ESA Climate Change Initiative (CCI)

Greenland Ice Sheet (GIS) Essential Climate Variable (ECV)

System Specification Document (SSD)

Prime & Science Lead: Louise Sandberg Sørensen
DTU Space, Copenhagen, Denmark

Technical Officer: Anna Maria Trofaier
ESA ECSAT, Didcot, United Kingdom

Consortium:
Asiaq Greenland Survey (ASIAQ)
DTU-Space, Department of Geodynamics (DTU-S)
DTU-Space, Department of Microwaves and Remote Sensing (DTU-N)
Danish Meteorological Institute (DMI)
ENVironmental Earth Observation IT GmbH (ENVEO)
S&T Norway AS (S&T)
Technische Universität Dresden (TUDr)
The Geological Survey of Denmark and Greenland (GEUS)
The Niels Bohr Institute (NBI)
Shepherd Space (SS)

Signatures page

Prepared by	Louise Sandberg Sørensen Science Lead, DTU-S	May 22, 2026
Issued by	Daniele Fantin, Project Manager, S&T	Daniele Fantin Digitally signed by Daniele Fantin Date: 2026.05.22 10:02:02 CEST
Approved by	Anna Maria Trofaier ESA Technical Officer	Anna Maria Trofaier Digitally signed by Anna Maria Trofaier Date: 2026.06.04 10:11:14 +01'00'

Table of Contents

Change Log

Issue	Author	Affected Section	Reason	Status
1.0	S&T	All	Draft version	
1.1	GEUS	5	Updated MFID	
1.2	ENVEO, S&T	3	Updated IV	Reseased to ESA on 6 July 2021
2.0	All	All	Update for CCI+ Phase 2	Released to ESA on 1 December 2023
3.0	All	All	Update for CCI+ Phase 3	Released to ESA

Acronyms and Abbreviations

Acronyms	Explanation
ATBD	Algorithm Theoretical Basis Document
C3S	Copernicus Climate Change Service
CCI	Climate Change Initiative
CS2	CryoSat-2
CSR	Center for Space Research, University of Austin
DEM	Digital Elevation Model
(D)InSAR	(Differential) Interferometric Synthetic Aperture Radar
DMI	Danish Meteorological Institute
DTU-N	DTU Microwaves and Remote Sensing Division
DTU-S	DTU Geodesy and Earth Observation division
ETAD	Extended Timing Annotation Dataset
ECV	Essential Climate Variable
ENVEO	ENVironmental Earth Observation IT GmbH
EO	Earth Observation
ESA	European Space Agency
GCOS	Global Climate Observation System
GCP	Ground Control Point
GEUS	Geological Survey of Denmark and Greenland
GFZ	Deutsche GeoForschungsZentrum
GIA	Glacial Isostatic Adjustment
GIS	Greenland Ice Sheet
GLL	Grounding Line Location
GMB	Gravimetry Mass Balance
GRACE(-FO)	The Gravity Recovery and Climate Experiment (Follow On)
InSAR	Interferometric Synthetic Aperture Radar
IPP	Interferometric Post-Processing
IV	Ice Velocity
JPL	NASA Jet Propulsion Laboratory
LOS	Line of sight
MFID	Mass Flux and Ice Discharge
NBI	Niels Bohr Institute, University of Copenhagen
NEGIS	North East Greenland Ice Stream
OT	Offset tracking
PROMICE	Danish Program for Monitoring of the Greenland Ice Sheet

RA	Radar Altimetry
RMS	Root Mean Square
SS	Shepherd Space
S1	Sentinel-1
S2	Sentinel-2
SAR	Synthetic Aperture Radar
SDS	Science and Data System
SEC	Surface Elevation Change
SLC	Single look complex
SMB	Surface Mass Balance
SOW	Statement of Work
TOA	Top of Atmosphere
TPROP	Technical Proposal
TU Dresden	Technische Universität Dresden
URD	User Requirement Document
TOPS	Terrain Observation by Progressive Scans
OT	Offset Tracking

1 Introduction

1.1 Purpose and Scope

This document contains the System Specification for the Greenland_Ice_Sheet_cci (GIS_cci) project for CCI+ Phase 3, in accordance with the contract and SoW [AD1 and AD2].

The purpose of the document is to specify the characteristics of the ESA GIS_cci ECV processing systems.

This SSD document builds on the Phase 2 SSD [RD1] of the 'Greenland_Ice_Sheet' CCI+ project.

1.2 Document Structure

This document is structured into an introductory chapter followed by five chapters describing the processing system for the CCI+ ECV parameters, which are:

- Surface Elevation Change (SEC and dSEC)
- Ice Velocity (IV) and SAR Line-of-Sight Velocity
- Gravimetric Mass Balance (GMB)
- Mass Flux and Ice Discharge (MFID)

Each ECV chapter provides:

- A specification of the purpose of an operational ECV production system and its intended use.
- An overview of the context of the system, defining all significant interfaces among system components and crossing the system's boundaries.
- A definition of the fundamental operations to be performed within the system to accept and process the inputs and to process and generate the outputs.
- A description of the major constraints of the system.
- A description of operational scenarios for the system including data sources, valid ranges of values, timing considerations, operator requirements, and special interfaces.
- Specification of the environmental characteristics of where the system will be installed.
- Specification of the growth, expansion, and capability characteristics, as well as limitations of the system.

1.3 Applicable and Reference Documents

Table 1-1: List of Applicable Documents

No	Doc. Id	Doc. Title	Date	Issue/ Revision/ Version
AD1	ESA/Contract No. 4000126523/19/I-NB - Greenland_Ice-Sheets_CCI+ and its Appendix 1 (incl CCN3)	CCI+ Phase 1 New R&D pm CCI ECVs for Greenland _Ice Sheet_cci (incl CCN3)	Cont: 2019.03.06 CCN3: 2022.12.05	-
AD2	ESA-EOP-SC-AMT-2021-53	Climate Change Initiative Extension (CCI+) Phase 2 - New R&D on CCI Essential Climate Variables -SoW (incl Annexes)	2022.06.10	Issue 1 Revision 2
AD3	DTU-ESA-GIS-CCI-P3-PROP-001	CCN5 FOR PHASE 3 (2025-2026) Technical, Management,Implementation and Financial Proposal (PROP)	2025.05.30	2.0
AD4	ST-DTU-ESA-GISCCI+-ATBD-001	Algorithm Theoretical Basis Document	2024.11.30	2.1

Note: If not provided, the reference applies to the latest released Issue/Revision/Version

2 Surface Elevation Change (SEC)

2.1 System overview

The SEC operational ECV production system takes as input all available elevation data from the Level-2 products of radar altimetry missions since 1991 and processes these to produce 5-year (and for CS2 2-year) mean rate of surface elevation change products over the Greenland ice sheets. Input data for the SEC production system consists of multi-mission radar altimetry data (Table 2.1), auxiliary correction data, ice sheet DEM, and ice mask. The SEC system is described in detail in the current ATBD [AD-4].

Table 2-1: Current Radar Altimetry Data Sets

Mission	Product Level	Baseline	Source	Operational Dates
CCI+ GIS				
ERS-1	L2	FDR4ALT Land-Ice v1.0	ESA	02-Aug-91 to 02-Jun-96
ERS-2	L2	FDR4ALT Land-Ice v1.0	ESA	13-May-95 to 04-Jul-03
ENVISAT	L2	FDR4ALT Land-Ice v1.0	ESA	09-Apr-02 to 18-Oct-10
CryoSat-2	L2i SIN	Baseline-E	ESA	16-Jul-10 to present
CryoSat-2	L2i LRM	Baseline-E	ESA	16-Jul-10 to present
Sentinel-3B	L2	Baseline-005	ESA	01-Mar-16 to present
Sentinel-3B	L2	Baseline-005	ESA	01-May-18 to present

2.2 Operational scenarios

The SEC processor works in four distinct modes: True repeat track, along track, cross-over and Plane-fit. The choice of mode depends on the data used as input. We refer here to the ATBD for further information.

