Surface Water and Ocean Topography (SWOT) Project

Science Data Product Granule Boundary and Sampling Definition

Initial Release

Document Custodians:

Curtis Chen JPL Algorithm System Engineer	Date	Roger Fjørtoft CNES HR Algorithm System Engineer	Date

Approved by:

Shailen Desai JPL Measurement System Engineer	Date	Nicolas Picot CNES Measurement System Engineer	Date

Concurred by:

Oh-Ig Kwoun JPL SDS Manager	Date	Hélène Vadon CNES SDS Manager	Date

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 Product Data Management System

See Sect 1.3 for how to cite this document

September 19, 2022
JPL D-102104

National Aeronautics and Space Administration
Jet Propulsion Laboratory
California Institute of Technology

CHANGE LOG

VERSION	DATE	SECTIONS CHANGED	REASON FOR CHANGE
Draft	2022-09-15	ALL	Approved for export (LRS-067559)
Initial Release	2022-09-19	ALL	Initial Release Approved for public release (CL#22-4860)

Table of Contents

Table of Figures

Table of Tables

List of TBC Items

Page	Section

List of TBD Items

Page	Section

1 Introduction

1.1 Purpose and Scope

This document describes the manner in which granule boundaries and sampling are defined for SWOT science data products. Because the definitions for low-rate (LR) and high-rate (HR) KaRIn data products are related, they are both considered here. The greatest emphasis of this document is on KaRIn data products, though it also describes the radiometer and nadir altimeter products for completeness and context.

The Product Description Documents (PDDs) for the various data products provide information on specific products that should be consistent with this document. This document is intended to give a high-level overview of the architectural design and the details of the relationships between the sampling and granule definitions of the different products, while the PDDs are intended to give specific information on individual products for product users. In the event of conflicts, the PDDs take precedence over this document.

1.2 Document Organization

Section 2 provides a set of definitions for use within the context of this document in order to allow precise interpretation of the content described here.

Section 3 provides a qualitative description of the general relationships between the granule boundary definitions and sampling schemes across the various KaRIn data products in order to provide the rationale for the definitions in subsequent sections.

Section 4 provides definitions for a set of idealized reference tile boundaries that are used to define the actual tile boundaries for a number of KaRIn data products.

The remaining sections give details on the sampling and granule boundaries of the various SWOT science products, emphasizing the relationships between the definitions for different products.

1.3 Citing This Document

Please cite this document as follows:

> JPL D-102104, "SWOT Science Data Product Granule Boundary and Sampling Definition," Jet Propulsion Laboratory Internal Document, 2022.

2 Definitions for Granule Boundaries and Sampling

2.1 Granule Size Definitions

A data product is a family or class of outputs from the SWOT science data system that provides a particular family or class of SWOT measurement information. The definition of a data product is documented in its corresponding PDD. In many cases, a data product involves only one kind of file, but in some cases, a data product comprises multiple kinds of files that capture different but related information. For example, the L2_LR_SSH product includes four files (Basic, WindWave, Expert, and Unsmoothed), and the L2_HR_RiverSP product includes two shapefiles (for reaches and nodes), each of which consists of multiple files that are parts of the shapefile format.

Data products over the continuous stream of SWOT mission data are additionally split into spatial or temporal granules, which represent the quanta for the sizes of the data product files that are distributed to users. In other words, a granule is a "chunk" of a data product that is placed into a file or a set of files when produced operationally by the SWOT science data system and made available for users to download. Each granule covers a predefined spatial or temporal extent. Different products have different granule sizes and definitions, the details of which are described in this document.

Granules are defined with respect to the terms below. For the purposes of defining granules, the following definitions are assumed in this document:

- *Cycle*: One orbital repeat period, which is approximately 21 days for the science orbit and 1 day for the calibration orbit.
- *Orbit* or *revolution*: Each turn of the SWOT spacecraft around the Earth is called a revolution-(approximately 100 minutes). There are 292 revolutions in a 21 day repeat cycle, numbered chronologically from 1 to 292. There are 14 revolutions in a 1 day repeat cycle, also numbered chronologically from 1 to 14.
- *Pass*: Half an orbit, split at approximately the southernmost and northernmost points in the orbit. Ascending and descending passes are those in which the spacecraft travels from the south to the north and from the north to the south, respectively. Details on the exact definitions of pass boundaries for the purposes of SWOT data product definitions are discussed later in this document. There are 584 passes in a 21 day repeat cycle, numbered chronologically from 1 to 584. There are 28 passes in a 1 day repeat cycle, numbered chronologically from 1 to 28. Odd and even numbered passes are ascending and descending, respectively.
- *Swath* or *full swath*: The cross-track extent of KaRIn observations at a given time, spanning both the left and right sides from nadir. The swath, from approximately 4˚ incidence angle on one side to approximately 4˚ incidence angle on the other side, is approximately 120 km wide. The far-range extent of the swath varies. The sampled extent of the swath is determined by KaRIn parameters and the viewing geometry, and the swath extent meeting minimum performance specifications is determined by the surface reflectivity and other factors.

- *Half swath* or *side*: One side (left or right) of the swath from nadir. A half swath is approximately 60 km wide from nadir to approximately 4° incidence angle. Note that KaRIn data between nadir and approximately 10 km (to either side) are not subject to performance requirements, but it is useful to include data all the way to nadir in product granules.
- *Left* and *right*: Left and right, in reference to half swaths or look directions, are defined in this document such that the cross product of the local up vector and the Earth-relative velocity vector points generally to the left. In other words, the terms "left" and "right" are defined in this document with respect to a viewer standing on the Earth at the nadir point and facing in the direction of the spacecraft Earth-relative velocity vector. Left and right in this context are thus independent of the spacecraft attitude (particularly the yaw state). However, which of the left or right sides is east or west of the nadir track depends on whether the pass is ascending or descending.
- *Tile*: A granule that covers a half swath in cross track and approximately 60 km along track to give a nearly square aspect ratio. Details of the tile definition are provided in this document.
- *Scene*: A granule that covers a full swath in cross track and approximately 120 km along track to give a nearly square aspect ratio. Details of the scene definition are provided in this document.

Note that the approximate sizes described above are exclusive of padding or overlap between granules that are introduced for processing purposes in some cases.

2.2 Data Sampling

The sampling of a data product refers to the manner in which individual data points within a product are arranged spatially and/or temporally. The terms "sampling" and "posting" are used interchangeably in this document. While the sampling may be related to the granule boundaries, the sampling does not necessarily need to be tied to or aligned with the granule boundaries. The arrangement of the valid data in the files of a data product may also be decoupled from the sampling and granule definition. For example, see Figure 1, in which different grids of data samples, illustrated by blue points, may be assumed for the same granule boundaries, illustrated by the red lines. The valid data may be arranged in the file in various ways as well.

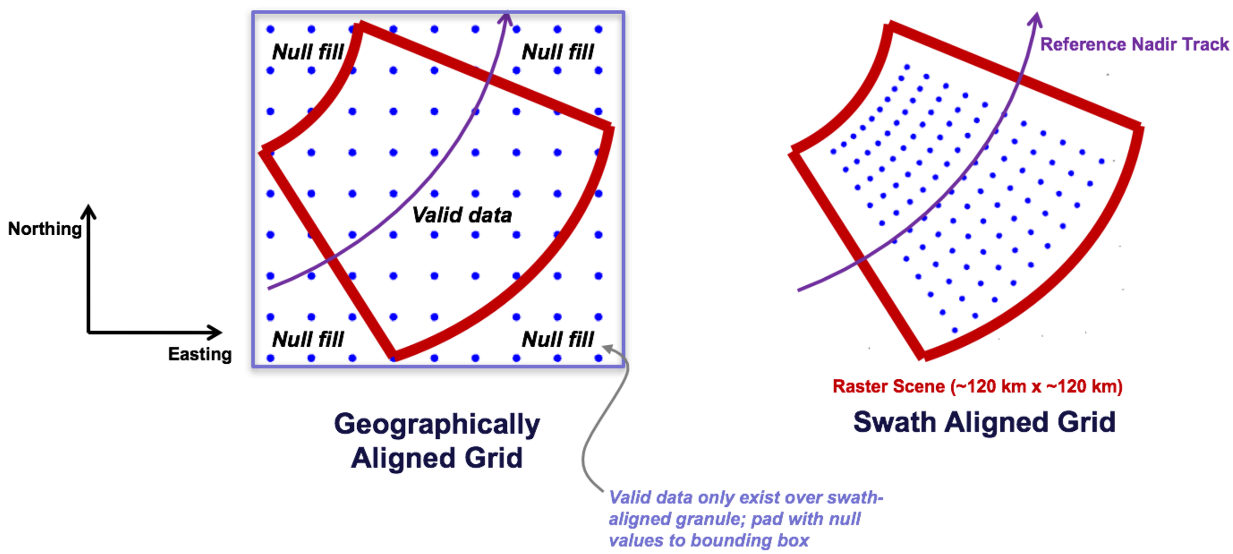

Figure 1. Illustration of granule boundaries in relation to sampling.

