

Crossover Analysis and Automated Layer-Tracking Assessment of the Extracted DEM of the Basal Topography of the Canadian Arctic Archipelago Ice-Cap

Mohanad Al-Ibadi¹, Jordan Sprick¹, Sravya Athinarapu¹, Victor Berger¹, Theresa Stumpf¹, John Paden¹, Carl Leuschen¹, Fernando Rodríguez¹, Mingze Xu², David Crandall², Geoffrey Fox², David Burgess³, Martin Sharp⁴, Luke Copland⁵, and Wesley Van Wychen⁵

¹Center for Remote Sensing of Ice Sheets, University of Kansas, USA, ²School of Computing and Informatics, Indiana University, USA, ³Geological Survey of Canada, ⁴Dept. of Earth and Atmospheric Sciences, University of Alberta, Canada, ⁵Dept. of Geography, Environment and Geomatics, University of Ottawa, Canada

Abstract—In 2014, as part of the NASA Operation IceBridge project, the Center for Remote Sensing of Ice Sheets operated a multi-beam synthetic aperture radar depth sounder/imager over the Canadian Arctic Archipelago (CAA) to generate digital elevation models (DEMs) of the glacial basal topography. In this work, we briefly describe the processing steps that led to the generation of these DEMs, algorithm improvements over previously published results, and assess the results from two different perspectives. First, we evaluate the self-consistency of the DEMs where flight paths cross over each other and two measurements are made at the same location. Secondly, we compare the quality of the outputs of the ice-bottom tracker before and after applying manual corrections to the tracker results; the tracker is an algorithm that we implemented to automatically track the ice-bottom. Even though the CAA ice-caps are mountainous areas, where the scenes often have ice and no ice regions, which makes the imaging complicated, the statistical results show good tracking performance and a good match between the overlapped DEMs, where the mean error of the crossover DEMs is 37 ± 9 m.

Keywords— Synthetic aperture radar imaging, SAR, ice, ice-bottom tracking, tomography, DEM

I. INTRODUCTION

The Canadian Arctic Archipelago (CAA) contains one-third of the global volume of land ice outside the ice sheets, but its contribution to sea-level change is largely unknown. A recent mass balance estimate indicated a loss of 61 ± 7 (Gt/yr) of ice, contributing 0.17 ± 0.02 mm/yr to sea-level rise [1]. The present work is to generate DEMs of the ice bottom to improve estimates of the present sea-level contribution from the CAA ice-caps and to supply boundary conditions to ice dynamics models that will enable understanding of the basal controls on the glaciers in order to predict future contributions to sea-level [2].

In this work, an airborne Multichannel Coherent Radar Depth Sounder (MCoRDS), which was developed at the University of Kansas [3], was used to collect data from the Canadian Arctic Archipelago (CAA) as part of the NASA Operation IceBridge 2014 arctic campaign. The radar was equipped with a 15 sensor array in the cross-track, where the center 7 sensors were used for transmit and all 15 sensors each received and recorded an independent channel of data. Here, we only used the data from

the center 7 sensors, which are mounted on the fuselage. The radar was operated in a time multiplexed multi-beam mode, where the first beam looked left (-30°) of nadir, then at nadir, and then right (30°) of nadir as shown in Fig. 1a. With a transmit Hanning window, most of the energy is contained in a 30° beamwidth. The transmitted up-chirp was centered on 195 MHz, had a 30 MHz bandwidth, and a pulse duration of $3 \mu\text{s}$. The corresponding range-resolution is about 5 m in ice. The azimuth SAR resolution is 2.5m. At a 1000m AGL flight altitude, the SAR aperture is about 480m for a surface target. The pulse repetition frequency (PRF) is 12 KHz with 13 presums for each beam for an effective recorded PRF of 307 Hz.

