

Self Consistent Bathymetric Mapping from Robotic Vehicles
in the Deep Ocean

by

Christopher N. Roman

B.S., Virginia Polytechnic Institute and State University (1997)
S.M., University of California San Diego (1999)

Submitted to the Joint Program in Applied Ocean Science & Engineering
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

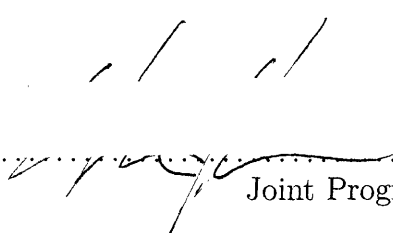
MASSACHUSETTS INSTITUTE OF TECHNOLOGY


and

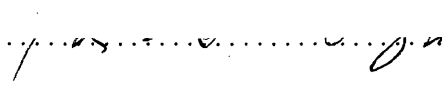
WOODS HOLE OCEANOGRAPHIC INSTITUTION


June 2005

© Woods Hole Oceanographic Institution 2005. All rights reserved.

Author 
Joint Program in Applied Ocean Science & Engineering
May 12, 2005

Certified by 
Hanumant Singh
Associate Scientist, WHOI
Thesis Supervisor

Accepted by 
Mark Grosenbaugh
Chairman, Joint Committee for Applied Ocean Science & Engineering



ARCHIVES

Self Consistent Bathymetric Mapping from Robotic Vehicles in the Deep Ocean

by
Christopher N. Roman

Submitted to the Joint Program in Applied Ocean Science & Engineering
on May 5, 2005, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Obtaining accurate and repeatable navigation for robotic vehicles in the deep ocean is difficult and consequently a limiting factor when constructing vehicle-based bathymetric maps. This thesis presents a methodology to produce self-consistent maps and simultaneously improve vehicle position estimation by exploiting accurate local navigation and utilizing terrain relative measurements.

It is common for errors in the vehicle position estimate to far exceed the errors associated with the acoustic range sensor. This disparity creates inconsistency when an area is imaged multiple times and causes artifacts that distort map integrity. Our technique utilizes small terrain “sub-maps” that can be pairwise registered and used to additionally constrain the vehicle position estimates in accordance with actual bottom topography. A delayed state Kalman filter is used to incorporate these sub-map registrations as relative position measurements between previously visited vehicle locations. The archiving of previous positions in a filter state vector allows for continual adjustment of the sub-map locations. The terrain registration is accomplished using a two dimensional correlation and a six degree of freedom point cloud alignment method tailored for bathymetric data. The complete bathymetric map is then created from the union of all sub-maps that have been aligned in a consistent manner. Experimental results from the fully automated processing of a multibeam survey over the TAG hydrothermal structure at the Mid-Atlantic ridge are presented to validate the proposed method.

Thesis Supervisor: Hanumant Singh
Title: Associate Scientist, WHOI

Acknowledgments

I've been fortunate to spend the last several years surrounded by a unique set of people from whom I've learned a great deal.

Hanu Singh, my advisor, I thank most for the opportunities and freedom he's provided. Although the word efficient wouldn't describe my graduate career, I'm fortunate to have seen all sides of getting real data from the ocean. Hanu has exuded a motivating and seemingly unsustainable level of excitement for this work since the day I met him. For my committee members, I'm grateful to Dana Yoerger for his suggestions, his motivation to work on problems that matter, and the entertainment provided by his willingness to expound on the dubious intentions of *the man*. John Leonard has always been supportive and provided sound advice over the past several years, more of which I should have listened to. I'm also grateful to Rob Sohn and Susan Humphris for the opportunity to tackle the TAG mound bathymetry. It was a great experience, thanks.

My office mates Oscar, Ryan, and Mike have been through it all. Oscar has been a great companion inside and outside of work. He will appear in any story I ever tell involving submarines, beer, and roasted pigs. From day one he's been a part of enjoying this place and the most usual suspect. You won't find a harder working or more competent bubba than Ryan. Ryan has been a great help in the office, at sea and a good friend otherwise. Oscar, Ryan, and I share the *honor* of being the original SEAbed trio. I think we all learned a lot more than just how, or how not, to build an AUV. Mike has provided consistently helpful consultation on working matters and is in all respects an ideal office mate. He also deserves acknowledgment as a gifted and kindred spirit in the fine art of backyard and hardware store aisle mechanical design.

The United States Deep Submergence Lab and (former) Trailer Park has been an interesting place to work. I'm glad Ann Stone always has a smile and the solution to my latest paperwork conundrum. John Howland is always interested and up for a chat. Working, sailing and diving with Matt Heintz has been fun. Andy Bowen, Cathy Offinger, Steve Gegg, Will Sellers, Chris Taylor, Jimmy Varnum, Pete Collins, and Tom Crook all get JASON to do what it does, and made the TAG cruise a good time and a thesizable experience. Fellow deep submergence junkies Brian, Brendan, and Rich have always brought interesting ideas, new perspectives and good beer to the table.

Down in the village Carlos, Esmail, John and Terry could always figure out how to build it in time. Neil McPhee has been a great help in getting SEAbed out the door. Terry Hammar always has something you can borrow.

Julia, Marsha and the WHOI education office couldn't do it better. They give a lot to this program and have great tolerance for those of us still unable to grasp the concept of registration forms.

Time spent living at 24 Millfield street with my John, Jim, Dirk, Charlie, Oscar, Mike, Mark and Akiko will always provide material for an entertaining or bizarre story. To borrow a line, I will remember you all as people that I've met.

And there are those that have made Woods Hole home. Joe and Doc Reddy are always good for a laugh. Rachel's always got a smile, and a sweet green bike. Fabian, dude. Jason says it like it is and always has a good time. Claudia will always raise the bar when it comes to cooking. Dave and Kate have yet to realize we're just using them for Max. Lara

and Henrik know how to throw a party. Chances are you'll run into Nick and Amy at the beach or on the water. Tim and Liz always have the tunes and an adventure.

I'm also thankful to the greater Falmouth windsurfing crew with whom I've shared the water. There is something great about knowing you'll always find a friend (or fool) at the beach, even on the most miserably windy and cold days.

I can't thank my parents, Stephen and Diane, enough. They've given more that I ever could have asked for.

A finally to my wife Stephanie, thanks for believing in me more than I believe in myself, and being the best part of these days and the days to come.

Thanks

Chris

This work was funded by the CenSSIS ERC of the Nation Science Foundation under grant EEC-9986821 and in part by the Woods Hole Oceanographic Institution through a grant from the Penzance Foundation.

Contents

1	Introduction	17
1.1	Motivation	17
1.1.1	Problem statement	18
1.2	Related research	18
1.2.1	Underwater Navigation	19
1.2.2	Sonar mapping	21
1.2.3	Simultaneous Localization and Mapping	24
1.2.4	Registration methods	26
1.3	Thesis breakdown	27
1.3.1	Outline of methods	27
1.3.2	Assumptions and restrictions	28
1.3.3	Contributions	29
1.3.4	Thesis structure	29
2	Acoustics for mapping	31
2.1	Introduction	31
2.2	Range determination	31
2.3	Outlier rejection and ping clean up	33
2.3.1	Inner ping	33
2.3.2	Over multiple pings	35
2.4	Summary	36
3	Sub-mapping SLAM bathymetry	39
3.1	Introduction	39
3.1.1	State vector and coordinate frames	39
3.2	Vehicle model	41
3.2.1	Navigation sensor measurements	42
3.3	Vehicle navigation	43
3.4	Sub-map creation	44
3.4.1	Dynamic map sizing	45
3.5	Relative pose measurements	52
3.6	Summary	53

4	Terrain registration	55
4.1	Introduction	55
4.2	Relative position measurements	55
4.2.1	Selection of methods	55
4.2.2	Sub-mapping specifics	56
4.3	Methods	57
4.3.1	Correlation	57
4.3.2	Iterative closest point matching	59
4.4	Measurement evaluation	67
4.4.1	Surface error	67
4.4.2	Principal component analysis (PCA)	68
4.4.3	Point based errors	69
4.5	Summary	71
5	Experimental results and validation	75
5.1	Introduction	75
5.2	Survey description	75
5.3	Complete maps	79
5.3.1	Composite surface errors	80
5.3.2	Terrain consistency checking	82
5.3.3	Preliminary maps	83
5.4	Sub-mapping results	85
5.5	Robustness to common errors	90
5.6	Additional map refinements	96
5.7	Summary	97
6	Conclusions	99
6.1	Introduction	99
6.2	Summary	99
6.3	Limitations & Future Work	100
6.3.1	Ground truth	100
6.3.2	Navigation	100
6.3.3	Terrain registration	101
6.3.4	Acoustic modeling	101
A	Relative pose transformations	103
A.1	Basic definitions	103
A.2	Additional relations	104
A.2.1	Head-to-tail	104
A.2.2	Inverse	105
A.2.3	Tail-to-tail	105
B	3D Point set matching	107
B.1	Surface gridding	107
B.2	PCA surface normal estimation	107
B.3	Point sampling methods	109

