

**Greenland Ice Mapping Project 2 (GIMP-2) Algorithm
Theoretical Basis Document**

A NASA MEASURES PROJECT

Ian Joughin

Ben Smith

Ian Howat

Ted Scambos

Version: 2.0

1 Introduction

Greenland's mass balance, the difference between annual snowfall and ice loss from melting and iceberg calving, has changed dramatically over the last two decades [Shepherd *et al.*, 2012]. One of the first large changes was the rapid disintegration and near-doubling in speed of Jakobshavn Isbrae, Greenland's largest outlet glacier [Thomas *et al.*, 2003; Joughin *et al.*, 2004]. Large speedups soon followed on Greenland's second and third largest outlet glaciers, Kangerdlugssuaq and Helheim [Howat *et al.*, 2005; Luckman *et al.*, 2006]. Over the same period nearly all of the glaciers along Greenland's southeast coast sped up by 50% or more [Rignot and Kanagaratnam, 2006]., with continued acceleration since then [Joughin *et al.*, 2010; Moon *et al.*, 2012].

The highly-variable dynamics of outlet glaciers suggest that recent Greenland observations provide only isolated snapshots of mass balance. Therefore, special care must be taken in how these and other mass-loss estimates are evaluated, particularly when extrapolating to the future, since short-term spikes could yield erroneous long-term trends. Rather than yielding a well-defined trend, recent results are significant in that they

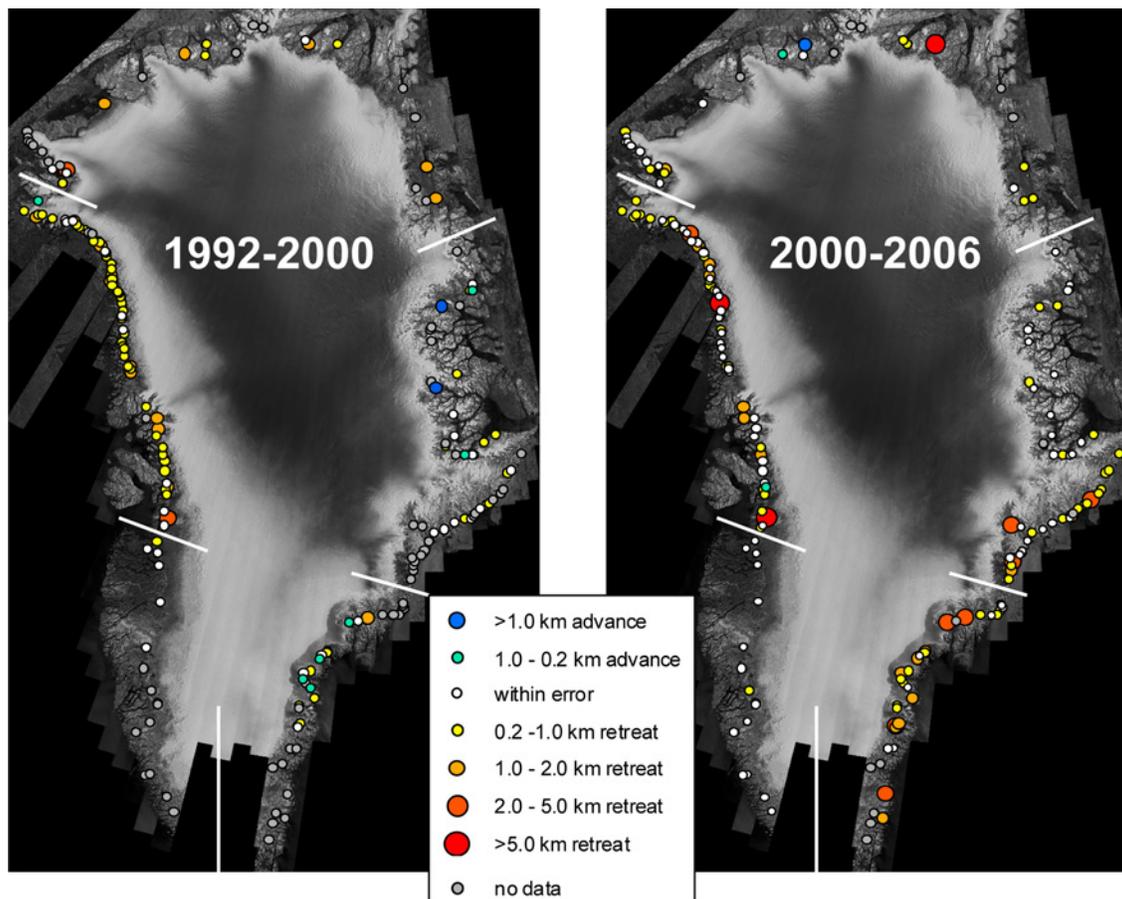


Figure 2.1. Retreat of Greenland's outlet glaciers (those with widths >2-km) over the periods from 1992-to-2000 and from 2000-to-2006. Note the significantly greater rates of retreat, particularly along the southeast coast, in the warmer 2000-to-2006 period (e.g., Figure 2). We will derive vector ice-front positions to continue the record of retreat shown here.

show Greenland mass balance can fluctuate rapidly and unpredictably. *Therefore, accurate estimates of Greenland's mass-balance and a firm understanding of the dynamics that drive mass balance requires annual-to-sub-annual observations of outlet glacier variability to avoid aliasing of this rapidly varying signal.*

Changes similar to those we see during the current warming may have occurred in Greenland during the warming in the 1930's. Unfortunately, we have only limited data with which to look back at this period. In order that current and future researchers not face similar limitations with the present warming, through the *Greenland Ice Mapping Project* (GIMP), we are producing a uniform set of Earth System Data Records (ESDRs) to measure and document the current state and near-term evolution of the Greenland Ice Sheet. In addition to providing products of utmost priority to the community currently trying to assess ice sheet stability, we also are providing an important baseline data for future generations.

1.1 Goals of This Document

This document provides an overview of the GIMP-1/2 products and descriptions of the algorithms used to produce them. Actual documentation of the products, including detailed descriptions of the formats, can be found along with the products at the National Snow and Ice Data Center (NSIDC).

2 Product Descriptions

The Greenland Ice Mapping Project is producing a comprehensive suite of geo-referenced products for the ice sheet, which includes:

- Annual to sub-annual, high-resolution (20-m), ice-sheet-wide SAR image mosaics (e.g., Figures 2.1 – NSIDC-0633) and 50-m coastal region mosaics from Sentinel (NSIDC-TBD).
- Annual ice-sheet wide velocity maps for Greenland derived using interferometric SAR (InSAR) and posted at 250/500-m intervals (e.g., Figure 2.2 – NSIDC-0478) and multi-year full ice sheet mosaic (NISC-0670).
- Sub-annual velocity maps for fast outlet glaciers at the temporal sampling allowed by the suite of sensors used to generate the estimates (e.g., ASTER, Landsat, and TerraSAR-X) and posted at 100-m intervals (NSIDC-0646 and NISDC-0481).
- Monthly (March-September) lower-resolution (250-m) ice-sheet-wide MODIS image mosaics (NSIDC-TBD).
- An all-Greenland Digital Elevation Model (DEM) posted at 30-m intervals, with both time-averaged and time-varying versions to enable elevation change measurement (Version 1 NSIDC-0645; and Version 2 NSIDC-0715).
- Greenland-wide 15-m rock and ice surface classification masks constructed from Landsat for the year 2000 (NSIDC-0714).
- Greenland-wide 15-m pan-sharpened Landsat 7 mosaic for the year 2000 with time-stamped pixels for change detection (NISDC-0713).

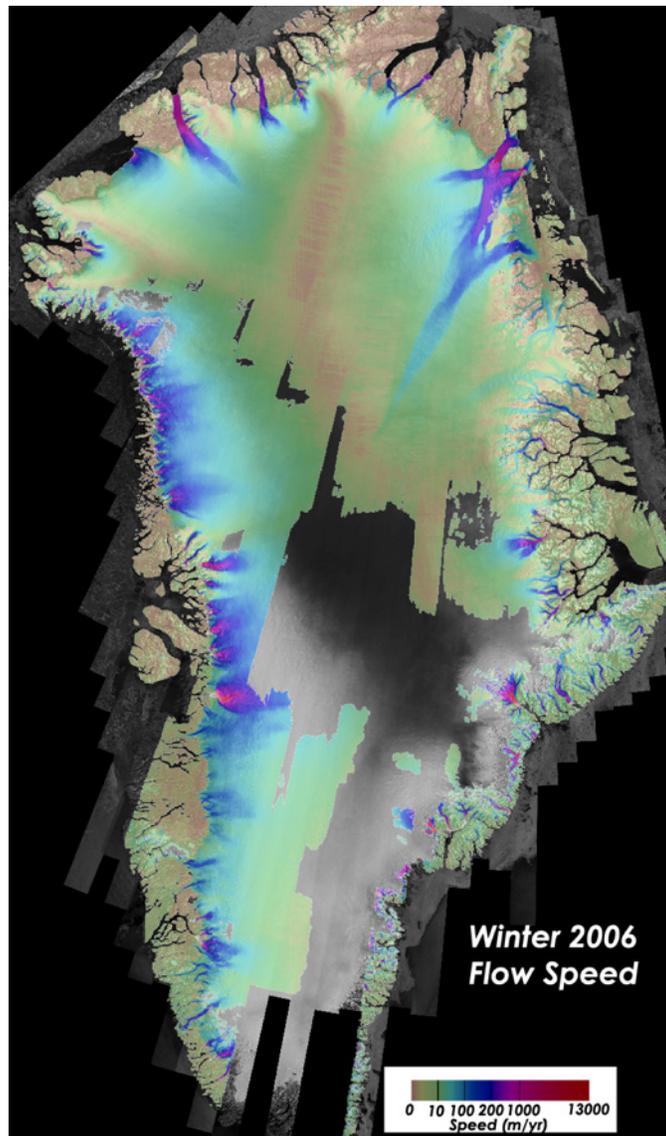


Figure 2.2. Map of flow speed for winter 2006 over a SAR image mosaics for the same period. Because the range of speeds varies from 0 to nearly 14-km/yr, a log color table is used. Irregular gaps in coverage correspond to regions where the interferometric coherence is poor. Despite these gaps, we obtain velocities on nearly all fast moving outlet glaciers, which represent the most critical regions for this study. Note at each point where speed is shown, we obtain estimates of both horizontal components of velocity. Both the velocity and image mosaics shown here are representative of the types of GIMP products.

- Annual vector ice-front positions for all Greenland glaciers (~180) wider than 2-km (NSIDC-0642).

2.1 SAR Image Mosaics

The GIMP project produces annual, high-resolution (20-m), ice-sheet-wide SAR image mosaics for the Greenland Ice Sheet (e.g., Figure 2.1). These products are distributed both as calibrated radar backscatter, σ_0 , and as grayscale values with a

contrast stretch optimized for visual interpretation. We produced new RADARSAT mosaics for the winters of 2000/1, 2005/6, 2006/7, 2007/8, and 2008/9, and 2012/13 [Joughin *et al.*, 2016]. For 2009/10, GIMP provides an uncalibrated mosaic for the 2009/2010 winter using ALOS L-band SAR data.

As an additional product, GIMP produces un-calibrated geo-referenced images at around the Greenland margin starting in 2015 using Sentinel 1A/B data. These data are available for 2015 through present (at 50-m resolution).

2.2 Annual Velocity Mosaics

Using the same data used to derive the SAR mosaics, GIMP generates annual velocity mosaics (e.g. Figure 2). Currently, the project has sufficient data for a minimum of 6 winters: 2000/1, 2005/6 through 2009/10, 2012/2013. From 2014/2015 onwards we rely on Sentinel 1A/B data to produce these products.

2.3 Multi-Year Velocity Mosaic

In addition to the annual velocity maps, GIMP provides a composite map that is an average of the annual maps. This product provides the greatest accuracy and coverage on the slow-moving interior regions, which provides an important constraint for ice sheet modeling studies. Caution should be exercised in interpretation of this product on fast moving glaciers with changing flow speed.

2.4 Velocity Time Series for Targeted Sites

Where the necessary data are available on fast moving glaciers that are changing rapidly, we generate velocity time series that sample the glacier as frequently as possible. During daylight periods, we use ASTER/LandSAT image pairs to measure velocity. We also use fine-resolution TerraSAR-X data at least 20 sites around Greenland, which DLR began acquiring in early 2009. Although not as frequent, these tiles also include a more extensive set of TerraSAR-X sites, which are only acquired in winter.

At each targeted site, we establish a rectangular grid so that all the estimates for that glacier are geo-referenced to the same area, creating a “stack” of velocity maps through time. Although some products are oversampled, we produce all of these products with a 100-m posting to accommodate the higher resolution of TerraSAR-X.

2.5 Greenland DEM

We have combined the best available data from several sources to produce a new high-resolution DEM for Greenland with a 30-m posting. Our original product relied on ASTER photogrammetry constrained by ICESAT and ERS altimetry. Although the resolution is limited (~250 m), this product represents the relatively smooth ice sheet surface with a 1-to-10 m error suitable for most glaciological applications. Our subsequent GIMP DEM version 2, is largely based on WorldView stereo imagery, greatly improving accuracy.

2.6 MODIS Mosaics

We produce a monthly (March-through-September) set of MODIS image 250-m resolution mosaics using cloud-free near-nadir regions from several individual images. The processing for these mosaics emulates the MODIS Mosaic of Greenland, although

with fewer images (similar to the recently-published MODIS Mosaic of Antarctica [Scambos *et al.*, 2007]). For each monthly product, we produce gray-scale ~200-m resolution Band 1 enhanced-resolution Greenland mosaics. We also produced a single, pan-sharpened, 250-m true color (Band 1-4-3) mosaic product.

2.7 Annual Ice Front Positions

Using the high-resolution image mosaics, we hand-digitize the annual positions of the ice fronts for all of Greenland's outlet glaciers with width of 2-km or more. Figure 2.1 demonstrates the utility of such products for detecting ice front retreat. These vector ice front positions are distributed both as ESRI shape files and ASCII text files.

3 SAR Image Mosaics

We have produced several SAR image mosaics as part of this project. In addition to the mosaics, the SAR data are used to generate velocity products. Thus, aspects of the SAR processing (e.g. geo-location) also are relevant to the velocity products.

3.1 Range-Doppler Processing

Fine azimuth resolution (~5 m) is achieved in SAR images by correlating the data with the expected Doppler history (reference function) for each pixel. The Doppler centroid is the frequency about which the Doppler bandwidth is centered. If the Earth did not spin, the Doppler centroid would be zero for a radar that images from broadside. The motion of the earth, however, shifts the Doppler spectrum to yield a non-zero, latitude-varying Doppler centroid. ERS-1 and 2 use yaw-steering to continually adjust their pointing to compensate for the Earth's motion in order to achieve a nearly zero Doppler centroid that exhibits only minor variation along track. ERS data are well suited to most SAR processing algorithms, which process an entire image using a fixed Doppler centroid. All the images we produce processed with a basic range-doppler algorithm to produce single-look complex images, which are subsequently multi-looked, calibrated, and mosaicked.

3.1.1 RADARSAT

RADARSAT operates with a nominally broadside imaging geometry so that the Doppler centroid varies significantly with latitude, particularly near the poles. This can degrade image quality when there is no compensation for the shift in Doppler centroid, limiting the portion of a data take that can be processed as single image (e.g., <200 to 300 km). Breaking a data take into several small frames makes later processing considerably more difficult. To avoid this limitation, a method is needed that allows along-track Doppler updates.

A typical range-Doppler processor creates an image as a sequence of patches. Simply updating the Doppler centroid for each patch is a non-trivial solution, because this changes the image geometry at each patch, leading to discontinuities at patch boundaries without proper compensation. Instead we used a simpler approach that allowed easy modification of an existing version of the Gamma Modular SAR processor (MSP) [Werner *et al.*, 2000]. Note more recent versions of the MSP processor include a similar

implementation. For consistency with past data sets, however, we use the version that we modified specifically for RADARSAT.

Doppler bandwidth is limited by the width of the antenna pattern. The pulse repetition rate of the radar is set to adequately sample this bandwidth. To maintain a consistent patch-to-patch geometry, we assumed a fixed Doppler centroid for the entire image, but with a reference function that has twice the bandwidth (length) imposed by the antenna pattern. This double bandwidth reference function, however, is not adequately sampled. To avoid aliasing, at each patch the reference function is band-pass filtered to use only an adequately sampled portion of the bandwidth centered about the true Doppler centroid. This has the effect of compressing the energy in the actual Doppler bandwidth, while maintaining the geometry of a fixed Doppler centroid throughout the image. The length of the expanded reference function is arbitrary. This double-length reference function allows the Doppler centroid to vary by roughly 1250 Hz along a RADARSAT image, allowing roughly 1000-km of SAR data to be processed as single image. The penalty is a reduction in processing efficiency because twice as many samples are lost at the end of the patches during the convolution. Large patch sizes can reduce this computational penalty to 10 to 20%.

3.1.2 ALOS

Because the Doppler does not vary significantly for ALOS, we process all ALOS data to single-look complex images using the Gamma MSP (version 20080630).

3.1.3 TerraSAR-X

The German Space Agency (DLR) does not provide raw signal data. Instead, we receive the single-look complex images that are produced by their standard processor, which include the precise science-quality orbits.

3.1.4 Sentinel 1A/B

Sentinel 1A/B are provided as a series of individual bursts for the SLC image. We use the Gamma ISP package to assemble these bursts into a single SLC, which we run through our processing chain in the same manner as other SLC products.

3.2 Geolocation

All of our geo-location of SAR data, including velocity products, is accomplished with the same basic set of routines to transform latitude, longitude, and elevation to the range/azimuth SAR image coordinates.

