University of Rhode Island

DigitalCommons@URI

Open Access Dissertations

2015

Factor Graphs and Submap Simultaneous Localization and Mapping for Microbathymetry

James Ian Vaughn

University of Rhode Island, vaughn.ian@gmail.com

Follow this and additional works at: https://digitalcommons.uri.edu/oa_diss

Recommended Citation

Vaughn, James Ian, "Factor Graphs and Submap Simultaneous Localization and Mapping for Microbathymetry" (2015). *Open Access Dissertations*. Paper 392.
https://digitalcommons.uri.edu/oa_diss/392

This Dissertation is brought to you for free and open access by DigitalCommons@URI. It has been accepted for inclusion in Open Access Dissertations by an authorized administrator of DigitalCommons@URI. For more information, please contact digitalcommons-group@uri.edu.

FACTOR GRAPHS AND SUBMAP SIMULTANEOUS LOCALIZATION AND
MAPPING FOR MICROBATHYMETRY

BY

JAMES IAN S. VAUGHN

A DISSERTATION SUBMITTED IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
IN
OCEAN ENGINEERING

UNIVERSITY OF RHODE ISLAND

2015

DOCTOR OF PHILOSOPHY DISSERTATION
OF
JAMES IAN S. VAUGHN

APPROVED:

Dissertation Committee:

Major Professor Christopher Roman

Stephen Licht

Richard Vaccaro

John W. King

Nasser H. Zawia
DEAN OF THE GRADUATE SCHOOL

UNIVERSITY OF RHODE ISLAND

2015

ABSTRACT

The utility of high-resolution bathymetric surveys important to many problems in oceanography is often limited by the poor navigation options available to Unmanned Underwater Vehicles. This thesis presents a novel method to integrate conventional dead-reckoning navigation with recent advances in factor graph Simultaneous Localization and Mapping (SLAM) algorithms. Dead-reckoning is represented in the factor graph as a new type of factor using a linearized state transition model derived from the Extended Kalman Filter commonly used for dead reckoning. The new factor graph submap-SLAM is faster and more scalable than prior methods and is shown to properly represent navigation uncertainty. This new method is used to evaluate change detection algorithms using surveys before and after excavation of the Monterrey A shipwreck. Factor graph submap-SLAM is shown to significantly reduce change detection artifacts caused by navigation error. Finally, a derivation using the Cramer Rao Lower Bound demonstrates that all of the navigation improvement provided by SLAM over dead reckoning results from the quality of submap matches. This result leads to a metric that may be evaluated online during a survey to assess the terrain matching potential and may be used in the future to optimize survey trajectories for post-processing.

ACKNOWLEDGMENTS

No thesis happens in a vacuum, and this dissertation would never have happened without the assistance of a great many individuals. My advisor, Chris Roman, has been patient throughout this lengthy process. His occasional advice and dedicated efforts to keep the distractions at bay for the last year have been essential to the final result seen here. The defense committee has also been very helpful. As a group, they have insisted that the scope of this dissertation remain tractable while asking good questions. Stephen Licht, in particular, provided some excellent guidance on how to actually finish a dissertation. Gopu Potty was especially accommodating of the scheduling antics that allowed the defense to happen. Richard Vaccaro, John King, and Steve Carey showed up to committee meetings and asked good, relevant questions. Not all students are fortunate enough to find such a fine committee among the university's busy faculty.

The data for this thesis, and the enormous volume of data not in here that was still a critical part of the process, are the result of ongoing collaboration with the Ocean Exploration Trust and its employees, contractors, students, and supporters. I've lost track of everybody I've been to sea with over the years, but it's taken a lot of people to run these cruises. Mark, Mary, Al, Sarah, several Mikes, and a couple of Dans are just a few who left particular impressions. Captain Pavel and his crew have been especially accommodating of scientists wanting to do difficult things with the ship. These surveys were all done by pilots, including Brennan, Bob, Todd, Eric, Alex, Josh, and many more. On the shore side, the science department including Katy Croff-Bell, Nicole Raineault, and Mike Brennan, have been very supportive in fostering collaboration with outside scientists. They have also answered a thousand queries about how some sensor records data, what that site was, and similar sorts of things. Being an engineer, I've never claimed to

know any actual ocean science, and they've tirelessly filled in the gaps. The data engineering team, Ethan Gold and Justin Lowe, have been especially tolerant of my nit-picking. Not many people would take kindly to some student arriving on their ship and promptly demanding the timeserver get replaced. Ethan and Justin not only take such things in stride, but are usually happy to grab a beer afterwards.

The Monterrey expedition was a massive collaboration with many different organizations taking part, including OET, URI, NOAA, BOEM, BSEE and Texas A & M University. The expedition relied on an impressive amount of private financial support. It was a privilege to be a part of such a fascinating interdisciplinary project. I can only hope this modest contribution will be of some use as the archaeological work continues some two years after the final Herc recovery.

Collecting data is exciting, but the bulk of the scientific process happens in the lab. I've been fortunate to share this time with many fellow students over the years, including Bryan, Brian, Conner, John, Jeanie, Mike, Ned, Scott and Will, who we always seem to forget in these lists. MS students come and go, so in my mind the past few years will be defined by the long-timers in the Roman Lab. Gabrielle Inglis, who went before, Clara Smart, who will come after, and Regina Yopak, who's just getting warmed up. Dave Casagrande, who graduated but didn't leave and remains a standard for qualitative analysis. It's been a long run, but you've all kept it interesting.

Finally, and most importantly, my family has been consistent in their support for continuing education from the beginning. John, Margaret, and Duncan in Savannah have been very supportive over the years, especially during the undergraduate degree that got me here in the first place. The Massachusetts connection, especially Henry Sr., Julie and until recently Laura, hosted many family events within convenient driving distance. Henry the third was an excellent housemate

for many years; I wish him the best of luck on his new endeavor. My sister, Laura, has provided both diversion and occasional accommodation on trips to UNH over the past few years. None of this would have started without Mom & Dad, my first and favorite teachers. They will be thrilled, no doubt, that their children are finally finished with formal education.

Financial support for this dissertation has been provided primarily by the Office of Naval Research, the National Science Foundation, and the National Oceanic and Atmospheric Administration's Office of Exploration Research.

TABLE OF CONTENTS

ABSTRACT	ii
ACKNOWLEDGMENTS	iii
TABLE OF CONTENTS	vi
LIST OF TABLES	x
LIST OF FIGURES	xi
CHAPTER	
1 Introduction	1
1.1 Overview	1
1.1.1 Mapping Sensors	5
1.1.2 The Navigation Problem	9
1.2 Data Sources	13
1.3 Simultaneous Localization and Mapping	15
1.3.1 Submap-SLAM	16
List of References	20
2 Representation of Dead Reckoning Odometry in Factor Graphs	26
2.1 Introduction	26
2.1.1 Problem Statement	28
2.1.2 Related Prior Work	30
2.2 The Complete 12 Degree-of-Freedom Model	33
2.3 EKF-based Approximation	34

	Page
2.3.1 State Transition Model	34
2.3.2 Factor Graph Representation	37
2.4 Properties of Conditional Factor Graph Dead Reckoning	40
2.4.1 Comparison to Augmented-State EKF	40
2.4.2 Correlation Time	44
2.5 Validation Under Realistic Conditions	47
2.5.1 Comparison with Dead Reckoning Data	48
2.5.2 Adding LBL Data	50
2.6 Conclusion	51
List of References	53
3 Factor Graphs for Change Detection	56
3.1 Introduction	56
3.2 Monterrey A	57
3.2.1 Site Overview	57
3.2.2 Surveys Conducted	59
3.3 Processing Single Surveys	61
3.3.1 Conventional Dead-Reckoning	61
3.3.2 Long Baseline Positioning	61
3.3.3 Submap-SLAM	64
3.3.4 Results	65
3.4 Environmentally-Induced Non-Bathymetric Changes	79
3.4.1 Tides	80
3.4.2 Atmospheric Pressure	81

	Page
3.4.3 Water Column Properties	82
3.4.4 Error Budget for Environmental Changes	84
3.5 Excavation-Induced Bathymetric Changes	85
3.6 Processing Multiple Surveys for Change Detection	88
3.6.1 Rigid Alignment	89
3.6.2 Non-Rigid Alignment With Anchor Nodes	90
3.6.3 Results	93
3.7 Conclusions	103
List of References	105
4 Predicting Submap-SLAM Performance from Seafloor Properties	109
4.1 Submap SLAM	109
4.2 Predicting SLAM Performance From Submap Properties	112
4.2.1 Probabilistic Models for SLAM	112
4.2.2 An Information View of Submap-SLAM	115
4.2.3 Submap-to-Submap Link Validity	118
4.3 Metrics to Evaluate Possible Link Quality	121
4.3.1 Autoregistration	121
4.3.2 Normal Projection as an Alternate Metric for Matching .	125
4.3.3 z -Variance as a Fast Metric for Matching Performance .	127
4.4 Comparison of Autoregistration and Normal Variance	128
4.4.1 Relationship with Rugosity	132
4.5 Implications of Submap-Matching Limitations for SLAM and Survey Planning	134

	Page
4.6 Summary	135
List of References	136
BIBLIOGRAPHY	139

LIST OF TABLES

Table		Page
1	Navigation sensors on ROV <i>Hercules</i>	14
2	Monterrey A Surveys	60
3	Summary of Depth Sensor Biases	85
4	Statistics for Change Surfaces at Monterrey A	104
5	Registerability metric processing time	132

LIST OF FIGURES

Figure		Page
1	Layout of sensors and survey patterns with ROV <i>Hercules</i> . . .	1
2	Scale comparison of shipboard and ROV MBES data	5
3	Raw data from a Blueview MB-1350-90 MultiBeam Echo Sounder	6
4	High-frequency multibeam and structured light data	8
5	Typical deep-ocean USBL navigation performance	12
6	ROVs <i>Hercules</i> and <i>Argus</i>	14
7	Bathymetric submap construction	17
8	Example Submap-SLAM Factor Graph	18
9	Example SLAM Dead Reckoning Graph	30
10	Simplified SLAM Dead Reckoning Graph	40
11	Variable Elimination Graphical Proof, Part 1	42
12	Variable Elimination Graphical Proof, Part 2	43
13	State transition matrix over time	46
14	Comparison of EKF and factor-graph for dead reckoning	49
15	Comparison of EKF and factor-graph for LBL processing	52
16	Context of Monterrey A Excavation	58
17	Monterrey A Photomosaics	59
18	Sound velocity stability at Monterrey A	63
19	Large Pre-Disturbance Bathy, DR Navigation	66
20	Large Pre-Disturbance Bathy, LBL Navigation	67
21	Large Pre-Disturbance Bathy, SLAM	68

Figure		Page
22	Error in Large Pre-Disturbance Survey	69
23	Small Pre-Disturbance Bathy, DR Navigation	70
24	Small Pre-Disturbance Bathy, LBL Navigation	71
25	Small Pre-Disturbance Bathy, SLAM	72
26	Error in Small Pre-Disturbance Survey	73
27	Post-Disturbance Bathy, DR Navigation	74
28	Post-Disturbance Bathy, LBL Navigation	75
29	Post-Disturbance Bathy, SLAM	76
30	Error in Post-Disturbance Survey	77
31	Self-Consistency Error Histograms	78
32	Tides during Monterrey A excavation	80
33	Barometric pressure near Monterrey A	81
34	Seawater properties over Monterrey A	83
35	Herc Captures of Artifacts That Changed	85
36	Log-histogram of <i>Hercules</i> positions at Monterrey A	86
37	Monterrey A Sediment Properties from Post-Disturbance Imagery	87
38	Anchor Nodes for Multi-Session Alignment	91
39	Multi-session Navigation Uncertainty	94
40	Measured Change Between Pre-Disturbance Surveys, DR	95
41	Measured Change Between Pre-Disturbance Surveys, LBL	96
42	Measured Change Between Pre-Disturbance Surveys, Single-Survey SLAM	97
43	Measured Change Between Pre-Disturbance Surveys, Multi-Session SLAM	98

Figure		Page
44	Change Between Large Pre- and Post- Disturbance Surveys, DR	99
45	Change Between Large Pre- and Post- Disturbance Surveys, LBL	100
46	Change Between Large Pre- and Post- Disturbance Surveys, Single-Map SLAM	101
47	Change Between Large Pre- and Post- Disturbance Surveys, Multi-Session SLAM	102
48	Change Detection Surface Histogram	103
49	An apparently registrable navigation artifact	111
50	Comparison of submap autoregistration for interesting and boring submaps	123
51	Autoregistration results for Monterrey A	128
52	Autoregistration results of Knidos F	129
53	Autoregistration results of the Kula summit	130
54	Autoregistration results of the south side of Mt. Dent	131
55	Autoregistration compared with z -variance and normal complement metrics	132
56	Rugosity at Kula Summit	133

CHAPTER 1

Introduction

1.1 Overview

Mapping is a fundamental tool for ocean exploration and research. Maps are important tools in geology [1], biology [2], archaeology [3, 4, 5], and many other fields. In addition to providing context for qualitative analysis of sites, the resulting bathymetric models are increasingly being used for quantitative analyses, including habitat classification [6], numerical modeling [7] and long-term monitoring [8]. Maps also have important roles in a number of policy applications including antiquities conservation [9], ocean spatial planning [10], and resource management [11].

Figure 1: The approximate sensor layout on Remotely Operated Vehicle (ROV) *Hercules* is shown on left. The camera footprints are shown as red and blue boxes and the MultiBeam Echo Sounder (MBES) sonar as a line of blue dots. The sensors remain fixed on the vehicle throughout a survey. Large areas are covered by moving the entire vehicle back and forth across the target area, as shown on right. The vehicle track is shown as a grey line running back and forth over the surface of the resulting survey, and an example MBES footprint is shown in cyan.

Mapping in the deep ocean presents a number of unique challenges. Imaging systems offer a trade-off between covering large areas and providing detailed data.

Placing a sensor closer to a surface gives greater measurement resolution over a smaller area. Surveying the deep ocean floor from the ocean’s surface limits a surveyor to a single choice of altitude. Higher resolution and greater flexibility are obtained by deploying a sensor in the water column. This approach is most commonly done by mounting a seafloor imaging sensor on an unmanned underwater vehicle (UUV). UUVs are broadly classified into two categories. Remotely operated vehicles (ROV) are controlled and powered through a tether to a support vessel. Autonomous underwater vehicles (AUV) are controlled autonomously by an onboard computer. This dissertation is directly applicable to high-resolution deep-ocean bathymetric surveys from both classes of UUVs in depths over 1000m.

To produce maps, bathymetric data are first collected by moving a seafloor sensor back and forth over the target area in parallel track lines. The final map is made by combining the seafloor measurements with the position of the sensor at the time each measurement was made. A variety of bathymetric sensors may be used and are introduced in greater detail in Section 1.1.1. GPS does not work underwater, making estimation of the sensor’s position and orientation by computing a navigation solution for the dive the greatest challenge in surveying from UUVs. Position data is provided by a combination of acoustic positioning sensors, which suffer from insufficient accuracy, and dead-reckoning, which produces a solution exhibiting steadily increasing error over time. These methods are further detailed in Section 1.1.2. The primary challenge in surveying from UUVs is to limit the error introduced by navigation.

Using previously-collected seafloor data to help find the location of a ship is a traditional navigation method dating back centuries [12, 13]. This method, known in its modern form as terrain-aided navigation, continues to be an area of active research with a broad range of applications [14, 15, 16, 17, 18]. UUVs often

survey areas where no high-quality maps are available, as producing such maps is the survey objective. The solution is to simultaneously estimate both a navigation solution and the resulting map. This approach is referred to in the literature as the Simultaneous Localization and Mapping (SLAM) problem or Concurrent Mapping and Localization (CML). The SLAM problem has been extensively studied over the last 30 years and a few relevant results in underwater mapping are described in Section 1.3.

A SLAM method may be evaluated using a number of criteria including quality of the final corrections, required data, and execution time. Map quality is directly influenced by the quality of the navigation, making evaluation of the map and navigation quality equivalent. Without GPS measurements, the navigation accuracy is difficult to evaluate, although external navigation data are used when available. Evaluating map quality is also non-trivial. This work relies principally on the Hausdorff distance self-consistency error metric of Roman & Singh [19] that measures the maximum error between overlapping survey lines. Different SLAM approaches require different sets of input data. On-line methods use only the data that have been collected so far. Off-line methods use all data for a given site and are necessarily restricted to post-processing. Real-time processing presents a number of challenges beyond requiring causal algorithms. Processing power on untethered UUVs is often limited by power constraints. Adding more computing power uses power that may otherwise be used to increase vehicle endurance. Processing surveys on the vehicle can reduce the maximum survey duration. The work presented here is concerned with producing the highest-quality map products possible from the collected data using post-processing methods.

One SLAM method for underwater bathymetric surveying is the Submap-SLAM method introduced by Roman [20]. This method groups consecutive pings

together into rigid “submaps” and uses standard SLAM techniques to refine an estimate of the position of each submap. The underlying approach is described in greater detail in Section 1.3.1. The original Submap-SLAM algorithm was implemented with an augmented-state Extended Kalman Filter (EKF). While effective, the EKF proved computationally prohibitive for large numbers of submaps. Sparse matrix methods using factor graphs have been a popular recent advancement for addressing large SLAM problems [21, 22, 23]. A new method to adapt the dead-reckoning constraints of Submap-SLAM to factor graphs is discussed in Chapter 2. Factor graph methods are a general approach to solving large probabilistic constraint networks and may be applied to a number of problems. The novel factor graph construction method is then applied to the practical problem of detecting bathymetric change during the excavation of the Monterrey A shipwreck in Chapter 3. Although few statistically significant changes are measured, the versatility of the factor-graph dead-reckoning constraints introduced in Chapter 2 is demonstrated with application to both Submap-SLAM and an uncalibrated long-baseline (LBL) acoustic navigation network. Submap-SLAM performance at Monterrey A suggests that not all areas of the sea floor are equally useful for submap matching. The implications of this for SLAM performance are studied further in Chapter 4. SLAM performance is found to depend principally on the quality of submap matches. An online algorithm to score different areas of seafloor for their usefulness to a later SLAM algorithm is introduced. Together, these studies demonstrate the broad applicability of factor graph methods to practical problems in deep-water bathymetric surveying and propose a metric to guide online survey planning in the future.

1.1.1 Mapping Sensors

A variety of sensors may be used to measure the shape of the seafloor from a UUV, including Multi-Beam Echo Sounders (MBES), stereo cameras and structured light systems. MultiBeam Echo Sounders (MBES) are the typical mapping sensor. They provide excellent coverage in a variety of water conditions and ranges but suffer from resolution limited to the order of centimeters. Stereo cameras produce very high-resolution surfaces and image data, but are limited by water quality and the poor propagation of light in water. Structured light systems are an emerging but unproven technology that attempts to provide the resolution of stereo systems with improved robustness to water quality issues.

Figure 2: Data from a 1.35MHz ROV-mounted MBES survey of the U-166 reduced to 2.5cm grid cells (colored) compared with data from a 30kHz ship-mounted MBES in 1300m of water reduced to 15m grid cells (Grey tiles). While using fundamentally similar methods and beam shapes, the differences in frequency make each system uniquely suited to surveying at a particular scale.

Acoustic-based systems have long been the standard tool for surveying from ships and UUVs. Unlike electromagnetic energy, acoustic energy propagates well through seawater. Although propagation is robust to particulates in the water col-

umn, it is strongly influenced by the sound velocity profile [24]. Most propagation effects can be corrected for with post-processing, and the short distances common in ROV surveying limit the impact of these phenomena. The long wavelength of acoustic waves in seawater (7.6mm at 200kHz) limits the resolution of acoustic sensors.

Acoustic absorption and scattering in seawater are frequency-dependent [25]. Multibeam systems may be constructed to operate at low frequency, long range, and limited resolution; high frequency, short range, and higher resolution; or anywhere in between. For example, the E/V *Nautilus* is equipped with a 30kHz Kongsberg EM302 MBES that can provide resolutions of 25-50m at ranges up to 8,000m. The Remotely Operated Vehicle ROV *Hercules* has a 1.35MHz Blueview MB1350-90 that measures with 0.025-0.050m resolution at ranges up to 15m. These systems output similar types of data at vastly different scales (Figure 2).

Figure 3: Raw ping data from the Blueview MB1350-90 mounted on ROV *Hercules*. Range rings (white) are in meters. The detected seafloor is shown with red marks.

MBES systems work by transmitting a narrow slice of acoustic energy in a transmit beam perpendicular to the sensor’s direction of motion. Anything in the

water column scatters acoustic energy back to the sensor. The sensor’s receive array then beams the incoming return into beams by beamforming slices parallel to the direction of motion. By examining the backscatter over time and receive direction, the sensor builds a picture of acoustic backscatter as a function of distance from the sensor and angle from centerline, as shown in Figure 3. The sea floor is then assumed to be the hard object with a solid return closest to the sensor, and its range is recovered using one of several bottom detection algorithms [26, 27] .

Each ping from a MBES produces a line of soundings along the seafloor. Most multibeam systems can measure a swath from 45° to 150° wide. While these soundings are often quite accurate, they convey information about a very small portion of the seafloor. This makes the SLAM problem much harder because it is very difficult to use a single ping to measure the location of the UUV based on previous measurements.

Although most commonly used to take seafloor images, stereo cameras also produce bathymetric data. By configuring two cameras with overlapping fields of view and simultaneous exposures, it is possible to triangulate portions of the seafloor and recover a distance from the camera to the seafloor. Stereo images may be processed using either dense or sparse methods. Dense methods attempt to match each pixel or image patch in one image to another image patch in the corresponding image. This method can produce extremely accurate, dense point-clouds with hundreds of thousands of data points per image. However, it often fails near edges, occlusions, and regions of poor texture. Sparse stereo processing methods identify individual features and match them between images. A variety of methods to identify and features may be used, including the Scale Invariant Feature Transform (SIFT) and Speeded Up Robust Features (SURF). Sparse stereo generates a much smaller number of points per image pair, often on the order of 10’s

to 100's, but there is often a high probability these features may be matched between multiple stereo pairs that overlap [28, 29, 30]. Cameras image a rectangular portion of the seafloor. Registering overlapping stereo image pairs is significantly aided by the large amount of information spread over a wide 2D region of the seafloor, in contrast to the single line of soundings from MBES systems.

Stereo cameras provide a wealth of information, but are severely limited by light's poor propagation through seawater. A number of effects degrade stereo imaging performance underwater. Absorption typically limits the maximum practical range of cameras with artificial illumination to altitudes less than 5m in most situations. Particulates in the watercolumn scatter light back at the camera and degrade image quality [31, 32, 33, 34]. The backscatter problem is particularly evident in coastal regions.

Figure 4: Single track lines from a 1.35MHz MultiBeam Echo Sounder (top) and a structured-light line scanner (bottom). 1.35 MHz is a very high multibeam frequency and may be considered typical of the highest-resolution MBES systems commercially available. Both datasets are shown on a 5mm grid. Calibration remains a key challenge for the structured-light sensor.

The application of structured light sensors to bathymetric measurement grew out of efforts to match the resolution of stereo techniques with a sensor that is more robust to backscatter. The structured-light sensor consists of a sheet laser and camera mounted on a rigid rig of known geometry. This arrangement minimizes the volume of water that is both illuminated and in the camera's imaging

frustum in an attempt to reduce the impact of backscatter. As the structured-light system relies strictly on geometry and backscatter to recover 3D structure, rather than recognizing common areas or features between images, the sensor can still produce highly-accurate results on surfaces with uniform texture. The density of measurements from the structured light sensors can be very high. Like MBES systems, the structured-light sensor on *Hercules* produces a single line of data under the vehicle [35].

1.1.2 The Navigation Problem

Navigation is the greatest challenge in underwater surveying. Radio waves do not propagate underwater, limiting the use of GPS to infrequent surface measurements. Deep-ocean surveys are typically conducted in water depths of more than 1km and may require hours to descend from the surface and its available GPS measurements. A variety of methods have been used to allow UUVs to navigate underwater.

One common solution is to integrate the vehicle’s velocity over time to produce a position estimate, known as “dead-reckoning” (DR). This method is similar to, but more accurate than, the well-known method of inertial navigation. In practice, a dead-reckoning navigation solution typically includes inertial measurements. Vehicle velocities are measured with a Doppler Velocity Log (DVL) rigidly mounted on the UUV. The DVL measurements must be rotated from a vehicle-oriented reference frame to a global reference frame. This transform requires accurate attitude measurement. Attitude error is often a larger source of error in the final position than the velocity measurements themselves [36, 37].

Error in the global velocity estimate from both attitude and velocity measurement adds up over time, resulting in a navigation solution that is extremely precise in the short term but becomes increasingly inaccurate over time. The posi-

tion given is relative to the start of the integration and is not geo-referenced. While it is possible to use inertial / DR solutions in the water column, several phenomena significantly degrade accuracy. Measuring velocity relative to the seafloor also requires the seafloor to be within range of the vehicle's DVL. Typical UUV-mounted DVLs have maximum ranges of 50-600m, far less than common deep-ocean survey depths of 1000-4000m. While it is possible to measure velocity relative to the water column, water is often moving due to currents. Descents to over 1000m typically take over an hour, which is enough time for these effects to degrade a surface fix.

Another navigation method is to use an acoustic positioning system. Two types of acoustic positioning are commonly in use. Long BaseLine (LBL) systems rely on ranges between the UUV and multiple seafloor-moored beacons and use trilateration to calculate a position. LBL systems can provide meter-level position over areas of many square kilometers or sub-meter accuracy over smaller areas. Deploying beacons is time consuming, expensive, and may be prohibited by law in ecologically- or archaeologically-sensitive areas. These limitations have motivated a strong interest in alternatives.

The other popular acoustic position method is Ultra-Short BaseLine (USBL) positioning. USBL systems use a much smaller transducer array to measure the position of the UUV relative to a support vessel. This relative position may be combined with the support ship's attitude and GPS-based position data to produce an estimate of the UUV's global position. While relatively easy to deploy and use, USBL positioning accuracy decreases with depth as the effect of small angular errors increases. A quarter degree error in measuring the depression angle to the vehicle introduces a positioning error of only 11cm at a depth of 25m, but nearly 11m at a depth of 2500m. In addition to accuracy limitations of the USBL array itself, angular error is also introduced by the need to measure the attitude

of the USBL receive array. In order to recover the position of the UUV in a geodetic reference frame, it is necessary to rotate the position vector of the UUV from a ship-oriented reference frame to a globally-oriented reference frame. This transform requires the ship’s roll, pitch, and heading at the time the ping was received. Measurement error and latency in processing and communicating the ship’s attitude to the USBL system all contribute significantly to error in the final position.

An example of USBL performance in the field at 2600m depth is shown in Figure 5. Fusing sensors with different error characteristics is a classic problem in robotics. One traditional solution is to use an Extended Kalman Filter (EKF) to fuse DR measurements with the USBL fixes [38]. The EKF solution can be corrupted by outliers, and biases in the USBL fixes can further corrupt the resulting solution as shown by the blue trackline in the top of Figure 5. Even if these issues were resolved, the USBL system may be too inaccurate for use at depth. A 2D-histogram of the distance between the USBL fixes and DR track is shown in on the bottom-left of Figure 5. These distances have a standard deviation of approximately 10m in each x and y . The solution is better-constrained in depth. The depth difference has a standard deviation of 0.5m and is shown on the bottom-right of Figure 5.

