Multi-Sensor Mapping for Low Contrast, Quasi-Dynamic, Large Objects

Vikrant Shah[1], Kristin Schild[2], Margaret Lindeman[3], Daniel Duncan[4], David Sutherland[5], Claudia Cenedese[6], Fiammetta Straneo[3] and Hanumant Singh[1]

Abstract— This paper proposes a systems level solution for addressing the problem of mapping large moving targets with slow but complicated dynamics with multiple sensing modalities. While this work is applicable to other domains we focus our efforts on mapping rotating and translating icebergs. Our solution involves a rigidly coupled combination of a line scan sensor - a subsurface multibeam sonar, with an area scan sensor - an optical camera. This allows the system to exploit the optical camera information to perform iceberg relative navigation which can directly be used by the multibeam sonar to map the iceberg underwater. This paper details the algorithm required to compute the scale of the navigation solution and corrections to find iceberg centric navigation and thus an accurate iceberg reconstruction. This approach is successfully demonstrated on real world iceberg data collected during the 2018 Sermilik campaign in Eastern Greenland. Due to the availability of iceberg mounted GPS observations during this research expedition we could also groundtruth our navigation and thus our systems level mapping efforts.

Fig. 1. R/V Adolf Jensen surveying the SF18-1 iceberg in the Sermilik fjord. Understanding melt rates of such icebergs such as the one pictured above is a critical parameter for modelling global climate change.

I. Introduction

This paper addresses the problem of mapping dynamic targets with unknown motion from a moving platform. An important application of this work is mapping icebergs in order to calculate their melt rates. Iceberg melt accounts for 30-50% of freshwater flux to the North Atlantic Ocean from the Greenland Ice Sheet[1]. As iceberg production accelerates due to climate change, the temporal and spatial distribution of this meltwater will be key to determining its effect on ocean circulation [2]. Existing estimates of iceberg melt rates depend on assumptions about iceberg geometry and poorly-constrained melt parameterizations [3], [4], [5]. However, iceberg melt could be measured directly by mapping the entire iceberg multiple times and estimating the volumetric difference. Acoustic sensors such as multibeam sonar are ideal for mapping the underwater iceberg geometry, which comprises about 90% of the iceberg volume. This approach has been demonstrated to be effective in direct measurements of underwater melt of a glacier face [5], but the dynamic nature of icebergs motion, with time varying x and y translation, and a rotational component, complicates the application of such methods.

Our main contributions are as follows:

1) We describe a system level approach to leverage information from 'area scan' camera sensors to improve mapping of 'line scan' multibeam sensors.
2) We introduce a method of mapping a dynamic floating target such as an iceberg or a floating vessel from both above the water and below the water. We demonstrate this method by mapping several icebergs from a ship.
3) We describe an approach to estimate relative navigation in a dynamic environment coupled with challenging contrast and illumination conditions which utilizes Contrast Limited Adaptive Histogram equalization [6] to enable distinctive feature detection.
4) We validate the iceberg motion estimate results by comparing against iceberg mounted GPS sensors and by considering the consistency of mapping over overlapping areas.

[1]Dept. of Electrical and Computer Engineering, Northeastern University, Boston, MA (email: shah.vi@husky.neu.edu, ha.singh@northeastern.edu)

[2]Climate Change Institute & School of Earth and Climate Sciences, University of Maine, Orono, ME (email: schild.km@gmail.com)

[3]Scripps Institution of Oceanography (email: lindeman@ucsd.edu, fstraneo@ucsd.edu)

[4]University of Texas Institute for Geophysics, Austin, TX (email: dduncan@ig.utexas.edu)

[5]Department of Earth Sciences, University of Oregon, Eugene, OR (email: dsuth@uoregon.edu)

[6]Woods Hole Oceanographic Institution, Woods Hole, MA (email: ccenedese@whoi.edu)

II. Background

Mapping icebergs and ice floes and the related problems of mapping sea ice and ship hulls have been well studied problem by field roboticists. Some of the notable works that have made strides in these related areas include [7], [8], [9], [10], [11], [12]. In this section, we discuss their merits and shortcomings and how they relate to our technique. While a multi-beam sonar [13] is a great tool for mapping, it does not easily lend itself to the Simultaneous Localization And Mapping (SLAM) problem [14]. This is in contrast to cameras where consecutive images can have very high overlap thus aiding localization by adding constraints to navigation. The fundamental issue is that multi-beam sonars are line scan sensors which allow us to obtain a 1D projection of a 2D slice of a 3D scene. Thus consecutive scans with a multi-beam sonar, in general, have minimal or no overlap with previous scans and thus are difficult to include within a SLAM framework. This serious limitation of multibeam sonar systems is compounded when we consider dynamic objects that may be translating and rotating.