C Sonar sensor offset refinement	111
D TAG survey 2	115

List of Figures

1-1	Relative contributions of mapping system errors	19
1-2	Acoustically created terrain map over an archaeological site	24
1-3	Sub-mapping conceptual sketch	27
2-1	Sample SM2000 sonar image	32
2-2	Dependence of returned pulse duration on incidence angle	34
2-3	Block diagram of sonar processing	34
2-4	Single ping range detection	35
2-5	Ping images showing outlier rejection	36
3-1	Flow diagram for the delayed state EKF	40
3-2	Vehicle coordinate system	41
3-3	Uncertainty of the sonar ping placement	50
3-4	Figures showing the changing map statistics	51
3-5	Vector diagram of sub-map coordinate frames	52
4-1	Sample sub-map regions	57
4-2	Correlation matching example	59
4-3	Point-to-point convergence	61
4-4	Normal space point sampling	64
4-5	Comparison of iterative closest point (ICP) convergence behavior	65
4-6	Confidence ellipse for the ICP registration covariance	66
4-7	Local minima for the ICP registration	67
4-8	Properties of the surface error metric	69
4-9	Sample sub-map regions	70
4-10	Example registration sequence	72
4-11	Example registration sequence, con't	73
5-1	TAG mount survey tracklines	77
5-2	The JASON ROV	77
5-3	TAG mount survey 1 tracklines in detail	78
5-4	Sonar sounding density and example sub-maps	79
5-5	Map-to-map distance measure	81
5-6	Incremental addition of sub-maps and a sub-map measurement	82
5-7	Overlapping sub-maps	84
5-8	Results for LBL filtering	85

5-9	Sub-mapping pose network	86
5-10	Map-to-map surface error for the sub-mapped terrain	87
5-11	Comparison of terrain maps	88
5-12	Zoom into an Ocean Drilling Program (ODP) re-entry cone	89
5-13	Comparison slices through complete maps	89
5-14	Results for a sparse version of survey 1	91
5-15	The effect of sensor offset error and heading bias	93
5-16	Appearance of terrain mapped with Doppler velocity log (DVL) bias	94
5-17	Results for simulated heading dependent bias	95
5-18	Surface errors before and after pose optimization	98
A-1	Coordinate frame sketch	103
B-1	Gridding sketch and Gaussian kernel	108
C-1	Short section of trackline used to determine the sonar sensor offsets	112
C-2	Map-to-map error for two different roll offsets	112
C-3	Reduction of error for the sonar sensor angular offsets	113
D-1	TAG mount survey 2 tracklines	115
D-2	Survey 2 dead reckoning (DR) surface error and map overlap	116
D-3	Survey 2 surface error after sub-mapping	117
D-4	Final terrain for survey 2	118

List of Tables

5.1	Summary of survey details	76
5.2	Navigation sensors	76

List of Acronyms

ADCP	acoustic doppler current profiler
AUV	autonomous underwater vehicle
DOF	degree of freedom
DR	dead reckoning
DVL	Doppler velocity log
EKF	extended Kalman filter
EXACT	A 300 kHz long base line acoustic navigation system
FOG	fiber optic gyroscope
GPS	Global Positioning System
ICP	iterative closest point
IEKF	iterated extended Kalman filter
IMU	inertial measurement unit
INS	inertial navigation sensor
LBL	long baseline
ML	maximum likelihood
ODP	Ocean Drilling Program
PCA	principal component analysis
PCD	point cloud data
ROV	remotely operated vehicle
SLAM	Simultaneous Localization and Mapping
USBL	ultra-short-baseline

Chapter 1

Introduction

1.1 Motivation

Acoustic measurement techniques have been used extensively to gather information about the topography of the sea floor. The favorable properties of sound propagation through water make acoustic range sensing possible over scales from centimeters to full ocean depth. Constructing a bathymetric map requires both a set of range measurements to the sea floor and the corresponding locations of the vessel or vehicle carrying the range sensor when the measurements were taken. The precision and accuracy of these two pieces of information dictates the fidelity of the resulting map. Thus, mapping is a coupled problem where inaccuracy in either range or position will corrupt the accuracy of the other during the creation of the map. In the limiting cases, a perfect range sensor will be limited by position or navigation errors and perfect navigation estimates will be limited by the range sensor accuracy. In any real mapping system, inaccuracies in both range sensing and navigation will be present, and efforts to improve the resulting product should therefore focus on the element contributing the greater amount of error to the final map.

As an example, consider that in recent years Global Positioning System (GPS) measurements have greatly improved ship-based sea floor mapping systems. Ships are now able to make maps all over the globe using accurate and repeatable navigation that was previously impossible to obtain. This major positioning advancement has improved large scale sea floor mapping accuracy to an extent that would not have been achievable by better sonar range measurements alone.

Bottom mapping from underwater vehicles, which operate at much closer proximity to the sea floor, offer the potential for much finer resolution and higher terrain accuracy than that achievable from surface-based surveys. Remotely operated vehicles (ROVs) and autonomous underwater vehicles (AUVs) are regularly outfitted with acoustic range sensors and are capable of flying survey patterns close to the bottom in rough terrain. Close proximity to the bottom avoids many of the acoustic limitations for ship-based surveys such as water depth and sound speed profiling. Vehicle-based mapping systems regularly support science, forensics, exploration archeology and military applications [4, 23, 118, 144, 145]. Unfortunately, good vehicle position information is still difficult to obtain underwater.

Although many methods of positioning underwater do exist, they are all limited by accuracy or scale. In comparison to the high sub-meter resolution of today's commercially

available vehicle-based range sensors, navigation remains the limiting factor when creating vehicle-based terrain maps. A single sonar ping, whether from a single beam or multibeam sonar system, represents a very accurate relative measurement between the sensor and the environment. The navigation limitation manifests itself as an inability to place the ping ranges in space to form an accurate representation of the environment in a single global coordinate frame. This thesis focuses on the navigation limitations of mapping algorithms and offers a solution designed to enforce consistency between the acoustic mapping sensor data and the navigation data. The end result is a terrain map constructed without the inconsistencies and mis-registrations that typically reduce the utility of maps created in navigationally-limited circumstances.

1.1.1 Problem statement

The map making process involves several steps which introduce error. The total mapping error diagrammed in Fig(1-1) symbolically shows the individual error contributions from navigation, sensor offsets, modeling and the mapping sonar itself. These divisions represent the steps required to take sensor measurements, in the sensor coordinate frame, and abstract them to a map. *Sonar* errors include all the factors related to obtaining a sensor relative measurement to the environment. *Sensor offsets* are the transforms between the vehicle frame and the mapping and navigation sensors that can only be directly measured with limited precision and are generally refined using the mapping data. *Modeling* errors are associated with the difference between the estimated and actual vehicle pose as a function of vehicle dynamics and navigation sensor noise. This primarily represents the vehicle frame attitude and depth estimation. Depth, pitch, roll and heading are measured from known environmental references and filtered with a vehicle model. *Navigation* errors are the potentially large scale $[x, y]$ positioning errors caused by dead reckoning navigation, heading sensor bias and deviation, and poor or unavailable ground-referenced position measurements.

Although all four pieces of the uncertainty contribute to the total error, vehicle-based mapping is currently navigation-limited. To reduce this limitation and move to a more equal distribution of errors this thesis focuses on the following tasks:

- creating additional vehicle positioning constraints by matching or registering sections of bathymetric data which have been viewed multiple times in a single survey,
- combining these constraints in a navigational framework to provide improved vehicle navigation estimates, and
- producing as a final product a dense surface terrain map of a natural sea floor with an associated error representation.

1.2 Related research

The tools and techniques used in the thesis have been adapted from the communities of robotics, acoustic underwater mapping, and graphical modeling. Although many of the individual components related to the goal of improved terrain mapping have been addressed,

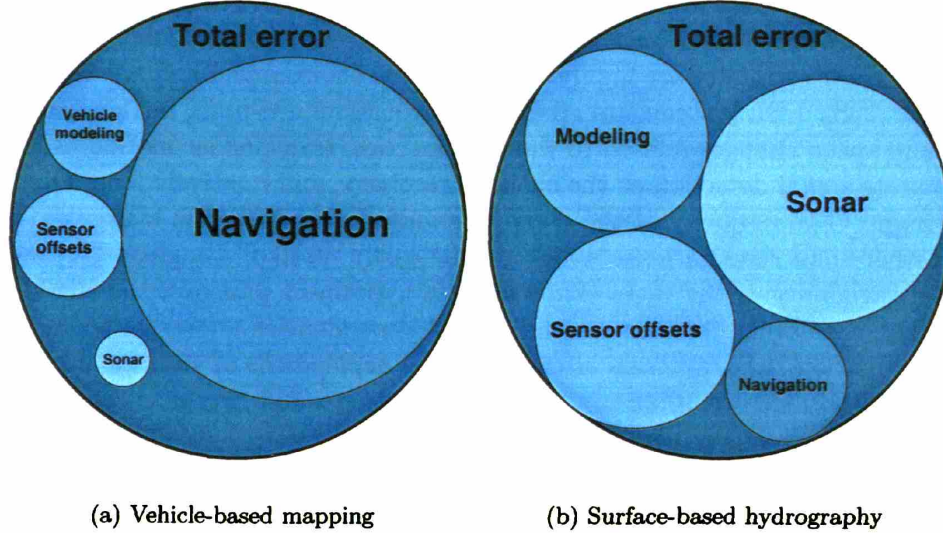


Figure 1-1: A comparison of the contributing errors for vehicle-based and ship-based mapping. (a) Proportional error sources for deep water mapping. (b) Error relations for surface-based mapping. Vehicle-based mapping is navigationally limited where as for surface-based surveys navigation is relatively well know in comparison to other potential error sources.

this thesis combines them for the first time into a robust algorithmic framework capable of handling unstructured seafloor mapping. The following sections of this chapter summarize the background and context for the concepts that serve as a building blocks for the presented mapping algorithm.

1.2.1 Underwater Navigation

The desire to create accurate acoustic and photographic maps with underwater vehicles has pushed the need for better underwater navigation systems and estimation techniques. Underwater positioning systems can be grouped according to methods which use fixed ground based references or those based on relative positioning through velocity integration. Each of these methods has its own associated error sources, and the choice of method is often dictated by the goals of the mapping effort. Ultimately, this thesis will focus on the usefulness of accurate DR navigation over short time scales.