Much like using a MBES from the surface, variations in sound speed throughout the water column may cause the USBL pings to travel along a curved path between the vehicle and receive array. Much of this distortion can be corrected using ray-tracing models. These models require accurate information about the sound velocity profile to function accurately. While maintaining an accurate sound velocity profile is often a priority during shipborne MBES surveying, during UUV operations it is typically only measured during the vehicle’s descent and ascent.

Figure 5: An example of deep-ocean USBL navigation performance at Tempus Fugit in 2600m of water. The navigation tracks were aligned by subtracting the mean position difference in x , y , and z .

Current sound-speed profiling instruments require a cable or thin wire back to the ship. The risk of the profiling instrument colliding with the UUV or that wire becoming entangled with the UUV is significant enough that deploying conventional profiling methods during a UUV dive is uncommon. Future profiling technologies that communicate acoustically and can be deployed in a manner that prevents collision with the primary UUV may provide better data in the future [39, 40].

Recently, dead-reckoning has been augmented with additional sensor measurements. Several variations on the concept of LBL-style acoustic range measurements have become a popular addition to acoustic communications. These systems provide a single range measurement from a GPS-equipped surface beacon. With the increased availability of highly-accurate timing sources, synchronous one-way travel time from a collection of surface beacons has also proven a useful method [41, 42, 43]. While this research is promising, surface buoys and other one-way navigation sources present another set of systems to maintain and recover.

1.2 Data Sources

Data for this thesis was collected using the ROV *Hercules* deployed from the E/V *Nautilus*. *Hercules* is part of a two-body ROV system that includes the depressor *Argus*. *Argus* provides both overhead lighting and a top-down view of *Hercules*. Both vehicles are shown in Figure 6.

All surveys are conducted by *Hercules*. The navigation sensors mounted on *Hercules* are summarized in Table 1 and include the usual ROV dead reckoning (DR) sensor suite [44]. Operational navigation during surveys uses a combination of Doppler Velocity Log (DVL) based dead-reckoning in real-time with DVLNAV [45] and the USBL system.

Hercules is equipped with several imaging sensors for mapping. Bathymetric mapping is done with a BlueView MB-1350-90 1.35MHz multibeam sonar with

(a) *Hercules* being launched

(b) *Argus* during launch

Figure 6: *Hercules* and *Argus* during launch. *Argus* usually flies above *Hercules* to provide overhead lighting and situational awareness. the two vehicles are connected by the yellow tether visible in both images.

INSTRUMENT	VARIABLE	SAMPLE RATE	PRECISION
Ixsea Octans North-seeking gyro	heading, pitch, roll	10Hz	0.1°-0.01°
RDI Doppler Velocity Log	body-frame velocity	4-10Hz	0.3% or better
Paroscientific Depth sensor	depth	2Hz	<2cm
Tracklink TL-5000 USBL	x, y, z	0.1Hz	10m

Table 1: Navigation sensors on ROV *Hercules*

a nominal 90° footprint. The BlueView is typically processed to produce 256 overlapping beams with a 1° beamwidth. All beams are arrayed in a fan under the vehicle with a fixed transmit beam. In addition, *Hercules* has a stereo pair of Allied Vision Tech GC1380 1.4 megapixel cameras. An experimental structured-light laser system with a 532nm green sheet laser and a third GC1380 camera is the most recent addition to the imaging suite. Although initial results suggest sub-millimeter precision is possible, concerns about calibration quality keep this sensor from being a major focus of this work. The mounting location and footprint of these sensors is shown in Figure 1.

Surveys are typically conducted with back-and-forth survey lines in a “mowing

the lawn” pattern. Survey line spacing is determined by the required overlap and is commonly 1-2m at a survey altitude of 3-5m. Survey speeds also vary from survey to survey, but are typically 10-20cm/s.

1.3 Simultaneous Localization and Mapping

Simultaneous Localization and Mapping (SLAM), methods attempt to overcome the limitations of current UUV navigation methods by using mapping data collected by the UUV to further constrain the navigation. SLAM is a form of terrain-aided navigation that does not require an *a priori* map of the environment. Terrain-aided navigation estimates UUV poses using available sensor data and odometry; SLAM solves for UUV poses and a map using the same data.

SLAM has been applied to many different problems since the mathematical foundations described in Smith and Cheeseman [46, 47]. Early research focused on identifying global features in the environment and directly estimating their locations as a map. These feature-based estimation problems have been solved with Extended Kalman Filters [48, 49, 50], Sparse Extended Information Filters [51], and, most recently, factor graphs [21]. The underwater environment does not usually provide the sort of easily-identifiable features that feature-based SLAM methods require. Globally-unique features are rare in natural seafloor settings.

Alternatively, pose-based SLAM methods use local features to estimate a constraint between two vehicle poses. The most successful examples of this underwater have been visual SLAM methods [8, 51, 52, 29, 53, 54]. These algorithms use local features from overlapping image sets to estimate the relative offset between non-sequential vehicle poses. Corner-based image features and descriptors, such as the Scale Invariant Feature Transform (SIFT) or Speeded-Up Feature Transform (SURF) are the most commonly used.

Natural underwater bathymetry does not usually include sharp corners suit-

able for use with gradient-based feature descriptors. Unlike photographic methods, a single measurement from a MBES or structured-light laser sensor can not provide a well-conditioned constraint between vehicle poses.

Two approaches have been used to overcome this limitation. The first technique, known as Submap-SLAM, is to group consecutive pings into submaps and will be introduced in Section 1.3.1. The second use a non-parametric estimation framework to address the non-Gaussian behavior of constraints between pings. Barkby’s BP-SLAM uses a Rao-Blackwellized particle filter to represent the non-Gaussian x, y position estimates that result from these weak constraints [55, 56]. BP-SLAM has proven fast and robust, but produces maps that are less self-consistent than those produced by Submap-SLAM. Particle filter methods become computationally-intensive as additional state variables are estimated. For this reason, BP-SLAM uses particles only for the x and y states and represents other states, including depth and altitude, using a Kalman filter. Finally, it is not immediately obvious how to couple the non-parametric BP-SLAM with parametric structure-from-motion techniques commonly used with visual SLAM. BP-SLAM is well-suited to real-time navigation but is less optimal for post-processing.

1.3.1 Submap-SLAM

Submap-SLAM addresses the limited information content of individual pings by grouping consecutive pings into rigid submaps. Submaps must be small enough that dead-reckoning error within a submap is negligible compared to sensor resolution. Each submap is represented as a pointcloud relative to the submap origin. The vehicle pose at the time of the first ping is used as this origin. Submaps are assumed to be rigid at creation time and are moved around by the SLAM algorithm but not updated internally. Overlapping submaps may then be registered to provide constraints between the origin of each submaps. The collection of

dead-reckoning (DR) constraints between submaps and relative submap registration constraints is then used to produce an improved navigation solution.

Figure 7: Construction of a bathymetric submap during a survey by *Hercules*. *Hercules*’ path is shown in red, with submap borders drawn in magenta. Consecutive pings are grouped together to provide enough information to constrain the registration process.

The idea of using overlapping bathymetric survey lines to improve navigation in an ad-hoc fashion dates back at least as far as the 1980s and pre-GPS shipboard multibeam surveying [57]. Additional survey lines remain a key quality-control method on hydrographic surveys [58]. Roman [20] extended this basic concept into a formal SLAM framework using an augmented-state Extended Kalman Filter (EKF) to solve the resulting constraint network. The work presented here adopts Roman’s submap assumptions and construction techniques. Like most EKF-based SLAM methods, the original Submap-SLAM algorithm was theoretically limited

to a small number of submaps. Early attempts [59, 53, 60] to adapt Submap-SLAM to scalable SLAM algorithms produced promising results despite using a poor representation of the DR constraints. A new method to add DR constraints to a factor graph based on a state-transition model is presented in Chapter 2. This new approach for loosely-coupling the DR and SLAM estimators closely follows the approximations made by the EKF DR filter and faithfully represents the conditional probabilities output by that filter while still minimizing the number of variables that must be solved by the final SLAM estimator. These DR constraints are applicable to other problems requiring a navigation factor graph for a UUV equipped with high-quality DR sensors.

Figure 8: Submap-SLAM breaks a survey into submaps. Measurements generate constraints between submap origins from dead-reckoning odometry (blue) and submap registration (green) are used refine the navigation solution. The example here shows a set of submaps for three lines from the Monterrey A post-disturbance survey. The full survey includes 123 submaps.

Unlike particle-filter methods, factor-graph Submap-SLAM may be readily

combined with other factor graph SLAM methods. Camera data has been combined with Submap-SLAM techniques to produce a single improved navigation solution for multiple sensors in [53, 23]. In addition to providing maps from multiple sensing modalities, such sensor fusion allows for multi-sensor reconstruction techniques that choose an optimal sensor source for each grid cell [54]. Extensive study of factor graphs by the robotics community extends beyond simply adding more sensor types. Anchor nodes [61] provide a mechanism for merging multi-session surveys. Multi-session surveys usually result from multiple survey efforts over a single, unchanging site. Chapter 3 uses anchor nodes to merge Submap-SLAM results from surveys before and after excavation of the Monterrey A shipwreck to detect changes made by the excavation.

Factor graph Submap-SLAM is significantly more computationally efficient than the augmented-state EKF Submap-SLAM of [21, 22], typically running in minutes instead of hours. Submap registration remains a computationally-intense task. The submap registration process could be run in parallel by processing each registration independently. With simple submap registration techniques and maps with tens of millions of soundings, like those presented here, factor graph Submap-SLAM can run all processing steps in an hour when running on commodity hardware.

Submap-SLAM does suffer from two drawbacks common to most parametric SLAM frameworks. First, the measurement models must be approximately linear and Gaussian [16]. Factor graph Submap-SLAM addresses non-linearities by using the non-linear factor graph introduced by Dellart and Kaess [21]. The fully-nonlinear measurement models are explicitly represented in the factor graph. The entire factor graph is then dynamically re-linearized during the optimization process. While the Gaussian noise approximation is usually acceptable for physical

sensors, it may be less applicable to registration outputs. The second significant drawback to factor graph SLAM is that it is not robust to poor data association. In the context of Submap-SLAM, poor data association results from registration errors. Common registration errors include matching submaps based on a local minima of the alignment metric instead of the global minima, matching two similar submaps that do not actually overlap, and underestimating the true uncertainty of the registration result. Even a single registration error can add more error to the final navigation solution than was present in the original dead-reckoning. A more complete discussion of this phenomena and a method to identify seafloor regions less likely to produce submaps that suffer from this effect is introduced in Chapter 4.

List of References

- [1] C. R. German, A. Bowen, M. L. Coleman, D. L. Honig, J. A. Huber, M. V. Jakuba, J. C. Kinsey, M. D. Kurz, S. Leroy, J. M. McDermott, B. M. de Lpinay, K. Nakamura, J. S. Seewald, J. L. Smith, S. P. Sylva, C. L. Van Dover, L. L. Whitcomb, and D. R. Yoerger, “Diverse styles of submarine venting on the ultraslow spreading mid-cayman rise,” *Proceedings of the National Academy of Sciences*, vol. 107, no. 32, pp. 14 020–14 025, 2010. [Online]. Available: <http://www.pnas.org/content/107/32/14020.abstract>
- [2] T. Bridge, T. Done, R. Beaman, A. Friedman, S. Williams, O. Pizarro, and J. Webster, “Topography, substratum and benthic macrofaunal relationships on a tropical mesophotic shelf margin, central great barrier reef, australia,” *Coral Reefs*, vol. 30, pp. 143–153, 2011, 10.1007/s00338-010-0677-3. [Online]. Available: <http://dx.doi.org/10.1007/s00338-010-0677-3>
- [3] D. F. Coleman, R. D. Ballard, J. B. Newman, T. S. Gregory, J. Howland, H. Singh, C. Roman, O. Pizarro, B. Foley, R. Eustice, A. Can, S. Webster, D. Piechota, C. Giangrande, C. Ward, R. Horlings, K. McBride, and M. J. Durbin, *Archaeological Oceanography*, R. D. Ballard, Ed. Princeton University Press, 2008.
- [4] D. Coleman, R. Ballard, and T. Gregory, “Marine archaeological exploration of the black sea,” in *OCEANS 2003. Proceedings*, vol. 3, Sept. 2003, pp. 1287–1291.

- [5] D. Coleman, J. Newman, and R. Ballard, "Design and implementation of advanced underwater imaging systems for deep sea marine archaeological surveys," *OCEANS 2000 MTS/IEEE Conference and Exhibition*, vol. 1, pp. 661–665, 2000.
- [6] O. Pizarro, S. Williams, and J. Colquhoun, "Topic-based habitat classification using visual data," in *OCEANS 2009 - EUROPE*, May 2009, pp. 1–8.
- [7] S. Grilli, C. O'Reilly, J. Harris, T. Bakhsh, B. Tehranirad, S. Banihashemi, J. Kirby, C. Baxter, T. Eggeling, G. Ma, and F. Shi, "Modeling of smf tsunami hazard along the upper us east coast: detailed impact around ocean city, md," *Natural Hazards*, vol. 76, no. 2, pp. 705–746, 2015. [Online]. Available: <http://dx.doi.org/10.1007/s11069-014-1522-8>
- [8] S. Williams, O. Pizarro, M. Jakuba, I. Mahon, S. Ling, and C. Johnson, "Repeated auv surveying of urchin barrens in north eastern tasmania," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, May 2010, pp. 293–299. [Online]. Available: <http://www.cas.edu.au/content.php/237.html?publicationid=659&displaypage=3>
- [9] M. L. Brennan, "The disarticulation of ancient shipwreck sites by mobile fishing gear: A case study from the southeast aegean sea," *INA Quarterly*, vol. 36(4), pp. 6–7, 2010.
- [10] M. Guarinello, E. Shumchenia, and J. King, "Marine habitat classification for ecosystem-based management: A proposed hierarchical framework," *Environmental Management*, vol. 45, pp. 793–806, 2010, 10.1007/s00267-010-9430-5. [Online]. Available: <http://dx.doi.org/10.1007/s00267-010-9430-5>
- [11] National Oceanic and Atmospheric Administration. "Extended continental shelf project." Webpage. 2012. [Online]. Available: <http://continentalshelf.gov/>
- [12] J. H. Moore, *The New Practical Navigator*, 2nd ed. Edmund M. Blunt, 1800.
- [13] N. Bowditch, *The American Practical Navigator*. National Imagery and Mapping Agency, 2002.
- [14] R. Enns and D. Morrell, "Terrain-aided navigation using the viterbi algorithm," *Journal of Guidance, Control, and Dynamics*, vol. 18, no. 6, pp. 1444–1449, Nov. 1995. [Online]. Available: <http://dx.doi.org/10.2514/3.21566>
- [15] S. Williams, G. Dissanayake, and H. Durrant-Whyte, "Towards terrain-aided navigation for underwater robotics," *Advanced Robotics*, vol. 15, no. 5, pp. 533–549, 2001. [Online]. Available: <http://dx.doi.org/10.1163/156855301317033559>

- [16] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT Press, 2006.
- [17] K. Anonsen and O. Hagen, “Terrain aided underwater navigation using pockmarks,” in *OCEANS 2009, MTS/IEEE Biloxi - Marine Technology for Our Future: Global and Local Challenges*, Oct 2009, pp. 1–6.
- [18] R. C. Kollanda, “Terrain relative navigation using bathymetric slope profile with a rao-blackwellized particle filter for long range auvs,” Master’s thesis, University of Rhode Island, 2012.
- [19] C. Roman and H. Singh, “Consistency based error evaluation for deep sea bathymetric mapping with robotic vehicles,” in *Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on*, May 2006, pp. 3568–3574.
- [20] C. Roman and H. Singh, “A self-consistent bathymetric mapping algorithm,” *Journal of Field Robotics*, vol. 24, no. 1-2, pp. 23–50, 2007. [Online]. Available: <http://dx.doi.org/10.1002/rob.20164>
- [21] F. Dellaert and M. Kaess, “Square Root SAM: Simultaneous localization and mapping via square root information smoothing,” *Intl. J. of Robotics Research (IJRR)*, vol. 25, no. 12, pp. 1181–1204, Dec 2006.
- [22] M. Kaess, A. Ranganathan, and F. Dellaert, “iSAM: Incremental smoothing and mapping,” *IEEE Trans. on Robotics (TRO)*, vol. 24, no. 6, pp. 1365–1378, Dec 2008.
- [23] C. G. Kunz, “Autonomous underwater vehicle navigation and mapping in dynamic, unstructured environments,” Ph.D. dissertation, Massachusetts Institute of Technology and Woods Hole Oceanographic Institution Joint Program, 2011.
- [24] F. Jensen, W. Kuperman, M. Porter, and H. Schmidt, *Computational Ocean Acoustics*. Springer, 2011.
- [25] Kinsler, Frey, Coppens, and Sanders, *Fundamentals of Acoustics*, 4th ed. John Wiley & Sons, 1999.
- [26] L. Yang and T. Taxt, “Multibeam sonar bottom detection using multiple subarrays,” in *OCEANS ’97. MTS/IEEE Conference Proceedings*, vol. 2, Oct 1997, pp. 932–938 vol.2.
- [27] L3 Communications SeaBeam Instruments. “Multibeam sonar theory of operation.” 2000. [Online]. Available: <http://www.ldeo.columbia.edu/res/pi/MB-System/sonarfunction/SeaBeamMultibeamTheoryOperation.pdf>

- [28] S. Lanser and T. Lengauer, "On the selection of candidates for point and line correspondences," in *In International Symposium on Computer Vision*. IEEE Computer Society Press, 1995, pp. 157–162.
- [29] I. Mahon, S. Williams, O. Pizarro, and M. Johnson-Roberson, "Efficient view-based slam using visual loop closures," *Robotics, IEEE Transactions on*, vol. 24, no. 5, pp. 1002–1014, Oct. 2008.
- [30] S. Negahdaripour and H. Madjidi, "Stereovision imaging on submersible platforms for 3-d mapping of benthic habitats and sea-floor structures," *Oceanic Engineering, IEEE Journal of*, vol. 28, no. 4, pp. 625–650, Oct. 2003.
- [31] J. Jaffe, K. Moore, J. McLean, and M. Strand, "Underwater imaging: Status and prospects," *Oceanography*, vol. 14, no. 3, pp. 66–76, 2001.
- [32] Y. Schechner and N. Karpel, "Clear underwater vision," in *Computer Vision and Pattern Recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE Computer Society Conference on*, vol. 1, 2004, pp. 536–543.
- [33] H. Singh, C. Roman, O. Pizarro, R. Eustice, and A. Can, "Towards high-resolution imaging from underwater vehicles," *The International journal of robotics research*, vol. 26, no. 1, pp. 55–74, 2007.
- [34] H. Zheng, B. Zheng, G. Ji, Z. Guo, and Y. Sun, "Image enhancement of underwater target detection by inhomogeneous illumination," in *OCEANS, 2012 - Yeosu*, May 2012, pp. 1–4.
- [35] G. Inglis, C. Smart, I. Vaughn, and C. Roman, "A pipeline for structured light bathymetric mapping," in *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, Oct 2012, pp. 4425–4432.
- [36] L. Whitcomb, D. Yoerger, and H. Singh, "Advances in doppler-based navigation of underwater robotic vehicles," in *Robotics and Automation, 1999. Proceedings. 1999 IEEE International Conference on*, vol. 1, 1999, pp. 399–406 vol.1.
- [37] G. Troni and L. L. Whitcomb, "Advances in in situ alignment calibration of doppler and high/low-end attitude sensors for underwater vehicle navigation: Theory and experimental evaluation," *Journal of Field Robotics*, vol. 32, no. 5, pp. 655–674, 2015. [Online]. Available: <http://dx.doi.org/10.1002/rob.21551>
- [38] M. S. Grewal and A. P. Andrews, *Kalman Filtering: Theory and Practice Using Matlab*, 2nd ed. Wiley & Sons, 2001.
- [39] B. Amaral, "The hydrodynamic optimization of a cable-flying vehicle," Master's thesis, University of Rhode Island, 2012.

- [40] L. Logan, “Performance analysis of an underwater wire flying profiling vehicle,” Master’s thesis, University of Rhode Island, 2014.
- [41] R. Eustice, L. Whitcomb, H. Singh, and M. Grund, “Experimental results in synchronous-clock one-way-travel-time acoustic navigation for autonomous underwater vehicles,” in *Robotics and Automation, 2007 IEEE International Conference on*, April 2007, pp. 4257–4264.
- [42] S. Webster, R. Eustice, H. Singh, and L. Whitcomb, “Preliminary deep water results in single-beacon one-way-travel-time acoustic navigation for underwater vehicles,” in *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, Oct 2009, pp. 2053–2060.
- [43] B. Crosbie, M. Anderson, E. Wolbrecht, J. Canning, and D. Edwards, “Synchronous navigation of auvs using whoi micro-modem 13-bit communications,” in *OCEANS 2010*, Sept 2010, pp. 1–5.
- [44] J. C. Kinsey, R. M. Eustice, and L. L. Whitcomb, “A survey of underwater vehicle navigation: Recent advances and new challenges,” in *IFAC Conference of Manoeuvring and Control of Marine Craft*, Lisbon, Portugal, September 2006, invited paper.
- [45] J. C. Kinsey and L. L. Whitcomb, “Preliminary field experience with the dvlnav integrated navigation system for oceanographic submersibles,” *Control Engineering Practice*, vol. 12, no. 12, pp. 1541 – 1549, 2004.
- [46] R. C. Smith and P. Cheeseman, “On the representation and estimation of spatial uncertainty,” *The International Journal of Robotics Research*, vol. 5, no. 4, pp. 56–68, 1986. [Online]. Available: <http://ijr.sagepub.com/content/5/4/56.abstract>
- [47] R. Smith, M. Self, and P. Cheeseman, “Estimating uncertain spatial relationships in robotics,” in *Robotics and Automation. Proceedings. 1987 IEEE International Conference on*, vol. 4, Mar 1987, pp. 850–850.
- [48] J. J. Leonard and H. F. Durrant-Whyte, “Simultaneous map building and localization for an autonomous mobile robot,” in *Intelligent Robots and Systems’ 91. Intelligence for Mechanical Systems, Proceedings IROS’91. IEEE/RSJ International Workshop on*. Ieee, 1991, pp. 1442–1447.
- [49] H. Durrant-Whyte and T. Bailey, “Simultaneous localization and mapping: part i,” *IEEE Robotics & Automation Magazine*, vol. 13, no. 2, pp. 99–110, 2006.
- [50] T. Bailey and H. Durrant-Whyte, “Simultaneous localization and mapping: part ii,” *IEEE Robotics & Automation Magazine*, vol. 13, no. 3, pp. 108–117, 2006.

- [51] R. M. Eustice, O. Pizarro, and H. Singh, “Visually augmented navigation for autonomous underwater vehicles,” *IEEE Journal of Oceanic Engineering*, vol. 33, no. 2, pp. 103–122, 2008.
- [52] O. Pizarro, R. Eustice, and H. Singh, “Large area 3-d reconstructions from underwater optical surveys,” *Oceanic Engineering, IEEE Journal of*, vol. 34, no. 2, pp. 150–169, Apr. 2009. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4918977
- [53] C. Kunz and H. Singh, “Map building fusing acoustic and visual information using autonomous underwater vehicles,” *Journal of Field Robotics*, vol. 30, no. 5, pp. 763–783, 2013. [Online]. Available: <http://dx.doi.org/10.1002/rob.21473>
- [54] G. Inglis, “Hybrid optical acoustic seafloor mapping,” Ph.D. dissertation, University of Rhode Island, 2013.
- [55] S. Barkby, S. B. Williams, O. Pizarro, and M. V. Jakuba, “A featureless approach to efficient bathymetric slam using distributed particle mapping,” *Journal of Field Robotics*, vol. 28, no. 1, pp. 19–39, 2011. [Online]. Available: <http://dx.doi.org/10.1002/rob.20382>
- [56] S. Barkby, S. B. Williams, O. Pizarro, and M. V. Jakuba, “Bathymetric particle filter slam using trajectory maps.” *I. J. Robotic Res.*, vol. 31, no. 12, pp. 1409–1430, 2012. [Online]. Available: <http://dblp.uni-trier.de/db/journals/ijrr/ijrr31.html#BarkbyWPJ12>
- [57] C. Nishimura and D. Forsyth, “Improvements in navigation using seabeam crossing errors,” *Marine Geophysical Researches*, vol. 9, no. 4, pp. 333–352, 1988. [Online]. Available: <http://dx.doi.org/10.1007/BF00315004>
- [58] *Field Procedures Manual*, Office of Coast Survey Std., 2014. [Online]. Available: http://www.nauticalcharts.noaa.gov/hsd/fpm/2014_FPM_Final.pdf
- [59] I. Vaughn, “Microbathymetry using self-contained navigation and simultaneous localization and mapping,” Master’s thesis, University of Rhode Island, 2012.
- [60] M. A. VanMiddlesworth, “Toward autonomous underwater mapping in partially structured 3d environments,” Master’s thesis, Massachusetts Institute of Technology, 2014. [Online]. Available: <hdl.handle.net/1721.1/87791>
- [61] J. McDonald, M. Kaess, C. Cadena, J. Neira, and J. J. Leonard, “Real-time 6-dof multi-session visual slam over large-scale environments,” *Robot. Auton. Syst.*, vol. 61, no. 10, pp. 1144–1158, Oct. 2013. [Online]. Available: <http://dx.doi.org/10.1016/j.robot.2012.08.008>

CHAPTER 2

Representation of Dead Reckoning Odometry in Factor Graphs

2.1 Introduction

Dead reckoning (DR) is a fundamental part of most SLAM methods. It is unusual in SLAM for any single measurement to provide a complete observation of the vehicle’s state. A DR solution is used to link vehicle states so that the partial observations can be combined to form a complete estimate. This process is conceptually similar to the “running fix” in traditional navigation, in which incomplete information about a vehicle’s position is advanced in time using integrated velocity information so that it may be combined with additional measurements that occur at a later time [1]. In SLAM methods, dead reckoning is used to advance an estimate of the vehicle’s state and the uncertainty of that estimate in time. With most sensors, SLAM requires only the pose, or position and attitude, of the vehicle and not the complete vehicle state.

Dead reckoning requires a full vehicle state that includes both pose and any other variables, such as velocity, required by the particular vehicle process model in use. As shown here, any linearizable process model may be used. A simple constant-velocity model presented in Section 2.2 and is used for this chapter. The DR filter requires that this full state is complete, meaning that additional states provide no information that would help predict the vehicle’s future motion [27]. Given the previous state, a complete state is conditionally independent of all other prior states. This conditional independence is the key property that may be exploited to efficiently solve the factor graph.

In contrast to traditional navigation methods, SLAM methods require updating an estimate of the vehicle’s trajectory over its complete history rather than just the most recent pose. This way, new information from another photograph,

bathymetric data, or some other observation may be used to constrain against and refine the pose estimate of previously-collected observations. A vehicle pose history is only necessary for sensors that make measurements relative to other poses or a landmark in the environment.