Three main radar processing steps were applied to form a 3D image of the scene. These are: pulse compression, synthetic aperture radar (SAR) processing, and array processing. These steps are detailed in Section 2.A. The results from these processing steps then go into an automatic layer tracker, which is an algorithm that automatically tracks the bottom (i.e. the bed) of the ice caps. This tracked layer, along with other inputs, were used to generate the digital elevation models (DEMs) of the ice-bottom. These are detailed in Section 2.B. In Section 3 we self-assess the generated DEMs by doing a crossover analysis, where the DEMs from crossing flight paths are compared. Also, we quantify the quality of the automated tracker by comparing the tracked layer before and after applying manual corrections by trained analysts. These results are discussed in Section 3, before we conclude our work in Section 4. In this work, our main contributions are: 1) generated the DEMs for all the data segments from the CAA (102 DEMs from 5 segments), where

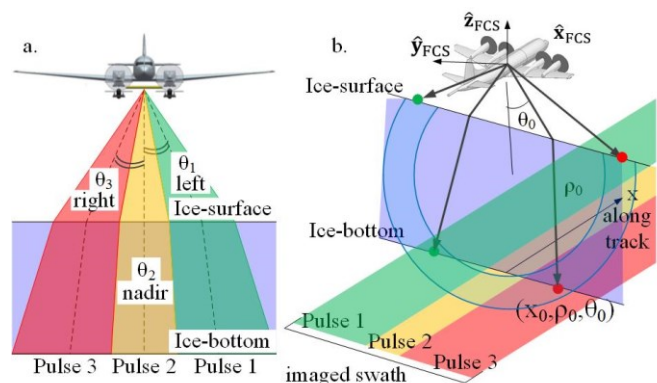


Fig. 1. a. Transmit configuration. b. SAR resolution.

we previously only had generated 7 DEMs and published the work in a separate paper, 2) performed crossover analysis to evaluate our results, and 3) modified the ice-layer tracker to better track the ice-surface and bottom and analyzed the new results.

II. ALGORITHM DESCRIPTION

A. Radar Processing and 3D Image Formation

Synthetic aperture radar (SAR) images are 2D images of the scene, where one axis represents the slow-time dimension (i.e. along-track) and the other axis represents the fast-time dimension (i.e. range-time). Pulse compression is used to resolve the targets in the range dimension by matched-filtering the received echoes (after being pre-conditioned). Then the data are focused in the along-track dimension using the frequency-wavenumber (f-k) migration algorithm. Each pixel in the SAR image contains direction of arrival (DoA) information from all targets at a specific along-track location (i.e. range-line) and a specific range-time (i.e. range bin). After these two steps we can view the scene as a thin sheet in the cross-track dimension, with a handful of unresolvable targets that lie in a constant range-time toroid around the radar, as shown in Fig. 1b. These targets can be resolved by estimating their DoAs using array-processing techniques, which is the third step in the 3D image formation process. Some other conditioning steps are also applied during these steps, such as motion compensation and receiver equalization.

In this work, the MULTiple Signal Classification (MUSIC) technique was used to estimate the directions of the echoes impinging on the array-antenna of the airborne radar [4]. We used MUSIC as a beamformer rather than an estimator (i.e. scan over DoAs from -90° to 90°) because the latter requires the exact number of sources or model number to be known (otherwise the tracker may track false targets or miss targets all together). Current efforts to estimate the model number using standard eigen-analysis of the data covariance matrices have failed due to a complicated eigenstructure that may be due to the time-bandwidth product of the array and multipath effects. The model number or assumed number of sources is two for each beam. In other words, the signal eigen-space for each beam is assumed to have a dimension of two. The output of the beamformer is a 3D image where the dimensions are along-track, range, and direction of arrival. The beamformer has the advantage that even when the signal eigen-space is not precisely estimated, there is still likely to be some reduction in the correlation between the actual source's steering vector and the null space of the signal eigen-space which can aid the ice bottom tracker even though it is not the steering vector with the lowest correlation due to errors in the signal eigen-space estimation. (In the absence of other information, the steering vector leading to the lowest correlation is usually taken to be the most likely DoA for the source.)

