqrt{\rho_0^2 + x^2}. \quad (8.54)$$

Equivalently, the phase history as a function of time is

$$\phi(t) = -\frac{4\pi}{\lambda} \sqrt{\rho_0^2 + (v_s t)^2}. \quad (8.55)$$

It is instructive to observe that the Doppler frequency f_D is related to the time derivative of the echo phase ϕ through

$$f_D = \frac{1}{2\pi} \frac{d}{dt} \phi(t), \quad (8.56)$$

so f_D can be derived by differentiating the phase history from Eq. (8.55) with respect to time:

$$f_D = \frac{1}{2\pi} \left(-\frac{4\pi}{\lambda} \frac{1}{2} \frac{2v_s^2 t}{\sqrt{\rho_0^2 + (v_s t)^2}} \right) \quad (8.57)$$

$$= -\frac{2v_s}{\lambda} \frac{v_s t}{\sqrt{\rho_0^2 + (v_s t)^2}} \quad (8.58)$$

$$= -\frac{2v_s}{\lambda} \frac{x}{\rho} \quad (8.59)$$

$$= \frac{2v_s}{\lambda} \sin \theta_{az}. \quad (8.60)$$

Here, θ_{az} is the azimuth angle of the look vector, which is positive for look vectors pointing forward of broadside. When x is positive, the target is behind the radar, giving a negative value of θ_{az} , so $\sin \theta_{az} = -x/\rho$. Therefore, Eq. (8.60) is equivalent to Eq. (8.43) for the straight, level trajectory assumed by Eq. (8.55).

The conclusion is that the phase history of the target contains the information of the target Doppler, as well as the variation in the target Doppler, over the radar flight track. The target Doppler—or, equivalently, the phase history—can be exploited to obtain much finer azimuth resolution than the antenna footprint.

8.2.2.2 Unfocused SAR Processing

The discussion of the phase history in Sect. 8.2.2.1 considers a single point target on the ground as observed over many pulse echoes along the radar flight track. However, a natural surface contains a vast number of individual targets that are distributed over the ground, and each of those targets contributes to many pulse echoes along the flight track. The desire to better distinguish different targets on the ground is of course what motivates the use of SAR processing to achieve fine resolution in azimuth.

Consider the scene illustrated in Fig. 8.9. The antenna footprint moves slightly from one pulse to the next, but over a limited number N_B of successive pulses (sometimes called a burst), the antenna footprint does not move very much compared to the overall along-track width of the footprint. Each echo contains contributions from all the targets on the ground within the antenna footprint, so the targets cannot be resolved in azimuth when considering any single echo. However, targets farther ahead of the radar have more positive Doppler frequencies, while targets farther behind the radar have more negative Doppler frequencies. Consequently, a simple way to

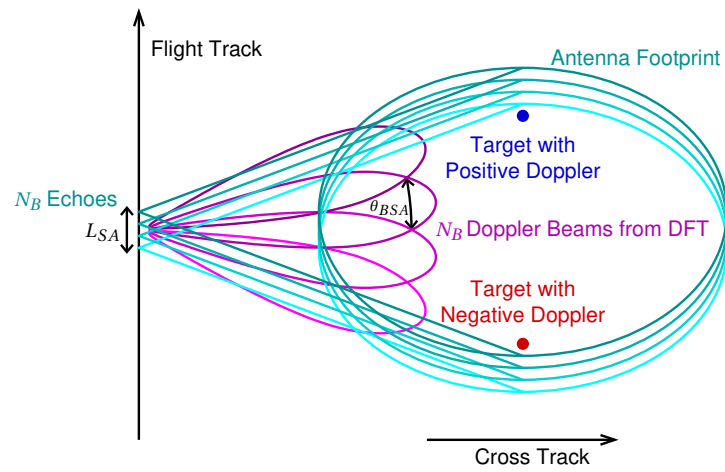


Figure 8.9: Illustration of successive pulse echoes whose Doppler can be used to resolve ground targets through unfocused SAR processing. Targets farther ahead of the radar have larger Doppler frequencies and vice versa (notwithstanding aliasing due to the sampling of the Doppler by the PRF). The N_B bins of the azimuth DFT over the N_B pulse echoes give the N_B magenta Doppler beams that are sensitive to targets at different Doppler frequencies or azimuth angles. The magenta lobe of each Doppler beam illustrates the sensitivity of that beam to targets as a function of the azimuth angle. Each cyan ellipse represents the footprint of the physical antenna on the ground for one of the N_B pulses.

obtain finer azimuth resolution is to compute the discrete Fourier transform (DFT) in azimuth over the N_B pulse echoes for each range bin after range compression. The frequency-domain result of the Fourier transform provides N_B Doppler bins, which are equivalent to angular bins through Eq. (8.60). Consequently, each Doppler bin from the DFT can be considered to come from a Doppler-synthesized antenna “beam” that is sensitive primarily to targets around the azimuth angle corresponding to the Doppler frequency of that DFT bin. The echo power in each of these Doppler bins can be mapped onto the ground with knowledge of the altitude and the velocity vector. The DFT may be implemented as a fast Fourier transform (FFT) for computational efficiency.

The DFT of each successive set of N_B echoes in along track provides another independent set of measurements that is slightly shifted in along track. The results from different sets of N_B echoes can be averaged together (possibly after resampling to align the measurements on the ground) to reduce noise. This technique is called unfocused SAR processing (by contrast, focused SAR processing is discussed below).

The set of targets within the antenna footprint span a Doppler bandwidth given by Eq. (8.48), and this bandwidth is sampled at the pulse repetition frequency (PRF) by the successive pulse echoes captured by the radar as it flies over. Note that the echoes give complex rather than real samples, so a PRF that is equal to the Doppler bandwidth B_{Dop} adequately (Nyquist) samples the Doppler spectrum over the antenna footprint. It is not necessary for the PRF to be twice Doppler bandwidth due to the complex sampling.

The distance Δx that the antenna footprint moves from one pulse to the next is the product of the radar speed v_s and the time T_p between pulses, which is called the pulse repetition interval (PRI). The PRI is the inverse of the PRF. Thus,

$$\Delta x = v_s T_p. \quad (8.61)$$

Analogous to Eq. (8.48), the azimuth angle θ_{Paz} that is sampled by the PRF can be expressed as

$$\text{PRF} = \frac{2v_s \theta_{\text{Paz}}}{\lambda}. \quad (8.62)$$

Assume that N_B is small enough that the distance $N_B \Delta x$ traveled by the radar does not appreciably change the viewing angles to the target. A DFT over N_B echoes therefore gives an azimuth angular resolution of θ_{Paz}/N_B . This is the beamwidth of a single Doppler beam from the unfocused SAR processing. The angular resolution can be multiplied by the range to obtain the azimuth resolution δ_{xu} of the unfocused SAR on the ground:

$$\delta_{xu} = \frac{\theta_{\text{Paz}}}{N_B} \rho_0 \quad (8.63)$$

$$= \frac{\lambda \text{PRF}}{2N_B v_s} \rho_0 \quad (8.64)$$

$$= \frac{\lambda}{2N_B v_s T_p} \rho_0. \quad (8.65)$$

If the PRF is approximately equal to the Doppler bandwidth B_{Dop} (i.e., critical sampling), then θ_{Paz} is approximately equal to θ_{Baz} , and, comparing Eqs. (8.41) and (8.63), the unfocused SAR resolution is about N_B times finer than the real-aperture resolution.

The unfocused SAR resolution can also be derived from an aperture synthesis perspective. With N_B pulses, the unfocused SAR forms a synthetic antenna aperture along the flight track whose length L_u is

$$L_u = N_B \Delta x. \quad (8.66)$$

Analogous to Eq. (8.39), the effective beamwidth θ_{BSA} of a synthetic aperture is

$$\theta_{BSA} = \frac{\lambda/2}{L_{SA}}, \quad (8.67)$$

where L_{SA} is the length of the synthetic aperture, and the factor of 2 arises because of the round-trip travel of the radar signal. Using Eq. (8.66) with $L_{SA} = L_u$, the unfocused SAR resolution is expressed as

$$\delta_{xu} = \theta_{BSA} \rho_0 \quad (8.68)$$

$$= \frac{\lambda/2}{N_B \Delta x} \rho_0 \quad (8.69)$$

$$= \frac{\lambda}{2N_B v_s T_p} \rho_0, \quad (8.70)$$

which yields the same result as Eq. (8.65). The conclusion is that extending the synthetic aperture length L_u enhances the azimuth resolution proportionately.

For the unfocused SAR approach described above, however, the aperture length L_u cannot be so long that the targets move from one Doppler bin to the next over the length of the aperture. That is, if the radar moves along the flight track by more than the azimuth resolution δ_{xu} , the targets will be smeared across multiple resolution cells. Therefore, the maximum synthetic aperture length is $L_u \leq \delta_{xu}$. Substituting this into Eq. (8.69) using Eq. (8.66), the finest possible unfocused SAR resolution δ_{xumin} , which is the same as the longest useful unfocused SAR aperture L_{umax} , is

$$\delta_{xumin} = L_{umax} \quad (8.71)$$

$$= \sqrt{\frac{\lambda \rho_0}{2}}. \quad (8.72)$$

As described above [see Eqs. (8.52), (8.54), and (8.60)], the Doppler arises from the change in target phase from pulse to pulse. This phase variation is a direct consequence of the range variation from pulse to pulse at the scale of the radar wavelength. However, unfocused SAR processing assumes that the range to the target does not change so much over the synthetic aperture that the target shifts from one range resolution cell to the next. Such a shift would also limit the azimuth resolution of unfocused SAR processing.

8.2.2.3 Focused SAR Processing

Focused SAR processing overcomes the limitations of unfocused SAR processing in order to achieve considerable improvements in azimuth resolution, albeit at the expense of greater processing complexity. The fundamental idea of focused SAR processing is to account for the full variation in the Doppler phase (and range) of a target so that the length of the synthetic aperture is limited only by the illumination of the target by the antenna as illustrated in Fig. 8.10.




Figure 8.10: Illustration of focused SAR. Resolution is obtained by coherently processing a synthetic aperture from the full span in along track over which the target is within the antenna beam. Targets at different locations on the ground are associated with different synthetic apertures.

Consider the phase history of a target from Eq. (8.54), which can be rewritten as

$$\phi(x) = -\frac{4\pi\rho_0}{\lambda} \sqrt{1 + \frac{x^2}{\rho_0^2}}. \quad (8.73)$$

The azimuth antenna beamwidth θ_{Baz} of a SAR is usually relatively narrow, so x is usually much smaller than ρ_0 , and the square root in Eq. (8.73) can be approximated by a binomial expansion, giving

$$\phi(x) \approx -\frac{4\pi\rho_0}{\lambda} \left(1 + \frac{x^2}{2\rho_0^2} \right) \quad (8.74)$$

$$\approx -\frac{4\pi}{\lambda} \rho_0 - \frac{2\pi x^2}{\lambda \rho_0}. \quad (8.75)$$

The first term in Eq. (8.75) that is independent of x is simply the range-dependent phase of the target from Eq. (8.10). The second term describes a quadratic variation in the target phase as a function of the radar position x along the flight track. Therefore, the phase history of the target behaves very much like a chirp waveform [see Eq. (8.13)]. That is, the signature of the target in slow time (over the many pulse echoes for which the target is in the antenna footprint) is analogous to the signature of the target in fast time (over the many delay samples of each pulse echo) before range compression.

Suppose for simplicity that the range variation of a target over the synthetic aperture is less than the range resolution. The 2-D SAR data after range compression but before azimuth compression might look like the convolution of the target distribution on the ground with a set of chirp-like waveforms in azimuth. These azimuth chirps would extend over a series of pulses around $\pm D_{az}/2$ of the target location; the antenna footprint width D_{az} corresponds to the pulse width τ_p in the analogy between azimuth compression and range compression. The azimuth chirps would have a chirp slope of $-2/(\lambda\rho_0)$ with respect to the radar along-track position x or $-2v_s/(\lambda\rho_0)$ with respect to the (slow) time t .

Consequently, focused SAR processing might entail a matched filter operation in azimuth that is implemented in a manner similar to that described in Sect. 8.1.3 for range compression. Namely, after the range compression of the raw 2-D SAR data, the set of samples over many pulse echoes from along-track locations x in a given range bin corresponding to ρ_0 would be correlated (in the azimuth dimension) with an azimuth chirp that corresponds to the expected phase history for targets at that range. From the second term of Eq. (8.75), this azimuth reference function would be defined as

$$r_a(x) = \begin{cases} \exp\left(-j \frac{2\pi}{\lambda \rho_0} x^2\right) & \text{if } -D_{az}/2 < x < D_{az}/2 \\ 0 & \text{otherwise} \end{cases}. \quad (8.76)$$

The azimuth chirp would have different parameters than the chirp of the transmit pulse, but it would have qualitatively similar features to the example shown in Fig. 8.4.

The resolution δ_{xf} of a focused SAR can be derived from the synthetic aperture length L_{SA} as described above for an unfocused SAR. That is, recalling Eqs. (8.67) and (8.68),

$$\delta_{xf} = \frac{\lambda/2}{L_{SA}} \rho_0. \quad (8.77)$$

The synthetic aperture length L_{SA} subtends an angle θ_{SA} from the target called the processing beamwidth:

$$L_{SA} \approx \theta_{SA} \rho_0. \quad (8.78)$$

Note that Eqs. (8.77) and (8.78) can be used to compute the resolution even when Earth and orbit curvature are significant.

For the simple case of a straight flight track over a flat Earth surface, however, the synthetic aperture length L_{SA} is equal to the antenna footprint width D_{az} , and the processing beamwidth can be chosen to equal the antenna beamwidth θ_{Baz} . Therefore,

$$\delta_{xf} = \frac{\lambda/2}{D_{az}} \rho_0. \quad (8.79)$$

Substituting the relationships between the antenna footprint width D_{az} , the antenna beamwidth θ_{Baz} , and the antenna along-track length L_x from Eqs. (8.40), (8.41), and (8.42), the focused SAR azimuth resolution δ_{xf} can be expressed only in terms of L_x :

$$\delta_{xf} = \frac{\lambda/2}{\theta_{Baz} \rho_0} \rho_0 \quad (8.80)$$

$$= \frac{\lambda/2}{\frac{\lambda}{L_x} \rho_0} \rho_0 \quad (8.81)$$

$$= \frac{L_x}{2}. \quad (8.82)$$

This result is often counterintuitive: It indicates that the azimuth resolution of a focused SAR does not degrade with range and actually becomes finer as the antenna gets *smaller*. This can be understood by considering that the resolution becomes finer as the synthetic aperture length L_{SA} grows, and the synthetic aperture length increases with the antenna beamwidth θ_{Baz}

because the target is then within the antenna footprint for longer. The beamwidth is inversely proportional to the antenna length L_x (see Fig. 8.10). Similarly, the synthetic aperture length grows as the range increases, thereby offsetting the degradation of the azimuth resolution with range that would be exhibited for if the angular resolution were fixed.

Note that the focused-SAR azimuth resolution δ_{xf} can also be derived from the Doppler bandwidth B_{Dop} by noting that the focused-SAR resolution δ_{tf} in slow time is

$$\delta_{tf} = \frac{1}{B_{\text{Dop}}}, \quad (8.83)$$

and the spatial resolution δ_{xf} is related to the temporal resolution δ_{tf} through radar velocity v_s :

$$\delta_{xf} = v_s \delta_{tf}. \quad (8.84)$$

Therefore, substituting Eqs. (8.39), (8.48), and (8.83) into Eq. (8.84),

$$\delta_{xf} = \frac{v_s}{B_{\text{Dop}}} \quad (8.85)$$

$$= \frac{v_s}{2v_s \theta_{\text{Baz}} / \lambda} \quad (8.86)$$

$$= \frac{v_s}{2v_s (\lambda / L_x) / \lambda} \quad (8.87)$$

$$= \frac{L_x}{2}. \quad (8.88)$$

This gives the same result as Eq. (8.82).

Note that the resolution as described in this section is defined in terms of the width of the main lobe of the SAR PTR in the absence of noise. This follows the conventions in the SAR imaging community. Fine resolution does not necessarily imply that small features can be discerned in a practical sense if the data are too noisy, however. The resolution with respect to the width of the PTR is nonetheless a useful performance metric because it describes the number of independent realizations of the measurement that are averaged when spatially smoothing an image over a given length scale. These independent realizations are called looks. If a ground-projected SAR image is smoothed with a kernel whose effective widths in the along-track and cross-track directions are W_x and W_y , the number of independent looks N_L is approximately the number of SAR resolution cells covered by the smoothing kernel:

$$N_L \approx \frac{W_x W_y}{\delta_{xf} \delta_y}. \quad (8.89)$$

Here, the kernel effective widths W_x and W_y are defined in the FWHM sense (see Sect. 3.1.17) and are assumed to be greater than the corresponding resolutions δ_{xf} and δ_y .

The description of focused SAR processing above is simplified and idealized in several ways in order to give an intuitive overview of the basic principle without having to deal with details that are important for the actual implementation but would not necessarily provide much additional insight to typical users of SWOT data. Actual SAR processing typically must also account for the following:

- *Range migration.* Range migration refers to the shifting of a target from one range bin to another over the synthetic aperture due to the normal variation of the range history. Because of range migration, range compression and azimuth compression cannot be treated completely separably as in the simple description above, especially because the phase history depends on the broadside range ρ_0 . Multiple algorithms exist for handling range migration. These algorithms often involve tradeoffs in processing precision and computational efficiency, although the details are beyond the scope of this document.
- *Exact phase history.* The simple description above assumed a quadratic approximation to the phase history, but the exact phase history must often be used in high-precision applications. For spaceborne SAR systems, the phase history usually accounts for the curvature of the orbit and the Earth surface.
- *Realistic antenna gain variations and pointing.* The simple description above implicitly assumed uniform antenna weighting within the antenna footprint and complete rejection of echoes outside the antenna footprint. The gain of a realistic antenna falls off gradually from the boresight over the main lobe, however, and it also exhibits sidelobes. Echoes from portions of the antenna pattern that correspond to Doppler frequencies that alias with respect to the sampling of the PRF cause artifacts called azimuth ambiguities in the SAR images. Non-ideal antenna pointing also influences the way that azimuth compression is undertaken. These details have implications on the characteristics of the azimuth PTR.

The considerations above often influence the choice of the synthetic aperture length L_{SA} (or equivalently, the processed Doppler bandwidth B_{Dop} or the processing beamwidth θ_{SA}) during azimuth compression. The synthetic aperture length can be adjusted during ground processing by appropriately choosing the limits of the azimuth reference function in Eq. (8.76). It may be desirable to shorten the synthetic aperture length relative to the physical antenna beamwidth, equivalently filtering out higher Doppler frequencies during azimuth compression and thereby sacrificing azimuth resolution, in order to suppress aliased Doppler energy at the edges of the Doppler band if the PRF is not very high.

8.2.2.4 Azimuth Presumming

As described above, the Doppler spectrum is shaped by the antenna pattern in azimuth. The Doppler bandwidth B_{Dop} therefore scales with the antenna azimuth beamwidth θ_{BAZ} , as quantified by Eq. (8.48). The Doppler spectrum is sampled by the PRF, so the PRF must be sufficiently large to avoid unwanted aliasing, implying that the PRF must also scale with the Doppler bandwidth. The instrument data rate increases with the PRF. Therefore, while a wider Doppler bandwidth affords the potential for finer azimuth resolution, which can be exploited for additional spatial averaging, the wider bandwidth also necessitates increased resources for data storage and downlink.

If the full Doppler bandwidth is not required, the raw data may be filtered and resampled to a lower azimuth sample rate within the instrument in order to reduce the data downlink volume. This process is called azimuth presumming (“pre” because it occurs in real time within the instrument before data storage and downlink, and “summing” because a low-pass filter might

be implemented as a simple summation). Effectively, the raw data are low-pass filtered in the along-track direction in order to reduce the Doppler bandwidth of the data. Once the Doppler bandwidth has been reduced, the sample rate can be reduced through decimation or resampling (the resampling and filtering operations can equivalently be combined). The result is a series of echoes, which are each a function of range or fast time, with one echo per presum output time step. The output sample rate in azimuth is often called the presum PRF, which is typically lower than the PRF by a factor of 2 or more.

The presum data cannot be processed to the full resolution of the original data in that the presum filter involves a loss of information. However, the edges of the Doppler band sampled by the PRF that are removed by the presum filter may be of limited utility for many purposes anyway due to the gradual falloff of the antenna pattern in azimuth. Those portions of the Doppler spectrum are often associated with the shoulders of the antenna main lobe in azimuth, and they typically have lower SNR and increased Doppler aliasing compared to the central portion of the Doppler band that is retained after presumming.

8.2.2.5 Target Motion

It is also worth noting that SAR processing usually assumes that targets on the ground are stationary (in an Earth-fixed frame). However, if the target is moving, its actual phase history will differ from the phase history assumed by the reference function used for azimuth compression at the true location of the target. A moving target may therefore appear blurred or shifted in the SAR image. Shifts of moving targets are typically in the azimuth direction along a contour of constant range because an additional linear phase variation due to the target line-of-sight motion makes the phase history of the moving target match the phase history of some other stationary target at the shifted location.

8.2.2.6 SAR Output Coordinates

SAR images can be represented in various coordinate systems because range compression and azimuth compression naturally involve choices in how the data should be referenced. However, samples of a single 2-D SAR image cannot be geolocated in 3-D space without extra information such as prior knowledge of the surface height in order to find the intersection in 3-D space between the range, the broadside direction (or Doppler), and the ground surface.

Ultimately, a single-channel SAR generally gives the power and phase over a 2-D image. The phase arises from the constant term in Eq. (8.75), which is equivalent to the phase from Eq. (8.10). That is, during SAR azimuth compression, the phase variations over the aperture from the second term of Eq. (8.75) are compensated by the azimuth reference function, leaving only the first term, which depends on the broadside range ρ_0 .

8.3 Interferometric SAR Topographic Mapping

While individual 2-D SAR images do not provide 3-D information, a pair of 2-D images that observe the same scene on the ground from different locations can be used to measure the surface

topography. The concept of combining the information from two 2-D images in order to obtain 3-D information is familiar to many, with stereo vision and associated stereo mapping techniques standing out as well-known examples. Stereo techniques typically rely on the contrast of the image reflectivity in order to discern differences in the images with viewing geometry in order to estimate distance. Because SAR images are coherent, however, they lend themselves to interferometric techniques, which exploit the phase of the observations. Radar interferometry has a number of advantages over stereo mapping techniques in Earth remote sensing applications: radar instruments offer a good combination of fine spatial resolution and broad coverage; they can make measurements during the day or the night and are robust to cloud cover; they can be very sensitive to the surface topography with a baseline length that is feasible on a single spacecraft. This section describes the fundamentals of interferometric SAR topographic mapping from a simplified, introductory perspective. See [44] and [45] for additional details and references on interferometric SAR techniques.

8.3.1 Cross-Track Interferometry

Suppose that SAR images are acquired from each of two antennas that are separated by a cross-track (i.e., horizontal) baseline as illustrated in Fig. 8.11. Let B be the baseline length, and let ρ_1 and ρ_2 be the ranges to the target from each of the two antennas. Let $\Delta\rho$ be the difference in range between the two antennas:

$$\Delta\rho = \rho_2 - \rho_1. \quad (8.90)$$

The center of the baseline can serve as a convenient reference point on the radar. Let ρ_0 and θ_{look} be the range and the look angle to the target from this reference point. Because the baseline length is usually orders of magnitude smaller than the range ($B \ll \rho_0$), the look vectors from the two antennas and from the baseline center to the target are very nearly parallel. Consequently,

$$\Delta\rho \approx B \sin \theta_{\text{look}}, \quad (8.91)$$

and

$$\rho_0 \approx \frac{\rho_1 + \rho_2}{2}. \quad (8.92)$$

While the range to the target from each antenna is measured, the height of the surface is not known *a priori* (this is what the system is intended to measure), so the altitude h_s of the radar relative to the surface is not known either. That is, in an individual SAR image, the location of the target in the dimension into the page in Fig. 8.11 is estimated through the target Doppler or phase history. The range measurement places the target somewhere on a contour of constant range as illustrated in Fig. 8.11, but different locations along the range contour correspond to different target heights. In order to solve for the 3-D location of the target, the interferometric phase is used to estimate $\Delta\rho$, which can then be used with Eq. (8.91) to estimate the look angle θ_{look} . Knowledge of θ_{look} and ρ_0 therefore provide the altitude h_s of the radar relative to the target given that

$$h_s = \rho_0 \cos \theta_{\text{look}}. \quad (8.93)$$

The cross-track position y_t of the target is similarly found from

$$y_t = \rho_0 \sin \theta_{\text{look}}. \quad (8.94)$$




Figure 8.11: Illustration of the measurement geometry for cross-track SAR interferometry. Two antennas that are separated by a cross-track baseline acquire images of the same scene on the ground. The illustration is not to scale, and for a realistic system, the baseline would be much smaller in proportion to the range such that the look vectors would be almost perfectly parallel, and the range contours for the two antennas would be much more closely aligned with each other.

However, in order obtain to height measurements that are sufficiently precise, the estimate of the range difference $\Delta\rho$ must also be very precise—much more so than afforded by the radar range resolution, which is typically meter scale. Interferometry entails the use of the SAR image phase to measure $\Delta\rho$ to a small fraction of the SAR wavelength.