The total number of variables to be estimated presents a practical performance limit on the maximum size of a dataset. Individual states are collections of variables that describe the vehicle at a specific moment in time. SLAM requires only the vehicle’s pose variables, or position and attitude, at the time of each photograph, submap origin, or other SLAM-related point in time. Each pose to be stored requires six variables. The complete DR states include pose and additional variables such as velocity. The constant-velocity model used here has 12 variables per state. In addition, conventional DR estimation requires a state at the time of each measurement, including high-rate DR sensors such as the compass, DVL, and depth sensor.

Reducing the total number of variables to be estimated may be accomplished through two methods. The first, demonstrated here, is to store only SLAM variables in the SLAM estimator by combining a sequence of DR measurements into a single constraint. This method can reduce the number of variables to estimate by one or two orders of magnitude. The second approach is to store only a subset of the variables required by DR for each state. With the 12-state constant-velocity model, this can at most reduce the number of variables to estimate by a factor of two. More importantly, it will be shown in Section 2.4.1 that the number of variables per state cannot be reduced without sacrificing the sparsity that factor graph exploit to outperform EKF algorithms. Although these performance gains are modest, eliminating all non-pose variables from each state allows the factor graph to be built without regard to the internal details of the DR filter. This

property is especially attractive when using a proprietary or commercial DR filter with internals that cannot be divulged for business or regulatory reasons. While not addressed in this chapter, reducing the size of each state is a problem worthy of future research. This chapter focuses on reducing the number of states represented in the SLAM estimator.

In this chapter, SLAM states are full vehicle states, including pose and velocity, that must be maintained to compute a SLAM navigation solution. The relationship between consecutive SLAM states must be estimated using integrated velocity measurements. The traditional UUV DR Kalman filter [2, 3, 4, 5] requires additional DR states at the time of every measurement. This filter is modified to combine all of these measurements into a single linearized constraint, effectively marginalizing out the unwanted DR-only states between consecutive SLAM states. Every SLAM state must be directly represented in the final optimization as either a delayed state in an augmented-state Kalman filter or node in a factor graph.

Estimating only SLAM states could provide an efficiency improvement as absolute measurements are a very large fraction of the total measurements. With the ROV *Hercules*, the attitude sensor, DVL, and pressure-based depth sensor produce approximately 20 samples per second. While the multibeam sonar can run at up to 15Hz, its pings are combined into submaps that usually require one SLAM state no more often than once per minute.

2.1.1 Problem Statement

The goal of the SLAM process is to produce an estimate of the vehicle's pose over time. The pose at a given time $p(t)$ is defined as the 6-tuple of position and attitude in a global reference frame.

$$\mathbf{p}(t) = [x, y, z, \theta, \phi, \psi]^T$$

DR estimation requires additional state variables. The vehicle model (Section

2.2) used here adds linear and rotational velocities to produce a complete 12-element state $\mathbf{x}(t)$.

$$\mathbf{x}(t) = [x, y, z, \theta, \phi, \psi, \dot{x}, \dot{y}, \dot{z}, \dot{\theta}, \dot{\phi}, \dot{\psi}]^T$$

Other DR filters may include other state variables such as sensor offsets and biases. The vehicle state must obey the Markov property where future states depend only on the current state and not prior states.

Conventional Kalman filtering computes a predicted state at the time of each measurement. Using only states needed by the SLAM estimator requires combining multiple DR measurements into a single constraint that faithfully represents the underlying uncertainty. Terrestrial SLAM often avoids this problem by using purely relative wheel odometry constraints that are conditionally independent between any two times. Wheel odometry is unavailable on survey UUVs. Instead, an underwater vehicle uses a mixture of relative (DVL) and absolute (attitude, depth) constraints between consecutive DR states that must be approximated in the final SLAM framework. The goal of this chapter is to develop and validate such a representation.

SLAM states will be denoted $\mathbf{x}_k$ with index k arranged by time. The set of all DR states, including SLAM and non-SLAM states will be indexed with i and written as $\mathbf{x}_i$. The time of SLAM state k will be denoted t_k , and the index of that SLAM state in the complete set of states as i_k . The states $\mathbf{x}_i$ define a Markov chain such that each state $\mathbf{x}_i$ only depends on the state before it $\mathbf{x}_{i-1}$. A graphical representation of this is shown in Figure 9. As the SLAM problem attempts to estimate the state of the vehicle at the time of sonar pings that correspond to submap origins, only those nodes need to be estimated by the final SLAM problem.

The goal of this section is to devise a method to produce constraints between consecutive SLAM states $p(\mathbf{x}_{k+1}|\mathbf{x}_k)$ that incorporates all the necessary informa-

Figure 9: The portion of a SLAM graph created by DR. Only the large nodes are of interest to the resulting SLAM problem. Submap-SLAM only requires estimating the first ping of each submap, shown circled in Blue.

tion about the states $\mathbf{x}_i$ for $i_k < i < i_{k+1}$ without requiring the final SLAM factor graph to explicitly represent them.

2.1.2 Related Prior Work

The original submap-SLAM method of [6] used an augmented-state EKF that maintained the dense covariance matrix between the current vehicle state estimate and each prior pose as a joint normal PDF. While effective, this approach does not exploit the sparsity inherent to DR navigation. This original study noted a practical upper limit to map size of approximately 100 submaps, at which point the computational cost of maintaining the EKF became prohibitive. The approach presented here also uses an EKF for DR, but exploits sparsity in the Markov assumption of the vehicle dynamics model to efficiently build a factor graph. A more detailed comparison between augmented-state EKFs for DR and the presented algorithm is given in Section 2.4.1.

Like dead reckoning, inertial navigation also integrates measured states to ar-

rive at a position estimate. Inertial navigation systems are of interest for many robotic systems. Factor graphs have been applied to the inertial navigation problem by Indelman et. al. [7]. This method builds factors between nodes by integrating raw IMU rotation rate and acceleration measurements. The standard non-linear factor graph of [8] may then be used to solve for the smoothed navigation solution. Accelerometer and gyro biases may also be represented as nodes in the graph. Indelman’s method requires that integrable measurements are available between every navigation node in the factor graph. The high-rate IMU data used by Indelman is not available on the *Hercules* system, and the DVL reports samples 10-100 times slower than a standard IMU. As a result, it is impossible to guarantee that a velocity measurement will be available between any two pings. A similar approach to sensor integration was used to generate DR constraints by Eustice, Singh, and Whitcomb [9] for use with one way travel time acoustic measurements every 5 or 20 seconds. As with Indelman’s work, this approach requires the time between positions to be estimated to be large relative to the interval between integrated DVL measurements. In contrast, the method presented here uses an EKF to provide estimated links even if the DVL measurement rate is relatively slow.

The same process model used in the EKF could also be used to build a factor graph. The original presentation of factor graphs for SLAM in [8] presents a method to include the non-linear measurement models presented in Eustice [4] in a factor graph. The process noise commonly used in Submap-SLAM [10, 6, 4] was singular and cannot be inverted to find a measurement information matrix. A simple modification could be used to build a similar factor graph directly from the raw measurements. Such a graph does not reduce the number of nodes as the method to be introduced in this chapter does, but permits non-linearities to be handled in the factor-graph framework. One method to reduce the number of

factors in use is the smart factors method of Carlone [11]. Smart factors divide the complete set of nodes into target and support variables. In the formulation of this chapter, SLAM poses are target variables and all other intermediary DR states are support variables. Smart factors are able to re-linearize DR factors as new SLAM information becomes available at the cost of still loading the support variables and related factors in memory. Non-linearity arises principally from attitude error, and the excellent attitude measurement available from modern fiber optic gyrocompasses limits this performance loss.

A simpler approach to reducing the number of nodes is to assume that the DR EKF provides relative measurements of some variables and absolute measurements of others. x, y position are often assumed to be incremental measurements while roll, pitch and depth are represented as absolute measurements. Heading may be represented either way depending on the quality of the heading sensor [12]. This approach will be shown to be an approximation of the method presented here. The conditions under which this approximation is valid are explored in Section 2.4.2.

Finally, Kunz presents an ad-hoc method of interpolating sensors onto the times of DVL measurements [13, 14]. Relative pose constraints are constructed by integrating these DVL measurements similarly to [15]. Interpolation is not without its own issues. Factor graphs assume that individual measurements are statistically independent, which may not be true of two interpolated measurements. Interpolating a slow 1Hz depth sensor onto a fast 10Hz DVL measurement will produce highly correlated depth samples. This will have the effect of adding more depth information to the factor graph than is provided by the sensors and over-constraining the resulting solution in depth. Depth error is already correlated (see Section 3.4), but processing methods should avoid unnecessary correlation. The simplest approach to avoiding this issue is to process each measurement directly

without interpolation.

2.2 The Complete 12 Degree-of-Freedom Model

Dead-reckoning in Roman's Submap-SLAM [6] used the 12-state model fully explained in Eustice's Visually Augmented Navigation [16, 4]. Given the previously-introduced state vector, the continuous-time constant-velocity process model used in the original submap-SLAM [6] may be described as

$$\dot{\mathbf{x}}_t = f(\mathbf{x}_t, \mathbf{u}_t) + \mathbf{w}_t.$$

The constant-velocity process model assumes that the body-frame velocity and angular rates remain constant. The thruster inputs are considered as noise that perturbs the system. Although it ignores vehicle dynamics and functions as an integrator, this simple model is adequate for a slowly-moving vehicle and may be applied to a large number of vehicles [4]. This model may be linearized and discretized as in [4] to produce a final time-varying model to predict the next state $\mathbf{x}_{i+1}$ given state $\mathbf{x}_i$ as

$$\hat{\mathbf{x}}_{i+1} = A_i \mathbf{x}_i + B_i \mathbf{u}_i + w_i \quad (1)$$

where w_i is the independent, time-varying process noise with zero mean and covariance matrix Q_i .

Similarly, a measurement $\mathbf{z}_i$ may be represented as a potentially non-linear function of the state vector $\mathbf{x}_i$ as $\mathbf{z}_i = h(\mathbf{x}_i) + m_i$ where m_i is zero mean, independent measurement noise with covariance R_i . A description of non-linear measurement models and their Jacobians for common UUV DR instruments is given in the appendices of [4]. The measurement model may be linearized as

$$\mathbf{z}_i = h(\bar{\mathbf{x}}_i) + J_{h_i}(\mathbf{x}_i - \bar{\mathbf{x}}_i) + m_i$$

where $\bar{\mathbf{x}}_i$ is the current predicted state value at the time of measurement $\mathbf{z}_i$.

2.3 EKF-based Approximation

The Extended Kalman Filter (EKF) is a well-established method for fusing data into a complete DR solution [4, 17, 5]. It will be shown that the output of an EKF can be represented as a probabilistic linear model over the set of output states from the EKF in Section 2.3.1. This model reproduces exactly the output of the EKF, including the same linearization errors. Multiple state transitions will be combined to produce a single, linearized probabilistic state update between SLAM poses $\mathbf{x}_k$ and $\mathbf{x}_{k+1}$. This single transition has the form of a linear process model and may be represented in a factor graph [8]. The resulting factor graph requires the complete state including velocities, unlike the classic augmented-state EKF representation [6]. These velocities are necessary to preserve the Markov property, which is not required in the state-history portion of an augmented-state EKF representation. A more detailed theoretical comparison of the two is explored in Section 2.4.1. The resulting factor graph will be compared with prior techniques for a selection of navigation-related problems in Section 2.5.

2.3.1 State Transition Model

The goal of this section is to develop a state transition model between consecutive SLAM states based on the DR EKF. Traditional methods of building factor graphs use measured wheel odometry to advance the state from time k to time $k + 1$. This method provides a model for the vehicle motion and uncertainty between consecutive poses as follows.

$$\mathbf{x}_{k+1} = \bar{f}(\mathbf{x}_k, \bar{\mathbf{u}}_k) + \bar{\mathbf{w}}_k$$

Although wheel odometry is unavailable for survey UUVs, the resulting model is very attractive. Producing such a model from the DR requires combining a number of update / predict cycles from the dense DR EKF to produce a single state transition between SLAM states.

The state transition function of the dense DR EKF $f(\mathbf{x}_i, \mathbf{u}_i)$ uses a model of vehicle dynamics to predict the vehicle's state at time $i + 1$. The linearized nature of the DR filter suggests a linear state transition model of the form

$$\mathbf{x}_{k+1} = \bar{\Phi}_k \mathbf{x}_k + \bar{\mathbf{u}}_k + \bar{\mathbf{w}}_k$$

where $\bar{\Phi}_k$, $\bar{\mathbf{u}}_k$, and $\bar{\mathbf{w}}_k$ are generated based on the predict/update cycles to the DR EKF between time k and $k + 1$. These values are derived as follows.

Every predict or update operation of the DR filter produces a new state mean and covariance estimate. Some of these are predicted state mean/covariances and some are the result of measurement updates. In practice, prediction and update steps occur as required by the individual asynchronous sensor rates. Multiple prediction steps may be taken if no navigation measurements are available between pings, or multiple update steps may occur if two different types of measurements are available at the same time step. These may be treated as a series of predict / update steps [18]. Even if that were not the case, the following derivation works with an arbitrary ordering of predict / update steps. Without loss of generality, it can be assumed that the DR filter's outputs follow the classical predict / update cycle.

Prediction steps are already based on a linearized state transition model. The state transition model from the previous complete state estimate $\mathbf{x}_i$ to the next state prediction $\tilde{\mathbf{x}}_{i+1}$ is copied from equation 1 as

$$\tilde{\mathbf{x}}_{i+1} = A_i \mathbf{x}_i + B_i \mathbf{u}_i + \mathbf{w}_i \tag{2}$$

with $\mathbf{w}_i \sim \mathcal{N}(\mathbf{0}, Q_i)$.

Update steps are less obvious. The Joseph form of the covariance update suggests a state transition matrix $I - K_i H_i$, where K_i is the Kalman Gain and H_i is the linearized measurement model. The update state is then commonly written

in terms of an innovation $\mathbf{y}_i = \mathbf{z}_i - H_i \tilde{\mathbf{x}}_i$, giving

$$\begin{aligned}\mathbf{x}_i &= \tilde{\mathbf{x}}_i + K_i \mathbf{y}_i \\ \mathbf{x}_i &= \tilde{\mathbf{x}}_i + K_i (\mathbf{z}_i - H_i \tilde{\mathbf{x}}_i) = (\tilde{\mathbf{x}}_i - K_i H_i \tilde{\mathbf{x}}_i) + K_i \mathbf{z}_i \\ \mathbf{x}_i &= (I - K_i H_i) \tilde{\mathbf{x}}_i + K_i \mathbf{z}_i.\end{aligned}\tag{3}$$

Recall that the Joseph form of the update to covariance matrix P is

$$P_i = (I - K_i H_i) \tilde{P}_i (I - K_i H_i)^T + K_i R_i K_i^T$$

where R_i is the measurement covariance of measurement $\mathbf{z}_i$. It may be shown that the Joseph form follows directly from equation 3. The covariance update propagates the covariance $\tilde{P}_i$ of $\tilde{\mathbf{x}}_i$ through the transform $I - K_i H_i$ and then adds the uncertainty of the transformed measurement $K_i \mathbf{z}_i$. The update step from $\tilde{\mathbf{x}}_i$ to $\mathbf{x}_i$ may thus be represented as the following linearized update step:

$$\mathbf{x}_i = (I - K_i H_i) \tilde{\mathbf{x}}_i + K_i \mathbf{z}_i + \mathbf{v}_i$$

which may be written in the more conventional form as

$$\mathbf{x}_i = \Phi_i \hat{\mathbf{x}}_i + \Gamma_i \mathbf{u}_i + \mathbf{v}_i$$

with

$$\begin{aligned}\Phi_i &= I - K_i H_i & \Gamma_i &= K_i \\ \mathbf{u}_i &= \mathbf{z}_i & \mathbf{v}_i &\sim \mathcal{N}(\mathbf{0}, K_i R_i K_i^T).\end{aligned}$$

Multiple predict and update steps may be combined into a single state transition in the obvious way by substituting in for the predicted and updated states. For example, combining equations 2 and 3 gives a single-step linearized state transition

model from $\mathbf{x}_i$ to $\mathbf{x}_{i+1}$ as

$$\begin{aligned}
\mathbf{x}_{i+1} &= (I - K_{i+1}H_{i+1}) \tilde{\mathbf{x}}_{i+1} + K_{i+1}\mathbf{z}_{i+1} + \mathbf{v}_{i+1} \\
&= (I - K_{i+1}H_{i+1}) (A_i\mathbf{x}_i + B_i\mathbf{u}_i + \mathbf{w}_i) + K_{i+1}\mathbf{z}_{i+1} + \mathbf{v}_{i+1} \\
&= (A_i - K_{i+1}H_{i+1}A_i) \mathbf{x}_i + (B_i\mathbf{u}_i - K_{i+1}H_{i+1}B_i\mathbf{u}_i + K_{i+1}\mathbf{z}_{i+1}) \\
&\quad + ((I - K_{i+1}H_{i+1}) \mathbf{w}_i + \mathbf{v}_{i+1})
\end{aligned}$$

These combine to give the transition model between states, without an intermediate prediction step, as simply

$$\mathbf{x}_{i+1} = \bar{\Phi}_i \mathbf{x}_i + \bar{\mathbf{u}}_i + \bar{\mathbf{w}}_i \quad (4)$$

where

$$\bar{\Phi}_i = (A_i - K_{i+1}H_{i+1}A_i) \quad \bar{\mathbf{u}}_i = B_i\mathbf{u}_i - K_{i+1}H_{i+1}B_i\mathbf{u}_i + K_{i+1}\mathbf{z}_{i+1}$$

$$\bar{\mathbf{w}}_i = (I - K_{i+1}H_{i+1}) \mathbf{w}_i + \mathbf{v}_{i+1}.$$

Assuming that the process noise $\mathbf{w}_i$ and measurement noise $\mathbf{v}_i$ are independent, the distribution of $\bar{\mathbf{w}}_i$ is

$$\bar{\mathbf{w}}_i \sim \mathcal{N}\left(\mathbf{0}, (I - K_{i+1}H_{i+1}) Q_i (I - K_{i+1}H_{i+1})^T + K_{i+1}R_{i+1}K_{i+1}^T\right).$$

Although presented here for a single predict / update cycle, the process may be continued for any sequence of predict / update cycles between SLAM states. Once a state transition model between SLAM states has been computed, a new DR link may be added to the factor graph.

2.3.2 Factor Graph Representation

Factor graphs are a mathematical tool to represent probabilistic models over large numbers of variables. Variables to be estimated are represented as “nodes” and constraints between these variables are represented as “factors” [19]. Factors

encode a conditional probability constraint between two or more nodes connected to that factor. Nodes that are not connected by a factor are conditionally independent given every other node in the graph. This conditional independence may be exploited to efficiently optimize the node values for a given set of constraints [19, 8, 20, 21, 22, 23]. Any Bayesian network, such as those that arise from modeling robot navigation, can be represented as a factor graph [19, 8].

For the SLAM problem, vehicle positions to be estimated are added as nodes and measurements that constrain these nodes are added as factors [8]. Nodes are assumed to be continuous Gaussian random variables. Factors may be non-linear, but must be linearized about the current estimate for all nodes they are connected to before the graph can be optimized. A factor is typically “linearized” by finding a Gaussian approximation to the underlying PDF. In general, factors may be re-linearized during the optimization.

In submap SLAM as presented here, full states with 12 variables each are added as nodes for each SLAM state. DR factors are added between consecutive nodes, and relative pose factors are added for each submap measurement. Building the DR factors requires encoding the state transition model between SLAM states $\mathbf{x}_k$ and $\mathbf{x}_{k+1}$ as a linearizable factor. The state transition model was built using measurement updates linearized about the EKF DR solution. Each vehicle state represented as a node in the factor graph $\mathbf{x}_k$ is linearized about the DR estimate of that state $\bar{\mathbf{x}}_k$ as

$$\mathbf{x}_k = \bar{\mathbf{x}}_k + \delta_k$$

Substituting the linearized factor graph states into the full state transition model gives

$$\begin{aligned}\mathbf{x}_{k+1} &= \bar{\Phi}_k \mathbf{x}_k + \bar{\mathbf{u}}_k + \bar{\mathbf{w}}_k \\ \bar{\mathbf{x}}_{k+1} + \delta_{k+1} &= \bar{\Phi}_k (\bar{\mathbf{x}}_k + \delta_k) + \bar{\mathbf{u}}_k + \bar{\mathbf{w}}_k\end{aligned}$$

$$\begin{aligned}\bar{\mathbf{x}}_{k+1} + \delta_{k+1} &= (\bar{\Phi}_k \bar{\mathbf{x}}_k + \bar{\mathbf{u}}_k) + \bar{\Phi}_k \delta_k + \bar{\mathbf{w}}_k \\ \delta_{k+1} - \bar{\Phi}_k \delta_k &= (\bar{\Phi}_k \bar{\mathbf{x}}_k + \bar{\mathbf{u}}_k) - \bar{\mathbf{x}}_{k+1} + \bar{\mathbf{w}}_k.\end{aligned}$$

Noting that the linearization points obey the noiseless state transition model

$$\bar{\mathbf{x}}_{k+1} = \bar{\Phi}_k \bar{\mathbf{x}}_k + \bar{\mathbf{u}}_k$$

gives the final result

$$\delta_{k+1} - \bar{\Phi}_k \delta_k = \bar{\mathbf{w}}_k. \quad (5)$$

Equation 5 gives several interesting results. First, only the linearization points $\bar{\mathbf{x}}_k$ and $\bar{\mathbf{x}}_{k+1}$, state transition matrix $\bar{\Phi}_k$ and covariance of $\bar{\mathbf{w}}_k$ need to be transferred from the DR EKF to the factor graph. Second, these values are created in the DR filter and cannot be updated by the factor graph. The factor graph DR links cannot be re-linearized without re-running the DR EKF. Permanently accepting the EKF’s linearization error is the key approximation that permits eliminating a large number of DR measurements from the factor graph. Non-linearity in DR arises from errors in angle. While a factor-graph representation with all the intermediate states allows a more thorough treatment of non-linearity, angular errors are typically small when high-quality attitude sensors, like those found on the ROV *Hercules*, are used.

Once a mathematical formulation for the factor graph has been found, an existing factor graph solver may be used to optimize the resulting graph. GTSAM is a factor graph SLAM implementation developed by Georgia Tech that includes many useful SLAM-related nodes and factors and uses methods well-described in the literature [8, 21, 23]. An excellent tutorial is also available [24]. The DR factors described above require defining full state nodes with 12 variables per node. The combined DR factor described in Equation 5 must also be implemented. The resulting graph may be solved using one of GTSAM’s standard optimizers. The

Figure 10: The final DR-only factor graph after intermediate states and factors have been combined into a single factor (red) between each pair of consecutive SLAM states. Compare the simplicity of this graph to the large number of nodes and factors in Figure 9.

results in this study used the Levenberg-Marquardt optimizer with the default settings. An example final factor graph with combined DR factors is shown in Figure 10.

2.4 Properties of Conditional Factor Graph Dead Reckoning

2.4.1 Comparison to Augmented-State EKF

In addition to reducing the number of states that must be stored in the factor graph, it would also be useful to reduce the size of each state to just the position and attitude variables. The assumptions required by factor graphs to preserve a sparse structure makes this non-trivial. The augmented-state EKF representation of [6] stores only vehicle poses in its state history. The mathematical basis for this provides some insights into the difference between factor-graph and augmented-state representations. Augmented-state EKFs explicitly represent the joint PDF over all poses in the navigation history. While storing the full PDF relaxes some of the assumptions that must be made about relationships between poses, a fully-dense representation imposes significant performance issues. Maintaining the correlation between the current state and all previous states is worse in practice than the

often-cited $O(N^3)$ cost of EKF updates in augmented-state EKF methods.

A factor graph representation of odometry requires a number of strong assumptions. The most important is the Markov property, which requires that the probability of each state given the previous state is independent of all other previous states. Intuitively, this requires that each state contains all the information required to apply the process model and advance the state. The Markov property guarantees that each node in the factor graph $\mathbf{x}_k$ is connected only to the previous state $\mathbf{x}_{k-1}$ and the next state $\mathbf{x}_{k+1}$. The structure of the resulting graph results in a sparse factor graph adjacency matrix and keeps the Jacobian used to optimize the factor graph sparse [8]. This sparsity is the key property required to efficiently optimize a factor graph [22, 23, 19, 25].

Without the Markov property, the relationship between states cannot be expressed concisely. This can be demonstrated on a factor graph with only odometry nodes. The grouping of state variables into multivariate graph nodes is largely arbitrary. A full vehicle state is usually drawn as one node, but in Figure 11a each state is broken into a pose node and an another node for the additional state variables such as pose rates. The odometry factors, shown as black circles, are defined over both the pose and “additional” nodes but remain the same internally. Variable node elimination is a standard procedure in factor graphs [19]. Eliminating a node involves taking the product over all factors involving that node and then summing out the variable to be eliminated. This procedure produces a new factor adjacent to all the nodes connected by factors to the eliminated node. A single step of this procedure is shown in Figure 11b with the new factor shown in red. The procedure is continued in Figures 12a and 12b. Once the final rate node has been eliminated in Figure 12c all odometry information is included in a single factor that spans all pose nodes. This factor could instead be generated from the PDF that

(a) Initial factor graph with states broken up into poses (top) and rates (bottom). Nodes in the same column form the same state node in the usual factor graph. Standard odometry factors are shown in black.

(b) Factor graph after eliminating the first rate node. Variable elimination creates a new factor spanning all nodes with factors to the eliminated node, shown in red.

Figure 11: Rate variables can be eliminated using standard variable elimination techniques, but the resulting factor graph dense. Proof concluded in Figure 12.

would result from running the augmented-state EKF using only DR measurement updates. While it may be possible to approximate this factor, the general trend in the literature has been to find ways to preserve exact sparsity [3, 16, 26, 27, 25]. The same approach is used here.

The lack of sparsity in an augmented-state EKF causes other performance issues. An augmented-state EKF state vector for Submap-SLAM may be broken down into two components. The first is the current vehicle state estimate. This full state is used to generate the DR track and provide the initial estimate of each submap origin. The vehicle state estimate is independent of survey length and

(a) Modified factor graph after eliminating two rate nodes.

(b) Factor graph after eliminating three rate nodes. Note how the final pose node depends on all three previous pose nodes through the red factor.

(c) Final factor graph. All pose factors depend on each other through the red factor.

Figure 12: Conclusion of the graphical proof begun in Figure 11. Observe that all nodes are connected by a common factor in the final diagram.

commonly has 12 elements $[6, 10, 4]$. The second part of the state vector is the set of pose estimates for each previous submap origin. This portion of the state vector grows with each new submap. For N submaps, storing the submap origins requires $6N$ state vector elements.

The augmented-state representation of the SLAM problem estimates the full covariance matrix between all N submap origins as well as between each of the N submap origins and the current vehicle state element. These relationships are stored in the state covariance matrix. This representation requires constantly updating the covariance between the current state estimate and every previous pose. Every prediction and update step requires updating the $12 \times 6N$ part of

the covariance matrix between the current vehicle state and every pose in the history. Prediction steps reduce the correlation between the current state and every previous pose. Measurement updates have a variety of effects, depending on the measurement model. Maintaining this cross-covariance with every single prediction and measurement step is computationally expensive. For maps near the 100-submap or so limit suggested by Roman [6], maintaining this $12 \times 6N$ cross-covariance for each of thousands of depth, attitude, and DVL measurement updates is more costly to apply than the several hundred submap-to-submap measurements that require inversion of the full covariance matrix.