Fixed Reference

Satellite based GPS, which can be used for accurate position estimation on land, does not work subsea due to the rapid attenuation of electromagnetic radiation in water. The closest analog underwater is long baseline (LBL) navigation [90] which uses bottom tethered acoustic beacons that are fixed at known locations. The round trip time of flight measurements between an acoustic transponder on the vehicle and the beacons can be used to triangulate the vehicle position in two and three dimensions. Typically operating at frequencies be-

tween 9 and 15 kHz, LBL systems can produce ground referenced position estimates with bounded error in deep water over kilometer scale ranges. Unfortunately, the actual error statistics for these estimates are highly coupled to the environment and are difficult to characterize [12] [138]. LBL systems are affected by acoustic multi-path, terrain-caused shadow zones, sound speed changes related to water properties, transponder motion caused by currents, accurate signal detection at the acoustic receivers, and relatively long, $O(1\text{sec})$ signal times of flight. The majority of these errors manifest themselves as biases and patterned outliers rather than random noise which can be easily filtered. Bingham [11] investigated the spatial variability of these errors using a hypothesis grid over the survey area and suggests that the ability to estimate the spatial dependencies allows for more robust and accurate navigation. LBL systems also require the deployment of additional infrastructure that makes quickly surveying an area difficult. In deep water a typical 3 beacon LBL net can take 24 hours to deploy and survey in.

Even given these drawbacks the benefit of a long range ground referenced position measurement is compelling enough to make LBL a standard navigation tool to produce position estimates accurate to between 1 and 10 meters. LBL performance can be improved, with the penalty of reduced range, by increasing the acoustic frequency. Systems like EXACT, which operates at 300 kHz, produce $O(1\text{cm})$ errors in position over ranges less than 200 meters [146]. The EXACT system has been used for underwater mapping [122,144] and to provide ground truth for DR navigation tests [139,140]. Over the shorter ranges this system is less susceptible to the bias and inaccuracy associated with sound speed errors, multi-path, and transponder motion.

Ship-based ultra-short-baseline (USBL) acoustic systems, which use an acoustic array to provide range and bearing measurements, in combination with surface GPS measurements can also generate navigation fixes in deep water [79,87]. The accuracy of these measurements however, is related by the angular resolution at the receiving array and translates to a position accuracy of $O(1\%)$ of the water depth. These systems do not require external beacons to be deployed, but do need a measurement of the water column sound velocity profile.

Overall, LBL and USBL provide useful data for working in the deep ocean, but the frequency dependent acoustic attenuation of sound in seawater [135] will always be a limiting factor in obtaining direct position measurements of high accuracy over long ranges.

Relative positioning

Velocity integrated navigation, commonly known as dead reckoning (DR), is the most frequently used method to navigate underwater vehicles. It requires no infrastructure external to the vehicle and relies principally on measurements of vehicle heading and ground relative velocity. The performance of DR navigation is directly proportional to the quality of the heading and velocity measurements [139], which each have their own inherent error sources. Heading measurements from magnetic compasses are often contaminated by random noise, heading dependent bias (deviation) and low bandwidth [53]. Fiber optic gyroscopes (FOGs) generate heading measurements that are much higher quality and bias free, but are currently only available from expensive inertial measurement units (IMUs) [98].

Velocity measurements for underwater vehicles typically come from an acoustic doppler

current profiler (ADCP) or DVL operating in a bottom lock mode. These bottom relative velocity measurements are typically accurate to better than $O(1\%)$ of the instrument velocity [73]. The combination of heading and DVL measurements has been used extensively for DR navigation and is generally expected to produce integrated position measurements accurate to 1% of the distance traveled [18, 139, 140] when heading dependent bias is minimal or nonexistent.

An additional complication to DR accuracy is the rotational offset between the attitude and velocity sensors. Raw velocity measurements are obtained in the DVL coordinate frame and need to be merged with heading measurements recorded in the heading sensor's coordinate frame. Although the coordinate frame offset can be roughly measured for an initial guess, any remaining error in the offsets will contribute to a growing deterministic bias in the integrated position estimates. When LBL measurements are available Kinsey [71, 73] has proposed methods to estimate this offset online using adaptive estimation techniques. When LBL measurements are not available a systematic way of determining this full offset has not been presented.

Although DR navigation is ultimately limited by time dependent error growth, the accurate short term navigation possible from precise navigation sensors is worth taking advantage of. The terrain mapping algorithm presented in this thesis will utilize this short term accuracy to construct small bathymetric maps over short time scales.

1.2.2 Sonar mapping

Acoustic mapping in the ocean has a long history of accomplishments and motivations [89, 135]. Starting with single beam ship-based acoustic soundings and progressing through evolutions of sonar design and positioning advancements, the achievable limits of sea floor mapping accuracy have been continually pushed. The thesis incorporates contributions from sea floor mapping efforts that can be broken down into the areas of hydrographic surveying, terrain aided navigation and vehicle-based acoustic mapping.

Surface-based hydrographic surveys

The roots of sea floor exploration and map making lie within the hydrographic community. This group's charter to provide the best possible maps to end users for navigation, exploration, and science has motivated considerable technological development. Multibeam sonar systems, capable of imaging swaths of the sea floor up to multiple times the water depth in width [6, 32, 33, 38, 88] have become standard tools for bathymetric mapping. Typically, multibeam measurements combined with ship's navigation, usually dead reckoned prior to GPS, are used to create tracks of bathymetric data that can be merged into a single map. A detailed error accounting for such systems is discussed by Hare [51]. The difficulty in merging crossing and overlapping tracks due to inconsistencies in the common areas has been a long standing problem. Nishimura [102] addressed this issue and suggested a 2D correlation measurement to determine a $[x, y]$ shift that minimizes the depth error between two overlapping sections of bathymetry. This method was used to constrain kilometer scale errors between crossing tracklines. A similar approach to remove errors was also presented by Kamgar-Parsi [67, 68]. More recently similar correlation measurements have been put

into a larger sparse matrix minimization [78]. This minimization utilizes the initial trackline positions and free surface gravity measurements as constraints. The resulting solution shifts individual tracklines that are assumed to be rigid. On a broader scale the compilation of many different surveys that have potentially different navigation errors has also been addressed [61]. In this work a Monte Carlo method was used to perturb depth estimates within the appropriate navigation related error limits to generate a composite map which shows a reduced variance in the predicted depth.

The hydrographic community has also investigated robust and automated ways to deal with the tremendous amount of the data generated by modern sonars systems [21]. Shallow water ship-based surveys can range in data size between 10^6 and 10^{10} individual soundings and surpass the capacity for interactive data filtering. These data sets typically consist of either beam ranges or points that have been transformed into 3D Cartesian space. In data sets this large the problem of outlier detection is significant as spurious soundings can easily corrupt the integrity of a trackline map or an entire survey. The common outlier rejection techniques [21, 22, 54, 55, 82] attempt to reject spurious ranges or points inconsistent with the surrounding data, either within an individual ping or in a preliminary map. It is worth noting that the outlier problem for surface based surveys is often more significant than in vehicle-based surveys due to the longer acoustic path length to the bottom and water column scatterers.

The transformation from individual soundings to a map has typically been done using various gridding techniques. These algorithms typically use a weighted sum of the sounding within a neighborhood of regularly spaced grid points. A more advanced gridding technique [20, 21] attempts to estimate the true depth at known points using the influence of neighboring soundings. This method also supports multiple depth hypotheses at a given location as a measure of robustness to outliers and systematic bias in the data. Other works on bathymetric gridding have focused on using adaptively generated Delaunay based triangular meshes rather than regularly spaced grids [22]. Triangular meshes are offered as a solution to the “low pass” effect that occurs with gridding algorithms and have the ability to adjust for density of the soundings on the sea floor.

Terrain aided navigation

There has been much interest in terrain aided navigation for underwater vehicles. The majority of this work has focused on the idea of generating a vehicle position estimate given an *a priori* map of the environment [24, 34, 104] rather than creating a map of the environment with which to navigate. These methods assume some type of onboard mapping sensor, typically a multibeam sonar, and some vehicle DR navigation capability. Carpenter [24] suggests the idea of using “local” or “short term” navigation to create small patches of bathymetry that can be matched to a larger known map. The most common method for obtaining a terrain match and a vehicle position measurement is correlation. Using a correlation measure alleviates the need to identify distinct targets in the environment and relies on more basic shape information. Carpenter has used bathymetric contours and a Hausdorff distance measure to determine matches and translational shifts between small sub-maps and a base map [24]. Nygren [104] proposed a correlation measure between a base map and the local terrain as measured by an acoustic array. These methods represent

the seafloor with contour lines or as a 2D height map. Although it is alluded to, none of these methods develop a framework for the simultaneous construction of and navigation with a terrain map.

The majority of terrain navigation algorithms use Kalman filters to merge the ground relative correlation measurements with the vehicle DR navigation. The Kalman filter requires a position estimate and a corresponding measurement uncertainty. The majority of these methods have not fully addressed this measurement uncertainty. The most complete treatment, by Nygren [104], relates the bathymetric error between the measured local terrain and prior maps using a Gaussian error assumption. Assuming the depth measurements are independent over the matching area a Gaussian likelihood is created as a function of correlated position and used to estimate the covariance of the terrain match.

Particle filter methods [48, 70] have also been suggested for underwater terrain navigation. Bachmann and Williams [3, 141] suggest that under typical operating conditions a vehicle instrumented with only a single beam echo sounder can improve its DR navigation significantly with a Rao-Blackwellized particle filter. These methods rely on the availability of a prior terrain map of the environment and use the discrepancy between the measured depth at the vehicle location and the map depth for particle resampling.