Assume that the instrument operates by transmitting from one antenna and receiving from both antennas, as is the case for KaRIn, although other designs are possible as well. Let ρ_1 be the range to a given target for the antenna that both transmits and receives, and let ρ_2 be the range to the same target for the antenna that only receives. From Eq. (8.10), the target phase in the first image is simply

$$\phi_1 = -\frac{4\pi}{\lambda}\rho_1. \quad (8.95)$$

In the second image, the signal travels along the path from first antenna to the target then back to the radar along the slightly different path to the second antenna. The echo delay, which gives rise to the phase, therefore is related to the total round-trip propagation distance $\rho_1 + \rho_2$. The target phase in the second image is thus

$$\phi_2 = -\frac{2\pi}{\lambda}(\rho_1 + \rho_2). \quad (8.96)$$

Denote the two complex SAR images as s_1 and s_2 . Each is a slant-plane image that is a function of both the range ρ and the along-track (azimuth) position x . Assuming that these images have not been spatially averaged, the images are called single-look, complex (SLC) images, where the term “look” here refers to the number of averaged samples. Let the magnitudes of the two SLC images be A_1 and A_2 so

$$s_1 = A_1 \exp(j\phi_1), \quad (8.97)$$

and

$$s_2 = A_2 \exp(j\phi_2). \quad (8.98)$$

A 2-D interferogram s_{21} can be formed from the two by images by multiplying one image (pixel wise) by the complex conjugate of the other image (the choice of which channel is conjugated is a matter of convention and bookkeeping):

$$s_{21} = s_1^* s_2 \quad (8.99)$$

$$= A_1 A_2 \exp(j(\phi_2 - \phi_1)) \quad (8.100)$$

$$= A_1 A_2 \exp(j\Delta\phi) \quad (8.101)$$

$$= A_{21} \exp(j\phi_{21}), \quad (8.102)$$

where the superscripted asterisk (“*”) again denotes complex conjugation and A_{21} and ϕ_{21} are the interferogram magnitude and phase, respectively. The interferogram phase is defined

$$\phi_{21} = \Delta\phi \quad (8.103)$$

$$= \phi_2 - \phi_1. \quad (8.104)$$

That is, the phase ϕ_{21} of the interferogram is the phase difference $\Delta\phi$ between the two SLC images. The interferogram phase, of course, is only defined modulo 2π radians.

The interferogram phase ϕ_{21} can be expressed in terms of the range to each antenna by substituting Eqs. (8.95) and (8.96) into Eq. (8.104):

$$\phi_{21} = -\frac{2\pi}{\lambda}(\rho_2 - \rho_1) \quad (8.105)$$

$$= -\frac{2\pi}{\lambda}\Delta\rho. \quad (8.106)$$

From Eq. (8.91), the range difference $\Delta\rho$ can be written in terms of the look angle θ_{look} or, equivalently, the scalar product between the baseline vector \vec{B} and the unit look vector \hat{l} , so

$$\phi_{21} \approx -\frac{2\pi}{\lambda}B \sin\theta_{\text{look}} \quad (8.107)$$

$$\approx -\frac{2\pi}{\lambda}\vec{B} \cdot \hat{l}, \quad (8.108)$$

where the baseline vector \vec{B} is the vector from one antenna phase center to the other. Note that Eq. (8.107) assumes that the baseline vector is purely horizontal, while Eq. (8.108) more generally allows for an arbitrary baseline orientation.

In other words, the interferogram phase ϕ_{21} , which is the phase difference between the individual SLCs, arises from the number of wavelengths by which the lengths of the paths to the target from the two antennas differ. The interferogram phase therefore gives a very precise estimate of $\Delta\rho$ and therefore θ_{look} . For the idealized case here, the target height can then be found as a function of the interferogram phase by rearranging and combining Eqs. (8.93) and (8.107):

$$h_s \approx \rho_0 \cos\left(\sin^{-1}\left(-\frac{\lambda}{2\pi B}\phi_{21}\right)\right). \quad (8.109)$$

This expression assumes that the interferogram phase ϕ_{21} has been unwrapped, meaning that an appropriate integer multiple of 2π radians has been added or subtracted in order to resolve the many-to-one mapping between possible target heights and observed interferometric phase values that can only be measured modulo 2π radians (see Sect. 8.3.5). Note that the expression above as well as others in this section involve approximations and simplifications because the intent of this document is to provide general, introductory insight into the measurement approach, not to give a detailed accounting of high-fidelity processing. The actual processing to generate KaRIn products avoids such approximations and simplifications and instead uses more advanced, albeit somewhat more complicated, algorithms.

The sensitivity of the height to the interferometric phase is

$$\frac{dh_s}{d\phi_{21}} = \pm \frac{\lambda\rho_0 \sin\theta_{\text{look}}}{2\pi B \cos\theta_{\text{inc}}}. \quad (8.110)$$

Note that this equation is derived from a more general height reconstruction model, not directly from the simplified model of Eq. (8.109). Therefore, Eq. (8.110) accommodates Earth curvature in that θ_{look} and θ_{inc} are not necessarily assumed to be identical. However, Eq. (8.110) does assume a purely horizontal baseline. The sign is left ambiguous to account for the conventions of the interferogram.

From Eq. (8.94), $\rho_0 \sin\theta_{\text{look}}$ in the numerator of Eq. (8.110) can be replaced by y_t . The implication is that the phase is more sensitive to the target height at near range than at far range. Equivalently, a constant offset in the interferometric phase corresponds to a tilt or a linear trend in the height as a function of the cross-track distance.

8.3.2 Interferometric Coherence

The interferometric coherence γ is often used as a measure of the level of random noise in an interferogram. This quantity is the complex correlation coefficient between the two SLCs comprising the interferogram, defined as

$$\gamma = \frac{\langle s_1^* s_2 \rangle}{\sqrt{\langle |s_1|^2 \rangle \langle |s_2|^2 \rangle}}, \quad (8.111)$$

where the angle brackets (“ $\langle \cdot \rangle$ ”) denote statistical expectation. The coherence usually varies spatially over the interferogram.

The magnitude of the coherence is a value between 0 and 1, with smaller values indicating noisier interferometric data. That is, if $|\gamma| = 0$, then the two SLCs are uncorrelated, and the interferogram phase is a random value that is uniformly distributed between 0 and 2π rad (or, equivalently, between $-\pi$ and $+\pi$ rad) and contains no information. On the other hand, if $|\gamma| = 1$, then two SLCs are identical except for a phase shift, and there is no random noise in the interferogram at all.

In practice, an estimate of the coherence is usually computed by spatially averaging the 2-D interferogram and SLC power images over a local neighborhood in order to estimate the statistical expectations in Eq. (8.111).

The terms “coherence” and “correlation” are often used interchangeably. Either term might also be used in reference to the magnitude of the complex correlation coefficient rather than to the complex value itself; the intent is usually clear in context.

Mechanisms that give reductions in the correlation magnitude are called sources of decorrelation. See Sect. 8.4 for discussion on common sources of decorrelation.

8.3.3 Coregistration

The 2-D interferogram is formed from the element-wise product of the two 2-D SLCs after one SLC has been conjugated, as described by Eq. (8.102). Before forming the interferogram in this manner, however, the two SLCs must be coregistered so that a pixel in one image contains echoes from the same targets on the ground as the corresponding pixel in the other image.

For example, the nonzero baseline may cause enough of a range shift between the two images that the echoes from a given target appear in different range bins in the different images. If the interferogram were formed from such images before coregistration, the pixels being multiplied would not be correlated, and the interferogram phase would be meaningless. That is, if the SLCs are not properly coregistered, the resulting misregistration causes a partial or complete loss of coherence and a commensurate increase in random noise in the interferometric phase.

Consequently, in preparation for interferogram formation, one or both SLCs may need to be resampled so that the two SLCs are sampled on a common slant-plane grid when the interferogram is formed. In the context of airborne SAR processing, this resampling operation is equivalent to motion compensating the two SLCs to a common flight track given that the common slant-plane grid implies a common effective flight track. For the example in Fig. 8.11, the common slant-plane grid might be the grid associated with ρ_0 and consequently with the baseline center. The resampling often involves both a shift and a range-dependent stretch. Because the

SLC image data are complex and band limited, the resampling is often accomplished with use of a truncated, weighted sinc interpolator.

Coregistration usually involves the use of a reference surface with some assumed elevation profile. This is because the range shift $\Delta\rho$ depends on the look angle θ_{look} , which, for a given range ρ_0 , depends on the target height relative to the radar (see Fig. 8.11). The target height is, of course, what the system seeks to measure, so the target height is generally not known perfectly. This is acceptable as long as the accuracy of the height of the reference surface is sufficient that residual misregistration due to the difference between the reference surface height and the true surface height does not cause any more decorrelation than can be tolerated. In most cases, a rough estimate of the reference surface is good enough to coregister the SLCs.

8.3.4 Flattening and Reference Locations

The interferometric phase from Eq. (8.106) is zero when the range difference $\Delta\rho$ between the two antennas is zero. The range difference is zero for targets whose look vectors from the baseline center are perpendicular to the baseline itself, per Eq. (8.108). If the baseline is perfectly horizontal, then the interferometric phase is proportional to $\sin\theta_{\text{look}}$ as expressed by Eq. (8.107).

For the KaRIn near-nadir viewing geometry, $\sin\theta_{\text{look}}$ has a nearly linear relationship with the target cross-track distance y_t , so even if the ground surface has a constant height, the unwrapped, absolute interferometric phase ϕ_{21} (see Sect. 8.3.5) varies almost linearly across the swath given that the baseline length B and the wavelength λ are constant. Height variations in the surface cause deviations in ϕ_{21} from the approximately linear phase progression that would be expected for a flat surface.

It is often desirable to remove the phase variations that correspond to a reference surface, which is often (but not always) flat, so that the remaining phase variations correspond to the surface height above the reference surface. This may be necessary, for example, before smoothing the complex values so that the values add coherently. This flattening operation is accomplished by computing what is known as the flat-Earth or flattening phase ϕ_f , which is the interferometric phase that would be expected if the true surface were the same as the assumed reference surface, and removing it from the absolute or unflattened interferometric phase ϕ_{21} in order to obtain the flattened interferometric phase ϕ_{21f} :

$$\phi_{21f} = \arg(\exp(j\phi_{21}) \exp(-j\phi_f)). \quad (8.112)$$

Here, $\arg(z)$ is the phase of the complex value z .

A conceptual example of flattening is shown in Fig. 8.12. The unflattened interferogram, which still contains the flat-Earth phase, is shown in panel (a). The modulations of the flat-Earth phase in the unflattened interferogram are due to the height variations in the surface. The flattened interferogram in (b) is obtained when the flat-Earth phase in (c) is removed from the unflattened interferogram in (a). The true topographic height is shown normalized in (d). Several other aspects of this figure are worth noting:

- Layover (see Sect. 8.2.1.3) is evident for the hump at the upper left of panel (b). All of the humps exhibit foreshortening in that they appear to lean toward the left.




Figure 8.12: Conceptual example of interferogram flattening for a scene with five notional humps in the height: (a) unflattened, absolute interferogram phase; (b) flattened interferogram phase; (c) flat-Earth phase removed from (a) to obtain (b); (d) ground-plane image of the normalized surface height. The phase images in (a), (b), and (c) are shown in the slant plane with range increasing to the right. The phase is shown wrapped between $-1/2$ and $+1/2$ cycles. The ground-plane image in (d) has cross track increasing to the right. Foreshortening and/or layover are evident for all humps in the slant-plane images. Phase wrapping is evident even for the flattened phase of the upper left hump in (b).

- Phase wrapping (see Sect. 8.3.5) is evident in panels (a), (b), and (c). Even the flattened phase must be unwrapped before height reconstruction is possible.
- The increased sensitivity of the phase to height variations at near range relative to far range [see Eq. (8.110)] can be seen when considering that the hump at the lower left of panel (b) has a greater range of phase values than the hump at the upper right even though its height in panel (c) is smaller.
- The non-uniform mapping between cross track and slant range is visible both in the tighter fringe spacing at the left of panel (c) relative to the right and in the difference in the appearance of the humps in the ground-plane image of panel (d) as compared to the slant-plane flattened phase in panel (c).

Note that an interferogram from a real spaceborne sensor would likely have many more fringes across the swath for its flat-Earth phase than shown in Fig. 8.12(c); the images in the figure are only for illustration.

In some cases, interferogram flattening may involve the use of a reference surface that is not flat. For example, a reference digital elevation model (DEM) might be used to compute the flattening phase, which would then have variations that correspond to the surface topography. In the ideal case that the true topography is identical reference DEM, the flattened interferogram would have zero phase everywhere. Differences between the true topography and the reference DEM would leave residual phase variations that are related to the errors in the reference DEM.

Whenever the interferogram is flattened, knowledge of the reference surface must be maintained so that the flattened phase can be interpreted. That is, the absolute (unflattened) phase can be related to the absolute height of the surface topography through knowledge of the imaging geometry. The flattened phase can be related only to the height of the surface relative to the reference surface used for flattening, so the reference surface must accompany the flattened phase if absolute referencing is desired.

Knowledge of the reference surface is often stored in the form of a reference location for each sample of the 2-D interferogram. The reference location is the absolute location in 3-D space to which the flattened interferogram phase is referenced. If the flattened phase is zero after phase unwrapping (see Sect. 8.3.5), the target geolocation would be equivalent to the reference location. The reference location includes not only the height, but the along-track and cross-track coordinates as defined by the combination of range, Doppler, and interferometric phase (see Sect. 8.3.6).

While interferogram flattening may be performed as its own processing step, it can also be incorporated into the formation of the individual SAR images.

8.3.5 Phase Unwrapping and Ambiguity Resolution

The interferogram phase is only defined over an interval of 2π rad, as described in Sect. 8.1.2. Therefore, while the phase is very sensitive to changes in the surface topography [see Eq. (8.110)], it is ambiguous in the sense that targets at very different locations might give the same wrapped phase. Consequently, the phase must be unwrapped before targets can be geolocated absolutely. Phase unwrapping is the process of inferring the integer number of cycles (multiples of 2π rad)

that must be added to the wrapped interferogram phase to obtain an unwrapped, absolute phase that can unambiguously give the target height h_s through Eq. (8.109).

Define the ambiguity height h_a as the magnitude of the change in target height that corresponds to a change in the interferometric phase of 2π rad. The ambiguity height can found by multiplying the right-hand side of Eq. (8.110) by 2π and taking the absolute value:

$$h_a = \frac{\lambda \rho_0 \sin \theta_{\text{look}}}{B \cos \theta_{\text{inc}}}. \quad (8.113)$$

Strictly, this is a linear approximation rather than an exact expression, but it is sufficient for the tutorial discussion in this document. Because the target cross track y_t is $\rho_0 \sin \theta_{\text{look}}$ [see Eq. (8.94)], which appears in the numerator of Eq. (8.113), the ambiguity height scales almost linearly with cross track.

If a reference DEM gives knowledge of the true surface height to within $\pm h_a/2$, then it is usually straightforward to compute the unwrapped phase by simply finding the integer number cycles that makes the unwrapped interferometric phase ϕ_{21} give a reconstructed height that is within $\pm h_a/2$ of the reference DEM. This also assumes that there is no layover, which might give multiple solutions since the range contour might intersect the true surface at multiple locations.

If the reference DEM is not necessarily accurate to $\pm h_a/2$ everywhere, phase unwrapping is often accomplished by first performing a relative spatial unwrapping followed by resolution of the absolute ambiguity. That is, first, a spatial unwrapping step solves for the relative phase relationship over contiguous groups of pixels. The result of the spatial unwrapping is a set of defined regions such that within each region, the pixels should have the correct relative phase relationship with respect to each other, but all the pixels within a region may still be offset by the same integer number of cycles from the correct unwrapped phase. The absolute ambiguity is then resolved (i.e., the integer number of cycles is estimated) for each region.

Typically, the spatial unwrapping step is based on the assumption that the phase is slowly varying enough spatially that the differences in phase from one pixel to the next are usually within $\pm\pi$ rad. The pixel-to-pixel differences in phase can then be integrated from one pixel to the next over the region to obtain the unwrapped phase up to an unknown, constant number of cycles for the region.

The ambiguity resolution step then takes advantage of the fact that the difference between the true surface height and the reference DEM height is usually more well behaved when considering a spatially large region than when considering an individual, noisy pixel. In other words, the estimation of a single integer number of cycles over a whole, spatially extensive region makes the estimation process less sensitive to small-scale errors in both the reference DEM and the measurement. Moreover, if a region spans a large cross-track extent, portions of the region that are farther from nadir and that consequently have larger ambiguity heights can allow the correct ambiguity to be determined for regions that are closer to nadir and that might be more difficult to unwrap otherwise.

If features in the data can be matched to ground-plane maps, then the ambiguity resolution can also be aided by the horizontal geolocation since an unwrapping error affects both the height and the cross-track position of a target.

8.3.6 Height Reconstruction and Geolocation

Once the interferometric phase has been unwrapped, height reconstruction and geolocation give the 3-D position of the target. Height reconstruction is typically combined with geolocation, which entails the solution of the absolute target position in 3-D space based on the range ρ_0 , the Doppler f_D , and the interferometric phase ϕ_{21} .

To build intuition on the geolocation process, it is helpful to consider that a contour of constant range is a sphere centered at the radar with radius ρ_0 . A contour of constant Doppler is a cone whose symmetry axis is the radar ECEF velocity vector such that $\vec{v}_r \cdot \hat{l}$ in Eq. (8.43) is constant. A contour of constant unwrapped interferometric phase is a cone whose symmetry axis is the baseline vector \vec{B} such that $\Delta\rho$ and $\vec{B} \cdot \hat{l}$ are constant in Eqs. (8.106) and (8.108). The intersection of the range sphere, the Doppler cone, and the interferometric phase cone give the target location. Note that this model is more general than is assumed for most of the discussion in earlier sections of this chapter; this model does not necessarily assume that the radar velocity is horizontal or that the baseline is purely in the cross-track direction. It does, however, assume that the range is large compared to the baseline length ($\rho_0 \gg B$) and that the motion of the radar while the signal is in flight is negligible. These latter approximations are not necessarily made in the KaRIn processing software, though. The generalized 3-D geolocation model is illustrated in Fig. 8.13.

For high-precision applications, the target geolocation must account for Earth curvature and non-ideal spacecraft velocity and attitude. Such details are beyond the scope of this document, but while the mathematical constructions are not necessarily trivial, the geolocation process generally amounts to solving for the three coordinates of the target position in space based on the three observations of range, Doppler, and interferometric phase.

For the simple case of a broadside viewing geometry with a purely horizontal flight track, the Doppler cone has a half angle of 90° and therefore reduces to a plane normal to the velocity vector for this special case. The half angle of the Doppler cone is indeed nearly (but not necessarily exactly) 90° for all cases relevant to KaRIn. The range contours shown in Figs. 8.5, 8.7, and 8.11 represent the circular intersections of different range spheres and the Doppler cone if the Doppler cone is taken to be the plane of the page. These circular range-Doppler intersections are equivalent to $C_{\rho D}$ in Fig. 8.13. The intersection of the interferometric phase cone with that circular range-Doppler contour gives the target position.

Errors in the interferometric phase affect the estimate of the angle of the phase cone, thereby moving the reconstructed target location along the circular range-Doppler contour as illustrated in Fig. 8.14. As a result, phase noise (or phase unwrapping errors, as described in Sect. 8.3.5) cause both height errors and cross-track location errors that are correlated with one another. For small errors where the curvature of range-Doppler contour is negligible, the cross-track error δ_c is related to the height error δ_h by

$$\tan \theta_{\text{inc}} \approx \pm \frac{\delta_h}{\delta_c}, \quad (8.114)$$

where the incidence angle θ_{inc} is relative to the local vertical direction. The sign depends on the swath side and the convention for defining the cross-track direction.

Note that errors in the knowledge of the orientation of the baseline vector \vec{B} can be particularly important. Suppose that the true baseline roll angle is zero, but the baseline roll angle




Figure 8.13: Geolocation diagram for the general case in which neither the velocity vector nor the baseline vector are purely horizontal, and the antenna pointing is arbitrary. Angles are exaggerated relative to a typical KaRIn case for illustration purposes. The point T is the target location. The points S_I , S_B , and S_{ZD} represent the radar locations along the flight track when the antenna illuminates the target, when the target is at broadside, and when the target is at zero Doppler, respectively. B is the interferometric baseline at the time of target illumination. C_ρ represents a spherical contour of constant range ρ that is centered at S_I and that goes through T . $C_{\rho g}$ is the intersection of C_ρ with an assumed flat ground surface that contains the target. C_{Dg} is the intersection of the Doppler cone containing the target with the flat ground surface. $C_{\phi g}$ is the intersection of the interferometric phase cone containing the target with the flat ground surface. $C_{\rho D}$ is the intersection of the range sphere C_ρ with the Doppler cone containing the target. $C_{\rho\phi}$ is the intersection of the range sphere C_ρ with the interferometric phase cone containing the target.




Figure 8.14: Simplified geolocation diagram showing the sensitivity of the target position to errors in the interferometric phase. Phase noise in the interferogram causes undesired variation in the angle of the phase cone and therefore the look angle θ_{look} to the reconstructed target location. This gives errors in both the height and the cross-track location that are correlated with one another. If the diagram were to scale, the baseline length would be much less than the width of the line indicating the look vector, so only a single look vector is shown.




Figure 8.15: Diagram illustrating the effect of an error in the knowledge of the baseline roll angle.

is believed to be some non-zero angle δ_θ . The interferometric phase measures the angle of the phase cone relative to the true baseline, so if height reconstruction proceeds with the incorrect baseline orientation, the estimated look angle to the target has an angular error of δ_θ . This would give an incorrect estimate for the target location just as a phase error would. The target location is very sensitive to such baseline roll errors because the target location error is the product of the angular error δ_θ and the range ρ_0 . That is, the look vector of length ρ_0 acts as a very long lever arm such that small angular errors in the direction of that lever arm give large errors in the target position. This is illustrated in Fig. 8.15.

Another intuitive way of thinking about errors in the knowledge of the baseline roll is that such an error effectively rotates the entire measurement frame. The height-reconstructed surface is therefore also rotated about the radar location by the angle of the roll knowledge error. This results in an erroneous cross-track tilt with a slope of approximately δ_θ in the measurement. A cross-track shift occurs as well.

Of course, the range, Doppler, and interferometric phase give the target location relative to the radar. Therefore, errors in the knowledge of the position of the radar translate directly to errors in the target geolocation. Specifically, errors in the knowledge of the radar height give corresponding errors in the estimated target height. Errors in the knowledge of the horizontal position of the radar give errors in the horizontal geolocation of the target, but the system requirements for horizontal accuracy are often much looser than for vertical accuracy given the limited horizontal resolution of the radar.

8.4 Interferometric Decorrelation and Looks

The interferometric coherence γ , defined in Sect. 8.3.2, warrants special attention because it both characterizes the level of random noise in an interferometric measurement and provides additional measurement information in its own right.

The total coherence γ_{tot} is often taken to be the product of several separable sources of decorrelation, as in

$$\gamma_{\text{tot}} = \gamma_{\text{snr}} \gamma_{\text{geo}} \gamma_{\text{vol}}, \quad (8.115)$$

where γ_{snr} , γ_{geo} , and γ_{vol} are the so-called noise, geometric, and volumetric correlation terms that are described below in this section. Each term has a magnitude between 0 and 1. Other sources of decorrelation also contribute to the total decorrelation, though they are often smaller contributors and are therefore not discussed in this document.

An estimate γ_{est} of the total coherence γ_{tot} is usually computed from interferometric data by averaging over spatial windows to approximate the statistical expectations in Eq. (8.111), giving

$$\gamma_{\text{est}} = \frac{\sum_k s_{1k}^* s_{2k}}{\sqrt{\sum_k |s_{1k}|^2 \sum_k |s_{2k}|^2}}, \quad (8.116)$$

where the index k usually encompasses the set of interferogram samples around some 2-D neighborhood around the estimate γ_{est} . The coherence is usually estimated as a 2-D image.

The averaging window for the coherence estimate γ_{est} is often (but not necessarily) the same as the window over which the complex interferogram values or the image power values are averaged in order to reduce noise. If there are N_L independent samples in the averaging window, the estimates are said to be taken over N_L independent looks.

The subsections below provide a brief introduction to largest sources of decorrelation and the effect of decorrelation on phase noise. A more thorough discussion is given in [44].

8.4.1 SNR Decorrelation

The SNR decorrelation γ_{snr} arises because the two SAR images s_1 and s_2 are each associated with additive noise. The two images are acquired by physically different receiver channels, so this additive noise is independent between the two images and is therefore completely uncorrelated.

The SNR decorrelation γ_{snr} is sometimes also called noise decorrelation or thermal decorrelation. The latter nomenclature arises because the main source of additive noise in the SAR images is usually thermal noise in the radar receiver electronics.

With each of s_1 and s_2 as the sum of correlated signal and uncorrelated noise, the resulting SNR decorrelation γ_{snr} is [46]

$$\gamma_{\text{snr}} = \frac{1}{1 + \frac{1}{\text{SNR}}}. \quad (8.117)$$

Therefore, as the SNR grows toward infinity, γ_{snr} approaches 1. Conversely, as the SNR goes to zero, γ_{snr} also goes to zero.