Table 2-2: User-defined input for the SEC system

User input	Format	Note
START TIME	ddmmyyyy	Start date of the chosen SEC period
END DATE	ddmmyyyy	End date of the chosen SEC period
GRID SIZE	X [km]	Grid cell size of final SEC grid
CORR. LENGTH	Y [km]	Correlation length of dh/dt data. This controls the smoothness of SEC grid
MIN #DATA	Z [number]	Number of data points required in each grid cell/segment.
DEM	Name of DEM	Choice of DEM to use in the relocation
Model	#	Defines which model to solve for in the least squares regression
Flags	binaries	1 for use this parameter. 0 do not use this parameter. Number of relevant flags depends on mission data.

The dSEC processor estimates gridded surface elevation change using a spatio-temporal state-space model, combining spatial correlations and temporal dynamics. The model is fitted by maximum likelihood, producing continuous dSEC fields and associated uncertainties.

Table 2-3: User-defined input for the dSEC system

User input	Format	Note
Satellite	Name	Satellite mission used (e.g. CryoSat-2)
Grid Size	X [km]	Grid resolution of the output
Start time	mmyyyy	Start date of the analysis period
End time	mmyyyy	End date of the analysis period

2.3 Hardware and software platform

The SEC operational ECV production system runs on most UNIX servers with adequate memory and disk storage space. The current SEC processing system is developed and tested in Mac-OS, and is transferred for operational runs to a server system running ax86_64 GNU/Linux UNIX operating system.

2.3.1.1 Hardware

Table 2-3: Processing Hardware at for the SEC processor

Model	Supermicro 848A-R1K62B
Processor	4 X AMD Opteron 6376 (64 cores), 64 bit
Memory (RAM)	256 GB
Local Hard Drive	30TB
Network Attached Storage	>100TB

2.3.1.2 Operating system

The system currently runs on the UNIX operating system x86_64 GNU/Linux. The system has been tested and runs on other Unix/Linux systems as well, including Mac OS X.

2.3.1.3 Tools and libraries

The SEC processor is written in Python. Python 3.5 packages: numpy, tables, scipy, matplotlib, os, sys, multiprocessing, shapefile, netCDF4, rasterio. The dSEC processor uses an R-module named RTMB.

2.4 Future concerns and developments

We are currently working on porting all the processing of the SEC solutions to run on the HPC.

2.5 References

Sørensen, L.S., et al. (2015). Envisat-derived elevation changes of the Greenland Ice Sheet, and a comparison with ICESat results in the accumulation area. Remote Sensing of Environment, v160, p56-62. DOI: 10.1016/j.rse.2014.12.022

Simonsen, S.B., and Sørensen, L.S., (2017). Implications of changing scattering properties on Greenland Ice Sheet volume change from CryoSat-2 altimetry. Remote Sensing of Environment, v190, p 207-216.

Sørensen, L.S., et al. (2018). 25 years of elevation changes of the Greenland Ice Sheet from ERS, Envisat and CryoSat-2 radar altimetry. Earth and Planetary Science Letters, vol 495, p 234-241. DOI: 10.1016/j.epsl.2018.05.015

Andersen, Natalia H., Sebastian B. Simonsen, Karina Nielsen, Mai Winstrup, Baptiste Vandecrux, Hui Gao, Beata Csatho, Anton Schenk, and Louise Sandberg Sørensen. "A State-Space Model for Monitoring Greenland Ice Sheet Surface Elevation Change from CryoSat-2." *EGUsphere* 2025 (2025): 1-25.

3 Ice Velocity

3.1 SAR Ice Velocity (SAR IV)

3.1.1 System overview

In this section, we describe the high-level processing system for deriving ice velocity maps from repeat-pass SAR data, applying both offset tracking and SAR interferometry (InSAR) techniques. The ENVEO SAR Software Package (ESP) Version 2.1 is used for the generation of ice velocity maps. ESP 2.1 applies incoherent and coherent offset tracking for mapping displacements in range and azimuth on repeat-pass SAR SLC data. It also measures line-of-sight displacements with InSAR and combines these to derive ice velocity from crossing-orbit measurements. The IV processor supports all common space-borne SAR missions, including Sentinel-1, Radarsat-2, ERS 1/2, ENVISAT, ALOS, TerraSAR-X and SAOCOM. Auxiliary data needed are a DEM and (optionally) an ocean/ice/land mask. An external a priori ice velocity map is also used for co-registration and calibration purposes. Figure 3-1 shows the high-level processing line for the IV production. It includes the following three modules:

- MODULE IV: Within this module, SAR data and the Extended Timing Annotation Dataset (ETAD), as well as further auxiliary files, are imported into the system and velocity maps are generated for pairs of repeat-pass SAR data of the same track. This module applies r to both offset tracking and InSAR processing. The outputs are IV maps in map projection as defined by the user. Each IV product includes a quality map.
- MODULE MERGE: this module combines the IV products from different tracks and image pairs. The outputs are temporally and spatially merged large/regional-scale ice velocity maps, including uncertainty estimates.
- MODULE VAL: the validation module is an associated module, which enables the inter-comparison and validation of the generated ice velocity map with in-situ velocity measurements or external ice velocity maps. The output is statistical information on the inter-comparison and scatterplots.
- MODULE IV-DATABASE: This module connects the velocity data archive to an online portal/toolbox for easy access, data visualisation, and further analysis, including the derivation of velocity profiles and calculation of ice discharge along pre-defined gates.

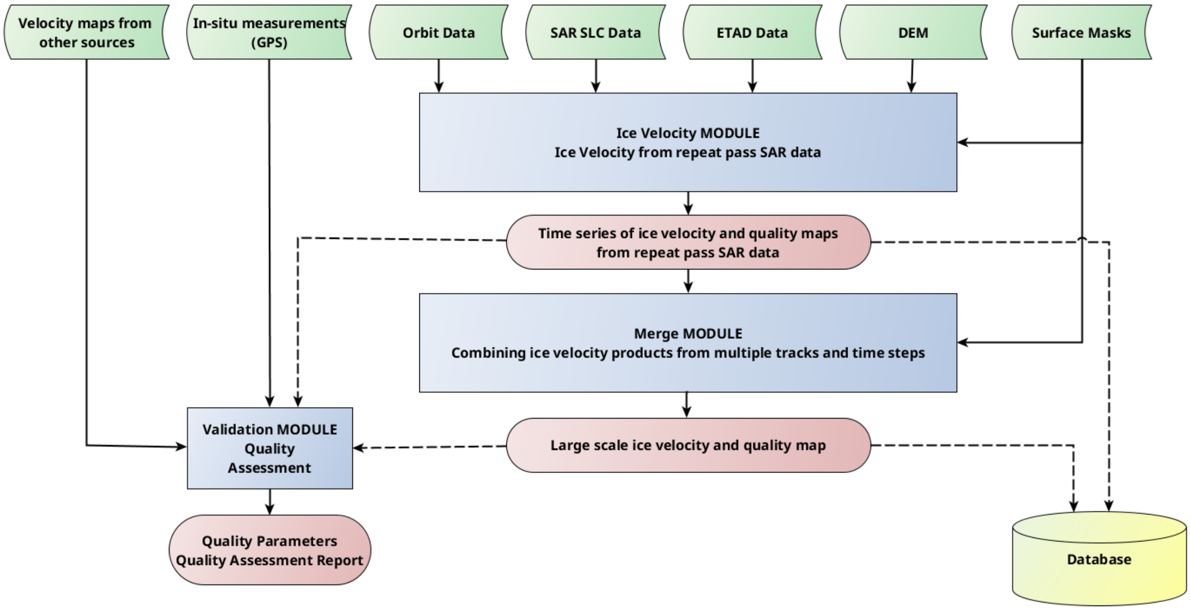

Figure 3-1: High-level flow chart of the IV processing system. Green – input data, Blue – processing modules, red - product and intermediate products.

MODULE IV

Offset-tracking

Figure 3-2 shows a high-level flow line of the ice velocity generation module with offset tracking. For Sentinel-1 Interferometric Wide (IW) swath single look complex (SLC) data, the processing is done on the burst level. The module has access to the SAR data archive, ETAD data and orbit data, as well as the DEM and optionally the surface masks. The processing is done track by track, for each track, image pairs are selected according to the time step. Using image geometry and a digital elevation model, the local shift between the image pair is calculated, which is considered in the displacement calculation using incoherent (or coherent) offset tracking. Besides the displacements in the slant range and azimuth, the quality of the matching is calculated. After outlier removal, the velocity of the 3 components is calculated using a DEM as input. Optionally, the surface masks are applied to mask the ocean and unglaciated areas.