2.4.2 Correlation Time

Prior work often assumes that some state variables may be represented as purely relative constraints while others may be represented as purely absolute constraints [12]. It will be shown in this section that these two extremes are approximations of the state transition model. The state transition approach provides a method to evaluate the validity of those approximations. A demonstration of the conditions for which this approximation is valid for data collected with the ROV *Hercules* is also included.

Not all vehicle state parameters are equally observable with a standard suite of DR sensors. Attitude, depth and velocity are measured directly, while a process model uses velocity information to build relative constraints between poses. The process model will propagate the effect of a single measurement to multiple neighboring states. The final state estimate at a given time will include a great deal of information from the most recent measurements and some information from older measurements. The state estimate after a measurement correlates not only with that measurement, but also many previous measurements in proportion to the age of each sample.

As a more concrete example, consider the depth estimate. Depth is measured directly by a pressure sensor and depth rate is measured by the DVL. The DVL-estimated z -velocities combine with the process model to provide a estimate of depth based on previous depth sensor measurements. The final EKF-estimated depth x_i may be written as a weighted average of this estimate $\hat{x}_i$ and the current measurement z_i . Note that the measurement model is simply $H_i = 1$ as the state x_i is simply the depth. Assuming that the Kalman gain is k_i , the standard state update may be re-written as a weighted sum as

$$x_i = (1 - k_i)\hat{x}_i + k_iz_i.$$

Every depth estimate is correlated with all previous depth estimates and, by extension, all previous depth measurements. The magnitude of this correlation decreases over time in relation to the Kalman gain. This observation makes intuitive sense; the previous depth sample is more useful in filtering the current depth measurement than one from 10 minutes ago. State variables that are observed through their derivatives, such as x, y position, are highly correlated with their previous values. With only relative DR measurements, the current position estimate continues to depend heavily on the initial position from which DR is started.

The previous analysis may be extended to the full state vector. Recall from Equation 4 that the transition from one SLAM state to another is also a weighted sum as

$$\mathbf{x}_i = \bar{\Phi}_{i-1}\mathbf{x}_{i-1} + \bar{\mathbf{u}}_{i-1}.$$

The effect of previous state $\mathbf{x}_{i-1}$ is moderated by the state transition matrix $\bar{\Phi}_{i-1}$. New measurements are added as the composite pseudo-measurement $\bar{\mathbf{u}}_{i-1}$. State variables that depend heavily on values from the previous state, such as x, y position, should have values close to one in the relevant rows of $\bar{\Phi}_{i-1}$ and a corresponding entry of $\bar{\mathbf{u}}_{i-1}$ that is close to the difference between the previous state

Figure 13: Diagonal coefficients of the state transition matrix $\bar{\Phi}$ between two submaps of Monterrey A. x and y have the same state transition characteristics, as do roll and pitch.

and the next state. State variables that are well-measured should have very small values in $\bar{\Phi}_{i-1}$ and values in $\bar{\mathbf{u}}_{i-1}$ that are close to the estimated value.

A useful simplification is to look at the diagonal elements of $\bar{\Phi}_k$ as the time between $\mathbf{x}_k$ and $\mathbf{x}_{k+1}$ increases. These values for a subset of states are shown in Figure 13. These elements determine how heavily a state variable at time k depends on its previous value at time $k - 1$. The diagonal elements of $\bar{\Phi}_k$ that correspond to the x, y states are one, indicating that the next position estimate depends heavily on the current estimate and that the elements of $\bar{\mathbf{u}}_k$ are the increment between positions. Similarly, those states with absolute measurements available become de-correlated from states far enough in the past. The time it takes for the current state to become de-correlated with the previous SLAM state depends on the measurement covariance and update rate for the sensor in question. The diagonal entry of $\bar{\Phi}_k$ for the roll and pitch states, which are measured at 10Hz with an assumed standard deviation of 0.02° , dies to 0.1% after 3.3 seconds. The entry for heading, which is measured at the same 10Hz but has a standard

deviation of 0.30° , does not pass the same threshold until 9.4 seconds. The depth sensor samples at a much-slower 2Hz and is assumed to have a standard deviation of 5cm. In Figure 13, the $\bar{\Phi}_k$ entry for depth at 51 seconds remains above the 0.01% threshold.

The results in Figure 13 call into question the idea that depth and attitude may be approximated as 4-DOF absolute factors while position may be approximated as a 2-DOF relative factor. That approximation is only valid when the entry on the $\bar{\Phi}_k$ diagonal is close to zero for attitude and depth and close to one for position. As may be seen in Figure 13, that is likely not the case for SLAM states less than 30 seconds apart. The exact time required for the approximation to become valid depends on the specific sensor sampling rates and variances. Submap origins are often a minute or more apart, but other poses that may be useful in the factor graph framework are often much closer together in time. Stereo images are taken by *Hercules* every three seconds. The LBL cycles used in Section 2.5.2 are less than five seconds apart, with some LBL range measurements being reported only milliseconds apart. The more general solution presented here allows the DR between these closely-related states to be faithfully transferred to the factor graph.

2.5 Validation Under Realistic Conditions

It is not trivial to validate the representation of DR in the odometry. SLAM converges well enough that simply testing whether a set of submap matches improves overall error metrics may not demonstrate the validity of any underlying mathematical premise. A previous version of this work improved map quality with Submap-SLAM despite using purely relative DR constraints [10].

The EKF, in both augmented-state and conventional form, remains a standard filter. Although the EKF is a filter while the factor graph approach produces a smoothed result, the EKF and factor graph results should be very similar. Both

methods should produce not only similar trajectory estimates but also comparable uncertainty estimates. As a smoother is able to use measurements before and after the current time step, the factor graph outputs should be expected to often have lower uncertainties. EKF and factor graph solutions to a DR problem are compared in Section 2.5.1. A better test is how the methods compare when additional measurements are added. The two methods are again compared within the context of LBL navigation in Section 2.5.2. Special emphasis is placed on comparing the covariance of the resulting estimates.

2.5.1 Comparison with Dead Reckoning Data

Dead reckoning is the simplest of all possible test cases. If the factor graph cannot closely approximate the uncertainty structure of the EKF-derived dead-reckoning, more complex factor graphs are unlikely to function correctly. For this test, an EKF is run at the same time as the state transition estimator. The state transition estimates are placed into a GTSAM factor graph as described in Section 2.3.2. The resulting GTSAM factor graph is then solved with a Levenberg-Marquardt algorithm to find the most likely trajectory estimate. The resulting trajectory and covariance estimates are compared in Figure 14.

Note the excellent agreement between state estimates. In the absence of additional constraints, the factor graph should reproduce the dead reckoning results as seen here. The uncertainty graph is also very similar. As expected, the factor graph solution shows slightly less uncertainty in the depth and attitude dimensions. For these state variables, information from absolute measurements from the depth sensor and fiber-optic gyrocompass is propagated both forward and backward in time. This non-causal processing slightly increases the amount of information available at any given timestep and consequently reduces the uncertainty.

The uncertainty for the x and y dimensions is very slightly higher for the factor

(a) Trajectory estimate

(b) Trajectory standard deviation estimate

Figure 14: Dead reckoning trajectory estimate (top) and estimated uncertainty of that trajectory (bottom) over the Monterrey A post-disturbance survey for EKF (blue) and GTSAM (red). The uncertainty spike at 22 minutes results from a known data recording issue with DVLNAV at the start of every hour.

graph than the EKF solution. No satisfactory explanation for this phenomenon has yet been found. Over-estimating the uncertainty is preferable to under-estimating it. Overconfidence in a solution makes the factor graph inconsistent and may impede data association [27]. Overconfidence may also suggest statistical significance when analyzing the resulting surfaces even where none exists.

2.5.2 Adding LBL Data

Dead reckoning is a good initial test case, but the optimizer’s result when additional factors are added is far more important. LBL navigation data collected at Monterrey A provides a set of factors that are independent of the idiosyncrasies of submap matching. The goal of this test is to process the same set of measurements with a conventional EKF and the hybrid EKF / factor graph method presented in this chapter. A complete discussion of LBL processing is deferred until Chapter 3.

The EKF solution uses the usual DR EKF with additional measurement updates for each LBL range. Ranges to beacon j at position $\mathbf{b}_j$ were added using a non-linear measurement model of the form

$$h(\mathbf{x}, b_j) = \|(\mathbf{x} \oplus \mathbf{x}_{vs}) - \mathbf{b}_j\|$$

where $\mathbf{x}_{vs}$ is the position of the LBL transponder in the vehicle body-frame and $\|\cdot\|$ gives the magnitude of the vector contained within. The beacon positions recovered from the field calibration were used directly. $\oplus$ is the compounding operator of Smith, Self and Cheeseman [28, 29]. Given the position of the vehicle $\mathbf{x}$ in some local reference frame and the position of the sensor in the vehicle body coordinate frame, $\mathbf{x}_{vs}$, the $\oplus$ operator combines the two coordinate transforms to give the position of the sensor in the local reference frame.

To ease comparison, the factor graph was built to match the EKF as closely as possible. In Chapter 3, where the navigation solution itself is of more interest, a more sophisticated model will be used. Unlike the EKF, the factor graph can

further refine its estimate of the beacon positions. For this test, very strong priors were placed on all beacon positions. Both EKF and factor graph solvers used the same measurement model, measurement covariance, and samples.

The resulting positions and standard deviations, again from the Monterrey A post-disturbance survey, are shown in Figure 15. The factor graph solution (red) is smoother because GTSAM is a non-causal smoother rather than a filter. Information from LBL measurements propagates forward and backward in time rather than suddenly updating the position estimate when a measurement occurs.

The position uncertainties are also very similar. Although the relationship between positions is set by the DR, all absolute position measurements from LBL ranges are processed by the factor graph. The factor graph’s acausal processing should result in a slightly lower x, y position uncertainty than the EKF’s filtering approach. This may be seen in the x and y uncertainties of Figure 15. Unlike x, y position, depth is better constrained by the pressure sensor included in the DR sensors than it is by the new LBL measurements. The factor graph depth value is more strongly influenced by the DR links than it is by the LBL data. The depth uncertainty should show similar behavior to the DR only case of Section 2.5.1. The depth uncertainty in Figure 15 matches that in Figure 14.

2.6 Conclusion

This chapter introduced a new method of loosely coupling a DR EKF to a factor graph for SLAM using a state transition model derived from the linearized EKF. Although the focus has been on Submap-SLAM, this method is applicable to any problem that uses a factor graph to improve the navigation of a DR-navigated UUV. With additional measurements and factors, the same DR factors may be applied to visual SLAM or one-way travel time acoustic navigation in addition to Submap-SLAM.

(a) Trajectory estimate

(b) Trajectory standard deviation estimate

Figure 15: LBL-navigated trajectory estimate (top) and estimated uncertainty of that trajectory (bottom) over the Monterrey A post-disturbance survey for EKF (blue) and GTSAM (red). The uncertainty spike in depth at 22 minutes results from the same data logging glitch seen earlier.

Even though the LBL factor graphs included thousands of nodes closer together in time than the depth sensor correlation time, the final pose estimates

and uncertainties were similar to the EKF solution. Once a factor graph has been formed, many of the standard factor graph-based algorithms from the literature may be applied to Submap-SLAM. One practical problem in bathymetric surveying is merging or comparing multiple surveys. A standard factor graph algorithm to this problem, anchor nodes, is applied to Submap-SLAM to quantify changes made during a shipwreck excavation in Chapter 3.

List of References

- [1] N. Bowditch, *The American Practical Navigator*. National Imagery and Mapping Agency, 2002.
- [2] J. C. Kinsey and L. L. Whitcomb, “Preliminary field experience with the dvlnav integrated navigation system for oceanographic submersibles,” *Control Engineering Practice*, vol. 12, no. 12, pp. 1541 – 1549, 2004.
- [3] R. Eustice, O. Pizarro, H. Singh, and W. H. Ma, “Visually augmented navigation in an unstructured environment using a delayed state history,” in *Proc. IEEE Int. Conf. on Robotics and Automation*, 2004, pp. 25–32.
- [4] R. Eustice, “Large-area visually augmented navigation for autonomous underwater vehicles,” Ph.D. dissertation, MIT/WHOI Joint Program, 2005.
- [5] B. Jalving, K. Gade, K. Svartveit, A. Willumsen, and R. Sorhagen, “Dvl velocity aiding in the hugin 1000 integrated inertial navigation system,” *Modeling Identification and Control*, vol. 25, no. 4, pp. 223–236, 2004.
- [6] C. Roman and H. Singh, “A self-consistent bathymetric mapping algorithm,” *Journal of Field Robotics*, vol. 24, no. 1-2, pp. 23–50, 2007. [Online]. Available: <http://dx.doi.org/10.1002/rob.20164>
- [7] V. Indelman, S. Williams, M. Kaess, and F. Dellaert, “Factor graph based incremental smoothing in inertial navigation systems,” in *Information Fusion (FUSION), 2012 15th International Conference on*, July 2012, pp. 2154–2161.
- [8] F. Dellaert and M. Kaess, “Square Root SAM: Simultaneous localization and mapping via square root information smoothing,” *Intl. J. of Robotics Research (IJRR)*, vol. 25, no. 12, pp. 1181–1204, Dec 2006.
- [9] R. M. Eustice, H. Singh, and L. L. Whitcomb, “Synchronous-clock, one-way-travel-time acoustic navigation for underwater vehicles,” *Journal of Field Robotics*, vol. 28, no. 1, pp. 121–136, 2011. [Online]. Available: <http://dx.doi.org/10.1002/rob.20365>

- [10] I. Vaughn, “Microbathymetry using self-contained navigation and simultaneous localization and mapping,” Master’s thesis, University of Rhode Island, 2012.
- [11] L. Carlone, Z. Kira, C. Beall, V. Indelman, and F. Dellaert, “Eliminating conditionally independent sets in factor graphs: A unifying perspective based on smart factors,” in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, May 2014, pp. 4290–4297.
- [12] M. A. VanMiddlesworth, “Toward autonomous underwater mapping in partially structured 3d environments,” Master’s thesis, Massachusetts Institute of Technology, 2014. [Online]. Available: hdl.handle.net/1721.1/87791
- [13] C. G. Kunz, “Autonomous underwater vehicle navigation and mapping in dynamic, unstructured environments,” Ph.D. dissertation, Massachusetts Institute of Technology and Woods Hole Oceanographic Institution Joint Program, 2011.
- [14] C. Kunz and H. Singh, “Map building fusing acoustic and visual information using autonomous underwater vehicles,” *Journal of Field Robotics*, vol. 30, no. 5, pp. 763–783, 2013. [Online]. Available: <http://dx.doi.org/10.1002/rob.21473>
- [15] N. Brokloff, “Matrix algorithm for doppler sonar navigation,” in *OCEANS ’94. ’Oceans Engineering for Today’s Technology and Tomorrow’s Preservation.’ Proceedings*, vol. 3, Sep 1994, pp. III/378–III/383 vol.3.
- [16] R. Eustice, H. Singh, W. Hole, and W. Hole, “Visually navigating the rms titanic with slam information filters,” in *in Proceedings of Robotics: Science and Systems*, 2005, pp. 57–64.
- [17] Y. Geng, R. Martins, and J. Sousa, “Accuracy analysis of dvl / imu / magnetometer integrated navigation system using different imus in auv,” in *Control and Automation (ICCA), 2010 8th IEEE International Conference on*, June 2010, pp. 516–521.
- [18] M. S. Grewal and A. P. Andrews, *Kalman Filtering: Theory and Practice Using Matlab*, 2nd ed. Wiley & Sons, 2001.
- [19] D. Koller and N. Friedman, *Probabilistic Graphical Models*. MIT, 2009.
- [20] F. Dellaert, “Square Root SAM: Simultaneous location and mapping via square root information smoothing,” in *Robotics: Science and Systems (RSS)*, 2005.

- [21] M. Kaess, A. Ranganathan, and F. Dellaert, “iSAM: Incremental smoothing and mapping,” *IEEE Trans. on Robotics (TRO)*, vol. 24, no. 6, pp. 1365–1378, Dec 2008.
- [22] M. Kaess, V. Ila, R. Roberts, and F. Dellaert, “The Bayes tree: Enabling incremental reordering and fluid relinearization for online mapping,” Computer Science and Artificial Intelligence Laboratory, MIT, Tech. Rep. MIT-CSAIL-TR-2010-021, Jan 2010.
- [23] M. Kaess and F. Dellaert, “Covariance recovery from a square root information matrix for data association,” *Journal of Robotics and Autonomous Systems (RAS)*, vol. 57, pp. 1198–1210, Dec 2009.
- [24] F. Dellaert, “Factor graphs and gtsam: a hands-on introduction,” Georgia Tech, Tech. Rep. GT-RIM-CP&R-2012-002, 2012, retrieved 25 Nov 2015. [Online]. Available: <https://research.cc.gatech.edu/borg/sites/edu.borg/files/downloads/gtsam.pdf>
- [25] S. Thrun, Y. Liu, D. Koller, A. Y. Ng, Z. Ghahramani, and H. Durrant-Whyte, “Simultaneous localization and mapping with sparse extended information filters,” *The International Journal of Robotics Research*, vol. 23, no. 7-8, pp. 693–716, 2004. [Online]. Available: <http://ijr.sagepub.com/content/23/7-8/693.abstract>
- [26] R. M. Eustice, O. Pizarro, and H. Singh, “Visually augmented navigation for autonomous underwater vehicles,” *IEEE Journal of Oceanic Engineering*, vol. 33, no. 2, pp. 103–122, 2008.
- [27] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT Press, 2006.
- [28] R. C. Smith and P. Cheeseman, “On the representation and estimation of spatial uncertainty,” *The International Journal of Robotics Research*, vol. 5, no. 4, pp. 56–68, 1986. [Online]. Available: <http://ijr.sagepub.com/content/5/4/56.abstract>
- [29] R. Smith, M. Self, and P. Cheeseman, “Estimating uncertain spatial relationships in robotics,” in *Robotics and Automation. Proceedings. 1987 IEEE International Conference on*, vol. 4, Mar 1987, pp. 850–850.

CHAPTER 3

Factor Graphs for Change Detection

3.1 Introduction

Many scientific investigations are concerned with how the seafloor changes in response to various processes. Repeated surveys of the same site can show how the site changes over time. A wide range of processes have been characterized by repeatedly mapping the same site, including mud volcanoes [1], volcanic eruptions [2], changes in coral reefs [3], and the impact of seafloor trawling on archaeological sites [4]. This chapter examines changes at the Monterrey A shipwreck site during a 2013 excavation.

Temporal analysis of a study site consists of individual surveys measuring site bathymetry at separate times. These surveys form a time series once each survey has been processed using available navigation. A comparison of dead-reckoning (DR), long baseline acoustic navigation (LBL) and Submap-SLAM navigation for each of the individual surveys at Monterrey A is presented in Section 3.3.

Maps from these surveys may be presented together for qualitative analysis or differenced to begin a quantitative analysis. Measured bathymetric change may result from changes to the actual bathymetry or from measurement biases. In addition to the accumulated DR error, depth sensor measurements can be biased by various environmental effects, such as tides. Error budgets and corrections for some of the effects are discussed in Section 3.4. Changes made to the site during excavation are extensively documented through video footage, photomosaics, and numerous log entries. A selection of the these documented changes and their relevance to bathymetric change are discussed in Section 3.5.

Once known measurement errors have been corrected for, individual surveys may be compared to quantitatively look for changes. To do this, the surveys

of Section 3.3 must first be referenced to a common datum. Even once rigidly aligned, the final surface for each survey is still affected by navigation errors that may bias or warp the resulting seafloor estimate and produce numerous artifacts upon differencing. Submap-SLAM is extended to perform a non-rigid alignment of surveys in Section 3.6.

Aligning surveys into a common datum is also useful to merge multiple overlapping surveys of the same site. Such a situation could arise if multiple vehicles map the same area or if a single vehicle conducts multiple surveying sessions. AUV batteries run out, ROV support ships get dragged off station by heavy currents, and vehicles of all types experience unexpected mechanical failures. The roughly 36-hour survey of the inner crater of the Kick'em Jenny volcano by the ROV *Hercules* in 2014, for example, was split into eight parts. Using the overlap between surveys to build a single map of the crater also requires a non-rigid alignment like that presented in Section 3.6.

3.2 Monterrey A

3.2.1 Site Overview

The Monterrey A shipwreck lies in 1330m of water approximately 170 nautical miles (315km) southeast of Galveston, TX. The site was excavated over a five-day period by a group of collaborators from the National Oceanographic and Atmospheric Administration, the Bureau of Ocean Energy Management, the Bureau of Safety and Environmental Enforcement, Texas A&M University, the Ocean Exploration Trust, and the University of Rhode Island. The ROV *Hercules* completed extensive stereo camera, structured-light, and high-frequency multibeam surveys before, during, and after the excavation. The excavation itself was extensively documented as demanded by current best-practices for deep-water archaeological excavation [5]. A spatial log-histogram of ROV position provides insight into where

time was spent on the wreck during excavation. Documented changes provide a useful method for verifying where actual change occurred.

(a) Location of Monterrey A

(b) Central artifact pile, including stove

Figure 16: The approximate location of the Monterrey A shipwreck sites in the Gulf of Mexico is shown on the left. On the right is a video capture showing the large artifact pile near the center of the ship.

Preservation and analysis of the recovered artifacts continues two years after the original expedition. Early results suggest the site is an early 19th-century two-masted ship that was unusually large for the time and region. Dates on some recovered artifacts show the vessel was lost no earlier than 1815. The presence of a cannon and a large quantity of small arms, combined with the political turmoil context of the Western Gulf of Mexico in the early 19th century, makes it possible the vessel was engaged in privateering or arms smuggling. The nearby presence of two wrecks from a similar period, Monterrey B & C, suggests that Monterrey A was lost in a storm or some other cataclysmic event.

The vessel appears to have sunk mostly intact and settled to the seabed upright. Photomosaics from the pre-disturbance and post-disturbance ROV surveys are shown in Figure 17. Most of the vessel's wood has rotted away. The outline visible in the photomosaics is the copper hull sheathing used to prevent biofouling and infestation by *taredo* molluscs, along with portions of the remaining hull still preserved by sediment. A large pile of metal artifacts, including a cannon and

(a) Monterey A, Pre-Disturbance

(b) Monterey A, Post-Disturbance

Figure 17: Photomosaics of Monterey A before and after excavation. Most excavation work focused on the aft-starboard quarter of the vessel. Small changes in illumination and camera position produced changes in appearance that are particularly visible on the bottom edge of the hull. Image courtesy of Clara Smart.

stove, sits in the middle of the wreck. The large anchor visible on the photomosaics and bathymetry surveys is located near the wreck’s bow. Although items were taken from all over the wreck site, most excavation work focused on a debris pile in the aft-starboard quarter of the vessel. The interim archaeological results are described in greater detail in [6].

3.2.2 Surveys Conducted

After an initial visual reconnoiter around the Monterey A site, four LBL beacons were deployed. A patch test was conducted to survey in the lever-arm and angular offsets of the Blueview MB-1350-90 multibeam on ROV *Hercules* while the beacon-to-beacon calibration was conducted. Four cinder blocks painted with

distinctive patterns were also deployed on site to provide reference datums for later processing. Only two are visible in the surveys seen here.

SURVEY	START TIME	END TIME	PINGS
Pre-Disturb., Complete	2013-07-19 22:29:02	2013-07-20 00:58:05	171,819
Pre-Disturb., Large	2013-07-19 22:29:02	2013-07-20 00:09:00	116,617
Pre-Disturb., Small	2013-07-20 00:09:00	2013-07-20 00:58:05	55,201
Mid-Disturbance	2013-07-21 21:52:06	2013-07-21 22:31:21	34,411
Post-Disturbance	2013-07-23 09:38:08	2013-07-23 10:56:33	92,118

Table 2: Surveys conducted at Monterrey A. Note that the pre-disturbance surveys may be treated as one large survey or two smaller ones.

A thorough pre-disturbance survey of the site was conducted before any attempts at artifact recovery. At the time of the pre-disturbance survey, the addition of the reference blocks was the only change made to the site. This initial survey was conducted as two complete sets of survey lines over the wreck. The first set of northeast / southwest lines covered the wreck and a possible debris field to the northwest. A second set of survey lines was run perpendicular to the first. Both sets of lines may be processed together as a single survey, referred to as the “Complete Pre-Disturbance” survey. Alternatively, the two sets of lines may be processed independently to evaluate the change detection algorithm in the absence of any change to the site. The first set of northeast / southwest lines is referred to as the “Large Pre-Disturbance” survey and the second set of northwest / southeast lines are the “Small Pre-Disturbance” survey. These two pre-disturbance surveys are used here as a control dataset with no change.

Two additional surveys were conducted following artifact recovery. A quick survey was conducted “Mid-Disturbance” after approximately 26 hours of excavation. Following another 14 hours of *Hercules* bottom-time, artifact recovery was considered complete and a final “Post-Disturbance” survey was conducted. A summary of all the surveys is given in Table 2. The number of multibeam pings in

each survey is a useful proxy for survey size. Given the changes detected in the post-disturbance survey, analysis of the mid-disturbance survey was determined to be redundant.

3.3 Processing Single Surveys

3.3.1 Conventional Dead-Reckoning

Dead-reckoning (DR) by integrating DVL velocity measurements is the most commonly-available navigation source for such maps. LBL beacon deployment is too expensive and time-consuming for routine use. In the past six years and 449 *Hercules* dives, the four dives at Monterrey A are the only ones for which LBL navigation was available. DR thus remains a standard navigation method for UUV bathymetric surveying. The bathymetric results for Monterrey A using DR navigation are shown in Figures 19, 23, and 27 for the large pre-disturbance, small pre-disturbance, and post-disturbance surveys.

3.3.2 Long Baseline Positioning

A Sonardyne Ranger Long Base Line (LBL) acoustic positioning network of four seafloor-mounted beacons was deployed on July 19th 2013. LBL beacon networks are typically calibrated by driving the ship in a pattern around each beacon while ranging to it. To save time, the pressure depth sensor and two-way acoustic communications of Sonardyne Mk 5 Compatt beacon was used to perform a partial calibration. As in the usual calibration procedure, the beacons pinged between themselves to measure ranges between each pair of beacons. The distance between each pair was measured with 10 independent pings. The overall location of the network was fixed by measuring the location of one beacon with the LinkQuest TrackLink 5000-series ultra-short-baseline (USBL) navigation system usually used by *Hercules*. This beacon position was assumed to be fixed in all subsequent field processing. The combination of acoustic ranges, depth measurements, and

a single arbitrarily positioned beacon constrains the array shape but not the array orientation. This orientation is typically recovered by measuring the position of two beacons from the surface. As one beacon position was arbitrarily fixed, the second was measured by driving *Hercules* from the fixed beacon to a second beacon and noting the distance and bearing with *Hercule's* DVL-based real-time dead-reckoning package.