The latitude and longitude coordinates are converted to ground-range coordinates of the SAR frame using an algorithm developed by Li

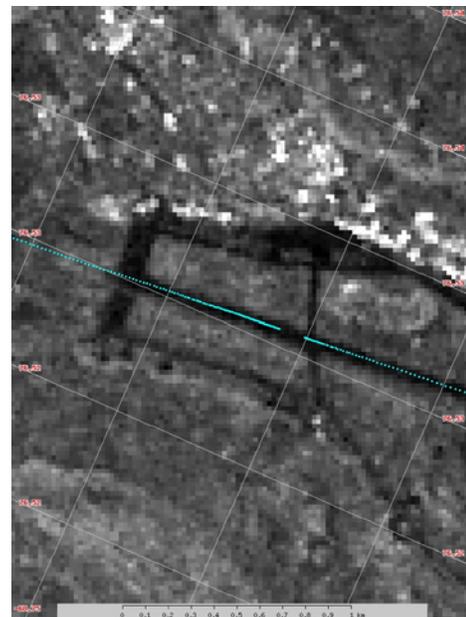


Figure 3.1 Oversampled SAR image mosaic showing runway at Thule, Greenland. Cyan-color points show a GPS runway survey performed with the NASA Airborne Terrain Mapper (ATM). The results suggest subpixel ($< 20\text{-m}$) absolute geo-location error. Note for scale the image is $\sim 2\text{km}$ wide.

[1993]. This algorithm initially transforms the latitude and longitude to an intermediate Cartesian coordinate system. It then uses the image center point, corner points, and an additional point to solve for the coefficients of a quadratic mapping from the intermediate coordinates to the ground-range coordinates. The ground range coordinates are then converted to slant-range coordinates of the SAR image using the known elevation of the point. Further details of this transformation are described by Joughin [1995].

This mapping was originally based on the image corner coordinates output by the SAR processor, which later proved inaccurate for long, non-frame swaths. Instead, we adapted the algorithm by tiling the image into several subregions, each with an individual mapping similar to that just described, but using the coordinates of the sub-region to define the transformation for that portion of the image. The algorithm uses an initial coarse mapping based on the image corner coordinates to find the approximate coordinates, which are used to determine the appropriate sub-region coordinate transform for the final coordinate transformation.

Geolocation errors of 10s to 100s of meters often are introduced by errors in the satellite-along-track time relative to the state-vectors, δ_a , and errors in the range-delay times, δ_r . Thus, the SAR images also require geometric calibration to improve location accuracy.

In our early processing of the RADARSAT mosaics we noticed along-track location inconsistencies of up to about 70 meters between adjacent tracks. To fix this, we cross-correlated patches in adjacent, overlapping images of the 2005/6 data at 15-m resolution to determine the relative timing errors. We used these relative values of δ_a to adjust the satellite timing to bring all of the tracks into relative agreement and then corrected by an additional term so that the mean shift (for 54 data takes) was zero. For the other C-band mosaics, we used a similar registration procedure except that we determined values of δ_a by aligning the data for each track with the 2005/6 mosaic. The ALOS-PALSAR along-track positions were consistent from track to track, but required a uniform correction of $\delta_a=0.062$ seconds to align the mosaicked results with independent data (e.g., LandSat and TerraSAR-X).

Comparison of the RADARSAT data with other independent images indicated that there was an approximately constant (for a given year) range delay. To fix this error, we located several easily identifiable features that were visible in both the 2005/6 mosaic and other well-geolocated data (e.g., WorldView and TerraSAR-X). Using these features, we determined that a correction in δ_r equivalent to 70 m was required to align the 2005/6 RADARSAT data with independent data. After evaluating all of the RADARSAT data, we used a fixed δ_r correction for each year ranging from 66 to 70 m. For the ALOS data no range correction was required.

After application of the range and along-track timing corrections, all of the RADARSAT mosaics were internally consistent to within better than the 20-m pixel spacing of the final mosaics. Comparison with other data (Landsat, WorldView, TerraSAR-X) indicates mean errors are also less than 20-m. Thus, the remaining source of error is the GIMP DEM used for terrain correction [Howat *et al.*, 2014]. For a nominal incidence angle of 38° , the horizontal range-dependent location error at a point should be

a factor ~ 1.25 ($1/\tan[38^\circ]$) greater than the corresponding elevation error [Curlander and McDonough, 1991]. For the ice sheet the GIMP-elevation error is 8.5 m [Howat et al., 2014], indicating sub-pixel (<20 m) location errors for ice-covered regions. For ice-free terrain the elevation errors assessed in different areas of Greenland vary substantially (8–40 m), indicating location errors of up to 50 m. In areas with severe slopes such as mountainous regions along the southeast coast, the DEM errors may be substantially larger.

In general, the DEM-dependent errors are common to all RADARSAT mosaics, so relative positions are consistent to better than 20 m. An exception occurs where glaciers have thinned by 10s of meters over the 12-year period over which the images were ACQUIRED. In such regions horizontal location errors are similar ($\sim 1.25x$) in magnitude to the corresponding elevation change.

While the RADARSAT mosaics are relatively self-consistent, location differences between the RADARSAT mosaics and ALOS-PALSAR mosaic can be substantially larger. Such differences arise because the RADARSAT images were acquired along descending orbits (satellite moving toward the equator in the northern hemisphere) and the ALOS-PALSAR data were acquired along ascending orbits (satellite moving toward the pole). As a result, points in the mosaics were imaged from opposing sides by the two satellites. Thus, if the location error due to terrain distortion is ΔX in a RADARSAT mosaic, the corresponding error in the ALOS-PALSAR image will be roughly in the opposite direction in the ALOS-PALSAR image, so the relative error between the images will be $\sim 2\Delta X$, assuming the elevation error is the same for both images. In addition, because of this opposite-side viewing, areas that are shadowed in the RADARSAT image will be laid-over or foreshortened in the ALOS-PALSAR image and vice versa. Foreshortening occurs when radar-facing slopes are such that slant-range pixels project to a large area on the ground, causing the radar-side of a mountainous feature to be compressed in slant-range coordinates. Layover occurs when slopes are such that points that are farther in ground range are nearer when mapped into slant-range coordinates than other points that are actually closer in ground range (i.e., a peak may be closer to the radar due to its height, even though it farther way horizontally than a lower elevation point). Although the terrain correction can fix much of this distortion, differences can be large (10s of meters) in areas of extreme slopes.

3.3 Calibration

The MSP processor produces uncalibrated SAR images. Calibrated SAR images typically are distributed as radar backscatter coefficient, σ^o , values, which are the ratios of the reflected power per unit area to the power incident on the surface [El-Darymli et al., 2014]. Since many SAR processors produce imagery with no knowledge of the topography, σ_o typically is computed as though the Earth were flat (i.e., zero elevation relative to the ellipsoid). Because ground-range resolution varies strongly with surface slope, the true σ_o has a slope-dependent component. Hence, for a surface with uniform scattering properties (i.e., constant σ^o), the radar derived σ^o will appear brighter for slopes facing the radar while it will appear darker for slopes facing away. Although the mosaicking algorithm uses a DEM, we have chosen not to perform a slope correction. This avoids potential spurious variations in backscatter due to errors in, and limited

resolution of, the GIMP DEM, and is consistent with the way the RADARSAT Antarctica Mapping Project (RAMP) and Modified Antarctic Mapping Mission (MAMM) mosaics were produced [Jezek, 1999; Jezek *et al.*, 2003]. For the relatively flat interior of the ice sheet, this approach produces values relatively close to the actual values. For the mountainous areas at the coast, it produces an image in which it is easier to distinguish topographic features (i.e., it adds a shaded-relief-like effect).

After application of the antenna pattern correction, the imagery was calibrated as

$$(1) \sigma^o = a \cdot DN - b.$$

Here the antenna-pattern-corrected output of the processor is given by a DN (digital number), and a and b are processor-dependent calibration coefficients, which were initially unknown. Because σ^o values have a large dynamic range, the backscatter coefficients are provided in dB (i.e., $10\log_{10}[\sigma^o]$).

To determine the backscatter coefficients, we calibrated the processor output using the MAMM mosaic of Antarctica [Jezek *et al.*, 2003] to compute the parameters in Equation (1). The MAMM mosaic was calibrated using processor parameters determined from global surveys of targets with known backscatter properties, and is thus a good reference surface to calibrate our processor. To do this, we processed uncalibrated RADARSAT images for seventeen areas around Antarctica, which we multi-looked (averaged) 4 pixels in range and 6 pixels in azimuth. After correction for the antenna pattern, we extracted data from areas of overlap with the MAMM mosaic, smoothing both data sets to 400-m resolution. For each area we performed a least-squares fit to extract the calibration parameters. Based on this analysis, we used the mean values of $a=0.03663$ and $b=0.0058$ to calibrate the Greenland RADARSAT mosaics. With these values the mean difference between the MSP processed data and the MAMM mosaic was 0.0 with a standard deviation of 0.7 dB. Some of this variance likely reflects the fact that we could not ensure that we used the same images as those that went into the MAMM mosaic.

We used a single set of parameters and antenna pattern to calibrate the RADARSAT data. As a consequence, potential variability in the instrument performance or antenna pattern with time could affect our radiometric precision [Srivastava *et al.*, 2007]. To assess the extent to which such drift might be present in our data, we examined average radar backscatter coefficients for three 12kmx12km regions of exposed rock where we expect natural variation in backscatter to be small. For these regions, the data were consistent to within ~ 1.5 dB with a decreasing trend of about 0.11 dB yr^{-1} (-0.09 , -0.13 , -0.12 dB yr^{-1} at 3 sites). We did not attempt to correct for this apparent drift. The MAMM mosaic that we used as our reference was collected at the same time as the 2000/1 Greenland mosaic. Hence, the 2000/1 mosaic should have the least uncertainty in an absolute sense inasmuch as the reference MAMM mosaic is correct.

Because we produced only a single mosaic and did not have a calibration reference for ALOS-PALSAR, we did not calibrate the L-band data. Nonetheless, as an image product it is still useful for studying changes such as glacier retreat.

3.4 Mosaicking

We mosaic the various strips of SAR imagery into a nearly seamless mosaic. In producing the annual mosaics, we select the imagery to satisfy two criteria: 1) the image are derived from approximately the same period each year; 2) the images for a given year are collected as close in time as possible, ideally from the same acquisition cycle. Nearly all RADARSAT cycles have some missing tracks, in which case we use the data for that track from the nearest (in time) available cycle.

The mosaicking code starts by defining an output grid for the specified region. Then for each image, the routine determines the corresponding bounding rectangular sub-region that contains the image. The code then loops through each point in this sub-region and maps the output coordinates to the SAR image coordinates using the geo-location procedures described above. If the result corresponds to a valid location in the image, a simple bi-linear interpolation is used to determine the value for that point.

With this mode, where overlapping images correspond to the same point in the output grid, the results are averaged together. To improve appearance, a feathering weight is applied to the points near image edges. All weights are summed in a weight buffer, and used to normalize the weighted average at the final step before the image is output. In slow moving, featureless regions in the interior, this produces a cosmetically pleasing result, since image boundaries are feathered and averaged where there is overlap. On fast moving glaciers, however, movement of the glacier between the collection of overlapping images can produce a “blurred” effect. (The alternative would be to only use a single image, which would introduce a discontinuity in fast moving regions.)

For the Sentinel 6- and 12-day velocity maps, since we mix ascending and descending coverage, rather than averaging we only use one image at each output point. The algorithm gives priority to descending images and uses the first image in the 6- or 12-day sequence.

4 Velocity Maps

Our SAR velocity maps are derived using a combination of speckle tracking and conventional interferometry. This is a processing intensive task, with several hundred Gigabytes of raw SAR data being processed to produce each velocity map.

4.1 SAR Processing

Drift in the instrument azimuth pointing from pass to pass produces differences in the Doppler centroid from image to image, which is particularly acute for RADARSAT because it does not use yaw-steering to achieve precise pointing. When this happens, the portions of the Doppler spectrum that do not overlap are not coherent and effectively act as noise when creating an interferogram or speckle tracking. To overcome this limitation,

the Gamma MSP processor was modified so that each image in an interferometric pair is processed with knowledge of the Doppler history of the other image so that only the overlapping portions of the Doppler spectra are retained. A linear function was used to represent the along-track variation of the Doppler centroid for each image, but a higher order polynomials could be substituted where needed. This additional filtering is typically taken care of during interferogram generation, but we included it in the RADARSAT processing since the data are also used for speckle tracking. This processor has been used to produce viable interferograms from data with as little as 20% spectral overlap.

4.2 Speckle Tracking

Speckle tracking takes advantage of the fact that SAR speckle, which for most applications is a source of noise, is coherent from image to image under the same conditions that produce interferometric phase coherence. Because the speckle is nearly independent from pixel to pixel, its sharply peaked correlation function allows matching through cross-correlation, which allows velocity estimate to be derived even in areas of the ice sheet that are otherwise featureless.

The cross-correlation matching operation to estimate the range and azimuth offsets can be performed using the complex or detected-amplitude images, with each method providing different advantages. For low correlation regions, the complex cross-correlation function is more strongly peaked. Reasonable matches can be achieved with complex correlation down to about 0.2 with a relatively small patch size (e.g., 24-by-24 pixels). A larger box size is needed to achieve a match with low-correlation amplitude data. The disadvantage to complex matching is that phase gradients (i.e., the interferometric phase) across the patches being matched can reduce or even eliminate the correlation peak, making it difficult to achieve matches in regions with high shear or steep topography. Amplitude matches are unaffected by the phase and can provide good matches in such regions.

To retain the advantages of both types of matches, we apply a matcher that uses a hierarchal approach. At each point, a complex match is attempted first. To reduce the impact of phase gradients, the local phase gradient is estimated using the interferogram and removed from one of the patches. Although the patch size used for the matches is 48-by-48 pixels, a Hanning window function is applied to the data, which reduces the size to approximately 24-by-24 pixels. The use of a small patch size helps minimize the effects of phase gradients. The algorithm decides whether to accept or reject the match based on an empirically determined correlation threshold of 0.18 and a limit on the range of acceptable offset values.

When a complex match fails, an amplitude match is attempted using 64-by- 64 pixel patches. A correlation threshold of 0.07 is used to accept these matches. If the match fails, then a third attempt is made using amplitude cross-correlation, but with a patch size of 192-by-192 pixels. If this fails, the matcher gives up and records a no-match value.

Before matching, any Doppler carrier on the data is estimated and removed for each patch and the data are oversampled by a factor of two. These steps avoid aliasing that can lead to biases in the estimates. After each cross-correlation, the peak is

oversampled by a factor of 10. The effect of the combined oversampling operations leads to a match resolution of 0.05 pixels. The matches are performed every 24-by-24 pixels intervals in the single-look images.

Once the matches have been completed, they are run through a program to cull out bad matches. At each point, the median for a surrounding 9-by-9 box is computed. Points that differ from the median by some threshold are discarded. The data are then smoothed with a moving average filter with dimensions selected by the user based on the quality of the data and resolution requirements. For fine-beam RADARSAT data, typical filter dimensions are 4-range-by-6-azimuth offset pixels, which yields a resolution of roughly 700 m. Averaging can reduce noise and quantization error in many cases to below 0.01 pixels. Once the culling is complete, small holes are filled via interpolation.

In the culling process the offset sample variance is estimated for a box surrounding each pixel. To remove the effects of trends in the data, a plane is fit to the data in each box and then subtracted from the data. The variance estimate is then reduced to account for the smoothing by the moving average filter. This is straightforward for the 24-by-24 complex matches computed on a 24-by-24 pixel grid, because the estimates are independent and the variance is reduced by one over the number of samples averaged. Since amplitude matches overlap each other on the 24-by-24 pixel grid, the effective number of samples averaged, N_{eff} , is less than the actual number averaged, N_{avg} . To account for this difference, the effective number averaged is computed as

$$(4.1) N_{eff} = \max \left(\sum_{i=1}^{N_{avg}} \frac{24^2}{n_i^2}, 1 \right),$$

where n_i is the size in points of the box used for the match. A minimum of at least one effective look is assumed. If all the pixels are complex matches, then $N_{eff} = N_{avg}$.

The offset variance estimates described yield local error estimates, but fail to resolve longer wavelength errors. This is a reasonable characterization of the range-offset errors. For the azimuth offsets, however, there are longer wavelength errors that appear as streaks across the azimuth offset estimates with along-track variability of several kilometers. These “streaks” are related to ionospheric effects [Gray *et al.*, 2000]. Locally these errors can be significantly larger (e.g., 0.1 to 0.2 pixels at C-band) than the image-wide estimate (0.02 to 0.04 pixels). It is difficult, however, to reliably estimate the spatial pattern of the errors. As a result, these errors are accounted for by using an estimate of the streak variance for the entire image, which is added to the sample variance estimate. This value is larger than most other errors so that the weighting will tend to de-emphasize azimuth-offset data in areas where more reliable range-offset or phase data are available.

4.3 Fast Tracking

In fast moving regions, our speckle tracking procedure can fail for one of 3 reasons. First, the extreme motion produces offsets that exceed the search window of our normal speckle tracker. Second, large strain rates produce strong variation that our culling procedure interprets as noise, and as a result, records a no-data value. Finally, poor interferometric correlation fails to yield a match. In the latter case, there may be

features in the images (e.g. crevasses) that could provide a match, in which case, the poorly correlated speckle may act as a source of noise.

For fast moving regions, we apply a special fast-tracking procedure. First, we match over a larger window, which can accommodate large displacements. Second, we increase the tolerance in the culler, so that a larger range of motion can be accommodated, at the expense of greater noise. Finally, in addition to tracking the speckle, we also match on the amplitudes after smoothing with a 3-by-3 boxcar filter to reduce the speckle and improve matches from features. The two separate (unsmoothed and smoothed image) matches are compared and the best match is selected. At the completion of this fast matching procedure, the results are merged with those from the normal procedure.