A number of approaches have been developed to overcome this limitation. One of the earlier works in this area employs collections of scans organized into local submaps [15] which can be aligned to constrain the navigation for geological mapping applications. This approach assumes that the target being mapped is static as this is a fundamental requirement for the submaps to be consistent. Building on this, [7] uses a method of circumnavigating an iceberg and then exploiting the consistency between the start and end of the surveys to estimate constant velocity iceberg motion model parameters. The method presented here is similar in the circumnavigation aspect, however this technique solves for an estimate of the exact motion of the iceberg computing the actual time varying velocities of the iceberg. Kimball extended his own [8] by adding a Doppler Velocity Log (DVL) to establish iceberg relative navigation and allowing for a more general motion model. However the accuracy of the final map is still dependent on the accuracy of the DVL aided navigation which can be quite limiting for large surveys [16]. Additionally, the authors test their approach only on simulated data. Others have implemented the above scheme on an AUV platform to generate a 3D model of the underside of an iceberg [9]. Their method uses multiple overlapping runs to generate the model in contrast to our approach which requires minimal overlap.

Perhaps the most comprehensive and relevant piece of work is a study using a three camera setup to estimate a photogrammetric solution to estimate the iceberg motion while a Remotely Operated Vehicle (ROV) performed a multibeam survey of the iceberg [10]. The estimated iceberg motion was removed from the ROV tracks and the multibeam model was assembled in an iterative process involving manually cleaning the models by applying motion corrections. Our approach differs in that we do not require two separate survey platforms making it resource and, more importantly, time efficient. Additionally it employs a monocular camera system instead of a multi-camera system. Since our approach directly estimates the iceberg relative navigation of the multibeam sonar, we can use an automatic algorithm to generate the final model instead of a manual iterative process.

Another notable approach, in this case with respect to mapping sea ice, is reported in [11] where the authors introduce a terrain orientation measurement factor in the pose graph to account for the heading of ice floes which is then measured by a ship which is docked in the floe. Without external measurements of these factors, they add limited constraints to the optimization problem. In [12] the author introduces the concept of surface elements or surfels which help ensure self consistency in multibeam data for pose graph optimization in the ship hull inspection problem. However, the limitation of this approach is that it requires multiple overlapping surveys of a given region and it cannot handle dynamic motion.

We would like to point out that even though this work is motivated by iceberg and vessel type applications, it can be equally applicable for any combination of area scan and line scan sensors. An apt example is that of mapping and motion estimation of asteroids and comets as described in [17].

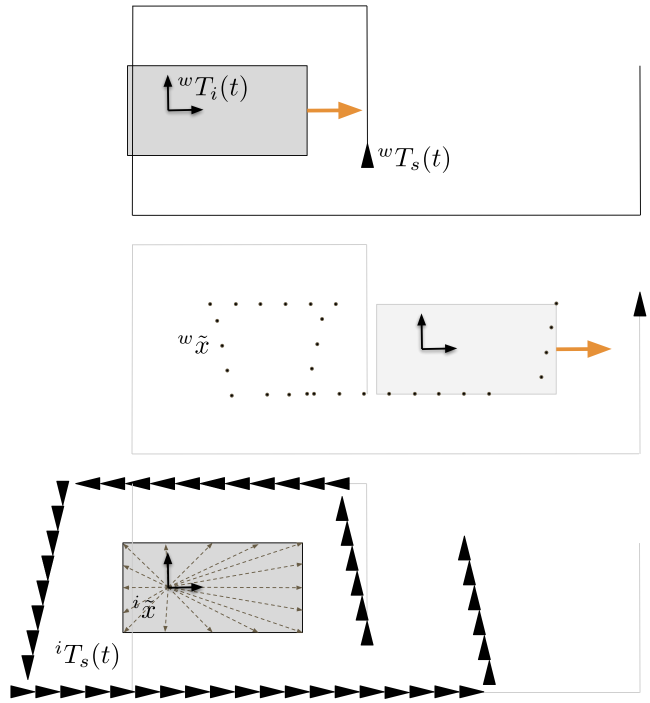

Fig. 2. A simple example for demonstrating the distortions. Top sketch, consider an iceberg moving left to right at constant velocity and a ship going around at 5x speed to map it. Middle sketch, resulting map of the iceberg shown by dots which are incrementally mapped by the ship. Bottom sketch, shows the calculated path of the ship which would generate the correct map of the iceberg from the same measurements if the iceberg was stationery.

III. Problem Statement

Our goal is to map a moving target using a line-scan sensor such as a multibeam sonar. In order to understand the effect of a moving target, we start with a simplified example. Figure 2 shows the process of mapping a 2D target using a point-scan sensor. As shown in fig 2(a), the target moves from left to right at a constant velocity and the the sensor circumnavigates the target at 5x the speed of the target. Figure 2(b) shows the resulting distorted map of the target where the sections of the target in the same direction as the sensor get elongated, sections in the opposite direction are shortened and those in perpendicular directions taper. One can imagine that time varying 2D translations accompanied with even small rotation of the target quickly makes the map much more complex to the extent that when a rotation is included there is no point wise correction that can be applied to the distorted map to restore the original map. Finally Figure 2(c) shows the simulated path of the sensor that would result in the corrected map assuming that the target is stationery.