The entire beacon deployment and calibration procedure was completed in approximately nine hours. With only a single vehicle-mounted transponder available, measuring the beacon positions by directly ranging between E/V *Nautilus* and the beacon network would have required many additional hours as well as mounting a transponder on the ship and surveying its ship-relative position. The abbreviated calibration procedure thus saved a great deal of time at the cost of some uncertainty in the overall geodetic position of the site and the final orientation of the beacon array. The science objectives required only a positioning system with local precision, but not geodetic accuracy. Uncertainty in array orientation is far more important. Positions returned by the LBL system appeared to be rotated relative to the real-time dead-reckoning during the pre-disturbance survey. This was reduced for real-time positioning on subsequent dives by applying an ad-hoc 4.9° rotation about the origin beacon.

Another solution is to use SLAM techniques to solve for the beacon positions. This is a well-established technique known in the literature as synthetic baseline navigation and is available commercially [7, 8, 9]. The LBL beacons are points that may be solved for as classic SLAM features. LBL beacon responses may be uniquely identified by a combination of center frequency and encoding scheme, ensuring perfect data association.

With a factor graph, the location of each beacon is added as a node in the

Figure 18: Sound velocity in meters per second as measured by the CTD on *Hercules* while excavating Monterrey A. The stability of the benthic environment at 1330m reduced variation between surveys.

graph. The field calibration described above was used as a weak prior on the horizontal position of each beacon as well as an initial position for the graph optimizer. The tide-corrected depths observed by each beacon were used as a prior to constraint beacon depth in the final graph. Ten pings between each beacon pair were collected during calibration. These ranges were added as range factors.

Uncertainty in the range measurements depends on both the measurement error in the inter-beacon range measurements as well as changes in the sound velocity. The calibration pings were used to estimate a timing uncertainty of 0.003 ms. An analysis of CTD-based sound velocity measurements when *Hercules* was on the wreck site (Figure 18) shows that the sound velocity uncertainty is 0.056m/s. Assuming independence these give a per-measurement range uncertainty of approximately 0.95 cm. This result agrees with the measurement error of “better than 1.5cm” published by the manufacturer [10]. Calibration ranges were computed from the average of 10 independent range measurements. Calibration factors were

thus assumed to have a much smaller standard deviation of 0.025cm.

In order to keep the factor graph numerically well-conditioned, it was necessary to initialize the beacon positions with weak priors while placing a very strong prior on the initial position of *Hercules*. In effect, this pins the survey datum to the field-computed LBL position at the start of the survey and solves for the beacon position that minimizes the overall error from that starting position. While effective at reducing the overall navigation error, this approach does not place all three surveys into a common reference frame. Noise in each survey slightly perturbs the beacon locations and thus the overall navigations solution. This accumulated error produces position biases in the target survey area between surveys. Instead, using the LBL data to refine the navigation solution and then applying the same rigid alignment technique as the DR surveys was found to be more effective.

3.3.3 Submap-SLAM

Submap-SLAM offers an alternative method of improving the navigation solution without the external navigation beacons required by LBL. The SLAM solution presented here does not use the LBL and relies only on DR sensors and the multi-beam sonar. The uncertainty of the first pose in the survey must remain very small to keep the factor graph numerically well-conditioned. This conditioning was ensured by placing a strong prior on the first position node based on the field-calculated LBL position at the start of the survey. If no external navigation is available, then an arbitrary initial position may be used.

Submaps were broken every 750 pings. The large pre-disturbance survey had 156 submaps, the small pre-disturbance survey had 74 submaps, and the post-disturbance survey had 123 submaps. Each survey had 40, 37, and 77 valid submap to submap links, respectively.

3.3.4 Results

The results of applying each of the three navigation methods to each of the three surveys are shown on the following pages. Bathymetry maps are shown on a purple-to-red colormap shifted in depth to match each survey’s vertical datum. A Hasdorff-distance self-consistency metric [11] is used to show the error of each submap on a white-to-yellow-to-red colormap draped over the underlying bathymetry for each survey in Figures 22, 26, and 30. Lower errors are shown as lighter colors and higher errors are shown as red colors. Portions of the survey with large self-consistency errors also have a high depth uncertainty. A histogram of self-consistency error for each navigation method is provided for each survey in Figure 31. Submap-SLAM modestly reduced the error of the two pre-disturbance surveys and significantly improved the self-consistency of the post-disturbance survey.

Figure 19: Large pre-disturbance survey, DR-only

Figure 20: LBL-navigated large pre-disturbance survey

Figure 21: SLAM renavigated large pre-disturbance survey

(a) Large pre-disturbance error map, DR navigation

(b) Large pre-disturbance error map, LBL navigation

(c) Large pre-disturbance error map, Submap-SLAM navigation

Figure 22: Roman & Singh's error metric for the large pre-disturbance survey. All three figures have matched colorbars.

Figure 23: Small pre-disturbance survey, DR-only

Figure 24: LBL-navigated small pre-disturbance survey

Figure 25: SLAM renavigated small pre-disturbance survey

(a) Small pre-disturbance error map, DR navigation

(b) Small pre-disturbance error map, LBL navigation

(c) Small pre-disturbance error map, Submap-SLAM navigation

Figure 26: Error in the small pre-disturbance survey for each navigation method as measured with Roman & Singh's error metric. Color bars are matched.

Figure 27: Post-disturbance survey, DR-only

Figure 28: LBL-navigated post-disturbance survey

Figure 29: SLAM renavigated post-disturbance survey

(a) Post-disturbance error map, DR navigation

(b) Post-disturbance error map, LBL navigation

(c) Post-disturbance error map, Submap-SLAM navigation

Figure 30: Self-consistency error in the post-disturbance survey. The colors are matched between all three subfigures. The SLAM results show significantly less error than DR and LBL navigation.

Figure 31: Self-consistency error histograms for each survey. Better navigation should reduce the self-consistency error and push the histogram further to the left. Submap-SLAM moderately improves the self-consistency error for the two pre-disturbance survey and significantly reduces error in the post-disturbance survey.

3.4 Environmentally-Induced Non-Bathymetric Changes

Although already corrected in the maps presented previously in Section 3.3.4, field processing of the Monterrey A pre- and post-disturbance surveys placed the post-disturbance survey an average of 0.58m deeper than the pre-disturbance survey. The depth of a survey is measured with the pressure-based depth sensor on the survey UUV. Pressure-based depth sensors are simple, inexpensive, reliable, and generally quite accurate compared to other methods.

A number of physical processes determine the pressure at depth. Slow variation in these physical processes can produce system biases both during and between surveys. The hydrostatic pressure at depth z is given in terms of gravity $g(\cdot)$ and water density $\rho(\cdot)$ as in [12, 13]

$$p(z) = p(z_0) + \int_{z_0}^z g(h) \cdot \rho(h) dh. \quad (6)$$

This function must be inverted to convert the pressure measurement to a depth value. Variation of each term adds uncertainty in the depth measurement. Not all sources of variation are practically relevant.

The most significant source of depth bias at Monterrey A was found to be tidal variation in the height of the sea surface z_0 . The tidal model examined in Section 3.4.1 accounted for 56 of the measured 58 centimeters of depth bias between the pre- and post-disturbance surveys. Atmospheric pressure at the sea surface $p(z_0)$ has also been proposed as a source of error. It is shown to produce up to 3cm of variation in Section 3.4.2. Changes in water density as a function of depth $\rho(h)$ can also change the measured pressure at a given depth. The standard Fofonoff equation that ignores seawater properties introduced an error of approximately 1.47m to all surveys. As with x/y position, however, change detection is more concerned with the difference between survey vertical datums. Section 3.4.3 will show that error between surveys introduced by changes in the watercolumn is on

the order of 3.5cm although the analysis is limited by the small quantity of available data.

The only remaining term in Equation 6 is the gravity term. Although gravity can vary with location, the gravitational anomaly is dominated by a dependence on latitude [12]. Changes in gravity over time are small compared to other sources of error [14].

3.4.1 Tides

Figure 32: Tides during Monterrey A excavation. Survey times are indicated with red boxes. The average value over the entire time on-site was removed from each dataset to produce a tide anomaly value within the excavation time window. This avoids the need to correct for changes in datum with timescales significantly longer than the expedition. The average tide anomaly for each survey is marked with a blue “x.”

The largest source of depth bias between surveys was found to be tides. The conventional solution is to use a tide gauge. NOAA tide gauges in the Gulf of Mexico are concentrated on the coastline. No data was found within 100km of the survey site. Instead, a tidal prediction at the Monterrey A site for the duration of the excavation is shown as the blue line in Figure 32. The prediction is interpolated from an hourly output of the TPXO 7.2 tidal inversion computed with the `Matlab` Tidal Model Driver [15, 16]. The computed tide model accounts for 96.6% of depth

bias between the pre- and post-disturbance surveys.

One hypothesized method to correct for tidal changes is to use the survey ship's GPS to measure sea surface height directly. A 6-hour moving average filter was used to eliminate noise from waves and other short-term phenomena and the filtered altitude measurement is shown in green in Figure 32. Although the GPS height shows changes at similar time scales, it does not match the observed difference in survey depth biases as closely as the predicted tide. This could be due to a number of factors, including earth tides. Changes in *Nautilus*'s draft as ballasting is adjusted and other unrecorded phenomena add further errors. Given these limitations and the close agreement between the well-established TPXO 7.2 model and the observed depth biases, a TPXO-derived tidal correction was applied to all depth measurements.

3.4.2 Atmospheric Pressure

Figure 33: Variation in depth due to pressure near Monterrey A. Pressure change is presented as variation about the mean to highlight differences. Neglecting differences in density, one millibar of pressure is approximately one centimeter of depth. These data are from a fixed oil platform about 100km from Monterrey A.

Atmospheric pressure sets the initial force on the column of water above the depth sensor. Corrections for atmospheric pressure may be required to meet IHO requirements for shallow-water surveying [17]. No barometric pressure was

recorded aboard *Nautilus* during the Monterrey A excavation. Data from the closest Nation Data Buoy Center (NDBC) station was used for this analysis instead. Station 42361 is a fixed drilling platform operated by Shell International. Pressure for the duration of the Monterrey A excavation is shown in Figure 33a. Although variance within a survey is small, pressure effects could result in up to five centimeters of depth bias between surveys.

For future surveys, it is helpful to understand the magnitude of these effects more generally. Although a complete analysis using a large number of stations is beyond the scope of this work, a rudimentary analysis of station 42361 over all of 2013 should provide crude insight into the observed variation. A histogram of all pressures recorded by station 42361 during most of 2013 is shown in Figure 33b. The 7,387 hourly samples have an average pressure of 1015.8 millibar and a standard deviation of 4.4 millibar. High rates of change in barometric pressure are associated with weather phenomena that may prevent survey activities [18].

3.4.3 Water Column Properties

The final potentially-significant source of error in depth sensor measurement comes from changes in the density of the water column above the survey site. Sea-water density varies with temperature, salinity, and pressure. When integrated over a 1,330m water column, these errors could be significant. The SBE49 conductivity, temperature, depth (CTD) sensor on *Hercules* routinely logs whenever Herc is in the water. Each descent and ascent thus provides a single crude CTD cast. Unlike a conventional CTD/rosette, little effort is made to run the cast at consistent speed. Although the SBE49 is a pumped CTD with internal automated algorithms to reduce salinity spiking and other common CTD artifacts, the data provided is often of modestly lower quality than a dedicated CTD/rosette [19, 20]. Of the five dives performed during the cruise (H1273-H1277), the ascent of H1273

and descent of H1274 had data gaps that prevented their use as CTD casts.

(a) CTD Casts over Monterrey A

(b) Depth error from ignoring seawater density

Figure 34: Seawater properties and their effect on depth sensor measurements above Monterrey A based on eight casts taken during ROV descents and ascents.

The eight best casts are shown overlaid in Figure 34a. Temperature and salinity are as reported by the instrument. Seawater density was computed using the TEOS-10 equations [21] as implemented by the `Matlab` version of the Gibbs Seawater Toolbox [22]. Water density is quite stable, particularly below 200m

depth. Depths used by the real-time navigation software on *Hercules* are not corrected for any water column properties. Figure 34b shows the difference between a pressure to depth conversion that integrates density changes and a latitude-only correction similar to [12]. Mean error values were computed for one meter bins for each cast. The per-bin mean over all eight casts is shown in the middle figure. The standard deviation within a bin is shown in the right graph of Figure 34b.

As usual, the absolute error of 147cm at 1330m is less important than how that error changes between surveys. The standard deviation between all eight casts is approximately 3.6cm. This suggests the effect of changes in water column properties between surveys is on the order of ± 5 cm. The quality of this analysis is limited by the small number of casts used. Below 200m, the error is extremely stable. The bin-to-bin change in error has a standard deviation under 1mm from 200m down to the site depth of 1330m. This also accounts for the leveling-off of the standard deviation graph seen in Figure 34b. Given that most of the changes in seawater properties occur in the top few hundred meters, diurnal heating and cooling effects are likely a major source of variation. The eight casts available do not provide enough data to fully characterize this daily behavior.

For these reasons, 3.6cm is only an approximation of the error introduced by density changes. No correction was applied to compensate for this phenomena. In the absence of casts close to survey start and end times, an attempt to compensate for such modest errors may add more bias than it removes.

3.4.4 Error Budget for Environmental Changes

A summary of the various depth sensor biases is given in Table 3. Errors within the time span of the excavation are described as relative errors and may affect the depth bias between change detection surveys. Absolute errors are those errors that are stable over the time of the excavation. While not relevant for

ERROR SOURCE	RELATIVE ERROR	ABSOLUTE ERROR
TIDES	58 cm	NOT INVESTIGATED
BAROMETRIC PRESSURE	3cm	5cm
WATER PROPERTIES	5cm	147cm

Table 3: A summary of the magnitude and effect of various depth sensor bias sources. Absolute error refers to an approximate bias over the time of the entire excavation and affects the global position estimate of the survey. Relative error gives the magnitude of effects between the surveys at Monterrey A.

the surveys presented here, these errors may be significant if Monterrey A is re-surveyed in future years. A model-based tidal was applied to the depth values used to process the surveys shown here. No other corrections were applied. While the tidal correction removed the majority of depth bias, any change analysis must be robust to a small shift in depth datum.

3.5 Excavation-Induced Bathymetric Changes

(a) Wood experiments deployed at Monterrey A (b) Stern of Monterrey A, showing demijohns

Figure 35: Items on the Seafloor that changed at Monterrey A during the survey. The wood experiments (left) were deployed forward of the bow datum brick, while one of the demijohns at the stern (right) was recovered.

Some 40 hours of excavation work was conducted over four dives between the pre- and post-disturbance surveys. During this time, artifacts were removed, experiments deployed, and a suction excavator displaced sediment. Some of the larger ceramics recovered, such as one of the demijohns in Figure 35b, were large

enough to be visible in the bathymetry surveys. Similarly, two wood degradation experiments (Figure 35a) were deployed off the site's stern. These should also be visible as a bathymetric change. Suction excavator impacts are more difficult to quantify. In general, disturbance to the site may be expected to correlate with where *Hercules* spent time during the excavation. Even lifting a jar is a time-consuming process with a robotic manipulator. A 2D histogram of *Hercules* positions during the excavation is shown in Figure 36.

Figure 36: A 2D histogram of *Hercules* positions at Monterrey A as measured by the LBL shown with a logarithmic scale. Changes made during the excavation should be strongly correlated with where *Hercules* spent time.

Sediment properties play a key role in how any excavation will appear in a bathymetric map. The surveys were conducted within tens of hours of the excavation holes and no significant bottom currents were observed on-site. It is unlikely that any excavated holes were filled in by current-driven activity. Although sediment push-cores were taken, they were analyzed for bulk chemistry, meiofauna, and hydrocabons but not grain size. The mechanical properties of the soil cannot

(a) Imprint of *Hercules* on the seafloor

(b) Datum brick 6.5m forward of the bow

(c) Datum brick 2.8m aft of the heavily-excavated stern

Figure 37: Images of seafloor disturbance after excavation. (a) A portion of the post-disturbance survey showing the effect of resting ROV *Hercules* on the seafloor. Note the shadow-creating hard ridge, which suggests the sediment has at least some adhesion. (b,c) Datum bricks from the bow and stern of the site. These were placed immediately prior to the start of the excavation, and may indicate how much sediment stirred up by the excavation landed on structures. Scales are approximate.

be verified with any great accuracy. Impressions of the bottom of the *Hercules* ROV on the seafloor may provide some qualitative sense of sediment properties. One such artifact is shown in Figure 37a. The crisp lines and well-defined ridge, visible from its shadow, suggest the soil may be of moderate grain size or smaller and has at least some adhesion. Much of the on-site digging used *Hercules*' suction excavation pump. This tool exhausts removed sediment out the back of *Hercules* onto the site. Exhausted sediment formed a thin film over the site near excavated areas, as may be seen in the photomosaics of Figure 17. A layer of sediment may be seen in the post-disturbance photomosaic on the datum brick 2.8m from the heavily-excavated stern section (Figure 37c). Compare this to the pristine datum brick 6.5m from the barely-disturbed bow shown in Figure 37b. While it is difficult to say anything conclusive about the sediment properties, the visible adhesion and lack of current suggests that any holes dug during the excavation likely did not experience sudden collapse. While the suction excavation tool did leave a layer of sediment on portions of the site, as seen in the datum brick, the thickness of that layer is likely small. These properties combine to suggest that *Hercules* was capable of excavating well-defined, detectable depressions.

3.6 Processing Multiple Surveys for Change Detection

Each of the surveys processed in Section 3.3 is referenced to its own datum. The last field-computed LBL fix before the start of the survey is used to initialize the DR filter. That fix also provides an origin latitude and longitude which is used to convert the survey's local x/y coordinates into a geodetic reference frame. The origin of each survey is thus determined by a single LBL navigation fix. This navigation is insufficiently accurate to reference two surveys to a common datum for change detection.

The goal of the alignment process is to estimate the transform between each

survey’s reference frame and a common reference frame. The reference frame of the first survey is used as the common datum. Rigid alignment estimates a simple offset between additional surveys and the common datum using each survey’s bathymetry surface. Non-rigid alignment represents the offset between each survey’s reference frame and the common reference frame as an anchor node [23]. The anchor nodes are used to merge the factor graphs for all relevant surveys into a single graph that can be minimized.

3.6.1 Rigid Alignment

Rigid alignment is the extension of submap matching techniques to entire surveys. As with submap alignment, a sum-of-squared differences (SSD) error metric was minimized to find the offset between maps. Any other pointcloud alignment technique may also be used. Each map remains rigid throughout the registration and differencing procedure. The large pre-disturbance survey was used as the basemap for both differencing operations. The difference between the small and large pre-disturbance surveys is presented as a baseline case in Figures 40, 41, and 42 for the DR, LBL, and single-survey SLAM navigation solutions. The difference between the post-disturbance and large pre-disturbance surveys for the same selection of navigation solutions is shown in Figures 44, 45, and 46. Navigation methods that produced more self-consistent maps in Section 3.3.4 produce difference surfaces with more clearly-identifiable bathymetric changes.

The rigid alignment can only solve for the offset between surveys. It does not exploit the significant overlap between surveys to refine the navigation solution for either survey.

3.6.2 Non-Rigid Alignment With Anchor Nodes

The Submap-SLAM framework already provides a mechanism to represent uncertainty within a survey and refine a survey’s positioning based on links within the survey. Single survey SLAM estimates a map surface using constraints from dead-reckoning and how well the submaps of a survey align with each other. Multi-session SLAM estimates a map surface using the same dead-reckoning and intra-survey submap constraints while also using constraints based on how well submaps align between the two surveys. These additional constraints between surveys provide further information for the SLAM navigation solution and produce a different map surface than single-survey SLAM would. It is for this reason the alignment process is described as non-rigid. Adding submap matches between surveys introduces two practical problems: a significant increase in the size of the SLAM estimation problem and the uncertain relationship between the datum of each survey.

Submap links between surveys influence the navigation solution of both surveys. The non-rigid alignment procedure thus requires solving the combined SLAM problem for both surveys simultaneously. The surveys at Monterrey A are already slightly larger than the 100-submap performance limit of EKF-based Submap-SLAM [24]. The combined large pre-disturbance and post-disturbance factor graph has 279 vehicle states to estimate. Although beyond the recommended limit for EKF SLAM [24], factor graphs have been demonstrated with thousands of states both in Section 3.3.2’s LBL solution and the literature [25, 26, 27].

Uncertainty in the transform between survey datums is the second major problem in survey merging. The transform between each survey and the common reference frame are included in the factor graph as anchor nodes and estimated during the SLAM process. This method has been demonstrated for the full six

degree of freedom visual SLAM problem [23]. The registration method used does not provide additional attitude information, so each reference frame is assumed to be well-constrained in attitude and only the three-dimensional position offset (x, y, z) is estimated.

Figure 38: Anchor nodes represent the transform between each reference frame and a common reference frame. They are used to both transform the results to the common datum and to for submap-to-submap factors between surveys.

A relative pose measurement between state i in survey j and state k in survey ℓ gives a relative pose constraint $\mathbf{z}_{ik}$. Adding this constraint to the factor graph requires computing the expected relative pose $\mathbf{p}_{ik}$ based on each of these poses in their respective survey reference frames $\mathbf{p}_{ji}$ and $\mathbf{p}_{\ell k}$. The pose $\mathbf{p}_{ji}$ may be transformed to the common reference frame g using the transform $\mathbf{T}_{gj}$ as

$$\mathbf{p}_{gi} = \mathbf{T}_{gj} \oplus \mathbf{p}_{ji}.$$

The compounding operator of Smith, Self and Cheeseman [28, 29], denoted $\oplus$, is used to combine coordinate transforms to find state i in the common reference frame via the transform $\mathbf{T}_{gj}$. The relative $\mathbf{p}_{ik}$ between this pose $\mathbf{p}_{gi}$ and pose k

from survey ℓ is thus

$$\mathbf{p}_{ik} = \mathbf{p}_{gj} \ominus \mathbf{p}_{gk} = (\mathbf{T}_{gj} \oplus \mathbf{p}_{ji}) \ominus (\mathbf{T}_{g\ell} \oplus \mathbf{p}_{\ell k}).$$

Relative pose constraints between poses in different surveys require four nodes: the vehicle states $\mathbf{x}_i$ and $\mathbf{x}_k$, which include poses, and the transforms $\mathbf{T}_{gj}$ and $\mathbf{T}_{g\ell}$ between each survey and the common reference frame. These transforms are the “anchors” for each survey and are estimated using anchor nodes. Relative pose measurements within a survey do not require anchor nodes and may be used exactly as in the single survey case.

In practice, it is convenient to use the reference frame of one of the surveys at the common reference frame. The anchor node for this survey is still included in the factor graph with a very strong prior to keep it close to the identity transform. The anchor nodes for additional surveys are added with a prior constraint based on the rigid alignment. Existing submap-to-submap measurements within a survey are unaffected by the presence of other surveys referenced to their own datums. Relative pose measurements between surveys further refine these transforms and reduce their uncertainty. The anchor transform for each survey is applied to the estimated navigation solution after the SLAM process has finished. Uncertainty in this transform introduces uncertainty in the navigation solution that would not be present in the single-survey case.

The position uncertainty estimate for the merged large and small pre-disturbance surveys is shown in Figure 39. The DR uncertainty with a rigid alignment is shown in green, a rigid alignment of the single-survey SLAM solutions in red, and the multi-session alignment with anchor nodes is shown in blue. The additional uncertainty introduced by the anchor transforms is clearly visible in the depth uncertainty. The second DR survey uncertainty (the dashed green line) does not include any additional uncertainty from datum merging and is likely

overconfident. Non-rigid alignment adds some information to the individual surveys and reduces their x/y uncertainty slightly. The primary benefit of merging the factor graphs is an improved estimate of the offset between surveys. Overall, navigation uncertainty in the common reference frame is less than 15cm in both x and y .

3.6.3 Results

Change detection results are only meaningful if the change is large relative to the depth uncertainty. Total propagated uncertainty (TPU) is not available for these surveys, so all changes less than $\pm 2.5\text{cm}$ were assumed to be statistically insignificant based on a combination of depth and multibeam uncertainty. Areas that did not change by at least this amount are shown in white. The colormap is centered about the average change value for each survey to allow for the depth sensor errors of Section 3.4. The large pre-disturbance survey was used as the base map in all cases. Areas of the map that are shallower in the survey being compared are shown in green and areas that are deeper are shown in red. Green corresponds to material that was added between surveys, and red for material that was removed. Both colors saturate for changes above $\pm 25\text{cm}$ in order to show small changes.

(a) Navigation uncertainty for the pre-disturbance surveys

(b) Navigation uncertainty for the large pre-disturbance and post-disturbance survey

Figure 39: Uncertainty in the merged multi-session surveys. The large pre-disturbance survey took place during the first 100 minutes. On each graph, the second survey begins at the 100 minute mark. The strictly DR uncertainty (green) resets to the initial value of 10cm in the horizontal. The plotted value does not include uncertainty in the transform to the common datum. The rigid SLAM result (red) does include that additional uncertainty but horizontal uncertainty within each survey is bounded by SLAM. Submap measurements between surveys significantly reduces the anchor node uncertainty for the multi-session results (blue) and achieves horizontal navigation uncertainty similar to SLAM results for the large pre-disturbance survey. Depth is well-constrained for both surveys by the depth sensor, and differences in depth uncertainty result from the anchor node z uncertainty. Time between the first and second surveys has been removed from the graph to ease comparison.

Figure 40: Measured change between pre-disturbance surveys, DR navigation and rigid alignment.

Figure 41: Measured change between pre-disturbance surveys, LBL navigation and rigid alignment.

Figure 42: Measured change between pre-disturbance surveys, Single-Survey Bathymetry SLAM navigation and rigid alignment.

Figure 43: Measured change between pre-disturbance surveys, multi-session Bathymetric SLAM.

Figure 44: Change between large pre-disturbance and post-disturbance survey, as measured with DR navigation and rigid alignment.

Figure 45: Change between large pre-disturbance and post-disturbance survey, as measured with LBL navigation and rigid alignment.

Figure 46: Change between large pre-disturbance and post-disturbance survey, as measured with single-survey Submap-SLAM and rigid alignment.

Figure 47: Change between large pre-disturbance and post-disturbance survey, as measured with multi-session Submap-SLAM.

3.7 Conclusions

Figure 48: Histograms of the change detection surfaces in Figures 40-47. The mean of each histogram has been set to zero to aide comparison. Changes less than 2.5cm are considered statistically insignificant and are shown in grey.

Navigation error introduces noise into the difference between surveys that can easily obscure the small changes made to the site between surveys. The histogram of values in the difference surfaces provides a method to evaluate the quality of the change detection results. These histograms for both the difference between the two pre-disturbance surveys, between which no changes were made, and between the pre-disturbance and post-disturbance surveys are shown in Figure 48. The mean survey difference was subtracted from each change surface before computing the histogram to account for differences in how the various survey alignment methods treat changes in the vertical datum. In the absence of change, shown in Figure 48a, the ideal histogram would appear as an extremely sharp peak at 0cm. The histogram for changes made during the survey, in Figure 48b, should also be clustered close to 0cm as the area affected by changes is much smaller than the unchanged area (see Figure 47). The mean μ and standard deviation σ in centimeters are given in Table 4 for the “no-change” surfaces between the two pre-disturbance surveys and the “change” surfaces between the large pre- and post-disturbance surveys.