B. Layer-Tracking Algorithm and DEM generation

Since manual tracking of the ice-bottom is impractical on a large scale, we have implemented an (automated) sequential tree-reweighted (TRW-S) algorithm for extracting the ice-bottom surface as well as a browser to visualize the 3D images

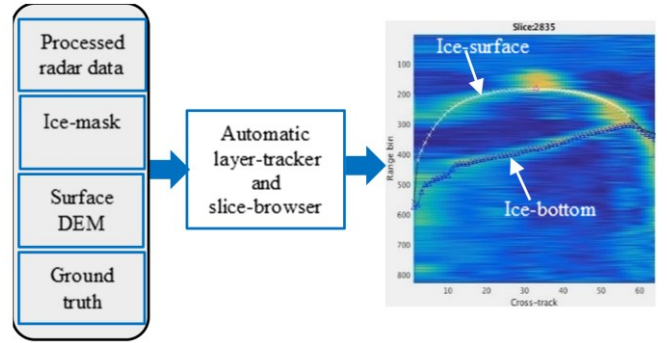


Fig 2. Input/output of the ice-layer tracker.

[5,7]. The word ‘layer’ here refers to the ice-surface or ice-bottom.

There are four main inputs to the layer-tracker (see Fig. 2):

1. The processed radar data is the 3D image where the ice-surface and ice-bottom need to be tracked.
2. The ice-mask [6] is a binary raster that is used to determine at each DoA whether there is ice or not. This is useful to force the automatic layer-tracker to alter the cost calculation accordingly and forces the ice-surface and ice-bottom to merge where there is no ice.
3. For each angle of incidence in the 3D image, the a priori surface DEM from ArcticDEM is used to find an estimate of the range-time to the ice-surface. Aside from providing the surface location, it is used to 1) extract image intensity properties of the surface scattering that are used to track the unknown ice-bottom and 2) to perform a rough calibration of the radar steering vectors by adjusting the radar-estimated DoA to match the surface-DEM-derived DoA. The averaged adjustment is then applied to the ice bottom for which there is no a priori information.
4. Ground truth points indicate where the ice-bottom layer should pass through. These are human labeled and only available for the nadir direction. The tracker does not assume these points are perfect, but the cost function is lower for layers that pass through them.

The tracker [7] was modified in several ways to improve its performance. The tracker optimization is based on a message passing algorithm in which each pixel in the image passes a “cost” message to its neighboring pixels to the left and right (direction of arrival dimension) and forward and backward (along-track dimension) in every iteration loop as shown in Fig. 3a. In [7], although messages are passed in all directions, in each of the two dimensions there is a preferential direction where the current iteration message is propagated while in the opposite direction the previous iteration message was propagated. This causes a strong bias towards the side of the image that the preferential direction starts from. For example, when the preferential direction is left to right, the left-most side of the image had a stronger effect on the result than all other columns because its message would be passed all the way across the image in a single iteration. [7] dealt with this by propagating messages from the left to right and then right to left

and from up to down and then down to up in each loop. The issue with this solution was that the most extreme directions of arrival (far-left and far-right), where the signal quality is worst, were being given too much influence. Since we have ground truth at nadir and the signal quality is often best at nadir, the preference direction was changed to be outward from nadir. So on the left side of the image, the preferential direction is toward the left and on the right side of the image, the preferential direction is toward the right (green arrows in Fig. 3). In this way, the nadir column asserts the greatest influence.

The cost message includes two components, a unary cost for each pixel that is independent of the layer and a binary cost which depends on how the layer changes from pixel to pixel (Fig. 3b). The unary term includes several components that we modified to improve the tracking. The first term is the template energy which was set by the mean squared difference between a template peak waveform and the image intensity surrounding the pixel in question. A single template was used on all pixels and was not scaled with the pixel intensity: a peak waveform that exactly matches the template intensity scores the lowest cost of zero while a peak waveform with a larger peak value than the template (and presumably higher quality) would score a higher cost because the mean squared difference is used. Instead of using the mean squared difference, we now use the negative of the correlation between the template and the waveform around the image pixel to set the cost.