Note also that a given data sample may have different coordinates for different purposes. For example, a sample of the L2_HR_PIXC pixel-cloud (PIXC) product is associated with a reference location as well as a measurement-based location.

This document is focused on granule boundaries more so than sample arrangement. Sampling is discussed mainly insofar as it is related to the granule boundaries. The product description documents for different products provide more details on data sampling.

3 General Relationships Between Granule Definitions

3.1 Granule Relationships

This section provides a high-level overview of the relationships between granules of different science data products before discussing the details of granule boundary definitions. The granule sizes are as follows:

- *Pass Granules*: KaRIn LR data products are produced on pass-sized granules, with passes defined in Section 2.1. However, as described later in this document, the exact boundaries between passes are not necessarily precisely the same between native-grid and fixed-grid LR data files.
- *Tile Granules*: KaRIn HR Single Look Complex (SLC) and Pixel Cloud (PIXC) data products are produced on tile-sized granules whose definitions are related as described later in this document. The tile boundaries are approximately fixed geographically. The nominal tile boundaries are tied to the LR fixed grid-sample definition.
- *Scene Granules*: KaRIn HR raster products are produced on scene-sized granules. A scene is nominally defined to correspond to a two-by-two set of tiles and are fixed geographically.
- *Continent-Pass Granules*: KaRIn HR river and lake single-pass (SP) products are produced on continent-pass-sized granules. These represent the intersection of a pass with a predefined continent.
- *Basin Granules*: KaRIn HR river and lake average products are produced on basin-sized granules. These correspond to all data from a given time span (nominally an orbit repeat cycle) that are contained within a predefined basin area.

3.2 Granule Overlap

The granules for some products overlap with the previous and next granules along track. The overlapping of granules allows simplifications in the processing software, which often requires input over a greater spatial extent than the extent of the output it produces. Overlap of the granules for specific data products is adopted when favored by the trade between the simplicity of the software and simplicity of the data products. In this document, the amount of overlap and the manner in which the overlap region is defined are described for each data product where applicable. In the absence of any discussion of overlap for a given data product, the granules should otherwise be assumed not to overlap.

3.3 Temporal vs. Spatial Granule Boundaries

A natural distinction exists between data products whose contents are more closely linked to SWOT observation time and those whose contents are more closely linked to spatial location of the observations. That is, some data products align more closely with the temporal spans over which the data were collected, regardless of the spacecraft position or attitude. Where measurements of the Earth are made, data samples in these temporally based products are typically geolocated based on the spacecraft state at the measurement times rather than processed to geographically fixed coordinates. Temporally defined products include:

- MOE/POE
- ATTD_RECONST
- L2_RAD_GDR, L2_RAD_IGDR, L2_RAD_OGDR
- L2_NALT_GDR, L2_NALT_IGDR, L2_NALT_OGDR
- L1B_LR_INTF
- L2_LR_SSH (native-grid file)

On the other hand, some data products have sample locations and/or granule boundaries that are tied to geographic definitions. Such geographic definitions may be related to time if the spacecraft position and attitude are known, but their spatial attributes remain their fundamental defining quantities. Note that because the granule boundaries and the sample posting of a product are separate (see Figure 1), it is possible for the granule boundaries to be geographically defined while the data samples within those granules are natively defined based on the spacecraft state at the measurement time. Products with geographically defined boundaries include:

- L1B_HR_SLC
- L2_HR_PIXC
- L2_HR_PIXCVec
- L2_HR_RiverSP
- L2_HR_RiverAvg
- L2_HR_LakeSP
- L2_HR_LakeAvg
- L2_HR_Raster
- L2_LR_SSH (3 fixed-grid files)

Note that the L2_LR_SSH product, which contains multiple files, has temporally defined boundaries for one file and spatially defined boundaries for the other files.

The distinction between temporal and spatial boundaries is considered in the following sections.

3.4 Temporal Pass Boundary Definition

For the data products that are temporally linked (see the previous section) and produced on pass-sized granules, it is not important for the granule boundaries to coincide precisely with specific locations relative to the planet. For example, the granule boundaries do not need to be at exactly the northernmost or southernmost points in the orbit. The granule boundaries should simply be close enough to the latitude extrema that the pass definition is intuitive, the naming conventions are unambiguous, etc. It is most important that the granule boundaries be defined in a way that the definitions can be applied consistently across the development, processing, and user communities.

The granule boundaries for the temporal, pass-based products are therefore defined by the times of pass boundaries in the Orbit Revolution File (ORF) produced by the CNES Mission System. For the purposes of granule boundary definition, the boundaries for each pass are defined exactly once and are not changed once defined (for example, reprocessing will not result in a change of the times that mark the beginning and end of any given pass). The following definition for temporal pass-sized granules is therefore adopted:

Definition: The following data products are produced in instrument telemetry download granules whose boundaries are defined in time by the first and last available measurements in those data. Each telemetry download usually spans ~ 1.5-2 hours or less, but may be longer if a telemetry downlink was missed.

- L2_RAD_OGDR
- L2_NALT_OGDR

Definition**:** The following data products are produced in pass-sized granules whose boundaries are defined in time by the times of pass boundaries from the ORF:

- L2_RAD_GDR, L2_RAD_IGDR
- L2_NALT_GDR, L2_NALT_IGDR
- L1B_LR_INTF
- L2_LR_SSH (native-grid file)

To be precise, a data sample should fall within a particular pass granule if the time tag of the data sample lies between the spacecraft event times (SCET) from the ORF that mark the beginning and end of that granule. For data products whose samples can be associated with time tags for different types of events, such as KaRIn transmit and KaRIn receive times, the data product definition indicates which specific event is to be used for partitioning the data into granules.

Definition**:** The starting times of as-flown (not predicted) ascending and descending passes from the daily updates of the ORF are used to define the temporal granule boundaries of associated temporal-pass science products.

It is desirable for boundaries of temporally defined KaRIn granules to be independent of the product generation latency. Provided that the latency is on the order of days or longer (not hours), the ORF supports this, as it is based on MOE. The use of reconstructed rather than predicted pass starting times ensures that the granule boundaries are uniquely defined.

3.5 Geographically Fixed Along-Track Sample Locations

Data products whose granule boundaries are geographically linked are tied to the swath-aligned, geographically fixed sampling used for the Basic file of the L2_LR_SSH data product.

In this geographically fixed grid, the sample spacing along the nadir track of the reference orbit on the reference ellipsoid (nominally the WGS84 ellipsoid) is constant. However, there is no closed-form analytical method for determining the positions of evenly spaced points along an ellipsoid (the solution involves elliptic integrals), so care must be taken to avoid inconsistencies between quantities computed by different entities, lest accumulated errors in integrating paths along the ellipsoid give significant differences in sample locations or granule boundaries if different computational methods are assumed.

In order to ensure consistency in the computation of geographically fixed sample locations and granule boundaries, the following approach is adopted.

The CNES Mission System, which is responsible for defining the reference orbit, provides the latitude, longitude, and heading (defined below) of spacecraft nadir locations on the WGS84 ellipsoid spaced every 125 m along the reference track for the first ascending pass and the first descending pass. The samples are arranged such that one sample falls exactly on the equator. The spacing of 125 m allows both edges and centers of 250 m pixels to be precisely located.