Note that Eq. (8.117) assumes that the SNR is the same between the two channels. If it is not, then the SNR decorrelation γ_{snr} is

$$\gamma_{\text{snr}} = \frac{1}{\sqrt{\left(1 + \frac{1}{\text{SNR}_1}\right)\left(1 + \frac{1}{\text{SNR}_2}\right)}}, \quad (8.118)$$

where SNR_1 and SNR_2 are the SNRs for each of the two SAR images.

Note that in this context, the SNR describes the ratio of the echo power to the receiver noise, but the receiver noise is not the only contributor to phase noise in the interferogram. This is because the randomness of the surface itself can cause decorrelation and phase noise in the interferogram. Consequently, the interferogram phase noise could be significant even in an ideal case with infinite SNR. Stated differently, the total coherence γ_{tot} can be less than 1 even if the SNR decorrelation γ_{snr} is ideally 1 if the geometric and/or volumetric correlation are less than 1.

8.4.2 Geometric Decorrelation

Geometric decorrelation, also known as baseline decorrelation or spatial decorrelation, is due to the slight difference in incidence angle θ_{inc} between the two SAR images s_1 and s_2 . While the difference in incidence angle is minute, the SAR phase is of course very sensitive to such small differences (this is what makes interferometry possible).

Suppose that the baseline length is zero such that the two SAR images are acquired from precisely the same location. In this case, the two SAR images would be identical, notwithstanding receiver noise, which is separately captured by γ_{snr} . As the baseline length increases, the two images become increasingly dissimilar. Consider a particular interferogram resolution cell. Because the interferogram phase ϕ_{21} is approximately proportional to $\sin\theta_{\text{look}}$ in Eq. (8.107), the phase varies somewhat over the cross-track extent of the resolution cell. The phase variation here is exactly that which gives rise to the flat-Earth phase described in Sect. 8.3.4. However, the flat-Earth phase cannot be removed within a single resolution cell because all echoes on the ground over that cell contribute to a single sample of the interferogram. During flattening, the flat-Earth phase for the center of the resolution cell is applied to that sample. The echoes from different locations on the ground in cross track that are within the same resolution cell therefore do not add completely coherently in the interferogram. This leads to the decorrelation captured by γ_{geo} .

If the range PTR of each SAR image is a sinc function (see Sect. 8.1.3), the geometric decorrelation γ_{geo} is given by

$$\gamma_{\text{geo}} = \begin{cases} 1 - |B_{\perp} / B_{\perp c}| & \text{if } |B_{\perp} / B_{\perp c}| < 1 \\ 0 & \text{otherwise} \end{cases}, \quad (8.119)$$

where the perpendicular baseline B_{\perp} is the component of the baseline B that is perpendicular to the look vector \vec{l} , and the critical perpendicular baseline $B_{\perp c}$ is given by

$$B_{\perp c} = \frac{\lambda \rho \tan \theta_{\text{inc}}}{\delta \rho}. \quad (8.120)$$

Here, δ_ρ is the slant-range resolution. If the baseline B is purely horizontal, then $B_\perp = B \cos \theta_{\text{look}}$.

Therefore, when the perpendicular baseline B_\perp is zero, the geometric correlation γ_{geo} is 1 (perfect correlation). As B_\perp increases, thereby increasing the difference in θ_{look} and θ_{inc} for the two SAR images, the geometric correlation γ_{geo} decreases because the differences between the two images grow. The correlation drops to zero and remains at zero as B_\perp exceeds a critical maximum value $B_{\perp c}$.

The critical baseline $B_{\perp c}$ is proportional to $\tan \theta_{\text{inc}}$, so it is smaller near nadir. Consequently, γ_{geo} is also typically smaller (there is greater decorrelation) at near range than far range. This can be understood by considering that if the slant-range resolution δ_ρ is fixed, then the cross-track resolution $\delta_y = \delta_\rho / \sin \theta_{\text{inc}}$ [from Eq. (8.38)] after projecting δ_ρ onto the ground is coarser at steeper incidence angles, so there is more phase variation over a resolution cell.

Note that geometric decorrelation can be minimized (γ_{geo} maximized) through an interferometric processing technique called range spectral filtering (also sometimes called Prati filtering) [47]. This involves filtering the range spectra of the two SAR images differently so that the opposite ends of the spectra are trimmed. The trimmed spectra maximize the spectral overlap and hence the correlation when the wavenumbers of each spectrum are projected onto the ground. A deeper discussion of this topic is beyond the scope of this document, however.

8.4.3 Volumetric Decorrelation

Volumetric decorrelation is somewhat similar in principle to geometric decorrelation in that the two SAR images s_1 and s_2 become more dissimilar with increasing baseline length if there is a vertical as well as a cross-track distribution of the scattering from the Earth surface. That is, geometric decorrelation arises because the echoes from different cross-track locations on a flat scattering surface have different interferometric phases over a resolution cell. Volumetric correlation arises because the echoes from different depths or vertical locations over a 3-D scattering volume have different interferometric phases over the resolution cell.

Radar signal penetration into ice, snow, or vegetation are classical examples of cases in which the radar signal is scattered from a 3-D volume rather than just a 2-D surface. However, variations in the height of a 2-D scattering surface over a resolution cell—such as from waves on the ocean surface—also give rise to volumetric decorrelation. The assumption here is that the horizontal projection of the range-Doppler resolution cell (or perhaps the spatial averaging window) is similar to or larger than the wavelength of the ocean waves.

Volumetric decorrelation becomes significant when the vertical variations of the scattering surface have a magnitude that is an appreciable fraction of the ambiguity height h_a [see Eq. (8.113)]. The magnitude $|\gamma_{\text{vol}}|$ of the volumetric decorrelation can then be computed from

$$|\gamma_{\text{vol}}| = \left| \int f(h) \exp\left(-j2\pi \frac{h}{h_a}\right) dh \right|, \quad (8.121)$$

where h is the surface height and $f(h)$ is the probability density function (PDF) of the scattering as a function of the height. It is assumed that $f(h)$ integrates to 1. Because the complex values of interferogram signals at different heights h do not all have the same phase, the resulting complex sum typically has a magnitude less than 1.

Note that the height in Eq. (8.121) can be referenced arbitrarily given that a constant height offset only affects the phase (not the magnitude) of the integral.

If the scattering PDF $f(h)$ has an assumed form with a single free parameter, then knowledge of $|\gamma_{\text{vol}}|$ can be used in an inversion for the unknown value. For example, if $f(h)$ is assumed to be a Gaussian with an unknown standard deviation σ_f , then an estimate of $|\gamma_{\text{vol}}|$ can be used to estimate σ_f . Note that the mean of $f(h)$ does not affect the magnitude of the volumetric decorrelation because a change in the mean is equivalent to a constant offset in h .

The estimate of the magnitude of the volumetric decorrelation $|\gamma_{\text{vol}}|$ can be obtained from the magnitude of the observed total decorrelation $|\gamma_{\text{tot}}|$ by solving Eq. (8.115) for $|\gamma_{\text{vol}}|$:

$$|\gamma_{\text{vol}}| = \frac{|\gamma_{\text{tot}}|}{|\gamma_{\text{snr}}| |\gamma_{\text{geo}}|}. \quad (8.122)$$

This assumes that the magnitudes of the SNR decorrelation and the geometric decorrelation (and any other significant sources of decorrelation) can be estimated from models, calibration data, and/or measurement information. This is often feasible because the instrument noise and hence the SNR can often be estimated from internal calibration data, and the geometric decorrelation depends mainly on system parameters that are reasonably well known.

8.4.4 Phase Noise

Precise expressions for the statistics of the interferometric phase noise are somewhat involved mathematically and depend on many details of the data processing. However, a simple, well-known expression for the standard deviation σ_ϕ of the phase noise that is useful under many conditions of practical interest is the Cramer-Rao bound [48]:

$$\sigma_\phi \approx \sqrt{\frac{1 - \gamma_{\text{tot}}^2}{2N_L \gamma_{\text{tot}}^2}}. \quad (8.123)$$

This expression applies for moderately high correlation γ_{tot} and a significant number of independent looks N_L .

Note that the number of independent looks N_L is not necessarily the same as the number of pixels of the 2-D interferogram that have been spatially averaged if the 2-D interferogram is oversampled, as is often the case.

The standard deviation σ_h of the random height noise can be found by multiplying σ_ϕ by the height sensitivity from Eq. (8.110):

$$\sigma_h = \left| \frac{dh}{d\phi_{21}} \right| \sigma_\phi \quad (8.124)$$

$$= \left| \frac{\lambda \rho_0 \sin \theta_{\text{look}}}{2\pi B \cos \theta_{\text{inc}}} \right| \sigma_\phi. \quad (8.125)$$

Note that this estimate of the level height noise does not necessarily capture random errors due to other mechanisms such as layover or “surfboard” errors (see Ch. 10), however.

Chapter 9

KaRIn Design Overview

This chapter gives a high-level description of the KaRIn measurement architecture. The information here is intended to help users of SWOT KaRIn products understand the characteristics of the data in the context of how the measurements are made.

9.1 KaRIn Hardware Design

The SWOT science objectives necessitate exquisite measurement accuracy, which in turn brings about a number of very difficult engineering challenges in the implementation of the instrument hardware. The KaRIn hardware design therefore incorporates many novel features to overcome these challenges. However, many of the details of the hardware design can be abstracted when considering the ultimate SWOT science products from a user perspective. Therefore, this section provides only a brief overview of the KaRIn hardware design, focusing on aspects that are likely to be most relevant to the science interpretation of the data.

A model view of the KaRIn hardware is shown in Fig. 9.1. Note that the nadir direction is pointed toward the top of the page in Fig. 9.1. A mechanical structure called the KaRIn metering structure provides a rigid backbone for the instrument. This structure gives rise to the KMSF coordinate frame (see Sect. 3.3.6). Key KaRIn instrument hardware parameters are summarized in Table 9.1.

Note that Table 9.1 contains nominal values. Some parameters, such as the PRF and the transmit pulse width, vary by a few percent over the orbit as the spacecraft altitude above the Earth surface changes (see Fig. 4.1).

9.1.1 KaRIn Antennas

The KaRIn antennas are designed to form four different beams, one for each combination of the two (left and right) sides of the measurement swath on the ground and the two ends of the interferometric baseline. This is because the two sides of the measurement swath on the ground cannot be observed simultaneously with a common transmit pulse, as targets on either side of nadir would have echoes with the same range and the same Doppler as each other (see Fig. 8.13). The targets to either side would therefore not be distinguishable from one another in the received data. Consequently, the antenna beams must illuminate only one side of the swath at a time and




Figure 9.1: Design model of KaRIn. Note that the illustration is oriented so that the nadir-looking direction (KMSF + z) is pointed toward the top of the page, while it would be pointing down toward the Earth in flight. The nadir module is shown for context, but it is separate from the KaRIn instrument.

Parameter	Value
Center frequency	35.75 GHz
Wavelength	0.83858 cm
Peak transmit power (at end of life)	1500 W
Transmit pulse width	6.4 μ s
PRF (per polarization)	4420 Hz
Pulse Bandwidth	200 MHz
ADC real sample rate	300 MHz
Baseline length	10.042 m
Reflectarray length (azimuth or along track)	5 m
Reflectarray width (elevation or cross-track)	0.26 m
Antenna 3 dB beamwidth in elevation	3°
Antenna 3 dB beamwidth in azimuth	0.1°
Antenna boresight off-nadir angle	$\pm 2.65^\circ$
Antenna peak gain (one way)	48 dB
Instrument mass	600 kg
Instrument average power consumption	800 W
Data rate, LR only	14 Mbps
Data rate, LR and HR	330 Mbps

Table 9.1: Key KaRIn instrument design parameters.

reject any echoes from the opposite side that may have ambiguous range-Doppler signatures with respect to the desired echoes on the illuminated side. The beams for the two sides use opposite electromagnetic polarizations (horizontal and vertical; see Sect. 3.1.15) in order to provide further isolation between the echoes of the two sides of the swath.

The four antenna beams are thus designated the H+, H−, V+, and V− beams, with the letter corresponding to the polarization (H or V) that is used for the swath side that is observed, and the sign (+ or −) corresponding to whether the observation is made from the +y or the −y (in KMSF) end of the interferometric baseline. Measurements of the two swath sides are then interleaved on a pulse-by-pulse basis. KaRIn nominally transmits from only one of the KMSF +y or −y sides at a time, depending on which of its redundant high power amplifiers (HPAs) is in use. The inactive HPA is a cold spare, meaning that it is powered on only if needed and only by ground command. While only one of the +y or −y sides transmits, both always receive for each swath side. The use of the four beams is illustrated in Fig. 9.2.

Thus, as KaRIn transmits pulses at a rate of approximately 8840 Hz, the first pulse relative to some arbitrary reference might be directed toward the H side of the swath using the H+ beam. Echoes from that pulse would be captured by both the H+ and H− beams. The second pulse would then be directed toward the V side of the swath using the V+ beam. Echoes from that pulse would be captured by both the V+ and V− beams. The pattern would then repeat, with “odd-numbered” pulses used for observations of the H side of the swath and “even-numbered” pulses used for observations of the V side of the swath. The effective PRF for each swath side is therefore approximately 4420 Hz, or half the combined pulse rate.




Figure 9.2: Illustration of the interleaved KaRIn measurements between the H and V sides. Notice that each of the four beams originates from a different feed. The arrow direction(s) for a given beam indicate whether the beam is used for both transmit and receive (bidirectional arrow) or for receive only (unidirectional arrow). This illustration assumes the nominal yaw state (0°) and transmit from the $+y$ HPA (but note that the science phase began with the $-y$ HPA in use). The spacecraft velocity direction is into the page.

	Nominal yaw (0°)	Flipped yaw (180°)
Left	H	V
Right	V	H

Table 9.2: KaRIn polarization by swath sides and yaw state.

Antenna designation	Beam
A1F1	V–
A1F2	H–
A2F3	H+
A2F4	V+

Table 9.3: KaRIn antenna designations.

Note that the mapping between the H and V polarizations and the left and right sides of the measurement swath depends on the yaw state of the spacecraft (see Sect. 3.1.8, Fig. 3.2, and Table 9.2).

The KaRIn antenna design implements the H+, H–, V+, and V– beams through the use a pair of reflectarrays, with one at each of the $\pm y$ ends of the interferometric baseline, and a set of four feeds, with two feeds on each of the $\pm y$ sides of the central KaRIn module near the mast roots (see Figs. 9.1 and 9.2). The H+ and V+ feeds illuminate the +y reflectarray to form the H+ and V+ beams, while the H– and V– feeds illuminate the –y reflectarray to form the H– and V– beams. These four combinations of reflectarray and feed are also sometimes denoted A1F1, A1F2, A2F3, and A2F4, where A1 and A2 refer to the –y and +y reflectarrays, and F1, F2, F3, and F4 refer to the V–, H–, H+, and V+ feeds, respectively (see Table 9.3).

Each planar reflectarray comprises an array of passive elements that each reflect incoming electromagnetic waves with a particular phase shift. The phase shifts of the array elements are designed to vary over the face of the reflectarray so that the reflection of a wavefront from a feed is focused into a tight antenna beam with the desired beam characteristics. This is analogous to the manner in which a parabolic reflector focuses the illumination from a feed in a typical dish antenna, but the reflectarray employs phase-shifting elements over a flat array surface rather than relying on the physical curvature of a reflector surface. The two feeds (H and V) illuminating a given reflectarray are offset from one another in the KMSF z direction and arranged so that their beams have boresights pointed in different directions in order to illuminate the swath sides to the left and right of nadir on the Earth surface.

An alignment mechanism at the center of each reflectarray allows for minor, single-axis adjustments to the orientation of the reflectarray relative to its feeds. The alignment mechanisms are used only during KaRIn commissioning, before KaRIn calibration, to optimize the alignment of the +y and –y beams with each other.

The KaRIn masts and reflectarrays are folded when the instrument is stowed to fit inside the fairing of the launch vehicle during launch. They are deployed shortly after launch, before the KaRIn radar electronics are powered on for the first time in space. After being deployed, the

antennas cannot be brought back to the stowed configuration.

9.1.2 KaRIn Radar Electronics

The KaRIn electronics on the central metering structure generate the transmit pulse, capture the received echoes, and provide for the overall command, control, and timing of the radar. During operation, the same KaRIn transmit and receive hardware is used for both the H and V swath sides, but there are parallel receive chains with separate hardware for the $+y$ and $-y$ channels.

The transmit pulse is created as follows: A chirp generator creates the transmit waveform from an on-board digital representation of the chirp. An up-converter next shifts the frequency of the chirp from a lower, working frequency to the final Ka-band carrier frequency through multiple stages of mixing with the synchronous oscillator outputs of a frequency synthesizer. A high-power amplifier (HPA) then amplifies the chirp to the necessary power level for transmission to the Earth. The HPA consists of an extended interaction klystron (EIK) and an accompanying high-voltage power supply (HVPS). KaRIn has two HPAs for redundancy. During normal operations, one HPA is powered on, while the other serves as a cold spare. The two HPAs are designated HPA+ and HPA– because they transmit from the $+y$ and $-y$ antenna feeds, respectively.

KaRIn began its commissioning activities with HPA+ but switched to HPA– very early in the mission. HPA– remained in use for the remainder of KaRIn commissioning (including antenna alignment) and throughout the entirety of the subsequent calibration phase and the transition to the science orbit.

A duplexer at the radar front end contains RF switches that alternately direct transmit pulses from the HPA in use to the H and V feeds on the transmitting side (H+ and V+ for HPA+ and H– and V– for HPA–).

The duplexer also switches the radar receiver electronics between the H and V feeds for both the $+y$ and $-y$ antenna channels. A dual-channel down-converter amplifies and synchronously shifts these signals down in frequency from Ka band to baseband. Analog-to-digital converters (ADCs) then digitize the quadrature (I/Q) baseband signals from each channel, and a set of field programmable gate arrays (FPGAs) perform extensive, real-time digital processing on the data (see Sect. 9.2). The data are then packetized and sent to the spacecraft solid-state recorder (SSR) for eventual downlink.

Flight software (FSW) running on a space-qualified, radiation-hardened computer processor controls and coordinates the operations of the radar. A power distribution unit (PDU) fed by spacecraft DC power conditions and provides DC power to the KaRIn digital electronics.

In addition to switching between antenna feeds, the duplexer also has switch paths that allow KaRIn to record internal, dynamic calibration or “loopback” signals. These signals are captured by using the KaRIn receive electronics to digitize and store the transmit pulse directly from the transmit hardware in addition to the echoes of the transmit pulse from the Earth surface. The loopback pulses provide information on the drift of the KaRIn analog hardware over orbital and sub-orbital time scales so that such variations can be compensated in later ground processing.

The KaRIn analog electronics that implement chirp generation, frequency synthesis, up-conversion, and down-conversion are collectively called the hyperbox (HB). The hyperbox and duplexer together are called the radio frequency unit (RFU). The digital electronics are sometimes referred to as the KaRIn Digital Electronics Subsystem (KDES). On each of the $\pm y$ sides,

the mast and reflectarray are together called a deployable antenna assembly (DAA).

9.1.3 Other KaRIn Hardware

KaRIn thermal hardware maintains the temperature stability of the electronics and structures, thereby allowing the instrument to meet its overall stability requirements. A panel on the $+y$ side of the KaRIn module (see Fig. 9.1) cools the instrument by allowing excess heat from the electrical power dissipation of the KaRIn electronics to escape into space through thermal radiation. Spacecraft yaw flips (see Sect. 3.1.8) are hence performed periodically so that this radiator, which faces the $+y$ direction, always faces away from rather than toward the sun.

Because the KaRIn measurement is very sensitive to variations in the orientation of the baseline, the instrument has its own inertial reference unit (IRU) with high-performance fiber-optic gyros that sense angular rates of the KaRIn body frame (KMSF) relative to an inertial frame. The KaRIn gyro data (along with the spacecraft star-tracker data) are used to reconstruct the estimate of the KaRIn attitude during ground processing, and the attitude knowledge is used to compensate for deviations from ideal spacecraft attitude control. The reconstructed attitude is provided in the ATTD_RECONST product.

9.2 On-Board Processing

The digitized data from the KaRIn ADCs is too voluminous to downlink, so the KaRIn digital electronics process the data in real time to make the downlink data volume more manageable. The measurement objectives and targets are different for oceanography and terrestrial hydrology, so the on-board processor (OBP) ultimately produces separate data streams that are tailored to these two science objectives. A low-rate (LR) stream involves the on-board generation of interferograms at approximately 250 m sampling and 500 m resolution. This processing enables a dramatic reduction in data volume and is suitable for wide, relatively flat targets such as oceans. A high-rate (HR) stream provides data in a rawer form that can be processed to finer resolution (approximately 5 m in azimuth and 0.75 m in slant range) during ground processing. Such data are needed to identify features such as rivers and lakes with widths around 100 m, which would be too small to observe with the LR data. LR data are downlinked continuously (including over land), while HR data are downlinked only a fraction of the time when the spacecraft is over desired target areas (see Sect. 9.3). This section provides a basic overview of the KaRIn on-board processing. For additional details, see [49] and [50].

9.2.1 LR OBP

KaRIn LR on-board processing entails the following steps:

- Range compression (see Sect. 8.1.3). The baseband echo data for each channel are correlated with range compression reference functions (RCRFs) that are stored on board. The convolution operations are implemented via frequency-domain multiplication using fast Fourier transform (FFT) techniques.

- Coregistration of the $+y$ and $-y$ channels (see Sect. 8.3.3). A sinc interpolator is used to non-uniformly resample one channel so that it is approximately aligned in range with the other channel for each of the H and V swath sides.
- Unfocused SAR azimuth compression (see Sect. 8.2.2). The OBP forms nine Doppler beams for each channel by taking an along-track discrete Fourier transform over apertures of nine consecutive pulse echoes (per side) for every range bin.
- Interferogram formation (see Sect. 8.3.1). For each beam on each side, one channel is multiplied by the complex conjugate of the other channel.
- Flattening (see Sect. 8.3.4). The range-dependent phase variation expected for an assumed, locally spherical Earth surface is removed from the interferogram by multiplication with the complex conjugate of a complex signal whose magnitude is constant and whose phase represents the so-called “flat-Earth” phase.
- Averaging and decimation. The flattened interferograms for each beam and side are averaged and decimated in both range and azimuth. The range averaging kernel is range dependent in order to give approximately constant resolution in cross-track. The azimuth averaging is over the interferogram samples of a given beam, side, and slant-range bin over multiple successive apertures. The resolution after averaging is approximately 500 m in both the along-track and cross-track directions, and the decimated data are sampled at approximately 250 m in both dimensions. The data are oversampled by roughly a factor of 2 in order to allow for additional filtering that can better reject narrowband spectral artifacts related to spacecraft disturbances that would alias into the desired signal band upon further decimation. A Blackman-Harris window is used for azimuth averaging, and a Parzen window is used for range averaging.

The processing algorithms for the steps above are designed to fit the available resources for real-time, on-board computing. As such, the algorithms are not necessarily those that might be selected for ground processing with fewer computational constraints and with the benefit of precise, ground-processed reconstructions of the spacecraft attitude and ephemeris. Rather, the primary objective of the on-board algorithms is to reduce the data volume for downlink without introducing additional, rapidly varying random error, which would be irreversible in later ground processing. The on-board processing does introduce systematic (slowly varying) biases in the interferometric phase that would cause errors if left uncompensated, but corrections for these effects are applied during later ground processing. The OBP does not compute surface heights or geolocations.

9.2.2 HR OBP

KaRIn HR on-board processing entails the following steps:

- Doppler centroid removal (see Sect. 8.2.2.1). A phase term that is constant per pulse and piecewise linear from pulse to pulse is applied in order to shift the sampled (but possibly

aliased) Doppler frequency to zero. This is needed for later azimuth presumming. Information on the phase that was removed is downlinked so that the phase can be re-inserted during ground processing.

- Range filtering and resampling. The 300 MHz ADCs oversample the signal by a factor of 1.5 (see Table 9.1), so the complex data are filtered and resampled to the 200 MHz bandwidth of the signal in order to minimize the downlink data volume.
- Azimuth presumming. Presumming is equivalent to applying a low-pass filter and resampling to a lower sample rate in azimuth for each range bin in order to reduce the downlink data volume. Presumming therefore reduces the Doppler bandwidth and removes noise. The nominal presum factor is 2.125, so after presumming, the echo data for each channel are effectively sampled at 2080 Hz rather than the 4420 Hz of the PRF. The reduction in Doppler bandwidth is irrecoverable, so the ultimate height estimates from the HR data stream have inherently higher random error than the eventual height estimates from the LR data stream when both are spatially averaged to the same resolution. However, the HR data can be processed to finer resolution (as defined by the width of the PTR) because they are not averaged on board.
- Block floating-point quantization (BFPQ). BFPQ encoding reduces the data volume by storing only a limited number of bits for each sample [51]. Blocks of data are scaled by a common multiplier so that the bits that are stored best capture the dynamic range of the data. The KaRIn HR data from the OBP are encoded with 3 bits for the mantissa of each real sample. While BFPQ encoding is lossy as a data compression technique, the quantization error in the interferometric data is relatively modest—particularly when compared to other noise sources. Note that a given target on the Earth surface is spread over the depth of the transmit pulse in range and over the presumed width of the antenna illumination in azimuth, so the quantization errors associated with that target are spread over a very large number of samples and tend to average out upon SAR focusing.