Another term in the unary cost is the bottom location or bin in the nadir direction which is a priori information available at each along-track position. In [7] this was constructed to ensure the ice bottom layer passed beneath this bottom bin. The bottom bin cost was modified to force the ice bottom layer to pass within a 20-pixel neighborhood of the bottom bin rather than strictly beneath. This improved results when the bottom bin was too low due to errors in the human labeling of the nadir bottom bin. Although the bottom bin tends to be fairly accurate, the previous tracker implementation would allow the bottom layer to pass far beneath the bottom bin with no cost penalty.

The final term that was modified is a surface repulsion term that increases the ice bottom layer cost if it approaches the ice surface. This was necessary because the layer tracker has no other way to choose one or the other layer to track and, furthermore, the ice surface is generally stronger and more consistent. Combined in this term, but not fully described in [7], was an ice mask term that overrode this term when there was no ice present. In this case the ice bottom tracker was supposed to be equal to the ice surface to indicate zero ice thickness. The issue is that at the boundaries between ice and no-ice, the ice bottom often gradually separates from the ice surface. To accommodate this, the surface repulsion term was modified to gradually increase away from no ice edges and another term was added to attract the layers together as they approach a no ice region. This modification helped remove artificial discontinuities at the no-ice/ice transitions.

The binary term is used to control the smoothness of the surface. This allows the tracker to work through low quality regions where several pixel choices might be equally weighted based on image intensity only. The binary cost increases the

likelihood that the pixel that produces the smoothest surface overall will be the lowest cost. In [7], the smoothing was set to produce a constant range layer in the cylindrical coordinate system of the image. The problem is that a constant range layer in the cylindrical coordinate system is a tube in Cartesian space. The smoothing term was modified to set to lowest cost an ice bottom layer with the same range-slope as the ice surface. Although this is still not a flat surface, it is generally flatter and more realistic than the tube and was simpler computationally than calculating a flat ice bottom in Cartesian space which must account for ice refraction from a non-flat ice surface layer.

After the layer tracker is run, the layers are visualized by viewing the layer overlaid on the radar image in three dimensions along with a corresponding satellite image map and ice mask that shows the flight track and image pixel locations. Where needed, the bottom layer is manually corrected by adding additional ground truth points to the 3D image and corrections are made to the ice mask. The tracker is then rerun in small neighborhoods with this new information to correct issues. Poor data quality areas are also tagged so that they are not included in the final output.

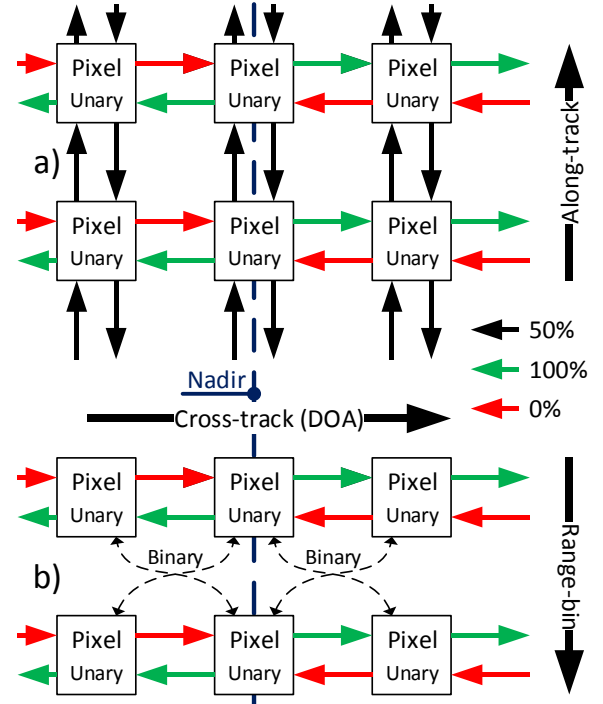


Fig. 3. Block diagram showing how TRW-S algorithm works.

III. RESULTS AND DISCUSSION

In this section, we assess our results from two perspectives. First, we self-assess our results by generating the DEMs of the overlapped areas from the crossing tracks (crossovers), and second, we present and discuss several statistics of the error between the output of the automatic layer-tracker before and after manual correction (MC). By MC we mean the manual correction by adding more ground truth, fixing the ice mask, and the data quality labeling step.