The heading is defined as the angle, measured positive toward the east from true north, of the projection of the Earth-relative spacecraft velocity vector into the local horizontal frame. This is illustrated in Figure 2. That is, let X_{YS}, Y_{YS}, Z_{YS}, be unit vectors along the axes of the Track Compensation Reference Frame (TCRF, often called the Yaw Steering or YS frame), where Z_{YS} is in the local geodetic up direction, Y_{YS} is in the direction of the cross product of Z_{YS} and the Earth-relative spacecraft velocity vector, and X_{YS} completes the right-handed coordinate system. Therefore, X_{YS} is in the direction of the component of the Earth-relative velocity vector in the local horizontal frame (normal to Z_{YS}). Let E and N be unit vectors in the local east and north directions at the nadir point. The heading angle ϕ_h can be linked to TCRF by the following relations: $\cos\phi_h = X_{YS} \bullet N$ and $\sin\phi_h = X_{YS} \bullet E$, with the dot operator ($\bullet$) indicating the scalar product of two vectors. It is thus straightforward to determine X_{YS} and Y_{YS} from ϕ_h or to determine ϕ_h from X_{YS} and Y_{YS}.

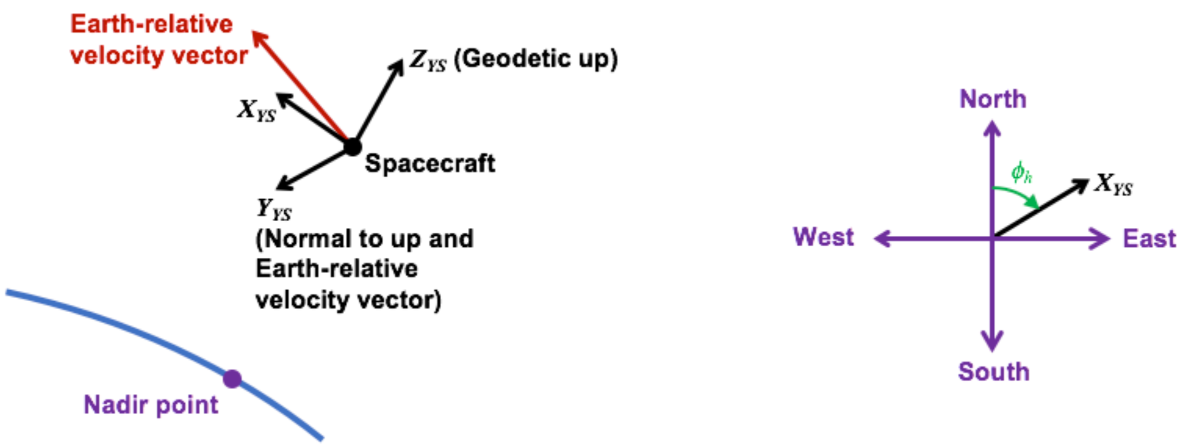

Figure 2. Illustration of the definition of the heading with respect to the TCRF coordinate system.

Note that because the Earth pole position is not known precisely in advance (i.e., when defining the reference orbit), the rotating Earth frame above is assumed to be the Terrestrial Intermediate Reference Frame (TIRF), which is a nominal Earth frame that does not include the small variations of the pole position.

The samples of latitude, longitude, and heading every 125 m following the reference nadir track along the ellipsoid starting in either direction from the equator are used to define granule boundaries and sampling grids as described below.

Much of the discussion in this document assumes the nominal orbit (21 day repeat period). However, the granule boundaries and sampling grids for the fast-sampling orbit (1 day repeat period) as well as orbit-transition or drifting-orbit periods are defined in the same or a very similar manner as for the nominal orbit. In the case of orbit transitions and drifting orbits, this document assumes that numbering conventions for revolutions and passes are defined as part of mission operations such that the entire phase is analogous to a single repeat period of one of the repeat-orbit phases. Since the orbit is not repeating in this case, the reconstructed MOE orbit is used in place of the reference orbit for definitions in this document during such phases. Short-latency processing of geographically fixed products during orbit transitions and drifting orbits is not planned.

4 Reference Tile Boundaries

4.1 Reference Tile Boundary Definitions

Tile boundaries used for the L1B_HR_SLC, L2_HR_PIXC, and L2_HR_PIXCVec products are swath aligned and approximately fixed geographically. The actual tile boundaries are allowed to vary slightly (generally less than 500 m on the ground) from their ideal fixed geographic boundaries in order to align the samples with the native sampling grid of the SLC product. It is useful to first define reference tile boundaries, which would apply if the spacecraft attitude were ideal and the actual orbit followed the reference orbit exactly, then to define the actual tile boundaries for non-ideal cases.

Definition: Let the reference along-track boundaries of tile k be defined by points T_{k-1} and T_k, which are two locations that fall on and are separated by a nominal tile width $L_T = 64$ km along the reference nadir track over the ellipsoid. Let X_{YSk-1} and X_{YSk} be the along-track directions of the TCRF at T_{k-1} and T_k, defined by the heading angles ϕ_{hk-1} and ϕ_{hk} at these locations. The reference tile boundaries in the along-track direction are the planes defined by the equations $X_{YSk-1} = 0$ and $X_{YSk} = 0$ that go through T_{k-1} and T_k, respectively. Points T_{k-1} for the first tile of the pass and point T_k for the last tile of the pass coincide with the geographic beginning and end of the pass, where $\phi_h = 90^\circ$, in order to avoid gaps in coverage; the lengths of these tiles generally differ from the 64 km nominal length.

Definition: Let the reference cross-track boundaries of a given tile be defined by the reference nadir track and a line $W_T = 64$ km from nadir in the direction (left or right) of the half-swath covered by the tile. The 64 km cross-track distance is defined along a spherical approximation to the ellipsoid where the sphere is tangent to the ellipsoid at the nadir point defining the cross-track direction. The sphere radius R_a is 6378.137 km. The boundary is projected in the local up direction relative to the ellipsoid, not relative to the approximating sphere. See Figure 3.

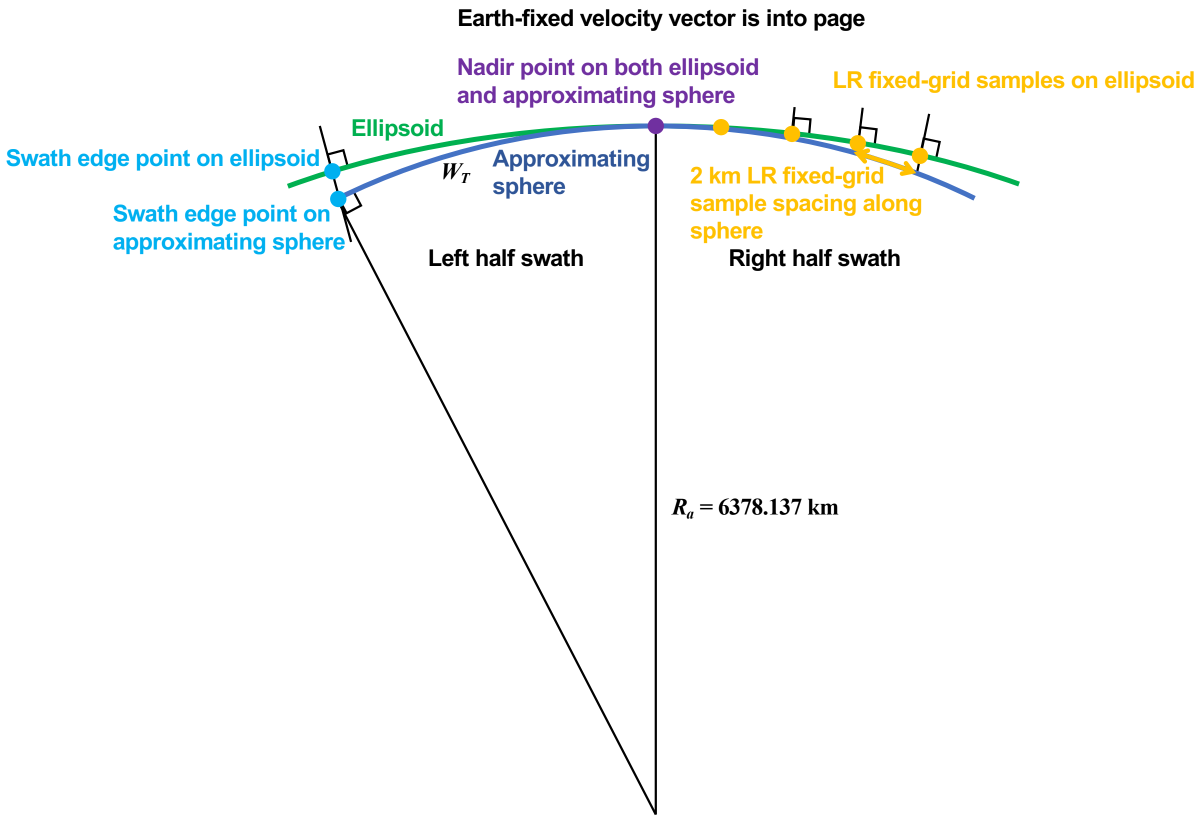

Figure 3. Illustration of the local spherical approximation to the ellipsoid in the cross-track direction.