Note that SAR data before interferogram formation behave statistically in a manner that is very similar to white noise. It is only after the correlation between the two interferometric channels is considered that the useful information is revealed. This is why on-board LR processing extends through interferogram formation before averaging is performed. This is also why typical data compression approaches do not work well on single-channel SAR data.

9.2.3 OBP Doppler-Centroid Estimation, Tables, and Processing Intervals

Both the LR and HR OBP algorithms rely on estimates of the Doppler centroid (see Sect. 8.2.2.1). For the LR data stream, the Doppler centroid estimate determines the directions toward which the nine beams are pointed (the center beam is pointed toward the Doppler centroid). For the HR data stream, the Doppler centroid determines the center of the Doppler band that is preserved by the azimuth presumming.

The OBP estimates the Doppler centroid by averaging pulse-pair products. That is, each range bin of a given pulse echo (before filtering or presumming) is multiplied by the complex conjugate of the preceding pulse echo such that the phase of the complex result is an estimate of

the pulse-to-pulse phase difference. These complex values are averaged both in range and over successive pulse-pair products within a block of data in time or along-track in order to reduce noise in the estimates. The pulse-pair phase estimate ϕ_{pp} can then be related to the Doppler centroid f_{DC} through

$$f_{DC} = \frac{\phi_{pp}}{2\pi} \text{PRF}, \quad (9.1)$$

although the estimate of the Doppler centroid is necessarily wrapped (aliased) into the span from $-\text{PRF}/2$ to $+\text{PRF}/2$. The values of the pulse-pair phase estimates ϕ_{pp} that were used for on-board computations are downlinked in the KaRIn telemetry. The sign conventions may depend on the data source, so users should consult the appropriate documentation that is specific to the data.

Separate estimates of the Doppler centroid are used for LR and HR on-board processing. The OBP uses the same general algorithm for both the LR and HR estimates, but the parameters used for the estimation algorithm are allowed to differ between the two. This is because HR data are typically collected over land, which generally has a lower SNR than ocean, so the HR estimate entails more averaging.

The parameters of several OBP algorithms depend on the spacecraft altitude above the Earth surface. These algorithms include LR coregistration, flattening, and range averaging. In addition, the data window position (DWP), which determines the span of HR range bins to be stored for each pulse echo, must be adjusted in order to account for the variations in the topography of the Earth surface. KaRIn handles these variations by storing a set of tables on board with pre-computed profiles of parameters (e.g., the PRF) that should be used for the radar data collection and the on-board processing. These profiles cover the orbit repeat cycle and are given as functions of “table time,” which is the time relative to the start of the repeat cycle for the assumed reference trajectory. The OBP steps through the tabulated profiles by advancing the table time as the actual clock time changes. The OBP makes table-time adjustments as necessary to keep its location along the reference trajectory assumed by the tables synchronized with the actual spacecraft position as reported by real-time information from DORIS (see Sect. 4.4.4). The values of some on-board tables are adjusted periodically as needed by ground operators.

On occasion, the KaRIn estimate of table time might become too different from the real-time ephemeris provided by DORIS or from the actual spacecraft position over the Earth. This might happen, for instance, after maneuvers or if DORIS updates become unavailable. In such events, the KaRIn data may contain artifacts, or the data may be completely invalid. For example, the captured data window may not even contain the echo from the surface. Data during these periods would typically be flagged in the KaRIn products. In order to get back to a nominal state, the KaRIn estimate of table time would need to be re-synchronized with the actual spacecraft position through a ground command. Such ground commands are called “U R Here” (i.e., “you are here”) or URH commands. These are planned operationally to occur shortly after maneuvers in order to minimize data loss.

Table 9.4 summarizes the definitions, approximate durations, and equivalent along-track ground distances of important OBP processing intervals. The transmit repetition interval (TRI) is the time between successive transmit events, which are directed towards opposite swath sides and at opposite polarizations. The pulse repetition interval (PRI) is the time between successive transmit events to a given swath side or at a given polarization. The aperture length (AL) is the number of pulse echoes for a given swath side over which the OBP performs unfocussed

Name	Definition	Duration	Ground Distance
TRI	Transmit repetition interval [1/(8840 Hz)]	113 μ s	0.723 m
PRI	Pulse repetition interval per side [2 TRI]	226 μ s	1.45 m
AL	Aperture length for LR [9 PRI]	2.04 ms	13.0 m
DW	DWP boundary [36 PRI]	8.14 ms	52.0 m
PB	Azimuth pixel posting boundary [18 AL]	36.7 ms	234 m
AV	Calibration averaging and PRF boundary [20 PB]	733 ms	4.68 km
D (LR)	LR Doppler estimation block [5 AV]	3.67 s	23.4 km
D (HR)	HR Doppler estimation block [25 AV]	18.3 s	117 km

Table 9.4: KaRIn OBP processing interval durations and equivalent along-track distances on the Earth surface. These values assume a nominal PRF (per swath side) of 4420 Hz and nadir-point along-track speed of 6.389 km/s.

SAR processing for the LR data stream. The DWP is allowed to vary only at DW boundaries every 36 PRIs. The LR OBP averages the interferogram samples in azimuth for each Doppler beam and range bin over 18 ALs, which define an azimuth pixel posting boundary (PB). Loopback signals are averaged over blocks of 20 PBs, which define a calibration averaging boundary, or AV. The PRF is only allowed to change at AV boundaries. Finally, the Doppler centroid is estimated over along-track blocks of data that are integer numbers of AVs. Since the beginning of the calibration phase, the number of AVs used for the Doppler-centroid estimate has been 5 for LR and 25 for HR; different values were used during commissioning, but, as of the beginning of the science phase, there is no plan to change these values again.

9.2.4 FPGA Reconfigurations

Much of the KaRIn on-board processing occurs in reconfigurable field-programmable gate arrays (FPGAs), which are electronic chips whose firmware implements the special-purpose, real-time processing described in Sects. 9.2.1 and 9.2.2. The FPGAs are somewhat susceptible to upset by ionizing radiation in the space environment, so they are reset (reconfigured) periodically in order to clear problems that may have arisen due to radiation events (also known as single event upsets or SEUs). KaRIn data are not available while the FPGA reconfigurations are occurring, however. The reconfigurations typically take around 90 s and occur several times per day. The reconfigurations are usually done while the spacecraft is at high latitudes (near $\pm 78^\circ$) in order to avoid coverage gaps at lower latitudes. FPGA reconfigurations therefore usually occur near pass boundaries.

9.2.5 Radiation Hits

While the FPGA reconfigurations described above are intended to mitigate the effects of radiation, radiation events (“hits”) can nevertheless introduce problems. Radiation effects can manifest in different ways in the KaRIn data depending on what part(s) of the FPGAs are hit. In some cases, inconsistencies in the data cause the ground processing to discard or to flag the affected

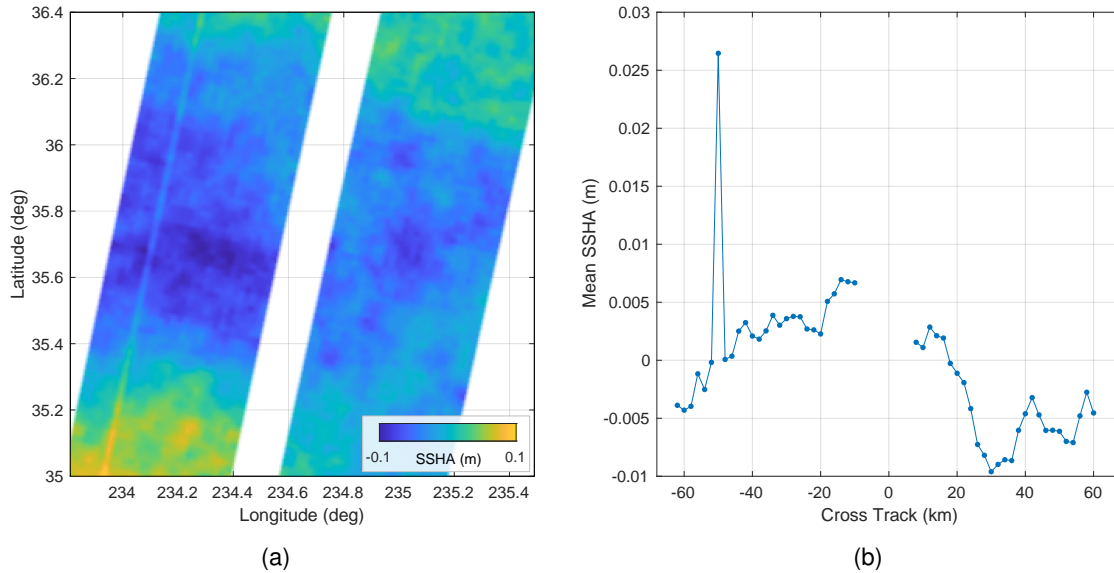


Figure 9.3: Example of the effect of a radiation hit on L2_LR_SSH data: (a) SSHA over a limited, arbitrarily chosen along-track extent; (b) mean of the SSHA in (a) over along track as a function of cross track. Constant and linear terms in cross track have been removed manually from the SSHA in this example in order to better illustrate the effect of the radiation hit. The radiation hit in this case causes a spike in the SSHA of nearly 3 cm at -50 km cross track that persists in along track. Only *good* measurements are shown. This case comes from forward-processed (CRID PIC0) data for cycle 010, pass 039, acquired on 2024-01-26.

data. This may cause coverage gaps. In other cases, the data may contain artifacts such as aphysical along-track “spikes” in the LR height measurements that are not flagged. The effects, whether missing data or artifacts, often persist until the next FPGA reconfiguration (see Sect. 9.2.4), so they can affect several consecutive passes. As FPGA reconfigurations usually occur at high latitudes near pass boundaries, and no data are available during an FPGA reconfiguration, the effects of a radiation hit may seem to disappear between one pass and the next.

An example of the effect of a radiation hit on the SSHA from the L2_LR_SSH product (see Sect. 6.9) is shown in Fig. 9.3. Note, however, that the number of spikes, their cross-track locations, and their magnitudes vary from one event to another.

Radiation events have not been observed to affect HR data directly (recall that there is much more extensive LR processing within the KaRIn FPGAs than HR processing), although it is possible that LR data that are affected by radiation events may affect the crossover corrections (see Sect. 9.4.2) upon which HR products rely.

9.3 HR Mask

SWOT cannot collect KaRIn HR data continuously because doing so would exceed the downlink capacity of the mission. Consequently, HR data are only collected over selected target areas. HR data collection is controlled by a so-called “HR mask” or “downlink mask” that is built offline on the ground then uplinked to the instrument and stored on board. Similar to other KaRIn on-board tables, the HR mask contains a series of times that HR data collection should be started and stopped relative to the progression of the spacecraft along its reference track within a repeat cycle (see the description of table time in the Sect. 9.2.3).

The cumulative duration of HR data collection must not exceed the storage and downlink capacity of the spacecraft. The HR mask is also constrained such that the minimum time between starting HR data collection and stopping HR data collection is 2 s. The minimum time between stopping HR data collection and starting the next HR data collection is 4 s. The KaRIn instrument can support either continuous HR data collection (but the downlink is insufficient for this) or no HR data collection at all.

The HR mask can be changed occasionally by the operations team. For example, different HR masks can be used to optimize the data collection for seasonal changes in the Earth surface. An HR mask is typically built through collaboration of the SWOT operations team and science leads. While target areas for HR data collection may be selected geographically (i.e., based on their latitude and longitude), the mask must eventually be translated into a set of HR start and stop times that meet the constraints described above. Special HR masks can also be used in the event of problems like the extended loss of a downlink station (see Sect. 5.5). See the [Mission Events](#) website for a timeline of changes to the HR mask.

KaRIn is able to collect HR data anywhere along the orbit, including ocean, Arctic, and Antarctic regions. The HR mask used for the calibration phase of the mission included many ocean and ice areas for calibration and checkout purposes. It is not possible to collect data for only one swath side or for only a portion of the swath in cross track.

9.4 KaRIn Ground Processing Design

The SWOT science products and the architecture of the SWOT operational ground-processing system that generates those products are described in Chs. 6 and 7. This section provides some additional background on the rationale for that architecture of products and processing implementation in the context of the SWOT KaRIn measurement.

9.4.1 KaRIn LR Processing

The starting point of KaRIn LR ground processing is the flattened, averaged interferograms (and power images) for each of the nine beams formed by the OBP (see Sect. 9.2).

While the OBP performs several of the early steps in a typical interferometric processing chain, however, the OBP has limited real-time knowledge of important parameters such as the spacecraft position, velocity, and attitude. Moreover, the OBP is constrained by the types of algorithms that can be implemented in real time onboard the spacecraft, with limited storage for surface maps and other such auxiliary information. Consequently, the primary objective of the

OBP is to reduce the data volume for downlink without introducing extra decorrelation, not to avoid introducing biases in the interferogram phase. The LR ground processing compensates for the approximations assumed by the OBP and removes biases introduced by the OBP in order to achieve the accuracy required of the SWOT data products.

The LR processing accomplishes this task by effectively simulating the measurement that the OBP would have made for an assumed reference surface on Earth. This simulation is possible because the details of the OBP algorithms and parameters are known by the LR ground processing, so the ground processing can compute faithful estimates of what the OBP would have produced had the true Earth surface matched the reference surface. The simulation also uses ground-reconstructed estimates of the spacecraft ephemeris and attitude that are generally superior to the real-time estimates available to the OBP. The difference between the downlinked interferogram phase from the OBP and the simulated phase (for the assumed reference surface) computed on the ground then provides information on the height of the true surface above or below the reference surface. In effect, the imperfect OBP flattening (see Sect. 8.3.4) is undone, and the LR ground processing flattens the interferogram with respect to the LR reference surface. The interferometric phase is assumed to be within $\pm\pi$ radians of the reference surface (no explicit phase unwrapping is performed).

Once the interferograms for each of the nine beams have been thus corrected, height reconstruction and geolocation can proceed for each beam. Because each beam has its own Doppler, the sampling grid of each beam is different from those of other beams. The geolocated measurements for each beam are therefore resampled to a common grid, which is defined as the native grid of the central beam. Once the beams are on this common (but not necessarily regular) grid, the beam measurements are averaged in order to reduce noise. This is called beam combining. Spatial averaging and resampling to a regular, geographically fixed grid follow.

Corrections for various geophysical effects are computed from models and/or input information (for example, radiometer measurements of the wet troposphere delay) so that the measurement represents the true height of the Earth surface.

The LR processing also uses the nine-beam echo power from the OBP in order to estimate the surface reflectivity and the wind speed. Additionally, the LR processing estimates the volumetric decorrelation (see Sect. 8.4.3) in order to infer significant wave height from the KaRIn data.

Additional details of the KaRIn LR processing algorithms can be found in the relevant KaRIn ATBDs.

9.4.2 Crossover Calibration

The KaRIn measurement is extremely sensitive to errors in the knowledge of the baseline roll as well as variations in the interferometric phase due to drifts in the radar dual-channel receiver, as described in Sect. 8.3.6. Such errors in the roll knowledge and/or phase are usually indistinguishable from each other, so they are often treated together and called roll/phase errors. Roll/phase errors result in cross-track tilts that may vary in the along-track direction.

Thanks to the high-performance sensors for estimating the spacecraft attitude and the excellent stability of the KaRIn instrument, variations in roll/phase errors in the along-track direction at spatial scales on the order of thousands of kilometers or shorter can be minimized or compensated during KaRIn LR processing on the ground. However, slow variations in the roll/phase

error at scales of thousands of kilometers or longer remain. These residual tilts are large (typically meter scale) compared to the KaRIn measurement precision. Consequently, KaRIn ground processing relies on crossover calibration in order to estimate these residual roll/phase errors.

Crossover calibration entails the empirical estimation of residual roll/phase errors in the KaRIn data by comparing ground-processed LR height measurements from ascending and descending passes that are near each other in time and over the same geographic area on the Earth surface. At crossover points, where the nadir tracks of ascending and descending passes intersect, the same area on the Earth surface is observed from different viewing geometries such that the cross-track direction is different with respect to ground coordinates for the ascending and the descending passes.

For example, suppose (hypothetically) that the nadir tracks from a pair of ascending and descending passes were to cross each other at a 90° angle. In this case, the cross-track direction for one pass would be the along-track direction for the other pass, and vice versa. If one assumes that the surface height has not changed significantly in the time between passes and that the along-track variations of the KaRIn errors vary much more slowly spatially than the swath width, the orthogonality of the viewing geometries would allow one to solve for the roll/phase errors in each of the passes. In reality, the passes do not cross each other at 90° , but the roll/phase errors can nevertheless be estimated based on the same principles.

The KaRIn crossover calibration processing thus examines pairs of ascending and descending passes in the processed LR data at ocean crossovers to estimate the roll/phase error in each pass. With the set of corrections from many crossovers, where each correction corresponds to a particular crossover location and observation time of one of the passes at that crossover, a profile of corrections over time is created. A key assumption is that the variations in the necessary corrections are slow and well behaved in the along-track direction. This is consistent with the power-law form of the SWOT oceanographic requirements for height accuracy.

The crossover processing uses only crossovers over the ocean because LR heights are not sufficiently well behaved over land (or ice) to obtain reliable roll/phase estimates from crossovers there. The processing also only uses ascending and descending passes that are sufficiently near each other in time lest variations in the ocean surface over time confound the roll/phase estimation.

Because the processing does not use crossovers over land, the corrections from the surrounding ocean are interpolated to obtain the corrections over land. These interpolated segments are the portions of the crossover correction that are applicable to most hydrology products. The corrections for inland hydrology therefore typically have larger errors than those for oceanography. Corrections near the middle of long spacecraft tracks over land and/or frozen areas at high latitudes usually have the largest errors. Even in those areas, however, the crossover corrections allow the residual roll/phase errors to be reduced to the level of the SWOT requirements, which are much smaller than the uncompensated roll/phase errors.

Note that the crossover corrections are reported in the L2_LR_SSH product but are not actually applied to the KaRIn surface height measurements. Consequently, users themselves must apply the reported crossover corrections to the reported height measurements from the L2_LR_SSH product (or compute and apply their own corrections) if they desire to avoid the meter-scale cross-track tilts due to residual roll/phase errors. The crossover corrections are applied to the heights in HR products (and also separately reported in the HR products for

information), so users of HR products do not need to apply the corrections themselves.

In addition to roll/phase errors, which manifest as linear cross-track tilts in the height, the crossover processing also estimates constant and quadratic errors. The quadratic errors, which might arise due to varying baseline dilation (baseline length changes), are computed in a manner similar to the linear roll/phase errors. The constant (as a function of cross track) corrections are estimated with respect to SWOT nadir-altimeter data. Therefore, after applying crossover corrections to the KaRIn data, the absolute height of the KaRIn measurements is effectively forced to be consistent with the absolute height of the SWOT nadir-altimeter measurements.

The number and geographic distribution of crossover points is greater for the science orbit than for the calibration orbit, so corrections from the crossover approach described above are generally somewhat better for the science orbit than for the calibration orbit.

9.4.3 KaRIn HR Processing

KaRIn HR processing begins with presumed HR data that have been BFPQ encoded (see Sect. 9.2). The HR processing begins with BFPQ decoding, followed by SAR image formation, which results in high-resolution SLCs for each channel. The SAR processing uses a reference DEM in order to coregister the two SLCs and to reference their phases so that an interferogram formed from the two SLCs is flattened with respect to the DEM.

The HR processing then proceeds to interferogram formation at the full resolution of the SLCs. The interferogram is next spatially averaged in azimuth in order to reduce phase noise in the interferograms and to obtain similar resolution in both the along-track and cross-track directions. The interferograms and power images after this initial stage of azimuth multilooking are coined “rare” quantities by analogy to cooking.

The HR processing continues with a classification algorithm, which uses the radar backscatter of the surface in order to detect water (i.e., to distinguish water from land), assuming that water at the steep SWOT incidence angles is usually brighter (more reflective) than land. The algorithm also classifies dark water (see Sect. 10.6) based on external global information about the prior probability of water on the surface. That is, pixels where water is not detected but where water is expected based on the external water-probability map are classified as dark water.

After classification, the processing adaptively averages water pixels in order to further reduce noise while preserving the ability to observe small water features. The data after adaptive averaging are called “medium” quantities, again referring to a cooking analogy.

Phase unwrapping and geolocation occur on these medium quantities to give the measured surface heights and horizontal locations of each slant-plane pixel that is detected as water. The phase unwrapping algorithm applies a complicated set of heuristics that first spatially unwrap detected water pixels then resolve the absolute ambiguity of each spatially unwrapped region. Crossover corrections are applied to the geolocated data in order to compensate for the large roll/phase errors that would otherwise be present in the measurements (see Sect. 9.4.2).

The HR processing entails the computation of geophysical corrections similar to those of the LR processing (see Sect. 9.4.1), although the corrections of the HR processing are not identical to those of the LR processing. For example, the SWOT radiometer cannot provide reliable estimates of wet tropospheric path delay over land, so the HR processing uses model-based wet-troposphere corrections exclusively. The HR processing naturally does not use ocean-specific

geophysical models such as ocean tide corrections.

Once pixel-level HR processing is complete, pixels with detected water are assigned to river and lake vector objects that represent river reaches and lakes.

Rivers are defined by a prior river database (PRD) that is an input to the SWOT HR processing (see Sect. 6.12). The HR processing assigns pixels to river nodes, and thus also to reaches, based on both the connectedness of pixels classified as water and the proximity of these pixels to the PRD centerlines. The processing then estimates quantities such as water surface area, WSE, and slope for each reach. River discharge is computed as well.

During lake processing, pixels are assigned to lake objects in a manner that is similar to how pixels are assigned to reach objects during river processing. However, whereas river processing gives information only on objects defined in the PRD, lake processing gives information both for lake features that correspond to objects that are defined in the prior lake database (PLD) as well as for lake features that are not defined in the PLD. Note that there is not always a one-to-one mapping between observed lake features and objects in the PLD because lakes may split, merge, be underobserved, etc. Pixels that are assigned to rivers are usually excluded from assignment to lakes (except for objects such as reservoirs that are in both the PRD and the PLD), so the lake processing occurs after river processing is complete. Lake storage change is computed for lakes in the PLD.

After pixels are associated with reach or lake objects, additional averaging over each object can give more precise horizontal geolocations, albeit at the expense of the ability to resolve features smaller than the object over which the averaging occurred. This is called height-constrained geolocation and results in what is sometimes known as the “well-done” layer. Such extensive averaging may be desirable because of the relationship between height errors and cross-track errors [see Eq. (8.114)], in which small height errors may be associated with large cross-track errors.

Raster processing resamples SWOT pixel-level measurement quantities to a regular grid in order to produce a convenient data product for surface water that does not correspond closely to river or lake features. Examples include floods and wetlands. The raster processing respects the classifications of the slant-plane pixels, however, so it involves more than simply regridding the data.

After a significant volume of SWOT HR data have been accumulated, HR processing produces a floodplain DEM product. This processing examines the historical relationship between water height and extent in order to infer the topography of land near water. That is, the processing assumes that the land at the edge of a detected water features has the same elevation as the water at the time of that observation. Therefore, the floodplain is mapped only over the extent to which it is inundated at some point in the set of SWOT observations since launch.

Aside from the floodplain DEM computation, the only HR processing that extends over multiple passes is the simple computation of statistical metrics (mean, minimum, median, and maximum) of quantities for river and lake vector objects over cycles. This occurs at the ends of the river and lake processing chains. All other HR algorithms are designed to operate in a pipelined, operational processing architecture (see Sect. 7.1); they do not jointly process multiple SWOT observations from different passes simultaneously. That is, while many of the algorithms within the HR processing are quite complicated, they generally operate on only one SWOT overpass at a time. In this context, crossover processing (which also uses many cycles of data at a time)

is not included in the set of HR algorithms, though the crossover corrections are inputs to HR processing.

Additional details of the KaRIn HR processing algorithms can be found in the relevant KaRIn ATBDs.

Chapter 10

Phenomenology

This chapter gives examples of interesting phenomenology that users might encounter in SWOT data products. In this context, the term “phenomenology” refers to the characteristics of the way that SWOT senses the Earth surface and their resulting manifestations in the SWOT observations. The intent of this chapter is to provide users with knowledge of SWOT phenomenology so that they are better able to interpret the measurements. This chapter does not necessarily attempt to catalog all applicable phenomenology, however. Rather, it focuses on classes of phenomenology into which additional insight would be most likely to help users understand the SWOT data.

Note that evolutions in the ground processing of SWOT data are expected over the life of the mission. Changes in processing algorithms and parameters may affect the phenomenology somewhat, although major changes in the high-level conclusions are not expected. The particular realizations of phenomenology in some of the example cases shown may change significantly, however.