A. DEM Crossovers

The MC layers are used for all the results in this section; most importantly this includes the manual corrections. We have a total of 20 crossovers, but due to space limitations we have only illustrated two. Fig. 4 shows two crossover examples over ice which are also representative of the types of error patterns seen in the 20 crossovers. Each example has four plots, which are the zoomed version of the original DEMs with the intersection marked in red along with an inset showing the error map, and a sample slice from each DEM at the intersection of the two flight lines showing the ice-surface and ice-bottom after applying the automatic tracker (detailed in Section 2.B) in addition to manual corrections. The flight line for each DEM is shown in black. The location of the 3D slice at the intersection of the flight lines is shown in the error map and marked to indicate the left and right DoA portions. In the slice view, the left and right DoA portions are marked in red and white, respectively. The difference between the DEMs is the crossover error since, from a self consistency standpoint, the two DEMs would ideally be the same and the difference would be zero. Here a slice from the 3D image is the cross-track image at a particular range line, i.e. a 2D cross section from the 3D image. Table 1 shows the statistics of the overlapped DEM errors for the crossovers in Fig. 4. Over all 20 crossovers, the mean absolute error is 37 m, the median absolute error is 12 m, and the root mean square-error (RMSE) is 37 ± 9 m. If we assume Gaussian statistics for the errors in each of the individual DEMs, then the RMSE for the crossovers that is shown here should be $\sqrt{2}$ larger than the RMSE of the individual images. Fig. 5 shows the change of the average RMSE over all 20 crossovers as the largest errors are removed. This plot shows that the lower 70% of the errors have an RMSE of 10m. Table 2 gives the same results as in Table 1, but with the largest 10% of errors removed. The resultant reduction in error statistics is larger than would be expected for Gaussian distributed errors. These results tell us that there is usually a good match between the overlapped DEMs, but there are a few large errors that are causing the mean statistics to be large (i.e. a heavy tail distribution due to outliers). From examining Fig. 4B, if the tracker fails to track the correct surface, a whole region may have a very large error and create a heavy tail distribution.

Table 1: Statistics from the error of the overlapped DEMS.

Fig 3.	a	b
Mean Error [m]	21	89
Median Error [m]	13	38
RMSE [m]	30	142

Table 2: Statistics from the error of the overlapped DEMS when the largest 10% of the errors were removed.

Fig 3.	a	b
Mean Error [m]	15	60
Median Error [m]	11	32
RMSE [m]	20	92

Table 3: Statistics of the layer-tracker errors (measured in range-bins).

Average mean error	Average median error	Average RMSE
4.7	1.2	13.6

The DEM errors can result from several causes. We divide these into a few categories: **1) Poor data quality**, due to shadowing and weak backscatter. Different flight paths and improved instrument parameters may improve this category, but these causes cannot be changed in post processing. For example, in Fig. 4B, there seems to be no bottom signal in the right-most angular bins of the 3D slice example of frame 21. In this case, the ice-bottom layer is likely wrong since there is no or very little signal to track. **2) Errors due to suboptimal array processing.** The MUSIC method is known to be suboptimal to MLE [8] and we assume a fixed model order of 2 even though the scene in general may have more or less than 2 scattering sources. The beamformer was setup to scan through 64 DoA bins with uniform sampling in wavenumber. This DoA sample spacing limits the accuracy of the DoA resolution. **3) Tracking errors.** The tracking algorithm nominally follows the path with the largest peak correction in *each DoA bin (slice column)*. Parametric DoA estimation approaches usually search for the DoA with the largest results in *each range bin (slice row)*. The tracker also does not perform an exhaustive search of all paths since this is an np-hard problem and therefore the ice-bottom result may not be the best regardless of other deficiencies.