There have been far fewer attempts to use feature-based map matching methods for terrain aided navigation. Sistiaga [123] has suggested using an attribute vector to describe the local geometry of features defined as morphologically invariant points. These points are taken from the difference between a low resolution base map and a smoothed version of a vehicle-based high resolution map. Majumder [85, 86] has used a feature-based sums of Gaussians method in a Bayesian framework for terrain aided navigation. By modeling feature locations as 2D Gaussian random variables he was able to construct a feature map over a grid of the sea floor. The sums of Gaussians environmental model provides a more complex representation of the environment than single points while maintaining the attractive computational properties of Gaussian descriptors. This method is also able to side step the data association issues required by most Kalman filter type algorithms. The vehicle navigation was propagated over time also with a Gaussian model. To create map relative position measurements a correlation technique was used to match the currently visible features to features stored in the map.

Vehicle-based mapping

Efforts to evaluate the mapping accuracy of vehicle-based underwater surveys have been somewhat limited. Stewart [128, 129] was the first researcher to use land robotic techniques for mapping with uncertain sensors and apply them to underwater acoustic mapping. Using the occupancy grid methods developed by Moravec [95] and Elfes [35] he attempted to model how the navigational and mapping sensor uncertainties contributed to a terrain map. Although this method was able to produce useful maps, its major disadvantage was its own honesty. Since the contributing factors to the mapping error (sensor and navigation) were combined into a single sensor model prior to representation in the occupancy grid, very uncertain navigation data would “blur” what would have otherwise been high resolution mapping sensor data. The resulting maps had soft edges and a “low-passed” look to them. Additional work on the occupancy grid concept for an extension to 3D [94] and the associ-

ation of specific sensor measurements to individual cells in the grid [131, 132] suggests some possible improvements to this limitation, but the method is still hindered by unfavorable scaling in large environments and when a large number of sensor readings are taken.

Exploiting the accuracy of the EXACT 300kHz navigation system Singh [121, 122] looked at the effect the mapping sensor to vehicle frame offsets have when combining data from multiple tracklines. In this case the EXACT system was able to reduce the $[x, y]$ navigation uncertainty to a small enough size that the sensor offset error was the dominant error contributor to the map, Fig(1-1). The sensor offsets were determined by minimizing a measure of the surface variance.




Figure 1-2: A photomosaic and bathymetric map created over an archaeological site [122]. Corresponding objects are indicated. This bathymetric was made using the EXACT LBL system capable of centimeter level precision over ranges of < 200 meters.

Using standard 9 kHz LBL navigation Jakuba [63] has been able to create maps of the rugged terrain found a hydrothermal vent sites from sonar data collected with an AUV. This work mentions errors which cause difficulty in merging tracklines into a correctly registered composite survey. Although the mis-matches in overlapping tracklines are indicated, a systematic method for removing the registration error is not presented.

The author [114] has shown that accurate composite terrain maps can be assembled by combining acoustic range images taken from multiple vantage points. In this work the complications associated with navigation error were limited by assuming a few discrete sensor vantage points, and more effort was expended on the creation of an accurate small scale scene. Here range image registration techniques were used to obtain refined estimates of the sensor vantage points and create a composite scene.

1.2.3 Simultaneous Localization and Mapping

In recent years the Simultaneous Localization and Mapping (SLAM) community within robotics has focused on the coupled problem of mapping an area while concurrently deriving improved position estimates from the map. SLAM algorithms have been shown to greatly improve robotic mapping in applications where the robot navigation is poor and the mapping sensor accuracy is high. These situations are similar in nature to the deep sea mapping problem where accurate navigation is hard to obtain. Algorithmically, the

attractive feature of this methodology is that it provides a common framework for manipulating navigation and mapping uncertainty. The specific solutions to the SLAM problem differ according to the types of environmental measurements they utilize and the manner in which they fuse the measurements with additional data, such as navigation. The following sections review many of the current SLAM techniques that are relevant for subsea mapping.

Environmental representations

The seminal paper by Smith [126] framed the SLAM problem as a probabilistic estimation problem and started what have become known as *feature-based* solutions. These methods attempt to identify and track the location of specific features in the environment. Feature locations, typically described using Gaussian approximations, are added to a filter state vector and represent the “map” of the environment. For this type of solution the mapping sensor measurements must be assigned to individual features currently in the state vector or declared as new features and added to the state vector. This data association problem can be a source of divergence for these algorithms [99]. Feature-based methods have been proposed to navigate AUVs using range and bearing data from active beacons or passive sonar targets in the environment [100, 101, 125, 130, 142]. The previously described terrain aided navigation by Majumder [85, 86] is a feature-based method using natural landmarks.

Featureless approaches do not extract specific features from the mapping sensor measurements and instead use the raw sensor measurements directly. This is commonly done with sensors that map a section of the environment at once. Lu [83] proposed one of the original featureless SLAM approaches using laser range scans of a 2D environment. Numerous SLAM algorithms continue to use 2D and more recently 3D laser scanning [103] to provide a representation of the environment and relative position measurements. Featureless methods usually associate an individual scan or a set of scan locations to a pose kept in a state vector.

Solution methods

Proposed SLAM frameworks to integrate the mapping and navigation data include the extended Kalman filter (EKF) [126], particle filters [93], sparse information filters [36, 133], junction tree filters [107] and constraint networks [83]. Each of these approaches have advantages and disadvantages in the context of bathymetry mapping, where the ability to retain and update old vehicle positions is desirable.

A delayed state version of the recursive EKF solution provides an iterative formulation for mapping and retaining knowledge of prior platform positions [39, 81]. This solution is subject to severe limits due to computational growth if additional methods are not used to reduce the $O(n^2)$ computational burden related to a dense covariance matrix at each measurement update [42, 47, 80]. The EKF solutions are also subject to linearization errors as the constraints between delayed states are linearized only when they are incorporated into the recursion [26]. If the trajectory of delayed states is deformed significantly, such as with a large loop closure, the constraints may no longer be accurate and bias the solution. However, this iterative solution has a possible real time implementation that could produce adjusted navigation information useful to the surveyor trying to ensure complete coverage of a survey area. Recently, the attractive sparseness properties of the information matrix have

been utilized in methods for reducing the computation of similar linear Gaussian iterative methods [42] [133]. In the context of underwater photographic mapping Eustice [36] [37] has addressed the scale problem for delayed state filters represented in the dual information form. The sparseness properties of the information representation have allowed the number of delayed states to be extended from 10's to 1000's.

Alternatively, if the navigation problem is treated as a more general network of possibly nonlinear constraints derived from mapping data that link previous vehicle positions to one another, several other possible solutions exist. In a feature-less scan-matched representation of the environment that assumes independent Gaussian measurement errors between scans, a maximum likelihood (ML) solution for the pose locations can be formulated. This solution takes the form of a linear problem involving a constraint matrix [83]. Extensions of this methodology have been used for very large maps [49]. Frese [41] presents a constraint based multigrid solution designed as an incremental mapping approach to achieve $O(n)$ update computation and retain the ability to relinearize pose constraints during the solution process. Bosse's [15] [16] ATLAS solution keeps all the constraint information in a relative framework and uses a non-linear least squares solution to resolve the resulting network for the pose positions. To avoid potential problems with overconfidence in network based solutions associated with unknown cross correlations Schlegel [117] advocates a pose network solution based on Covariance Intersection (CI). All of these approaches require an accurate initial guess for the solution and a correct network topology of links.

1.2.4 Registration methods

The ability to reorient ones self when given access to a set of maps requires that a registration, or relative transform, can be determined between maps portraying common portions of a larger scene. In the context of underwater navigation, map registration offers the possibility to recognize previously visited locations and reset navigation errors that have been accumulating over time. Two and three dimensional registration techniques have been actively researched in the fields of computer vision and graphical modeling, and are now being applied liberally in the field of robotics to create relative measurements of position.

All of the registration methods utilized in the terrain navigation methods described in Section 1.2.2 assume a 2D height map to represent the terrain and perform registrations. To move toward more general terrain matching it is necessary to consider methods which can work in full 3D. The close proximity to the sea floor provided by vehicles, will increase mapping resolution, but also decrease the ratio of viewing distance to scene relief. From vehicles there will be more extreme angles of incidence between the sea floor and the sonar beams, and an increased risk of occlusions in highly featured areas.

A significant body of work surrounds 3D registration techniques used to construct volumetric representations of objects and scenes scanned with laser range finders. Laser range finders produce an "image" of highly accurate ranges. The majority of the techniques to register the 3D point clouds constructed from range images are based on the iterative matching methods originally proposed by Besl [10] and Chen [27]. Improvements to the basic methods have addressed computation [116], robustness [43], scale [112], surface attributes [44] and solution methodology [46,92]. More recently there has been application of 3D modeling and registration in outdoor scenes [2, 76, 127, 137] and robotics [75, 120] [103].

In the past, the primary difference between the modeling and robotics applications was whether the sensor location was assumed to be known. More recently this has changed with modeling work that assumes no *a priori* knowledge of the object orientation within the sensors view [58].

The vast majority of registration methods based on point sampled surfaces, and the associated techniques for surface normal estimation, [56,66,91,108] use principal component techniques as a measure of robustness to sensor noise. However, with laser scanners the level of assumed sensor noise relative to the feature size and sampling density in the scanned scenes is small. There have been only a few attempts [25,114,134] attempts to use similar registration methods for sonar sensing, and a systematic approach to handling sonar related errors has not been presented. A broad survey of processing techniques related to acoustic imaging has been presented by Murino [96,97].

This thesis will also apply both 2D correlation and 3D registration techniques to the mapping sonar data when developing terrain based relative pose measurements.

1.3 Thesis breakdown

1.3.1 Outline of methods

The main objective of this thesis is to demonstrate that creating a feedback path that enforces consistency between the terrain mapping data and vehicle navigation data will produce more self consistent and accurate bathymetric maps. The proposed framework creates this feedback path by using small terrain sub-maps created over short time scales with a vehicle navigation estimate generated by dead reckoning. The registration of these sub-maps creates relative position measurements between the current and past vehicle states. These measurements are then fused into a SLAM navigation framework based on a delayed state EKF [81]. When sub-maps are created they are attached to a snap shot of the vehicle state, which is then stored in the delayed state vector and used as a local origin for the bathymetry in the sub-map. A schematic of the proposed algorithm is shown in Fig(1-3).