The output of this module is velocity maps (E, N, Z) in m/d per track and time step, and the quality map.

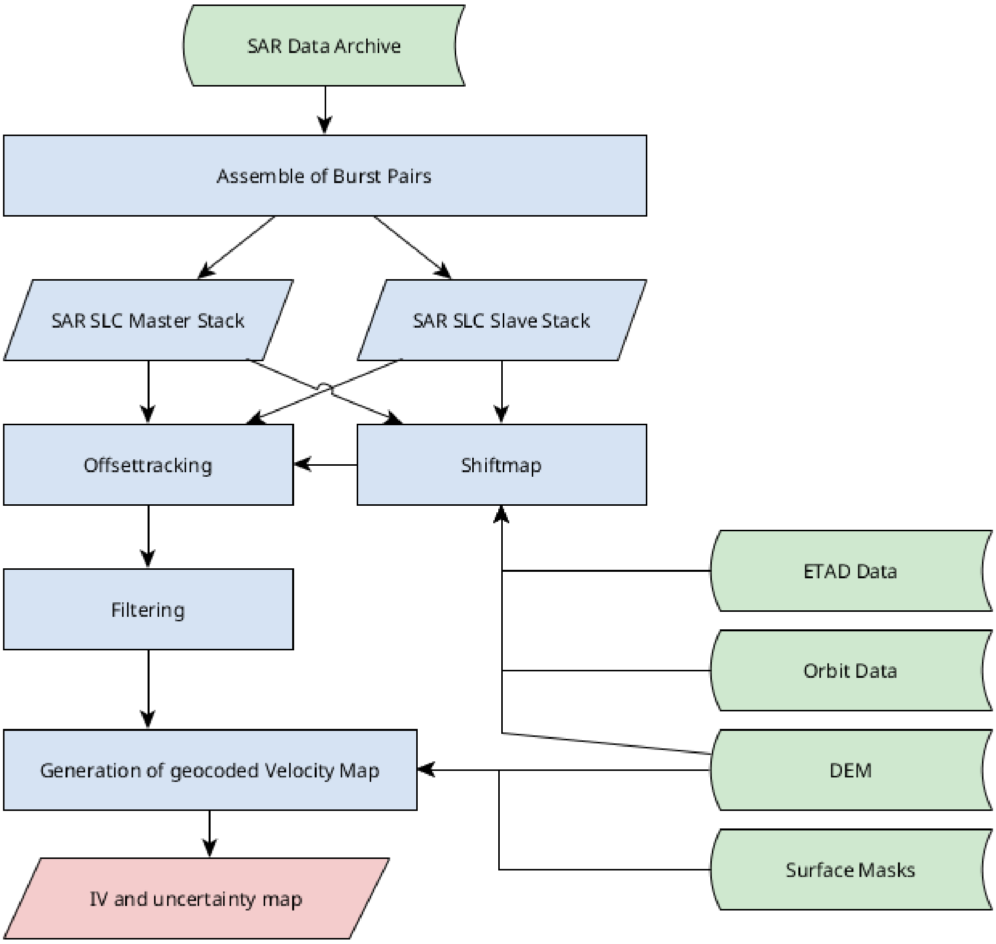

Figure 3-2: Processing steps for ice velocity generation using offset-tracking.

InSAR

Figure 3-3 describes the flowline of InSAR processing for ice velocity. This processing line is dedicated to Sentinel-1 acquisitions in TOPS mode. Similar to the offset-tracking module, the InSAR processing on Sentinel-1 data is performed at the burst level up to the phase unwrapping step. Burst-level processing includes co-registration, interferogram formation, flat-earth and topographic phase removal and coherence estimation. If working with Stripmap acquisitions, as in the case of TerraSAR-X or SAOCOM, these processing steps are directly applied to the whole scene and the processing line is adapted by simply skipping the debursting step.

Burst interferograms are mosaicked at the deburstring step to make a single large-scale interferogram for each master-slave pair. Phase unwrapping and phase calibration are performed at the debursted level. Once calibrated, the phase can be converted into line-of-sight (LOS) velocity and is then geocoded.

The InSAR module uses the SAR data archive in SLC format as input. For Sentinel-1 data, only data with a 6-day temporal baseline are used to avoid aliasing and decorrelation caused by fast ice motion or changing surface conditions due to melt or snowfall. The InSAR module uses precise orbits provided by ESA ~20 days after the acquisition, as well as the ETAD data. A digital elevation model must also be given as input to the module for co-registration, estimating the topographic phase and geocoding.

Optionally, an external a priori ice velocity map (from e.g., offset tracking) can be inserted in the module. It is then used for the co-registration and calibration steps: at the co-registration step, it is used to update the co-registration map with local displacements; at the calibration step, it is used for selecting slow-motion ground control points, which are used as reference values.

The output of this module is geocoded maps of LOS velocity. These maps are produced for 6-day and 12-day repeat pass data from all available tracks. In the MERGE module, corresponding to the inversion step, these images are used to derive a velocity map. The inversion step combines LOS velocities from crossing orbits to calculate the ice flow velocity components. The merged output is a 3-D velocity map (X, Y, Z) in meters per day.

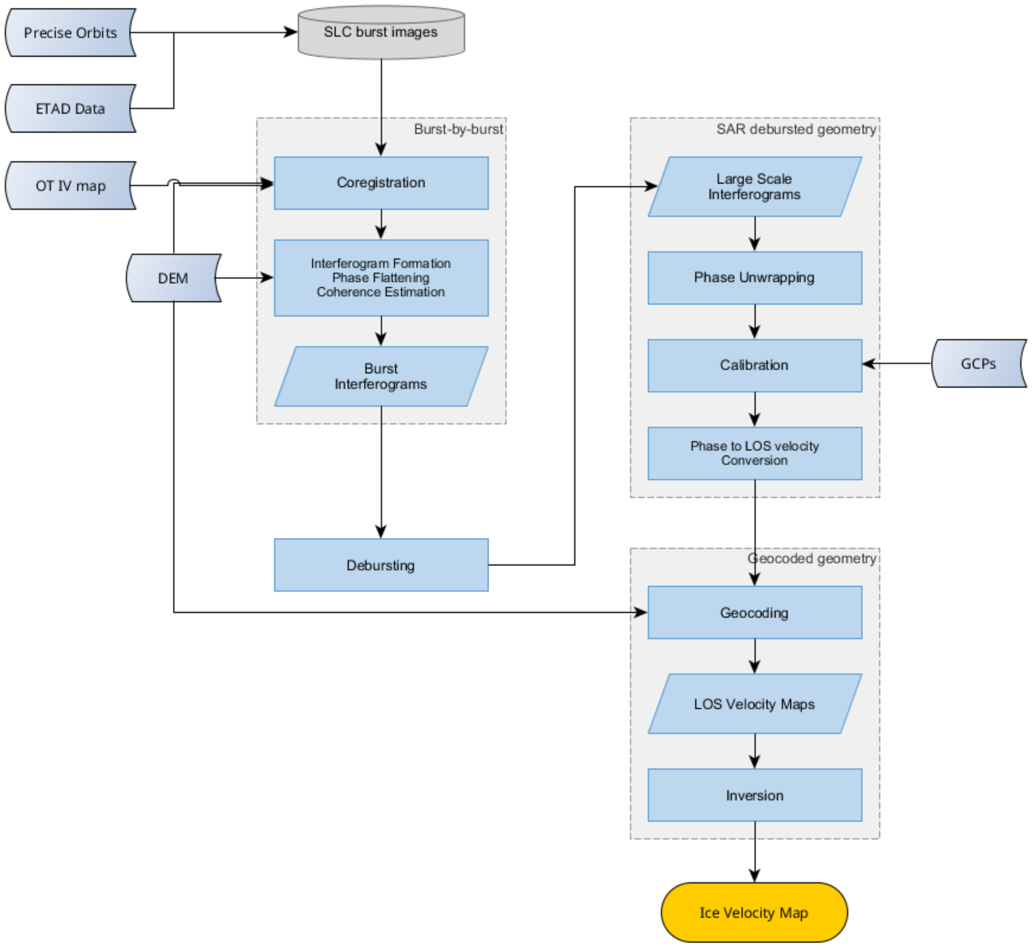

Figure 3-3: Processing steps for ice velocity generation using InSAR applied to Sentinel-1 IW TOPS SLC data. SLC – Single-Look Complex. OT – Offset Tracking. LOS – Line-of-sight. GCP – Ground Control Point.