Mathematically, this problem can be formulated as a problem of estimating the shape ${}^{i}\tilde{x}$ in a reference frame attached to the iceberg as depicted in Figure 2(c). However, what we are able to measure is

$$ {}^{w}\tilde{x} = {}^{w}T_s(t)^{s}x \tag{1} $$

where ${}^{w}\tilde{x}$ represents the distorted shape in the world frame shown in Figure 2(b) and ${}^{w}T_s(t)$ represents the estimated pose of the sensor in the world frame with time. Some of the previous works (e.g. [7]) directly manipulate ${}^{w}\tilde{x}$ to approximate ${}^{i}\tilde{x}$, while we find an estimate of ${}^{i}\tilde{x}$ directly. If we pre-multiply eq 1 with ${}^{i}T_w(t)$ the inverse of the position of the iceberg in the world frame we get

$$ {}^{i}T_w(t)^{w}\tilde{x} = {}^{i}T_w(t)^{w}T_s(t)^{s}x \tag{2} $$

the left hand side reduces to ${}^{i}\tilde{x}$ to give

$$ {}^{i}\tilde{x} = {}^{i}T_w(t)^{w}T_s(t)^{s}x \tag{3} $$

The approach, as described in [10], estimates the motion of iceberg ${}^{i}T_w(t)$ and that of the ROV ${}^{w}T_s(t)$, however solving for ${}^{i}\tilde{x}$ requires a manual iterative process. Eq 3 can be further reduced to

$$ {}^{i}\tilde{x} = {}^{i}T_s(t)^{s}x \tag{4} $$

where ${}^{i}T_s(t)$ represents the position of the sensor relative to the iceberg. Thus if we can measure or calculate the position of the sensor relative to the iceberg we can derive the shape of the iceberg as if it was measured while stationery. Using a DVL [8] attempted to estimate ${}^{i}T_s(t)$, however, their approach only provides a 3DOF estimate while this approach provides a 6DOF pose estimate using an optical camera. Figure 3 shows how the 6DOF relative pose estimate is incorporated to generate the corrected models. In the following sections we describe in detail the methodology used to calculate the ship's iceberg relative motion and estimate the iceberg shape.

IV. Algorithm

The goal of our processing algorithm is to negate the iceberg motion. Algorithm 1 shows the high level steps required to accomplish such a task. In this section we will discuss each step in detail.

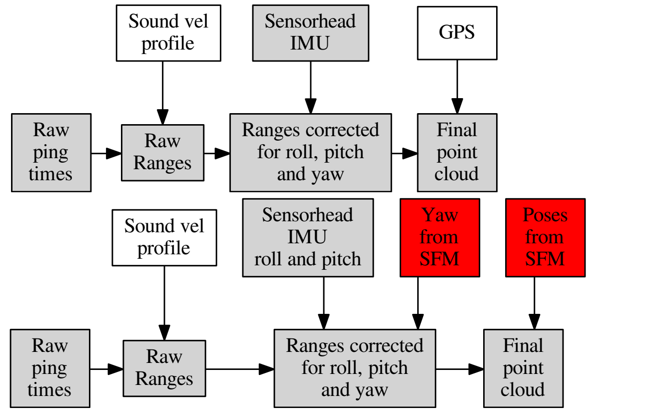

Fig. 3. Top flowchart shows the dataflow for generating a standard point cloud from a ship mounted multibeam. Bottom flowchart shows the dataflow for incorporating the iceberg centric yaw and position estimates into the corrected point cloud.

Algorithm 1: Estimate Iceberg-centric navigation

Input: Images of target with timestamps and GPS locations, lat/lon/heading used for multibeam pings

Output: Corrected lat/lon/heading to reprocess multibeam pings

Calculate monocular SFM solution pose for each image with respect to the iceberg ${}^{i}T_{c_n}$;

Fit a plane $\mathcal{N}$ to the set of poses ${}^{i}T_{c_n}$;

Find 3D rotation $\mathcal{R}_{3D}$ to rotate $\mathcal{N}$ to the XY plane;

Rotate all poses ${}^{i}T_{c_n}$ by $\mathcal{R}_{3D}$;

Find scale $\mathcal{S}$ by comparing translation of first n poses with the UTM coordinates ${}^{UTM}X$ of the images;

Scale all poses ${}^{i}T_{c_n}$ by $\mathcal{S}$;

Find best 2D rotation $\mathcal{R}_{2D}$ that minimises error between first n poses ${}^{i}T_{c_n}$ and ${}^{UTM}X$;

Rotate all poses ${}^{i}T_{c_n}$ by $\mathcal{R}_{2D}$;