Figure 1-3: *The basic concept behind the sub-mapping algorithm. The vehicle has flown the green trajectory above the bottom and the survey swath has been broken into a series of sub-maps. The reference frames along the trajectory indicate the vehicle positions where the sub-maps were started. The red regions indicate where the maps cover a common area of the seafloor and the potential exists to establish a link, shown in red, which constrains the relative position of the previously visited positions.*

The survey of SLAM algorithms presented in Section 1.2.3 suggests there are many potential options for a framework to combine the sub-map measurements and the vehicle navigation estimation. The choice of the EKF based solution is based on the following observations concerning many of the SLAM options.

- An important distinction between this application and many land-based applications is that underwater the surveyor can design the vehicle trajectory to avoid the need for closing large loops. This is often not possible in land-based applications where the vehicle trajectory is constrained by the environment, such as in a building. As a result we consider this application to be less prone to linearization and link proposal issues associated with closing large loops.
- The focus of the problem is on accurately mapping a specific area of interest on the sea floor rather than covering expansive amounts of terrain. Knowing this it is not necessary to penalize a choice of method based on a scale limitation. Experience suggests that maintaining up to 100 prior poses will suffice for a developmental and useful solution.
- It is desirable to maintain a potentially real-time implementation. As such, batch methods requiring all of the data are less desirable.

The relative pose measurements between sub-maps are obtained using sequential 2D and 3D registrations techniques. Terrain maps are stored using all the original mapping data and the registrations are performed without extracting distinct features from the mapping sensor measurements. By retaining all of the dense mapping data in the sub-maps the ability to extract additional geometric information when needed is preserved. The desire to accommodate a 3D registration is motivated by some of the limitations found in applying 2D methods to vehicle mapping in highly featured areas [62] [63].

1.3.2 Assumptions and restrictions

The algorithm and procedures presented here are considered applicable to a broad variety of applications requiring AUV and remotely operated vehicle (ROV) bathymetric surveys. To this regard the following list of conditions applies to the methods developed within this thesis.

- Due to the relatively short ranges between the vehicle and the bottom, it is assumed that the speed of sound is constant. Although in the proximity of a hydrothermal vent system this assumption can be easily violated, there are relatively few instances where a sonar will image the bottom directly through a large amount of hydrothermal fluid. Additionally, in such a complex spatially varying environment it is not realistic to consider that sufficient water property data could be taken for an accurate sound speed correction.
- Over the course of a survey the terrain being covered is considered static. There is no explicit accounting for the possibility of a changing environment.

- The terrain is considered unstructured and natural. Man made targets or beacons in the environment are not explicitly formulated in this algorithm. The algorithm requires a minimum amount of 3D terrain richness or structure consistent with what would be expected at geologically or archaeologically interesting sites. Obvious limitations to this method exist over large flat and featureless areas of the sea floor.
- Over the course of a survey all sensor positions with respect to the vehicle are assumed to be constant. A procedure to determine the static offset between the vehicle frame and the mapping sonar using short term navigation and mapping data will be presented.
- The primary navigation information used in the presented algorithm is derived from on-board sensor data. This method assumes the vehicle platform is instrumented with sensors sufficient to generate a dead reckoning position estimate. This would require at least 2D bottom relative velocity, vehicle heading and pressure depth.

1.3.3 Contributions

The main contributions of this thesis are as follows.

- For the first time a delayed state SLAM algorithm is applied to bathymetric mapping and real world results which show a clear improvement in mapping accuracy are given.
- A demonstrated improvement to 3D registration performance based on a point selection technique that incorporates properties of sonar mapping data is shown.
- A robust error metric to visualize artifacts in bathymetric maps is developed.

1.3.4 Thesis structure

The remaining chapters of this thesis are broken down to cover the individual aspects of the presented mapping algorithm. **Chapter 2** covers the basic aspects of acoustic range sensing and how they relate to mapping. The core of the SLAM algorithm is covered in **Chapter 3**. This chapter reviews the basic delayed state EKF and covers the specifics for this problem, including the vehicle modeling and the sub-map handling. **Chapter 4** develops methods for registering the small sections of acoustically mapped terrain generated by the EKF. **Chapter 5** presents experimental results for two surveys over a hydrothermal vent site. This sub-mapping method is compared to more standard mapping techniques and examples are given to show the robustness and failures of the bathymetric sub-mapping algorithm. **Chapter 6** concludes with a summary and suggests directions for future work and further improvement.

Chapter 2

Acoustics for mapping

2.1 Introduction

This chapter describes the sonar range sensing details relevant to the proposed sub-mapping algorithm. Based on the argument presented in Chapter 1, that the leading order cause of error in vehicle-based bathymetric maps is navigation related, the treatment of the acoustic range sensor itself is intentionally simple. The acoustic data is reduced down to a set of ranges and “confidences” for each sonar ping that are used for all subsequent processing. The ranges are defined for each beam using the peak returned amplitude and the confidence measure is based on the duration of the backscattered return windowed around the determined range. The final set of ranges is produced after automated data cleaning steps remove outliers from a preliminary set of proposed ranges.

2.2 Range determination

Oceanographic sonars used for vehicle-based mapping typically operate at frequencies greater than 100 kHz and trade off increased range resolution at the expense of sensing distance. The high acoustic frequency places mapping sonars in the rough surface scattering regime where incoherent contributions from individual bottom scatterers are primarily responsible for producing the backscattered acoustic energy. The transition to rough surface scattering from specularly directed scattering occurs when the incident acoustic wave length is proportional to surface shape excursions, or roughness, over the size of the beam footprint [89, 135]. For typical vehicle surveys flown between 15 and 50 meters in altitude the foot print size will be $O(1\text{m})$ and the wavelengths will be sub-centimeter for frequencies greater than 150kHz. Within this scattering regime the grazing angle dependence on the back scatter strength should be less significant than with spectral scatter and the duration of the return pulse should be proportional to the interaction length with the bottom [89, 135].

The sonar modeling only assumes that a high frequency pulse of short time duration τ is sent with a scanning single beam or a multibeam sonar, and that the return signal will be discretely sampled. For a multibeam system the sampled beam $s_\theta[k]$ at pointing angle

θ is taken as the magnitude of the complex beam formed signal

$$s_{\theta}[k] = \left| \sum_{n=1}^N x_n[k] \exp^{-ju(\theta,n)} \right|, \quad (2.1)$$

where, N is the number of head elements and $u(\theta, n)$ is the appropriate phase correction for each element of the receiving array. A sample beamformed ping is shown in Fig(2-1). The range to the bottom r along a beam is determined from the time of the peak amplitude for the returned signal assuming a constant sound speed. This detection method will be more accurate for beams near normal incidence and less accurate for beams incident with the bottom away from normal [50]. A sketch of the beam geometry is shown in Fig(2-2(a)). Due to the rough surface scattering, the side lobe interference created by high intensity specular scattering know to corrupt the near normal beams [1] has not been noticed. However, away from normal incidence the longer interaction length with bottom increases the probability that scatters off the beam axis will contribute to the return at times different than scatters on the beam axis. Phase based range detection methods for multibeam sonars, [50, 74, 143], can be applied to improve this performance, however the accuracy will still be limited by the seafloor roughness properties [13, 84] that affect phase coherence.

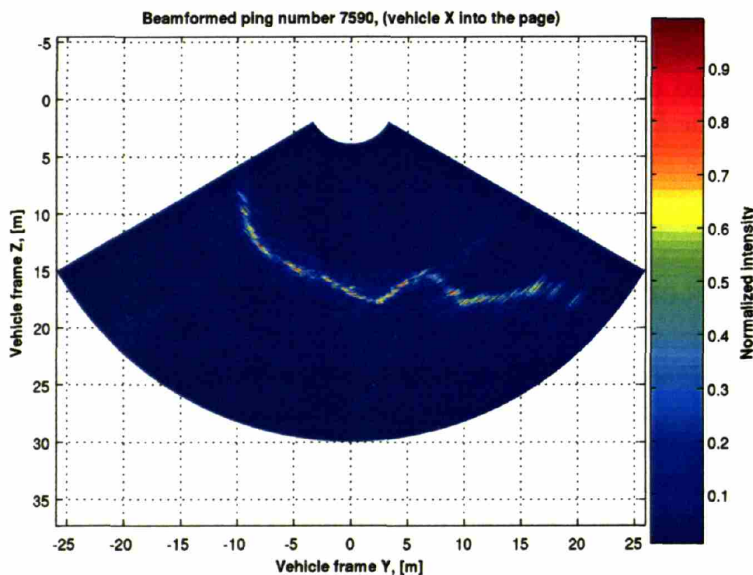


Figure 2-1: Sample SM2000 beam formed sonar image drawn in a Cartesian coordinate frame. The color scale indicates the beamformed amplitude normalized to the maximum returned amplitude. This ping is oriented as if the vehicle, located at $[0, 0]$, is flying into the page with a steep terrain rise to port. Note that the downward slope on the right is poorly imaged.