4.4 Interferogram Generation

Interferograms are produced with the Gamma software [Werner *et al.*, 2000] with minor modifications to accommodate the along-track variation in Doppler. Because displacements of several pixels can occur over 24-days, additional modifications were made to co-register the images using offsets on a 1-km grid rather than parameterizing each component of the offset field as a plane.

To remove the 2π -ambiguities in the interferometric phase, the data are unwrapped with a variation [Joughin, 1995] on the branch-cut algorithm [Goldstein *et al.*, 1988]. Many fast-moving areas in 24-day data cannot be unwrapped. Furthermore, the interferogram can be broken into several unwrapped but isolated regions. These disconnected regions all have relative phase offsets with respect to each other, which must be resolved. These ambiguities can be removed in many cases using the speckle-tracked range offsets, which are an absolute but noisier and lower-resolution estimate of the same range difference as the phase. This method is used in regions where there are sufficient data to overcome the errors in the range-offset data. Areas where an accurate estimate of the ambiguity cannot be obtained are discarded. Tests reveal that this method occasionally misses the true ambiguity by 1 to 2 multiples of 2π . For 24-day data, however, the error is relatively small (i.e., 1 to 2 m a^{-1}) and is often reduced further during baseline estimation.

After the unwrapping procedure, there may be several holes in the data. Interpolation is used to reduce smaller holes. This procedure begins by locating all the pixels in a given hole. A decision is then made whether to fill the hole based on a user selected threshold for the area of the hole. Holes are filled using the weighted sum of pixels on the hole border. The weights are determined as the inverse of the squared distance between the point being filled and each border point. Large holes are left unfilled.

The final step in preparing the interferograms consists of baseline estimation. Baseline parameters are determined using a least-squares fit to several control points of known velocity and height in each interferogram [Joughin *et al.*, 1996].

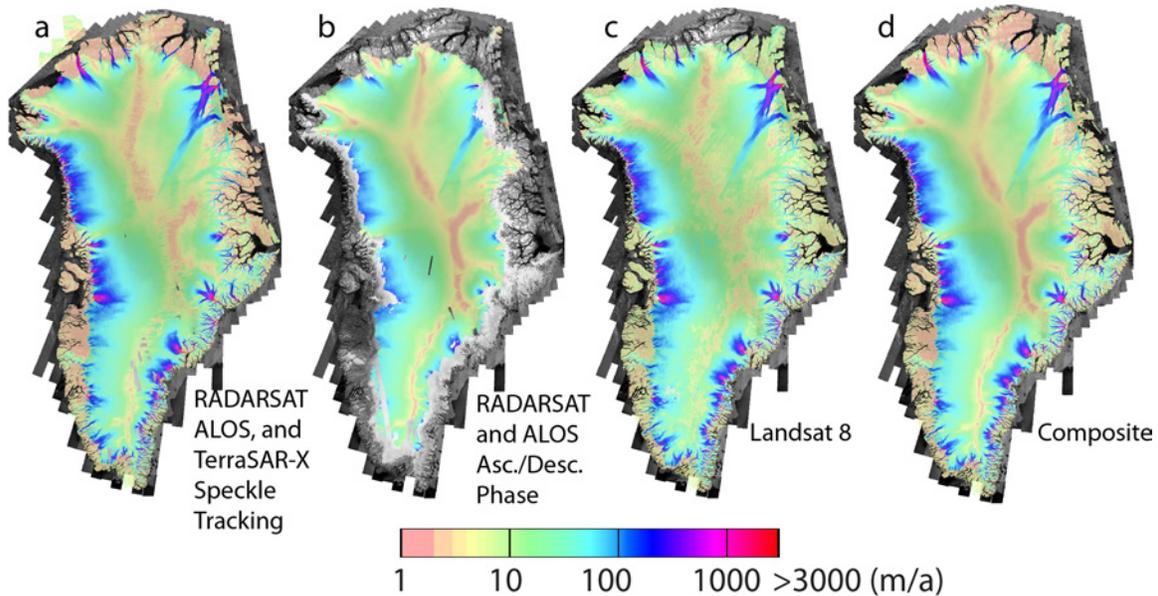


Figure 4.1. Ice sheet velocity estimates from a) speckle tracking b) interferometric phase in interior regions where phase can be unwrapped, c) Landsat 8, and d) and error weighted combination of all

4.5 Velocity Mosaicking

Once the interferograms and speckle-tracked offset fields are computed, they are used to estimate velocity. The velocity map is made by combining the phase information from overlapping ascending and descending pairs [Joughin *et al.*, 1998], by combining phase with azimuth shift data [Joughin, 1999], and finally by combining range and azimuth offset data to yield velocity [Gray *et al.*, 1998]. The mosaicking process begins by setting up output buffers for each of the velocity components. A polar stereographic projection with 45° rotation and standard latitude of 70° is used for the output grid. Data are posted at 0.5 km resolution. At any point in the output there may be data from three potential sources: phase/phase, phase/offset and offset/offset. For most of the annual maps, we have only descending RADARSAT data, so phase/phase estimates are not computed. In contrast, ALOS data are collected along ascending orbit tracks. If we acquire these data, then we will produce at least some phase/phase estimates in a composite product that averages results from all years.

The program begins by estimating the velocity using the phase/phase data. For each ascending image, the program loops through the descending images to find the areas of overlap. The velocity in these regions is estimated using the method described by Joughin and others [1998]. Each component of the result is then weighted and added to the output buffer. The weighting factors for the two velocity components, which are described below, are accumulated in separate buffers.

The phase/offset estimates are determined next, with velocity determined using the interferogram and accompanying azimuth-offset data. The range component is computed using the phase along with the slope, which is used to compensate for the vertical component of motion [Joughin *et al.*, 1996]. There are geometric effects in the azimuth data that are unrelated to surface motion. A 3-parameter model (similar to a plane fit) for the along-track variation of the non-surface-motion-related azimuth offsets

is solved using a least-squares fit to the control data. The result is used to correct the azimuth-offset data so that only a simple scale factor is needed to solve for the azimuth velocity component. The resulting velocity estimate is still in the radar coordinates, so a rotation is applied to transform the data to the output grid. The results are then weighted and summed in the output buffers.

Finally, the velocity is determined using both the range and azimuth offset data (offset/offset). This procedure is similar to the phase/offset estimate, except that the range difference is determined using offsets rather than phase. A separate baseline fit is used for the range-offset data. This fit uses a slightly different baseline model since the constant to determine absolute values is different for the range offsets and phase. This separate baseline fit also helps mitigate against any small systematic differences that may exist between the offset and phase estimates of the range difference.

As mentioned above, individual estimates are weighted and summed in an output buffer. As a result, for the v_x component the data are summed as

$$(4.2) \quad v_x = \sum_{i=1}^{N_{\text{samp}}} w_i v_{x,i}.$$

If the data are all statistically independent, then a minimum variance estimate is obtained by using weights, w_i , that are proportional to the inverse variance, $\sigma_{x,i}^{-2}$. Thus, with modifications described below, equation (4.2) is used to determine the weights. First, the variances for each velocity component are determined in the radar coordinates. These results are used to determine the variances and weighting factors for the data in the output coordinates. The weights are accumulated in separate buffers so the results can be used to renormalize the final results, forcing the final weights to sum to one.

The variances for the offset data are estimated by the culling program as described above. In the absence of other sources of error, the phase variance is determined by the interferometric correlation. Other sources of error such as tropospheric water vapor [Goldstein, 1995] are more difficult to characterize and are often larger than phase noise due to decorrelation. As a result, a nominal value of $\pi/2$ radians is used for the phase variance, which in most cases overstates the error.

The minimum variance estimate assumes that the averaged samples are independent. When phase/phase and phase/offset estimates are summed, the data are not independent. The same is true for phase/offset combined with offset/offset. To account for the possibility of double averaging, the variances of the phase estimates are doubled for the phase/phase estimates, since in most cases, the corresponding phase/offset data will also be estimated. Similarly, during the phase/offset estimation, the program checks whether a velocity estimate has already been made for the current estimate. If so, it assumes it was a phase/phase estimate and doubles the variance estimate. Doubling the variances in this fashion does not perfectly account for double averaging since the program does not keep track of the full history of what data have been summed at each point. When the program errs, it is more likely to assume that double averaging takes place so that the error tends to be over- rather than under-estimated. A similar procedure

is used to avoid double averaging the phase/offset and offset/offset combinations. There are cases where the phase/phase estimates can result in double averaging (e.g., one ascending pair with two overlapping descending images). Avoiding these errors would require an additional buffer to track the history of each image used in the velocity estimate. For the present, this feature has not been incorporated as such cases are reasonably rare with the amount of data typically available.

While the weighting method described above is designed to achieve a minimum variance estimate, it may be sub-optimal with respect to other factors. In particular, a discontinuity at a data-take boundary is a non-physical result and can lead to problems when attempting model inversions. As a result, additional weighting is employed to “feather” the data and redistribute local errors over a wider range. As the velocity is estimated for a data- take or data-take pair, the result and initial weights are saved in a temporary buffer. An additional weighting function is used to apply a linear taper from the edge of the data to some distance from the edge. For example, if the feather length is 20, then pixels on the edge are weighted by 0, pixels within 20 pixels of the edge are weighted linearly from 0 to 1, and interior pixels by 1. The feathering weights are used to update the initial weights in the temporary buffers, and the results are added to the weight buffers.

All SAR-only mosaics (NSIDC-0478) are produced as just described. The selected glacier velocities (NSIDC-0481) are produced in the same way except that in most cases a single TerraSAR-X image is used, so there is no averaging.

4.6 Landsat for Full Ice Sheet Mapping

The methods we apply to Landsat data for outlet glaciers are described below. Here we describe a simpler set of Landsat-8 algorithms we use to augment our SAR-derived all ice sheet mosaics. We computed velocities by applying a simple cross correlation procedure to high-pass filtered Landsat 8 data [Ahn and Howat, 2011; Fahnestock et al., 2016]. Unlike the SAR data, the Landsat images are reprojected to the output coordinate system prior to matching. We scale the displacements by the appropriate latitude dependent length-scale correction factor for the NSIDC polar stereographic projection. After this rescaling, the matching procedure ideally should provide absolute displacements directly from the offsets, but the uncertainties in the image registration are too large to produce sufficiently accurate results. Thus, we use a least-squares procedure to fit a plane to each scalar offset field, using the control points described below. After applying this correction, the offsets only need to be scaled by the time interval between images to produce an estimate of the velocity.

For the Landsat 8 data we use a similar procedure as that for the SAR data to estimate the matching error from neighborhood statistics. When we examined the data more closely, however, we found there were longer wavelength errors of comparable or greater magnitude to those detected by the neighborhood statistics, which are likely related to sensor errors or atmospheric effects. Since we use the radar data as control points (see below), we have a well sampled set of control points that generally sample the full scene. As a result, the residual errors from the parameter fits provide a good estimate for the average error for the full scene averaged over all length scales, but with no detail

on the spatial variability. On the other hand, while the neighborhood statistics do not estimate all sources of error, they do at least provide information on the spatial distribution of the errors related to the matching procedure. Thus, we used the following procedure to combine these error estimates. First, we compute the scene-wide average variance from the neighborhood statistics and subtract this value from the variance for residual error from the parameter fit. This difference provides a scene-wide estimate of the longer-wavelength errors. Then at each point, we add this long wavelength error to the spatially varying neighborhood statistics. The average variance of this result is identical to the scene-wide residual, but it conveys more information about the spatial variability of the error. Finally, we add to this error the uncertainty associated with the fitted parameters, which as for the baseline error model is derived from the covariance matrix from the least-squares fitting procedure.

In the mosaicker, we combine the Landsat 8 velocity as part of an error-weighted average. Figure 5.1c shows the Landsat 8 mosaic and Figure 5.1d, shows the final composite of radar and optical data for our multi-year composite product (NSIDC-0670).

5 Velocity Time Series from Optical Data

Frequent observations from optical sensors during periods with sufficient daylight provide data for measuring velocity on glacier surfaces with visible features. Our suite of algorithms produce ice-surface-velocity maps of major outlet glaciers of the Greenland Ice Sheet using optical (i.e. radiance) images from a combination of freely available satellite imagery. Since a primary motivation for the creation of this dataset is the detection and measurement of sub-seasonal variations in ice dynamics and discharge to complement InSAR-derived maps, we focus on obtaining the highest possible temporal resolution over the sensor era (1999 to present). Here we describe the data sources and feature-tracking algorithm used in this project.

5.1 Source Data

We used data from Landsat 7 Enhanced Thematic Mapper-Plus (ETM+) (1999-present), Landsat 8 Optical Land Imager (OLI, 2013-present), The Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) aboard NASA's Terra satellite (1999-present). In addition, we used SPOT5 ortho-imagery collected between 2006 and 2012 and released by the French Space Agency (CNES) under the IPY Spirit program.

5.1.1 ASTER

Launched in 1999, the ASTER instrument acquires 15-m images in the visible to near-infrared (VNIR) bands at latitudes below 84 degrees North, with a 16-day repeat interval. Raw radiance image granules are calibrated and orthorectified by the Land Processes Distributed Active Archive Center (LP DAAC) as the AST14OTH product [Abrams *et al.*, 2002]. The horizontal accuracy of 15-m AST14OTH products is cited as +/- 60 m. The LP DAAC distributes the AST14OTH products as geotifs in UTM/WGS 84 projection.

For GIMP, we obtained all available cloudless AST14OTH products over marine-terminating Greenland outlet glaciers using the USGS GLOVIS and WIST ordering

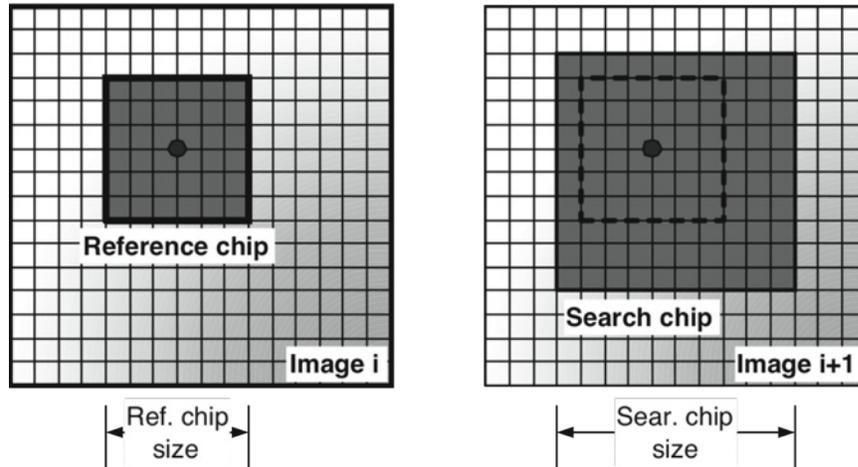


Figure 5.1. Reference and search chip in area-based matching. Dashed box in search chip illustrates overlay of reference chip toward reference chip to compute correlation coefficient.

systems. Each image was re-projected to the GIMP project Polar Stereographic projection.

5.1.2 Landsat ETM+ and OLI

Launched in 1999, the Enhanced Thematic Mapper Plus (ETM+) aboard the LANDSAT 7 satellite, collects 15-m resolution panchromatic (Band 8) radiance images with a 16-day repeat cycle. The USGS made the entire EMT+ image database freely available in January 2009. The USGS orthorectifies the raw data using a standard terrain correction (Level 1T) that provides systematic radiometric and geometric accuracy by incorporating ground control points while employing a Digital Elevation Model (DEM) for topographic accuracy. Geodetic accuracy of the product depends on the accuracy of the ground-control points and the resolution of the DEM used. Ground-control points used for Level 1Terrain correction come from the GLS2005 data set. DEM data used for terrain correction include SRTM, NED, CDAD, DTED, and GTOPO 30. The data are distributed by the USGS as geotifs in UTM/WGS 84 projection.

The Scan-Line Corrector (SLC) within the ETM+ instrument failed in May 2003. As a result, all images after that time have offset scans, appearing as strips of zero radiance that widen toward the edge of the image. Our feature tracing algorithm can utilize this corrupted imagery with minimal degradation to the quality of the derived velocity products.

Launched in 2013, the Optical Land Imager (OLI) aboard Landsat 8 is a substantial upgrade from the ETM+, featuring 12-bit, as opposed to 8-bit, radiometric precision and a higher signal-to-noise ratio. These qualities enable velocity extraction at both low-light levels, extending the collection season, and over the low-contrast interior of the ice sheet. We document these abilities in Jeong and Howat [*Jeong and Howat, 2015*]. The OLI data is geometrically processed using the same procedure as ETM+.

For GIMP, we obtained all available cloudless Landsat products over marine-terminating Greenland outlet glaciers using the USGS GLOVIS ordering system. Each image was re-projected to Polar Stereographic projection used by GIMP.

5.1.3 SPIRIT

Orthorectified SPOT5 imagery over Greenland were produced from images acquired in 2007 and 2008 as part of the SPOT 5 Stereoscopic Survey of Polar Ice: Reference Images and Topographies (SPIRIT) program by CNES. A description of dataset production and validation is given in [Korona *et al.*, 2009]. The SPIRIT image data is distributed in UTM projection, WGS84 datum and posted at 5 m pixels. We obtained all available SPIRIT products over Greenland, which we re-projected to Polar Stereographic coordinates.

5.2 Automated Repeat-Image Feature Tracking Overview

We produce velocity estimates using a suite of Automated Repeat-Image Feature Tracking (RIFT) algorithms applied to the optical data just described. All automated RIFT algorithms employ image-to-image matching based on cross-correlation of spatial variations in image intensity values, a procedure for which there is abundant existing literature and is similar in many regards to the speckle tracking procedures. We briefly summarize the process here. The two-dimensional (horizontal) velocity of the glacier surface is measured by tracking the displacement of visible features that move with the ice flow (i.e. crevasses, rock debris, melt features, etc.) between two, precisely geolocated images. The glacier velocity is calculated by multiplying the x-y image pixel offset of features between images by the image resolution over the time separation of the images.