4.2 Reference Tile Along-Track Length Selection

The nominal along-track length of the tile is chosen based on the following criteria (some are desirable but not strictly required):

- The nominal tile length should approximately match the half-swath cross-track width.
- The nominal tile length should be an integer multiple of the 2 km sampling of the L2_LR_SSH fixed-grid data.
- The number of tiles per pass should be even so that scenes can be defined as 2 x 2 sets of tiles and so that the equator falls on a tile edge rather than a tile center along the nadir track (the tile edge does not coincide with the equator due to the orbit inclination, of course).
- The number of tiles per pass should be a multiple of 4 if scenes are to be defined as 2 x 2 sets of tiles and the number of scenes is to be even so that the equator falls on a scene edge rather than a scene center.
- While the first and last tiles of the pass necessarily have different lengths than the nominal tile length because the length of the pass is not generally an integer multiple of the nominal tile size given the constraints above, the first and last tiles should have lengths that differ from the nominal tile length by no more than 10%. It would be preferable if the first and last times are slightly smaller than the nominal tile size rather

than slightly larger so that they do not drive resource usage, although this condition may not be achievable without simply increasing the nominal tile size such that the nominal size is the driver.

An along-track tile length L_T of 64 km fulfills the above criteria best, at least in a qualitative sense. Options for other tile along-track lengths are shown in Table 1. This gives 308 along-track tiles per pass (616 total, including left and right sides), or 616 x 584 = 359,744 tiles per cycle (including both left and right sides). This global number of tiles assumes that tiles are defined over the entire planet surface, even over the ocean where HR data would typically be collected infrequently (if at all for some areas). However, it is assumed that the tile definition should be complete enough that it can support any realization of the HR mask.

The southernmost and northernmost tiles, which extend to the pass ends, have lengths of 68.4 and 68.8 km, respectively. Since these tiles are larger than other tiles, they drive sizing calculations that are related to the maximum tile size (e.g., memory usage). Note that due to the eccentricity of the orbit, the lengths of the nadir track on either side of the equator are not identical.

Table 1. Number of along-track tiles and first and last tile sizes for different choices of along-track tile length.

Tile Length (km)	Num Along-Track Tiles/Pass	First Asc Tile Length (km)	Last Asc Tile Length (km)	Even Num of Scenes/Pass?
56	352	60.37	60.75	Y
58	340	58.37	58.75	Y
60	328	80.37	80.75	Y
62	318	64.37	64.75	N
64	**308**	**68.37**	**68.75**	**Y**
66	298	92.37	92.75	N
68	290	68.37	68.75	N
70	282	60.37	60.75	N
72	274	68.37	68.75	N
74	266	92.37	92.75	N
76	260	56.37	56.75	Y
78	252	110.37	110.75	Y

4.3 Reference Tile Cross-Track Width Selection

It is not essential that the tile length and width be identical, but it is convenient. However, if desired, the tile cross-track tile width W_T could be changed without changing the along-track tile length L_T. The cross-track tile width may need to be increased if the expected swath width, accounting for cross-track orbit variations, is larger.

The nominal value of W_T = 64 km accounts for a useful swath up to 60 km from the actual nadir point, and, for products defined with respect to the reference nadir track, 2.5 km of cross-track orbit variation, leaving 1.5 km margin or pad for far-range data that may be useful despite not being subject to performance requirements.

Note that for the L1B_HR_SLC, L2_HR_PIXC, and L2_HR_PICVec products, whose boundaries are defined with respect to the actual nadir location (see below) rather than the reference tile boundaries described here, the selection of W_T = 64 km accommodates 4 km beyond the required 60 km outer swath edge. However, these products usually include all data samples that are downlinked and successfully processed, even if they extend beyond 64 km from the nadir track. This is possible because the KaRIn HR data window is defined in slant range, and the mapping from slant range to distance along the Earth surface depends on the local topography.

4.4 Reference Tile Boundary Specification

The latitude, longitude, and heading are distributed for all tile boundary points T along the reference nadir track that defines tile boundaries for the first ascending and first descending passes. Including the points at the geographic pass boundaries, 309 points T are needed per pass to specify the along-track boundaries for the 2 x 308 = 616 tiles in the pass.

With the along-track tile boundaries defined for the first orbit (i.e., the first ascending and the first descending pass) of the cycle, the along-track tile boundaries for the remaining orbits are computed by shifting the longitudes of the points from the first orbit.

Definition: A given point T_{kn} defining the kth along-track tile boundary location for orbit n, with $n > 1$, is identical to the corresponding point T_{k1} for the kth tile of the first orbit ($n = 1$) after adding a constant shift $\Delta\Theta_{\mathrm{lon}}$, given by $\Delta\Theta_{\mathrm{lon}} = (n\text{-}1)(-25.890410959^\circ)$, to the longitude of T_{k1} (modulo 360˚).

Note that the actual locations of the ascending nodes of the reference orbit are not evenly spaced due to the higher order terms in the Earth gravity field assumed when generating the reference orbit. However, assuming symmetry of the tile boundary points T across all orbits in the repeat cycle simplifies the tile boundary definition and allows a smaller set of points to be distributed for users. The deviation of the ascending nodes of the reference orbit from the ideal, symmetric case assumed by the tile definition is at most around 200 m along the equator, which is small compared to likely deviations of the true orbit from the reference orbit and small compared to the 2 km sampling of L2_LR_SSH fixed-grid data.

The tiles are numbered based on their pass number (1 to 584) in the repeat cycle, the along-track tile number in the pass (1 to 308), and whether the tile is on the left or right side of nadir.

Definition: Tiles are named with the following convention: *PPP_TTTC*, where *PPP* is the three-digit, zero-padded pass number (001 to 584), *TTT* is the three-digit, zero-padded along-track tile number (001 to 308), and *C* is the character 'L' or 'R' designating the left or right sides, respectively. The along-track tile number is defined chronologically.

For example, tiles 001_001L and 001_001R would be the first left-side and right-side tiles of the first pass of the cycle (recall from the definitions in Section 2.1 that “left” and “right” are defined here with respect to the spacecraft velocity vector, not the attitude, and are hence independent of yaw flips). As pass 1 is ascending (odd numbered), these tiles would be the southernmost tiles in the pass. The northernmost tiles in this pass would be 001_308L and 001_308R. The next tiles in the cycle chronologically would be the northernmost tiles in the first descending pass: 002_001L and 002_001R.

Note that numbering the tiles based on their latitude (e.g., sequentially from south to north) was also considered, but the chronological numbering is motivated by simplicity, noting that pass numbers, cycle numbers, etc. are already numbered chronologically. Note that sequentially numbered passes do not have sequential longitudes. Moreover, near the southern and northern extremes of the orbit, many tiles have nearly identical latitudes when the nadir track is oriented toward the east ($\phi_h \approx 90^\circ$). Because latitude-based numbering would still not give an easy, intuitive mapping between tile number and latitude given the variation in heading of the nadir track, simplicity of the tile numbering convention is perhaps a more important consideration.

5 L1B_HR_SLC, L2_HR_PIXC, and L2_HR_PIXCVec Product Sampling and Granule Boundaries

5.1 L1B_HR_SLC, L2_HR_PIXC, and L2_HR_PIXCVec Sampling

The actual boundaries of tiles are related to the sampling defined by the SLC processor, so the SLC sampling is described here. The sampling of the SLC product is defined on a deskewed grid related to the spacecraft position and velocity at the times of the KaRIn presummed HR data. Let **r** be the spacecraft position vector and **v** be the spacecraft Earth-relative velocity vector at the time *t* corresponding to a given HR presummed echo (the distinction between the transmit time and the receive time is important for the SLC processor but not central to this description). Let **u** be the local geodetic up vector at **r**. Let the broadside plane be defined as the plane containing both **u** and **u** x **v** (the cross symbol denotes the vector cross product). The broadside plane is therefore normal to the horizontal component of **v**.