10.1 Speckle

Speckle is an intrinsic characteristic of SAR images that results in random variations that appear as multiplicative or “salt and pepper” noise [52]. An example is shown in Fig. 10.1. It also occurs in other contexts as well [53]. Speckle arises because the various scattering features within a typical resolution cell on the Earth surface all contribute to the received echo, but they do not necessarily all add coherently in phase. That is, suppose that a resolution cell contains a large number of point-like scattering elements that are randomly distributed spatially over the surface; this is usually an excellent model for many natural surfaces on Earth. Because the range to each scattering element varies at scales smaller than the radar wavelength, the phase of the echo contribution from each scattering element is random with respect to the phases of other scattering elements. Consequently, the echo contributions from the various scattering elements may add either constructively (in phase) or destructively (out of phase). The total echo from the resolution element therefore has a random phase that is uniformly distributed between 0 and 2π rad. The expected value of the power of the total echo is related to the surface reflectivity. However, the variance of the power is equal to the mean power, so individual realizations of power from different resolution cells with similar mean backscatter are very noisy.




Figure 10.1: Example of speckle on the Willamette River. The image shows the relative (arbitrarily scaled) power from the $-y$ SLC channel in the slant plane without any spatial averaging. Slant range increases to the right, and azimuth increases down the page, so north is approximately downward and slightly to the right. Because the resolution and sampling in azimuth are much finer than in range, the image appears stretched vertically. The river is the bright feature that runs from the upper left of the image to the lower right. It is approximately 100 m wide at this location. Note the multiplicative noise in the power due to speckle. The image is centered near $+44.234^\circ$ latitude and -123.163° longitude. It is 150 pixels wide and 300 pixels tall and is a small portion of reprocessed (CRID PGC0) L1B_HR_SLC tile 232L from cycle 482, pass 013, acquired on 2023-04-06.

More formally, the statistical distribution of the complex echo for a given SLC SAR sample of a typical, rough surface can usually be modeled very well by a circular Gaussian or complex normal distribution, which has independent, zero-mean Gaussian distributions for the real and imaginary parts of the echo. The distribution of each of the real and imaginary parts has a variance that is half the expected power P of the echo such that the variance of the complex distribution is equal to the expected power of the echo. The power of the echo therefore has an exponential distribution whose mean is P , and the phase of the echo has a uniform distribution between 0 and 2π rad. The magnitude (i.e., the square root of the power) has a Rayleigh distribution. Consequently, a large fraction of the samples in a single-look SAR image have echo magnitudes that are much smaller than the average echo magnitude. These samples are hence very susceptible to corruption by the additive receiver thermal noise.

Even without additive receiver thermal noise, speckle makes the power of an SLC image difficult to use for inferring the surface backscatter without additional averaging to reduce the variance. Spatial averaging (“taking looks” or “multilooking”) is hence usually performed for this purpose. Note that because the speckle-induced variance of the power estimates increases with the mean, increased radar transmit power or higher surface reflectivity do not necessarily improve the relative precision of the backscatter estimates.

Speckle might sometimes be considered a nuisance, but it comes about because SAR images are coherent, and this is the same property that enables interferometry. While the phase in a given SAR image is effectively random from one resolution cell to the next, the two SAR images from an interferometer such as KaRIn are highly correlated with each other. That is, the two SAR images exhibit nearly the same random variations in phase spatially such that the phase difference between the two images (i.e., the interferometric phase) contains information about the imaged surface. Interferogram samples that have low magnitudes often have very noisy phase, however, so spatial averaging of the interferogram samples is usually necessary to obtain useful phase estimates.

Speckle drives several aspects of the KaRIn data processing, both on board the spacecraft and on the ground. The different stages of spatial averaging in both LR and HR KaRIn data (beginning in the OBP for the LR data) are designed around the statistics of speckle in the data. Water detection during HR processing must account for the apparent brightness variations due to speckle, especially when dealing with features at scales near the resolution of the HR data. Finally, speckle influences the ultimate error statistics of virtually all KaRIn measurement quantities.

10.2 Layover

The physical causes of layover with respect to the SAR imaging geometry in general are described in Sect. 8.2.1.3. Layover impacts KaRIn measurements in ways that are somewhat different from other SAR systems, however, given the steep incidence angles of the KaRIn measurement and the fact that the primary water targets of interest are typically much brighter than the land that would often lay over into the water in the images. A more detailed assessment of the impact of layover on the KaRIn measurement is given in [54].

Recall from Sect. 8.2.1.3 that layover occurs when the local surface slope is greater than the incidence angle θ_{inc} . Because the open ocean is generally wide and flat, layover is usually not a concern there (except to the extent that the surfboard effect, described in Sect. 10.3, is similar to

layover). On the other hand, geometric layover is ubiquitous over land, as even gentle undulations in the topography of the land often exceed the KaRIn incidence angle, which is less than 1° at near range and about 4° at far range. Furthermore, small height variations due to vegetation, the banks of rivers, and other such features can contribute to geometric layover even when the land is otherwise relatively flat. For example, with a 1° incidence angle, a river bank that rises 1 m above the water surface would lay over across 57 m of the water surface in the cross-track direction toward near range.

Fortunately, while layover occurs everywhere, its impact is limited. When assessing the impact of layover, it is useful to distinguish the case that land lays over into a water body of interest from the case that a different water feature lays over into the water body of interest. These two cases are called land-water layover and water-water layover in this document. Land-land layover, in which land lays over into other land, is not discussed here given that KaRIn is generally intended to observe water (but note that land-land layover might result in bright features in the data that are falsely detected as water). Land-water layover is much more common than water-water layover because inland water occupies a small fraction of the surface area for typical scenes of interest. Most land surfaces, and vegetation in particular, are much less reflective than water, however, so while the water echo and the land echo cannot be distinguished from each other where land-water layover occurs, the much stronger water echo still usually dominates. Finally, because typical topographic variations over land are a significant fraction of the small KaRIn ambiguity height (see Sect. 8.3.5), the phase of the land echo that competes with the desired water echo is somewhat randomized. Consequently, a significant portion of the error due to land-water layover tends to average out along with other random noise (e.g., from sources such as decorrelation) upon spatial averaging of the KaRIn measurements. Layover therefore does increase the KaRIn measurement errors of inland water features, but the errors are usually not so great as to make the measurements unusable.

Perhaps counterintuitively, areas with the most rugged topography are not necessarily the places with the largest errors due to layover. As the topographic variations become larger than the ambiguity height, the phase contributions from layover become more evenly distributed between 0 and 2π rad such that they are better eliminated with spatial averaging. On the other hand, areas with relatively little topographic variation may in fact exhibit larger errors due to layover because such areas may give more of a bias that persists after spatial averaging.

Naturally, the effects of land-water layover vary a great deal with the specific location and viewing geometry. Conditions that affect the relative reflectivity of water and land (e.g., wind, soil moisture, foliage density, etc.) can also affect the impact of layover on KaRIn height measurements significantly.

The discussion above generally assumes that water is more reflective than land. Layover errors may be considerably worse for dark water (see Sect. 10.6) and water near bright land (see Sect. 10.9).

Where water-water layover occurs, the errors due to layover tend to be much worse than those due to land-water layover. In fact, by definition, the two water features that lay over into each other overlap in the slant plane and are indistinguishable in the KaRIn measurement. Therefore, there is no way that a single KaRIn height estimate for that location in the slant-plane can be correct for both of the water features. Moreover, water-water layover can cause phase unwrapping problems as well (see Sect. 10.10); phase unwrapping problems may affect sur-

rounding water areas that are not necessarily in water-water layover. Errors due to water-water layover are simply an unavoidable consequence of the KaRIn viewing geometry. They usually do not occur so often as to be intractable, however.

Note that the principles of modeling the error due to layover are relatively simple, but computing the effect at a particular location requires precise knowledge of the surface topography and the surface reflectivity for both the water and the surrounding land. Such knowledge is usually insufficient to give accurate predictions of layover-induced errors for specific KaRIn observations (e.g., at the pixel level). On the other hand, statistical characterizations may be feasible.

10.3 Surfboard Effect

The surfboard effect is a pseudo-random error caused by the nonlinear mapping of ocean waves into the slant plane [55]. The effect was so named by the authors of [55].

Ocean waves (surface gravity waves) are of course everywhere in the ocean, and they almost always have heights that are orders of magnitude larger than the KaRIn measurement precision. They are consequently an important potential contributor to measurement error in spaceborne altimetry. In order to minimize random errors caused by waves when estimating the SSH, the KaRIn measurements are spatially averaged over scales that are longer than the typical wavelengths of ocean waves.

Recall from Sect. 8.2.1.3, however, that variations in the height of a surface imaged by a SAR system cause foreshortening and possibly layover in the slant-plane SAR image. Foreshortening and layover become more significant as the local surface slopes approach or exceed the incidence angle of the SAR. For the steep SWOT viewing geometries, the image distortions due to foreshortening and layover can be significant even for the slopes associated with ocean waves. That is, ocean waves map in a nonlinear manner from the ground plane into the slant plane during the KaRIn measurement process such that a hypothetical, spatially monochromatic wave on the surface would not appear as a monochromatic wave in the SAR images; the SAR images would include spectral components (wavenumbers) that do not exist on the ground. Because of these nonlinearities in the slant-plane sampling of the surface, the observation may contain more energy at low wavenumbers than exists in the true wave field. These low-wavenumber spectral components may not be removed very well by the spatial averaging (low-pass filtering) applied during the generation of KaRIn data products. The residual error is known as surfboard error.

The specific manifestation of the surfboard effect depends on the realization of the wave field. The statistical behavior depends on the amplitude, wavelength, and direction of the waves (i.e., it depends on the full 2-D wave spectrum). Surfboard errors tend to increase rapidly with increasing SWH. They usually appear as pseudo-random errors that have a relatively flat along-track spectral shape up to some upper wavenumber falloff that may be lower than the filter cutoff of the KaRIn processing. Examples of SSHA along-track spectra showing the surfboard effect are given in Fig. 10.2. The surfboard errors are visible as the “humps” or “shoulders” in the blue curves, which show spectra without additional cross-track averaging beyond that performed on-board the spacecraft. The surfboard effect increases the SSHA error variance so that it is larger than what would be predicted from the interferometric coherence alone. Additional quantification of surfboard errors is discussed in [56].




Figure 10.2: Examples of along-track spectra of KaRIn SSHA showing surfboard errors. The upper and lower panels are for cases with 4.1 and 6.1 m SWH, respectively. In each panel, the vertical axis is the SSHA power spectral density on a logarithmic scale with one order of magnitude per vertical division. The horizontal axis is the along-track wavenumber in cycles per kilometer. The blue and orange curves show the spectra at the native 500 m cross-track resolution of the unsmoothed LR data and after smoothing in cross-track to 7.5 km resolution, respectively. The purple and red curves are the expected noise floors corresponding to the orange and blue curves based on the interferometric correlation and the along-track filter characteristics. The dashed black lines show the KaRIn instrument-level error allocation at $1.667 \text{ cm}^2/(\text{cycle}/\text{km})$, assuming 7.5 km cross-track averaging. The additional error above the white-noise floor in the blue curves (indicated by the dashed gray ovals) are caused by the surfboard effect. The spectra are computed from forward-processed (CRID PIA1) L2_LR_SSH Unsmoothed data over cycles 567 and 568, collected on 2023-06-29 through 2023-07-01.

10.4 Sea State Bias

The sea state bias (SSB) is a well-known effect in nadir altimetry (see, for example, Sect 1.5.1 of [57]). It is a bias in the radar-derived height measurement of the SSH due to the skewness of the heights of the ocean waves, an electromagnetic bias (EMB or EM bias), and a tracker bias.

The SSB for the SWOT nadir altimeter is similar to the SSB for other similar altimeters; a great deal of work has been devoted to the characterization of and correction for the SSB in this context, yet it remains an active field of research because it is often the largest source of residual error. For the KaRIn interferometric measurement on the other hand, the investigation of SSB effects is in its infancy as of the beginning of the SWOT science mission.

KaRIn does not rely on retracking as nadir altimeters do, so tracker bias is not directly applicable for KaRIn. KaRIn measurements are expected to be biased by wave skewness and EMB. However, the EMB for KaRIn may differ due to the off-nadir incidence angles of the KaRIn measurement. The EMB generally arises because the troughs of ocean waves are more reflective than the crests, thereby biasing the radar measurement toward lower SSH than the mean of the physical surface height. The differences in reflectivity are sensitive to incidence angle, however, so while models for the EMB from nadir altimetry provide a good starting point for the EMB in KaRIn observations, adjustments for KaRIn are under investigation.

Moreover, the KaRIn measurement uses SAR processing, which is sensitive to the motion of the ocean waves. Motion effects can couple to the interferometric measurement as described in Sect. 10.11. As different portions of ocean waves move differently, KaRIn SSB and motion effects are related. Research on this topic is in progress.

Evolutions in the understanding of and the correction for KaRIn SSB effects are likely throughout the duration of the SWOT mission and beyond.

10.5 Rain Attenuation

Ocean and inland-water scenes may sometimes exhibit very low KaRIn echo power and hence poor SNR due to attenuation of the radar signal from rain. An example of rain attenuation over the Pacific Ocean is shown in Fig. 10.3. The image shows the estimate of the backscatter coefficient σ° . Rain attenuation is evident as the dark areas with low backscatter estimates. Note that the ocean surface itself is not necessarily less reflective, but the KaRIn estimate of σ° is low because the signal is attenuated by rain in the atmosphere above the surface, and the KaRIn data processing to compute σ° from the signal power does not have sufficient information to compensate for the rain attenuation. The processing therefore gives an erroneously low σ° estimate. The example of Fig. 10.3 is for LR data over the ocean, but rain may affect HR data over land as well.

While there is a rain flag in the L2_LR_SSH product, its temporal and spatial resolution are limited, so it should not be relied upon to definitively indicate the presence of rain in the SWOT data. Rain attenuation in HR data may appear similar to dark water (see Sect. 10.6) and may be flagged as such.

Note that the term “rain” is used loosely in this context to refer to precipitation (condensed water in perhaps different forms) more generally. The KaRIn radar signals are not as strongly




Figure 10.3: Example of attenuation due to the presence of rain in KaRIn data over the Pacific Ocean. The image shows the estimate of σ^0 for cycle 577, pass 026 around $+8.6^\circ$ latitude and 239.8° longitude on 2023-07-10. The along-track direction increases down the page, while the cross-track direction increases toward the right. No estimate is reported close to nadir (black vertical bar in the middle of the image) or at far range on either of the left and right sides. The image covers approximately 600 km in along track and 130 km in cross track. The data come from a reprocessed (CRID PIB0) L2_LR_SSH product with 2 km resolution. Errors in the absolute radiometric calibration are likely present in this image; the image is for illustration of the relative spatial variations in reflectivity only.

affected by clouds or water vapor because the attenuation increases very rapidly as the characteristic size of the scattering objects (e.g., the rain-drop diameter) becomes a greater fraction of the radar carrier wavelength.

10.6 Dark Water

KaRIn performance relies on water to be bright, or highly reflective of the transmitted radar signal. This is required to achieve adequate SNR for both ocean and inland hydrology measurements. Additionally, water detection for HR data over land assumes that water is brighter than the surrounding land. Water—especially inland water—can sometimes be dark, however. Here, “bright” and “dark” refer to the average, relative, backscatter of the radar signal from water surfaces over some limited, local area for a given observation. The radar backscatter is quantified by the NRCS or σ° (see Sect. 8.1.4); bright areas have higher relative values of σ° while dark areas have lower relative values.

One common reason for water to be dark in KaRIn observations is that the water surface is too smooth and the radar reflections are therefore too specular. This is often associated with low-wind conditions near the water surface. When the water surface is “glossy,” it behaves much as a mirror would for visible light. At nadir, where the transmitted radar signal would be incident upon the mirror-like water surface at normal incidence, the echo would be very strong. However, at off-nadir angles, such as those of the KaRIn measurement, a specular reflection would cause the transmitted signals to bounce away from the radar rather than back toward it. The radar echoes would therefore be low, and the water would appear dark in the SAR images. Roughness of the water surface causes the surface to scatter the radar signal in a somewhat more diffuse manner so that the energy is spread over a small span of angles that includes the direction back toward the radar. The water surface would then appear bright in the SAR images.

The relationship between surface reflectivity and wind speed is illustrated in Fig. 10.4. The curves in this figure show the notional backscatter coefficient (NRCS or σ°) as a function of incidence angle θ_{inc} for a variety of wind speeds. As the wind speed decreases, the scattering becomes more specular. That is, for lower wind speeds, the curve of σ° vs. θ_{inc} becomes more sharply peaked near zero, going to higher values at normal incidence ($\theta_{\text{inc}} = 0$) but falling off more rapidly with incidence angle. If the falloff is too rapid compared to the incidence angles of the KaRIn measurements, the water appears dark.

Of course, if the surface is too diffuse, the scattered signal is spread across such a wide span of angles that the echo energy that returns to the radar is again very low. This is part of why land surfaces, which are usually very rough compared to the Ka-band wavelength, are usually darker than water surfaces. These relationships are fundamental to the HR classification algorithms (see Sect. 9.4.3). Note also that the HR processing algorithms attempt to identify the presence of dark water based on prior maps of the probability of water; that is, if water is not detected in the SWOT observation in an area where water would have been expected based on external information, the processing classifies the area as dark water.

An example of dark water due to specular scattering is shown in Fig. 10.5. The panels in this figure show the same area on the ground for KaRIn observations on four different days. The Connecticut River, Quabbin Reservoir, and specular ringing from nadir echoes (see Sect. 10.7) are labeled in panel (a). While the wind conditions are not necessarily identical across the entire




Figure 10.4: Example behavior of the NRCS (i.e., σ_0 or σ°) as a function of incidence angle for different wind speeds at Ka band over inland water. The wind speed is indicated by the legend. This plot is intended to illustrate the notion that low wind speeds correspond to specular scattering that can give low backscatter at KaRIn incidence angles. The solid, vertical black lines indicate the approximate KaRIn incidence angles at 10 and 60 km cross track. The curves come from a model that is used here only for conceptual visualization purposes; the values should not necessarily be taken to be representative of Ka-band backscatter in any quantitative sense. They may not even be reliable qualitatively beyond the idea that water becomes more specular at low wind speeds.

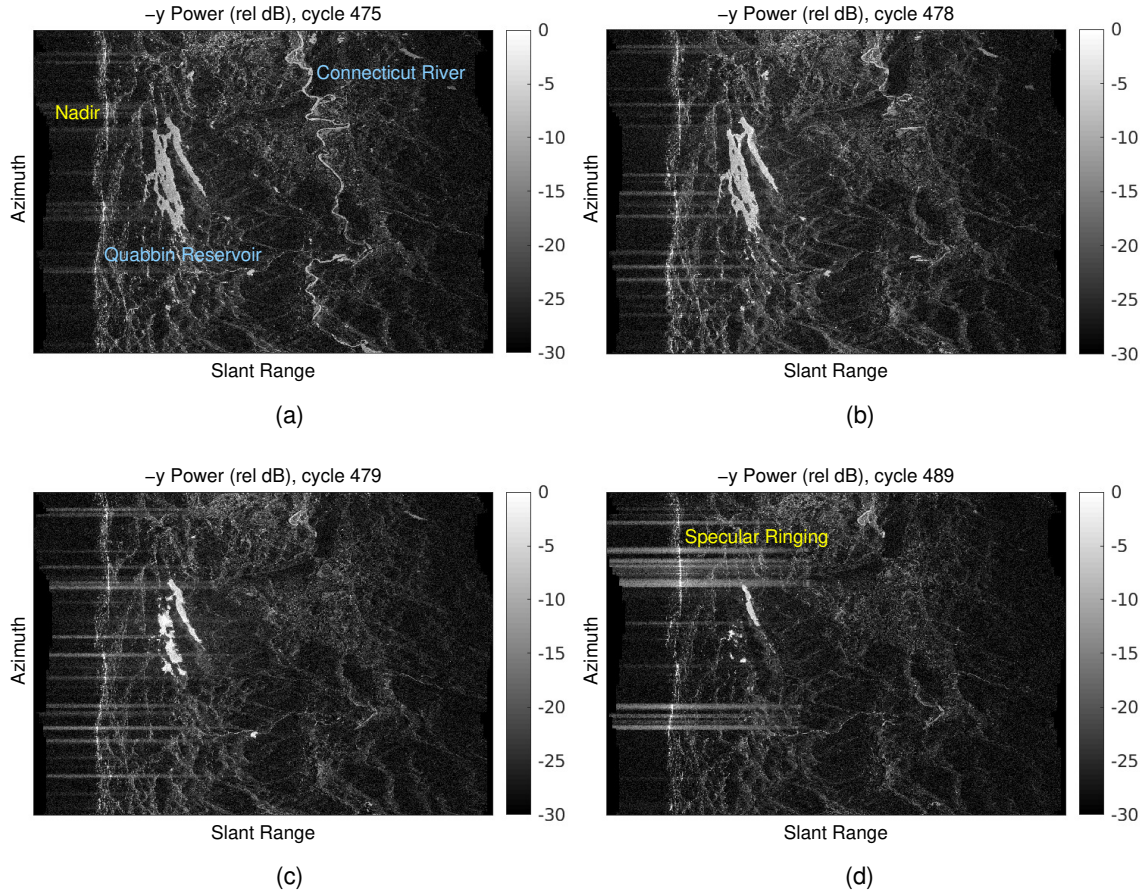


Figure 10.5: Examples of dark water on the Connecticut River near Amherst, Massachusetts from different days. All panels show the relative (arbitrarily scaled) $-y$ power in decibels for the same area on the ground (pass 009, tile 229L). Panels (a), (b), (c), and (d) show data for cycles 475, 478, 479, and 489 on 2023-03-30, 2023-04-02, 2023-04-03, and 2023-04-13, respectively. For all panels, slant range increases toward the right, and azimuth increases down the page, so north is approximately downward and slightly to the right. The tile is approximately 70 km in each dimension. The Connecticut river is bright and is visible approximately two thirds of the way toward far range in (a), but it is mostly dark in the other images. The Quabbin Reservoir is visible as the large, bright feature toward the left (near-range side) of the tile in (a) and (b), but it is partially dark in (c) and mostly dark in (d). Specular ringing (see Sect. 10.7) is also readily apparent in all panels. The data come from reprocessed (CRID PGC0) L1B_HR_SLC files with 7 looks in azimuth.

area, the panels are arranged such that the scattering generally appears least specular in panel (a) and most specular in panel (d). In panel (a), both the river and the reservoir are moderately bright, and some specular ringing (see Sect. 10.7) is evident. In panel (b), the river is dark while the reservoir, which is observed at a lower incidence angle because it is closer to the nadir track, is somewhat brighter than in (a). There is also more specular ringing in (b) than in (a), which indicates that echoes at nadir are stronger in (b). The reservoir is partially bright and partially dark in panel (c), possibly due to spatial variations in the wind or other factors that contribute to roughness of the water surface. Panel (d) exhibits highly specular scattering, in which both the river and the reservoir are almost completely dark while the specular ringing is the strongest of the four cases shown here, indicating the brightest returns at nadir.

One may observe that the brightness of the land in Fig. 10.5 is approximately the same in all four panels, indicating that the differences in the brightness of the water from one panel to the next is not due to rain attenuation (see Sect. 10.5).

Note that scattering from the ocean is usually not as specular as scattering for inland water can be when low-wind conditions are present because surface gravity waves in the ocean usually contribute to the surface roughness as well. However, the ocean can be specular due to phenomena such as slicks.

10.7 Specular Ringing

Specular ringing (also called nadir ringing or nadir contamination) occurs when the range sidelobes of bright, specular targets near nadir are strong enough to contaminate other targets in the swath. The term “ringing” refers to the decaying response to a strong, impulsive signal. Specular ringing is evident in Fig. 10.5, especially in panel (d).

Recall from Sect. 8.1.3 that the response of a point target after range compression has sidelobes that are much lower than the peak of the main lobe but that extend a significant distance from the peak in either direction in range, as illustrated in Fig. 8.4(b). However, if a target is very bright, both the main lobe and the sidelobes of its response are scaled proportionately. Therefore, a target that is much brighter than neighboring targets may have sidelobes that are as bright as or brighter than the neighboring targets. For example, if the sidelobes of the target response are 20 dB lower than the peak of the main lobe, but a particularly strong target is 30 dB brighter than its neighbors, the sidelobes of the bright target would still be 10 dB brighter than the main lobes of the neighbors.

Areas near nadir may be especially bright because of specular reflections from flat surfaces such as water (see Sect. 10.6 for a discussion of specular reflections). Therefore, the sidelobes from bright nadir echoes may be stronger than the main-lobe echoes of targets further off nadir that are not as bright. Note that the bright echoes from nadir do not necessarily need to come from point-like targets. For example, a water feature at or near nadir may have a finite spatial extent such that it is not necessarily very similar to a single point-like target. However, for most purposes, the water feature would behave like a collection of point-like targets that are distributed spatially over the water surface. Each of those targets might be bright and have its own sidelobes. Collectively, the sidelobes of all the individual targets of the water feature might contaminate areas far from nadir if they are bright enough.

Contamination from specular ringing may cause errors in water detection, which in turn may cause problems for phase unwrapping and assignment of pixels to rivers and lakes. Specular ringing may also increase the height errors for affected water surfaces in the swath as the desired water echoes compete with unwanted nadir sidelobes, which generally have a different interferometric phase than that of the desired water surface.

For the parameters of the KaRIn transmit chirp (see Table 9.1), the sidelobes of a target might stretch up to approximately 1 km from the peak of the main lobe in slant range. Therefore, a target at nadir has sidelobes that overlap with targets all the way out to approximately 40 km cross track given the steep incidence angle (the entire KaRIn swath occupies only about 2 km in slant range). Fortunately, the sidelobe levels relative to the main lobe fall off rapidly with distance from the main lobe, so specular ringing is usually somewhat more contained for all but the brightest of nadir targets.