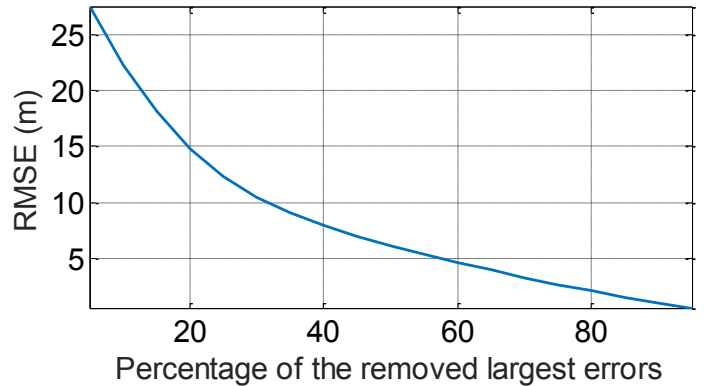


Fig. 5. Change of the average RMSE as we remove a percentage of the largest errors.

B. Layer-Tracking Assessment

Here we assess the ice-bottom tracking results by looking at statistics of the error (measured in range-bins) between the MC output, which has been manually corrected, and the result with no MC (NMC) which is the direct output from the automatic layer-tracker. Results are averaged over all 102 frames in this CAA dataset and each frame contains approximately 3332 slices. The reference layer is the MC ice-bottom layer. Thus, when the error is positive it represents the number of range-bins the NMC tracking result is above the MC tracking result. Note that the layer data from the first and last 5 DOA-bins, which correspond to the near-grazing angles, were removed in this comparison and in the generated DEM crossovers, as shown in the slice plots in Fig. 4. At these angles, the effective array aperture is very small and the received echoes are weak.

Table 3 shows the mean, median, and RMSE of the absolute error. Based on our previous published results, the old algorithm

had a mean error of 11.9 range bins over seven test frames whereas the new algorithm has a mean error of 2.3 range bins over the same test frames. Fig. 6 shows the cumulative distribution function of these errors. We see that $\sim 60\%$ of the errors are 0 (i.e. identical), 87% of the errors are within 5 range-bins from the MC results, and $\sim 96\%$ of the errors are within 25 range-bins from the MC results. These errors arise from different factors related to the error types mentioned in Section 3.A. From a qualitative inspection, the largest errors occur when the wrong layer is tracked which often means a few manually placed ground truth points allow the algorithm to track the correct layer.

These results show a good tracking capability, but with limitations where data quality is poor. In some cases, the MC result will also have errors, even for a trained analyst, especially in the places where the data quality is low (e.g. due to weak target echoes).




Fig. 6. Cumulative distribution function of the distance, measured in range-bins, between the manually corrected and non-manually corrected outputs of the layer-tracker.

IV. CONCLUSION

An automatic ice-layer tracker was implemented to track the ice-surface and ice-bottom based on the 3D images formed from the collected data, the ice-mask, the surface DEMs, and the human-labeled ground-truth. The DEMs of the basal topography of the Canadian Arctic Archipelago were then extracted using the tracked ice-bottom. Finally, these results were evaluated by determining the error statistics from the overlapping DEMs, and by comparing the tracked ice-bottom to the manually corrected one. We found that there is a good match between these overlapping DEMs, where, the mean error of the crossover DEMs is 37 ± 9 m, while the average error of the automatically tracked ice-layers, relative to the manually corrected tracked layers, is 13.6 range-bins, with 4.7 range-bins average absolute mean error.

ACKNOWLEDGMENT

We acknowledge support from NASA (NNX16AH54G), NSF (1443054), Canadian Space Agency, NSERC, and the State of Kansas. DEMs provided by the Polar Geospatial Center under NSF OPP awards 1043681, 1559691 and 1542736.