The maximum range resolution for a sonar transmitting a fixed length pulse of time duration τ is determined by the along axis depth of a scattering volume in which the returns from individual scatters can no longer be distinguished. This is the well know range

resolution cell

$$\Delta r = \frac{c\tau}{2}, \quad (2.2)$$

where c is the sound speed. To consider an acoustic return with finite samples taken at range spacings of dr a range error model of the form

$$\sigma_r = \sqrt{\left(\frac{dr}{2}\right)^2 + \left(\frac{c\tau}{4}\right)^2} \quad (2.3)$$

has been developed to characterize range error [50, 51]. These simple error estimates do not however consider the direct effects of incidence angle, bottom type, surface roughness and range on a ping-to-ping basis. As an alternative to the more complex statistical error modeling this would require, a confidence based approach is used instead to indicate returns with potentially poor range detection properties. For the peak amplitude detection method range inaccuracy will increase as the duration of the returned signal increases and a single peak in the backscattered energy becomes less distinct [17]. As a measure of return duration, D , the second moment of the returned pulse windowed around the determined range is calculated as

$$D = \left(\frac{\sum_{k=-w}^w k^2 s[k + k^*]}{\sum_{k=-w}^w s^2[k + k^*]} \right)^{1/2}, \quad (2.4)$$

where w is the number of samples specifying one half of a window width and the maximum return occurs at sample k^* . This measure serves the purpose of indicating beams that have interacted with the bottom away from normal incidence, been corrupted by side lobe interference or for any other reason lack a distinctive unimodal return peak. The histogram in Fig(2-2) shows that the duration D correlates with beam incidence quite well. The surface normals used to verify this relation were estimated from a 3D constructed terrain map and the surface normal estimation method described in Appendix B.2. In the subsequent processing the calculated duration for each range will not be used directly as a measure of range variance, but as an indicator of potential accuracy by which particular beam ranges are included in or excluded from the mapping process.

2.3 Outlier rejection and ping clean up

The sonar data processing is designed to automatically determine the beam ranges and durations, and then remove all returns suspected as range outliers. The outlier rejection will produce a final set of beam angles and ranges. Since the sonar ranges will be used to create small sub-maps, prior to creating a single composite map, the data cleaning is setup to operate on the range returns directly instead of a 3D point cloud. The steps in the range detection and data cleaning are outlined in Fig(2-3) and described below.

2.3.1 Inner ping

Within a single ping outlier rejection is accomplished with an amplitude threshold followed by a median filter based rejection. A minimum amplitude threshold is dynamically set to

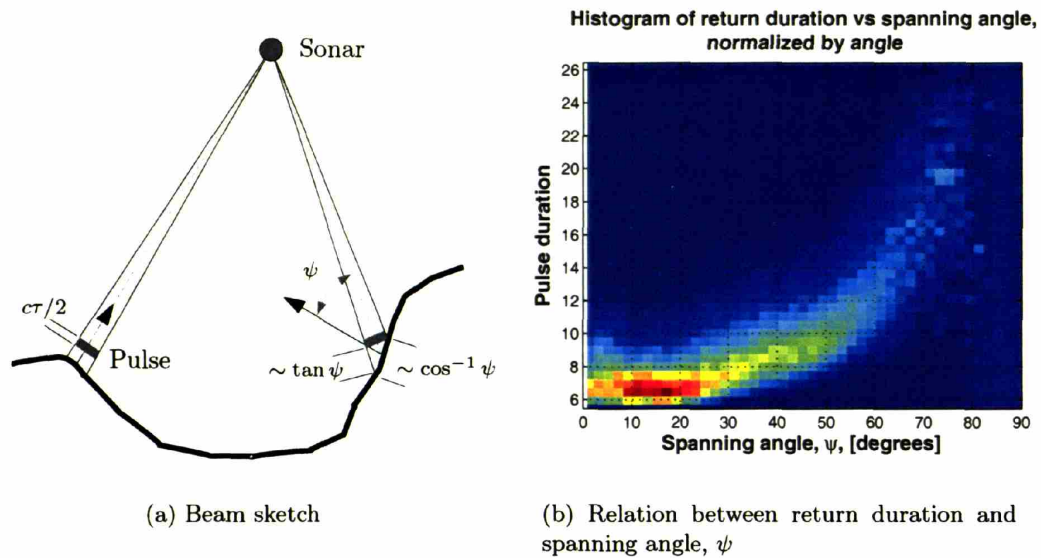


Figure 2-2: Angle of incidence dependence on back scattered return. (a) Sketch showing normal and grazing incidence. Away from normal incidence the sonar pulse will have a more interaction with the bottom and create a backscattered signal of longer duration. The duration will increase proportional to $\tan \psi$, where ψ is the spanning angle between the surface normal and the beam axis. The across track foot print will grow proportional to $\cos \psi$. (b) 2D histogram showing the relationship between the returned pulse duration and spanning angle. This was determined using surface normals calculated from mapping data. This graph has been normalized for each spanning angle bin.

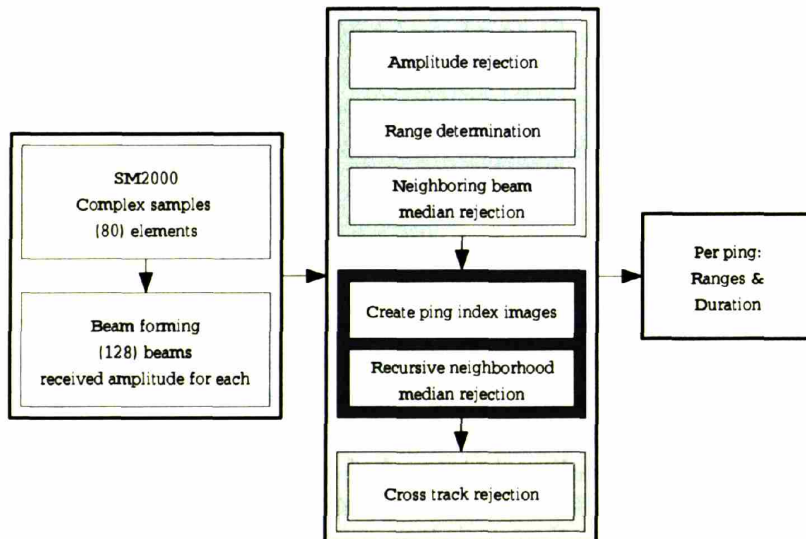


Figure 2-3: Steps in the automated sonar data processing.

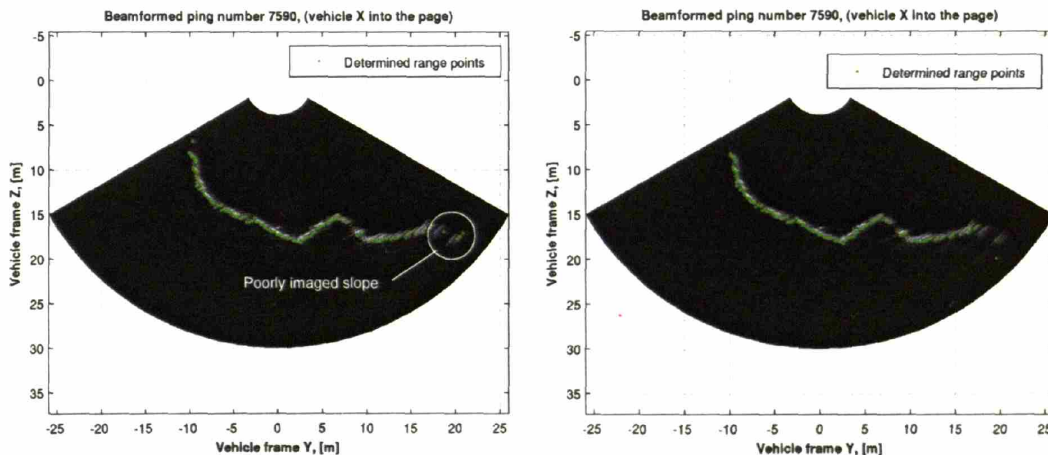
remove ranges from beams with little return energy. The time varying gain for the sonar

$$TVG = A \log r + 2Br + C \quad (2.5)$$

is used to account for spreading and attenuation. The remaining amplitude fluctuation between beams can be attributed to the bottom backscatter coefficient and the ensonified area. The dynamic threshold starts from an initial value and identifies ranges which fall below. If the number of ranges below the threshold exceeds a specified number the threshold is reduced, otherwise those range returns are removed. The initial guess for the threshold value can be related to beam pattern side lobe level relative to the main lobe, and the peak returned amplitude across the ping. For the data presented here the threshold was started at 22% of the peak amplitude returned over the ping.

Median rejection is accomplished by calculating the median range for a specified number of beams to each side of a selected beam. If the difference between the selected beam range and the median is greater than a threshold, the beam range is removed. This rejection is done from the inner beams to the outer.

These two simple checks are able to remove the significant fraction of range outliers within a single ping. An example range detected ping is shown in Fig(2-4).



(a) Beam ranges detected by amplitude

(b) Cleaned up beam ranges

Figure 2-4: Example of intra ping median rejection. (a) Single sonar ping with the return ranges indicated. (b) The same ping after outlier rejection. Note that a few range values on the poorly imaged slope to the right have been removed and many beams with low returned amplitude do not have a range defined.

2.3.2 Over multiple pings

The outlier rejection within a single ping will fail when a group of ranges are incorrect in a similar way. This can occur if another acoustic instrument contaminates a ping or if

large number of beams do not hit the bottom and instead pick up noise. To account for this neighboring pings in time are also used for median rejection. A range image using adjacent pings can be created for this purpose, Fig(2-5(a)). A median range image can also be calculated using a neighborhood of range pixels surrounding each pixel. Outlier ranges are identified by differencing the range image and the median image, and finding the returns that exceed a threshold. The image in Fig(2-5(b)) shows the kinds of outliers this method will detect.