METHOD	NO-CHANGE μ	CHANGE μ	NO-CHANGE σ	CHANGE σ
DR	-1.0	-1.5	3.8	3.3
LBL	-1.0	-0.7	4.2	3.1
RIGID-SLAM	-0.5	-1.1	3.4	2.7
MULTI-SESSION	-0.0	-0.3	3.1	2.5

Table 4: Statistics for the changes surfaces. The “No-Change” columns show the statistics for the histograms in Figure 48a from the difference surfaces between the two pre-disturbance surveys. The “change” columns give statistics for the difference surfaces between the large pre- and post-disturbance surveys. All values in cm, with μ and σ being the sample mean and standard deviation. Measured “change” due to navigation error overwhelms the actual change due to excavation, resulting in the higher average change for the no-change case.

All changes less than ± 2.5 cm are considered statistically insignificant and are shown in white on the change maps. The value of the histogram outside this threshold is particularly important. In the region beyond ± 2.5 cm, the LBL and DR solutions both have similar histogram values for both sets of change maps. The multi-session SLAM algorithm resulted in modest improvement over the rigidly-aligned SLAM surfaces.

Multi-session SLAM proved modestly effective at improving the quality of the change detection surfaces. SLAM algorithms typically assume a fixed environment. The changes at Monterrey A were small enough that this assumption remained valid. Most changes were small relative to the size of a submap, and a manual review of the proposed submap links was sufficient to reject invalid links. If enough of the site has changed that very few valid submap links between surveys are found, then multi-session SLAM will give the same solution as rigid-alignment SLAM algorithm. Even if the site changes significantly between surveys, single-session SLAM can still provide significant performance improvements over costly and time-consuming LBL beacon network.

The excellent quality of the SLAM solution at Monterrey A was in part a result of the large number of high-quality submap-to-submap matches available.

Submap-SLAM will revert to the DR navigation solution if insufficient submap-to-submap measurements are available. The quality of submap-to-submap matches is dependent on the properties of the submaps being matched and thus on the properties of the seafloor being imaged. Predicting the value of a particular piece of seafloor for registration allows a surveyor to not only select a navigation approach but also to adjust the survey plan to improve navigation performance. The relationship between seafloor properties and Submap-SLAM performance and a method to predict submap-to-submap matching performance in real-time during a survey are explored in Chapter 4.

List of References

- [1] T. Feseker, A. Boetius, F. Wenzhofer, J. Blandin, K. Olu, D. R. Yoerger, R. Camilli, C. R. German, and D. de Beer, “Eruption of a deep-sea mud volcano triggers rapid sediment movement,” *Nat Commun*, vol. 5, pp. –, Nov. 2014. [Online]. Available: <http://dx.doi.org/10.1038/ncomms6385>
- [2] P. Nomikou, M. Parks, D. Papanikolaou, D. Pyle, T. Mather, S. Carey, A. Watts, M. Paulatto, M. Kalnins, I. Livanos, K. Bejelou, E. Simou, and I. Perros, “The emergence and growth of a submarine volcano: The kameni islands, santorini (greece),” *GeoResJ*, vol. 12, no. 0, pp. 8 – 18, 2014. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S2214242814000047>
- [3] S. Williams, O. Pizarro, M. Jakuba, I. Mahon, S. Ling, and C. Johnson, “Repeated auv surveying of urchin barrens in north eastern tasmania,” in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, May 2010, pp. 293 –299. [Online]. Available: <http://www.cas.edu.au/content.php/237.html?publicationid=659&displaypage=3>
- [4] M. L. Brennan, D. Davis, R. D. Ballard, A. C. Trembanis, J. I. Vaughn, J. S. Krumholz, J. P. Delgado, C. N. Roman, C. Smart, K. L. C. Bell, M. Duman, and C. DuVal, “Quantification of bottom trawl damage to ancient shipwreck sites,” *Marine Geology*, 2015, in Revision.
- [5] D. F. Coleman, R. D. Ballard, J. B. Newman, T. S. Gregory, J. Howland, H. Singh, C. Roman, O. Pizarro, B. Foley, R. Eustice, A. Can, S. Webster, D. Piechota, C. Giangrande, C. Ward, R. Horlings, K. McBride, and M. J. Durbin, *Archaeological Oceanography*, R. D. Ballard, Ed. Princeton University Press, 2008.

- [6] C. E. Horrell and A. A. Borgens, “Archaeological investigations of an early nineteenth-century shipwreck in the gulf of mexico,” in *Offshore Technology Conference*, 2014.
- [7] M. Larsen, “Synthetic long baseline navigation of underwater vehicles,” in *OCEANS 2000 MTS/IEEE Conference and Exhibition*, vol. 3, 2000, pp. 2043–2050 vol.3.
- [8] E. Olson, J. Leonard, and S. Teller, “Robust range-only beacon localization,” *Oceanic Engineering, IEEE Journal of*, vol. 31, no. 4, pp. 949–958, Oct 2006.
- [9] IxBlue. “Ramses product brochure.” Web Page. Retrieved 26 Oct 2015. [Online]. Available: <https://www.ixblue.com/sites/default/files/downloads/ixblue-ps-ramses-02-2014-web.pdf>
- [10] “Fusion6g lbl web page.” Accessed 27 Oct 2015. 2015. [Online]. Available: <http://www.sonardyne.com/products/positioning/fusion-6g.html>
- [11] C. Roman and H. Singh, “Consistency based error evaluation for deep sea bathymetric mapping with robotic vehicles,” in *Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on*, May 2006, pp. 3568 –3574.
- [12] N. P. Fofonoff and R. C. Millard, “Algorithms for computation of fundamental properties of seawater,” UNESCO, Tech. Rep., 1983. [Online]. Available: http://www.jodc.go.jp/jodcweb/info/ioc_doc/UNESCO_tech/059832eb.pdf
- [13] P. M. Saunders and N. P. Fofonoff, “Conversion of pressure to depth in the ocean,” *Deep Sea Research and Oceanographic Abstracts*, vol. 23, pp. 109–111, 1976.
- [14] S.-C. Han and F. J. Simons, “Spatiospectral localization of global geopotential fields from the gravity recovery and climate experiment (grace) reveals the coseismic gravity change owing to the 2004 sumatra-andaman earthquake,” *Journal of Geophysical Research: Solid Earth (1978–2012)*, vol. 113, no. B1, 2008.
- [15] G. D. Egbert and S. Y. Erofeeva, “Efficient inverse modeling of barotropic ocean tides,” *J. Atmos. Oceanic Technol.*, vol. 19, no. 2, pp. 183–204, Feb. 2002. [Online]. Available: [http://dx.doi.org/10.1175/1520-0426\(2002\)019<0183:EIMOBO>2.0.CO;2](http://dx.doi.org/10.1175/1520-0426(2002)019<0183:EIMOBO>2.0.CO;2)
- [16] G. D. Egbert, S. Y. Erofeeva, and R. D. Ray, “Assimilation of altimetry data for nonlinear shallow-water tides: Quarter-diurnal tides of the northwest european shelf,” *Continental Shelf Research*, vol. 30, no. 6, pp. 668 – 679, 2010, tides in Marginal Seas - A special issue in memory of Prof Alexei Nekrasov. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0278434309003094>

- [17] J. S. Byrne and V. E. Schmidt, “Uncertainty modeling for auv acquired bathymetry,” *U.S. Hydrographic Conference (US HYDRO)*, 16-19 April 2015 2015. [Online]. Available: http://ccom.unh.edu/sites/default/files/publications/AUV_uncertainty_model_final.pdf
- [18] W. P. Crawford, *Mariner’s Weather*. Norton Nautical Books, 1978.
- [19] *SBE 911plus CTD Sensor*, Sea-Bird Electronics, 13431 NE 20th Street Bellevue, Washington 98005 USA, 2014, accessed 9 Nov 2015. [Online]. Available: http://www.seabird.com/sites/default/files/documents/911plusbrochureDec14_0.pdf
- [20] *SBE 49 FastCAT CTD Sensor Manual*, Sea-Bird Electronics, 13431 NE 20th Street Bellevue, Washington 98005 USA, 2015, accessed 9 Nov 2015. [Online]. Available: http://www.seabird.com/sites/default/files/documents/49_017.pdf
- [21] IOC, SCOR, and IAPSO, “The international thermodynamic equation of seawater 2010: Calculation and use of thermodynamic properties.” UNESCO” Intergovernmental Oceanographic Commission, Manuals and Guides No. 56, 2010. [Online]. Available: http://www.teos-10.org/pubs/TEOS-10_Manual.pdf
- [22] T. J. McDougall and P. M. Barker, *Getting started with TEOS-10 and the Gibbs Seawater (GSW) Oceanographic Toolbox*. SCOR/IAPSO WG127, 2011. [Online]. Available: <http://www.teos-10.org/software.htm>
- [23] J. Mcdonald, M. Kaess, C. Cadena, J. Neira, and J. J. Leonard, “Real-time 6-dof multi-session visual slam over large-scale environments,” *Robot. Auton. Syst.*, vol. 61, no. 10, pp. 1144–1158, Oct. 2013. [Online]. Available: <http://dx.doi.org/10.1016/j.robot.2012.08.008>
- [24] C. Roman and H. Singh, “A self-consistent bathymetric mapping algorithm,” *Journal of Field Robotics*, vol. 24, no. 1-2, pp. 23–50, 2007. [Online]. Available: <http://dx.doi.org/10.1002/rob.20164>
- [25] F. Dellaert and M. Kaess, “Square Root SAM: Simultaneous localization and mapping via square root information smoothing,” *Intl. J. of Robotics Research (IJRR)*, vol. 25, no. 12, pp. 1181–1204, Dec 2006.
- [26] M. Kaess, A. Ranganathan, and F. Dellaert, “iSAM: Incremental smoothing and mapping,” *IEEE Trans. on Robotics (TRO)*, vol. 24, no. 6, pp. 1365–1378, Dec 2008.
- [27] M. Kaess and F. Dellaert, “Covariance recovery from a square root information matrix for data association,” *Journal of Robotics and Autonomous Systems (RAS)*, vol. 57, pp. 1198–1210, Dec 2009.

- [28] R. C. Smith and P. Cheeseman, “On the representation and estimation of spatial uncertainty,” *The International Journal of Robotics Research*, vol. 5, no. 4, pp. 56–68, 1986. [Online]. Available: <http://ijr.sagepub.com/content/5/4/56.abstract>
- [29] R. Smith, M. Self, and P. Cheeseman, “Estimating uncertain spatial relationships in robotics,” in *Robotics and Automation. Proceedings. 1987 IEEE International Conference on*, vol. 4, Mar 1987, pp. 850–850.

CHAPTER 4

Predicting Submap-SLAM Performance from Seafloor Properties

4.1 Submap SLAM

Simultaneous Localization and Mapping (SLAM) methods are a broad class of techniques that use mapping sensors to improve an estimate of both the environment and a robot’s trajectory. Bathymetric Submap-SLAM does this with bathymetric sonars by grouping consecutive multibeam sonar measurements into small patches of terrain known as submaps. These submaps may then be registered to provide additional navigation constraints beyond the initial dead-reckoning navigation solution. Submap-SLAM has been demonstrated to work well in areas where the seafloor has enough structure for submaps to be registered accurately [1]. Submap methods can be extended to exploit recent advances in Factor-Graph SLAM methods for large data sets [2] and work well in conjunction with other graph-based SLAM methods for fusing other data sources, principally stereo cameras [3, 4]. Submap-SLAM only works in regions where the seafloor provides enough information to register submaps correctly. If the seafloor is too flat, registration may produce large errors outside the estimated uncertainty of the registration measurement. These errors are analogous to a data-association failure in feature-based SLAM, and often result in navigation solutions that are dramatically worse than the dead-reckoning solution.

This chapter presents several metrics that can be used to predict how “registerable” a piece of seafloor is and how useful it will be for reducing the final navigational uncertainty. Registerability refers to how well-suited a particular area of the seafloor is to being uniquely registered when seen multiple times. Flat regions will not register well, if at all, while regions with enough structure to well-constrain the registration process are highly registerable.

The resulting metrics may then be used to guide further data collection to minimize the final map uncertainty. In contrast to the active SLAM methods [5, 6], this does not require running the SLAM algorithm online. Instead, the focus is on providing a highly efficient metric that may be employed on vehicles with even the most basic computing hardware available.

Real-time navigation only needs to be good enough to keep the vehicle near its survey track. Many practical applications may use external acoustic navigation, such as Ultra Short Baseline (USBL) positioning, to achieve the required accuracy. Any computing power used to refine the navigation solution beyond what is operationally useful uses power that may be better spent on path-planning or increased run time. For these reasons, this chapter focuses on simple metrics that may be computed directly from the data while still allowing real-time planning software to optimize a survey path that benefits later post-processing.

The requirements of the final metric are as follows. First, it must reflect the effects of the bathymetry on the quality of the final navigation solution. Secondly, the metric must be causal and use only data that has already been observed. The metric should be computationally efficient to save power. Additionally, the registerability metric should reflect the properties of the underlying bathymetry rather than artifacts introduced by the measurement process. Finally, the metric must be insensitive to navigation error.

These represent the general requirements for a metric to be computed online from the dead-reckoning the navigation solution. Insensitivity to navigation artifacts is particularly important. Dead-reckoning navigation errors may introduce artifacts that look like highly-registrable features. An online metric should be robust to such errors.

The rest of this chapter is organized as follows. The Cramer-Rao Lower-Bound

Figure 49: A navigation artifact that appears to be registrable. In this 2014 survey of the Kick'em Jenny volcano, the corner of one line caught the edge of the crater wall. Due to a navigation error, that corner appears by itself almost 3.5m from the actual crater wall. In that context, it appears to be a well-defined spire that should match well rather than the simple artifact it actually is. A 2.5m grid is draped over the bathymetry for x/y scale.

(CRLB) for Submap-SLAM is used in Section 4.2 to show that the information in Submap-SLAM may be reduced to a dead-reckoning part and a submap matching part, thus demonstrating the intuitive result that all navigation improvement comes from submap registration measurements. This result motivates the submap autoregistration metric presented in Section 4.3.1, in which a submap is registered to itself. The autoregistration metric provides the most direct measure of how useful a submap is likely to be to the SLAM solution. Autoregistration requires running the full registration algorithm and may be too computationally intense for some applications. A normal-based metric, similar to normal space occupancy metric of [1], is presented Section 4.3.2. The normal metric is found to be more computationally-intense than autoregistration. A more practical method that measures the slope-corrected z -variance of a portion of seafloor is presented in Section 4.3.3. Both normal and z -variance metrics are compared to the autoregistration

metric in Section 4.4. Finally, discussion is concluded in Section 4.5 with the z -variance metric shown to meet the requirements outline above.

4.2 Predicting SLAM Performance From Submap Properties

4.2.1 Probabilistic Models for SLAM

A Bayesian model for Submap-SLAM is a prerequisite for further analysis. In keeping with the model introduced in the previous chapter, the vehicle's state at time t is represented as

$$\mathbf{x}_t = \left[x, y, z, \theta, \phi, \psi, \dot{x}, \dot{y}, \dot{z}, \dot{\theta}, \dot{\phi}, \dot{\psi} \right]^T.$$

The vehicle trajectory from the initial time until time step t is represented as the set of states $X_t = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_t\} = \{\mathbf{x}_{1:t}\}$. Consecutive bathymetric measurements are grouped into submaps. Each submap is represented as a cloud of points relative to an origin for each submap. Submaps are kept small enough that navigational error within a submap is smaller than the resolution of the mapping sensor. For economy of representation, only the submap origins are represented as $\mathbf{x}_t$, where the submap time is the time of the first ping in that submap. The key approximation of Submap-SLAM is that the SLAM problem may be reduced to estimating the submap origin poses. These submap origin poses may then be used to reconstruct either a complete map or estimated vehicle trajectory for the entire survey.

Following the convention of the Kalman Filter literature, all sensor measurements are denoted $\mathbf{z}$. These measurements are classified into one of three groups, and these groups are identified with their indices. The first class are direct measurements of the vehicle state, typically heading, pitch, roll and depth. These measurements are filtered and interpolated to the time indices of interest and denoted $\mathbf{z}_i$.

An Augmented State Extended Kalman Filter (EKF) is used to collect measurements from the vehicle's dead-reckoning sensors into a dead reckoning measurement $\mathbf{z}_{i,i+1}$ modeled as a 6-dimensional multivariate Gaussian constraint with covariance $R_{i,i+1}$ between consecutive states i and $i + 1$ [2]. The portions of the vehicle's state that can be measured directly, principally heading, pitch, roll and depth, are interpolated to the time of each multibeam ping and collected into $\mathbf{z}_i$. The error for these measurements is assumed to be Gaussian with covariance R_i . Together, constraints between poses $\{z_{i,i+1}\}$ and absolute constraints on individual poses $\{z_i\}$ contain all information from the dead-reckoning sensors. Solving for the trajectory $\mathbf{x}_{1:N}$ that minimizes the error in these observations gives the dead-reckoning measurement.

Submap-SLAM improves the final navigation solution by adding submap registration measurements between overlapping submaps. Let the set of all pairs of submaps that actually overlap be represented as a set of index pairs $M = \{j, k\}, j \neq k$ such that j, k are in M if and only if submaps j and k overlap for the true location of their origin. Registering these submaps gives a measurement $\mathbf{z}_{j,k}, j \neq k$ with covariance matrix $R_{j,k}$.

All the measurements, including odometry, absolute navigation sensors, and submap registration measurements are combined into an estimation framework in order to estimate a navigation solution with lower variance than a purely dead-reckoning-based solution. In this case, the log-likelihood function of the measurements, with normalizing constant η , is proportional to

$$\begin{aligned} \log p(\text{all } z | X_{1:N}) &= \log \eta + \sum_{i=0}^{N-1} \|(\mathbf{x}_{i+1} \ominus \mathbf{x}_i) - \mathbf{z}_{i,i+1}\|_{R_{i,i+1}} \\ &\quad + \sum_{i=0}^N \|\mathbf{x}_i - \mathbf{z}_i\|_{R_i} + \sum_M \|(\mathbf{x}_j \ominus \mathbf{x}_k) - \mathbf{z}_{j,k}\|_{R_{j,k}}. \end{aligned}$$

where $\|\cdot\|_R$ denotes $\|\cdot\|_R$ is the Mahalanobis distance with covariance R , or $\|\cdot\|_R =$

$\cdot^T R^{-1} \cdot$. The $\ominus$ operator is the relative pose difference operator of Smith and Cheeseman [7]. Alternately, these constraints may be used to construct a linear least-squares problem where the error for each measurement is weighted by the square-root information matrix. Such constructs are the basis of Factor-Graph SLAM methods. Following the derivation in the Factor Graph Smoothing and Mapping paper of [8], the problem may be linearized in matrix form as

$$\hat{X} = \underset{X}{\operatorname{argmin}} ||AX - b||^2$$

$$= \underset{X}{\operatorname{argmin}} \left\| \begin{bmatrix} F_2 & G_2 & & \cdots & & & \\ & F_3 & G_3 & & \cdots & & \\ & & F_4 & G_4 & \cdots & & \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ & & & & \cdots & G_{N-1} & \\ & & & & \cdots & F_N & G_N \\ \hline K_1 & & & & \cdots & & \\ & K_2 & & & \cdots & & \\ & & K_3 & & \cdots & & \\ & & & K_4 & \cdots & & \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ & & & & \cdots & K_{N_1} & \\ & & & & \cdots & & K_N \\ \hline H_{1,N} & & & & \cdots & & J_{1,N} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ & & & H_{4,N-1} & \cdots & J_{4,N-1} & \end{bmatrix} X - \begin{bmatrix} \mathbf{z}_{1,2} \\ \mathbf{z}_{2,3} \\ \mathbf{z}_{3,4} \\ \vdots \\ \mathbf{z}_{N-2,N-1} \\ \mathbf{z}_{N-1,N} \\ \hline \mathbf{z}_1 \\ \mathbf{z}_2 \\ \mathbf{z}_3 \\ \mathbf{z}_4 \\ \vdots \\ \mathbf{z}_{N-1} \\ \mathbf{z}_N \\ \hline \mathbf{z}_{1,N} \\ \vdots \\ \mathbf{z}_{4,N-1} \end{bmatrix} \right\|^2.$$

As in Dellart and Kaess's 2006 introduction to Factor Graph SLAM [8], the Jacobians are weighted by multiplying the measurement covariance matrices from the left to simplify the problem to a simple least-squares solution. F_i and G_i are the weighted Jacobians of the of the dead-reckoning measurement model with respect to $\mathbf{x}_{i-1}$ and $\mathbf{x}_i$, respectively. Similarly, K_i is the Jacobian of the absolute measurement model with respect to $\mathbf{x}_i$, and $H_{j,k}$ and $J_{j,k}$ are the Jacobians of the measurement model between submaps j and k with respect to $\mathbf{x}_j$ and $\mathbf{x}_k$,

respectively. The topmost partition contains the constraints generated between consecutive poses by dead-reckoning. The middle partition contains constraints resulting from absolute measurements of vehicle state, e.g., heading. The bottom partition results from registration measurements between submaps. There are no landmarks in this construction. All states in X are vehicle poses at some time index of interest.

4.2.2 An Information View of Submap-SLAM

The Cramer Rao Lower Bound (CRLB) is an important tool in estimation theory that provides a lower bound for the variance of an unbiased estimate $\hat{X}_{1:N}$ of a true quantity $X_{1:N}$ from the log-likelihood $\log p(X_{1:N})$. For this estimation problem, the true trajectory $X_{1:N}$ is treated as the parameter vector. The Cramer-Rao Lower bound is the inverse of the Fisher Information matrix of a given estimator. Recall that the linearized estimation problem of the previous section was pre-multiplied with the measurement covariance matrices to simplify the problem. Since each measurement is independent, the linearized problem may be modeled as having additive noise that is Independent and Identically Distributed (IID). IID linear estimation problems have the well-known to have an information matrix [9] given by

$$I(X_{1:N}) = A^T A.$$

This is the limit for the linearized estimation problem, and not the true non-linear problem initially introduced. Thus, although the linearized problem attains the CLRB, the non-linear version does not. The linearized problem, however, provides the appropriate insights with much simpler math.

Expanding $A^T A$ by partitioning the A matrix into the dead-reckoning / absolute measurement piece D and submap matching piece M gives

$$A^T A = \begin{bmatrix} D^T & M^T \end{bmatrix} \begin{bmatrix} D \\ M \end{bmatrix} = (D^T D + M^T M).$$

$D^T D$ is the Fisher Information Matrix for the dead-reckoning navigation problem. This result agrees with the property that Fisher Information of independent measurements simply adds. The contribution from submap links to the Fisher Information is $M^T M$. Conceptually, the “larger” this matrix product, the more information has been obtained from submap matching and the lower the uncertainty in the final navigation estimate. The magnitude of a matrix is not itself a simple concept. The matrix M is structured as

$$M = \begin{bmatrix} M_1 \\ M_2 \\ \vdots \\ M_L \end{bmatrix}$$

with each block of rows M_ℓ giving the results of a single submap-to-submap registration. Applying the same partitioning method to the rows of matrix M shows that the Fisher Information of the submap measurements as a whole is the sum of the individual measurements m_ℓ . For L submap-to-submap measurements, this is

$$M^T M = \sum_{\ell=0}^L M_\ell^T M_\ell.$$

Each row of the matrix M comes from a single submap-to-submap match. Taking each row-block M_ℓ , the resulting products $M_\ell^T M_\ell$ are outer products and, for N poses to estimate, are $6N \times 6N$. Each product $M_\ell^T M_\ell$ gives the information gain from each submap match. More submap matches increases the total information. This observation agrees with the obvious intuition.

Recall that each row-block of the matrix A was multiplied through on the left by its square-root information matrix in order to reduce the Mahabois distance error metric to a simple least-squares criterion. The weighted Jacobian matrices

$H_{j,k}$ and $J_{j,k}$ are composed of the product of the square-root information matrix for the link, $R_{j,k}^{-1/2}$ as well as the unweighted Jacobians $\bar{H}_{j,k}$ and $\bar{J}_{j,k}$ as

$$H_{j,k} = R_{j,k}^{-1/2} \bar{H}_{j,k} \quad J_{j,k} = R_{j,k}^{-1/2} \bar{J}_{j,k}.$$

The outer product $M_\ell^T M_\ell$ may thus be expanded as

$$M_\ell^T M_\ell = \begin{bmatrix} (R_{j,k}^{-1/2} \bar{H}_{j,k})^T \\ (R_{j,k}^{-1/2} \bar{J}_{j,k})^T \end{bmatrix} \begin{bmatrix} R_{j,k}^{-1/2} \bar{H}_{j,k} & R_{j,k}^{-1/2} \bar{J}_{j,k} \end{bmatrix}$$

$$= \begin{bmatrix} \bar{H}_{j,k}^T R_{j,k}^{-T/2} R_{j,k}^{-1/2} \bar{H}_{j,k} & \bar{H}_{j,k}^T R_{j,k}^{-T/2} R_{j,k}^{-1/2} \bar{J}_{j,k} \\ \bar{J}_{j,k}^T R_{j,k}^{-T/2} R_{j,k}^{-1/2} \bar{H}_{j,k} & \bar{J}_{j,k}^T R_{j,k}^{-T/2} R_{j,k}^{-1/2} \bar{J}_{j,k} \end{bmatrix}.$$

Noting that $R_{j,k}^{-T/2} R_{j,k}^{-1/2} = R_{j,k}^{-1} = R_\ell^{-1}$ and collecting the raw Jacobians into a matrix $\bar{M}_{j,k} = [\cdots \bar{H}_{j,k}^T \cdots \bar{J}_{j,k}^T \cdots]^T = \bar{M}_\ell$, the information content of a single link ℓ from pose j to k can be expressed as

$$M_\ell^T M_\ell = \bar{M}_\ell^T R_\ell^{-1} \bar{M}_\ell.$$

The raw Jacobian matrices transform the information matrix from the measurement reference frame into constraints on the two poses j and k . The effect of this transform on the magnitude of the information added by the link is therefore minimal. The information added to the estimation problem by a link is therefore dominated by the covariance of that link.

This derivation of the information added by a single pose-to-pose constraint assumes an omniscient submap registration algorithm that can register any submap pair presented that actually overlaps. Sensor noise, small overlap area, and other issues prevent the majority of overlapping submaps from generating valid links. Overlap that fails to result in a valid link measurement removes the measurement $\mathbf{z}_{j,k}$ from the pose graph, reduces the overall amount of Fisher information, and increases navigation error in the final map. Failed submap links are therefore a missed opportunity to use available environment data to refine the navigation solution. Failed submap links, where the submap registration pipeline misses an opportunity to make a valid link, can not degrade the navigation solution below the quality of the dead-reckoned solution. Invalid links, where the registration process produces a link with significant errors in uncertainty, are a far more significant problem.

4.2.3 Submap-to-Submap Link Validity

Invalid links pose a much larger problem than links with high variance, as they can corrupt the navigation solution beyond the uncertainty represented in the SLAM optimization framework. Invalid links typically occur in regions of poor registerability. A reasonable definition is required before discussing how submap registration performance metrics may be used to avoid invalid links.