The same basic workflow is used to measure displacement using automated procedures for locating the same features within two images. First, images are co-registered and cropped so that they have the same dimensions and overlap in space. Next, a square subset of pixels from the first image (the reference chip; see Figure 5.1) is extracted and compared sequentially to all subsets of pixels within a larger region of the second image (the search chip; see Figure 5.1). As the reference chip moves within search chip, the algorithm calculates the correlation coefficient at each position, producing a map of correlation intensity. The coordinates of the peak in correlation gives the conjugate point of the center of the reference chip in the search chip, and its 2-dimensional offset between the images.

The correlation coefficient is defined as:

$$\rho = \frac{\sigma_{RS}}{\sigma_R \cdot \sigma_S}$$

$$\sigma_R = \sqrt{\frac{\sum \sum (g_R(i,j) - \bar{g}_R)^2}{nm - 1}}$$

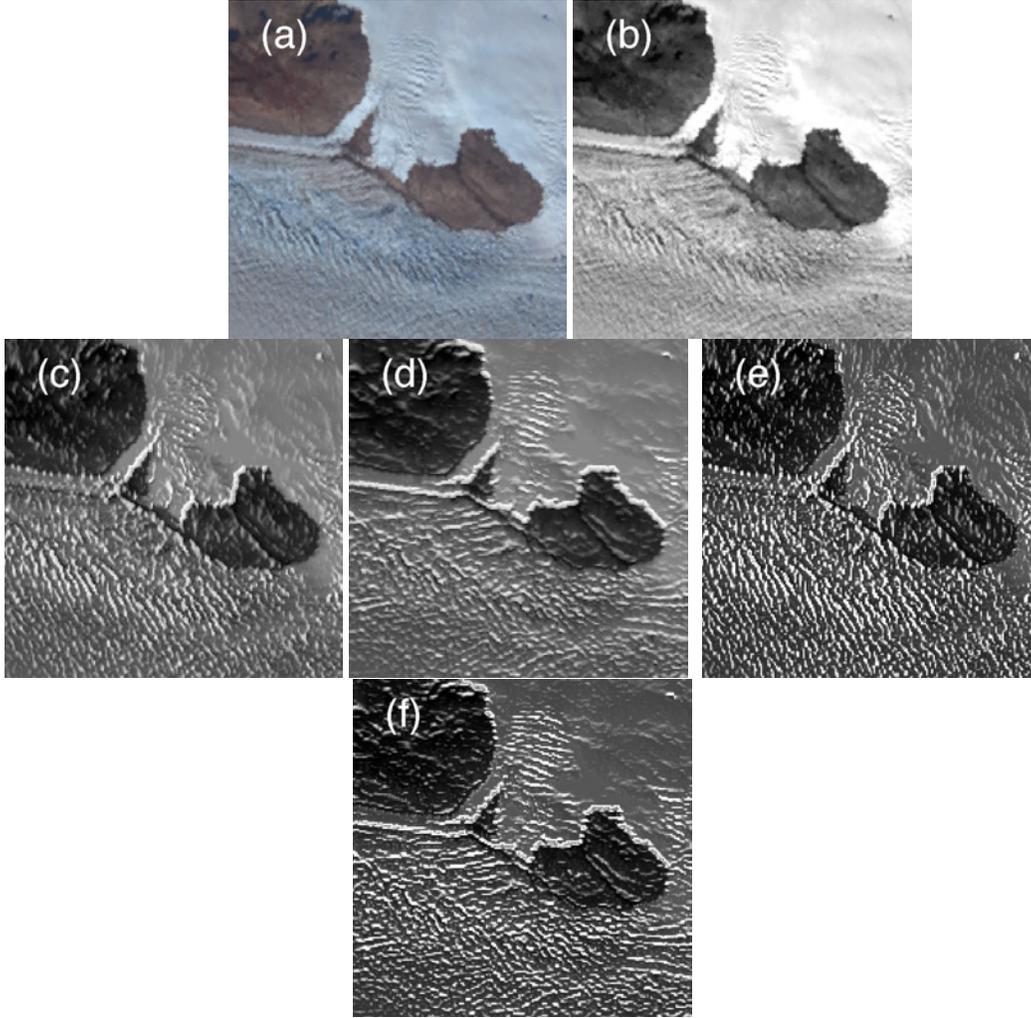


Figure 5.2. Original (a) and enhanced images: (b) principle component intensity image, (c) 0° edge enhanced, (d) 90° edge enhanced, (e) 0° edge+high pass filtered (e) (f) and 90° edge+high pass filtered (f).

$$\sigma_s = \sqrt{\frac{\sum \sum (g_s(i,j) - \bar{g}_s)}{nm - 1}}$$

$$\sigma_{RS} = \frac{\sum \sum g_R \ i \ j - \bar{g}_R \cdot g_S \ i \ j - \bar{g}_S}{nm - 1}$$

where, subscript R and S indicate reference and search chip. The pixel value is denoted as, $g(i,j)$, at position i,j and g_R and g_S represent the mean of the reference and search chip values, respectively. The values n and m provide the dimensions of the reference chip.

To determine the location of the peak in correlation within the search chip, a quadratic surface is fitted to the grid of correlation values, yielding a position as a

fraction of the pixel. For example, the equation for fitting a surface to a 3x3 grid of correlation values would be:

$$r_i = a \cdot x_i^2 + b \cdot y_i^2 + c \cdot x_i y_i + d \cdot x_i + e \cdot y_i + f - \rho_i$$

where, ρ is correlation coefficient, (x,y) is row and column positions in 3x3 chip, r is the residual, and a - f are the coefficients of the quadratic surface.

Calculating the correlation coefficients and finding the peak correlation is a relatively straightforward procedure. Several factors, however, make the practical application of this procedure much more complicated. First, there is a trade-off between the reference chip size in the uniqueness of the match and the impact of errors in image co-registration. Also, the appropriateness of the chip sizes will depend on the speed of the glacier, so that the feature does not reside outside the search chip, and the spatial scale of the feature being tracked. These effects make the appropriate chip sizes spatially dependent, impeding the general applicability of the correlation parameters. In addition, images typically must be enhanced using spatial filtering to bring out trackable features and to suppress shadows. Again, this filtering is often site and/or image specific. In addition, some minimum threshold for acceptable peak correlation must be specified, below which the match is rejected. This number also is specific to a particular image pair or location, and so care must be taken to choose the correlation threshold that filters the maximum number of spurious matches, without rejecting correct matches. Finally, spurious matches must be culled. Typically, *a-priori* information, such as the direction and maximum velocity of glacier flow, are used to eliminate spurious matches. Nearest-neighbor and non-local mean filters are employed, which assume some level of smoothness in the velocity field. Prior to Landsat 8, RIFT was only successful in the bare-ice or crevassed regions of ice sheets where there was adequate feature definition. However, the 12-bit precision of OLI is able to resolve matches in the texture of the ice sheet surface, allowing successful RIFT application to the ice sheet interiors.

5.3 The Multiple image/Multiple Chip (MIMC) Algorithm

Here we describe our method (Multiple Image/Multiple Chip, MIMC) for efficiently tracking the displacement of glacier surface features between a large number of coregistered image pairs, minimizing manual, empirical and local metrics. Successful application of MIMC requires that: A) images are cloud-free or masked; B) image pairs are precisely co-registered or that that ice-free or otherwise stationary features are visible; and C) crevasses and other surface features are visible in each image. In general RIFT can only be applied to bare-ice zone or regions of high surface strain rates (e.g. where crevasses are present), such as outlet glaciers. MIMC is documented in detail in Ahn and Howat [2011] and Jeong and Howat [2017]. Here we provide the basic components of the algorithm.

5.3.1 Image Enhancement

For each image pair, the MIMC algorithm calculates the pixel displacement field between five sets of sequentially filtered images, where each filter enhances a particular aspect of the image to improve matching (see Figure 5.2). First, matching is carried out between the raw grayscale intensity images. For multi-band imagery, MIMC generates a

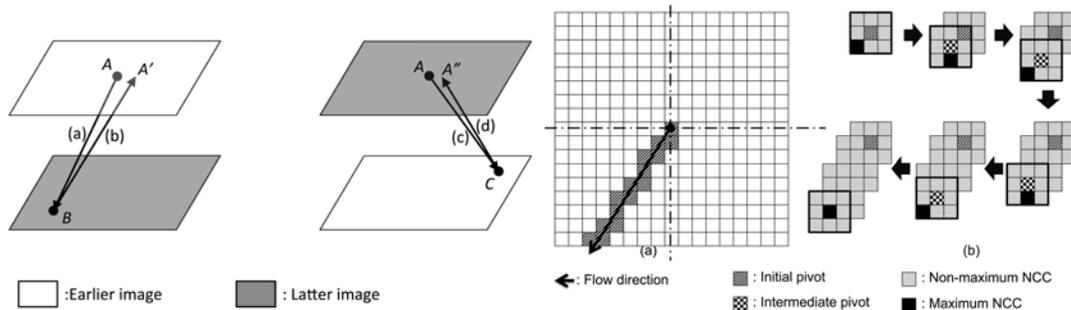


Figure 5.3. (left) Schematic illustration of quadramatching, (a): Original forward, (b): Original backward, (c): Swapped original, (d): Swapped backward. (right) Normalized cross correlation peak searching algorithm using dynamic linear constraint, (a): Initial pivots lie under the extent of flow vector, (b): iterative searching for the cross correlation peak.

grayscale intensity image from a forward principle component rotation (PCR) of the bands, also known as dimension reduction transformation [Gonzales and Woods, 2002]. The PCR reduces saturation and the visibility of fog/thin clouds and maximizes the intensity information from each band into a single value. Matching is then performed again on the intensity images following the application of 3x3 0° and 90° directional edge (Prewitt) filters. This filter enhances crevasse edges. Finally, each directionally filtered image is high-passed filtered to enhance smaller-scale features (such as crevasses) and suppress larger-scale variations induced by shadowing, topography, surface melt and fog. Examples of these filters are provided in Figure 5.2.

5.3.2 Quadramatching and Dynamic Linear Constraint Algorithms

Our algorithm builds upon the redundant solution approach of Multi-Image Multi-Chip (MIMC) RIFT method presented in Ahn and Howat [2011] with the addition of iterative forward and reverse matching that we term Quadramatching (QM). Scambos et al. [Scambos et al., 1992] proposed adding a redundancy to the match solution by swapping the reference and search chips from the initial, “forward” match and using these to solve for an additional “reverse” match. In the study, the closer the solutions of the forward and reverse matches, the greater the confidence of the displacement solution.

For the QM approach, we obtain further redundancy in the match solution by iteratively reversing both the direction of the search and the order of the image pair (pair swapping), resulting in four solutions that are combined to provide the match location and confidence. The QM procedure is illustrated in 5.3 Initially, point A on the first image (I_0) is matched to point B on the second image (I_1). A new reference chip is the extracted around point B on I_1 , and the search is repeated on I_0 , giving the backwards match to point A'. This location is referred to as the “original pairing” result. Next, the order of images I_0 and I_1 are reversed and the forward-backward search procedure is repeated, giving the “swapped pairing” match at location A''.

Although those four sets of measurements are from the same image pair, they are not identical measurements. The differences between the solutions can be identified by first defining the initial match (forward matching from the original pairing) as:

$$\textit{original forward} := f(I_0^A, I_1) = \vec{B} \quad (5.1)$$

where the first and the second arguments on the left side are the reference chip and the search window (i.e. image to find the matching point). The superscript indicates the coordinates of the image chip and the subscript is the source of the image chip or the search window (0 is the earlier, 1 is the latter). Thus, Equation (5.1) states that the original forward solution, point B in I_1 , is determined from a reference chip around point A in I_0 and a search chip in I_1 . In the same convention, the three other matches are described as:

$$\textit{original backward} := f(I_1^B, I_0) = \vec{A}' \quad (5.2)$$

$$\textit{swapped forward} := f(I_1^A, I_0) = \vec{C} \quad (5.3)$$

$$\textit{swapped backward} := f(I_0^C, I_1) = \vec{A}'' \quad (5.4)$$

Since each match solution has a different reference chip (I_0^A , I_0^B , I_1^A , and I_0^C), they yield displacement vectors with different origins. Respectively, their displacements are:

$$\vec{d}_{OF} := \overrightarrow{AB} \quad (5.5)$$

$$\vec{d}_{OB} := \overrightarrow{BA}' \quad (5.6)$$

$$\vec{d}_{SF} := \overrightarrow{AC} \quad (5.7)$$

$$\vec{d}_{SB} := \overrightarrow{CA}'' \quad (5.8)$$

with subscripts O , S , F , and B for original pair, swapped pair, forward match, and backwards match, respectively. This procedure is applied to each set of chip size and filter combinations in MIMC, thus increasing the number of displacement solutions by a factor of four.

Glacier speed can change abruptly in response to stress perturbations caused by, for example, calving front retreat and variations in basal water pressure. Measurement of these speed variations is a typical objective for RIFT applications. The direction of flow, however, tends to remain relatively constant through time; since direction is primarily

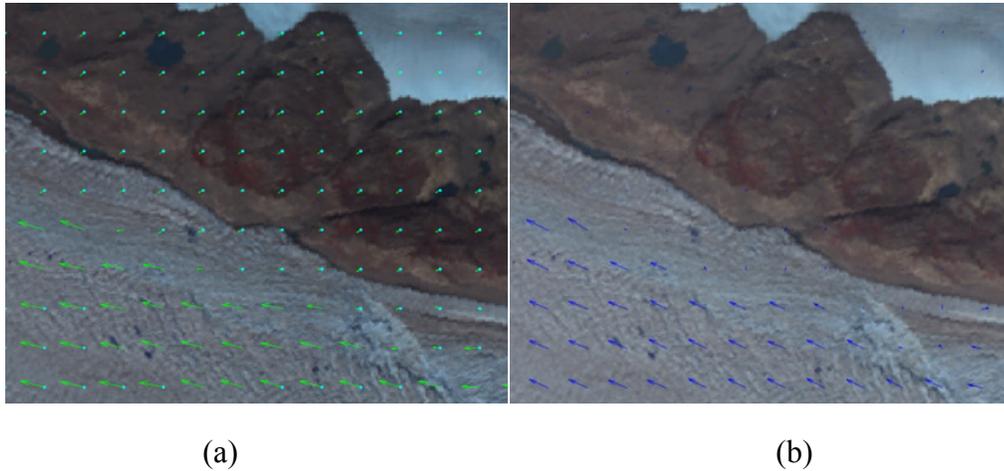


Figure 5.4. Displacement (a) without registration and (b) after registration.

determined by the surface slope and large variations in flow direction that would require large changes in the ice thickness field. Therefore, if the direction of flow is known from an existing velocity map, or can be estimated from a digital elevation model or numerical ice flow model, it can be used to both constrain the area of NCC matching and to prevent spurious matches without biasing the measurement result. Examples of this approach include Scambos et al. [Scambos et al., 1992] and Fallourd et al [2010] that made use of existing ice flow measurements as *a priori* information to limit the location of the matching solution to within a defined rectangular boundary. Another major advantage of using such a directional constraint on the match solution is the large reduction in the number of NCC computations per match area. This is especially important when utilizing redundant match approaches such as MIMC and QM.

A concern in the use of an existing velocity field as *a priori* information for constraining the matching procedure is that errors in that velocity field, or change in surface flow between the times when the *a priori* field was constructed and the imagery to be used for RIFT was obtained, could influence the result, reducing the accuracy of the solution. To mitigate this effect, we use a non-deterministic approach, illustrated in Fig. 5.3, to constrain the search for the NCC peak within the matching procedure. First, the *a priori* displacement vector, starting at the reference origin, is projected onto the search image and the pixels that intersect this vector are chosen as initial “pivots” from where the NCC peak search begins. For each initial pivot pixel, the NCC field is calculated for the search image chip centered on the pivot and the 8 surrounding pixels. The location of the maximum NCC value among these 9 solutions is chosen as the new “intermediate” pivot point for the next iteration. The NCC fields for this intermediate pivot and the surrounding pixels are calculated so that the pivot at the next stage is the location of the maximum NCC among the series. The procedure is then repeated until the point of maximum NCC is the intermediate pivot itself (the center of the 3-by-3 cell) so that there is no update in the pivot location. This iteration is then repeated for all initial pivots.

This procedure ensures that each iteration, starting from initial pivot, will converge to a local maximum in NCC space only within or near the known direction of flow. Moreover, by constraining the search area to approximate a vector, the number of NCC calculations is drastically reduced. The length of the vector of initial pivot points is bounded by the magnitude of the expected maximum displacement and the coregistration error. Here, the extent of each search grid ($L_{i,j}$) was determined by:

$$\vec{L}_{i,j} = s \cdot \left(\frac{\vec{v}_{i,j}}{r \cdot t} \cdot 365 \right) + c \quad (5.9)$$

where s is a scale factor (chosen to 1.8 in this study based on the maximum expected fluctuation of the flow speed for the test glaciers), $v_{i,j}$ is velocity (unit: m/yr), r is spatial resolution of the image, t is the length of temporal baseline of a pair (unit: day) and c is the maximum expected image coregistration error. Increasing s and/or c reduces the linear constraint, allowing for more deviation in flow direction but increasing the processing time.

5.3.3 Image Co-registration

Once the displacement field has been generated and spurious matches filtered by the clustering method, the MIMC technique improves image co-registration by selection of displacement vectors that are over off-ice regions, either manually or by specifying an off-ice mask.

Figure 5.4 shows the registration approach we apply. Apparent displacement shows (Figure 5.4a) errors in land areas that are not moving. A calculated correction offset in x-y direction that remove displacement in land area is applied to calibrate the displacement results (Figure 5.4b). Note that registration error might not be resolved by a single x-y shift. If land area is distributed well, different transformations can be applied such as conformal or affine transformations. Visual inspection of velocity in land area is used to validate the registration results.