REFERENCES

- [1] A. Gardner, G. Moholdt, B. Wouters, G. Wolken, D. Burgess, M. Sharp, J. Cogley, C. Braun, C. Labine, "Sharply increased mass loss from glaciers and ice caps in the Canadian Arctic Archipelago," *Nature*, vol.473, pp.357-360, 2011.
- [2] A. Gilbert, G. Flowers, G. Miller, B. Rabus, W. Van Wychen, A. Gardner, L. Copland, "Sensitivity of Barnes Ice Cap, Baffin Island, Canada, to Climate State and Internal Dynamics," *Journal of Geophysical Research – Earth Surface*, vol.121, Issue 8, 2016.
- [3] F. Rodríguez-Morales, S. Gogineni, C. Leuschen, J. Paden, J.Li, C. Lewis, B. Panzer, D. Gomez-García Alvestegui, A. Patel, K. Byers, R. Crowe, K. Player, R. Hale, E. Arnold, L. Smith, C. Gifford, D. Braaten, C. Panton, "Advanced Multifrequency Radar Instrumentation for Polar Research," *IEEE Trans. Geoscience and Remote Sensing*, Vol. 52, No. 5, May 2014.
- [4] J. Paden, T. Akins, D. Dunson, C. Allen, P. Gogineni, "Ice-Sheet Bed 3-D Tomography," *Journal of Glaciology*, vol.56, no.195, pp. 3-11, 2010.
- [5] M. Al-Ibadi, J. Sprick, S. Athinarapu, T. Stumpf, J. Paden, C. Leuschen, F. Rodríguez, M. Xu, D. Crandall, G. Fox, D. Burgess, M. Sharp, L. Copland, and W. Van Wychen, "DEM Extraction of the Basal Topography of the Canadian Arctic Archipelago Ice Caps via 2D Automatic Layer-Tracker," *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pp. 965-968, 2017.q
- [6] Arendt, A., et al., 2012, Randolph Glacier Inventory [v2.0]: A Dataset of Global Glacier Outlines. Global Land Ice Measurements from Space, Boulder Colorado, USA, Digital Media, 2012.
- [7] M. Xu, D. Crandall, G. Fox, and J. Paden, "Automatic Estimation of Ice Bottom Subsurfaces from Radar Imagery," *IEEE International Conference on Image Processing (ICIP)*, 2017.
- [8] P. Stoica and A. Nehorai, "MUSIC, Maximum Likelihood, and Cramer-Rao Bound," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 37, no. 5, pp. 720-741, 1989.