Submap registration borrows from the point cloud registration literature and is typically described as estimating a relative pose between two submaps. SLAM frameworks require a probability density function (PDF) over the relative pose $\mathbf{z}_{j,k}$ between the two submaps in question. The usual SLAM registration pipelines solve this by assuming that the estimated relative pose measurement $\hat{\mathbf{z}}_{j,k}$ is the mean of a Gaussian PDF with an estimated or assumed covariance matrix. The goal of submap relative pose estimation is to produce an estimate of the PDF over

the relative pose between submaps conditioned on the bathymetry itself, $\hat{p}(\mathbf{z}_{j,k})$. This PDF is usually assumed to be Gaussian, that is

$$\hat{p}(\mathbf{z}_{j,k}) \approx \mathcal{N}(\hat{\mathbf{z}}_{j,k}, \hat{R}_{j,k}).$$

Ideally, this PDF is centered around the true value $\tilde{\mathbf{z}}_{j,k}$. Given any two true poses $\tilde{\mathbf{x}}_j$ and $\tilde{\mathbf{x}}_k$, the true relative pose offset is given through use of Smith, Self and Cheeseman’s compounding operator [7, 10] as

$$\tilde{\mathbf{z}}_{j,k} = \tilde{\mathbf{x}}_j \ominus \tilde{\mathbf{x}}_k.$$

Intuitively, a link may be said to be “valid” if the true measurement is reasonably probable given the estimated measurement and associated uncertainty model. Similarly, a link is “invalid” if the true relative pose is extremely improbable under the resulting PDF. The ideal registration process produces an estimate $\hat{\mathbf{z}}_{j,k}$ that closely agrees with this true value and is independent of navigation measurements. A “valid” link is an estimated measurement and uncertainty model such that the probability of the true value $\tilde{\mathbf{z}}_{j,k}$ given the estimated PDF is acceptably high. That is:

$$\text{A link } \mathbf{z}_{j,k}^{\wedge} \text{ is: } \begin{cases} \text{“Valid” iff } & \hat{p}(\tilde{\mathbf{z}}_{j,k}) \geq \tau \\ \text{“Invalid” iff } & \hat{p}(\tilde{\mathbf{z}}_{j,k}) < \tau \end{cases}$$

where τ is some acceptable threshold value. If the estimated PDF is assumed to be Gaussian, then it is expedient to re-write the threshold as a Mahalanobis distance, i.e.,

$$\text{A link } \mathbf{z}_{j,k}^{\wedge} \text{ is: } \begin{cases} \text{“Valid” iff } & \|\tilde{\mathbf{z}}_{j,k} - \hat{\mathbf{z}}_{j,k}\|_{\hat{R}_{j,k}} \leq \tau' \\ \text{“Invalid” iff } & \|\tilde{\mathbf{z}}_{j,k} - \hat{\mathbf{z}}_{j,k}\|_{\hat{R}_{j,k}} > \tau' \end{cases}$$

where τ' is the original threshold τ transformed appropriately.

The uncertainty estimate is therefore an integral part of submap registration. Most submap registration algorithms minimize a cost function over the relative

pose between submaps to arrive at an optimal estimate of relative pose as

$$\hat{\mathbf{z}}_{j,k} = \underset{\mathbf{z}_{j,k}}{\operatorname{argmin}} J(\mathbf{z}_{j,k}).$$

A number of different cost functions and minimization procedures may be used. The Iterated Closest Point (ICP) family of algorithms use a wide variety of point cloud to point cloud error metrics and several different optimization methods to arrive at an estimate of the relative pose offset [11]. For 2D or 3D bathymetric matching, the simpler Sum of Squared Differences (SSD) metric has also proven popular [1, 2, 3]. Estimating the uncertainty of a link has received less attention in the literature. Most implementations either assume a typical value or compute an uncertainty specific to the terrain underlying each link by fitting an uncertainty model to the cost function. For example, several implementations fit a quadratic to the cost function surface near the estimated peak and use that as the measurement covariance matrix $\hat{R}_{j,k}$. In this way, the link uncertainty model reflects the terrain shape and the corresponding degree of certainty in the match.

Invalid submap links come from a number of possible sources. One common cause of invalid links is attempting to register two submaps that do not actually overlap. Even when submaps do not actually overlap, the cost function often has a local minima that may be accepted by the registration algorithm. Errors in the submaps, either from measurement noise or a systematic error like a lever-arm offset, can also cause local minima in the cost function to surpass the true global minimum. This issue is well-known and has been documented since the early attempts at correcting shipboard SeaBeam Multibeam data during the pre-GPS era [12, 13].

These issues in submap matching are analogous to errors in data association with feature-based SLAM methods. As with feature-based methods, data association errors are not modeled by the SLAM framework and typically introduce large

errors in the final estimated trajectory. Identifying and filtering out these data association errors is the most significant practical problem in Submap-SLAM. The only widely-available tool for terrain-based navigation corrections to bathymetry, MB-System’s `mbnavadjust` program, therefore recommends that most links be verified by a human. Although `mbnavadjust` is not a SLAM algorithm, the manual recommends the operator take great care to avoid adding invalid links [13].

In practice, the true relative pose offset $\tilde{\mathbf{z}}_{j,k}$ is unknown and this validity metric cannot be evaluated. Links must therefore be classified as either “valid” or “invalid” based on observable link characteristics. Invalid links are errors not adequately modeled by the estimator model and are thus similar to outliers. The estimation frameworks most commonly employed for bathymetric SLAM are not robust to erroneous inputs, so even a single invalid link can significantly degrade the final SLAM result. Invalid links most commonly occur when a local minima corrupt the matching process [12, 13]. Focusing survey effort and crossing lines on parts of the survey with good registration properties helps reduce the likelihood of generating invalid links. An efficient means of identifying such regions is an important first step in devising a survey method to plan the vehicle’s path accordingly.

4.3 Metrics to Evaluate Possible Link Quality

4.3.1 Autoregistration

As shown in Section 4.2.2, any improvement in the final navigation solution from SLAM depends on high-quality, low-covariance links between submaps. The quality of valid links and the percentage of invalid links varies significantly with submap content. The most direct way to observe this effect is to register a submap to itself and examine the resulting cost function and covariance structure. Registering a submap to itself, or “autoregistration,” is a best-case scenario for any

registration algorithm. There is no noise, distortion, sampling irregularities, or edge effect issues that can confuse or corrupt the match. As autoregistration uses only standard output from SLAM-ready registration algorithms it may be used with a wide variety of registration methods. In the following section, a sum of squared differences (SSD) metric is used [2]. An example of autoregistration comparing highly-structured and uninteresting submaps is presented in Figure 50.

Autoregistration may be considered an analogy to the variance or autocorrelation of the underlying submap. For correlation-based submap registration techniques, the cost function of registering a submap to itself is the submap’s autocorrelation function. Fitting an error model to the autocorrelation function measures the “width” of the autocorrelation peak. A narrow autocorrelation peak implies a wide spectral density function [14]. Autocorrelation has a long history in signal processing to predict time-of-arrival estimation performance in sensing and communications [9, 15].

Autoregistration produces an estimate of the uncertainty in matching a link to itself in the form of a covariance matrix. For future analysis, it is helpful to reduce this covariance matrix to a scalar score for each submap. In this section, a geometric argument will be used to show that the largest eigenvalue of the measurement covariance matrix is a natural choice for an uncertainty score.

It is helpful to visualize the uncertainty of a submap relative pose measurement. In one dimension, uncertainty in a Gaussian is often represented as error bars at some confidence bound from the mean. For a multivariate Gaussian, these error bars are extended to multidimension ellipsoids. Although this extension works in any number of dimensions, visualization on the page is simplified by considering only the 2×2 covariance matrix of the submap measurement corresponding to the x and y terms of the submap pose measurement. In this case, the uncertainty

Figure 50: A comparison of highly-structured and mostly-flat submaps for registration quality. Note how the uncertainty, shown as white concentric ovals in subfigures (c) and (d), are much larger for the poorly-registerable submap. These submaps are from different sections of the large pre-disturbance survey of Monterey A at 1300m depth.

bounds are represented as ellipses. Examples of such error ellipses are shown in white in Figure 50. The natural choice of a score to represent the uncertainty of a given Gaussian is the radius of the circle that contains the error ellipse. Geometrically, this is the semi-major axis of the error ellipse. This is also the variance along the direction with the most variation. Principal component analysis (PCA) is a standard method to find the direction of maximum variation in a dataset by looking at the eigenvalue / eigenvector decomposition of the data covariance matrix [16]. Applying the same method, the variance of the principle component of the measurement covariance matrix is the maximum eigenvalue of the covariance matrix. Intuitively, the seafloor does not constrain position in all directions equally.

A sloping seafloor constrains navigation well in the up-slope / downslope direction and poorly in the along-slope direction. The score used is uncertainty in the worst possible direction.

It is also useful to consider an error in meters, rather than squared meters, so the final score uses the square root of the largest eigenvalue of the covariance matrix. This figure behaves like a standard deviation on link uncertainty. Lower values indicate less uncertainty in the submap match and therefore more information available to improve the navigation solution through SLAM methods.

Once the performance of submap autoregistration has been reduced to a scalar, it is possible to plot this value over the area of a map to visualize the effect of terrain variation. To do this, the autoregistration score for each submap was assigned to each sounding that made up that submap. This value was then gridded using a standard gridding algorithm in QPS's DMagic software and draped on the bathymetry. The remaining metrics of this chapter will be presented in a similar fashion. The autoregistration score is an error with smaller values indicating better registration performance. To ease visual comparison with metrics in the upcoming sections, the autoregistration score is present with a colormap from brown to red representing poor registerability to white over areas that are highly registerable.

Visualizations of the autoregistration metric are shown for Monterrey A in Figure 51b, Knidos F in Figure 52b, the Kula mud volcano summit in Figure 53b, and Mt. Dent in Figure 54b. Several interesting phenomena are apparent. The autoregistration uncertainty is smaller over areas of variable bathymetry, for example, over the ship remains at Monterrey A and Knidos F. The error is also larger in submaps that occurred at the end of a survey line. The Monterrey A survey track runs parallel to the ship. The Knidos F site was surveyed from two directions 90° apart, so line ending submaps occur at all four edges of the site.

The Kula site was surveyed from left to right for the image orientation seen in Figure 53. The poor registration performance of these submaps is an artifact of their shape rather than the underlying bathymetry. Autoregistration shows how both the terrain properties and the shape of the submaps affect the registration uncertainty.

Autoregistration represents a “gold-standard” figure-of-merit for submap matching performance, as it requires actually carrying out the matching procedure, which is computationally expensive and likely prohibitive for real-time applications. This motivates the need for a metric that is easier to compute. In cases where computing power is not strictly restricted, autoregistration does provide high-fidelity results at modest computation cost, requiring an average of only 4.939 seconds per submap over all four datasets presented on a Core i5 desktop. By comparison, each submap required approximately 67 seconds of data collection, showing that autoregistration runs far faster than real-time.

4.3.2 Normal Projection as an Alternate Metric for Matching

An alternative registerability metric examines the normal vector for each point in a submap relative to the normal vector of a best-plane fit to the submap. Specifically, it looks for the length of the component of a point’s normal vector that is orthogonal to the submap’s normal vector. Submaps with a great deal of 3D structure will have many points with a normal vector that is different from the normal vector among all points. The point-wise normal vectors in planar submaps will generally agree strongly with the submap-wide normal vector. This method is similar to the normal-space occupancy metric used to construct the submaps in [1].

The component of a point i ’s normal vector $\hat{\mathbf{n}}_i$ that is orthogonal to the submap normal vector $\hat{\mathbf{n}}_s$ is indicative of how much an individual point’s slope

differs from the submap’s slope. This component, $\mathbf{v}_i$, may be computed as

$$\mathbf{v}_i = \hat{\mathbf{n}}_i - (\hat{\mathbf{n}}_i \cdot \hat{\mathbf{n}}_s)\hat{\mathbf{n}}_s.$$

The value of the normal metric for the N points with valid computed normals is then

$$v_s = \frac{1}{N} \sum_{i=1}^N \|\mathbf{v}_i\|.$$

Some care is required to account for possible sign differences as the signs of point cloud surface normals are often ill-defined [17]. Point-wise normal calculation was done using the Point Cloud Library [18] and requires looking at a neighborhood of points around the given point. Inspection of the point clouds showed three grid cells, or 7.5cm, to be an effective search radius for normal calculation with the BlueView MB1350-90 sonar data. Areas where multibeam bottom detection was poor, such as at the edge of the swath, had few enough points that normal estimation was noisy.

To make the comparison to the other metrics, the normal metric was computed for each individual submap. The value for a given submap was assigned to all soundings in that submap, and the result was gridded in DMagic using the same parameters as the autoregistration metric. The results are shown in Figures 51d, 52d, 53d and 54d. Computing the normal vector for the submap is quite quick. The per-point normal vectors can be quite time-consuming to compute even when a kD-tree is used for nearest-neighbor searching. An average of 18.43 seconds per submap was required to compute this metric. Although it still runs in real-time, it is far slower than autoregistration and does not lend itself to computation on a rolling-window basis.

4.3.3 z -Variance as a Fast Metric for Matching Performance

The simplest method of detecting 3D structure simply looks at the variance of the z -coordinate of every sounding in a submap. In practice, it is helpful to correct for the slope and other effects. This may be done by fitting a plane to the submap and looking only at the component of the sounding coordinates normal to that plane. Rather than actually fitting a plane, it is usually faster to use a Principal Component Analysis (PCA) of the same type used for normal estimation. Specifically, given a set of N submap soundings $\mathbf{x}_i, i \in \{1, \dots, N\}$ as 3×1 column vectors, one first computes the 3×3 covariances matrix of the soundings as

$$C_s = \frac{1}{N-1} \begin{bmatrix} \mathbf{x}_1 & \mathbf{x}_2 & \dots & \mathbf{x}_N \end{bmatrix} \begin{bmatrix} \mathbf{x}_1^t \\ \mathbf{x}_2^t \\ \vdots \\ \mathbf{x}_N^t \end{bmatrix}.$$

An eigenvalue decomposition is then used to perform a Principal Component Analysis of the pointcloud. The largest eigenvalue corresponds to the direction with the most variation in position and is usually the along-track direction of travel. The shortest direction has the least variation and is often used as the normal estimate [16, 17].

This metric is quite fast to compute. The dot product may be broken up into sums over the dot products of individual pings, making it easy to compute a rolling covariance matrix over a fixed number of pings. Reducing the per-sounding operations to a simple dot product and sum makes this method computationally cheap. The relatively expensive eigenvalue decomposition is only carried out over a small 3×3 matrix and is only evaluated once per submap.

Although it can be cheaply evaluated over a rolling window, for consistency with other metrics a single value was computed for each submap, assigned to each sounding in that submap and gridded in DMagic as with the other two metrics. The results are shown in Figures 51c, 52c, 53c and 54c.

One limitation of this method is that it is subject to effects from sampling density. The raw points are used as-is, and in an online implementation a fixed number of pings would be used to evaluate the metric. If the vehicle sits in one place or has a highly variable survey speed, the z variance may be evaluated over different spatial windows. In practice, surveys are run at consistent altitude and speed, which limits this effect. If it proves problematic, the point cloud being evaluated could be re-sampled to a consistent spatial sampling rate by gridding it.

4.4 Comparison of Autoregistration and Normal Variance

Figure 51: Autoregistration and metric results for Monterrey A. To show scale, 5m grid lines were draped over the survey surface.

Figure 52: Registerability metrics for Knidos F. Grid lines are draped on the bathymetry and spaced 5m apart.

The autoregistration, z -variance, and orthogonal-normal metrics for registerability are presented along with the bathymetry for Monterrey A in Figure 51, Knidos F in Figure 52, the Kula mud volcano summit in Figure 53, and finally for the Mt. Dent site in Figure 54.

The three registerability metrics generally agree. As expected, all three metrics heavily favor the region directly over the shipwreck remains in the Monterrey A and Knidos F datasets. These regions have a great deal of 3D structure and are expected to register well. The normal metric appears to approximate the autoregistration results more closely than the z -variance method does on Monterrey A, but the z -variance approach still correctly identifies the most useful area as the

Figure 53: “Registerability” metric results for the Kula Mud Volcano summit. 5m grid lines are draped over the surface for scale.

shipwreck itself. Many of the small, degenerate submaps seen in the autoregistration results for Monterrey A and Knidos F are not as strongly flagged as poor in the z -variance or normal metric maps. This suggests that these methods are more robust to submap creation issues than the autoregistration metric.

Unlike the shipwreck datasets, the Kula dataset does not include both obviously excellent and obviously poor regions of registerability. All three metrics agree that a region near the center of the map is highly registerable, but differ slightly on exactly which regions are most useful. As on the shipwreck datasets, the z -variance and normal metrics appear more robust to, although not quite immune from, edge effects.

(a) South side of Mt. Dent, bathymetry

(b) Mt. Dent autoregistration

(c) Mt. Dent Z-variance metric

(d) Mt. Dent normal metric

Figure 54: Autoregistration and metric results for a 2013 survey of a vent field on the south side of Mt. Dent. Navigation error corrupts the final bathymetric surface, but the registration metrics still pick out the vents themselves as registerable seafloor. 5m grid lines are draped over the surface for scale.

The Mt. Dent dataset is particularly interesting as it occurs on a hillside. The bathymetry shown is dead-reckoned and suffers from significant navigation error. All three metrics still pick out the vent field as the most registerable area of the survey. The normal metric appears to be a better approximation to the autoregistration results, especially at Monterrey A and Mt. Dent.

Scatter plots are shown comparing the autoregistration results to the z -variance results in Figure 55a and the normal metric in Figure 55b. Each point on the scatter plot is a single submap from the dataset indicated by its color. Although there is strong correlation within each dataset, there is much less correlation between datasets. Furthermore, the relationship between each metric and

the autoregistration often breaks down as the autoregistration becomes susceptible to invalid links. Many of these submaps occur at the end of a line.

The z -variance metric is dramatically faster than the other two methods, requiring an average of 1.7 milliseconds per submap compared with the autoregistration metric's average 4.959 seconds per submap (see Table 5). Combined with the ease of using a rolling window of submaps instead of requiring submaps and the ease of implementation, the z -variance approach is ideally suited to a real-time registerability prediction approach.

Figure 55: z -variance and normal metrics compared to the autoregistration metric for several different datasets.

Autoregistration	z -Variance	Normal
4.939 s	0.002 s	18.4332 s

Table 5: Average registerability metric processing time on an Intel Core i5-4670K desktop CPU over all four datasets

4.4.1 Relationship with Rugosity

Rugosity is a common metric used in habitat classification, geomorphology, and many other applications [19, 20, 21]. Rugosity is traditionally as the ratio of the length of a chain laid over the terrain to the total length of that path projected into a plane. With recent advances in underwater imaging, rugosity is

Figure 56: Rugosity metrics over Kula Summit with the normal metric for comparison. Compare with the registerability metrics of Figure 53

now commonly taken as the ratio of surface area to the surface area projected into a plane-of-best-fit over a given seafloor window [22]. The rugosity definition of Friedman et. al. over a mesh with triangles of individual area a_j and normal vector $\hat{\mathbf{n}}_j$ is

$$r = \frac{\sum_j a_j}{\sum_j a_j |\hat{\mathbf{p}} \cdot \hat{\mathbf{n}}_j|}$$

where $\hat{\mathbf{p}}$ is the normal vector for the plane-of-best-fit over the region of interest. Rugosity is thus the inverse of an area-weighted mean of the dot product of the triangle normal and the submap normal.

Rugosity is typically calculated over a small window. Values calculated using the maximum 12×12 window in QPS's Fledermaus software are shown in Figure

56c. This looks very similar to a gradient magnitude. For consistency with the rest of the analysis in this chapter, rugosity calculated over entire submaps is shown for the Kula dataset in Figure 56d. Note the similarity with the normal metric in Figure 56b.

Areas where the plane-of-best-fit normal tends to agree with the surface normal of an individual triangle will have a rugosity close to one. Areas where the two do not commonly agree will have higher rugosity. The normal metric was intended to measure exactly this. Both metrics look for how a local normal varies from the submap-wide normal, much as how variance is commonly used to look for the extent of variation about the mean. Rugosity and the normal metric are thus much like a variance defined over the normal vectors in that they measure how much the normal vectors of a submap vary. Submaps with a great deal of normal vector variation have long been expected to register better than those where the local normal vectors are close to the submap-wide normal vector [1]. This may be seen in the similarity to the autoregistration metric of Section 4.3.1.

4.5 Implications of Submap-Matching Limitations for SLAM and Survey Planning

The z -variance metric provides a causal, fast, approximation to the registerability that could be used a first step in on-line planning to optimize crossing and verification lines. These lines generate the most useful links between submaps and are an especially important part of the process of collecting data for use with SLAM. Placing these survey lines over the most useful part of the survey area is a potentially useful optimization, especially in areas with highly variable structure. Using regions with the most structure available may also prevent generating invalid links that introduce unmodeled error into the SLAM navigation solution. Most of the ROV *Hercules* surveys are conducted by humans who are quite adept at

optimizing cross line planning. These metrics are a useful first step in extending this capability to AUVs and other autonomous platforms.

In addition to the possibilities for on-line planning, some of these metrics may also be useful for survey planning before beginning a survey. Mapping data from previous surveys, although usually at a lower resolution, is often available during pre-survey planning. These metrics provide a principled way to select regions for crossing and verification lines.

These interest metrics are also intermediate steps towards quantifying the quality of coverage. Submaps that register well contain a lot of structure that can be seen at the scale of the submap. This means both that the resolution of the bathymetric sensor is high enough to measure fine-grained structure and also that the swath width is wide enough to capture the structure that is present. If the interesting features to be measured are smaller than the resolution of the bathymetric sensor, they will be blurred, poorly resolved, and contribute less to refining the navigation solution than if they had been measured properly. If the swath width is much narrower than features to be imaged, submaps will appear flat and register poorly. Narrow swaths also mean that more survey lines are required to cover the same area. More survey lines in turn mean that more distance must be covered in a given survey and that more dead-reckoning error will accumulate over the course of the survey. These results suggest there is an optimal survey altitude that depends on the characteristics of the sensor and the structure of the seafloor being measured. The results of this chapter do not attempt to find that altitude, and it remains a topic for future research.

4.6 Summary

Issues with submap matching in poorly-constrained areas has been a long-understood issue with using terrain to improve navigation [12]. The metrics pre-

sented in this chapter provide a way to evaluate the usefulness of surveyed areas in real-time for navigation refinement through SLAM and other methods. The z -variance metric has been shown to correlate strongly with an autoregistration metric derived from the CRLB for Submap-SLAM as performed with Factor Graphs. It is computationally cheap enough that almost any vehicle with real-time access to survey data may use it to predict which regions are most likely to benefit the final SLAM solution and works well on raw data that has not yet been partitioned into submaps. The z -variance metric also shows less influence from malformed submaps at the end of lines than autoregistration, although some influence is clearly apparent (e.g., Figure 54c). Finally, by processing pings as they occur without requiring a map in a global reference frame, the z -variance approach is relatively robust to navigation errors from using onboard dead-reckoning-only navigation. The proposed z -variance metric has or approximates many of the properties of the ideal metric outlined at the beginning of this chapter. Although subject to some limitations, it may provide a computationally cheap basis for online planning of crossing lines.

List of References

- [1] C. Roman and H. Singh, “A self-consistent bathymetric mapping algorithm,” *Journal of Field Robotics*, vol. 24, no. 1-2, pp. 23–50, 2007. [Online]. Available: <http://dx.doi.org/10.1002/rob.20164>
- [2] I. Vaughn, “Microbathymetry using self-contained navigation and simultaneous localization and mapping,” Master’s thesis, University of Rhode Island, 2012.
- [3] C. Kunz and H. Singh, “Map building fusing acoustic and visual information using autonomous underwater vehicles,” *Journal of Field Robotics*, vol. 30, no. 5, pp. 763–783, 2013. [Online]. Available: <http://dx.doi.org/10.1002/rob.21473>
- [4] G. Inglis, “Hybrid optical acoustic seafloor mapping,” Ph.D. dissertation, University of Rhode Island, 2013.

- [5] C. Leung, S. Huang, and G. Dissanayake, “Active slam using model predictive control and attractor based exploration,” in *Intelligent Robots and Systems, 2006 IEEE/RSJ International Conference on*, Oct 2006, pp. 5026–5031.
- [6] H. Carrillo, I. Reid, and J. Castellanos, “On the comparison of uncertainty criteria for active slam,” in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, May 2012, pp. 2080–2087.
- [7] R. C. Smith and P. Cheeseman, “On the representation and estimation of spatial uncertainty,” *The International Journal of Robotics Research*, vol. 5, no. 4, pp. 56–68, 1986. [Online]. Available: <http://ijr.sagepub.com/content/5/4/56.abstract>
- [8] F. Dellaert and M. Kaess, “Square Root SAM: Simultaneous localization and mapping via square root information smoothing,” *Intl. J. of Robotics Research (IJRR)*, vol. 25, no. 12, pp. 1181–1204, Dec 2006.
- [9] S. Kay, *Fundamentals of Statistical Signal Processing, Estimation Theory*. Prentice Hall PTR, 1993, vol. 1.
- [10] R. Smith, M. Self, and P. Cheeseman, “Estimating uncertain spatial relationships in robotics,” in *Robotics and Automation. Proceedings. 1987 IEEE International Conference on*, vol. 4, Mar 1987, pp. 850–850.
- [11] S. Rusinkiewicz and M. Levoy, “Efficient variants of the icp algorithm,” in *3-D Digital Imaging and Modeling, 2001. Proceedings. Third International Conference on*, 2001, pp. 145 –152.
- [12] C. Nishimura and D. Forsyth, “Improvements in navigation using seabeam crossing errors,” *Marine Geophysical Researches*, vol. 9, no. 4, pp. 333–352, 1988. [Online]. Available: <http://dx.doi.org/10.1007/BF00315004>
- [13] D. Caress, *mbnavadjust(1) manpage*, version 5 ed., June 2013.
- [14] K. S. Shanmugan and A. M. Breipohl, *Random Signals*. John Wiley & Sons, 1988.
- [15] S. M. Kay, *Fundamentals of Statistical Signal Processing, Detection Theory*. Prentice Hall PTR, 1998, vol. 2.
- [16] I. T. Jolliffe, *Principal Component Analysis*. Springer-Verlag, 1986.
- [17] R. B. Rusu, “Semantic 3d object maps for everyday manipulation in human living environments,” Ph.D. dissertation, Computer Science department, Technische Universitaet Muenchen, Germany, October 2009.
- [18] R. B. Rusu and S. Cousins, “3D is here: Point Cloud Library (PCL),” in *IEEE International Conference on Robotics and Automation (ICRA)*, Shanghai, China, May 9-13 2011.