***Definition*:** SLC samples are defined in rows on a deskewed grid such that the samples in each row have reference locations that are uniformly spaced in slant range and are located at intersections of the broadside plane and the surface of a reference digital elevation model (DEM).

Note that the vertical component of **v**, while much smaller than the horizontal component, is generally not zero, so the reference locations of SLC samples are not processed at truly zero Doppler even under ideal attitude conditions. Under nonideal attitude (particularly pitch) conditions, SLC samples may be illuminated at times and corresponding spacecraft positions differing appreciably from *t* and **r**. In such cases, the samples in a deskewed row may also be illuminated at different times and spacecraft positions than each other. (This is accommodated in the SLC processor.)

Note that **r** and **v** here are the estimates of position and velocity from the ephemeris information provided to the SLC processor. These do not follow the reference orbit perfectly. The reference locations therefore generally depend on whether the SLC processor is run with the POE or the MOE product.

Reference locations of samples in the L2_HR_PIXC product are derived from the reference locations of samples in the L1B_HR_SLC product. The reference locations of samples in the L2_HR_PIXC product are generally aligned with those of the L1B_HR_SLC product, but they do not coincide exactly due to the various steps in the PIXC processor. However, the actual geolocations of samples based on the SWOT measurements after height reconstruction do not generally coincide with these reference locations, though the primary differences are in the height and cross-track directions.

Samples in the L2_HR_PIXCVec product correspond exactly to samples in the corresponding L2_HR_PIXC product, although their geolocations differ based on the algorithm processing involved to generate them. That is, there is a one-to-one mapping between samples in the L2_HR_PIXC product and samples in the L2_HR_PIXVec product.

5.2 Actual L1B_HR_SLC, L2_HR_PIXC, and L2_HR_PIXCVec Tile Boundaries

Deskewed rows of samples in the L1B_HR_SLC, L2_HR_PIXC, and L2_HR_PIXCVec products are not split across different tiles, so the actual boundaries of their tiles are defined to approximate the reference tile boundaries described above. This is accomplished by defining the actual tile boundaries around whether the spacecraft position **r** that is used to set the sampling grid falls within the reference tile boundaries defined above. L1B_HR_SLC tiles contain additional padding at the tile along-track ends, which introduces tile overlap, in order to simplify the PIXC processing.

Definition: Let the reference along-track boundaries of L2_HR_PIXC tile k be defined by points T_{k-1} and T_k and the corresponding planes $X_{YSk-1} = 0$ and $X_{YSk} = 0$. The actual along-track boundaries of tile k of the PIXC product are defined such that a row of deskewed samples is included in tile k if the spacecraft position **r** that defined the row of samples is located between the reference tile boundary planes through the points T_{k-1} and T_k given by the equations $X_{YSk-1} = 0$ and $X_{YSk} = 0$.

The reference along-track boundaries of a given L1B_HR_SLC tile are defined similarly to the corresponding L2_HR_PIXC along-track tile boundaries, except that the planes defining the L1B_HR_SLC tile boundaries are shifted by P_{xSLC}, where nominally $P_{xSLC} = 4$ km, in order to provide the desired overlap of $\pm P_{xSLC}$ in the L1B_HR_SLC tiles.

Definition: The actual along-track boundaries of tile k of the L1B_HR_SLC product are defined such that a row of deskewed samples is included in tile k if the spacecraft position **r** that defined the row of samples is located between the boundary planes given by the equations $X_{YSk-1} = -P_{xSLC}$ and $X_{YSk} = +P_{xSLC}$ where X_{YSk-1} and X_{YSk} are specified in the local coordinate systems defined by the points T_{k-1} and T_k.

The KaRIn data from the left and right swaths are acquired at different times with different polarizations, and stored as separate data channels, so they are already naturally split. This split forms the basis of splitting the left and right tiles at a given along-track location, effectively dividing the left and right tiles by the true nadir track (not the reference nadir track).

Definition: The actual cross-track boundaries of an L1B_HR_SLC, L2_HR_PIXC, or L2_HR_PIXCVec tile are defined such that all KaRIn data corresponding to the left half-swath are contained in the left tile, and all data corresponding to the right half-swath are contained in the right tile.

Note that the actual cross-track boundaries of L1B_HR_SLC, L2_HR_PIXC, and L2_HR_PIXCVec tiles could vary from the cross-track boundaries of the reference tiles above by a few kilometers due to cross-track orbit variations.

Differences between the actual along-track L2_HR_PIXC and L2_HR_PIXCVec tile boundaries and the reference tile boundaries are expected to be small (considerably less than 1 km).

The actual L1B_HR_SLC, L2_HR_PIXC, and L2_HR_PIXCVec tile boundaries may change slightly between POE and MOE processing.

Because the samples in the L2_HR_PIXCVec product correspond exactly to samples in the corresponding L2_HR_PIXC product, their tile boundaries coincide exactly as well.

6 L2_HR_RiverSP and L2_HR_LakeSP Product Sampling and Granule Boundaries

6.1 L2_HR_RiverSP and L2_HR_LakeSP Sampling

Data elements in the L2_HR_RiverSP and L2_HR_LakeSP products are vector objects defined by the associated a priori databases and/or polygons of observed water bodies. Multiple positions representing different quantities (e.g., a priori knowledge of location vs. measured location) are associated with each object.

6.2 River and Lake SP Granule Boundaries

The L2_HR_RiverSP and L2_HR_LakeSP products are provided in full-swath, continent-pass granules. Because the SP products simply contain collections of vector objects, the granule boundaries are defined by the geographic extent of the set of vector objects that have the same continent identifier as that of the SP product and are observed on the pass number of the SP product.

Note that pixels in the L2_HR_PIXC product may be attributed to river or lake vector objects in a given SP granule even if the pixels are located outside the pass or continent boundaries corresponding to the SP granule if the water body spans the pass or continent boundary. (The continents are defined by the science team with respect to geographic features such that this scenario is very unlikely, but it is not strictly forbidden by the definitions.) Similarly, pixels that are located outside of a continent polygon may be attributed to a river or lake vector object that is inside the polygon, as the pixel attribution is performed without regard to the continent boundaries.

Both L2_HR_RiverSP and L2_HR_LakeSP products use the same continent boundary definitions, which are defined by the top-level basins of the public HydroBASINS database (https://hydrosheds.org/page/hydrobasins). Continents are defined by non-overlapping, geographically fixed polygons. Each river reach or lake from the a priori databases is associated with exactly one continent. The reference location of the river or lake from the database is used for this purpose rather than the SWOT-measured location so that the association of a river or lake with a continent does not depend on the SWOT measurement. The association of prior objects to continents can also therefore be computed offline.

Observed lakes and unassigned features in the L2_HR_LakeSP product are associated with continents depending on the largest area of intersection of the polygon defining their boundary.

Note that while continents can be defined relative to expected HR downlink mask characteristics, the definitions in this document are intended to be general enough that the algorithm processing behavior is fully specified even in the event that an arbitrary HR downlink mask is used. For example, the definitions in this document and the description of continent and basin polygons globally should specify the desired processing behavior and product characteristics for HR data collected over the ocean for Cal/Val or other purposes.

The cross-track extent of the L2_HR_RiverSP product is defined such that any reach whose centroid is within 80 km of the reference nadir track is always included in the product, although it may be null filled if too few detected water pixels assigned to it. The 80-km threshold is significantly wider than the 64-km reference cross-track extent of L2_HR_PIXC tiles for the following reasons:

- A reach whose centroid is 80 km from nadir may have nodes that are closer to nadir by half the reach length.
- The actual nadir track deviates from the reference nadir track by around 2 km.
- While 64 km can be considered a nominal cross-track width, the actual SWOT observation swath is defined by range, not cross-track extent, and it is possible for there to be useful pixels farther out in cross track than 64 km.
- The threshold is set conservatively large with the rationale that it is better to include null-filled entries in the product that indicate that no observation was made than to fail to provide data that could have been useful.