5.4 Postprocessing

A critical postprocessing step of any redundant match approach such as MIMC and QM is the selection of the best single displacement from the population of redundant matches that will nearly always contain spurious single solutions (Figure 5.5). We have developed a novel post-processing method, termed pseudosmoothing, which determines the most probable displacement from a population of redundant matches. Similar to the “voting cell” method in MIMC version1, pseudosmoothing is applied to clusters of multiple matching results (i.e. reference chip size, image filter, and QM) in each of the single grid. We apply a connectivity-based clustering algorithm rather than clustering the displacements into discrete cells as in the voting cell method. If the minimum Euclidian difference between the displacement d_i and elements in a cluster C_k is less than a threshold value, ρ_{min} , the displacement d_i is considered as a member of C_k . This criterion for the cluster C_k can be formularized as:

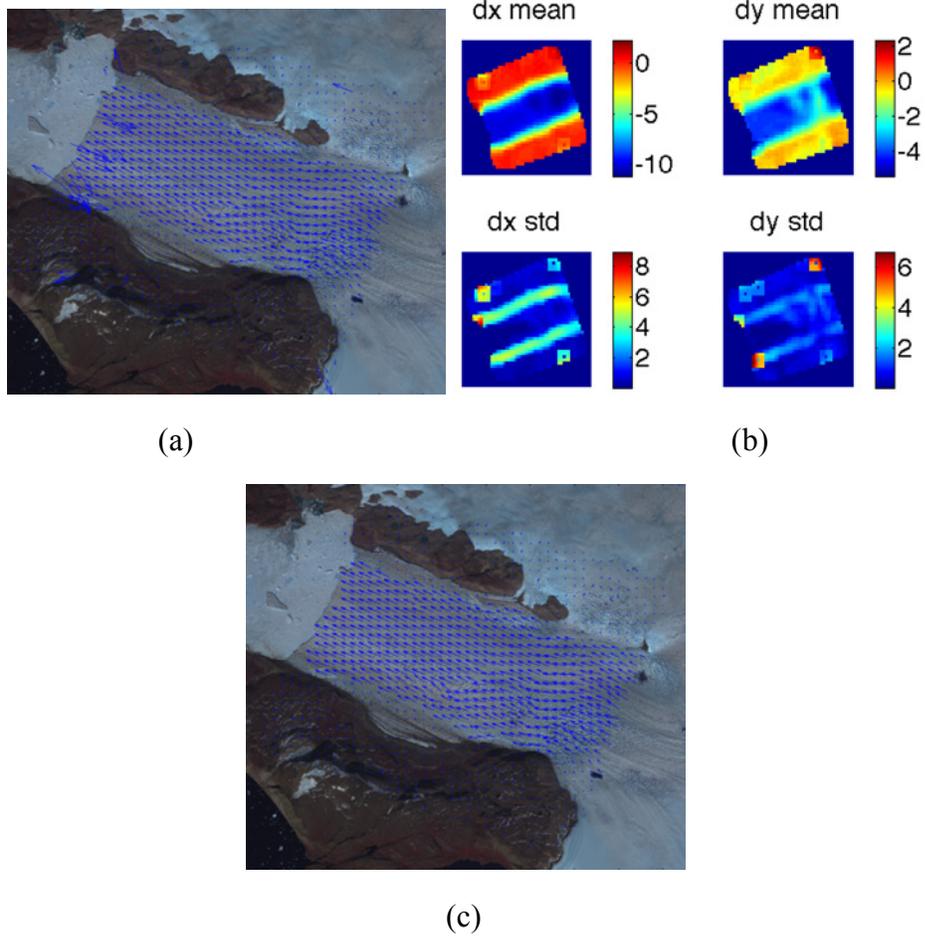


Figure 5.5. Neighboring filtering. (a) MIMC displacement results, (b) 5x5 mean and standard deviation grid for x-y direction (c) Neighbor filtered displacement results.

$$C_k := \left\{ \vec{d}_i \mid \left| \vec{d}_i, C_k \right| \leq \rho_{\min} \right\} \quad (5.10)$$

where the distance between the displacement and clusters are:

$$\left| \vec{d}_i, C_k \right| := \min \left(\left| \vec{d}_i, \vec{d}_j \right| \right), \quad \vec{d}_j \in C_k \quad (5.11)$$

therefore, for any clusters C_k and C_l ,

$$\left| C_k, C_l \right| > \rho_{\min}, \quad k \neq l \quad (5.12)$$

where

$$\left| C_k, C_l \right| := \min \left(\left| \vec{d}_i, \vec{d}_j \right| \right), \quad d_i \in C_k, \quad d_j \in C_l \quad (5.13)$$

We use C_k of equation 5.10 and a ρ_{min} of 1 pixel. Next, clusters of displacements are defined. A cluster is considered prominent if its number of samples is more than 60%

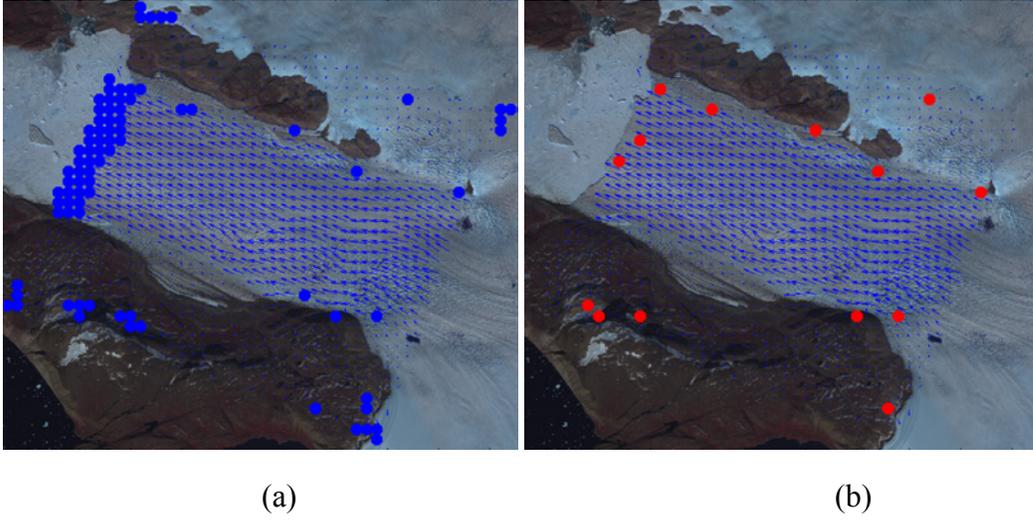


Figure 5.6. Gap filling procedure. (a) Neighbor filtering results with points that are removed from previous filtering. (b) Re-found displacements by comparing with mean grids (x-y direction).

of the total population. The mean displacement of the prominent cluster is considered the prominent displacement (d_0) in the corresponding grid. However, due to the possibility of increased spurious matching, not all grid points have a prominent cluster and corresponding d_0 s. For those locations, an intermediate displacement (d_I) is estimated instead.

The estimation of an intermediate displacement starts from finding the expected displacements (d_e) based on the given *a priori* velocity information. The expected displacement is obtained from:

$$\vec{d}_{e,(i,j)} = \frac{\vec{v}_{i,j}}{r \cdot 365} \cdot t \quad (5.14)$$

with the same variables as Equation (5.9). The difference in definitions of coordinate system between *a priori* and the image may require reversing the coordinate axes. When using *a priori* values for this calculation, seasonal and annual variability in flow speed should be considered. Therefore, a scale factor needs to be applied to d_e to compensate the variability. This factor is calculated by comparing d_e with the neighboring prominent displacements (d_0) or d_I in the earlier iterations. When there are a sufficient number of d_0 or d_I values around the location (i, j), d_I in the current iteration is calculated as

$$\vec{d}_{1,(i,j)} = \left| \frac{\hat{d}_0}{\hat{d}_{e,(i,j)}} \right| \cdot \vec{d}_{e,(i,j)} \quad (5.15)$$

where the hat denotes the averaged values in the neighbor of (i, j) . This iteration is repeated until all grids are filled with either d_0 or d_1 .

Equations (5.14) and (5.15) imply that d_1 at point (i, j) is not determined from the clustered displacements of the grid point, but from the neighboring d_0 or d_1 and the *a priori* displacement. The mean value of the closest cluster to d_1 for each grid point, therefore, is chosen as the initial displacement (d_2). In this way, the displacement at each grid point is determined from the mean values of grid clusters, which are either d_0 or d_2 .

The resulting d_2 is then adjusted through iterative and anisotropic weighted quadratic fitting. For a location (i, j) , the displacements (d_2) of the neighboring grid points are weighted using a negative exponential function with respect to distance and flow direction. To accommodate large lateral gradients (i.e. shear strains) in glacier flow, the neighboring displacements along the flow direction are given a greater weight than those in the across-flow direction. We term this postprocessing method "pseudosmoothing" because it finds a cluster's displacement close to the smooth value, but the resultant displacement is neither interpolated nor an average.

5.5 Gap filling, Interpolation and Smoothing

Filtering of spurious results leaves holes in the displacement grid. In this stage, MIMC displacements that have been removed from MIMC clustering and neighbor filtering are checked against x-y directional mean grids to fill gaps. Figure 5.6 illustrates points removed in the initial culling removed that this subsequent gap-filling procedure determined were actually valid. The closest MIMC displacement to x-y mean grids from 15 matches is searched and filled.

The gaps left up to this stage are filled using inverse distanced weighted interpolation (3x3 window) and moving average (3x3 window) (Figure 5.7). The weighting used is $(D^{-2})/\sum(D^{-2})$, where D is Euclidian distance between grid pixels [Shepherd, 1968].

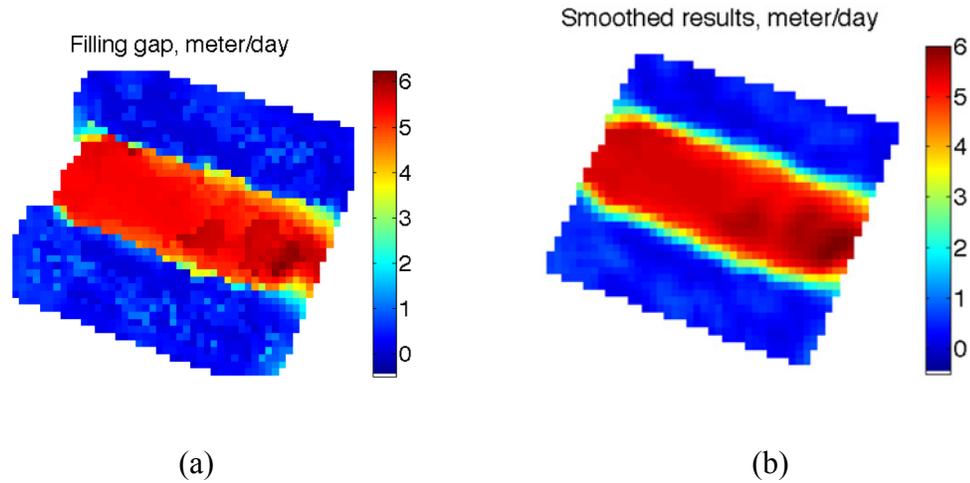


Figure 5.7. (a) Inverse distance weighted interpolation and (b) moving average

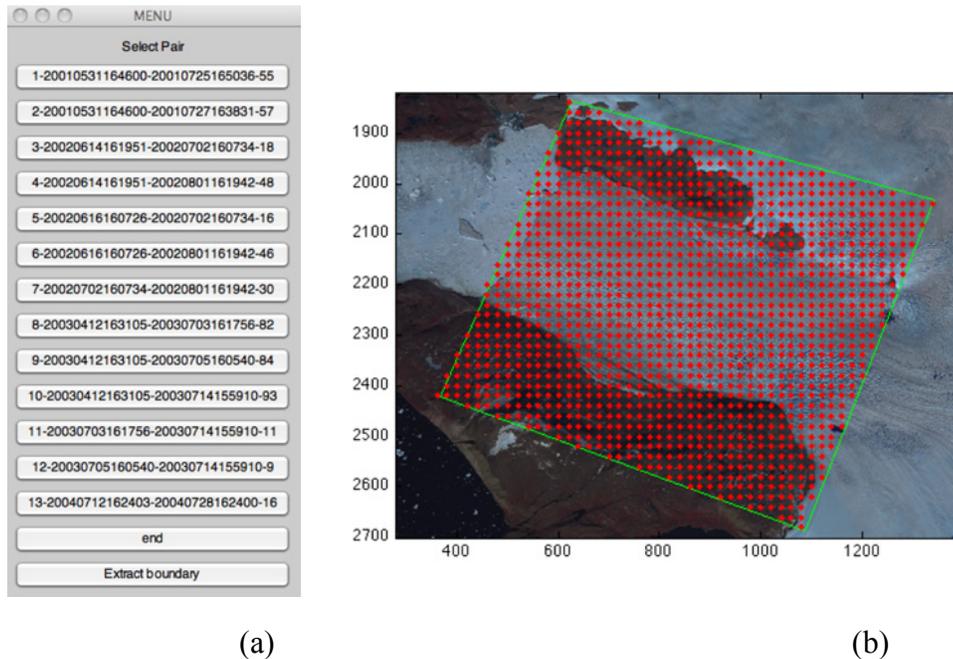


Figure 5.8(a) List of ASTER pair for velocity calculation in Alison Gl. Date, orbit and time separation are listed. (b) Grid posts (red points) generated within boundary (green line)

5.6 Example Applications and Analysis

Here we present results from application of MIMC to AST140TH image pairs for 5 outlet glaciers along Greenland's west coast: Alison glacier (07/02/2002-/08/01/2002, 30-day image separation), Igdlulik (07/02/02-08/01/02, 30-day image separation), Kong Oscar (07/08/2004-07/26/2004, 18-day image separation), and Upernavik North and South (06/12/2003-07/28/2003, 46-day image separation). We use ASTER (AST14DMO) orthorectified image products generated by the LPDAAC.

For each glacier, a region of interest (ROI) polygon is defined and appropriate pairs for velocity computation are automatically selected based on overlap with the ROI and the separation measured in days.

Figure 5.8a shows a pair list for Alison glacier (10~120 day separation). Once an image pair is on display, screen digitizing extracts boundary. Figure F5.8b illustrates grid posts generated. The post interval is 20 pixels (300 meter) and total number of posts is 1251. Land area is included for analysis on registration accuracy.

Comparison with conventional correlation matching

To gauge MIMC's performance, the results are compared with conventional DEFINE correlation matching results. Three bands (RGB) and three chip sizes (30x30, 40x40 and 50x50) are used, which totals 9 pairs of matches. Figure 5.9 shows displacement results with conventional correlation matching (30x30 Red band) and MIMC results for four glaciers.

5.7 Error Budget

Two major sources of uncertainty in displacement calculations are: 1) image

matching ambiguity and 2) registration error. Dietrich et al [2007] analyzed the

uncertainty in matching based on steepness of parabola fitting into correlation, estimating

less than 0.1 pixel for stationary reference targets and 0.1-0.3 pixel for glacier features.

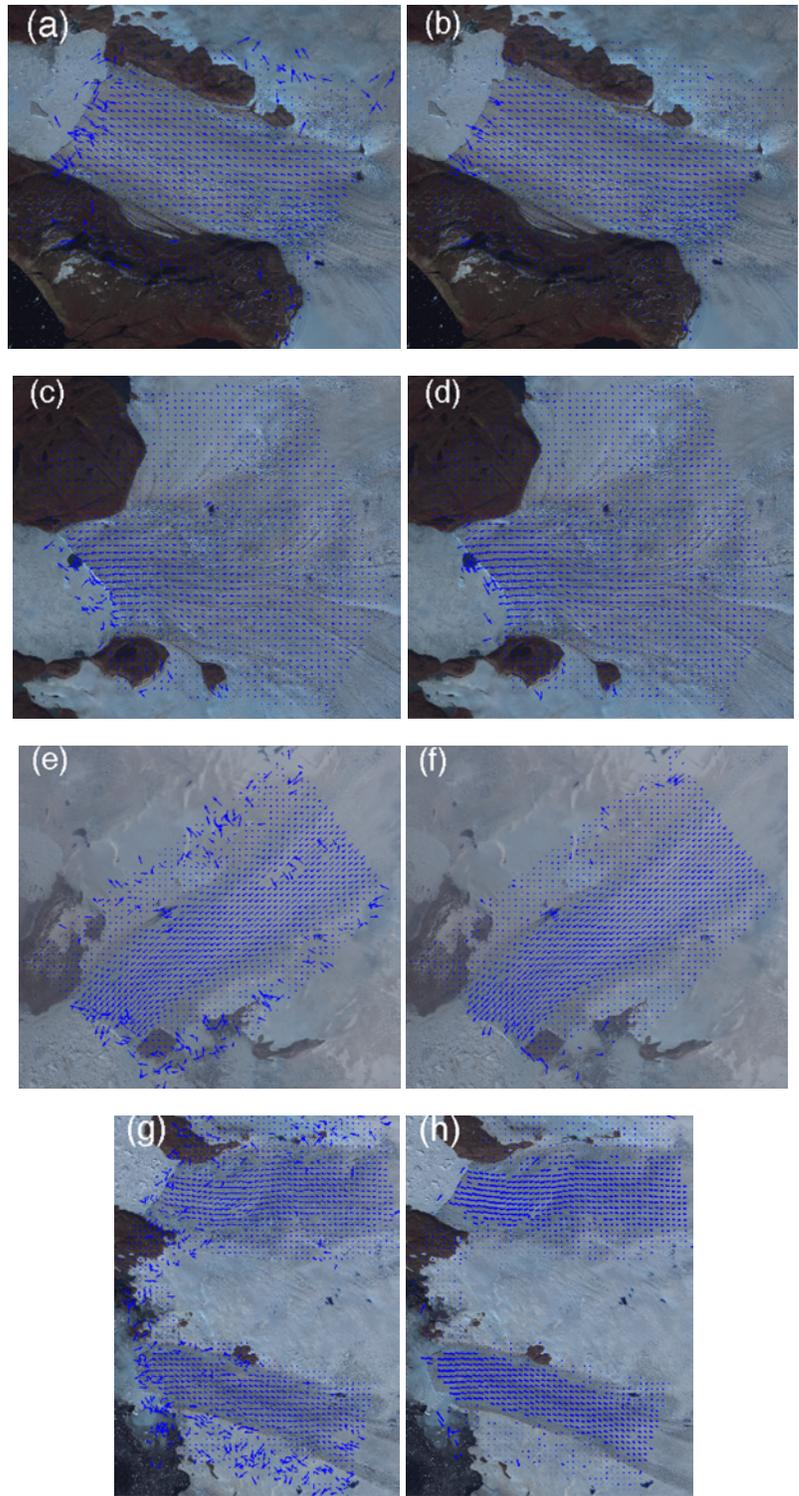


Figure 5.9 Red band 30x30 reference chip size displacement result and MIMC displacement results without any filtering for Alison (a, b), Igdlulik (c d), KongOscar (e, f), Upernavik (g, h) Glaciers