- [19] K. Briggs, A. Lyons, E. Pouliquen, L. Mayer, and M. Richardson, “Seafloor roughness, sediment grain size, and temporal stability,” in *Underwater Acoustic Measurements: Technologies & Results*, 2005. [Online]. Available: <http://www.dtic.mil/dtic/tr/fulltext/u2/a455132.pdf>
- [20] J. T. Anderson, D. Van Holliday, R. Kloser, D. G. Reid, and Y. Simard, “Acoustic seabed classification: current practice and future directions,” *ICES Journal of Marine Science: Journal du Conseil*, vol. 65, no. 6, pp. 1004–1011, 2008. [Online]. Available: <http://icesjms.oxfordjournals.org/content/65/6/1004.abstract>
- [21] R. D. Hobson, “Surface roughness in topography: a quantitative approach,” in *Spatial Analysis in Geomorphology*. Har, 1972, pp. 221–245.
- [22] A. Friedman, O. Pizarro, S. B. Williams, and M. Johnson-Roberson, “Multi-scale measures of rugosity, slope and aspect from benthic stereo image reconstructions,” *PLoS ONE*, vol. 7, no. 12, p. e50440, 12 2012. [Online]. Available: <http://dx.doi.org/10.1371/journal.pone.0050440>

BIBLIOGRAPHY

- “Fusion6g lbl web page.” Accessed 27 Oct 2015. 2015. [Online]. Available: <http://www.sonardyne.com/products/positioning/fusion-6g.html>
- Amaral, B., “The hydrodynamic optimization of a cable-flying vehicle,” Master’s thesis, University of Rhode Island, 2012.
- Anderson, J. T., Van Holliday, D., Kloser, R., Reid, D. G., and Simard, Y., “Acoustic seabed classification: current practice and future directions,” *ICES Journal of Marine Science: Journal du Conseil*, vol. 65, no. 6, pp. 1004–1011, 2008. [Online]. Available: <http://icesjms.oxfordjournals.org/content/65/6/1004.abstract>
- Anonsen, K. and Hagen, O., “Terrain aided underwater navigation using pockmarks,” in *OCEANS 2009, MTS/IEEE Biloxi - Marine Technology for Our Future: Global and Local Challenges*, Oct 2009, pp. 1–6.
- Bailey, T. and Durrant-Whyte, H., “Simultaneous localization and mapping: part ii,” *IEEE Robotics & Automation Magazine*, vol. 13, no. 3, pp. 108–117, 2006.
- Barkby, S., Williams, S. B., Pizarro, O., and Jakuba, M. V., “A featureless approach to efficient bathymetric slam using distributed particle mapping,” *Journal of Field Robotics*, vol. 28, no. 1, pp. 19–39, 2011. [Online]. Available: <http://dx.doi.org/10.1002/rob.20382>
- Barkby, S., Williams, S. B., Pizarro, O., and Jakuba, M. V., “Bathymetric particle filter slam using trajectory maps.” *I. J. Robotic Res.*, vol. 31, no. 12, pp. 1409–1430, 2012. [Online]. Available: <http://dblp.uni-trier.de/db/journals/ijrr/ijrr31.html#BarkbyWPJ12>
- Bowditch, N., *The American Practical Navigator*. National Imagery and Mapping Agency, 2002.
- Brennan, M. L., “The disarticulation of ancient shipwreck sites by mobile fishing gear: A case study from the southeast aegean sea,” *INA Quarterly*, vol. 36(4), pp. 6–7, 2010.
- Brennan, M. L., Davis, D., Ballard, R. D., Trembanis, A. C., Vaughn, J. I., Krumholz, J. S., Delgado, J. P., Roman, C. N., Smart, C., Bell, K. L. C., Duman, M., and DuVal, C., “Quantification of bottom trawl damage to ancient shipwreck sites,” *Marine Geology*, 2015, in Revision.

- Bridge, T., Done, T., Beaman, R., Friedman, A., Williams, S., Pizarro, O., and Webster, J., “Topography, substratum and benthic macrofaunal relationships on a tropical mesophotic shelf margin, central great barrier reef, australia,” *Coral Reefs*, vol. 30, pp. 143–153, 2011, 10.1007/s00338-010-0677-3. [Online]. Available: <http://dx.doi.org/10.1007/s00338-010-0677-3>
- Briggs, K., Lyons, A., Pouliquen, E., Mayer, L., and Richardson, M., “Seafloor roughness, sediment grain size, and temporal stability,” in *Underwater Acoustic Measurements: Technologies & Results*, 2005. [Online]. Available: <http://www.dtic.mil/dtic/tr/fulltext/u2/a455132.pdf>
- Brokloff, N., “Matrix algorithm for doppler sonar navigation,” in *OCEANS ’94. ’Oceans Engineering for Today’s Technology and Tomorrow’s Preservation.’ Proceedings*, vol. 3, Sep 1994, pp. III/378–III/383 vol.3.
- Byrne, J. S. and Schmidt, V. E., “Uncertainty modeling for auv acquired bathymetry,” *U.S. Hydrographic Conference (US HYDRO)*, 16-19 April 2015 2015. [Online]. Available: http://ccom.unh.edu/sites/default/files/publications/AUV_uncertainty_model_final.pdf
- Caress, D., *mbnavadjust(1) manpage*, version 5 ed., June 2013.
- Carlone, L., Kira, Z., Beall, C., Indelman, V., and Dellaert, F., “Eliminating conditionally independent sets in factor graphs: A unifying perspective based on smart factors,” in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, May 2014, pp. 4290–4297.
- Carrillo, H., Reid, I., and Castellanos, J., “On the comparison of uncertainty criteria for active slam,” in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, May 2012, pp. 2080–2087.
- Coleman, D. F., Ballard, R. D., Newman, J. B., Gregory, T. S., Howland, J., Singh, H., Roman, C., Pizarro, O., Foley, B., Eustice, R., Can, A., Webster, S., Piechota, D., Giangrande, C., Ward, C., Horlings, R., McBride, K., and Durbin, M. J., *Archaeological Oceanography*, Ballard, R. D., Ed. Princeton University Press, 2008.
- Coleman, D., Ballard, R., and Gregory, T., “Marine archaeological exploration of the black sea,” in *OCEANS 2003. Proceedings*, vol. 3, Sept. 2003, pp. 1287–1291.
- Coleman, D., Newman, J., and Ballard, R., “Design and implementation of advanced underwater imaging systems for deep sea marine archaeological surveys,” *OCEANS 2000 MTS/IEEE Conference and Exhibition*, vol. 1, pp. 661–665, 2000.
- Crawford, W. P., *Mariner’s Weather*. Norton Nautical Books, 1978.

- Crosbie, B., Anderson, M., Wolbrecht, E., Canning, J., and Edwards, D., “Synchronous navigation of auvs using whoi micro-modem 13-bit communications,” in *OCEANS 2010*, Sept 2010, pp. 1–5.
- Dellaert, F., “Square Root SAM: Simultaneous location and mapping via square root information smoothing,” in *Robotics: Science and Systems (RSS)*, 2005.
- Dellaert, F. and Kaess, M., “Square Root SAM: Simultaneous localization and mapping via square root information smoothing,” *Intl. J. of Robotics Research (IJRR)*, vol. 25, no. 12, pp. 1181–1204, Dec 2006.
- Dellaert, F., “Factor graphs and gtsam: a hands-on introduction,” Georgia Tech, Tech. Rep. GT-RIM-CP&R-2012-002, 2012, retrieved 25 Nov 2015. [Online]. Available: <https://research.cc.gatech.edu/borg/sites/edu.borg/files/downloads/gtsam.pdf>
- Durrant-Whyte, H. and Bailey, T., “Simultaneous localization and mapping: part i,” *IEEE Robotics & Automation Magazine*, vol. 13, no. 2, pp. 99–110, 2006.
- Egbert, G. D. and Erofeeva, S. Y., “Efficient inverse modeling of barotropic ocean tides,” *J. Atmos. Oceanic Technol.*, vol. 19, no. 2, pp. 183–204, Feb. 2002. [Online]. Available: [http://dx.doi.org/10.1175/1520-0426\(2002\)019\(0183:EIMOBO\)2.0.CO;2](http://dx.doi.org/10.1175/1520-0426(2002)019(0183:EIMOBO)2.0.CO;2)
- Egbert, G. D., Erofeeva, S. Y., and Ray, R. D., “Assimilation of altimetry data for nonlinear shallow-water tides: Quarter-diurnal tides of the northwest european shelf,” *Continental Shelf Research*, vol. 30, no. 6, pp. 668 – 679, 2010, tides in Marginal Seas - A special issue in memory of Prof Alexei Nekrasov. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0278434309003094>
- Enns, R. and Morrell, D., “Terrain-aided navigation using the viterbi algorithm,” *Journal of Guidance, Control, and Dynamics*, vol. 18, no. 6, pp. 1444–1449, Nov. 1995. [Online]. Available: <http://dx.doi.org/10.2514/3.21566>
- Eustice, R., Whitcomb, L., Singh, H., and Grund, M., “Experimental results in synchronous-clock one-way-travel-time acoustic navigation for autonomous underwater vehicles,” in *Robotics and Automation, 2007 IEEE International Conference on*, April 2007, pp. 4257–4264.
- Eustice, R., “Large-area visually augmented navigation for autonomous underwater vehicles,” Ph.D. dissertation, MIT/WHOI Joint Program, 2005.
- Eustice, R., Pizarro, O., Singh, H., and Ma, W. H., “Visually augmented navigation in an unstructured environment using a delayed state history,” in *in Proc. IEEE Int. Conf. on Robotics and Automation*, 2004, pp. 25–32.

- Eustice, R., Singh, H., Hole, W., and Hole, W., “Visually navigating the rms titanic with slam information filters,” in *in Proceedings of Robotics: Science and Systems*, 2005, pp. 57–64.
- Eustice, R. M., Pizarro, O., and Singh, H., “Visually augmented navigation for autonomous underwater vehicles,” *IEEE Journal of Oceanic Engineering*, vol. 33, no. 2, pp. 103–122, 2008.
- Eustice, R. M., Singh, H., and Whitcomb, L. L., “Synchronous-clock, one-way-travel-time acoustic navigation for underwater vehicles,” *Journal of Field Robotics*, vol. 28, no. 1, pp. 121–136, 2011. [Online]. Available: <http://dx.doi.org/10.1002/rob.20365>
- Feseker, T., Boetius, A., Wenzhofer, F., Blandin, J., Olu, K., Yoerger, D. R., Camilli, R., German, C. R., and de Beer, D., “Eruption of a deep-sea mud volcano triggers rapid sediment movement,” *Nat Commun*, vol. 5, pp. –, Nov. 2014. [Online]. Available: <http://dx.doi.org/10.1038/ncomms6385>
- Fofonoff, N. P. and Millard, R. C., “Algorithms for computation of fundamental properties of seawater,” UNESCO, Tech. Rep., 1983. [Online]. Available: http://www.jodc.go.jp/jodcweb/info/ioc_doc/UNESCO_tech/059832eb.pdf
- Friedman, A., Pizarro, O., Williams, S. B., and Johnson-Roberson, M., “Multi-scale measures of rugosity, slope and aspect from benthic stereo image reconstructions,” *PLoS ONE*, vol. 7, no. 12, p. e50440, 12 2012. [Online]. Available: <http://dx.doi.org/10.1371/journal.pone.0050440>
- Geng, Y., Martins, R., and Sousa, J., “Accuracy analysis of dvl / imu / magnetometer integrated navigation system using different imus in auv,” in *Control and Automation (ICCA), 2010 8th IEEE International Conference on*, June 2010, pp. 516–521.
- German, C. R., Bowen, A., Coleman, M. L., Honig, D. L., Huber, J. A., Jakuba, M. V., Kinsey, J. C., Kurz, M. D., Leroy, S., McDermott, J. M., de Lpinay, B. M., Nakamura, K., Seewald, J. S., Smith, J. L., Sylva, S. P., Van Dover, C. L., Whitcomb, L. L., and Yoerger, D. R., “Diverse styles of submarine venting on the ultraslow spreading mid-cayman rise,” *Proceedings of the National Academy of Sciences*, vol. 107, no. 32, pp. 14 020–14 025, 2010. [Online]. Available: <http://www.pnas.org/content/107/32/14020.abstract>
- Grewal, M. S. and Andrews, A. P., *Kalman Filtering: Theory and Practice Using Matlab*, 2nd ed. Wiley & Sons, 2001.
- Grilli, S., OReilly, C., Harris, J., Bakhsh, T., Tehranirad, B., Banihashemi, S., Kirby, J., Baxter, C., Eggeling, T., Ma, G., and Shi, F., “Modeling of smf tsunami hazard along the upper us east coast: detailed impact around ocean

- city, md,” *Natural Hazards*, vol. 76, no. 2, pp. 705–746, 2015. [Online]. Available: <http://dx.doi.org/10.1007/s11069-014-1522-8>
- Guarinello, M., Shumchenia, E., and King, J., “Marine habitat classification for ecosystem-based management: A proposed hierarchical framework,” *Environmental Management*, vol. 45, pp. 793–806, 2010, 10.1007/s00267-010-9430-5. [Online]. Available: <http://dx.doi.org/10.1007/s00267-010-9430-5>
- Han, S.-C. and Simons, F. J., “Spatiospectral localization of global geopotential fields from the gravity recovery and climate experiment (grace) reveals the coseismic gravity change owing to the 2004 sumatra-andaman earthquake,” *Journal of Geophysical Research: Solid Earth (1978–2012)*, vol. 113, no. B1, 2008.
- Hobson, R. D., “Surface roughness in topography: a quantitative approach,” in *Spatial Analysis in Geomorphology*. Har, 1972, pp. 221–245.
- Horrell, C. E. and Borgens, A. A., “Archaeological investigations of an early nineteenth-century shipwreck in the gulf of mexico,” in *Offshore Technology Conference*, 2014.
- Indelman, V., Williams, S., Kaess, M., and Dellaert, F., “Factor graph based incremental smoothing in inertial navigation systems,” in *Information Fusion (FUSION), 2012 15th International Conference on*, July 2012, pp. 2154–2161.
- Inglis, G., “Hybrid optical acoustic seafloor mapping,” Ph.D. dissertation, University of Rhode Island, 2013.
- Inglis, G., Smart, C., Vaughn, I., and Roman, C., “A pipeline for structured light bathymetric mapping,” in *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, Oct 2012, pp. 4425–4432.
- IOC, SCOR, and IAPSO, “The international thermodynamic equation of seawater 2010: Calculation and use of thermodynamic properties.” UNESCO” Intergovernmental Oceanographic Commission, Manuals and Guides No. 56, 2010. [Online]. Available: http://www.teos-10.org/pubs/TEOS-10_Manual.pdf
- IxBlue. “Ramses product brochure.” Web Page. Retrieved 26 Oct 2015. [Online]. Available: <https://www.ixblue.com/sites/default/files/downloads/ixblue-ps-ramses-02-2014-web.pdf>
- Jaffe, J., Moore, K., McLean, J., and Strand, M., “Underwater imaging: Status and prospects,” *Oceanography*, vol. 14, no. 3, pp. 66–76, 2001.
- Jalving, B., Gade, K., Svartveit, K., Willumsen, A., and Sorhagen, R., “Dvl velocity aiding in the hugin 1000 integrated inertial navigation system,” *Modeling Identification and Control*, vol. 25, no. 4, pp. 223–236, 2004.

- Jensen, F., Kuperman, W., Porter, M., and Schmidt, H., *Computational Ocean Acoustics*. Springer, 2011.
- Jolliffe, I. T., *Principal Component Analysis*. Springer-Verlag, 1986.
- Kaess, M. and Dellaert, F., “Covariance recovery from a square root information matrix for data association,” *Journal of Robotics and Autonomous Systems (RAS)*, vol. 57, pp. 1198–1210, Dec 2009.
- Kaess, M., Ila, V., Roberts, R., and Dellaert, F., “The Bayes tree: Enabling incremental reordering and fluid relinearization for online mapping,” Computer Science and Artificial Intelligence Laboratory, MIT, Tech. Rep. MIT-CSAIL-TR-2010-021, Jan 2010.
- Kaess, M., Ranganathan, A., and Dellaert, F., “iSAM: Incremental smoothing and mapping,” *IEEE Trans. on Robotics (TRO)*, vol. 24, no. 6, pp. 1365–1378, Dec 2008.
- Kay, S., *Fundamentals of Statistical Signal Processing, Estimation Theory*. Prentice Hall PTR, 1993, vol. 1.
- Kay, S. M., *Fundamentals of Statistical Signal Processing, Detection Theory*. Prentice Hall PTR, 1998, vol. 2.
- Kinsey, J. C., Eustice, R. M., and Whitcomb, L. L., “A survey of underwater vehicle navigation: Recent advances and new challenges,” in *IFAC Conference of Manoeuvring and Control of Marine Craft*, Lisbon, Portugal, September 2006, invited paper.
- Kinsey, J. C. and Whitcomb, L. L., “Preliminary field experience with the dvl-nav integrated navigation system for oceanographic submersibles,” *Control Engineering Practice*, vol. 12, no. 12, pp. 1541 – 1549, 2004.
- Kinsler, Frey, Coppens, and Sanders, *Fundamentals of Acoustics*, 4th ed. John Wiley & Sons, 1999.
- Kollanda, R. C., “Terrain relative navigation using bathymetric slope profile with a rao-blackwellized particle filter for long range auvs,” Master’s thesis, University of Rhode Island, 2012.
- Koller, D. and Friedman, N., *Probabilistic Graphical Models*. MIT, 2009.
- Kunz, C. and Singh, H., “Map building fusing acoustic and visual information using autonomous underwater vehicles,” *Journal of Field Robotics*, vol. 30, no. 5, pp. 763–783, 2013. [Online]. Available: <http://dx.doi.org/10.1002/rob.21473>

- Kunz, C. G., “Autonomous underwater vehicle navigation and mapping in dynamic, unstructured environments,” Ph.D. dissertation, Massachusetts Institute of Technology and Woods Hole Oceanographic Institution Joint Program, 2011.
- L3 Communications SeaBeam Instruments. “Multibeam sonar theory of operation.” 2000. [Online]. Available: <http://www.ldeo.columbia.edu/res/pi/MB-System/sonarfunction/SeaBeamMultibeamTheoryOperation.pdf>
- Lanser, S. and Lengauer, T., “On the selection of candidates for point and line correspondences,” in *In International Symposium on Computer Vision*. IEEE Computer Society Press, 1995, pp. 157–162.
- Larsen, M., “Synthetic long baseline navigation of underwater vehicles,” in *OCEANS 2000 MTS/IEEE Conference and Exhibition*, vol. 3, 2000, pp. 2043–2050 vol.3.
- Leonard, J. J. and Durrant-Whyte, H. F., “Simultaneous map building and localization for an autonomous mobile robot,” in *Intelligent Robots and Systems’91. Intelligence for Mechanical Systems, Proceedings IROS’91. IEEE/RSJ International Workshop on*. Ieee, 1991, pp. 1442–1447.
- Leung, C., Huang, S., and Dissanayake, G., “Active slam using model predictive control and attractor based exploration,” in *Intelligent Robots and Systems, 2006 IEEE/RSJ International Conference on*, Oct 2006, pp. 5026–5031.
- Logan, L., “Performance analysis of an underwater wire flying profiling vehicle,” Master’s thesis, University of Rhode Island, 2014.
- Mahon, I., Williams, S., Pizarro, O., and Johnson-Roberson, M., “Efficient view-based slam using visual loop closures,” *Robotics, IEEE Transactions on*, vol. 24, no. 5, pp. 1002–1014, Oct. 2008.
- Mcdonald, J., Kaess, M., Cadena, C., Neira, J., and Leonard, J. J., “Real-time 6-dof multi-session visual slam over large-scale environments,” *Robot. Auton. Syst.*, vol. 61, no. 10, pp. 1144–1158, Oct. 2013. [Online]. Available: <http://dx.doi.org/10.1016/j.robot.2012.08.008>
- McDougall, T. J. and Barker, P. M., *Getting started with TEOS-10 and the Gibbs Seawater (GSW) Oceanographic Toolbox*. SCOR/IAPSO WG127, 2011. [Online]. Available: <http://www.teos-10.org/software.htm>
- Moore, J. H., *The New Practical Navigator*, 2nd ed. Edmund M. Blunt, 1800.
- National Oceanic and Atmospheric Administration. “Extended continental shelf project.” Webpage. 2012. [Online]. Available: <http://continentalshelf.gov/>

- Negahdaripour, S. and Madjidi, H., “Stereovision imaging on submersible platforms for 3-d mapping of benthic habitats and sea-floor structures,” *Oceanic Engineering, IEEE Journal of*, vol. 28, no. 4, pp. 625–650, Oct. 2003.
- Nishimura, C. and Forsyth, D., “Improvements in navigation using seabeam crossing errors,” *Marine Geophysical Researches*, vol. 9, no. 4, pp. 333–352, 1988. [Online]. Available: <http://dx.doi.org/10.1007/BF00315004>
- Nomikou, P., Parks, M., Papanikolaou, D., Pyle, D., Mather, T., Carey, S., Watts, A., Paulatto, M., Kalnins, M., Livanos, I., Bejelou, K., Simou, E., and Perros, I., “The emergence and growth of a submarine volcano: The kameni islands, santorini (greece),” *GeoResJ*, vol. 12, no. 0, pp. 8 – 18, 2014. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S2214242814000047>
- Field Procedures Manual*, Office of Coast Survey Std., 2014. [Online]. Available: http://www.nauticalcharts.noaa.gov/hsd/fpm/2014_FPM_Final.pdf
- Olson, E., Leonard, J., and Teller, S., “Robust range-only beacon localization,” *Oceanic Engineering, IEEE Journal of*, vol. 31, no. 4, pp. 949–958, Oct 2006.
- Pizarro, O., Eustice, R., and Singh, H., “Large area 3-d reconstructions from underwater optical surveys,” *Oceanic Engineering, IEEE Journal of*, vol. 34, no. 2, pp. 150 –169, Apr. 2009. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4918977
- Pizarro, O., Williams, S., and Colquhoun, J., “Topic-based habitat classification using visual data,” in *OCEANS 2009 - EUROPE*, May 2009, pp. 1 –8.
- Roman, C. and Singh, H., “Consistency based error evaluation for deep sea bathymetric mapping with robotic vehicles,” in *Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on*, May 2006, pp. 3568 –3574.
- Roman, C. and Singh, H., “A self-consistent bathymetric mapping algorithm,” *Journal of Field Robotics*, vol. 24, no. 1-2, pp. 23–50, 2007. [Online]. Available: <http://dx.doi.org/10.1002/rob.20164>
- Rusinkiewicz, S. and Levoy, M., “Efficient variants of the icp algorithm,” in *3-D Digital Imaging and Modeling, 2001. Proceedings. Third International Conference on*, 2001, pp. 145 –152.
- Rusu, R. B., “Semantic 3d object maps for everyday manipulation in human living environments,” Ph.D. dissertation, Computer Science department, Technische Universitaet Muenchen, Germany, October 2009.

- Rusu, R. B. and Cousins, S., “3D is here: Point Cloud Library (PCL),” in *IEEE International Conference on Robotics and Automation (ICRA)*, Shanghai, China, May 9-13 2011.
- Saunders, P. M. and Fofonoff, N. P., “Conversion of pressure to depth in the ocean,” *Deep Sea Research and Oceanographic Abstracts*, vol. 23, pp. 109–111, 1976.
- Schechner, Y. and Karpel, N., “Clear underwater vision,” in *Computer Vision and Pattern Recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE Computer Society Conference on*, vol. 1, 2004, pp. 536–543.
- SBE 911plus CTD Sensor*, Sea-Bird Electronics, 13431 NE 20th Street Bellevue, Washington 98005 USA, 2014, accessed 9 Nov 2015. [Online]. Available: http://www.seabird.com/sites/default/files/documents/911plusbrochureDec14_0.pdf
- SBE 49 FastCAT CTD Sensor Manual*, Sea-Bird Electronics, 13431 NE 20th Street Bellevue, Washington 98005 USA, 2015, accessed 9 Nov 2015. [Online]. Available: http://www.seabird.com/sites/default/files/documents/49_017.pdf
- Shanmugan, K. S. and Breipohl, A. M., *Random Signals*. John Wiley & Sons, 1988.
- Singh, H., Roman, C., Pizarro, O., Eustice, R., and Can, A., “Towards high-resolution imaging from underwater vehicles,” *The International journal of robotics research*, vol. 26, no. 1, pp. 55–74, 2007.
- Smith, R., Self, M., and Cheeseman, P., “Estimating uncertain spatial relationships in robotics,” in *Robotics and Automation. Proceedings. 1987 IEEE International Conference on*, vol. 4, Mar 1987, pp. 850–850.
- Smith, R. C. and Cheeseman, P., “On the representation and estimation of spatial uncertainty,” *The International Journal of Robotics Research*, vol. 5, no. 4, pp. 56–68, 1986. [Online]. Available: <http://ijr.sagepub.com/content/5/4/56.abstract>
- Thrun, S., Burgard, W., and Fox, D., *Probabilistic Robotics*. MIT Press, 2006.
- Thrun, S., Liu, Y., Koller, D., Ng, A. Y., Ghahramani, Z., and Durrant-Whyte, H., “Simultaneous localization and mapping with sparse extended information filters,” *The International Journal of Robotics Research*, vol. 23, no. 7-8, pp. 693–716, 2004. [Online]. Available: <http://ijr.sagepub.com/content/23/7-8/693.abstract>
- Troni, G. and Whitcomb, L. L., “Advances in in situ alignment calibration of doppler and high/low-end attitude sensors for underwater vehicle navigation:

- Theory and experimental evaluation,” *Journal of Field Robotics*, vol. 32, no. 5, pp. 655–674, 2015. [Online]. Available: <http://dx.doi.org/10.1002/rob.21551>
- VanMiddlesworth, M. A., “Toward autonomous underwater mapping in partially structured 3d environments,” Master’s thesis, Massachusetts Institute of Technology, 2014. [Online]. Available: <hdl.handle.net/1721.1/87791>
- Vaughn, I., “Microbathymetry using self-contained navigation and simultaneous localization and mapping,” Master’s thesis, University of Rhode Island, 2012.
- Webster, S., Eustice, R., Singh, H., and Whitcomb, L., “Preliminary deep water results in single-beacon one-way-travel-time acoustic navigation for underwater vehicles,” in *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, Oct 2009, pp. 2053–2060.
- Whitcomb, L., Yoerger, D., and Singh, H., “Advances in doppler-based navigation of underwater robotic vehicles,” in *Robotics and Automation, 1999. Proceedings. 1999 IEEE International Conference on*, vol. 1, 1999, pp. 399–406 vol.1.
- Williams, S., Pizarro, O., Jakuba, M., Mahon, I., Ling, S., and Johnson, C., “Repeated auv surveying of urchin barrens in north eastern tasmania,” in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, May 2010, pp. 293–299. [Online]. Available: <http://www.cas.edu.au/content.php/237.html?publicationid=659&displaypage=3>
- Williams, S., Dissanayake, G., and Durrant-Whyte, H., “Towards terrain-aided navigation for underwater robotics,” *Advanced Robotics*, vol. 15, no. 5, pp. 533–549, 2001. [Online]. Available: <http://dx.doi.org/10.1163/156855301317033559>
- Yang, L. and Tact, T., “Multibeam sonar bottom detection using multiple subarrays,” in *OCEANS ’97. MTS/IEEE Conference Proceedings*, vol. 2, Oct 1997, pp. 932–938 vol.2.
- Zheng, H., Zheng, B., Ji, G., Guo, Z., and Sun, Y., “Image enhancement of underwater target detection by inhomogeneous illumination,” in *OCEANS, 2012 - Yeosu*, May 2012, pp. 1–4.