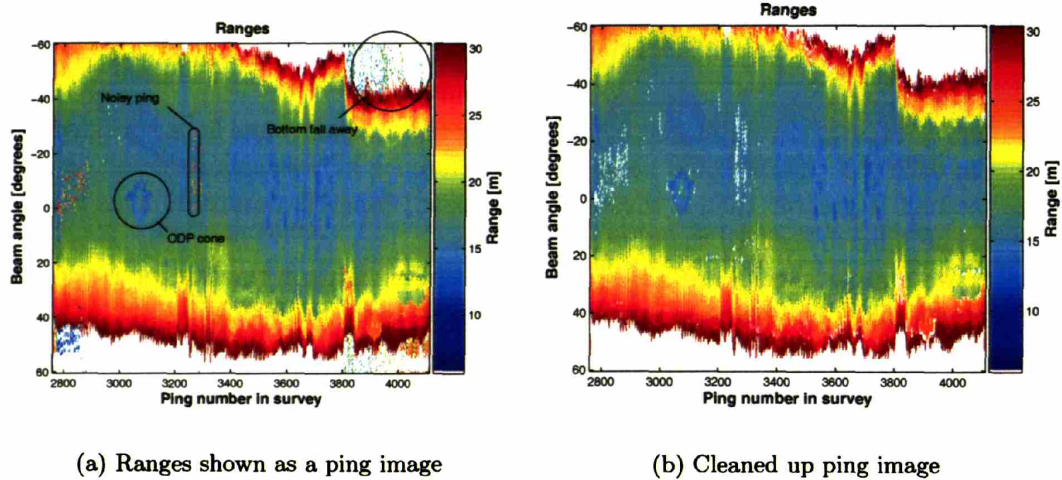


Figure 2-5: Results of the neighboring ping outlier rejection. (a) Successive pings can be shown as a range image in a beam angle and ping number space. The marked areas show regions with outlier ranges. Note the majority of outer beams do not have ranges defined. An Ocean Drilling Program (ODP) re-entry cone shows up clearly in the range data. (b) The same pings as (a) after the ping image median rejection. Numerous outliers have been removed without removing a large number of the good beam ranges.

As a final rejection step a cross track filter can be used. This check is made to ensure that the determined ranges for increasing beam angles away from nadir correspond to bottom points that are further outboard than the previous ones. This is useful for reducing range errors at the outer beams and applicable if the environment contains no overhanging features. The check for each side of the array is simply

$$r_i \sin(\theta_i) > r_{i-1} \sin(\theta_{i-1}), \quad (2.6)$$

where $\theta_i > \theta_{i-1}$ are the beam angles away from nadir, where $\theta = 0$.

2.4 Summary

This chapter has detailed the very general assumptions related to the acoustic range sensor requirements and processing for the proposed sub-mapping algorithm. The individual beam ranges to the sea floor are determined using amplitude only information in a manner appli-

cable to most commercially available sonar systems. Knowing that range detection will be poor away from normal incident, a simple returned pulse duration statistic is used to indicate potentially inaccurate ranges. An automated data cleaning process is used to remove outlier ranges and reduce the set of initially proposed ranges to the set that will be used for mapping. Since surface sampling redundancy can be build into surveys by overlapping tracklines, the thresholds for the data cleaning are set aggressively to remove questionable range returns that could cause error in the sub-map terrain registration process.

Chapter 3

Sub-mapping SLAM bathymetry

3.1 Introduction

The proposed sub-mapping algorithm is formulated around the delayed state extended Kalman filter [39, 81]. The delayed state filter is used to compute a dead reckoned vehicle trajectory from navigation sensor data and allow for updating the position estimates of previously visited vehicle locations. A *continuous-discrete* EKF [5] implementation is used to handle asynchronous navigation measurements and produce a causal estimate of the current vehicle position and attitude. The vehicle position and attitude estimate is used to project the mapping sonar data over a short time window and create local terrain sub-maps. The data within each sub-map will be referenced to a local origin declared to be the current vehicle pose at the time the sub-map is started. Sonar data will be added to a sub-map until one of several conditions is met to indicate the map's closure. A new map, with a new reference frame, is started immediately following the closure of the previous map.

The diagram in Fig(3-1) shows the data paths for the filter. The creation of bathymetry sub-maps requires knowledge of the current vehicle state and the range detected sonar data. Newly created sub-maps are stored and their reference frame origins remain in the delayed portion of the filter state vector. When a map is closed checks are made to determine possible overlap with the other maps in the catalog. Overlapping maps are pairwise registered to generate relative pose measurements between the sub-map origins stored in the delayed state vector. As this filter runs the origins of the sub-maps are updated using the relative pose measurements, and a network of links between previously visited vehicle poses is created.

The remaining sections of this chapter outline the specifics of the delayed state EKF for the problem of bathymetric sub-mapping. In particular, the constant velocity vehicle model sufficient to capture the slow dynamics of a broad class of marine vehicles is described, and the relevant issues related to sub-map generation are discussed.

3.1.1 State vector and coordinate frames

The proposed filtering algorithm and sub-map manipulation strategy is developed around the idea of reference frames and pose composition. A 6 DOF pose can be considered a coordinate transformation that represents the spatial relationship between two reference frames. The basic *head-to-tail* and *tail-to-tail* composition relations developed by Smith [126] will

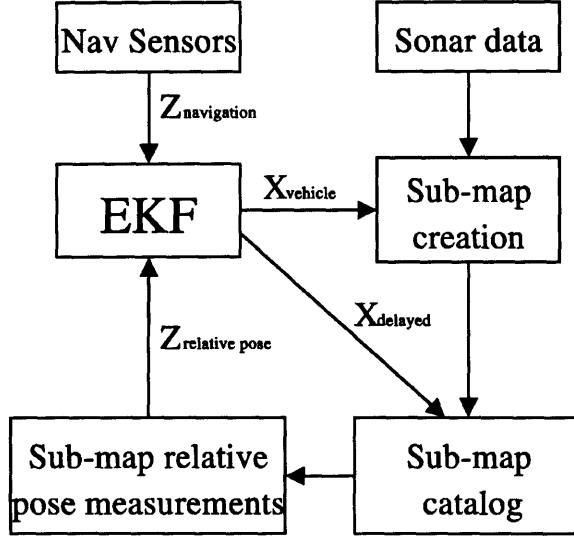


Figure 3-1: *The delayed state EKF block diagram. The proposed algorithm utilizes vehicle navigation data to create small bathymetric sub-maps. The sub-map origins will be held in the delayed state vector and used to create relative pose measurements that reduce navigation error.*

be used to manipulate these transformations. These composition relations (summarized in Appendix A) join together successive pose relations and propagate first order estimates of their uncertainty. The relevant coordinate frames for the filter are shown in Fig(3-2).

The filter state is written to represent the vehicle body frame position with respect to a local level integration frame, indicated by \mathbf{x}_v . All navigation sensor measurements relate to the vehicle body frame through individual sensor transforms that specify the static pose offsets of each sensor as physically mounted to the vehicle. The mapping sonar measurements are also related through the vehicle body frame and a sensor offset. All of the vehicle-to-sensor offsets are considered static. Additionally, we consider that an individual sensor, such as a north seeking heading sensor, will produce a sensor measurement with respect to its own sensor local level frame that may differ from the vehicle local level frame.

To accommodate sub-mapping the complete filter state vector, \mathbf{x}_{aug} , is partitioned into the current vehicle state \mathbf{x}_v and a delayed portion consisting of previously visited vehicle positions. The state vector in (3.1) shows the vehicle state and N delayed states serving as sub-map origins.

$$\mathbf{x}_{aug} = [\mathbf{x}_v^\top, \underbrace{\mathbf{x}_{s_1}^\top, \dots, \mathbf{x}_{s_N}^\top}_{\text{delayed states}}]^\top \quad (3.1)$$

This state vector will grow in length as new sub-maps are created and delayed states representing their locations are added to the filter. The notation for the delayed states indicates that the delayed state, \mathbf{x}_{s_i} , marked by subscript s describes the transform from the local level origin to the origin of sub-map i . The covariance matrix for the filter describes the covariance of the vehicle, $\mathbf{P}_{\mathbf{x}_v \mathbf{x}_v}$, the covariance of the sub-map origins, $\mathbf{P}_{\mathbf{x}_{s_i} \mathbf{x}_{s_i}}$, and all of the respective cross correlations $\mathbf{P}_{\mathbf{x}_v \mathbf{x}_{s_i}}$ and $\mathbf{P}_{\mathbf{x}_{s_i} \mathbf{x}_{s_j}}$.