Skvarca et al. [2003] conservatively estimates 0.5 pixels for correlation matching and 1.5

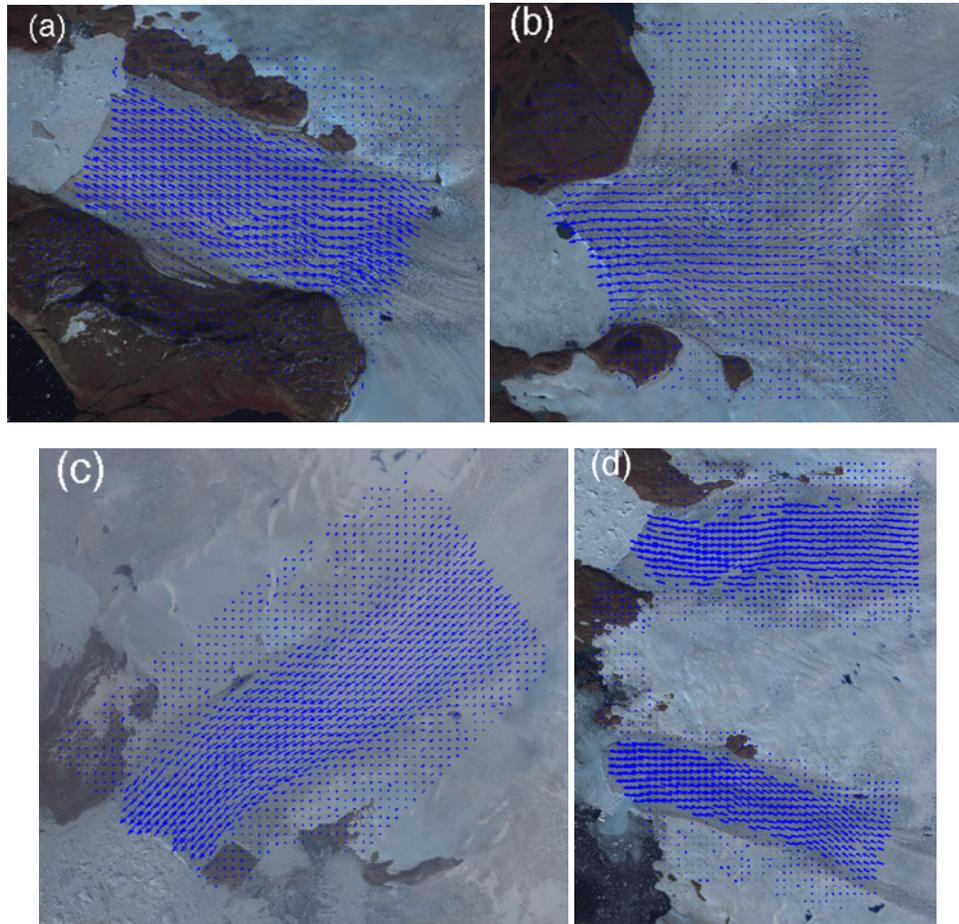


Figure 5.10. Final displacement results for (a) Alison, (b) Igdlulik, (c) KongOscar and (d) Upernavik glacier

pixel uncertainty in the registration process. The correlation matching uncertainty is difficult to quantify, since quadratic surface fitting uses solely values around peaks in correlation (3x3 or 5x5 neighbor). However, based on comparisons with off-ice targets, 0.3 pixels uncertainty can be considered as conservative estimate. Registration error results from the distribution of control points and the transformation model that is used. In this study, control points are measured on ice-free areas near the glacier channel and a simple x-y offset transformation model is used to remove registration error. Following this offset correction, off-ice displacements are typically less than 1 pixel. Matching ambiguity and registration errors therefore account for 1 pixel of displacement error. This error gives a speed uncertainty of about 1 m/day for a 15-m pixel resolution and the nominal ASTER repeat interval of 16 days.

Since errors in registration are typically much greater than the matching error, the registration method used in this study is applicable for almost all glaciers when even a small amount of ice-free land is visible and individual image pairs are only a few 10's of

km across. If the image pair covers a large area, however, the x-y offset approach used here might not provide a global registration. Higher order registration correction may need to be implemented considering the area of interest and the distribution of land areas.

5.8 ETM+ Scan Line Corrector Off Mode (SLC-off) Feature Tracking

For this project, we introduce a method of feature tracking using Landsat SLC-off data. For this method, we use a similar MIMC procedure as described above but apply the cross-correlation in the spatial rather than FFT domain. In this method, 5 pixels on either side of each stripe are masked and all striped regions are set to 0 value. We also use a ~50% larger search/reference chip size than for the FFT method.

The performance of the SLC-off feature tracking is assessed by comparing SLC-on estimates with displacements estimates from the same with simulated SLC-off striping applied to the same images. First, we applied our normal MIMC algorithms to an SLC-on pair from Kangerdlussuaq glacier (April 5th and 14th, 2001; Figure 5.11). Using these

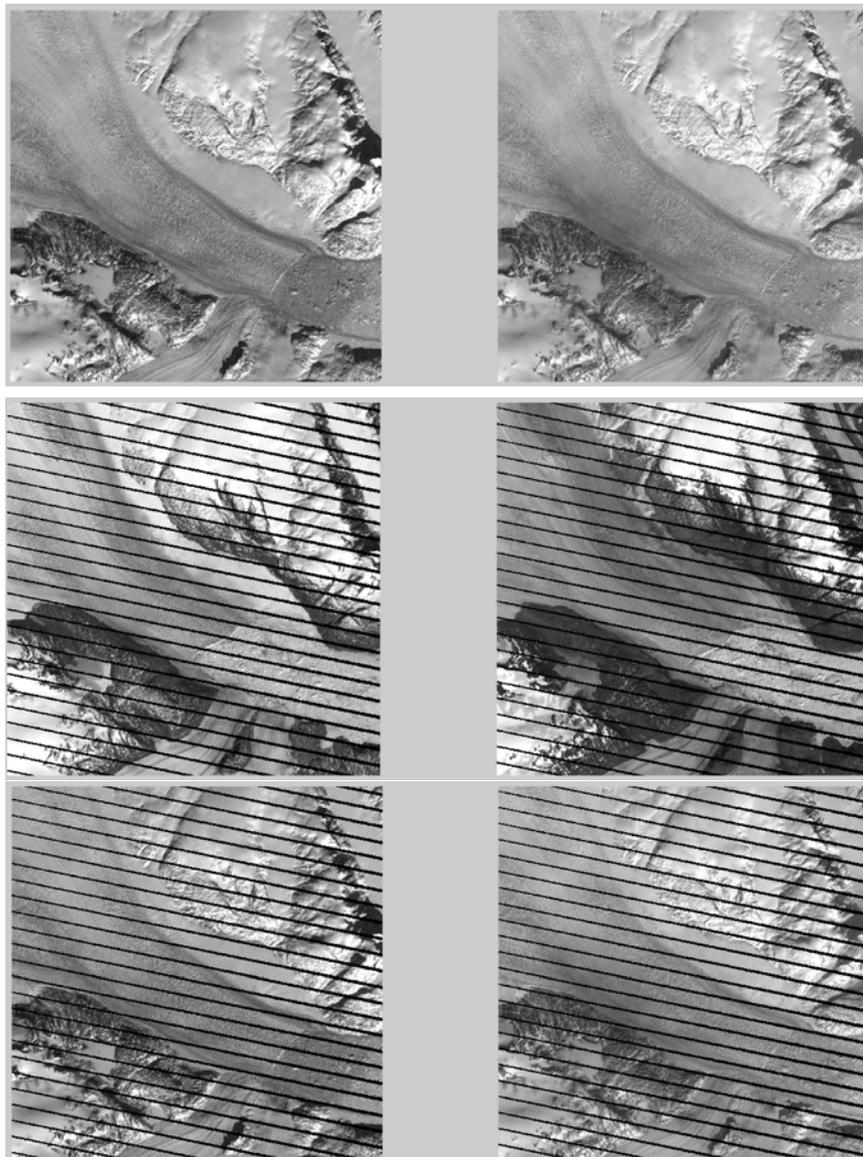


Figure 5.11. (top) Test SLC-on image from April 5th and 14th, 2001. (middle) SLC-off image pair with from the same pair from July 7, 2003 and August 1, 2003. (bottom) 2001 test pair with SLC-off striping applied.

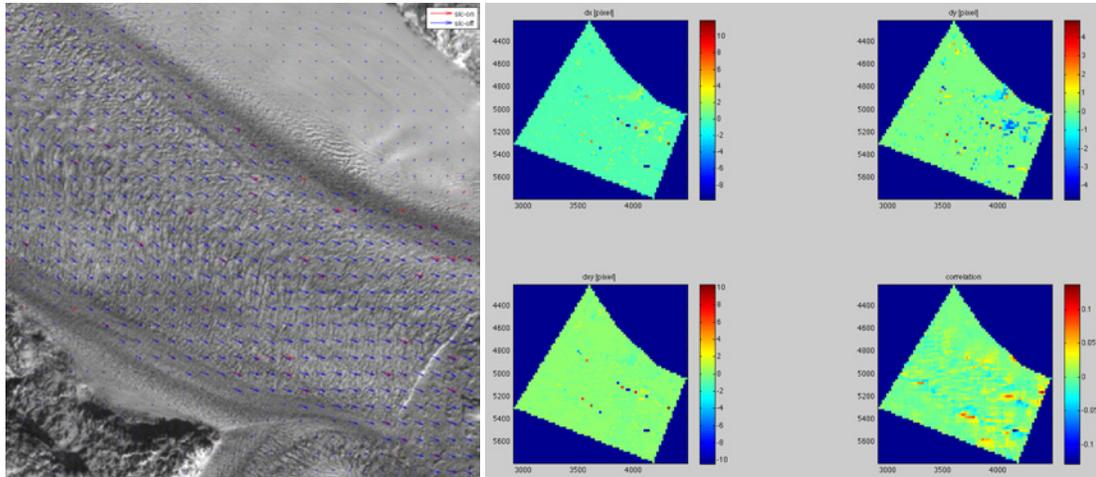


Figure 5.12. (left)(blue arrows) Simulated SLC-off velocity solutions for the test pair plotted with (red arrows) the SLC-on solutions. (right) Difference map of x and y displacements and correlations between simulated SLC-off and SLC-on runs.

same images, we created a synthetic SLC-off pair to which we applied our SLC-off algorithm. To simulate the SLC-off striping, we extracted and applied the stripe masks to the data for an SLC-off pair from the same path/row and the test images (July 7, 2003 and August 1, 2003).

The results of the SLC-on and simulated SLC-off tests are compared in Figure 5.12. Overall, the difference between the results is small, with most of the glacier exhibiting little or no difference in velocity magnitude or direction between test cases. Differences tend to occur at shear margins, where large spatial gradients in velocity produce aliasing of the striped data. These results suggest that spatial-domain feature tracking using the MIMC procedure is able to resolve glacier displacement from SLC-off imagery at a similar accuracy and coverage as SLC-on data. Spatial domain processing is much more computationally expensive than FFT processing, but still practical for typical multi-processor computers.

6 Landsat 7 ETM+/RADARSAT-1 Image Mosaic

Our first objective is to assemble an ice-sheet wide imagery mosaic to be used for mapping and land surface classification at the highest possible spatial resolution and within as narrow a time window as possible to enable change detection. South of $\sim 81.2^{\circ}\text{N}$, we use Landsat 7 ETM+ imagery orthorectified and distributed by the U.S. Geological Survey (<http://glovis.usgs.gov/>). Using 1 August 2000 as a target date, we selected imagery from July and August, as close in time as possible to 1 August for the years, in preferential order, of 2000, 1999, 2001, and 2002. All imagery were automatically filtered for clouds using the algorithm presented in Luo et al. [2008], adapted to Landsat by Oreopoulos et al. [2011], and visually inspected for quality. In some cases additional manual cloud masking was required. In order to increase the consistency of the grayscale between images, each digital number image was converted

to reflectance, including corrections for sun angle and distance using the parameters provided in the metadata. Multi-spectral bands 1 through 4 were pan-sharpened to 15-m posting using band 8 and a simple and fast additive method in which the band 8 image was down-sampled to 30 m and differenced from each multispectral band. The difference image was then up-sampled to 15-m using bilinear interpolation and added to the band 8 image.

The pan sharpened reflectance images were then re-gridded via cubic convolution and mosaiced to the reference grid. Where images overlapped, the pixel that was closest in time to the target date of 1 August 2000, was selected. No edge feathering was applied. The mosaiced images were then converted back to a byte precision digital number by linearly scaling the reflectance values to the global minimum and maximum for each band.

The USGS employs two levels of geo-registration processing for their imagery (see http://edcsns17.cr.usgs.gov/helpdocs/landsat/product_descriptions.html#nlaps_lpgs). First, Standard Terrain Correction (Level 1T) incorporates both ground control points and a DEM for terrain corrections. Geodetic accuracy depends on the accuracy of the ground control and the quality of the DEM and is better than 90 m. Imagery covering the periphery and margin of the ice sheet, where features are visible on the surface, are processed to L1T. For L1T imagery, the root-mean-square of the residual between the geo-location model and the ground control are provided in the imagery metadata and are typically on the order of several meters. Second, Systematic Correction (Level 1G) uses only the satellite ephemeris for geo-location, providing a 1σ geometric accuracy within 250 meters. Scenes over the featureless interior of the ice sheet are typically processed to L1G.

North of the maximum extent of Landsat we include synthetic aperture radar amplitude imagery mosaics acquired between October and December 2000 by the RADARSAT-1 satellite. These data were produced by the Applied Physics Lab at the University of Washington as part of GIMP [Joughin *et al.*, 2016]. The data are distributed at 20-m resolution and were up-sampled through bilinear interpolation to 15-m to match the resolution of Landsat band-8. We merged the RADARSAT and Landsat band-8 imagery by applying a stretch to the RADARSAT image so that the histograms of both datasets match where they overlap. As with Landsat, the primary source of geolocation error in the RADARSAT imagery is error in the DEM used for terrain correction and are similar in magnitude to the Landsat mosaic [Moon and Joughin, 2008].

The final image mosaic is distributed in tiles, with one image for each band, plus an index image in which each pixel gives the index number of its corresponding source image in an accompanying metadata file. The metadata file lists each Landsat scene identification number (scene ID) used in the mosaic for that tile, the acquisition time, and the root-mean-square control point registration error where available. The original scene ID, acquisition date and geo-location error for any pixel in an image can thus be obtained using the index image.

7 Land Classification Masks

Land classification masks are needed for co-registration of repeat imagery and elevation data, as ice surfaces can change with time while areas of exposed bedrock provide control. Further, the accurate delineation of ice boundaries provides a benchmark for measuring future ice margin changes. Landsat-7 ETM+ data are commonly used for mapping snow and ice, either manually, by tracing the margin with a computer mouse directly on the imagery, or automatically, from multi-spectral classification techniques [Rastner *et al.*, 2012]. Automatic methods are far more efficient and are effective for ice and snow that is free of surface debris. However, the drawbacks of automated, multi-spectral classification methods are that (1) they cannot differentiate between seasonal/ephemeral snow cover and glacial ice, (2) they fail at marine margins when dense packs of icebergs and sea ice are present, (3) much of the marginal ice of the Greenland Ice Sheet and surrounding glaciers is debris covered and (4) Landsat does not cover the most northern regions of the ice sheet. For these reasons, we abandoned multi-spectral mapping methods in favor of manual digitization of the panchromatic and pan-sharpened multispectral image mosaic presented in **Sect. 3**. Even with manual methods, the ice margin can be difficult to locate visually in areas of abundant debris and snow cover. Margins of debris-covered ice were identified by breaks in surface slope, emerging melt water streams, color differences and the presence of small melt water ponds typical of debris-covered glaciers. Similarly, glaciers were differentiated from perennial snowfields by visible crevassing, surface moraines, and the existence of a visible toe. Snowfields without these features were not classified as glaciers. Using the same method, we also digitized the coastline to produce an ocean mask, with the null of the ice and ocean masks being ice-free terrain (including freshwater lakes).

Uncertainty in these classification masks arise from three sources of error: (1) image pixel resolution, (2) image geo-registration and (3) erroneous selection or non-selection of pixels (i.e. mapping error). All error sources are expected to vary randomly in space, although there is likely a systematic component of error source (2) over distances equivalent to the size of a single image (e.g. 185 km for Landsat 7) due to errors in the registration model used to orthorectify the image, which typically is on the order of ± 5 m, or $1/3$ of a pixel for L1T-processed imagery.