The cross-track extent of the L2_HR_LakeSP product files is defined similarly to that of the L2_HR_RiverSP product, but there are also some important differences:

- L2_HR_LakeSP_Prior files contain all lakes that are present in the Prior Lake Database (PLD) and whose centroid is within 80 km of the reference nadir track. Therefore, reservoirs present in a L2_HR_RiverSP granule will likely also be present in the corresponding L2_HR_LakeSP_Prior file, although slight differences right at the granule edge are possible. Such features are unlikely to be observed by SWOT anyway, however. Unobserved PLD lakes within the granule yield void objects (no polygon, and fill values for the variables).
- L2_HR_LakeSP_Obs files provide observed water surfaces linked to PLD lakes. The same set of PLD lakes is initially considered (centroid within 80 km of the reference nadir track), but only actually observed water features are included (there are no void objects for unobserved PLD lakes).
- L2_HR_LakeSP_Unassigned files represent observed water bodies in the corresponding L2_HR_PIXC products that have not been linked to PLD lakes, nor previously to river reaches in the Prior River Database (PRD). Its cross-track extent is therefore not linked to PLD lakes and PRD reaches whose centroid is within 0-80 km.

However, there is a second set of cross-track extent configuration parameters with respect to the actual orbit that is used to truncate the detected water from the input L2_HR_PIXC product (for example to 7-65 km) in order to reduce the risk of artifacts in lake boundaries and other associated inaccuracies. This is motivated by the decreasing observability of water surfaces towards the swath edges (and very sparse pixels near-nadir) and the need to distinguish whether a lake (or unassigned feature) should be considered fully or partially observed. These thresholds on the ground-projected cross-track distance have direct impact on the polygons and values of the variables of the L2_HR_LakeSP_Unassigned, L2_HR_LakeSP_Prior and L2_HR_LakeSP_Obs files but not on the set of which PLD lakes is included in the L2_HR_LakeSP_Prior files (as unobserved PLD lakes are kept as void objects).

7 L2_HR_RiverAvg and L2_HR_LakeAvg Product Sampling and Granule Boundaries

7.1 L2_HR_RiverAvg and L2_HR_LakeAvg Sampling

Data elements in the L2_HR_RiverAvg and L2_HR_LakeAvg products are vector objects that correspond exactly to the objects in the L2_HR_RiverSP product and to the objects in the Prior file of the L2_HR_LakeSP product. The average products have different granule spatial and temporal boundaries, collecting observations of the river and lake vector objects over multiple passes.

7.2 L2_HR_RiverAvg and L2_HR_LakeAvg Granule Boundaries

L2_HR_RiverAvg and L2_HR_LakeAvg products are provided in basin-cycle granules. Because the Avg products simply contain collections of vector objects, the granule boundaries are defined by the geographic extent of the set of vector objects that have the same basin identifier as that of the Avg product and are observed in the cycle number of the Avg product.

Note that as with L2_HR_RiverSP and L2_HR_LakeSP products, pixels from the L2_HR_PIXC product may be attributed to river or lake vector objects in a given L2_HR_RiverAvg or L2_HR_LakeAvg granule even if the pixels are located outside the basin boundaries corresponding to the Avg granule if the water body spans the basin boundary. (It is expected that basins are defined with respect to geographic features such that this scenario is very unlikely, but it is not strictly forbidden by the definitions.)

Both L2_HR_RiverSP and L2_HR_LakeSP products use the same basin boundary definitions, which are defined by the second-level basins of the public HydroBASINS database. Basins are defined by non-overlapping, geographically fixed polygons. No basin spans multiple continents. There are up to 9 basins per continent. Each river or lake from the a priori databases is associated with exactly one basin by determining which polygon, if any, contains the reference location for the river or lake. The reference location of the river or lake from the database is used for this purpose rather than the SWOT-measured location so that the association of a river or lake with a basin does not depend on the SWOT measurement. The association of objects to basins can also therefore be computed offline.

As with SP products, the definitions here are intended to be applicable for any possible HR mask, including collection of HR data over the ocean.

8 L2_HR_Raster Product Sampling and Granule Boundaries

8.1 L2_HR_Raster Sampling

The samples of the L2_HR_Raster product are geographically fixed on a uniform Universal Transverse Mercator (UTM) grid. The grid is aligned with the northing and easting dimensions of the UTM zone at the center of the granule. The raster product granules are swath-aligned scenes, so the UTM grid for each scene encompasses a greater area than the data coverage of the scene, with null-filled values in void areas that are not covered by the granule, as in the left-hand diagram of Figure 1.

8.2 L2_HR_Raster Scene Boundaries

L2_HR_Raster scene boundaries are defined so that scenes correspond to 2 x 2 sets of tiles with geographically fixed boundaries as illustrated in Figure 4.

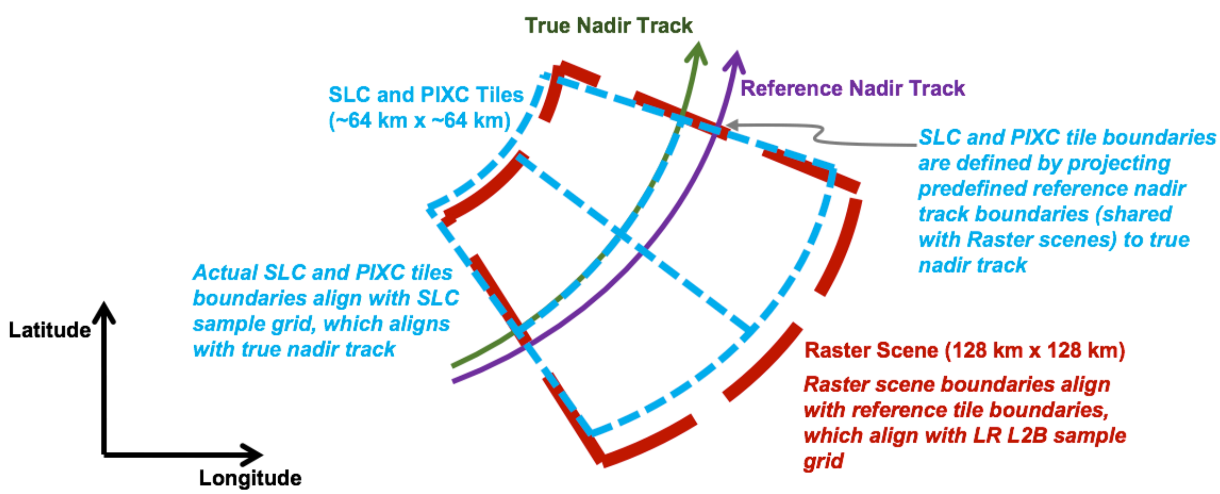

Figure 4. Illustration of the relationship between raster scene boundaries and SLC and PIXC tile boundaries.

Definition: The along-track boundaries of scene m coincide with the reference along-track boundaries defined by points T_{k-1} and T_{k+1} and the corresponding planes defined by the equations $X_{YSk-1} = 0$ and $X_{YSk+1} = 0$ for k odd.

Definition: The cross-track boundaries of a scene extend ±(W_S/2) from nadir. Let W_S = 128 km. The ±(W_S/2) cross-track distance is defined along a spherical approximation to the ellipsoid where the sphere is tangent to the ellipsoid at the nadir point defining the cross-track direction. The sphere radius R_a is 6378.137 km. The boundary is projected in the local up direction relative to the ellipsoid, not relative to the approximating sphere. See Figure 3.

Definition: Scenes are named with the following convention: *PPP_SSS*, where *PPP* is the three-digit, zero-padded pass number (1 to 584), and *SSS* is the three-digit, zero-padded along-track scene number (1 to 154). The scene number is defined chronologically.

An alternative that was considered for the scene numbering convention was to use only odd (or only even) numbers so that scene numbers are aligned with tile numbers within a pass. However, continuous, sequential numbering is simpler and provides greater consistency with other products such that it was more beneficial overall. Users can simply multiply the scene number by two in order to find a corresponding tile number if desired.

The first and last scenes of a pass are slightly larger than other scenes just as the tiles at the ends of the pass are slightly larger than other tiles.

Note that it is not essential that $W_S = 2T_S$. The nominal value W_S should be large enough that the scene encompasses all useful data over all actual nadir tracks with the possible exception of extreme outliers. The true, as-flown nadir tracks are expected to deviate from the reference nadir track by no more than a few kilometers. The value of W_S can be increased if the cross-track deviation of the true nadir track from the reference nadir track is large enough that the geographically fixed boundaries of the scene would miss useful data (for example if the KaRIn SNR is sufficient that a wider swath is expected).