Convert multibeam Lat/Lon to UTM coordinates;

Resample/interpolate poses ${}^{i}T_{c_n}$ to match multibeam timestamps;

Calculate corrected Lat/Lon/Heading from the resampled poses

The first step in the process is to find iceberg relative poses of the camera. In our approach, we use a monocular camera (GoPro Hero 5 Black) since it provides a convenient combination of fixed focal length images coupled with GPS sensors in a rugged package suitable in the marine environment. The camera captures images of the target at 1 Hz above water and is rigidly coupled to the multibeam pole hanging under water. During the survey the ship is steered around the iceberg while keeping it in the camera frame. Two different survey techniques were evaluated:

1) A circular path around the iceberg. This method works really well for the Structure From Motion

Fig. 4. Top picture shows a typical image of an iceberg with challenging illumination conditions (sun behind the target). Bottom picture shows the same image processed with Contrast Limited Adaptive Histogram Equalization (CLAHE) [6].

(SFM) algorithm since a large portion of the iceberg is always in view. However, the circular path is challenging for the multibeam. Since the yaw direction of the beam is continuously changing, it picks up pings from different areas of the iceberg without any obvious correlation. Thus pings cannot be registered to particular sections of the iceberg which would help to migrate them back to their corrected locations.

2) Straight line segments around the iceberg. This method simplifies the multibeam processing since all the pings in a given line segment can be tracked to a corresponding part of the iceberg. Additionally, any calibration errors between the multibeam, IMU and ship typically do not manifest within a particular line segment. This technique, however, creates additional challenges for the SFM solution because towards the ends of each line segment, the iceberg is observed in a very small part of the camera frame. When computing the full SFM solution, these sections of the dataset are the only thing that tie various line segments together. Thus, this might leave undesirable flexibility in the final bundle adjustment.

As shown in Section V later in the paper, we can handle both these cases within our algorithmic framework.

A. Image Pre-processing

The first step in most sparse photogrammetric applications is understanding the quality of feature detection. Our dataset includes images of icebergs which are very low contrast targets. The challenge is compounded by the lighting conditions where a number of images look into the sun, causing the target to be under exposed. In other cases cloudy days provide low lighting with very limited contrast where it is difficult to separate clouds from the actual icebergs. Some images also include fog where only parts of the iceberg are clearly visible. It was quickly evident that standard feature descriptors such as SIFT [18] and ORB [19] could not perform reliable registration on such low contrast iceberg images.

Borrowing from our previous work on low contrast underwater images, we used Zernike Moments [20] which worked reasonably well in these conditions but was very slow. However, pre-processing images with Contrast Limited Adaptive Histogram Equalization (CLAHE) [21], which we have also used for underwater marine imagery, allowed us to use the more standard (SIFT, ORB) feature descriptors. Additionally, pre-processing with CLAHE was found to be computationally more efficient by an order of magnitude than the Zernike Moment descriptors without any penalty in terms of performance.

Originally introduced for medical imaging [6], CLAHE is an enhancement on Adaptive Histogram Equalization where the original image is divided into equal sized context regions and a clip limit is used to clip and redistribute the clipped pixels equally among the other bins. The primary effect of the clip limit is to restrict the slope of the cumulative histogram to the clip limit. Figure 4 demonstrates the effect of CLAHE on a poorly illuminated iceberg image. Typically such non-linear image transformations are undesirable, however in this case, we find the resulting benefits outweigh any issues that such a nonlinear operation might create.

B. Structure From Motion and Bundle Adjustment

In order to solve for the monocular camera poses with respect to the iceberg we lean on the well established technique of feature based bundle adjustment [22] which solves for the camera poses along with 3D landmarks by minimizing the reprojection error as:

$$\min_{\hat{P}^i, \hat{\boldsymbol{X}}_j} \sum_{ij} d\left(\hat{P}^i \hat{\boldsymbol{X}}, x_i^j\right)^2 \tag{5}$$

We start with a calibrated camera model and perform an initial alignment of all the cameras ensuring that the loop closure is accurately incorporated. Following the initial alignment we perform a number of bundle adjustment steps to fine tune the camera calibration using landmarks with low co-variances which helps improve the overall co-variance on the landmarks. We use the Metashape package [23] as outlined above to obtain very accurate results for the relative poses. However, since we use a monocular camera system, we cannot directly estimate absolute scale [22]. In the following sections we describe how the scale is estimated.