Error source (1) contributes a random error of 1 pixel for each ice boundary pixel. The position of any point of the ice margin has an uncertainty of 21 m while the total error for a given area of ice is then $(8N)^{1/2}x^2$, where N is the number of boundary pixels and x is the pixel posting in meters.

Erroneous selection or non-selection of pixels can be due to debris cover, shadows, and misidentification by the operator, as well as the ambiguity of delineating an ice boundary at glacier fronts ending in packs of icebergs. Without ground control, delineation of the ice edge in areas of debris cover, terminal moraines and persistent snow cover is subjective. These errors are difficult to quantify. We estimated uncertainties due to ambiguity in the ice edge and operator error by comparing mappings done by three different operators over the same area. On average, each operator identified 24.21 km (1614 pixels) of ice margin over the common area, with a 660 m (44 pixels)

difference between the maximum and minimum mappings, giving an estimated error of $\pm 3\%$, which is similar to other comparisons [Paul *et al.*, 2013]. This error, however, is expected to vary widely by particular location and size of area considered.

Initial versions of the GIMP classification mask have been used and analyzed in two studies. Rastner *et al.* [2012] compared the version 1.1 GIMP classification to their own, semi-automated delineation of peripheral glaciers and ice caps, which also utilized Landsat 7 data. They found an overall difference in classified area of 6%. This difference was mostly due to misclassification of debris-covered margin in GIMP. That study incorporated the GIMP classification into their dataset for far northern regions, and their combined map has been included in the global Randolph Glacier Inventory [Pfeffer *et al.*, 2014]. Citterio and Ahlstrom [2013] compared the version 1.2 GIMP classification to glacier outlines mapped from aerial photography in the 1980's and were able to measure local changes in margin positions between the datasets. They also detected some classification errors. Errors detected in both of these studies have been corrected in the current version 2.0 of the mask, along with additional quality control by our team. Both the ice and ocean classification masks were used in the production of the Digital Elevation Model, described next.

8 Digital Elevation Model

Here we describe the GIMP version 2 DEM. For completeness, we include details of the version 1 DEM in an appendix.

9.1 Source Data and Quality Control

The primary data source are submeter-resolution, panchromatic stereoscopic imagery collected by the GeoEye-1 and WorldView 1,2 and 3 satellites operated by DigitalGlobe Inc. These data are distributed by the National Geospatial Intelligence Agency (NGA) and archived at the Polar Geospatial Center at the University of Minnesota through the NextView License program. Source imagery are subjected to user restrictions, but derived products, including Digital Elevation Models (DEMs), may be openly distributed.

The DEMs used here were extracted solely from in-track stereoscopic imagery, in which both images of the stereo pair were collected minutes apart along the same orbital pass. Imagery are collected in strips 12 to 17-km wide (i.e. the swath width) and up to 300 km long. The strips are segmented into scenes with $\sim 20\%$ overlap for distribution. Individual, overlapping scenes are then paired for stereo processing. We used the Surface Extraction from TIN-based Search-space Minimization software for producing DEMs (SETSM, Noh and Howat, 2015) on High Performance Computing (HPC) systems at the Ohio Supercomputer Center, University of California San Diego (Gordon) and the National Center for Supercomputing Applications (Blue Waters) using multiple OSU and NSF-supported allocations. Raster DEMs with a resolution of 8m were extracted and filtered based on the density of point matches, which is designed to remove erroneous surfaces due to clouds, waters, severe shadowing and other sources. The filtered scenes

were then mosaicked back into strips using the iterative slope regression method of Nuth and Kaab [2011]. DEM strips were then visually quality-controlled and scored on a scale of 1 through 5, with 1 given the highest quality. DEMs scored a 4 or 5 were not used and those scored a three were manually masked to remove erroneous surfaces. Most commonly, DEM strips scored a 3 were those in which a portion was cloudy but the majority was clear.

9.2 Mosaicking and Data Format

Quality controlled strips were then mosaicked using the Nuth and Kaab [2011] iterative slope regression method for coregistration of DEM's and registration to LiDAR altimetry, which was the same method used for the GIMP DEM v1. The GIMP DEMv2 is distributed in two versions of mosaic: **v2reg** and **v2fit**.

The objective of the **v2reg** mosaic is to provide the best estimate of elevation and its uncertainty at the time of data acquisition at each pixel, for the purpose of elevation change measurements. The edges of registered strips in the **v2reg** mosaics are not feathered or smoothed so that surface discontinuities will be present in areas of change. These mosaics are provided as both quarterly time series and a merged "best data" mosaic. The date of each pixel, given as the day since Y2K, and the 1-sigma error are provided as supplementary fields.

The objective of the **v2fit** mosaic is to provide the single best continuous surface for applications that use slope information, such as for supraglacial stream flow modeling, and for visual displays, at the cost of a less confident elevation in areas of change. These mosaics are produced by selecting adjoining strips that yield the best alignment and by feathering edges of strips through distance-weighted averaging. The complete mosaic is then registered to all available Operation IceBridge data. The **v2fit** mosaic is provided as a single mosaic using all available DEM data.

The procedure for **v2reg** mosaicking is as follows: All DEM strips are individually registered to Operation IceBridge Airborne Topographic Mapper (ATM) and Land, Vegetation and Ice Sensor (LVIS) LiDAR point cloud data. We use all years of data (beginning in 1993) for bare rock areas and data collected within 30 days of image acquisition for ice covered areas, as classified by the GIMP Land Classification Products. Strips are then indexed by quarter of the year of acquisition, using the month ranges in **Table 1**. For each quarterly mosaic, registered strips acquired during that quarter added to the mosaic sequentially, in order of ascending quality rank and registration error, with the best data remaining. Strips with 1-sigma registration errors of larger than 4 m are excluded and treated as unregistered strips. Registered strips are added in "underprint" mode, in which only empty pixels are filled, without data averaging, edge feathering or smoothing, so that each pixel provides the best estimate of elevation at that date. Once all registered strips have been added, unregistered strips from the same quarterly acquisition period are added to fill gaps through co-registration to registered data in the mosaic. Unregistered strips are added in order of quality rank and lowest coregistration error. Following coregistration, the edges of unregistered strips are warped using a linear weighted adjustment from the edge of overlap to 500 pixels into the added strip. The error at pixels in the unregistered strip is the root-sum of squares of the mean error of the registered strips used as coregistration reference and the coregistration error.

Once all strips with overlap to existing, registered data have been added through coregistration, remaining gaps are filled with available unregistered data. The strip with the newest data coverage is first added and strips with overlap to that initial “anchor” strip are then coregistered and added using the same procedure as above. Errors for unregistered clusters of data are assigned “inf”.

A single v2reg “best data” reference mosaic is created from the series of quarterly mosaics by selecting pixel with the lowest registration error. Following the creation of this mosaic, areas with unregistered ice-free areas in the quarterly mosaics are coregistered to the reference mosaic.

The procedure for **v2fit** mosaicking is as follows: The mosaic is initialized by adding the quality level 1 strip with the most spatial coverage to the mosaic grid. The overlapping strip with the best coregistration fit (lowest residual after fit) is then added if it contributes 100 or more pixels of new data to the mosaic. A linear, inverse-distance feathering is applied to the overlapping region of the new and mosaicked data. Strips are added sequentially in this way until no additional overlapping strips with adequate new data remain. If additional, non-overlapping strips exist, another cluster of coregistered data is initialized by again adding the highest quality rank strip with the most coverage and then adding additional overlapping data as before.

Once the mosaic is built, each cluster of coregistered data is registered to Operation IceBridge Airborne Topographic Mapper (ATM) and Land, Vegetation and Ice Sensor (LVIS) LiDAR point cloud data. We use all years of data (beginning in 1993) for bare rock areas and data collected between 2010 and 2015 for ice covered areas, as classified by the GIMP Land Classification Products. The error reported in *_v2fit_sigma.tif is the standard error of the fit to the LiDAR point cloud. If no registration data is available, the error value is NaN.

Following registration, the GIMP ocean masks are applied and the DEM is merged with the GIMP v1 DEM to fill data gaps. The data are merged by interpolating the difference between the GIMP v1 and v2 DEMs at the boundaries of gaps, over the area of the gap and adding that difference to the v1 data. The adjusted v1 data is then added to the pixel with missing data values.

9 MODIS Mosaics

We create high-resolution MODIS-based mosaic images using an image ‘stacking’ technique that provides increased spatial and radiometric resolution [*Scambos et al.*, 1999; 2007]. In this procedure, we combine a series of MODIS band 1 images (250 m resolution) having similar illumination geometry (solar elevation and azimuth) covering the entire ice sheet, and project them onto an over-sampled grid, typically a polar stereographic at 100 m gridding scale. High-pass filtering at a coarse scale (160 km ground spatial scale) creates images with near-uniform histograms that are free from earth curvature effects on illumination. Clouds and other image defects are masked, and the edges of these masks are ‘feathered’. A separate image stack of the number of scenes compiled for each final-image grid cell is also created, with fractional values applied to

the ‘feather edges’ of the masks. A sum of the image stack is created, and divided by the image ‘count’ grid, creating a smooth and uniform-contrast image of the ice sheet surface.

This image ‘super-resolution’ technique results in merged multi-scene images with spatial resolution generally 0.7 to 0.5 the original pixel size, for regions of the mosaic with 3 to 10 images [Scambos *et al.*, 1999]. Further, image radiometric resolution is improved by a factor close to the square root of the number of images contributing to a region. Re-stretching the images can reveal subtle surface flow or other structures. We plan to create monthly composite images of 5 to 15 scenes.

10 Ice Fronts

Using the ortho-rectified SAR mosaics, supplemented as needed with optical imagery, we will hand digitize the ice fronts of the roughly 200 glaciers with greater than 2-km. The methods for accomplishing this are described by Moon and Joughin [Moon and Joughin, 2008].

11 Appendix: GIMP Version 1 Digital Elevation Model

The quality of data over most of the Greenland Ice Sheet in global elevation datasets, such as GTOPO30 and the more recent GDEM, is too poor to be of use for glaciological applications. The standard DEM used in glaciological studies was created from a combination of satellite radar altimeter and aerial Photogrammetry [Bamber *et al.*, 2001] with a posting of up to 1-km. This DEM was enhanced to 625 m posting through photogrammetry by Scambos and Haran [2002]. While these DEMs are accurate to a few meters over the relatively flat interior of the ice sheet, they have poor resolution over the steeper margins and higher-relief periphery.

Our objective is to enhance DEM resolution and accuracy, particularly over the ice sheet margin and periphery, by integrating high-quality photogrammetric topography data into the existing low-resolution DEM and registering the DEM to elevations acquired by the Geoscience Laser Altimeter System (GLAS) aboard the Ice and Cloud Elevation Satellite (ICESat). Our approach follows the schematic shown in **Fig 3**. We focus on generating a continuous surface and we ignore temporal changes in ice elevation, which are over 100 m near the fronts of some rapidly retreating glaciers, and produce a DEM that approximates the mean elevation over the ICESat era (2003-2009). We first present each input dataset and then describe the procedure for merging them, followed by a description of errors and artifacts in the resulting DEM.

11.1 ICESat GLAS

All data are referenced to elevations obtained from by ICESAT GLAS between 2003 and 2009. We use the 633 products of the GLA12 release corrected for time-varying elevation biases, as estimated based on apparent variation of the mean-sea-surface height [Shepherd *et al.*, 2012]. Poor-quality returns were removed using

techniques developed for elevation-change estimation that identify the best-quality returns based on parameters that describe the shape and amplitude of the returned laser pulse [Shepherd *et al.*, 2012]. Elevations were corrected for detector saturation, and the time-varying bias correction should remove offsets associated with campaign-to-campaign variations in the shape of the transmitted pulse [Borsa *et al.*, 2014]. Elevations calculated in this way should be accurate to better than 0.1 m, or two orders of magnitude smaller than the expected DEM uncertainty.

11.2 Photo-Enhanced Bamber (PEB) DEM

The most widely used DEM for the entire ice sheet is that presented in *Bamber et al.* [2001], created from a combination of radar altimeter and stereo-photogrammetric data from the mid 1990's. These data were validated against airborne altimeter data, also from the mid-1990's, with a reported, ice-sheet wide 1σ error of ± 7 m and errors of several hundred meters at the coasts. This DEM was subsequently enhanced through photogrammetry with AVHRR imagery [Scambos and Haran, 2002], hereafter referred to as the Photo-Enhanced Bamber or PEB DEM which improved the effective spatial resolution and accuracy of the DEM by $\sim 30\%$, so that 1σ errors in the ice sheet interior, where slopes are $\sim 10^{-3}$ are ± 2 m. Errors in the marginal areas were equivalent to the original *Bamber et al.* [2001] DEM.

The PEB DEM was provided by the NSIDC in a spherical Lambert azimuthal projection at a posting of approximately 627 m. These data were re-gridded to EPSG 3413 and up-sampled to 30-m posting using bilinear interpolation. The re-gridded data were then co-registered to the ICESat GLAS point cloud using an iterative, 3-D conformal transformation [Noh and Howat, 2014]. This procedure results in residuals between the DEM surface and ICESat point cloud with a normal distribution and a mean of zero. Co-registration was performed on 25 km by 25 km tiles with 5 km of overlap. The co-registered tiles were then mosaiced with linear distance-weighted edge feathering. The root mean square (RMS) of the residuals between the PEB DEM and the ICESat point cloud following co-registration are given in **Table 1**. The total RMS error of ± 21.8 m is nearly three times higher than reported by *Bamber et al.* [2001] and *Scambos and Haran* [2002], likely due to the more extensive sampling by ICESat relative to the airborne altimetry used in those studies, especially over ice-free terrain where errors are much higher. The RMS errors over the interior ice sheet are more consistent with reported errors.

11.3 GDEM V2

The Global Digital Elevation Model (GDEM) is a global, 30-m posted DEM produced by the Ministry of Economy, Trade, and Industry (METI) of Japan and the United States National Aeronautics and Space Administration (NASA) [Slater *et al.*, 2011]. The GDEM is created by average-stacking individual stereo-photogrammetric DEM's acquired by the Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) between 2000 and 2010. Following an initial release in 2009, Version 2 was released in October 2011. The GDEM is distributed in $1^\circ \times 1^\circ$ tiles in

geographic projection. The distribution includes metadata giving the number of individual AST14DEM granules that were stacked to obtain each posted elevation. No information, however, is given regarding which scenes were used, so the time period of elevation measurements cannot be determined directly.

GDEM data quality is poor over much of the ice sheet owing to low-contrast surfaces on snow and ice. Additionally, artifacts due to shadows, clouds and blunders in the automated matching algorithm are abundant over all terrains. Following re-projection and gridding of the GDEM Version 2 to the GIMP grid, we applied a pyramiding standard deviation filter in which the DEM is smoothed to progressively finer resolutions and differenced from the native-resolution DEM. Pixels with differences exceeding 2.5σ of the mean are discarded. Since ice-covered terrain is substantially smoother than ice-free terrain, we apply this filter separately to the two land classifications, using the land classification masks. Following automated filtering, we manually masked blunders visible on a hillshade image of the DEM. These procedures removed nearly all data from above 1600 m elevation, which is approximately the average mass balance equilibrium line altitude. Following filtering and masking, GDEM covers 30% of Greenland's total area, and respectively 92% and 19% of its total ice-free and ice-covered terrain.

11.4 SPIRIT DEM

Photogrammetrically-derived DEMs over Greenland were produced from images acquired in 2007 and 2008 as part of the SPOT-5 Stereoscopic Survey of Polar Ice: Reference Images and Topographies (SPIRIT) program. A description of dataset production and validation is given in *Korona et al.* (2009). The SPIRIT DEM is distributed in UTM projection and referenced to the EGM96 Geoid and posted at 40 m. Two versions of each DEM, processed with different correlation parameters, are provided, along with data quality and interpolation masks. *Korona et al.* (2009) reports a slightly better precision and accuracy of SPIRIT DEM ($< \pm 5$ m) over ASTER DEM's based on validation experiments with ICESat.

For this project, we obtained all available SPIRIT DEM products over Greenland. Each DEM was re-projected to EPSG 3413 and the WGS-84 ellipsoid and up-sampled to 30 m. As advised in *Korona et al.* [2009], we use version 2 of each DEM and mask out all interpolated pixels. We then applied the same filtering and masking procedure as used for the GDEM.

Each individual SPIRIT DEM was then co-registered to overlapping regions of the filtered GDEM using the 3-D conformal transformation [*Noh and Howat, 2014*]. This provided a consistent registration between the SPIRIT and GDEM datasets to facilitate merging. Each individual SPIRIT DEM was then stacked into a single mosaic by taking the median elevation at each pixel, keeping track of the number of individual measurements. The resulting filtered SPIRIT mosaic covers 10% of Greenland's total area, and respectively 24% and 8% of its total ice-free and ice-covered terrain. The most continuous coverage is along the southwestern and southern coasts, with approximately

50% of the land and ice area covered in each tile, or most of the land and ice area below 1500 m elevation.

11.5 CNES Mean Sea Surface Height

Stereo-photogrammetric methods typically cannot resolve open water surfaces due to the lack of features, so that these surfaces are usually interpolated from the shoreline. This and the presence of icebergs result in spurious sea surface heights in stereo-photogrammetric DEMs. To ensure correct sea surface heights, we apply the ocean mask to the final DEM and replace those ocean surfaces with the CLS11 mean sea surface height product from the Centre National d'Etudes Spatiales (CNES). The CNES CLS11 is the 16-year mean of TOPEX/POSEIDON, ERS 1&2, GFO, JASON-1, ENVISAT altimeter measurements gridded to 1/3 of a degree [Schaeffer *et al.*, 2012]. We re-project these data to EPSG 3413 and up-sample them to the 30 m GIMP grid using bilinear interpolation.