MODULE MERGE

This module aims to produce a large/regional-scale ice velocity map using the output of MODULE IV as input. The generation of a regional ice velocity map requires the combination of IV results from several tracks. The long tracks are geocoded into a common map projection of the output grid, using the annotated imaging and orbit parameters and a digital elevation model as input. The merging combines the velocity components separately in m/d. The output of this module is a large-scale, regional or ice sheet-wide flow velocity and quality map.

MODULE VAL

The validation module aims for an independent validation of the ice velocity product using in-situ ice velocity data, as provided by GPS. The module inter-compares the velocity components separately but can also be applied to the velocity magnitude. Currently, the module requires manual interaction, mainly in the selection and pre-processing of the in-situ ice velocity data. Outputs are statistical parameters of the inter-comparison and scatter plots for visualisation. This module also allows us to inter-compare ice velocity products from other sources. The module takes different map projections into account and provides statistical parameters of the pixel-by-pixel inter-comparison, histograms, and spatial maps of the differences.

MODULE IV-DATABASE

The ice velocity products are linked to the IV database at cryoportal.enveo.at. The web-based portal allows easy access and simple analysis of all ice velocity products (single maps, monthly & annual). The current versions support the following analysis:

- Ice Velocity (IV) visualisation and download of data:
    - Visualise the full-time series of ice velocity profiles along pre-defined central flow lines of all major outlet glaciers.
    - Visualise the full-time series of velocity at fixed points along the central flowline.
- Mass Flux Ice Discharge (MFID) calculation:
    - For single glaciers
    - Sub basins
    - Basins
    - Ice sheet wide

3.1.2 Operational scenarios

1. Continuous observations of margins and the ice sheet interior with offset-tracking.

The track-by-track offset-tracking processing is applied in regions where continuous acquisitions of S1 IW SAR data are acquired. This processing scenario enables the monitoring of ice velocity variations with short time intervals. This observation scenario has been proposed by ENVEO and implemented by ESA for the Greenland Ice Sheet margins with continuous acquisitions since June 2015.

2. Production of annual (within C3S) and monthly ice sheet-wide ice velocity maps

Ice sheet-wide velocity maps for Greenland are generated on a monthly and annual basis and rely on the Sentinel-1 acquisition campaigns. They are performed in addition to the continuous acquisitions of S1 at the ice sheet margins.

3. Production of (winter campaign) ice velocity map of the ice sheet interior with crossing orbit InSAR

A similar observation scenario is planned for track-by-track InSAR processing. InSAR processing focuses on the ice sheet interior, where ice motion is slower, and conditions are stable from one acquisition to another. Data in wintertime are expected to have better coherence, and the coverage of overlapping ascending and descending tracks over the ice sheet interior is improved during the winter campaign. For this purpose, Sentinel-1C and 1D data with 6- and 12-day time intervals will be used. Additionally, from January-April 2026, 1-day repeat data is acquired using Sentinel-1C and Sentinel-1D. L-Band InSAR covers areas decorrelating in C-band, in particular in regions with strong velocity gradients due to reduced fringe frequency, enabling reliable phase unwrapping.

3.1.3 Hardware and Software Platform

The main processing at ENVEO is performed on the processing system described below.

3.1.3.1 Hardware

The main (Sentinel-1) ice velocity processing is done on 3 server machines and 18 virtual machines on a cluster (IKB cluster), which are connected to a mass storage of about 800 TB. The system applies OPENMP to support multiple CPUs and Cores. For development, visualisation, and quality control, an additional 8 workstations are used.

Table 3-1: Processing Hardware for the ENVEO SAR IV processor.

	Processing	**Development, Visualisation, Quality Control**
Model	GNU/Rocky Linux 9	GNU/Linux Fedora
Number of machines	5	8
Processor	Intel Xeon E5-2643 v4 @ 3.4 GHz AMD EPYC 7343 @ 4.0 GHz	AMD Ryzen 7
Memory (RAM)	128 GB	32 GB
Local Hard Drive	500 GB	1 TB
Network	Ethernet 10000baseT/Full	Ethernet 1000baseT/Full
Network Attached Storage	Ca 800 TB network storage	Ca 800 TB network storage

	Processing	
Model	GNU/Rocky Linux 9	
Number of nodes	18	
Processor	Intel Xeon CPU-X5670 2.93GHz 2 cores	
Memory (RAM)	10 GB	
Local Hard Drive	-	
Network	Ethernet 10000baseT/Full	
Network Attached Storage	Ca 800 TB network storage	

3.1.3.2 Operating system

The ENVEO ESP Ice Velocity processing system runs on common Linux operating systems. Currently, the tested systems are Rocky Linux 9, Fedora 43 and later releases, but the software will also work on other common Linux/Unix systems.

3.1.3.3 Tools and libraries

Tools & Libraries

GCC / OpenMP 11.5	http://gcc.gnu.org
cmake	http://cmake.org
PROJ.4	https://trac.osgeo.org/proj/
Python 3.9 (numpy, scipy, etc)	https://www.python.org/
GDAL (v 3.4)	http://www.gdal.org/
FFTW 3	http://www.fftw.org/
Libxml	http://www.xmlsoft.org/
NetCDF, HDF	http://www.unidata.ucar.edu/

Visualisation Tools

QGIS	http://qgis.org
Cryoportal Ice Velocity (developed by ENVEO)	https://cryoportal.enveo.at

3.1.4 Future concerns and developments

Further processor updates are ongoing to accommodate the ever-expanding data coverage and increasing temporal resolution. The ESP software has been successfully connected to and tested on cluster systems utilising several hundred cores. This capability is particularly valuable for large-scale campaign processing of big datasets, such as those encountered in polar ice sheet studies.

3.2 Optical Ice Velocity (Opt IV)

3.2.1 System Overview

In this section, we describe the processing system for deriving ice velocity maps from pairs of Sentinel-2 optical images by applying an offset tracking technique. The optical ice velocity processor developed by S&T consists of 6 modules, which are listed in Table 3-2. To run these modules, the user needs to configure the parameters from a master config file then run one script that flexibly strings together the IV modules, manages the data streams, and handles the instantiation of the processing steps. The high-level processing line for the IV production is shown in Figure 3-5.

Table 3-2: Name and purpose of each module composing the optical IV processor

Module name	Purpose
Product download	Queries and downloads S2 products within a given AOI and time range.
IV_pair	Finds all available pairs of S2 products between 7 and 15 days apart within the queried time range.
IV	Creates an IV map from a given S2 pair.
IV_merge	Creates a single stacked or averaged IV product of all the IV maps. Mosaicking tiles if needed.
IV_release	Prepares the final IV product in a zip file containing product description, quick look etc.