For example, scene 001_001 would be the first scene of the first pass of the cycle. As pass 1 is ascending (odd numbered), this scene would be the southernmost scene in the pass. The northernmost scene in this pass would be 001_154. Scene 001_154 coincides with reference tiles 001_307L, 001_307R, 001_308L, and 001_308R. The next scene in the cycle chronologically would be the northernmost scene in the first descending pass: 002_001.

9 L2_HR_FPDEM Product Sampling and Granule Boundaries

9.1 L2_HR_FPDEM Sampling

The gridded file of the L2_HR_FPDEM product is sampled on a uniform geodetic latitude-longitude grid with a sample spacing of approximately 1 arcsec. As a single granule of the gridded FPDEM product is built (interpolated) from many KaRIn observations from different passes (including possibly ascending, descending, left-looking, and right-looking viewing geometries), the FPDEM sampling is unrelated to the native KaRIn sampling.

The ungridded FPDEM product is a point cloud of water pixels from multitemporal KaRIn observations as mentioned above (from L2_HR_PIXC/PIXCVec products), with each point having its own height-constrained geolocation (latitude, longitude, and height) and time-tag.

9.2 L2_HR_FPDEM Granule Boundaries

The FPDEM product is given in non-overlapping 1° × 1° latitude-longitude granules.

10 L1B_LR_INTF Product Sampling and Granule Boundaries

10.1 L1B_LR_INTF Native-Grid Sampling

The L1B_LR_INTF product is sampled at the native sampling of the LR KaRIn data, which is determined by the KaRIn pulse timing, the state of the on-board processor, and the spacecraft position, velocity, and attitude. Each sample is associated with a measurement time (nominally the KaRIn transmit time) as well as a location, which is based on a reference Earth surface. Samples for the left and right sides are defined separately.

10.2 L1B_LR_INTF Temporal Pass Boundaries

The L1B_LR_INTF product is provided in granules with full-swath coverage. Temporally defined overlap with respect to the temporal pass boundaries described in Section 3.4 is specified via a constant parameter P_{tLRL1B}. The nominal value of P_{tLRL1B} is 3.92 s, which gives approximately ±25 km overlap.

Definition: Let t_{k-1} and t_k be the starting and ending times of pass k as defined in Section 3.4. An L1B_LR_INTF data sample is placed into the granule for pass k if the measurement time t of the sample satisfies $t_{k-1} - P_{tLRL1B} < t \leq t_k + P_{tLRL1B}$.

11 L2_LR_SSH Product Native-Grid Sampling and Granule Boundaries

The L2_LR_SSH product contains four files (Basic, Expert, WindWave, and Unsmoothed). The Basic, Expert, and WindWave files provide data on a fixed grid, while the Unsmoothed file provides data on a native grid. This section applies only to the native-grid file.

11.1 L2_LR_SSH Native-Grid Sampling

Like the L1B_LR_INTF product, the native-grid file of the L2_LR_SSH product is sampled at the native sampling of the LR KaRIn data, although the L2_LR_SSH native-grid data use the KaRIn-measured heights rather than a reference surface for geolocation. There is no extra "padding," as the sampling is determined based on the samples downlinked by the KaRIn instrument. There are no samples in the gap near nadir between the left and right half swaths because these samples would always be null when considering the native sampling grid; separate sampling grids are defined for the left and right sides.

11.2 L2_LR_SSH Temporal Pass Boundaries

The native-grid file of the L2_LR_SSH product is provided in granules with full-swath coverage and temporal pass boundaries, with no overlap, as described in Section 3.4.

The cross-track extent of the L2_LR_SSH native-grid data is determined by the geolocations of the native KaRIn samples.

12 L2_LR_SSH Product Fixed-Grid Sampling and Granule Boundaries

The L2_LR_SSH product contains four files (Basic, Expert, WindWave, and Unsmoothed). The Basic, Expert, and WindWave files provide data on a fixed grid, while the Unsmoothed file provides data on a native grid. This section applies only to the fixed-grid files.

12.1 L2_LR_SSH Fixed-Grid Sampling

The sampling grid of L2_LR_SSH fixed-grid data is swath aligned and fixed geographically such that data from a given pass number over different repeat cycles can be directly compared without resampling. Data from different passes must be resampled for direct comparisons, however. The along-track locations of the L2_LR_SSH fixed sampling grid are derived from the same set of nadir-track locations, spaced at 125 m along the ellipsoid and produced by the CNES Mission System, that are used to define reference tile boundaries. The cross-track sample locations are then tied to the along-track sample locations at nadir and defined on the ellipsoid such that their projections onto an approximating sphere are evenly spaced (see Figure 3).

***Definition*:** Let L2_LR_SSH fixed-grid samples along the reference nadir track be defined at a set of points B_k for each pass, where k is the along-track index of the sample. Two points B are located ±1 km from the equator along the nadir track over the ellipsoid, and successive points are located every 2 km thereafter along the nadir track over the ellipsoid until the geographic ends of the pass (where $\phi_h = 90^\circ$) in either direction from the equator are reached.

***Definition*:** Let L2_LR_SSH fixed-grid samples in the cross-track direction be defined for each sample B_k, which is located on the nadir track, such that cross-track samples are located at the intersection of the ellipsoid and the plane defined by the equation $X_{YSk} = 0$, where X_{YSk} is the along-track component of the TCRF frame at B_k. The samples are arranged such that one sample is at the reference nadir location B_k, and the projections of other samples in the local up-down direction (at the samples) are evenly spaced every 2 km along a sphere of radius R_a that is tangent to the ellipsoid at B_k. The approximating sphere radius R_a is 6378.137 km. Samples extend a distance of $\pm(W_B/2)$ from nadir along the approximating sphere in cross track. Let $W_B = 140$ km. This value of W_B is chosen to allow a conservative amount of orbit, topography, and/or parameter variation that might shift the locations of samples.

As described above, the TCRF frame is uniquely specified by the location of B_k and the Earth-relative heading ϕ_{hk} at that point.

An alternative description of the cross-track sampling can be stated as follows. Cross-track sample locations can be computed by first defining a temporary set of points that are evenly spaced 2 km apart on the approximating sphere. These temporary points are not in general on the ellipsoid, except along the nadir track. Computing the latitudes, longitudes, and heights (relative to the ellipsoid) of these points, then keeping only the latitudes and longitudes (discarding the height by forcing it to zero), has the effect of projecting the temporary locations along the local up direction onto the ellipsoid. These new points represent the fixed-grid sample locations.

At the ends of passes, there is necessarily a discontinuity in the along-track sample spacing because the length of a pass is not an integer multiple of the sample spacing. While pass-to-pass discontinuities could be avoided by defining continuous nadir locations for the entire repeat cycle, there would still be discontinuities for each repeat cycle, and a much larger set of nadir locations would need to be distributed to users, so the simple, symmetric approach described here has been adopted.

The definition of cross-track sample locations using a spherical approximation to the ellipsoid avoids the need for computing exact distances along the ellipsoid, which is computationally difficult and typically subject to approximations, which could result in accumulated error when integrating the distances over long distances. The spherical approximation is chosen so that the approximation involved is simple and well defined and can be computed simply and consistently.

A definitive set of points B_k for the first orbit of the repeat cycle is derived from data provided by the CNES Mission System distributed to users. Reference nadir points for subsequent orbits are computed by applying a longitude shift to the points for the first orbit, similar to the manner in which reference tile boundaries and scene boundaries are defined.

***Definition*:** A given point B_{kn} defining the kth nadir-track sample location for orbit n, with $n > 1$, is identical to the corresponding point B_{k1} for the kth nadir-track sample location of the first orbit ($n = 1$) after adding a constant shift $\Delta\Theta_{\text{lon}}$, given by $\Delta\Theta_{\text{lon}} = (n\text{-}1)(\text{-}25.890410959^{\circ})$, to the longitude of B_{k1} (modulo 360˚).

As described in the context of tile boundaries, the assumption of ideal, evenly spaced ascending nodes simplifies the L2_LR_SSH fixed-grid sampling definition. The variations in the ascending node spacing of the reference orbit (less than 200 m along the equator) are small compared to the expected orbit control error and the L2_LR_SSH fixed-grid sample spacing.

Because of Earth rotation, the nadir track curves in the cross-track direction. The along-track spacing of samples is therefore only constant where it is defined to be constant along the nadir track itself. The along-track spacing of samples at far range varies on the order of ±2 m over the orbit.

The spherical approximation used to define cross-track samples introduces very little horizontal deviation of the sample grid from 2 km (less than 1 cm of integrated error at the far edges of the swath).