Additionally, since the reference frame of the SFM poses is arbitrary, we fit a plane to the poses using Singular Value Decomposition (SVD) as described in [24]. Since our ship is going around the iceberg at an essentially constant sea level, this best-fit plane is an

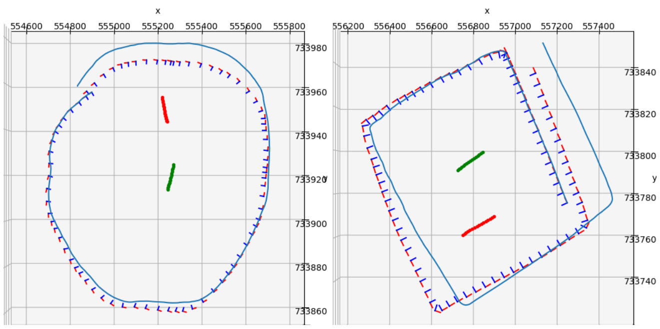

Fig. 5. In the above figures, the blue line represent the raw GPS tracks of the camera, RGB axis show every 50th camera pose from the SFM solution and the tracks in red and green show the measurements from the iceberg mounted GPS sensors. The plot on the left is for a circular survey of the SF18-2 iceberg while the one on the right is for a box style survey.

XY-plane parallel to the sea level. We rotate all the poses to this plane and set the first pose as its origin.

C. GPS

The multibeam surveys are referenced using GPS times and locations. Hence, we use the GPS location and time reference from the camera metadata to sync the multibeam data with the optical images at a millisecond level of accuracy. Since our camera runs at 1Hz we cannot use it to correct for the roll and pitch of the multibeam sonar. However, when circumnavigating an iceberg we expect the change in yaw to be very slow (a little more than 360°over the survey duration) allowing the camera to accurately capture yaw corrections.

D. Scale Estimation and Trajectory Alignment

Once we have two solutions from the GPS and SFM trajectories, the next step is to estimate the scale difference. We make the reasonable assumption that the motion of the iceberg is small relative to the ship velocity, thus we can expect the local trajectories between the GPS and SFM to match well. For most of our data sets we use the first 400 points (about 6 mins) for this computation. We then use the technique from [25] to match the scales using:

$$s = \left(\frac{\sum_{i=1}^{n} \left\| r_{1,i} - \frac{1}{n} \sum_{i=1}^{n} r_{1,i} \right\|^2}{\sum_{i=1}^{n} \left\| r_{2,i} - \frac{1}{n} \sum_{i=1}^{n} r_{2,i} \right\|^2} \right)^{\frac{1}{2}} \tag{6}$$

Once both the solutions are at the same scale, the optimal rotation R between them is estimated using [26] by first evaluating matrices P_1 and P_2 of the points around their respective centroids. The covariance matrix Q and its SVD decomposition is then calculated using:

$$Q = P_1^T P_2 \tag{7}$$

$$Q = USV^T \tag{8}$$

The optimal rotation R_{2d} is then:

$$R_{2d} = V \begin{bmatrix} 1 & 0 \\ 0 & det(VU^T) \end{bmatrix} U^T \tag{9}$$

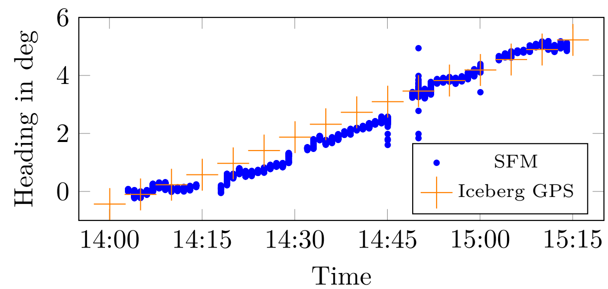

Fig. 6. Comparison of the iceberg rotation estimated by the SFM solution against the heading measured by iceberg mounted GPS sensors. The slope of the plot represents the iceberg rotational velocity and highlights the accuracy of our estimate.

Figure 5 shows an example of the result of the above sequence of operations on the SF18-2 iceberg datasets.

E. Synchronization and Interpolation

Once we obtain both the trajectories in the same scale, same plane and with correct orientations as outlined above, the only thing that remains to be calculated is the corrections for the multi-beam data. We do this by synchronizing the inertial measurement unit (IMU) and GPS data for the multi-beam with that from the camera and interpolating the results to give the corrections in the multi-beam frame of reference. These corrections are then processed as shown in Figure 3 to produce the corrected multi-beam surveys. Figure 7 shows some of the point clouds produced using the above algorithm. These are discussed in detail in Section V

V. Results

The primary objective of this work was constructing accurate iceberg geometries. The approach described above was used to process the data collected during the Sermilik 2018 campaign. More than ten icebergs were surveyed and mapped with the multi-beam sonar. Out of these ten icebergs, six also had a visual survey performed that was coincident with the multi-beam surveys. Finally, four of the icebergs had helicopter-deployed GPS sensors (two units per iceberg to enable rotational calculations). Unfortunately, one of the GPS sensors sank and two of the surveys had GPS data missing due to technical issues. So there was one complete GPS ground truth record on the SF18-2 iceberg. Also, some of these icebergs were surveyed multiple times a few days apart to enable comparative analysis. However, for the purposes of this paper, our focus was on demonstrating a robust automated technique that works for iceberg mapping which we could ground truth with available GPS data.