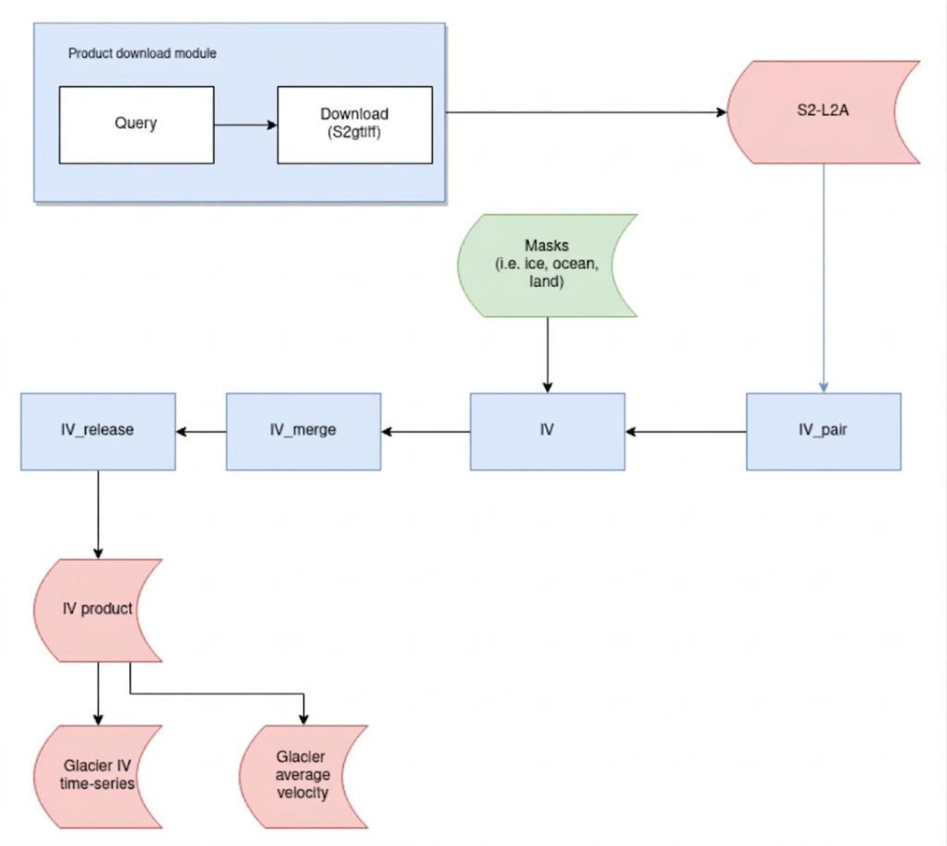

Figure 3-5: high-level flow chart of the optical IV processor. Green identifies input data, blue processing modules, grey specific processing tasks within a module, and red are intermediate and final products.

3.2.2 Product download module

The product download module takes as required input an AOI and a time range, from which it queries Sentinel-2 L2A images and downloads them to a local folder. Additional filtering criteria, such as cloud coverage, can be specified. The S2 product will be converted to the GeoTIFF format.

IV pair

All possible pairs of the retrieved S2 images that are at a minimum of seven and a maximum of 15 days apart will be found. These pairs will then be passed on.

IV

For each pair of S2 images, we will apply a feature tracking algorithm to calculate the north and east ice velocity. Furthermore, the co-registration offset will be measured, and the root-mean-squared (rms) will be extracted from the ice velocities.

IV Merge

The IV maps generated in the previous step will be averaged, stacked and mosaiced together in the case where the AOI covers multiple tiles. Both a time series and an averaged IV map will be created.

IV Release

The final output will consist of the IV time series and the averaged IV map. These will be delivered in a zip file, and NetCDF products together with a quick look image, comments, and product description as the CCI requirements have specified.

3.2.3 Operational Scenarios

The operation scenario relies on the availability of Copernicus Sentinel-2 A/B satellite acquisitions. Firstly, the observations for 9 pre-selected glaciers are retrieved based on temporal and spatial criteria through a network interface (e.g., SciHub or AWS) and stored locally. The production of the seasonal IV map is planned for the end of the summer season, after the end of the S2 satellite's observational season over Greenland.

Once the observational season is over, glacier-based IV maps are generated according to two processing scenarios:

- Production of IV time-series using observations with a 6-12-day repeat.
- Production of seasonal IV map.

The scenario has been proposed by S&T and approved by ESA.

3.2.4 Hardware and software platform

3.2.4.1 Hardware

Model	Custom-built
Processor	Intel Core i7 CPU @ 3.40GHz
Memory (RAM)	32 GB
Local Hard Drive	2 TB
Network Attached Storage	1x30TB

Figure 3-6: Processing Hardware for the IV Processor

3.2.4.2 Operating system

The processor operates on any Linux system.

3.2.4.3 Tools and Libraries

The Opt-IV processor requires Python 3.8 or later to run support scripts.

Linux Libraries required: gdal.

Python Libraries required(in addition to the Python standard library): click, h5py, loguru, netCDF4, numba, numpy, opencv, osr, pyproj, pytest, scipy, yaml, zipping.

3.3 SAR Uplift rate

The SAR uplift rate product is derived from an intermediate product generated in the InSAR IV processing chain at DTU. This processing chain is very similar to that of ENVEO, but is implemented in the DTU IPP (Interferometric Post Processor) processor. The uplift product is a time-series product that allows the study of small-scale velocity variations at a high temporal and spatial resolution. Whereas the spatial resolution is achieved by employing InSAR, the temporal resolution is achieved because each uplift map is derived from a single InSAR pair.

3.3.1 System Overview

3.3.1.1 Data selection and download

Due to the high temporal and spatial resolution, the uplift rate product is produced only for selected areas of interest (AOI) where dynamic phenomena can be expected, but in the future may be produced for larger regions in order to discover and track dynamic events on a larger scale. For a selected region of interest, IW SLCs covering the AOI are searched for in the Copernicus Dataspace Ecosystem (CDSE). Only SLCs from tracks which are acquired continuously (i.e. every 6 or 12 days) are considered. With the current Greenland acquisition scenario, most areas are continuously imaged. Once the data has been downloaded, each track is processed separately.

3.3.1.2 Processing Setup

To begin with, a reference SLC from the track is selected, and all bursts covering the AOI are identified. A reference (i.e. master) geometry is then defined based on these bursts and the master SLC timing and state vectors. Then, all possible 6- or 12-day pairs are identified, and each pair is processed individually, as described in the following section.

3.3.1.3 Interferogram generation and geocoding

The processing flow for a single pair is illustrated in. First, the two SLCs are coregistered to the reference geometry using SLC timing, precise state vectors, a digital elevation model (DEM) based on TanDEM-X, and a multiyear averaged (2016-2019) IV mosaic based on PROMICE ice velocity maps. The DEM is applied both to coregister and flatten the interferogram (thus removing the topographic phase component), whereas the IV is used only for coregistration, preserving the interferometric phase due to displacement. To reduce phase discontinuities at burst boundaries, burst overlap spectral diversity is used to refine the coregistration in the overlap region. The interferograms are multi-looked, and phase unwrapped using a Minimum Cost Flow algorithm, calibrated using a set of pre-selected slow-moving ground control points (GCPs), and converted to LoS velocity. The LoS velocity map is subsequently geocoded to a 100 m x 100 m spaced grid in EPSG 3413 projection (Polar stereographic, with a reference longitude of -45 and a latitude of the true scale of 70), along with ancillary data (Line-of-sight vector azimuth and elevation angles, displacement error estimate), which are output to individual files. The same output grid is used for all pairs so that the data can be easily stacked in a time series.