11.6 Data Merging

Following co-registration and stacking, the SPIRIT DEM mosaic was differenced from the GDEM and the differences were extrapolated across the grid using an inverse-distance interpolation. The extrapolated difference map was then added to the SPIRIT stack. The GDEM and SPIRIT DEMs were then merged under the following conditions at each pixel:

- (1) If there was a GDEM value, but no SPIRIT value, the pixel is assigned the GDEM value.
- (2) If there was a SPIRIT value, but no GDEM value, the pixel is assigned the corrected SPIRIT value.
- (3) If there were both GDEM and SPIRIT values, and the pixel is over ice-free terrain, the pixel is assigned GDEM value. This is due to the GDEM's higher spatial resolution.
- (4) If there were both GDEM and SPIRIT values, and the pixel is over ice, the pixel is assigned the average of the GDEM and SPOT values, weighted by the N number of observations, where N equals 1 for GDEM plus the number of individual SPIRIT DEMs used in the stack described in **Sect. 5.4**.

The merged GDEM and SPIRIT DEM (merged G&S) was then co-registered to the ICESat GLAS point cloud using the 3-D conformal transformation [Noh and Howat, 2014]. The RMS validation errors of the merged G&S DEM are given in **Table 1**. To

assess the improvement in validation score provided by the higher-resolution data, **Table 1** also gives the RMS errors for the PEB DEM exclusive to areas of overlap with the merged G&S DEM. On average, the merged G&S DEM improves validation score by a factor of 8 over the ICESat-registered PEB DEM.

To combine the merged G&S DEM and PEB DEM, the PEB DEM was first adjusted by differencing it from the merged G&S DEM and interpolating the differences across areas of no data in the merged G&S DEM. The difference was then added to the PEB DEM and the two DEM's were combined using the following rules at each pixel:

- (1) If there was a merged G&S DEM value, the pixel is assigned the merged G&S value.
- (2) If there was no merged G&S DEM value, the pixel is assigned the adjusted PEB DEM value.

An ocean mask (see Sect. 4.) is then applied and those pixels are replaced with the CLS11 sea surface heights, as described in Sect. 5.5. The final GIMP DEM thus provides an altimeter-registered, relief-enhanced version of the PEB DEM. The enhancement is most pronounced over regions of high relief on the margin and periphery of the ice sheet. Notably, whereas outlet glaciers are not clearly defined in the PEB DEM, the GIMP DEM resolves outlet glacier termini and fjord walls in detail.

11.7 Errors and Artifacts

The overall RMS of the differences between the GIMP DEM and ICESat elevation is ± 9.1 m, which is less than half that of the ICESat-registered PEB. The error on ice-free terrain (± 18.3 m) is over twice that of ice-covered terrain (± 8.5 m), which is to be expected considering the higher relief at the ice-free margin. We note that an unknown amount of this error can be attributed to differences in the geometries of the ICESat footprint, which has a typical diameter of 70 m, and the DEM pixels. The effect of this difference will increase with slope. Additionally, over ice, some amount of the validation error can be attributed to temporal variations in surface elevation, ranging from decimeters over the interior to 10's of meters over rapidly thinning outlet glaciers. Besides ice thinning, the advection of crevasses and other surficial expressions with ice flow contributes an unknown error. These validation errors should, therefore, be viewed as an upper bound for the true standard data error.

The largest validation errors exist for the most northern regions, for which little high-resolution data exist and coverage is mostly from the PEB DEM. Higher errors, exceeding ± 20 m, are also found in areas of extreme relief, such as the Geikie Peninsula (tiles 4-2 and 5-2), where gaps in high-resolution data coverage exist over steep mountain glaciers and icecaps.

RMS errors are the smallest for the flattest surfaces (e.g. the interior ice sheet), increasing with slope to a peak of ± 24 m at 2° . RMS error then decreases to ± 13 m for 5° slopes before increasing again. The peak in RMS error at 2° slope corresponds roughly

with the equilibrium line of the ice sheet and, thus, the boundary between the merged G&S and PEB DEM. Errors in both the PEB and merged G&S DEM's result in spurious, step-like transitions between the two. This effect results in the continuous zone of large errors running along the southeast margin, which is especially steep. A positive peak in mean and median errors of 2.1 m and 0.9 m, respectively, at 1.1° shows a positive bias in the GIMP DEM relative to ICESat over the area just inland of the snowline. Over steeper terrain, this bias becomes increasingly negative (i.e. the GIMP DEM reports increasingly lower elevations than ICESat), from -0.5 m at 5° to -1.5 at 25°. Since most of the coverage of surfaces with 1° slopes are from the PEB DEM, the positive bias could be explained by either slope-dependent errors in the PEB or thinning between the PEB and ICESat epochs, with neither effect completely compensated during co-registration. The cause of the negative bias over steeper terrains is unknown. Since these biases are spatially variable and are small (< 10%) relative to the random error, we do not correct for them.

Where merged G&S coverage exists above the snow line, the apparent surface is much rougher, with pitting resulting from blunders in the surface matching procedure used to generate the DEMs. These roughness features typically have amplitudes of several meters.

Rapid ice thinning and front retreat also cause DEM artifacts. Many fast-moving outlet glaciers thinned by 10's of meters, reaching over 100 m in some cases, during the data collection period. This thinning causes offsets between DEM surfaces acquired at different times and, when stacked, can result in spurious offsets and discontinuities in the surface. Additionally, ice-front retreat between date of the imagery used in construction of the ice cover mask and DEM data acquisition causes incomplete masking of the ocean boundary. For outlet glaciers, this often means that areas of dense icebergs remain in the DEM.

12 References

- Abrams, M., S. Hook, and B. Ramachandran (2002), *ASTER user handbook, v2: Advanced spaceborne thermal emission and reflection radiometer*, Jet Propulsion Laboratory.
- Ahn, Y., and I. M. Howat (2011), Efficient Automated Glacier Surface Velocity Measurement From Repeat Images Using Multi-Image/Multichip and Null Exclusion Feature Tracking, *Ieee T Geosci Remote*, 49(8), 2838–2846, doi:10.1109/TGRS.2011.2114891.
- Bamber, J. L., S. Ekholm, and W. B. Krabill (2001), A new, high-resolution digital elevation model of Greenland fully validated with airborne laser altimeter data, *J. Geophys. Res. Oceans*, 106(B4), 6733–6745, doi:10.1029/2000JB900365.
- Borsa, A. A., G. Moholdt, H. A. Fricker, and K. M. Brunt (2014), A range correction for ICESat and its potential impact on ice-sheet mass balance studies, *The Cryosphere*, 8(2), 345–357, doi:10.5194/tc-8-345-2014.

- Citterio, M., and A. P. Ahlstrøm (2013), Brief communication “The aerophotogrammetric map of Greenland ice masses,,” 7(2), 445–449, doi:10.5194/tc-7-445-2013.
- Curlander, J. C., and R. N. Mcdonough (1991), *Synthetic Aperture Radar: Systems and Signal Processing*, Wiley, New York.
- Dietrich, R., H. G. Maas, M. Baessler, A. Rülke, A. Richter, E. Schwalbe, and P. Westfeld (2007), Jakobshavn Isbræ, West Greenland: Flow velocities and tidal interaction of the front area from 2004 field observations, *J Geophys Res*, 112(F3), F03S21–, doi:10.1029/2006JF000601.
- El-Darymli, K., P. McGuire, E. Gill, D. Power, and C. Moloney (2014), Understanding the significance of radiometric calibration for synthetic aperture radar imagery, pp. 1–6, IEEE, Toronto.
- Fahnestock, M., T. Scambos, T. Moon, A. Gardner, T. Haran, and M. Klinger (2016), Rapid large-area mapping of ice flow using Landsat 8, *Remote Sens Environ*, 185, 84–94, doi:10.1016/j.rse.2015.11.023.
- Fallourd, R., E. Trouvé, C. A. Deledalle, and J. M. Nicolas (2010), Regularization of displacement fields for glacier monitoring, *geodesy.unr.edu*.
- Goldstein, R. (1995), Atmospheric Limitations to Repeat-Track Radar Interferometry, *Geophys Res Lett*, 22(18), 2517–2520, doi:10.1029/95GL02475.
- Goldstein, R. M., H. A. Zebker, and C. L. WERNER (1988), Satellite Radar Interferometry - Two-Dimensional Phase Unwrapping, *Radio Science*, 23(4), 713–720.
- Gonzales, R. C., and R. E. Woods (2002), *Digital Image Processing*.
- Gray, A. L., K. E. Mattar, and G. Sofko (2000), Influence of ionospheric electron density fluctuations on satellite radar interferometry - Gray - 2000 - Geophysical Research Letters - Wiley Online Library, *Geophys Res Lett*, doi:10.1029/2000GL000016/pdf.
- Gray, A. L., K. E. Mattar, P. W. Vachon, R. Bindschadler, K. C. Jezek, R. Forster, and J. P. Crawford (1998), InSAR results from the RADARSAT Antarctic Mapping Mission data: estimation of glacier motion using a simple registration procedure, vol. Geosci. and Rem. Sens. Proc., 1998, vol. 3, pp. 1638–1640, IEEE, Seattle.
- Howat, I. M., A. Negrete, and B. E. Smith (2014), The Greenland Ice Mapping Project (GIMP) land classification and surface elevation data sets, *The Cryosphere*, 8(4), 1509–1518, doi:10.5194/tc-8-1509-2014.
- Howat, I. M., I. Joughin, S. Tulaczyk, and S. Gogineni (2005), Rapid retreat and acceleration of Helheim Glacier, east Greenland, *Geophys Res Lett*, 32(22), L22502, doi:10.1029/2005GL024737.

- Jeong, S., and I. M. Howat (2015), Performance of Landsat 8 Operational Land Imager for mapping ice sheet velocity, *Remote Sens Environ*, 170, 90–101, doi:10.1016/j.rse.2015.08.023.
- Jeong, S., I. M. Howat, and Y. Ahn (2017), Improved Multiple Matching Method for Observing Glacier Motion With Repeat Image Feature Tracking, *Ieee T Geosci Remote*, 55(4), 2431–2441, doi:10.1109/TGRS.2016.2643699.
- Jezek, K. C. (1999), Glaciological properties of the Antarctic ice sheet from RADARSAT-1 synthetic aperture radar imagery, *Ann. Glaciol.*, 29(1), 286–290, doi:10.3189/172756499781820969.
- Jezek, K. C., K. Farness, R. Carande, X. Wu, and N. Labelle-Hamer (2003), RADARSAT 1 synthetic aperture radar observations of Antarctica: Modified Antarctic Mapping Mission, 2000, *Radio Science*, 38(4), –n/a, doi:10.1029/2002RS002643.
- Joughin, I. (1999), Tributaries of West Antarctic Ice Streams Revealed by RADARSAT Interferometry, *Science*, 286(5438), 283–286, doi:10.1126/science.286.5438.283.
- Joughin, I. R. (1995), Estimation of ice-sheet topography and motion using interferometric synthetic aperture radar, Seattle, March.
- Joughin, I. R., R. Kwok, and M. A. Fahnestock (1998), Interferometric estimation of three-dimensional ice-flow using ascending and descending passes, *Ieee T Geosci Remote*, 36(1), 25–37.
- Joughin, I., B. E. Smith, I. M. Howat, T. Moon, and T. A. Scambos (2016), A SAR record of early 21st century change in Greenland, *J Glaciol*, 62(231), 62–71, doi:10.1017/jog.2016.10.
- Joughin, I., B. E. Smith, I. M. Howat, T. Scambos, and T. Moon (2010), Greenland flow variability from ice-sheet-wide velocity mapping, *J Glaciol*, 56(197), 415–430.
- Joughin, I., R. Kwok, and M. Fahnestock (1996), Estimation of ice-sheet motion using satellite radar interferometry: Method and error analysis with application to Humboldt Glacier, Greenland, *J Glaciol*, 42(142), 564–575.
- Joughin, I., W. Abdalati, and M. Fahnestock (2004), Large fluctuations in speed on Greenland's Jakobshavn Isbræ glacier, *Nature*, 432(7017), 608–610, doi:10.1038/nature03130.
- Korona, J., E. Berthier, M. Bernard, F. Rémy, and E. Thouvenot (2009), SPIRIT. SPOT 5 stereoscopic survey of Polar Ice: Reference Images and Topographies during the fourth International Polar Year (2007–2009), *ISPRS Journal of Photogrammetry and Remote Sensing*, 64(2), 204–212.
- Li, S. (1993), *Development and Application of an Algorithm for a Two-Way*

Transformation Between the Geographical Coordinates and the ASF SAR Image Coordinates, Fairbanks.

- Luckman, A., T. Murray, R. de Lange, and E. Hanna (2006), Rapid and synchronous ice-dynamic changes in East Greenland, *Geophys Res Lett*, 33(3), L03503, doi:10.1059/2005GL025048.
- Luo, Y., A. P. Trishchenko, and K. V. Khlopenkov (2008), Developing clear-sky, cloud and cloud shadow mask for producing clear-sky composites at 250-meter spatial resolution for the seven MODIS land bands over Canada and North America, *Remote Sens Environ*, 112(12), 4167–4185, doi:10.1016/j.rse.2008.06.010.
- Moon, T., and I. Joughin (2008), Changes in ice front position on Greenland's outlet glaciers from 1992 to 2007, *J Geophys Res-Earth*, 113, F02022, doi:10.1029/2007JF000927.
- Moon, T., I. Joughin, B. Smith, and I. Howat (2012), 21st-Century evolution of Greenland outlet glacier velocities, *Science*, 336(6081), 576–578, doi:10.1126/science.1219985.
- Noh, M.-J., and I. M. Howat (2014), Automated Coregistration of Repeat Digital Elevation Models for Surface Elevation Change Measurement Using Geometric Constraints, *Ieee T Geosci Remote*, 52(4), 2247–2260, doi:10.1109/TGRS.2013.2258928.
- Nuth, C., and A. Kääb (2011), Co-registration and bias corrections of satellite elevation data sets for quantifying glacier thickness change, *The Cryosphere*, 5(1), 271–290, doi:10.5194/tc-5-271-2011.
- Oreopoulos, L., M. J. Wilson, and T. Varnai (2011), Implementation on Landsat Data of a Simple Cloud-Mask Algorithm Developed for MODIS Land Bands, *Geoscience and Remote Sensing Letters, IEEE*, 8(4), 597–601, doi:10.1109/LGRS.2010.2095409.
- Paul, F. et al. (2013), On the accuracy of glacier outlines derived from remote-sensing data, *Ann. Glaciol.*, 54(63), 171–182, doi:10.3189/2013AoG63A296.
- Pfeffer, W. T. et al. (2014), The Randolph Glacier Inventory: a globally complete inventory of glaciers, *J Glaciol*, 60(221), 537–552, doi:10.3189/2014JoG13J176.
- Rastner, P., T. Bolch, N. Moelg, H. Machguth, R. Le Bris, and F. Paul (2012), The first complete inventory of the local glaciers and ice caps on Greenland, *The Cryosphere*, 6(6), 1483–1495, doi:10.5194/tc-6-1483-2012.
- Rignot, E., and P. Kanagaratnam (2006), Changes in the velocity structure of the Greenland ice sheet, *Science*, 311(5763), 986–990, doi:10.1126/science.1121381.
- Scambos, T., and T. R. Haran (2002), An image-enhanced DEM of the Greenland ice

sheet, *Ann. Glaciol.*, 34, 291–298.

Scambos, T., G. Kvaran, and M. A. Fahnestock (1999), Improving AVHRR resolution through data cumulation for mapping polar ice sheets, *Remote Sens Environ*, 69(1), 56–66.

Scambos, T., M. Dutkiewicz, J. Wilson, and R. Bindshadler (1992), Application of image cross-correlation to the measurement of glacier velocity using satellite image data, *Remote Sens Environ*, 42(3), 177–186.

Scambos, T., T. M. Haran, M. A. Fahnestock, T. H. Painter, and J. Bohlander (2007), MODIS-based Mosaic of Antarctica (MOA) data sets: Continent-wide surface morphology and snow grain size, *Remote Sens Environ*, 111(2-3), 242–257, doi:10.1016/j.rse.2006.12.020.

Schaeffer, P., Y. Faugere, J. F. Legeais, A. Ollivier, T. Guinle, and N. Picot (2012), The CNES_CLS11 Global Mean Sea Surface Computed from 16 Years of Satellite Altimeter Data, *Marine Geodesy*, 35(sup1), 3–19, doi:10.1080/01490419.2012.718231.

Shepherd, A. et al. (2012), A reconciled estimate of ice-sheet mass balance, *Science*, 338(6111), 1183–1189, doi:10.1126/science.1228102.

Shepherd, D. (1968), A two-dimensional interpolation function for irregularly-spaced data, vol. ACM '68, pp. 517–524, New York.

Skvarca, P., B. Raup, and H. De Angelis (2003), Recent behaviour of Glaciar Upsala, a fast-flowing calving glacier in Lago Argentino, southern Patagonia, *Ann. Glaciol.*, 36(1), 184–188, doi:10.3189/172756403781816202.

Slater, J. A., B. Heady, G. Kroenung, W. Curtis, J. Haase, D. Hoegemann, C. Shockley, and K. Tracy (2011), Global Assessment of the New ASTER Global Digital Elevation Model, *Photogramm Eng Rem S*, 77(4), 335–349.

Srivastava, S. K., S. Cote, P. Le Dantec, R. K. Hawkins, and K. Murnaghan (2007), RADARSAT-1 calibration and image quality evolution to the extended mission, *Advances in Space Research*, 39(1), 7–12, doi:10.1016/j.asr.2006.02.027.

Thomas, R. H., W. Abdalati, E. Frederick, W. Krabill, S. Manizade, and K. Steffen (2003), Investigation of surface melting and dynamic thinning on Jakobshavn Isbrae, Greenland, *J Glaciol*, 49(165), 231–239.

Werner, C., U. Wegmuller, T. Strozzi, and A. Wiesmann (2000), Gamma SAR and interferometric processing software, edited by H. Sawaya-Lacoste, Gothenburg, Sweden.