Figure 5 shows the results of the iceberg centric navigation on the SF18-2 iceberg when performed on two different instances utilizing a circular and box style method for survey. The lines in red and green show the two iceberg mounted GPS tracks. The track in blue shows the raw GPS position of the camera while the

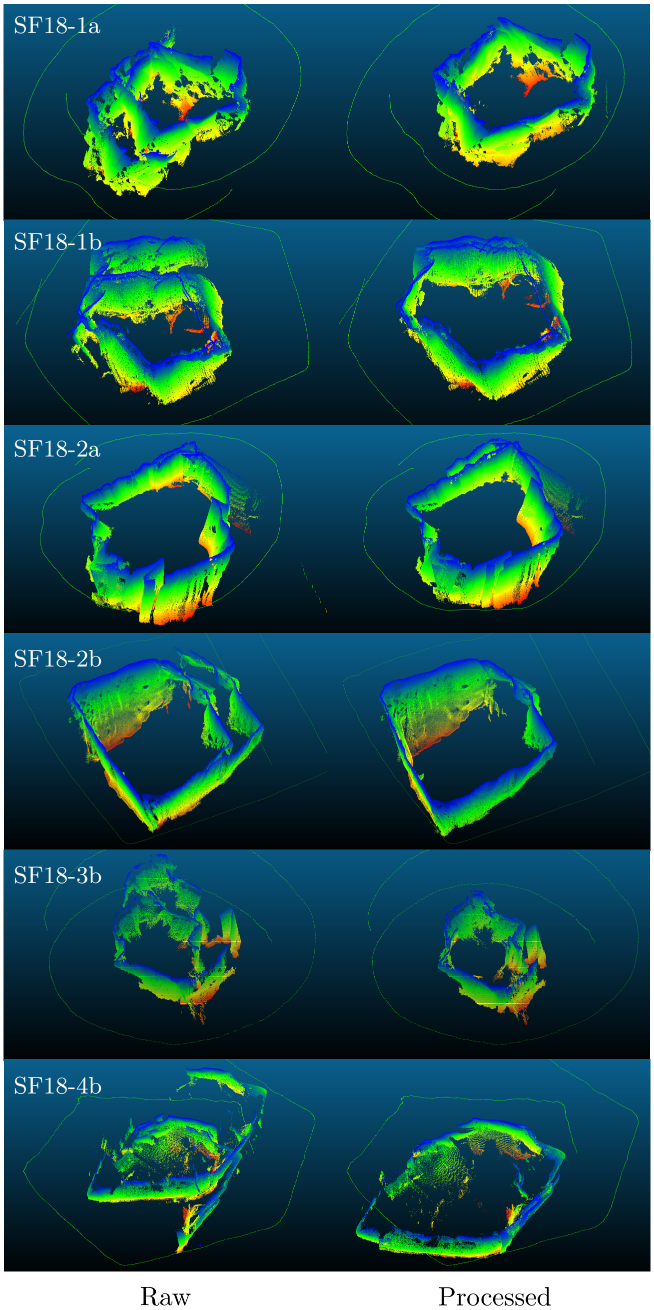

Fig. 7. The images on the left show raw multibeam point clouds of the icebergs with the ship's path as a green line. The images on the right show the point clouds after iceberg centric reprocessing.

RGB axes show the full 6-DOF pose from the SFM solution for every 50th camera image. On the left figure, which corresponds to the circular survey, it can be seen that the iceberg is translating north and rotating counter clockwise. On the right figure, corresponding to the box survey for the same iceberg, we see that the iceberg is translating northwest and very slowly rotating counter clockwise. The difference between the RGB pose track and the GPS track shows the corrections that need to be fed into the multibeam reprocessing to calculate the corrected sonar map of the iceberg.

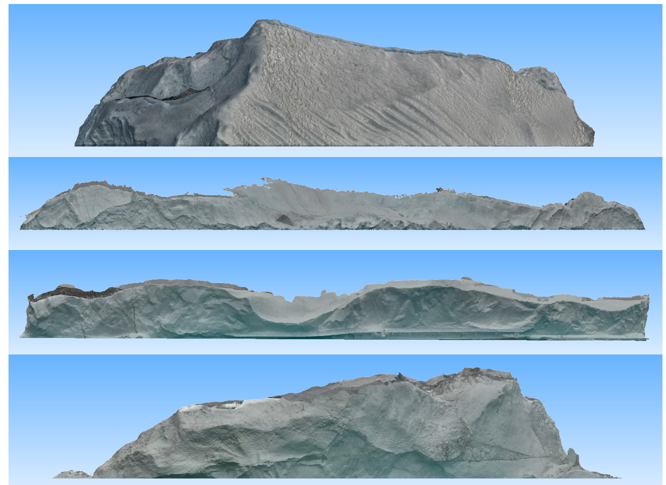

Fig. 8. The above images are snapshots of 3D meshes of the iceberg sails using dense reconstruction on the SFM pose solutions. Each of these sections representing a side of the iceberg is generated using 300-800 images. Images from top to bottom show models of SF18-1, SF18-2, SF18-3 & SF18-4 icebergs.