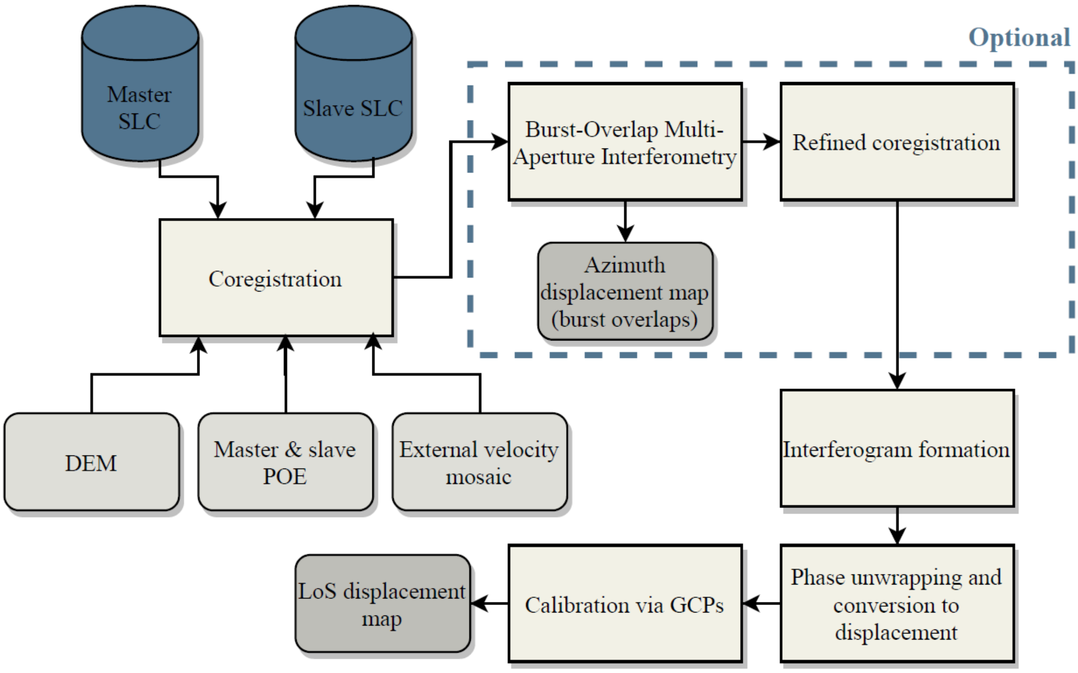

Figure 3-7: Processing flow to generate LoS-velocity for a single SLC pair

3.3.1.4 Time series generation

The uplift product is intended mainly to study small-scale, rapidly varying vertical displacements, which are generally of much smaller magnitude than the displacements due to the overall horizontal ice flow. The latter are, on the other hand, slowly varying in time, and are typically spatially correlated on a larger scale. To separate the two contributions, the pixel-wise temporal median of all the line-of-sight velocity maps is calculated, resulting in a single velocity map, which is then subtracted from all velocity maps in the time series. The residual LoS-velocities are then attributed to vertical displacement and converted to uplift rate by dividing by the sine of the elevation angle from ground to sensor. This allows the data to be analyzed directly, and the original line-of-sight velocity maps can be reconstructed multiplying by the sine of the elevation angle and adding the median velocity map. Both line-of-sight elevation and azimuth angles, as well as the median uplift rate, are provided with the data. This also allows for 3D-inversion if additional line-of-sight velocity maps from crossing tracks become available.

3.3.1.5 Packaging

The uplift rate maps are packaged as three-dimensional (time, x, y) variables in a NetCDF file following the CF Metadata conventions. In addition, the median LoS velocity map, the LoS azimuth angle map, and the LoS elevation angle map are provided as two-dimensional (x,y) variables in the same file, along with all metadata about the projection and dimension variables, such that the data can be read using common GIS and geoprocessing tools.

3.3.2 Operational scenario

The Uplift Rate product is a prototype product and will be generated for a few limited areas of interest.

3.3.3 Hardware and software platform

3.3.3.1 Hardware

The processing of individual pairs is run in parallel on the DTU Central HPC LSF cluster (https://www.hpc.dtu.dk/?page_id=2520), which comprises several different processing nodes, assigned by the HPC job scheduler. Time series generation and packaging are carried out on a local workstation (HP Z440 with 6-core Xeon E5-1650v4 3.6 GHz CPU, 64 GB memory, and 1TB disk)

3.3.3.2 Operating System

The HPC cluster runs Scientific Linux 7.9, whereas the local workstation runs Ubuntu 20.4.

3.3.3.3 Tools and libraries

Required Linux packages:

- **GCC / OpenMP**
- **libxml2**
- **lapack**
- **hdf5**
- **fftw3**
- **Python 3.X**

Required Python 3 libraries:

- **scipy**
- **numpy**
- **limxml**
- **dateutil**
- **netcdf4**
- **pil**
- **gdal**

4 Gravimetric Mass Balance (GMB)

4.1 DTU Space

4.1.1 System overview

The GMB system produces mass change time series of the Greenland Ice sheet, as well as for pre-defined drainage basins. Also, 5-year trend (water equivalent) maps are produced.

The system algorithms and methods are described in the [ATBD], while the system input and output data, including auxiliary data, are described in detail in the [PSD].

The overall system overview is shown in Figure 4.1.

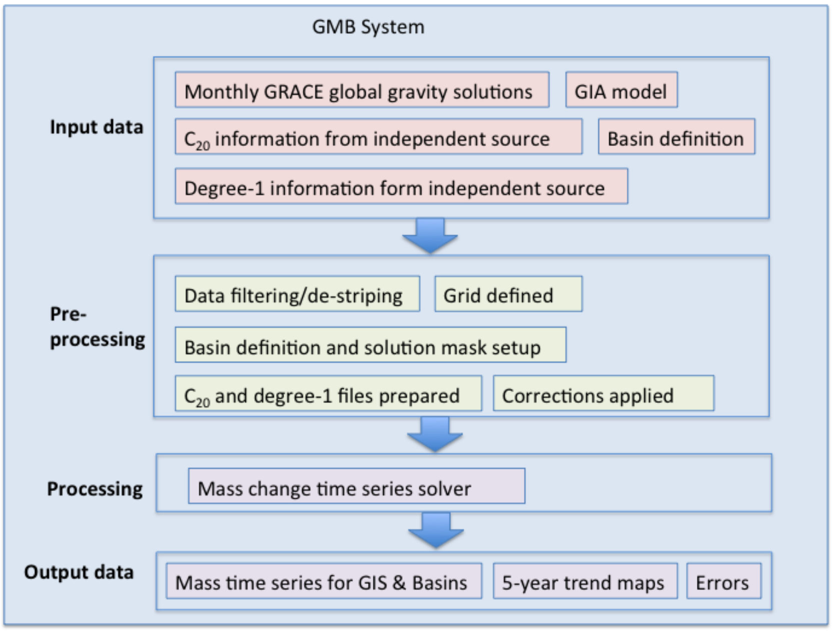

Figure 4.1: Overview of the GMB system at DTU Space. Everywhere in the picture where we say we process C_{20}, we now process both C_{20} and C_{30}.

The GIS CCI+ GMB system is a semi-automatic processing system implemented in Fortran and gnuplot. **Figure 4.1** gives a schematic overview of the GMB processing line, which consists of two main modules: the pre-processing and processing modules. Input data download is done semi-automatically. In the following, the two modules of the GMB processor are briefly described.

Pre-processing: Based on the selections in the Methods Module, GRACE monthly solutions and time series of auxiliary datasets are pre-processed, combined and corrected. This includes the conversion from Stokes-coefficients to mass equivalent coefficients (equivalent water height), the addition of degree-one coefficients, the replacement of degree C_{20} and C_{30} and the reduction of GIA. All monthly solutions are reduced to a specified reference value.

Processing (Mass change estimation): Mass change time series for every grid cell of the GMB gridded product and every basin of the GMB basin product is produced by integrating the product of the point mass inversion (Barletta et al., 2013) of each monthly solution. The statistical error characterisation of the time series is derived. Basin-averaged time series are used to calculate a mass balance estimate for each basin.

4.1.2 Operational scenarios

Data procurement:

Shortly after the availability of a new monthly solution is announced via the GRACE SDS newsletter, all necessary input data are semi-automatically collected via https with scripts.

The solution for the GRACE mission generated by the University of Texas Centre for Space Research (UTCSR) is provided via https://isdc-data.gfz.de/grace/Level-2/CSR/RL06/ . The solution for the GRACE-FO mission generated by UTCSR is provided via https://isdc-data.gfz.de/grace-fo/Level-2/CSR/RL06.3/ . The utilised ASCII format is called GRCOF2, and it is used for the harmonic coefficients in the official GRACE GSM products format.

The data volume depends on the maximum spherical harmonic degree lmax (for CSR: lmax=60, 96). The size of a single monthly solution is only about 0.7MB. The resulting volume of all GRACE data, comprising GRACE solutions from different processing centres for the period 2002 until the present and the associated background models, is clearly below 2GB.

User-defined parameters:

User input	Format	Note
DIST	Y [km]	Distance from solution area where signal is set to zero
LAMBDA	X [number]	Smoothing parameter used in the inversion
DEG 1	Filename	Filename of the chosen degree 1 file
C20 C30 CORR	Filename	Filename of the chosen C20-C30_file
DATA SOURCE	Name	Name of data source e.g CSR, GFZ, ITGS
GIA	Name of GIA	Choice of GIA model to use.