Specular ringing can be flagged with use of models for the target response during ground processing if very bright echoes are observed. Very strong targets may exhibit nonlinear distortions that limit the reliability of these flags, however. Moreover, the absence of available data where specular-ringing pixels are discarded may nonetheless cause measurement errors, particularly in HR area estimates. Mitigations for the impacts of specular ringing will likely continue to evolve.

It is also worth noting that nadir is not necessarily always the closest point to the radar when there is significant topography (see Sect. 8.2.1.3). Therefore, in a slant-plane image with range increasing to the right, such as shown in Fig. 10.5, the nadir echo is not necessarily the echo from the scene that is furthest to the left (at the shortest range). However, nadir is generally the location at which the look vector is parallel to the direction of gravity and is therefore where a flat water surface would give a strong, specular echo back in the direction of the radar.

Specular ringing usually does not occur over the ocean because the ocean is usually not as flat as inland water can be (e.g., due to the presence of surface gravity waves in the ocean). However, specular ringing over ice sheets is often apparent in KaRIn LR data. Specular ringing may appear somewhat different in LR data than in HR data due to the LR on-board processing performed. Additionally, targets near nadir can sometimes be so bright that they cause overflows in the fixed-point arithmetic in the KaRIn LR on-board processing, particularly in the inverse fast Fourier transform (IFFT) used for on-board LR range compression. These overflows can cause nonlinear artifacts to appear in the LR data. These artifacts are flagged in the L2_LR_SSH product if detected.

10.8 Dark Water Misclassification Due to Reference DEM Errors

As discussed briefly in Sects. 9.4.3, the classification algorithm used during HR processing attempts to identify dark water (see Sect. 10.6). Because dark water is not directly distinguishable from land in the SWOT observations, however, the classification algorithm relies on external maps of the *a priori* probability of water. That is, an area in the SWOT observation is classified as dark water if (bright) water was not directly detected there, but water was expected at that location based on information from the prior probability map. This algorithm relies on the reference DEM as well as information from the SWOT range measurements of nearby detected water in order to determine the 3-D spatial relationships between the ground-plane maps of the

prior water probability and the native slant-plane SWOT observations (see Sect. 8.2.1). Consequently, errors in the reference DEM and/or the SWOT range measurements can cause information from the maps of prior water probability to be misprojected from the ground plane to the slant plane. This may cause artifacts in the classification results. In particular, pixels that do not actually correspond to dark water in the L2_HR_PIXC product may be incorrectly classified as dark water. These classification errors may then propagate to other errors in the L2_HR_RiverSP, L2_HR_LakeSP, and L2_HR_Raster products. No classification is performed during the processing of LR data, so no corresponding issue arises for LR products.

Incorrect knowledge of the height of the ground surface relative to the altitude of the spacecraft causes a shift in cross-track or slant-range when converting between the ground plane and the slant plane, so the classification errors described above may result in an apparent smearing or “doubling” of water features in the cross-track direction in some HR products. That is, the water feature itself may be bright and properly detected and classified as water, but an additional set of pixels that is shaped like the water feature may also be incorrectly classified as dark water at a location that is shifted in slant range and/or cross track. There may hence be two “copies” of the water feature in the HR products—one classified as open (bright) water, and one classified as dark water. Alternatively, in a less extreme case, pixels that are incorrectly classified as dark water may appear along one side or the other of a water feature in slant range or in cross track, giving an erroneous buffer of dark water pixels at the edge of an otherwise correctly classified water feature. Erroneous dark water pixels may therefore affect the area results for rivers and lakes in the L2_HR_RiverSP and L2_HR_LakeSP products.

An example of erroneous dark-water classification due to reference DEM errors is shown in Fig. 10.6. Only *good* and *suspect* pixels and lake polygons are shown. All of the images in this example are shown in the ground plane (i.e., as functions of latitude and longitude). While the heights of the open-water pixels in panel (c) are reasonable, the pixels classified as dark water in panel (d) have a very wide range of heights (around ± 50 m). The height variations in panel (d) are clearly aphysical for a water surface and are caused by errors in the reference DEM, perhaps in conjunction with errors in the map of the prior water probability. Such errors are likely due to the rugged topography around the lake. Because of the height errors of the dark-water pixels, however, pixels at similar latitude-longitude coordinates in panel (a) project to different slant ranges. The lake-level processing of the slant-plane data then gives a shifted or doubled copy of the lake in cross track in panel (b). The western copy of the lake corresponds to the correctly detected and geolocated feature (class 4), while the eastern copy of the lake is the aberrant dark-water feature (class 5).

Note that if the errors in the reference heights of the dark-water pixels in Fig. 10.6 had been smaller, the dark-water copy of the lake would be less shifted; if the dark-water copy of the lake were shifted by less than the width of the lake, the artifact would appear as a buffer of dark-water pixels only along the eastern edge of the lake. River products may be affected in a similar manner.

The incorrectly classified dark-water pixels can cause errors in the surface area or extent of reach, node, and lake features. The fractional error can be especially large for small features. Future improvements to the ground algorithms may be able to mitigate the projection issues underlying these errors.




Figure 10.6: Example of dark water misclassification due to reference DEM errors: (a) geolocated classification values from the L2_HR_PIXC product for Kozjak Lake in North Macedonia; (b) the same classification values as from (a) but geolocated based on information from the L2_HR_PIXCVec product; (c) L2_HR_PIXC height and geolocation for open-water pixels (class 4) only; (d) L2_HR_PIXC height and geolocation for dark-water pixels (class 5) only. Polygons from the Obs file of the L2_HR_LakeSP product are shown in magenta in (b). Classes 3, 4, and 5 represent water near land, open (bright) water, and dark water. The data are from forward-processed (CRID PIC0) data of cycle 007, pass 402.




Figure 10.7: Example of bright land near the Willamette River. The image shows the relative (arbitrarily scaled) power from the $-y$ SLC channel in the slant plane after 7 looks in azimuth. The Willamette River runs from the upper left of the image to the lower middle. Fern Ridge Lake is the large, bright feature toward the upper right; the lake is approximately 7 km wide. Slant range increases to the right, and azimuth increases down the page, so north is approximately downward and slightly to the right. Many agricultural fields near the river are as bright as the water features, or nearly so. The city of Eugene, Oregon, near the upper left of the image, is also brighter than the typical, undeveloped land. Layover of the rougher land features toward the far right of the image is also evident. The image shows a portion of tile 232L from cycle 482, pass 013 (cropped to show detail). The data are from the reprocessed (CRID PGC0) L1B_HR_SLC product.

10.9 Bright Land

Land is typically much darker than water, but in some cases, land can be as bright as (or brighter than) water, even when the water is not dark. An example of bright land near the Willamette River is shown in Fig. 10.7. Many of the agricultural fields in the valley surrounding the river appear nearly as bright as the river. Moreover, the developed areas in the city of Eugene, Oregon are bright as well. From ground truth (not shown), the area is known to have been wet due to recent rains, but the area was not flooded. Rather, many areas had large puddles, which likely contributed to the brightness of the land. There may be other reasons for the higher-than-normal land reflectivity as well.

The reflectivity of the land varies from region to region at a given site and of course also varies from site to site. It can also vary with surface conditions and therefore with season. There are

few other Ka-band radars at incidence angles similar to those used by KaRIn, so land reflectivity in this regime is not well characterized.

Bright land can be a problem because it may be misclassified as water during HR processing. This can lead to erroneous estimates of the area and the water surface elevation in L2 HR products. River quantities may be especially affected because rivers are often long and narrow, which makes the process of accurately and robustly attributing water in the images to river features difficult. This makes them susceptible to errors if nearby bright land is incorrectly processed as part of the river.

Bright land can also worsen the impact of land-water layover because it would have an effect similar to water-water layover (see Sect. 10.2). Similarly, bright land can exacerbate phase unwrapping problems (see Sect. 10.10).

10.10 Phase Unwrapping Artifacts

Phase unwrapping, as described in Sects. 8.3.5 and 9.4.3, involves complicated processing algorithms to best solve what is fundamentally an underdetermined problem. The HR phase unwrapping algorithm usually works quite well for reasonably sized water features that are detected by SWOT, but it does sometimes select the incorrect 2π phase ambiguity. As the ambiguity resolution is performed for spatially unwrapped regions, the unwrapping errors result in misplaced “chunks” (spatially contiguous sets of pixels) of detected water. These incorrectly unwrapped regions typically have large horizontal shifts from their true locations in cross-track along with large height errors. Consequently, while phase unwrapping errors do not necessarily occur frequently, they are often quite glaring when they do occur.

The cross-track errors are typically near integer multiples of 750 m, which is the horizontal error associated with a phase error of one cycle. This value is approximately constant over the swath. The height errors are typically near integer multiples of the ambiguity height h_a , which is given by Eq. (8.113). The ambiguity height increases nearly linearly from nadir to far range. The cross-track error and the height error for a given unwrapping error are related through Eq. (8.114).

An example image of a phase unwrapping error is shown in Fig. 10.8. This example is from the Rhone River, where a portion of the river appears disconnected from the rest of the river. It is shifted by approximately 3 km in the cross-track direction toward far range. It also differs in height from the rest of the river by approximately 200 m (not shown in the image). This corresponds to four cycles of unwrapped phase error, or 8π rad. The feature occurs near the far side of the left swath side in an ascending pass; the incidence angle is 3.7° . Therefore, $1/\tan(\theta_{\text{inc}})$ is a factor of about 15, which is consistent with Eq. (8.114) given the errors in cross track and in height.

As is evident from Fig. 10.8, erroneously unwrapped regions can appear far out of family from their surroundings. It is important for users of HR data to be able to recognize such errors in order to distinguish them from observations of other nearby water features and to thus avoid misinterpreting the data.

Additionally, where phase unwrapping errors occur, pixels that belong to a given river or lake feature may be incorrectly excluded from the feature, so unwrapping errors may result in “holes” or “missing” portions of rivers and/or lakes.




Figure 10.8: Example of a phase unwrapping error in HR data from the Rhone River. The image shows HR pixel classifications in the ground plane based on the L2_HR_PIXC latitude and longitude coordinates of the pixel geolocations. Light blue and purple indicate pixels that are classified as interior water and as water near land, respectively (other classes are not shown). The centerline of the river from the prior river database is shown in red. The part of the river that should be near the tail of the yellow arrow, continuous with the rest of the river, appears near the head of the arrow instead. It is shifted by approximately 3 km toward far range in the cross-track direction because it was incorrectly unwrapped and placed on the wrong 2π phase ambiguity. The image is from forward-processed data (CRID PIA1) of cycle 523, pass 003, tile 235L, which was acquired on 2023-05-16.

Phase unwrapping errors are generally more common (in a relative sense) toward the near-range portion of the swath, where the ambiguity height h_a is smaller and where it is hence more difficult to determine the correct ambiguity in the presence of errors in the reference DEM. Phase unwrapping errors are also generally more common in areas of significant topographic variation (e.g., mountainous areas), for water bodies affected by water-water layover (e.g., near dams), and for small or narrow water features.

Phase unwrapping errors typically do not occur for LR data over the ocean because the height of the ocean surface is usually known to within $\pm h_a/2$.

10.11 Target Motion

As briefly discussed in Sect. 8.2.2, it is well known that moving targets appear shifted along contours of constant range in SAR images. This effect has implications for high-precision interferometric height measurements, however. This is because contours of constant interferometric phase are not necessarily aligned with contours of constant range on the surface (see Fig. 8.13). Therefore, when height reconstruction is performed for the target based on its apparent (shifted) location while its observed interferometric phase corresponds to its true location, the target height is estimated incorrectly. This effect is described in greater detail in Sect. 2.4.6 of [58].

The immediate consequence of target motion is a height error that scales with the line-of-sight (radial) component of the target velocity vector and becomes worse for squinted viewing geometries. Fortunately, the KaRIn antenna pointing is maintained very near broadside thanks to the excellent attitude control of the SWOT spacecraft and the stability of the KaRIn instrument, so the sensitivity to target motion is minimized, and the resulting height errors are limited. The SWOT error budget [59] includes an allocation for such errors for surfaces within the regime of the SWOT requirements.

In the case that water moves as a slab that translates uniformly, the motion-induced height errors are well behaved and have only low-order spatial variability. However, ocean waves present a more complicated pattern of spatially varying target motion that warrants special attention in the context of height estimation (see Sect. 2.5.2.3 of [58] for additional details). With the steep look angles of the KaRIn measurement, the target shifts are mainly due to the vertical component of the surface velocity. For the orbital motion of a surface gravity wave propagating in the along-track direction, the vertical component of the velocity is greatest in magnitude between the crests and troughs of the wave, and the sign of the vertical velocity is opposite for patches of the surface on either side of a crest (see Fig. 10.9). Those patches of water are therefore shifted in opposite directions. Consequently, the wave is distorted in the SAR image so that it appears more or less “peaky” than it is in reality. This phenomenon is known as wave bunching.

If the azimuth resolution of the SAR image is finer than the wavelength of the ocean waves, wave bunching may cause errors in the interferometric height estimates. Consider a case in which the observed, bunched (red) surface of Fig. 10.9 is uniformly sampled horizontally, and the heights are averaged. The result is not equal to the mean height of the true (blue) surface. Moreover, the observed surface contains wavenumbers that are not present in the true wave surface, so spatial averaging or low-pass filtering may not remove the resulting errors. However, if the averaging is weighted by the power of the echo, and if the power of the echo is proportional




Figure 10.9: Illustration of wave bunching in the context of height estimation for a hypothetical, sinusoidal ocean wave. When the true wave surface (blue) is propagating in the along-track direction, its spatially varying vertical velocity (green) causes it to appear distorted (red) in a SAR image. Note that Fig. 2.15 of [58] is based on the same source illustration.

to the area of the true surface that contributes to a resolution cell (neglecting EM bias), the error can be minimized.

The height estimates in the L2_LR_SSH product are computed from complex interferogram quantities that have been averaged to approximately 500 m resolution, which is longer than the wavelength of most (but not necessarily all) ocean waves. This complex averaging approximates the power-weighted averaging described above because the interferogram magnitude is closely related to the image power. While such averaging makes the height estimates in the L2_LR_SSH product susceptible to EM bias (see Sect. 10.4), it also makes the results less sensitive to wave bunching. On the other hand, the height estimates in KaRIn L2_HR_PIXC and L2_HR_Raster products are given at resolutions that may be much finer than the wavelengths of typical ocean waves. While this fine resolution may allow them to avoid EM bias, it also makes them susceptible to errors related to wave bunching. For this as well as other reasons, care must be taken when attempting to use KaRIn HR data over the ocean.

Note that the case illustrated in Fig. 10.9 is a simple, hypothetical case with only a single, sinusoidal wave component. If the wave spectrum contains many wavenumber components, the nonlinear distortion may make the signature of the error much more complicated than in the example shown.

Wave-related target motion generally increases with the wave amplitude and hence the SWH. Therefore, height errors related to wave motion are closely related to SSB, and the distinction between them is not crisply defined. Both SSB and the coupling of wave motion to height error are areas of continuing investigation for KaRIn data.

Finally, another effect related to target motion is blurring of the SAR image in azimuth. Such

blurring can occur if there are spatial variations in the target motion such that the target decorrelates over the azimuth integration time (approximately 2 ms for LR and 100 ms for HR). Equivalently, if different portions of a nominal resolution cell on the ground have different radial velocities, they experience different azimuth shifts, which collectively result in a blurring of the target in azimuth.

Motion blurring can be more significant for HR data than for LR data given the finer azimuth resolution of the former. Motion blurring is often called “coherence time smearing” in this context when the coherence time of moving water is shorter than the integration time required to achieve some nominal resolution δ_x . Such smearing would result in the water being blurred in azimuth over an extent that is wider than δ_x , while the stationary land around the water remains well focused. This may cause water features to appear wider in the along-track direction than they actually are.

Motion blurring in LR data may be less noticeable visually, but it may cause errors in the phase bias correction computed during ground processing (see Sect. 9.4.1). When the motion blurring is due to the wave motion, the effect may be worse for larger SWH, and it is again related to SSB.

10.12 LR Errors Over Land

The KaRIn LR measurement is tailored for oceanography, but the L2_LR_SSH product is available over land as well. As noted in Sect. 6.9, however, LR measurements over inland water may be unreliable for all but the largest of inland water features.

Figure 10.10 shows an example of a gross error in the water elevation reported in the L2_LR_SSH product over Lake Geneva (Léman). The SSHA from the L2_LR_SSH product is shown here because, over land, the reported SSHA is nearly equivalent to the WSE as defined in the HR products (see Sects. 3.1.24 and 3.1.25 and [6]). The crossover calibration reported in the L2_LR_SSH product has been applied in order to remove cross-track trends (see Sect. 6.9). The figure shows only data between 10–60 km from nadir on either side of the swath, so the lake is partially cut off in cross track in both passes shown. Measurements flagged as *bad* are not shown.

In the example of Fig. 10.10, the measurements over the lake are very different between the two passes shown. In the descending pass (264), the water elevation over the lake is roughly consistent with expectations (approximately 372 m relative to the geoid). On the other hand, in the ascending pass (335), the water elevation over the lake is 800–900 m higher than expected, with an along-track variation of more than 100 m over the lake surface. The results from pass 335 are clearly artificial.

The errors for pass 335 of this example arise because the OBP reference surface used for flattening (see Sect. 8.3.4) during KaRIn on-board processing (see Sect. 9.2.1) does not match the actual elevation of the lake. Recall from Sect. 9.2.3 that quantities such as the OBP reference surface are stored in tables that are parameterized in terms of “table time,” which is conceptually equivalent to location along the reference nadir track. The OBP reference surface is constant for each swath side (it does not vary with distance from nadir on a given side), and it is constrained with respect to how quickly it can vary in along track. The table is designed with the intent that

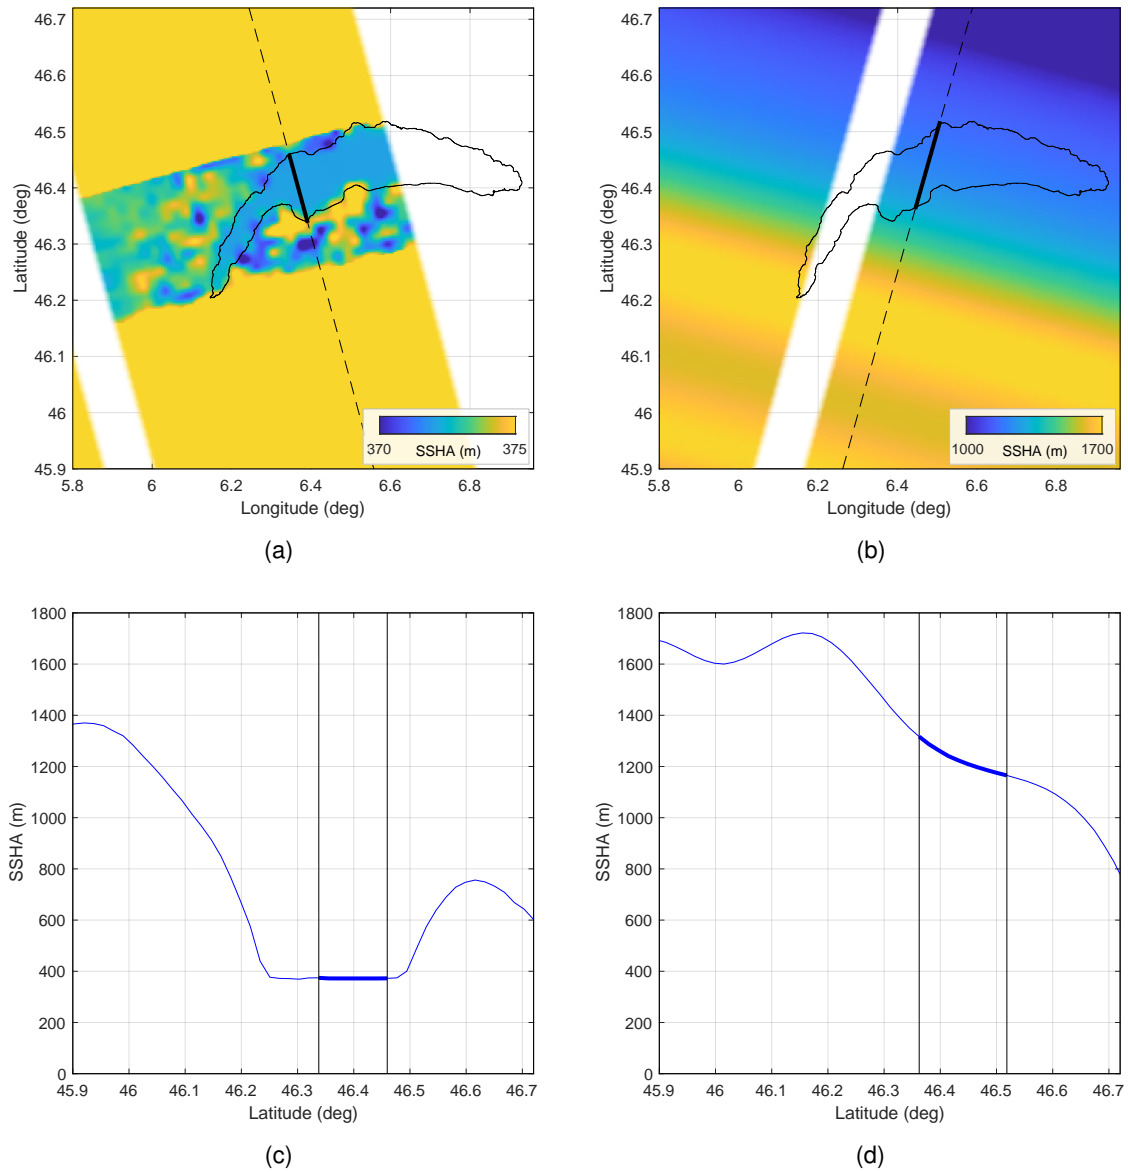


Figure 10.10: Example of a gross LR error over inland water: (a) SSHA from the L2_LR_SSH product for cycle 007, pass 264 over Lake Geneva; (b) SSHA over Lake Geneva for cycle 007, pass 335; (c) cut of the SSHA along the dashed line in (a); (d) cut of the SSHA along the dashed line in (b). Over land, the SSHA from the L2_LR_SSH product is approximately equivalent to the WSE from HR products. Note the vastly different color scales between (a) and (b); the color range is highly saturated for (a). The measurement in (a) is roughly consistent with the true WSE (around 372 m), while the measurement in (b) is erroneous by hundreds of meters and has an aphysical along-track trend of over 100 m across the lake. The polygon for Lake Geneva from the PLD is indicated in (a) and (b). The thick portions of the curves in (c) and (d) correspond to the lake as indicated by the thick, solid portions of the lines indicating the cuts in (a) and (b). These are 14 km and 18 km long, respectively. The data are from forward-processed (CRID PIC0) data.

the OBP reference surface agrees with the expected elevations of lakes whose along-track extent is greater than a given threshold. Between those large lakes, the OBP reference surface may vary greatly to follow the approximate, general topography of the land surface. For pass 335, the along-track extent of the lake (as defined by the offline software that generates the tables) is below this size threshold, so the OBP reference surface is not constrained to agree with the expected lake elevation. The OBP reference surface therefore varies over the lake surface in order to fulfill other optimization objectives in the table design. With the OBP reference surface differing from the true lake elevation by hundreds of meters, the resulting elevations in the L2_LR_SSH product show commensurate errors (recall from Sect. 9.4.1 that phase unwrapping does not occur in the LR ground processing).

In pass 264, the along-track extent of the lake is above the size threshold, so the OBP reference surface is close enough to the true lake elevation that the resulting LR measurement is reasonable. Large variations in the reported SSHA both before and after the lake in along track are evident for both passes 264 and 335 as the OBP reference surface varies. The variations differ between the two passes because the ground tracks differ.

While the reported elevations over the lake are reasonable over much of the lake surface for pass 264, close inspection nevertheless shows aphysical, meter-scale variations in elevation within 1–2 km of the shore. These variations are not unexpected in LR data over inland water because of the influence of the surrounding land on the LR measurement. Similar artifacts of smaller magnitude may persist over the lake surface up to several kilometers from the shore. Some of the artifacts may be related to the ground processing, and the nature of the artifacts may change in future versions of the L2_LR_SSH product.

Chapter 11

Geophysical Effects

This chapter provides a high-level description of the various geophysical effects that affect the SWOT NAlt and KaRIn measurements. Sections 11.1 and 11.2 describe geophysical effects that impact the radar signals and must be taken into account to recover the physical surface heights observed by SWOT. Section 11.3 describes geophysical phenomena that contribute to the observed surface heights and approaches that can be used to account for them.

The SWOT products provide corrections for these effects. Note of course that the provided corrections will have residual errors that may have both temporal and spatial structure. The models that are used to generate these corrections are subject to change in different versions of the products. This user handbook is not intended to identify the specific models that are used in the SWOT products. Users should consult the respective PDDs and product metadata to identify the models that are used in the corresponding product versions.

11.1 Media Delays

Radar measurements of range are necessarily affected by the speed of the radar signal through the medium in which the signal propagates. In particular, the mass in the troposphere and the ionosphere adds delays to the round-trip times of radar signals from the NAlt and KaRIn instruments.