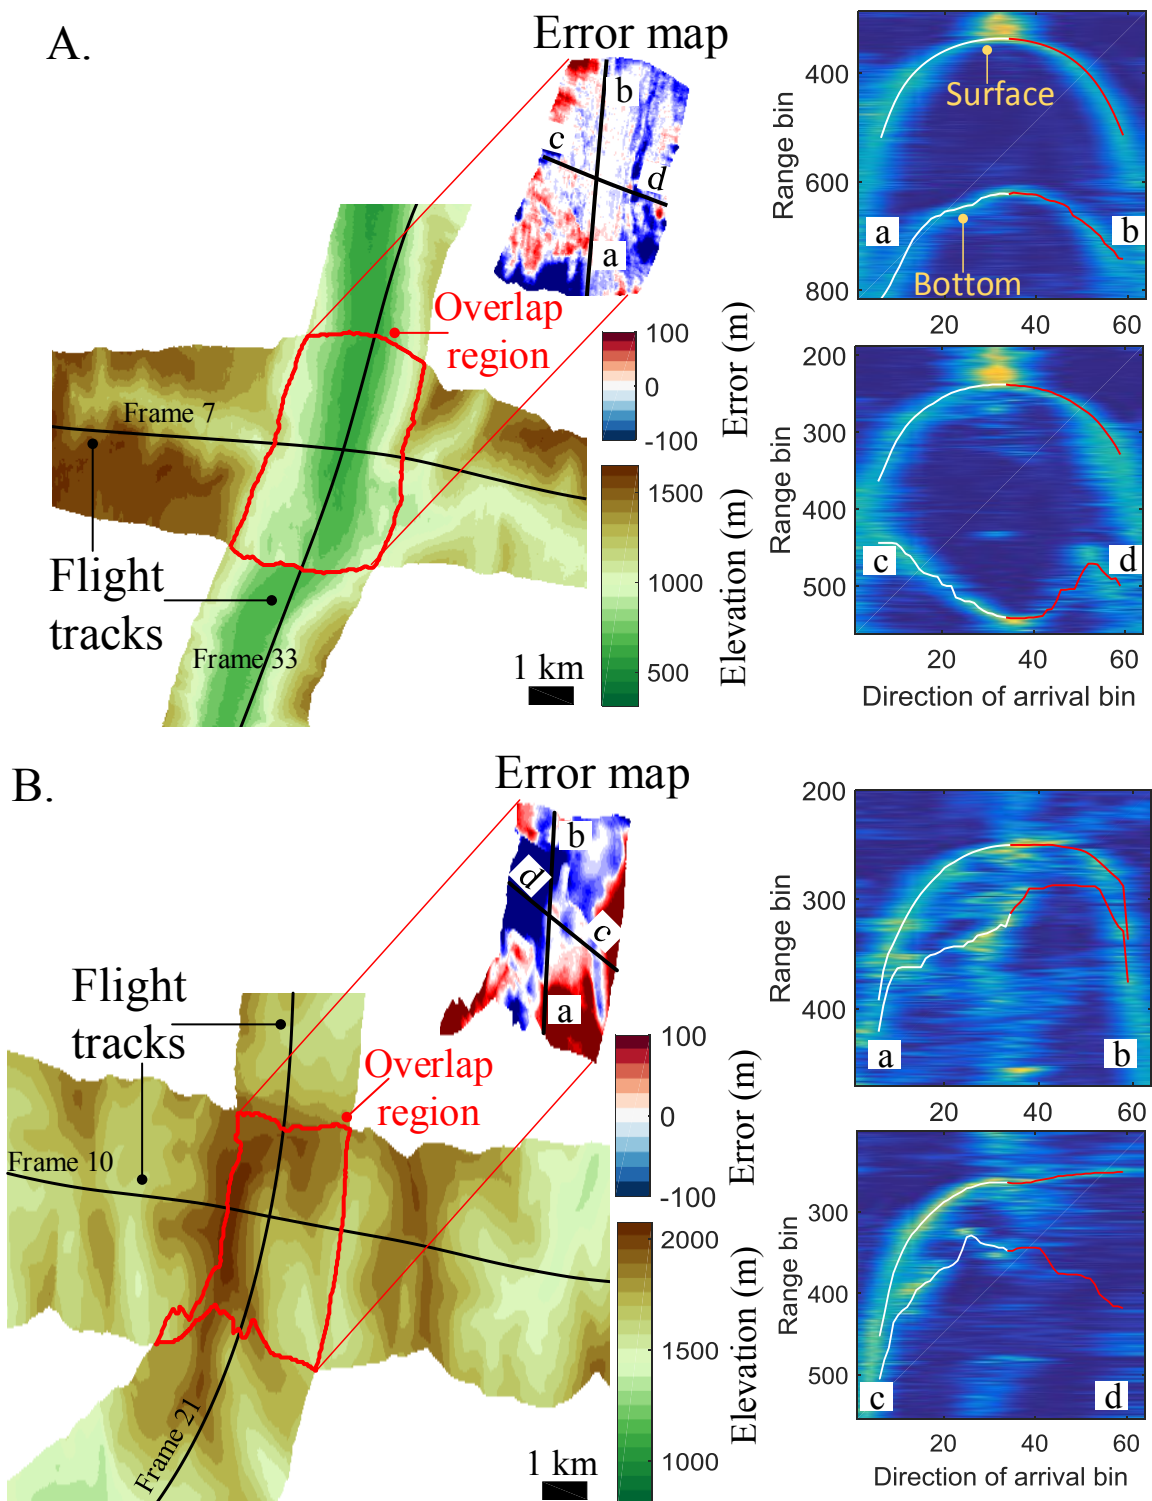


Fig. 4. Examples of a consistent (A) and inconsistent (B) crossover. The inconsistent crossover is caused by weak basal scattering for which the tracker fails. The DEM for each flight line is shown along with the difference between the two. Example slices from the 3D images used to construct the DEMs are also given.