Figure 6 shows a comparison of the SFM solution against the ground truth from the iceberg mounted GPS on the SF18-2b survey. The blue dots are obtained by subtracting the camera heading from the SFM solution compared to the multi-beam IMU. The + symbols are obtained from the heading between the two iceberg mounted sensors. Notice that the scale is in degrees and the difference between the two plots is a small fraction of a degree. Also, far more important than the absolute heading, the slope of the two plots, which represents the rotational velocity of the iceberg, is also extremely consistent. This implies that our method is able to correctly and accurately measure the iceberg's rotational velocity.

Figure 7 shows the multibeam point clouds for the six surveys that were conducted with coincident camera data. The images on the left shows the raw point clouds recorded by the multibeam in UTM coordinates. These point clouds are distorted because of the iceberg translation and rotation as described in Section III. The green line shows the GPS track of the ship. Finally, in the images on the right are the point clouds generated by reprocessing the raw multibeam data using our algorithm. It can be observed that the algorithm performs extremely well by comparing the overlap of the points at the beginning and end of the surveys. We note that we are not performing Iterative Closest Point (ICP) alignment on the raw data, so the final result is purely the outcome of running the raw data through our algorithm. Another observation worth noting is that in a couple surveys such as in the case of iceberg SF18-2a there is an offset at the beginning of the survey, This is because the visual data recording was only started partway into the survey as opposed to the very beginning.

In summary it is worth highlighting that our method was utilized across multiple sets of data in an automated manner based entirely on measures derived from the multisensor geometry. The measure of such algorithms for mapping in bathymetric, under-ice and ship hull inspection problems have always relied on examining the consistency between overlapping measurements. This is primarily due to our inability as a community to obtain ground truth measurements. By this measure our algorithm works extremely well. In addition, however, we were in a position to compare our results with independent ground truth in the form of multiple GPS sensors on the iceberg and in this case too our results were in remarkable agreement.

Figure 8 shows the iceberg sail 3D meshes generated using dense reconstruction on the 6-DOF pose alignment SFM solution of the iceberg visual data. We can observe that the reconstruction is able to capture the finest details on the mesh that might not even be visible on the raw images owing to the CLAHE pre-processing. Thus we were able generate very detailed models of both the parts above the water surface and the underwater sections.

VI. Conclusion and Future Work

In this paper, we presented a novel method of rigidly coupling an area scan sensor - an optical camera, and a line scan sensor - a multi-beam sonar, with both operating in domains where there respective performance is superior. A multi-beam sonar underwater coupled to a camera above water. With this approach, we were able to leverage the modalities for both the sensors, the camera for continuously tracking our sensing system's six degree of freedom pose with respect to a dynamic target and the multibeam sonar for providing high quality maps underwater. We describe a detailed methodology for correcting the navigation using this data and for correcting the resulting multibeam model. We were able to demonstrate the successful application of our technique on all the available datasets of icebergs which included a variety of shapes and sizes, had different kinds of translational and rotational motions, and in different environments comprising a wide range of visual conditions. We also made a comparison of our estimate of the iceberg motion to the ground truth obtained from GPS sensors mounted on the iceberg. We were able to show that our technique can provide iceberg motion results as accurate as iceberg mounted sensors without having to undertake risky manned helicopter landings on icebergs to install and retrieve the sensors.

Our systems level approach for mapping dynamic targets is better than the current state of the art which tries to either navigate relative to the iceberg underwater or estimates the iceberg motion and then iteratively corrects the multibeam model. In addition to being more accurate, our approach is resource and, more importantly, time efficient requiring just one circumnavigation around the target offshore. We hope to continue this work during successive trips to Greenland that would include a stereo camera setup which would reduce the uncertainty in the model scale estimation. We will be making our datasets and algorithms available to the community as soon as our scientific collaborators have published their results in the polar and physical oceanographic communities.

References

[1] E. M. Enderlin, I. M. Howat, S. Jeong, M. J. Noh, J. H. Van Angelen, and M. R. Van Den Broeke, "An improved mass budget for the Greenland ice sheet," Geophysical Research Letters, 2014.

[2] J. L. Bamber, A. J. Tedstone, M. D. King, I. M. Howat, E. M. Enderlin, M. R. van den Broeke, and B. Noel, "Land Ice Freshwater Budget of the Arctic and North Atlantic Oceans: 1. Data, Methods, and Results," Journal of Geophysical Research: Oceans, 2018.