System log

For each run of the system, a log file is created to ensure transparency for the user on what input parameters were chosen for the given run.

This log contains information on the parameters described in the table above.

4.1.3 Hardware and Software Platform

The GMB operational ECV production system will run on most UNIX servers with adequate memory and disk storage space. The current GMB processing system is operated and tested on a server system running a x86_64 GNU/Linux UNIX operating system.

4.1.3.1 Hardware

Table 4.1: Processing Hardware for the GMB processor

Model	Lenovo ThinkSystem SD630 V2
Processor	2x Intel Xeon Gold 6342 (24 core, 2.8 GHz)
Memory (RAM)	512 GB
Local Hard Drive	480 GB SSD
Network Attached Storage	100TB

4.1.3.2 Operating system

The system currently runs on the UNIX operating systemx86_64 GNU/Linux.

The system has been tested and run on other Unix/Linux systems as well, including Mac OS X and a Windows machine with Linux on a Virtual Machine.

4.1.3.3 Tools and libraries

The system requires:

i) gfortran

ii) gnuplot

iii) NetCDF v4 library

4.1.4 Future concerns and developments

The procedure is semiautomatic, but the need to keep track of the publication of new data releases and of changes in the technical reports requires a careful human assessment.

4.2 TU Dresden

4.2.1 System overview

The GMB system produces mass change time series of the Greenland Ice sheet, as well as for pre-defined drainage basins.

The system algorithms and methods are described in the [ATBD], while the system input and output data, including auxiliary data, are described in detail in the [PSD].

The overall system overview is shown in Figure 4.2. it is a semi-automatic processing system implemented in MATLAB. It consists of four main modules, as described below. The control flow between the modules is handled by an operator. Input data procurement is done automatically.

Adaption of mass change estimation methods (Methods Module): In this module, the methods used to estimate mass changes are adapted to the characteristics of the latest input data (GRACE solutions and auxiliary datasets) in order to minimise both leakage errors and signal noise. Based on a quality assessment of different GRACE releases, performed in the spherical harmonic and spatial domains, the release to be used in the final product generation is selected. Moreover, the handling of low-degree spherical harmonics, GIA model reduction and additional model corrections (e.g. residual oceanic signals) are specified. Finally, the sensitivity kernels implied by the specific mass estimation method are derived [Döhne et al. 2023].

The Methods Module comprises the most labour-intensive tasks in GMB processing line. Comprehensive tests are required to verify the effectiveness of methodological amendments and modifications of the setup, resulting from the input data characteristics. These tests demand the extensive support of an operator.

Pre-processing: Based on the selections in the Methods Module, GRACE monthly solutions and time series of auxiliary datasets are pre-processed, combined and corrected. This includes the conversion from Stokes-coefficients to mass equivalent coefficients (equivalent water height), the addition of degree-one coefficients, replacement of degree C_{20}, and C_{30} if applicable and the correction for GIA. All monthly solutions are reduced to a specified reference value.

Mass change estimation (Estimation Module): The sensitivity kernels derived in the Methods Module and the time series provided by the Pre-processor serve as main inputs to this module. Mass change time series for every grid cell and every basin is produced by integrating the product of each monthly solution and the corresponding sensitivity kernel. For the basin product, the statistical error characterisation of the time series is derived. To infer an overall accuracy measure, GRACE errors propagated to the mass balance estimate are complemented by leakage errors and errors in GIA models. For this purpose, the mass change estimation methods are applied to the model outputs.

Product validation (Validation Module): The Validation Module summarises all validation procedures focusing on the GMB basin product.

In the case of a successful validation, the GMB products are written in the final output formats. ASCII and NetCDF output files are written using MATLAB. After passing a final check, the files are manually uploaded to the web server hosting the data portal.

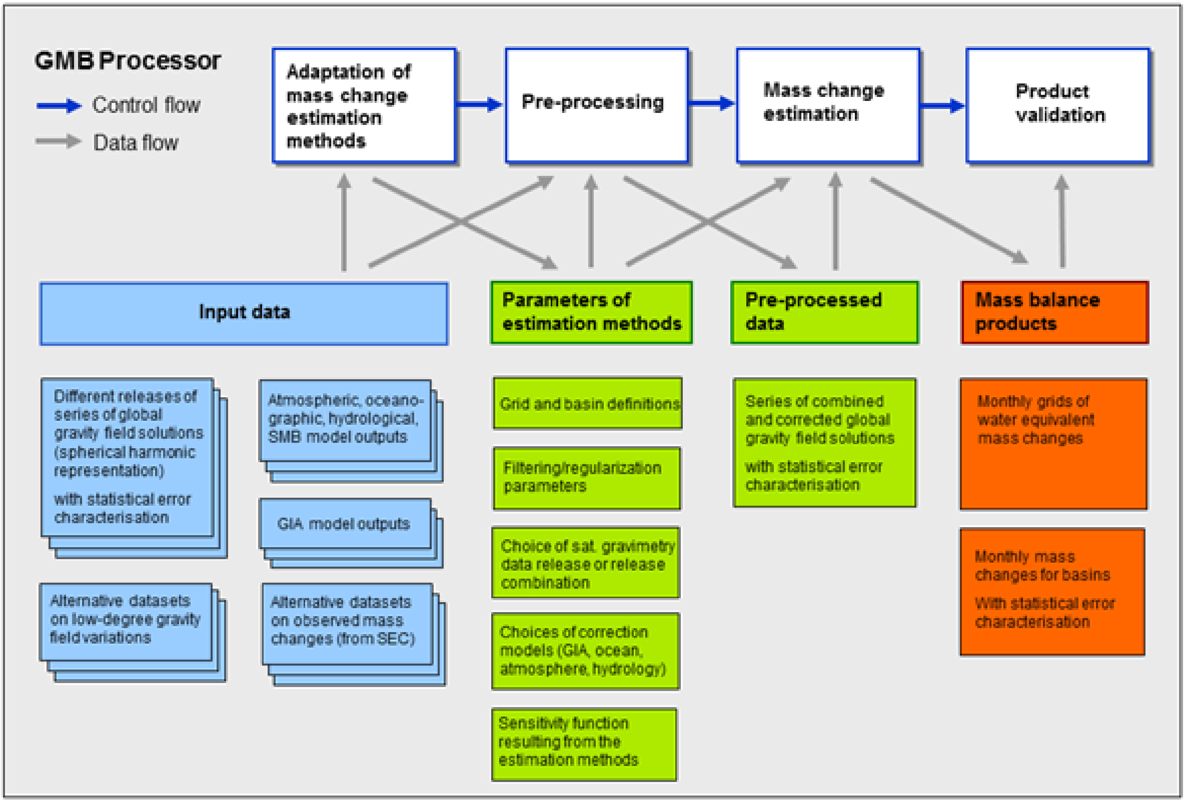

Figure 4.2: Overview of the GMB system at TU Dresden.

4.2.2 Operational scenarios

Data procurement: All regularly updated input datasets, i.e. GRACE monthly solutions and time series for both degree one and C_{20}, are automatically downloaded by means of a shell script executed by a weekly cron job.

4.2.3 Hardware and software platform

4.2.3.1 Hardware

The hardware used in Table 4.2 is capable of hosting the GMB processor and reliably generating actual and future versions of the product.

Table 4-2: Processing Hardware at TU Dresden for the GMB processor

Model	Virtual Machine hosted by TUDresden Center for Information Services and High Performance Computing, 16/32 Cores available
Processor	AMD EPYC 7513 32-Core Processor, 64 bit
Memory (RAM)	64 GiB
Local Hard Drive	200 GB
Network Attached Storage	1 TB

4.2.3.2 Operating system

The operating system in use is the Linux system Ubuntu 20.04 LTS (long-term support). More recent systems, such as Ubuntu 22.04 LTS or the upcoming 24.04 LTS, are also suitable to host the GMB data production system.

4.2.3.3 Tools and Libraries

The core GMB production system is implemented using the commercial numerical computing environment MATLAB version R204a (9.12). All processing steps and the generation of products in ASCII and NetCDF format are carried out using MATLAB. Additional tools are utilised to generate products in alternative formats, check the content and metadata of the products, and visualise the final results.