The range corrections for the media delays are computed and applied along the path of the respective radar signals. For the NAlt this is the same as the vertical correction. In contrast, KaRIn measurements are off nadir, and the corrections are computed and applied along the slant path. The KaRIn corrections also account for the differential delay between the two (+y and -y) antennas. However, the SWOT products always report the equivalent vertical troposphere and ionosphere range corrections. The vertical corrections for KaRIn are computed by applying obliquity factors to the applied slant-path corrections. For example, the relationship between vertical and slant-path delay corrections, τ_z and τ_p respectively, is

$$\tau_p = \frac{\tau_z}{\cos \theta}, \quad (11.1)$$

where θ is the angle with respect to the local vertical of the radar signal path. Note that θ varies between approximately 0.7° and 4° between 10 km and 60 km cross track.

Reporting the vertical range corrections allows users to easily swap out alternative range corrections as they choose. The above relationship suggests that swapping out vertical range corrections will result in an error in the slant-path correction of $\delta\tau_z\theta^2/2$, where $\delta\tau_z$ is the difference between the original and alternative vertical correction. At the outer edge of the KaRIn swath, the error in the slant-path correction is then 0.5% of the difference, or 0.5 mm for a difference of 100 mm in the vertical correction.

As described in Sect. 3.3.8 the sign convention of the range corrections is such that the corrections should be added to the uncorrected radar range measurements. Specifically, the range corrections due to media delays typically have negative values such that they shorten the uncorrected radar range when applied. Positive values effectively reflect the inherent noise in the corrections. The sign of these “range corrections” is opposite from what is often referred to as the “range delay,” where the range delay is typically treated as a positive quantity.

11.1.1 Dry Troposphere

The gases in the troposphere contribute to the index of refraction, where the refractive index depends on pressure and temperature. The vertically integrated range delay from these gases is a function of the surface pressure when hydrostatic equilibrium and the ideal gas law are assumed. The dry troposphere range correction, τ_{dry} , can then be computed from (e.g., [60])

$$\tau_{\text{dry}} = K \times P \times [1 + 0.0026 \times \cos(2\phi)], \quad (11.2)$$

where P is the surface atmospheric pressure, ϕ is the latitude, and K is -2.2771 mm/mbar.

SWOT products use, and provide, the dry troposphere correction as computed using surface pressure fields from the European Centre for Medium-range Weather Forecasts (ECMWF) operational analysis. The pressure fields, which are produced every 6 hours, are interpolated in time and space, and a model is used to account for the S1 and S2 atmospheric tides (e.g., [61]).

The NAlt products provide two dry troposphere corrections. The first is computed using ECMWF sea surface pressure fields, and the second is computed using surface pressure at the measurement altitude. The dry troposphere correction in the KaRIn LR products is computed using the ECMWF sea surface pressure fields, while in the HR products it is computed from the vertical integral upward from the target altitude of ECMWF three-dimensional fields for pressure, humidity, temperature, and cloud liquid water content.

The dry troposphere correction typically varies from -2.1 to -2.4 m. Errors in the ECMWF surface pressure fields are dependent on location and can be up to a few millibars, which translates to approximately 1 cm in the dry troposphere correction.

11.1.2 Wet Troposphere

Water vapor in the troposphere also contributes to the index of refraction and therefore to a delay in the radar signals. The NAlt and KaRIn LR products use, and provide, a wet troposphere correction from the onboard microwave radiometer (see Sect. 4.4.3). The radiometer brightness-temperature measurements effectively provide a direct measure of the wet troposphere content over the oceans and may also be valid over large inland water bodies. The benefit of the AMR measurements is the capability to capture short-wavelength variations in the wet troposphere

correction. However, the AMR wet troposphere corrections are subject to land, rain, and ice contamination.

All of the SWOT products also use, and provide, a wet troposphere correction computed from pressure, temperature, and humidity fields from the ECMWF operational analysis. The fields, which are produced every 6 hours, are interpolated in time and space. The NAlt products also provide two model wet troposphere corrections, the first at the altitude of the ECMWF DEM (effectively mean sea level over the oceans), and the second computed from the vertical integral from the measurement altitude of ECMWF three-dimensional fields for pressure, humidity, temperature, and cloud liquid water content. The model wet troposphere correction on the KaRIn LR products is computed at the altitude of the ECMWF DEM, while on the HR products it is also computed from the explicit vertical integral upward from the target altitude.

The wet troposphere correction typically varies from 0 to -0.45 m. The global standard deviation of the over-ocean differences between the radiometer and the ECMWF model wet troposphere correction is typically less than 1 cm. The AMR measurements are expected to have better accuracy and spatial resolution than the ECMWF model over oceans given the inherent limitation in the spatial and temporal resolution of the ECMWF fields. Meanwhile, the correction from the ECMWF model is best used over land.

The wet troposphere also attenuates the radar signal, thereby affecting the backscatter estimate. The SWOT products use a similar approach for the atmospheric attenuation as is used for the wet troposphere correction. The NAlt and KaRIn LR products use, and provide, an atmospheric attenuation correction from the radiometer. All of the SWOT products also use, and provide, an atmospheric attenuation correction computed from pressure, temperature, and humidity fields from the ECMWF operational analysis.

11.1.3 Ionosphere

Radar signals are delayed by the free electrons in the Earth's ionosphere. This delay is inversely proportional to the square of the radar signal frequency. The ionosphere correction, τ_{iono} , in meters, is expressed as

$$\tau_{\text{iono}} = -40.3 \times \frac{\text{TEC}}{f^2}, \quad (11.3)$$

where TEC is the total electron content in units of electrons/m², and f is the radar signal frequency in units of Hz. The TEC varies with local time, season, and solar cycle.

The dual-frequency NAlt enables a direct measurement of TEC, and subsequently the ionosphere correction at the Ku- and C-band frequencies. Along-track averaging is often used to reduce the noise in this measurement (e.g., [62]). The NAlt products provide unsmoothed and smoothed ionosphere corrections that are derived from the Ku- and C-band range measurements after correcting them for the frequency-dependent sea state bias. The NAlt products also provide an ionosphere correction that is computed from Global Ionosphere Maps (GIMs; see, for example, [63]) as a backup correction. The KaRIn LR and HR products use, and provide, an ionosphere correction that is computed from the GIMs.

The ionosphere TEC can vary from nearly zero to over 100 TECU, where 1 TECU is equivalent to 10^{16} electrons/m² (e.g., [62]). As such, the ionosphere correction typically varies from zero to -0.22 , -1.4 , and -0.03 m at the Ku, C, and Ka bands (13.6, 5.3, and 35.75 GHz), respectively. The

accuracy of the NAlt dual-frequency unsmoothed ionosphere correction is approximately 0.5 cm at significant wave heights of 2 m. The global standard deviation of the over-ocean differences between the NAlt dual-frequency and GIM model TEC is on the order of 5 and 10 TECU close to solar minimum and maximum, respectively. At solar maximum, this corresponds to 0.02, 0.14, and 0.003 m for the Ku-, C-, and Ka-band ionosphere corrections, respectively. However, these differences have significant systematic dependencies on local time, season, and solar cycle.

11.2 Sea State Bias Correction

As described in Sect. 10.4, the sea state bias (SSB) is an effect on radar-measured height due in large part to the difference in signal reflectivity between the peaks and troughs of ocean waves. Generally, the effect tends to bias the observed height low because ocean troughs reflect more of the radar signal than peaks. The SSB is dependent on the frequency of the radar signal.

Unlike the media corrections, the sea state bias correction is applied as a vertical correction to both the NAlt and KaRIn range measurements as it is a scattering effect rather than a propagation effect. The sign convention is the same as is used for the media corrections with the SSB correction typically having negative values.

The NAlt and KaRIn LR products use, and provide, SSB corrections based upon empirically derived models at each of the Ku-, C-, and Ka-band frequencies. The SSB correction is zero over land for both the NAlt and LR products. The KaRIn HR products do not apply a SSB correction, even if the HR data were collected over the ocean.

The SSB correction typically varies from 0 to -0.4 m at each of the Ku-, C-, and Ka-band frequencies. An often-used approximation of the accuracy of the empirically derived SSB models is 1% of the significant wave height. Note that empirical models for the SSB correction have a known difficulty in resolving the overall bias of the correction that can be at the level of a few centimeters. Errors in the relative bias between the Ku- and C-band SSB can subsequently also result in a bias error in the NAlt dual-frequency ionosphere correction.

11.3 Geophysical Models

There are various geophysical phenomena that contribute to variations in the surface heights observed by SWOT. This section describes the most significant of these effects along with a high-level description of the approaches used to model them. Some of these effects have temporal variations with shorter periods than the revisit time of the SWOT measurements. As such, they appear at aliased periods.

11.3.1 Geoid

The geoid is an equipotential surface of the Earth's gravity field that is closely associated with the mean sea surface. The height of the geoid above the reference ellipsoid (see Sect. 3.3.3) is sometimes referred to as the geoid undulation or height. The geoid height over the entire Earth ranges from approximately -107 m to $+86$ m, with a root mean square of 29.3 m, as shown in Fig. 11.1. The long wavelength components of the geoid height are typically determined from




Figure 11.1: Geoid heights in meters from the EGM2008 model [64]. This particular geoid model is shown for illustrative purposes only; the geoid model used when generating SWOT products may change from one processing version to another.

spherical-harmonic models of the Earth's gravity field. These spherical-harmonic models are typically determined from satellite gravity missions such as GRACE. However, very short wavelength changes in the geoid height occur, for example, over mountains, seamounts, trenches, and ridges. High-resolution surface gravity data are used to capture very short wavelength features in the geoid height but are generally only available for certain regions of the Earth. The best available geoid models typically combine satellite and surface gravity data (e.g., [64]).

The geoid serves as the reference surface for the water surface elevations reported in the KaRIn HR products. The geoid may also be relevant to some ocean applications. As such, the geoid height from a global model is provided in the NAlt, KaRIn LR, and KaRIn HR products.

By convention, the geoid height provided in SWOT products is in the mean-tide system in that it also includes the zero-frequency (i.e., time-independent) permanent tide height. Meanwhile, the source of the geoid, for example the EGM2008 geoid model [64], is often provided in a tide-free system. In SWOT products, the difference Δh between the mean-tide and tide-free geoid heights, h_{mt} and h_{tf} respectively, is modeled as a function of latitude θ_{lat} by

$$\Delta h = h_{mt} - h_{tf} \quad (11.4)$$

$$= (1 + k_2) H_{perm} \sqrt{\frac{5}{4\pi}} \left(\frac{3}{2} \sin^2 \theta_{lat} - \frac{1}{2} \right). \quad (11.5)$$

SWOT products use the permanent tide potential amplitude $H_{perm} = -0.31460$ m from [65] [66] with the Love number $k_2 = 0.3$. A plot of Δh as a function of latitude is shown in Fig. 11.2. The sign of Δh is such that the mean-tide geoid height (relative to the ellipsoid) is larger at the equator




Figure 11.2: Difference Δh between mean-tide and tide-free geoid heights.

than the tide-free geoid height. Therefore, at the equator, the height of a surface relative to the mean-tide geoid would be smaller than the height of the same surface relative to the tide-free geoid.

11.3.2 Mean Sea Surface

The Mean Sea Surface (MSS) represents the average height of the ocean surface above the reference ellipsoid (see Sect. 3.3.3). Typically, the MSS is determined by averaging global sea surface height measurements over sufficiently long time periods to remove seasonal and spurious sea surface height signals. MSS models are typically provided on two-dimensional grids with a resolution that is consistent with the altimeter and other data used to generate them. The MSS can be useful for data editing purposes, for the calculation of along-track and cross-track geoid gradients, and for the calculation of gridded gravity anomalies. It can also serve as a reference surface to which sea surface height data from different altimeter missions can be combined. The MSS closely follows the geoid and therefore has similar amplitudes, as shown in Fig. 11.3.

The MSS is only relevant to the oceans and is therefore only provided in the NAlt and KaRIn LR products. The MSS serves as the reference surface for the sea surface height anomalies reported in these products. By convention, the MSS provided in SWOT products is in the mean tide system, in that it also includes the zero-frequency (time constant) permanent tide height.

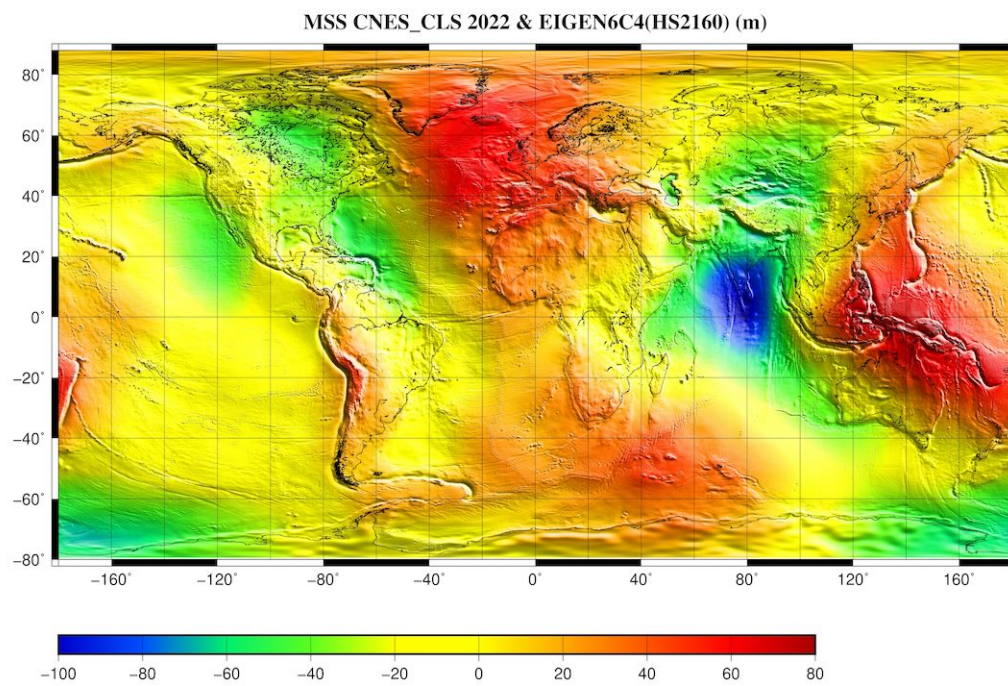


Figure 11.3: Mean sea surface from the CNES/CLS 2022 model (Credit CNES/CLS) [67]. This particular MSS model is shown for illustrative purposes only; the MSS model used when generating SWOT products may change from one processing version to another.

11.3.3 Mean Dynamic Topography

The Mean Dynamic Topography (MDT) represents the Mean Sea Surface referenced to a geoid and corrected for various geophysical effects such as tides. MDT models are typically provided on two-dimensional grids with a resolution that is consistent with the altimeter and other data used to generate them. The MDT provides the absolute reference surface for ocean circulation.

The MDT is only relevant to the oceans and is therefore only provided in the NAlt and KaRIn LR products. The MDT height ranges from -1.5 to $+1.8$ m, with a global root mean square of 0.75 m.

11.3.4 Tides

Tides are perhaps the most significant contributor to spatiotemporal variations in the SWOT observations of surface heights over the ocean, and they are a significant contributor over land as well. While tides are of scientific interest in themselves, their effects must be removed when studying more subtle variations in surface height. High-accuracy models are used to provide predicted tide heights in SWOT NAlt and KaRIn products so users can remove those effects as they choose. The sections below generally focus on the tidal effects on SWOT surface-height observations.

11.3.4.1 Luni-Solar Tides

The predominant tidal effects result either directly or indirectly from luni-solar gravitational forces acting on the Earth. These include a zero-frequency (i.e., time-independent) effect called the permanent tide. The time-dependent luni-solar tidal effects occur in three tidal bands, the semi-diurnal, diurnal, and long-period tidal bands. Each tidal band has rich spectral content, with periods ranging from 11 to 14 hours, 21 to 30 hours, and a few days to 18.6 years, respectively.

By convention, the models for tide heights provided in SWOT products have zero mean and exclude the contribution from the permanent tide. For this reason, the provided model of the geoid height is in the mean-tide system (see Sect. 11.3.1) and the models of the MSS are also considered to be in the mean-tide system. Users are cautioned to exercise care with the treatment of the permanent tide when comparing SWOT data to in-situ surveys. Refer to the end of this section for more details.

Luni-solar tidal effects are typically segregated into the following contributions:

- **Solid Earth Tide**, also referred to as the **Body Tide**: The solid Earth tide represents the direct response of the solid Earth crust to the tide-generating forces. The solid Earth tide can be accurately modeled by assuming an elastic (in phase) response of the solid Earth where the tide height is proportional to the tide-generating potential. The proportionality constant for solid Earth tide heights is referred to as the Love number. The solid Earth tide height model that is being used in NAlt and KaRIn products has been used in satellite altimetry products since the Topex/Poseidon mission. This model assumes that the Love number is frequency independent with a significant exception near the $K1$ (once per sidereal day) tidal frequency due to a resonance in the Earth's liquid core (e.g., [68]). With

this approach, the spectral and spatial content of the solid Earth tide height is the same as the tide-generating potential. Ocean tide models derived from historical altimetry data are therefore typically consistent with this solid Earth tide height model, in that the ocean tide models will absorb any errors in this solid Earth tide model.

Note that the current conventions of the International Earth Rotation Service [38] adopt a more rigorous solid Earth tide model that accounts for second order effects from additional frequency-dependent Love numbers as well as an anelastic (out of phase) response of the solid Earth to the tide-generating potential. The RMS difference between the altimetry and IERS solid Earth tide height models is less than 2 mm (e.g., [69]).

The solid Earth tide affects the solid Earth crust and therefore has similar amplitudes of up to 0.3 m over both ocean and non-ocean surfaces. As such, the NAlt, KaRIn LR, and KaRIn HR products all provide a model for the solid Earth tide height.

- **Ocean Tide:** The ocean tide represents the direct barotropic response of the oceans to the tide-generating forces. The spectral content of the ocean tides is the same as the tide-generating potential. However, the spatial structure of the ocean tide is complicated by the inherent dynamics of the oceans, especially in the diurnal and semi-diurnal tidal bands. In the long-period tidal band, the spatial structure gradually approaches that of the tide-generating potential with longer periods, especially greater than 6 months (e.g., [70]). As such, the long-period tides are often assumed to be in equilibrium with, or proportional to, the tide-generating potential. The most accurate ocean tide models are typically determined empirically from satellite altimetry data or from models that combine dynamic equations of motion with empirical altimetry and tide gauge data.

The ocean tide is only applicable to ocean surfaces and can have amplitudes of up to 5 m or more. The NAlt and KaRIn LR products provide two models for the diurnal and semidiurnal ocean tide height (e.g., [69], [71]). These SWOT products also provide a model for the long-period ocean tide height that assumes an equilibrium response to the tide-generating potential as well as a hydrodynamic model that accounts for small deviations from an equilibrium response.

- **Load Tide:** The load tide represents the response of the solid Earth crust to the load of the ocean tide. It is an indirect response to the tide-generating potential and therefore has the same spectral composition. However, the spatial structure of the load tide height is more closely aligned with the loading mass of the ocean tide. While the load tide has the largest amplitudes over the oceans, namely beneath the loading mass, it extends over non-ocean surfaces. Models of the load tide typically accompany models of the ocean tide for self-consistency, especially since ocean altimetry measurements observe the sum of the ocean and load tide height. The load tide height can have amplitudes of up to 0.1 m over ocean and non-ocean surfaces. Therefore, the NAlt, KaRIn LR, and KaRIn HR products all provide models for the load tide height that are consistent with the ocean tide models adopted for the SWOT ocean products.
- **Internal Tide:** Internal tides represent the indirect baroclinic response of the oceans to the barotropic ocean tide. They are generated by the barotropic tidal flow over the underlying

bathymetry (e.g., [72]). The internal tide heights can have amplitudes of a few centimeters at the ocean surface. The internal tide itself is often separated into coherent and incoherent components. The coherent component is stable and can be predicted using recent models. Meanwhile, the incoherent component is more difficult to predict due to the interaction of the barotropic ocean tide with the inherent ocean circulation. The internal tide is only applicable to ocean surfaces. The NAlt and KaRIn LR products provide a model for the coherent internal ocean tide height.

- **Geocentric Ocean Tide:** The term “geocentric ocean tide” has been used in satellite altimetry products to refer to the sum of the ocean and load tide heights:

$$\text{geocentric ocean tide} = \text{ocean tide} + \text{load tide height}. \quad (11.6)$$

SWOT products explicitly provide the geocentric ocean tide and load tide height separately. The ocean tide itself can be recovered from the difference between the provided geocentric ocean tide and load tide height.

NAlt and KaRIn LR products use models to provide the solid Earth tide height, the geocentric ocean tide height, the load tide height, and the internal tide height. KaRIn HR products use the same models to provide the solid Earth tide height and the load tide height.

As mentioned above, the solid Earth tide model in SWOT products excludes the permanent tide. In contrast, most ground positioning software packages that are used to compute coordinates from in-situ surveys likely adopt the solid Earth tide model from the IERS Conventions [38], which includes what they refer to as a “permanent deformation”. As such, care should be taken to properly account for this permanent deformation when comparing coordinates computed from those positioning software packages to heights provided in the SWOT products. If the positioning software package includes the permanent deformation in the background solid Earth tide model, then the computed coordinates exclude that deformation and are referred to as “conventional tide free” values. In this case, the permanent deformation should be added to the conventional tide free coordinates before comparing them to SWOT measurements. Specifically, the following representation of the permanent deformation Δh_{pd} should be added to the conventional tide free coordinates, where

$$\Delta h_{pd} = h_2 H_{\text{perm}} \sqrt{\frac{5}{4\pi}} \left(\frac{3}{2} \sin^2 \theta_{\text{lat}} - \frac{1}{2} \right). \quad (11.7)$$

The IERS conventions recommend using the permanent tide-potential amplitude $H_{\text{perm}} = -0.31460$ m from [65] [66] and the Love number $h_2 = 0.6078$. Note that Eq. (11.7) ignores a small latitudinal variation of the Love number. The effect of the permanent deformation is shown in Fig. 11.4. Note that there is also a transverse north-south component of the permanent deformation, with an amplitude of less than 2.5 cm, that is less relevant for comparisons to SWOT measurements as the impact is a small difference in latitude.

11.3.4.2 Pole Tide

Variations in the geocentric location of the Earth’s instantaneous rotation axis, or polar motion, introduces a differential centrifugal force that causes displacements of the solid Earth and




Figure 11.4: Permanent deformation Δh_{pd} to add to conventional tide-free coordinates.

oceans (e.g., [73]). These displacements are referred to as the pole tide as they can be treated in much the same way as the luni-solar tides. Observed polar motion, and therefore the centrifugal potential, is dominated by periodic variations with periods of approximately 433 days (the period of the Chandler wobble normal mode of the Earth) and 12 months (seasonal). The centrifugal potential has a similar spatial form as the diurnal luni-solar tidal potential. Observed polar motion also has a secular (long term non-periodic) motion. This drift must first be removed from observed polar motion before applying tidal theory to compute the pole tide (e.g., [74]).

The pole tide can be similarly segregated into the solid Earth (or Body) pole tide, the ocean pole tide, and the load pole tide. The solid Earth pole tide height can also be accurately modeled as being proportional (using Love numbers) to the differential centrifugal force. It has amplitudes of up to 0.01 m. The long periods of the two dominant polar motion frequencies allow the ocean pole tide to be modeled as being in equilibrium with, or proportional to, the differential centrifugal potential. The ocean pole tide also has amplitudes of up to 0.01 m (eg, [75]). The ocean pole tide is only relevant to ocean surfaces. The load pole tide can be derived from the assumed ocean pole tide. The load pole tide height has amplitudes that are an order of magnitude smaller than the ocean pole tide. The total pole tide height over the oceans can then have amplitudes of up to 0.02 m.

NAlt and KaRIn LR products use models to provide the sum of the solid Earth pole tide height, ocean pole tide height, and load pole tide height in a single variable, with the ocean pole tide height set to zero over land. KaRIn HR products use the same models to provide the sum of the solid Earth pole tide height and load pole tide height (but not the ocean pole tide height) in a single variable. Note, therefore, that KaRIn HR products omit the ocean pole tide height even

for HR data collected over the ocean. These pole tide heights are computed for SWOT products using predicted or observed polar motion from the International Earth Rotation Service. Tables 5.2 and 5.3 summarize which products use predicted versus observed polar motion.

11.3.5 Ocean Response to Atmospheric Forcing

The height of the ocean surface responds to variations in the atmosphere above it. These effects are captured by the related inverse barometer correction and the dynamic atmosphere correction. They are typically used to compute SSHA from SSH per Eq. (3.4).