[3] E. M. Enderlin, G. S. Hamilton, F. Straneo, and D. A. Sutherland, "Iceberg meltwater fluxes dominate the freshwater budget in Greenland's iceberg-congested glacial fjords," Geophysical Research Letters, 2016.

[4] T. Moon, D. A. Sutherland, D. Carroll, D. Felikson, L. Kehrl, and F. Straneo, "Subsurface iceberg melt key to Greenland fjord freshwater budget," Nature Geoscience, 2018.

[5] D. A. Sutherland, R. H. Jackson, C. Kienholz, J. M. Amundson, W. P. Dryer, D. Duncan, E. F. Eidam, R. J. Motyka, and J. D. Nash, "Direct observations of submarine melt and subsurface geometry at a tidewater glacier," Science, 2019.

[6] K. Zuiderveld, "Contrast Limited Adaptive Histogram Equalization," in Graphics Gems, 1994.

[7] P. Kimball and S. Rock, "Sonar-based iceberg-relative navigation for autonomous underwater vehicles," Deep-Sea Research Part II: Topical Studies in Oceanography, vol. 58, no. 11-12, pp. 1301–1310, 2011. [Online]. Available: http://dx.doi.org/10.1016/j.dsr2.2010.11.005

[8] P. W. Kimball and S. M. Rock, "Mapping of translating, rotating icebergs with an autonomous underwater vehicle," IEEE Journal of Oceanic Engineering, 2015.

[9] R. S. McEwen, S. P. Rock, and B. Hobson, "Iceberg Wall Following and Obstacle Avoidance by an AUV," in 2018 IEEE/OES Autonomous Underwater Vehicle Workshop (AUV). IEEE, 2018, pp. 1–8.

[10] A. Younan, F. Ralph, T. Ralph, and J. Bruce, "Overview of the 2012 Iceberg Profiling Program," in Arctic Technology Conference. St. John's, Newfoundland and Labrador, Canada: Offshore Technology Conference, oct 2016. [Online]. Available: http://www.onepetro.org/doi/10.4043/27469-MS

[11] C. G. Kunz, "Autonomous underwater vehicle navigation and mapping in dynamic, unstructured environments," Ph.D. dissertation, Massachusetts Institute of Technology and the Woods Hole Oceanographic Institution, 2012. [Online]. Available: https://dspace.mit.edu/handle/1721.1/71277

[12] P. N. V. Teixeira, "Dense, Sonar-based Reconstruction of Underwater Scenes," Ph.D. dissertation, Massachusetts Institute of Technology and Woods Hole Oceanographic Institution, 2019.

[13] H. Singh, L. Whitcomb, D. Yoerger, and O. Pizarro, "Microbathymetric mapping from underwater vehicles in the deep ocean," Computer Vision and Image Understanding, 2000.

[14] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," IEEE Transactions on Robotics, 2016.

[15] C. Roman and H. Singh, "Improved vehicle based multibeam bathymetry using sub-maps and SLAM," in 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS, 2005.

[16] L. Whitcomb, D. Yoerger, and H. Singh, "Advances in Doppler-based navigation of underwater robotic vehicles," Proceedings - IEEE International Conference on Robotics and Automation, 1999.

[17] N. Takeishi and T. Yairi, “Visual monocular localization, mapping, and motion estimation of a rotating small celestial body,” Journal of Robotics and Mechatronics, 2017.

[18] D. G. Lowe, “Distinctive image features from scale-invariant keypoints,” International Journal of Computer Vision, 2004.

[19] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, “ORB: An efficient alternative to SIFT or SURF,” in Proceedings of the IEEE International Conference on Computer Vision, 2011.

[20] O. Pizarro and H. Singh, “Toward large-area mosaicing for underwater scientific applications,” IEEE Journal of Oceanic Engineering, 2003.

[21] R. Eustice, O. Pizarro, H. Singh, and J. Howland, “UWIT Underwater Image Toolbox for Optical Image Processing and Mosaicking in MATLAB,” in Underwater Technology, vol. 2002-Janua. IEEE, 2002, pp. 141–145. [Online]. Available: http://ieeexplore.ieee.org/document/1002415/

[22] R. Hartley and A. Zisserman, Multiple view geometry in computer vision. Cambridge university press, 2003.

[23] Agisoft LLC, Metashape Professional Edition Ver 1.5, 11 Degtyarniy, St. Petersburg, Russia, 191144, February 2019.

[24] C. M. Shakarji, “Least-Squares Fitting Algorithms of the NIST Algorithm Testing System,” Journal of Research of the National Institute of Standards and Technology, 1998.

[25] B. K. P. Horn, H. M. Hilden, and S. Negahdaripour, “Closed-form solution of absolute orientation using orthonormal matrices,” Journal of the Optical Society of America A, 1988.

[26] W. Kabsch, “A solution for the best rotation to relate two sets of vectors,” Acta Crystallographica Section A, 1976.