L2_LR_SSH fixed-grid samples are indexed from the far-range edge of the left side to the far-range edge of the right side.

12.2 L2_LR_SSH Fixed-Grid Geographic Pass Boundaries

While the pass boundaries for temporally linked data products such as the L1B_LR_INTF product and the native-grid file of the L2_LR_SSH product are derived from the ORF as described above, it is more natural for the fixed-grid files of the L2_LR_SSH product to follow pass boundaries that are geographically defined. This is because it would be convenient for users to be able to rely on a given geographically fixed, swath-aligned sample to always fall within a particular pass. Moreover, the use of symmetric sampling grids between different passes would simplify the sampling grids for user analysis and interpretation, but it would result in discontinuities in the sampling grids at pass boundaries since the length of a pass is not an integer multiple of the 2 km sampling spacing. It is therefore important to tie the fixed-grid granule boundaries to the discontinuities in the sampling grids.

Overlap between fixed-grid granules would help minimize the inconvenience of these sampling discontinuities by extending the regular sampling of each pass partly into the spatial extent of the previous and next passes. That is, fixed-grid granules overlap in their spatial coverage, but the granules that overlap may each be sampled on their own uniform grids, so the samples do not align between the two granules in the overlap region (see

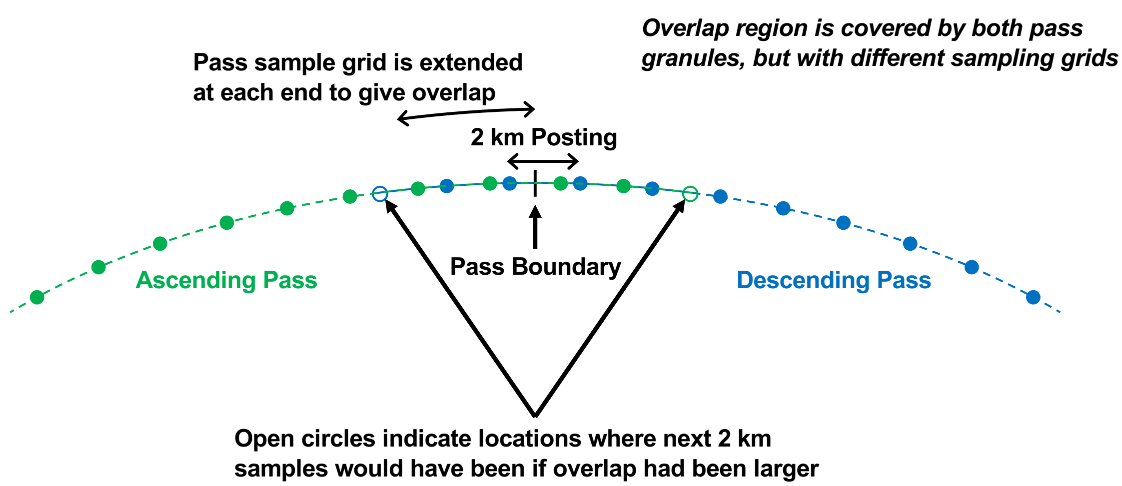

Figure 5). A user can then use the granule whose sampling grid is most convenient for evaluating data near the pass boundary, especially when the user performs his/her own interpolation near the pass boundary. Let the amount of overlap be given by a parameter P_{xLR2f}, where the nominal value of P_{xLR2f} is 5 km.

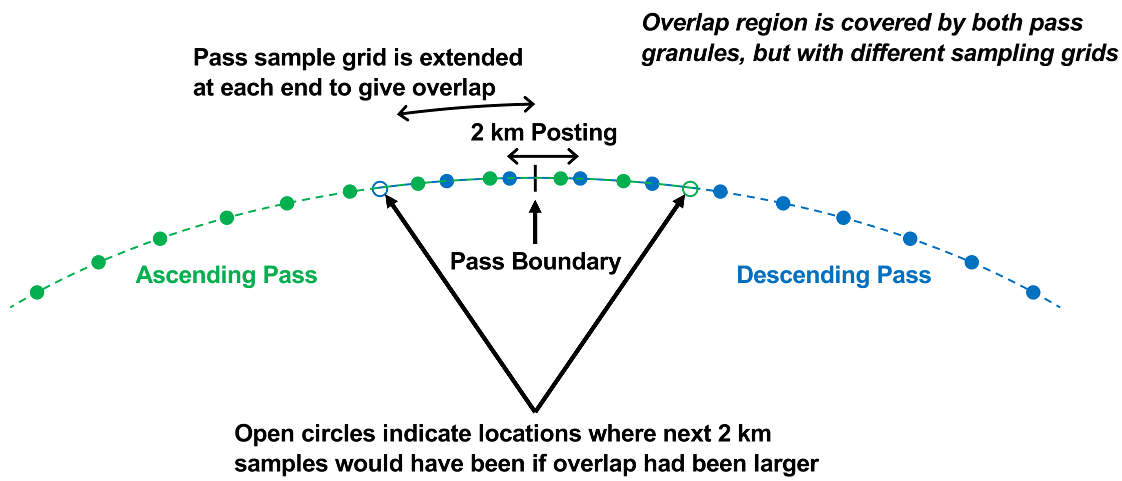

Figure 5. Illustration of the overlap between granules of L2_LR_SSH fixed-grid files, which have intrinsically different sampling grids from pass to pass.

Definition: Pass boundaries for fixed-grid L2_LR_SSH product files are defined by the extent of the nadir points B_k used to define the sampling grid for each pass. The first and last points B_k in each pass are chosen such that they extend the regular nadir sampling of the pass as far as possible but no farther than P_{xLR2f} from the geographic pass ends defined by the reference tile along-track boundaries.

The geographic L2_LR_SSH fixed-grid pass boundaries are hence independent of the ORF and can be determined a priori.

Note that the L2_LR_SSH fixed-grid pass boundaries are also therefore the subject to the simplification that passes beyond the first orbit are shifted in longitude by a constant angle $\Delta\Theta_{\mathrm{lon}}$, as described above.

13 Nadir Altimeter Product Sampling and Granule Boundaries

13.1 Nadir Altimeter Sampling

The nadir altimeter products (L2_NALT_GDR, L2_NALT_IGDR, and L2_NALT_OGDR) are sampled at the native sampling of the altimeter, which is determined by the altimeter timing and the spacecraft position and attitude. Each sample is associated with a measurement time as well as a location computed from the antenna pointing.

13.2 Nadir Altimeter Product Granule Boundaries

The L2_NALT_GDR and L2_NALT_IGDR products are provided in granules with temporal pass boundaries as described in Section 3.4. The L2_NALT_IGDR product is provided in granules that span the measurement times provided in each instrument telemetry download.

14 Radiometer Product Sampling and Granule Boundaries

14.1 Radiometer Sampling

The radiometer products (L2_RAD_GDR, L2_RAD_IGDR, and L2_RAD_OGDR) are sampled at the native sampling of the dual radiometers, which is determined by the radiometer timing and the spacecraft position and attitude. Each sample is associated with a measurement time as well as a location computed from the antenna pointing.

14.2 Radiometer Product Granule Boundaries

The L2_RAD_GDR and L2_RAD_IGDR products are provided in granules that include both left- and right-looking antennas with temporal pass boundaries as described in Section 3.4. The L2_RAD_OGDR product is provided in granules that span the measurement times provided in each instrument telemetry download.

15 MOE, POE and ATTD_RECONST Sampling and Granule Boundaries

15.1 Ephemeris and Attitude Sampling

The MOE and POE products provide orbit ephemeris state vectors at intervals of 10 seconds along the satellite orbit. The ATTD_RECONST product provides spacecraft orientation as quaternions with a sampling interval of 15.625 ms (i.e., 64 Hz).

15.2 Ephemeris and Attitude Product Granule Boundaries

The MOE, POE, and ATTD_RECONST products are provided as daily files, spanning 26 hours and centered at 12:00:00 (TAI) of each day.

Appendix A. **Acronyms**

AD	Applicable Document
CNES	Centre National d'Études Spatiales
HR	High Resolution
JPL	Jet Propulsion Laboratory
LR	Low Resolution
NASA	National Aeronautics and Space Administration
RD	Reference Document
SDS	Science Data System
SWOT	Surface Water Ocean Topography
TBC	To Be Confirmed
TBD	To Be Determined