11.3.5.1 Inverse Barometer Correction

At long periods, the oceans tend to respond hydrostatically to atmospheric pressure and can be modeled as an inverse barometer. An inverse barometer correction, IB, is provided in the NAlt and KaRIn LR products. This correction is computed as

$$\text{IB}(t) = b \times (P(\phi, \lambda, t) - \bar{P}(t)), \quad (11.8)$$

where P is the surface atmospheric pressure at latitude ϕ and longitude λ , $\bar{P}(t)$ is the time-varying average of the global surface atmospheric pressure over the oceans, and the constant b is -9.948 mm/mbar [76]. SWOT products use surface atmospheric pressure from the ECMWF operational analysis to compute the IB correction. These pressure fields are typically available at 6-hour intervals and are interpolated in both space and time. A model for the S1 and S2 atmospheric tides is subtracted from the surface atmospheric pressure before computing the inverse barometer correction (e.g., [61]). As noted in Table 5.2, the NAlt OGDR products compute the inverse barometer correction using forecast pressure fields, while all other NAlt and KaRIn LR products use analyzed pressure fields. The inverse barometer height typically ranges from -0.4 to $+0.9$ m, consistent with an average global over-ocean surface pressure of 1010.9 mbar.

11.3.5.2 Dynamic Atmospheric Correction

At short periods, the ocean response to wind and pressure is mostly dynamic. Barotropic models have been shown to have the capability to account for much of this signal.

NAlt and KaRIn LR products provide a dynamic atmosphere correction (DAC) to account for the dynamic ocean response to wind and pressure forcing. This correction is composed of the low-frequency inverse barometer correction and the high-frequency response to atmospheric pressure and wind forcing as computed from a barotropic ocean model (e.g., [77]). The ocean model is forced by surface winds and pressures from ECMWF after subtracting a model for the S1 and S2 atmospheric tides (e.g., [61]). The initial release of SWOT products used barotropic model output that was filtered to retain periods shorter than 20 days before forming the DAC. This 20-day cutoff is based upon the Nyquist period of the approximately 10-day exact orbit repeat period of the reference altimeter missions (Topex/Poseidon, Jason-series, and Sentinel-6). The cutoff period is subject to change for the SWOT products.

The DAC can be different on the various SWOT ocean products due to the latency of the ECMWF fields, with some products necessarily using forecast, instead of analyzed, ECMWF

fields to force the barotropic ocean model. Tables [5.2](#) and [5.3](#) summarize these differences. The DAC height typically ranges from -0.4 to $+0.9$ m.

Bibliography

- [1] L.-L. Fu, D. Alsdorf, R. Morrow, E. Rodriguez, N. Mognard, “SWOT: The Surface Water and Ocean Topography Mission,” Jet Propulsion Laboratory, JPL Publication 12-05, https://swot.jpl.nasa.gov/system/documents/files/2179_SWOT_MSD_final-3-26-12.pdf, Pasadena, CA, February, 2012.
- [2] JPL D-61923, Revision B, “SWOT Science Requirements Document,” Jet Propulsion Laboratory Internal Document, Pasadena, CA, 2018.
- [3] P. Vaze, S. Neeck, W. Bannoura, J. Green, A. Wade, M. Mignogno, G. Zaouche, V. Couderc, E. Thouvenout and F. Parisot, “The Jason-3 Mission: Completing the Transition of Ocean Altimetry from Research to Operations,” *Proceedings SPIE Remote Sensing, Sensors, Systems, and Next-Generation Satellites XIV*, vol. 7826, 2010, <https://doi.org/10.1117/12.868543>.
- [4] J.-F. Cretaux, K. Nielsen, F. Frappart, F. Papa, F. Calmant and J. Benveniste, “Hydrological applications of satellite altimeter: rivers, lakes, man-made reservoirs, inundated areas,” In: *Satellite Altimetry Over Oceans and Land Surfaces*, Stammer, D., Cazanave, A. (Eds.), CRC Press, Earth Observation of Global Changes, 2017.
- [5] C. Donlon et al., “The Copernicus Sentinel-6 mission: Enhanced Continuity of Satellite Sea Level Measurements From Space,” *Remote Sensing of Environment*, vol. 258, 2021, <https://doi.org/10.1016/j.rse.2021.112395>.
- [6] JPL D-56407, Revision C, “SWOT Product Description Document: Level 2 KaRIn Low Rate Sea Surface Height (L2_LR_SSH) Data Product,” Jet Propulsion Laboratory Internal Document, 2025.
- [7] JPL D-105502, “SWOT Algorithm Theoretical Basis Document: Level 2 KaRIn Low Rate Sea Surface Height (L2_LR_SSH) Science Algorithm Software,” Jet Propulsion Laboratory Internal Document, 2023.
- [8] JPL D-61923, Revision C, “SWOT Product Description Document: Level 1B KaRIn Low Rate Interferogram (L1B_LR_INTF) Data Product,” Jet Propulsion Laboratory Internal Document, 2025.
- [9] JPL D-105501, “SWOT Algorithm Theoretical Basis Document: Level 1B KaRIn Low Rate Interferogram (L1B_LR_INTF) Science Algorithm Software,” Jet Propulsion Laboratory Internal Document, 2023.

- [10] SALP-ST-M-EA-17043-CN, "SALP Products Specification: SWOT Nadir Altimeter (L2_NALT_GDR) User Products," Centre National d'Études Spatiales Internal Document, 2023.
- [11] JPL D-56417, Revision B, "SWOT Product Description Document: Level 2 Radiometer Geophysical Data Record (L2_RAD_GDR)," Jet Propulsion Laboratory Internal Document, 2022.
- [12] JPL D-105809, "SWOT Algorithm Theoretical Basis Document: Level 2 Radiometer Geophysical Data Record (L2_RAD_GDR) Science Algorithm Software," Jet Propulsion Laboratory Internal Document, 2023.
- [13] JPL D-56411, Revision C, "SWOT Product Description Document: Level 2 KaRIn High Rate Water Mask Pixel Cloud (L2_HR_PIXC) Data Product," Jet Propulsion Laboratory Internal Document, 2025.
- [14] JPL D-105504, "SWOT Algorithm Theoretical Basis Document: Level 2 KaRIn High Rate Pixel Cloud (L2_HR_PIXC) Science Algorithm Software," Jet Propulsion Laboratory Internal Document, 2023.
- [15] JPL D-56413, Revision C, "SWOT Product Description Document: Level 2 KaRIn High Rate River Single Pass Vector (L2_HR_RiverSP) Data Product," Jet Propulsion Laboratory Internal Document, 2025.
- [16] JPL D-105505, "SWOT Algorithm Theoretical Basis Document: Level 2 KaRIn High Rate River Single Pass (L2_HR_RiverSP) Science Algorithm Software," Jet Propulsion Laboratory Internal Document, 2023.
- [17] JPL D-56414, Revision C, "SWOT Product Description Document: Level 2 KaRIn High Rate River Average Vector (L2_HR_RiverAvg) Data Product," Jet Propulsion Laboratory Internal Document, 2025.
- [18] SWOT-TN-CDM-0673-CNES, Revision C, "SWOT Product Description Document: Level 2 KaRIn High Rate Lake Single Pass Vector Product (L2_HR_LakeSP)," Centre National d'Études Spatiales Internal Document, 2025.
- [19] SWOT-NT-CDM-1753-CNES, Revision A, "SWOT Algorithm Theoretical Basis Document: Level 2 KaRIn High Rate Lake Single-Pass Science Algorithm Software (L2_HR_LakeSP)," Centre National d'Études Spatiales Internal Document, 2025.
- [20] SWOT-TN-CDM-0676-CNES, Revision C, "SWOT Product Description Document: Level 2 KaRIn High Rate Lake Average Vector Product (L2_HR_LakeAvg)," Centre National d'Études Spatiales Internal Document, 2025.
- [21] SWOT-TN-CDM-0677-CNES, Revision C, "SWOT Product Description Document: Level 2 KaRIn High Rate Pixel Cloud Vector Attribute Product (L2_HR_PIXCVec)," Centre National d'Études Spatiales Internal Document, 2025.
- [22] JPL D-56416, Revision C, "SWOT Product Description Document: Level 2 KaRIn High Rate Raster (L2_HR_Raster) Data Product," Jet Propulsion Laboratory Internal Document, 2025.

- [23] JPL D-105507, “SWOT Algorithm Theoretical Basis Document: Level 2 KaRIn High Rate Raster (L2_HR_Raster) Science Algorithm Software,” Jet Propulsion Laboratory Internal Document, 2023.
- [24] SWOT-TN-CDM-1695-CNES, Revision A, “SWOT Product Description Document: Level 2 KaRIn High Rate Floodplain DEM Product (L2_HR_FPDEM),” Centre National d'Études Spatiales Internal Document, 2022.
- [25] JPL D-56410, Revision C, “SWOT Product Description Document: Level 1B KaRIn High Rate Single Look Complex (L1B_HR_SLC) Data Product,” Jet Propulsion Laboratory Internal Document, 2025.
- [26] JPL D-105503, “SWOT Algorithm Theoretical Basis Document: Level 1B KaRIn High Rate Single Look Complex Data (L1B_HR_SLC) Science Algorithm Software,” Jet Propulsion Laboratory Internal Document, 2023.
- [27] SWOT-IS-CDM-0658-CNES, Version 1.2, “SWOT Product Description Document: Precise and Medium-Accuracy Orbit (POE and MOE) Ephemeris Data Product,” Centre National d'Études Spatiales Internal Document, 2021.
- [28] SWOT-IS-CDM-0684-CNES, Version 1.3, “SWOT Product Description Document: Reconstructed Attitude (ATTD_RECONST) Product,” Centre National d'Études Spatiales Internal Document, 2022.
- [29] SWOT-IS-CDM-1073-CNES, Version 1.1, “SWOT Product Description Document: Satellite Center of Mass (SAT_COM) Product,” Centre National d'Études Spatiales Internal Document, 2021.
- [30] SWOT-IS-CDM-1508-CNES, “SWOT Product Description Document: Level 1 DORIS Payload Tracking Data (L1_DORIS_RINEX) Product in RINEX Format,” Centre National d'Études Spatiales Internal Document, 2020.
- [31] JPL D-56447, Revision A, “SWOT Product Description Document: Level 1 GPS Payload Tracking Data (L1_GPSP_RINEX) Product in RINEX Format,” Jet Propulsion Laboratory Internal Document, 2020.
- [32] JPL D-102104, “SWOT Science Data Product Granule Boundary and Sampling Definition,” Jet Propulsion Laboratory Internal Document, 2022.
- [33] Z. Altamimi, P. Rebischung, L. Métivier, and X. Collilieux, “ITRF2014: A new release of the International Terrestrial Reference Frame modeling nonlinear station motions,” *Journal of Geophysical Research*, vol. 121, pp. 6109–6131, 2016, doi: 10.1002/2016JB013098.
- [34] Z. Altamimi, P. Rebischung, X. Collilieux et al., “ITRF2020: an augmented reference frame refining the modeling of nonlinear station motions,” *Journal of Geodesy*, vol. 97, no. 47, 2023, doi: 10.1007/s00190-023-01738-w.
- [35] World Geodetic System 1984 (WGS 84), <https://earth-info.nga.mil/index.php?dir=wgs84&action=wgs84>.

- [36] H. Vermeille, "Direct Transformation from Geocentric Coordinates to Geodetic Coordinates," *Journal of Geodesy*, vol. 76, no. 8, pp. 451–454, Nov. 2002.
- [37] J. Y. Tien, B. Bachman Okhiro, S. X. Esterhuizen, G. W. Franklin, T. K. Meehan, T. N. Munson, D. E. Robison, D. Turbinder, L. E. Young, "Next generation scalable spaceborne GNSS science receiver," *Proceedings of the 2012 International Technical Meeting of the Institute of Navigation*, pp. 882–914, 2012.
- [38] G. Petit and B. Luzum, IERS Conventions 2010 (IERS Technical Note 36), Verlag des Bundesamts für Kartographie und Geodäsie, Frankfurt am Main, 179 pp, 2010, https://iers-conventions.obspm.fr/conventions_material.php.
- [39] SALP-NT-BORD-OP-16137-CN, "DORIS satellites models implemented in POE processing," Centre National d'Études Spatiales Internal Document, 2023, Available online at <https://ids-doris.org/documents/BC/satellites/DORISSatelliteModels.pdf>.
- [40] A. Couhert, J. Moyard, F. Mercier, S. Desai, F. Lemoine, S. Houry, D. Vincent and P. Perrachon, "Copernicus Sentinel-6 Michael Freilich: a New Laboratory for Precise Orbit Determination," in *44th COSPAR Scientific Assembly*, Athens, Greece, July 16–24, 2022.
- [41] E. H. Altenau, T. M. Pavelsky, M. T. Durand, X. Yang, R. P. d. M. Frasson and L. Bendezu, "The Surface Water and Ocean Topography (SWOT) Mission River Database (SWORD): A Global River Network for Satellite Data Products," *Water Resources Research*, vol. WRCS25408, 2021, <https://doi.org/10.1029/2021WR030054>.
- [42] SWOT-IS-CDM-1944-CNES, Revision C, "SWOT Auxiliary Data Description Document: Prior Lake Database (LakeDatabase)," Centre National d'Études Spatiales Internal Document, 2025.
- [43] J. C. Curlander and R. N. McDonough, *Synthetic Aperture Radar: Systems and Signal Processing*, (Wiley, New York, 1991).
- [44] 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*, vol. 88, no. 3, pp. 333–382, Mar. 2000.
- [45] R. Bamler and P. Hartl, "Synthetic Aperture Radar Interferometry," *Inverse Problems*, vol. 14, pp. R1–R54, 1998.
- [46] H. A. Zebker and J. Villasenor, "Decorrelation in Interferometric Radar Echoes," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 30, no. 5, pp. 950–959, Sep. 1992.
- [47] F. Gatelli, A. M. Guarnieri, F. Parizzi, P. Pasquali, C. Prati, and F. Rocca, "The Wavenumber Shift in SAR Interferometry," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 32, no. 4, pp. 855–865, Jul. 1994.
- [48] E. Rodriguez and J. M. Martin, "Theory and design of interferometric synthetic aperture radars," *IEEE Proceedings-F*, vol. 139, no. 2, pp. 147–159, Apr. 1992.

- [49] JPL D-79130, Revision B, “KaRIn: Ka-band Radar Interferometer On-Board Processor (OBP) Algorithm Theoretical Basis Document (ATBD),” Jet Propulsion Laboratory Internal Document, 2021.
- [50] JPL D-55533, Revision B, “KaRIn: Ka-band Radar Interferometer On-Board Processor (OBP) for Land Algorithm Theoretical Basis Document (ATBD),” Jet Propulsion Laboratory Internal Document, 2017.
- [51] R. Kwok and W. T. K. Johnson, “Block Adaptive Quantization of Magellan SAR Data,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 27, no. 4, pp. 375–383, Jul. 1989.
- [52] C. J. Oliver and S. Quegan, *Understanding Synthetic Aperture Radar Images*, (Artech House, Boston, 1998).
- [53] J. W. Goodman, *Statistical Optics*, Wiley-Interscience, New York, 1985.
- [54] M. Durand, C. Chen, R. P. d. M. Frasson, T. M. Pavelsky, B. Williams, X. Yang, and A. Fore, “How Will Radar Layover Impact SWOT Measurements of Water Surface Elevation and Slope, and Estimates of River Discharge?” *Remote Sensing of Environment*, vol. 247, Sept. 2020, doi: 10.1016/j.rse.2020.111883.
- [55] E. Peral, E. Rodriguez, and D. Esteban-Fernandez, “Impact of Surface Waves on SWOT’s Projected Ocean Accuracy,” *Remote Sensing*, vol. 7, pp. 14509–14529, 2015, doi: 10.3390/rs71114509.
- [56] E. Peral, D. Esteban-Fernández, E. Rodríguez, D. McWatters, J.-W. De Bleser, R. Ahmed, A. C. Chen, E. Slimko, R. Somawardhana, K. Knarr, M. Johnson, S. Jaruwatanadilok, S. Chan, X. Wu, D. Clark, K. Peters, C. W. Chen, P. Mao, B. Khayatian, J. Chen, R. E. Hodges, D. Bousalis, B. Styles, and K. Srinivasan, “KaRIn, the Ka-Band Radar Interferometer of the SWOT Mission: Design and In-Flight Performance,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 62, pp. 1–27, 2024, doi: 10.1109/TGRS.2024.3405343.
- [57] P. Escudier, A. Couhert, F. Mercier, A. Mallet, P. Thibaut, N. Tran, L. Amarouche, B. Picard, L. Carrere, G. Dibarboue, M. Ablain, J. Richard, N. Steunou, P. Dubois, M. Rio, and J. Dorandeu, “Satellite Radar Altimetry: Principle, Accuracy, and Precision,” in *Satellite Altimetry over Oceans and Land Surfaces*, D. Stammer and A. Cazenave, Ed., Boca Raton, CRC Press, 2018, pp. 1–69.
- [58] E. Rodriguez, D. Esteban Fernandez, E. Peral, C. W. Chen, J.-W. De Bleser, and B. Williams, “Wide-Swath Altimetry: A Review,” in *Satellite Altimetry over Oceans and Land Surfaces*, D. Stammer and A. Cazenave, Ed., Boca Raton, CRC Press, 2018, pp. 71–112.
- [59] JPL D-79084, Revision A, “SWOT Mission Performance and Error Budget,” Jet Propulsion Laboratory Internal Document, 2017.
- [60] J. Saastamoinen, “Atmospheric correction for the troposphere and stratosphere in radio ranging for satellite,” *Geophys. Monogr.*, vol. 15, 1972.

- [61] R. D. Ray and R. M. Ponte, “Barometric tides from ECMWF operational analyses,” *Ann. Geophys.*, vol. 21, pp. 1897–1910, 2003, doi:10.5195/angeo-21-1897-2003.
- [62] D. A. Imel, “Evaluation of the TOPEX/POSEIDON dual-frequency ionosphere correction,” vol. 99, no. C12, pp. 24895–24906, 1994, doi:10.1029/94JC01869.
- [63] A. Mannucci, B. Wilson, D. Yuan, C. Ho, U. Lindqwister, and T. Runge, “A global mapping technique for GPS-derived ionospheric total electron content measurements,” *Radio Science*, vol. 33, no. 3, pp. 565–582, 1998, doi:10.1029/97RS02707.
- [64] N. K. Pavlis, S. A. Holmes, S. C. Kenyon and J. K. Factor, “The development and evaluation of the Earth Gravity Model 2008 (EGM2008),” *Journal of Geophysical Research Solid Earth*, vol. 117, pp. 1978–2012, 2022, doi:10.1029/2011JB008916.
- [65] D. E. Cartwright and R. J. Taylor, “New computations of the tide-generating potential,” *Geophysical Journal International*, vol. 23, pp. 45–73, 1971, doi:10.1111/j.1365-246X.1971.tb01803.x.
- [66] D. E. Cartwright and A. C. Edden, “Corrected tables of tidal harmonics,” *Geophysical Journal International*, vol. 33, pp. 253–264, 1973, doi:10.1111/j.1365-246X.1973.tb03420.x.
- [67] AVISO, “Mean Sea Surface MSS_CNES_CLS2022 Description,” [Online], <https://www.aviso.altimetry.fr/en/data/products/auxiliary-products/mss/mss-cnes-cls2022-description.html>.
- [68] J. M. Wahr, “Body tides on an elliptical, rotating, elastic and oceanless earth,” *Geophysical Journal International*, vol. 64, no. 3, pp. 677–703, March 1981, doi: 10.1111/j.1365-246X.1981.tb02690.x.
- [69] R. D. Ray, “Precise comparisons of bottom-pressure and altimetric ocean tides,” *Journal of Geophysical Research Oceans*, vol. 118, pp. 4570–4584, 2013, doi:10.1002/jgrc.20336.
- [70] R. D. Ray and S. Y. Erofeeva, “Long-period tidal variations in the length of day,” *Journal of Geophysical Research Solid Earth*, vol. 119, pp. 1498–1509, 2014, doi:10.1002/2013JB010830.
- [71] F. H. Lyard, D. J. Allain, M. Cancet, L. Carrère, and N. Picot, “FES2014 global ocean tide atlas: design and performance,” *Ocean Science*, vol. 17, pp. 615–649, 2021, doi:10.5194/os-17-615-2021.
- [72] L. Carrère, B. K. Arabic, B. Dushaw, G. Egbert, S. Erofeeva, F. Lyard, R. D. Ray, C. Ubelmann, E. Zaron, Z. Zhao, J. F. Shriver, M. C. Buijsman, and N. Picot, “Accuracy assessment of global internal-tide models using satellite altimetry,” *Ocean Science*, vol. 17, pp. 147–180, 2021, doi:10.5194/os-17-147-2021.
- [73] J. M. Wahr, “Deformation induced by polar motion,” *Journal of Geophysical Research*, vol. 90, no. B11, pp. 9363–9368, 1985, doi: 10.1029/JB090iB11p09363.
- [74] J. C. Ries and S. D. Desai, “Conventional model update for rotational deformation,” in Fall AGU Meeting, New Orleans, LA, 2017, doi: 10.26153/tsw/2659.

- [75] S. Desai, J. Wahr, and B. Beckley, "Revisiting the pole tide for and from satellite altimetry," *Journal of Geodesy*, vol. 89, pp. 1233–1243, 2015, doi: 10.1007/s00190-015-0848-7.
- [76] C. Wunsch, "Bermuda sea level in relation to tides, weather and baroclinic fluctuations," *Rev. Geophys. Space Phys.*, vol. 10, pp. 1–49, 1972, doi:10.1029/RG010i001p00001.
- [77] AVISO, "Dynamic Atmospheric Correction: AVISO+," [Online], <https://www.aviso.altimetry.fr/en/data/products/auxiliary-products/dynamic-atmospheric-correction/description-atmospheric-corrections.html>.

Acknowledgements

A portion of the work described in this document was performed at the Jet Propulsion Laboratory, California Institute of Technology under contract with the National Aeronautics and Space Administration.

Appendix A

Acronyms

ADC	Analog-to-Digital Converter
AL	Aperture Length
AMR	Advanced Microwave Radiometer
ATBD	Algorithm Theoretical Basis Document
BFPQ	Block Floating-Point Quantization
CMR	Common Metadata Repository
CNES	Centre National d'Études Spatiales
CRID	Composite Release Identifier
DAA	Deployable Antenna Assembly
DAC	Dynamic Atmosphere Correction
DAD	Dynamic Auxiliary Data
DEM	Digital Elevation Model
DFT	Discrete Fourier Transform
DORIS	Doppler Orbitography and Radiopositioning Integrated by Satellite
DWP	Data Window Position
ECEF	Earth-Centered Earth-Fixed
ECMWF	European Centre for Medium-range Weather Forecasts
EMB	Electromagnetic Bias
FFT	Fast Fourier Transform
FPDEM	Floodplain Digital Elevation Model
FPGA	Field Programmable Gate Array
FSW	Flight Software
FWHM	Full-Width at Half Maximum
GCRF	Geocentric Celestial Reference Frame
GDR	Geophysical Data Record
GIM	Global Ionosphere Map
GNSS	Global Navigation Satellite System
GPS	Global Positioning System

GPSP	Global Positioning System Payload
HB	Hyperbox
HPA	High-Power Amplifier
HR	High-Rate
HVPS	High-Voltage Power Supply
IDS	International DORIS Service
IFFT	Inverse Fast Fourier Transform
IGDR	Interim Geophysical Data Record
ILRS	International Laser Ranging Service
IRF	Impulse Response Function
IRU	Inertial Reference Unit
ITRF	International Terrestrial Reference Frame
KaRIn	Ka-band Radar Interferometer
KDES	KaRIn Digital Electronics Subsystem
KMSF	KaRIn Metering Structure Frame
LEOP	Launch and Early Operations Phase
LLH	Latitude-Longitude-Height
LR	Low-Rate
LRA	Laser Retroreflector Array
MDT	Mean Dynamic Topography
MLI	Multilayer Insulation
MOE	Medium-accuracy Orbit Ephemeris
MSS	Mean Sea Surface
NAIt	Nadir Altimeter
NASA	National Aeronautics and Space Administration
NRCS	Normalized Radar Cross Section
OBP	On-Board Processor
OCM	Orbit Control Maneuver
OGDR	Operational Geophysical Data Record
PB	Posting Boundary
PCU	Processing Unit
PDD	Product Description Document
PDU	Power Distribution Unit
PGE	Product Generation Executable
PIXC	Pixel Cloud
PLD	Prior Lake Database
PO.DAAC	Physical Oceanography Distributed Active Archive Center
POE	Precise Orbit Ephemeris
PRD	Prior River Database
PRF	Pulse Repetition Frequency
PRI	Pulse Repetition Interval
PSF	Point Spread Function

PTR	Point Target Response
RCRF	Range Compression Reference Function
RCS	Radar Cross Section
RFU	Radio Frequency Unit
RINEX	Receiver Independent Exchange
SAD	Static Auxiliary Data
SAR	Synthetic Aperture Radar
SDP	Science Data Product
SEU	Single Event Upset
SLC	Single-Look Complex
SNR	Signal to Noise Ratio
SSALTO	Segment Sol multimissions d' ALTimétrie, d' Orbitographie et de localisation
SSB	Sea State Bias
SSH	Sea Surface Height
SSHA	Sea Surface Height Anomaly
SSR	Solid State Recorder
SWH	Significant Wave Height
SWORD	SWOT River Database
SWOT	Surface Water and Ocean Topography
TAI	International Atomic Time
TCN	Track, Cross-Track, Nadir
TDS	THREDDS Data Server
TEC	Total Electron Content
TECU	Total Electron Content Unit
TRI	Transmit Repetition Interval
TT	Terrestrial Time
UTC	Coordinated Universal Time
UTM	Universal Transverse Mercator
WSE	Water Surface Elevation