4.2.4 Future concerns and developments

At present, the GMB processor is not fully automated and needs to be operated manually. The input data volume and its update frequency, as well as the processor runtime, are moderate. Hence, manual operation is not considered a major drawback, and no full automation of the GMB processor is foreseen.

4.3 References

Barletta, V. R., Sørensen, L. S., & Forsberg, R. (2013). Scatter of mass changes estimates at the basin scale for Greenland and Antarctica. The Cryosphere, 7(5), 1411-1432.

Döhne, T., Horwath, M., Groh, A., & Buchta, E. (2023). The sensitivity kernel perspective on GRACE mass change estimates. Journal of Geodesy, 97(1), 11, https://doi.org/10.1007/s00190-022-01697-8

5 Mass Flux and Ice Discharge (MFID)

5.1 System overview

The mass flux and ice discharge algorithms have two primary steps

- Calculation of discharge gate locations.
- Calculation of discharge through the gates.

Gate locations are calculated once using a baseline velocity product and then remain fixed in space and time. Discharge is then calculated for each time when a velocity map exists.

Note - some of the text below comes from Mankoff et al. (2019), where this algorithm was first described.

Gates are algorithmically generated for fast-flowing ice (greater than 150 m yr-1) close to the ice sheet terminus, determined by the baseline-period data. We apply a 2D inclusive mask to the baseline data for all ice flowing faster than 150 m yr-1. We then select the mask edge, which is near the BedMachine ice mask (not including ice shelves), which effectively provides grounding line termini. We buffer the termini 10 km in all directions, creating ovals around the termini and once again down-select to fast-flowing ice pixels. This procedure results in gates 10 km upstream from the baseline terminus that bisect the baseline fast-flowing ice. We manually mask some land- or lake-terminating glaciers, which are initially selected by the algorithm due to fast flow and mask issues.

We select a 150 m yr-1 speed cut-off because slower ice, taking longer to reach the terminus, is more influenced by SMB. The choice of a 10 km buffer follows from the fact that it is near-terminus and thus avoids the need for (minor) SMB corrections downstream, yet is not too close to the terminus where discharge results are sensitive to the choice of distance-to-terminus value (Mankoff et al. 2019), which may be indicative of bed (ice thickness) errors.

Discharge is calculated at each gate by multiplying the provided velocity by the thickness of the ice at the gate location. Ice thickness is calculated assuming the BedMachine v6 basal topography as a constant bed elevation and linearly interpolated ice surface elevation based on the ESA CCI SEC relative to a fixed DEM.

5.2 Operational scenarios

The operational scenario for this product is to generate an estimate of total Greenland ice discharge for every timestep when any velocity product exists. If a velocity product does not cover all gates or all portions of all gates, velocity is linearly interpolated from the surrounding times, or forward- or backwards-filled for the last and first time, respectively.

5.3 Hardware and Software Platform

5.3.1 Hardware

The development environment is a common laptop. It currently has 32 GB of RAM and 5 TB of internal storage, but can be run on a significantly smaller system. We suggest at least 8 GB of RAM and 1 TB of storage.

5.3.2 Operating system

We use the Linux operating system, specifically Ubuntu 22..04.5 LTS, with the following CPU and architecture:

Architecture: x86_64
CPU op-mode(s): 32-bit, 64-bit
Byte Order: Little Endian

```
CPU(s):          8
On-line CPU(s) list: 0-7
Thread(s) per core:  2
Core(s) per socket:  4
Socket(s):          1
NUMA node(s):       1
Vendor ID:          GenuineIntel
CPU family:         6
Model:              142
Model name:         Intel(R) Core(TM) i7-8650U CPU @ 1.90GHz
Stepping:           10
CPU MHz:            758.158
CPU max MHz:        4200.0000
CPU min MHz:        400.0000
BogoMIPS:           4224.00
Virtualisation:     VT-x
L1d cache:          32K
L1i cache:          32K
L2 cache:           256K
L3 cache:           8192K
NUMA node0 CPU(s):   0-7
```

5.3.3 Tools and libraries

We use the following tools for this workflow:

This work was performed using only open-source software, primarily GRASS GIS (Neteler et al. 2012) and Python (Van Rossum and Drake Jr 1995), in particular the Jupyter (Kluyver et al. 2016), pandas (McKinney 2010), numpy (Oliphant 2006), statsmodel (Seabold and Perktold 2010), x-array (Hoyer and Hamman 2017), and Matplotlib (Hunter 2007) packages. The entire work was performed in Emacs (Stallman 1981) using Org Mode (Schulte et al. 2012). The parallel (Tange 2011) tool was used to speed up processing. We used proj4 (PROJ contributors 2018) to compute the errors in the EPSG 3413 projection. All code used in this work is available at https://github.com/GEUS-Glaciology-and-Climate/MFID .

5.3.4 Future concerns and developments

Development of this work is ongoing at https://github.com/GEUS-Glaciology-and-Climate/MFID as part of the ESA CCI R&D.

Concerns for the future of this work are mitigated by the fact that the entire process is documented, and the code is available for anyone to replicate and improve the product. However, replicating the work has proven non-trivial so far.

Nonetheless, concerns include the quality of the ice thickness data and the ability to generate SEC maps at a regular frequency and with enough spatial coverage to cover the gate locations.

5.4 References

Hoyer, Stephan, and Joseph J. Hamman. 2017. “Xarray: N-D Labeled Arrays and Datasets in Python.” *Journal of Open Research Software* 5 (April). https://doi.org/10.5334/jors.148.

Hunter, John D. 2007. “Matplotlib: A 2D graphics environment.” *Computing in Science & Engineering* 9 (3): 90–95.

Kluyver, Thomas, Benjamin Ragan-Kelley, Fernando Pérez, Brian Granger, Matthias Bussonnier, Jonathan Frederic, Kyle Kelley, et al. 2016. “Jupyter Notebooks – a Publishing Format for Reproducible Computational Workflows.” Edited by F. Loizides and B. Schmidt. IOS Press.

Mankoff, Kenneth D., William Colgan, Anne Solgaard, Nanna B. Karlsson, Andreas P. Ahlstrøm, Dirk van As, Jason E. Box, et al. 2019. “Greenland Ice Sheet Solid Ice Discharge from 1986 Through 2017.” *Earth System Science Data* 11 (2): 769–86. https://doi.org/10.5194/essd-11-769-2019.

McKinney, Wes. 2010. “Data Structures for Statistical Computing in Python.” In *Proceedings of the 9th Python in Science Conference*, edited by Stéfan van der Walt and Jarrod Millman, 51–56.

Neteler, M., M. H. Bowman, M. Landa, and M. Metz. 2012. “GRASS GIS: A Multi-Purpose Open Source GIS.” *Environmental Modelling & Software* 31: 124–30. https://doi.org/10.1016/j.envsoft.2011.11.014.

Oliphant, Travis E. 2006. *A Guide to Numpy*. Vol. 1. Trelgol Publishing USA.

PROJ contributors. 2018. *PROJ Coordinate Transformation Software Library*. Open Source Geospatial Foundation. https://proj4.org/.

Schulte, Eric, Dan Davison, Thomas Dye, and Carstin Dominik. 2012. “A multi-language computing environment for literate programming and reproducible research.” *Journal of Statistical Software* 46 (3): 1–24.

Seabold, Skipper, and Josef Perktold. 2010. “Statsmodels: Econometric and Statistical Modeling with Python.” In *9th Python in Science Conference*.

Stallman, Richard M. 1981. “EMACS the Extensible, Customizable Self-Documenting Display Editor.” *Proceedings of the ACM SIGPLAN SIGOA Symposium on Text Manipulation* -. https://doi.org/10.1145/800209.806466.

Tange, O. 2011. “GNU Parallel - the Command-Line Power Tool.”*;Login: The USENIX Magazine* 36 (1): 42–47. https://doi.org/10.5281/zenodo.16303.

Van Rossum, Guido, and Fred L. Drake Jr. 1995. *Python Reference Manual*. Centrum voor Wiskunde en Informatica Amsterdam.

End of document