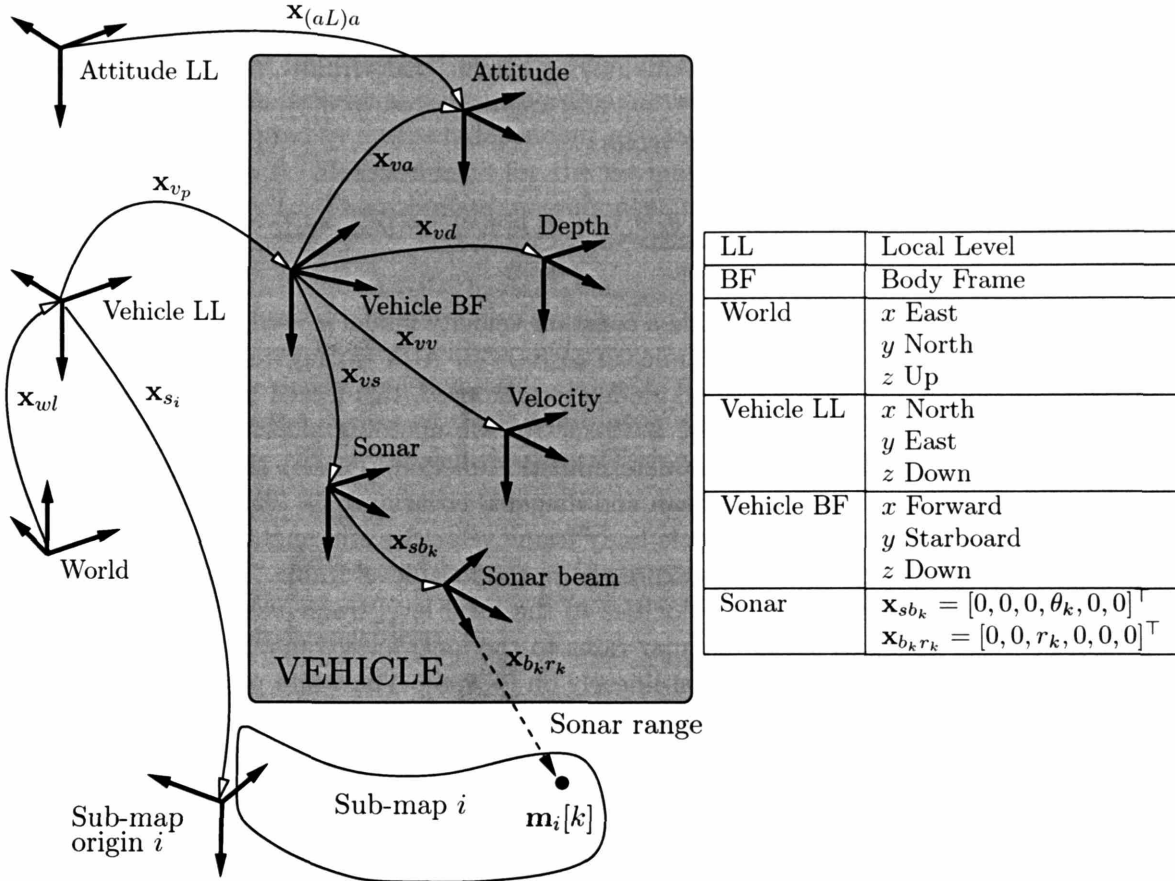


Figure 3-2: Coordinate system overview. This diagram illustrates the coordinate system convention used to model the vehicle and sensor frames. All transforms are parameterized over 6 DOFs. The static sensor offsets, $\{\mathbf{x}_{va}, \mathbf{x}_{vd}, \mathbf{x}_{vv}, \mathbf{x}_{vs}\}$, are measured with respect to the vehicle body frame. A procedure to refine the vehicle-to-sonar offset using the mapping data is given in Appendix C. To avoid excessive subscripting the vehicle state and sub-map origins will be written as \mathbf{x}_v and \mathbf{x}_s , respectively and the local level frame origin is implied. Transform \mathbf{x}_{sbk} is used to describe the angles for the individual sonar beams as a roll with respect to the sonar frame. The measured sonar ranges R are considered along the z axis of the rolled sonar beam frame. The k_{th} 3D point within sub-map i is written as $\mathbf{m}_i[k]$ and located at the end of the beam range vector \mathbf{x}_{brk} .

$$P_{aug} = \begin{bmatrix} P_{\mathbf{x}_v \mathbf{x}_v} & P_{\mathbf{x}_v \mathbf{x}_{s_1}} & \cdots & P_{\mathbf{x}_v \mathbf{x}_{s_N}} \\ P_{\mathbf{x}_{s_1} \mathbf{x}_v} & P_{\mathbf{x}_{s_1} \mathbf{x}_{s_1}} & \cdots & P_{\mathbf{x}_{s_1} \mathbf{x}_{s_N}} \\ \vdots & \vdots & \ddots & \vdots \\ P_{\mathbf{x}_{s_N} \mathbf{x}_v} & P_{\mathbf{x}_{s_N} \mathbf{x}_{s_1}} & \cdots & P_{\mathbf{x}_{s_N} \mathbf{x}_{s_N}} \end{bmatrix} \quad (3.2)$$

3.2 Vehicle model

The pose of the vehicle body frame is described by a six DOF parameterization with position and attitude variables measured in the local level reference frame. The angular conventions

follow those of Fossen [40], using a heading ψ , pitch θ and roll ϕ Euler angle sequence to take the vehicle local level frame to the moving vehicle body frame. Additional states for the vehicle body frame velocities, $[u, v, w]$, and angular rates, $[p, q, r]$, are also placed in the vehicle portion of the complete state vector.

$$\mathbf{x}_v = \underbrace{[x, y, z, \phi, \theta, \psi]}_{\text{position}} \underbrace{[u, v, w, p, q, r]}_{\text{velocity}}^\top = [\mathbf{x}_{v_p}, \mathbf{x}_{v_v}]. \quad (3.3)$$

To model the motion of the vehicle a constant velocity model is used. This simple model is sufficient to capture the slow dynamics of an ROV or AUV during a survey type mission. Although more complicated dynamic models can be used, this model has proven sufficient in experimentally demonstrating the bathymetric sub-mapping algorithm. The model is written in the form of a non linear deterministic function $\mathbf{f}(\mathbf{x}_v(t))$ that is perturbed by white process noise, \mathbf{w} , with zero mean and diagonal covariance \mathbf{Q} . The kinematic portion of the vehicle model relates the vehicle body frame velocities and angular rates to the time derivatives of the position variables expressed in the local level frame. The rotation matrix ${}^l_v\mathbf{R}(\phi, \theta, \psi)$ maps the body frame velocities to the local level frame velocities. The matrix $J(\phi, \theta, \psi)$ maps the body frame angular rates to the local level frame angular rates. Both ${}^l_v\mathbf{R}(\phi, \theta, \psi)$ and $J(\phi, \theta, \psi)$ depend non-linearly on $[\phi, \theta, \psi]$. The white process noise \mathbf{w} adds to the linear and angular acceleration terms and represents a probabilistic disturbance to the vehicle motion which accounts for the unmodeled vehicle thruster inputs that perturb the system from its current constant velocity. The complete continuous time model is expressed as

$$\dot{\mathbf{x}}_v(t) = \mathbf{f}(\mathbf{x}_v(t)) + \mathbf{w}(t) \quad (3.4)$$

$$= \underbrace{\begin{bmatrix} {}^l_v\mathbf{R}(\phi, \theta, \psi) \begin{bmatrix} u \\ v \\ w \end{bmatrix} \\ J(\phi, \theta, \psi) \begin{bmatrix} p \\ q \\ r \end{bmatrix} \\ \hline \mathbf{0}_{[6 \times 1]} \end{bmatrix}}_{\mathbf{0}_{[6 \times 1]}} + \underbrace{\begin{bmatrix} 0 \\ \vdots \\ \vdots \\ \vdots \\ \vdots \\ 0 \\ \hline \mathbf{w}_{[6 \times 1]} \end{bmatrix}}_{\mathbf{w}_{[6 \times 1]}} \quad (3.5)$$

where, $\mathbf{w}_{[6 \times 1]} = [w_1, w_2, w_3, w_4, w_5, w_6]^\top$. The details of the reference frame kinematics and additional information on underwater vehicle modeling can be found in Fossen [40].

3.2.1 Navigation sensor measurements

The formulation of the *continuous-discrete* [5] filter allows for asynchronous handling of navigation data produced by different sensors. The navigation sensor measurements are written as $\mathbf{z}[t_k]$ and the state prediction of the measurements is handled using non-linear measurement models of the form

$$\hat{\mathbf{z}}[t_k] = \mathbf{h}_n(\mathbf{x}_v[t_k], \mathbf{x}_{v(\text{sensor})}), \quad (3.6)$$

where $\mathbf{x}_{v(sensor)}$ is one of the static vehicle-to-sensor offsets drawn in Fig(3-2). These measurement models are implemented as *mixed-coordinate* functions that predict the sensor measurements in the individual sensor coordinate systems. The sensor measurements are assumed to be corrupted by a time independent zero mean Gaussian noise \mathbf{v} with covariance \mathbf{R} , where $\mathbf{E}[\mathbf{w}\mathbf{v}^\top] = 0$. Measurements for the navigation sensors are available at discrete times represented by t_k . The complete measurement model is then

$$\mathbf{z}[t_k] = \mathbf{h}_n(\mathbf{x}_v[t_k], \mathbf{x}_{v(sensor)}) + \mathbf{v}. \quad (3.7)$$

For this application the filter utilizes navigation measurements of the body frame velocities, surface relative depth, and three axis attitude. Although LBL position estimates will be used to evaluate the output of the sub-mapping algorithm, they are not incorporated directly into the filter. Thus, the vehicle state filtering within this framework is really dead reckoning integration.

3.3 Vehicle navigation

The EKF uses the continuous time prediction equations to move the state estimate forward incrementally from time t_{k-1} to t_k for the next navigation measurement or sonar ping. The first order EKF requires the Jacobian $\mathbf{F}_v(t)$ of the vehicle model $\mathbf{f}(\mathbf{x}_v(t))$ taken over all elements of the state vector. When the entire augmented state is considered, the delayed states are not affected by the vehicle model and their time derivatives are assigned to be zero. For N delayed states this is written as

$$\dot{\mathbf{x}}_{aug} = \begin{bmatrix} \dot{\mathbf{x}}_v \\ \mathbf{0}_{[6N \times 1]} \end{bmatrix} = \begin{bmatrix} \mathbf{f}(\mathbf{x}_v(t)) \\ \mathbf{0}_{[6N \times 1]} \end{bmatrix}. \quad (3.8)$$

Thus, the time derivative equations for the mean state vector and covariance take the form.

$$\dot{\mathbf{x}}_v(t) = \mathbf{f}(\mathbf{x}_v(t)) \quad (3.9a)$$

$$\dot{\mathbf{P}}_{aug}(t) = \begin{bmatrix} \mathbf{F}_v(t) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \mathbf{P}_{aug}(t) + \mathbf{P}_{aug}(t) \begin{bmatrix} \mathbf{F}_v(t) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}^\top + \begin{bmatrix} \mathbf{Q} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}. \quad (3.9b)$$

where,

$$\mathbf{F}_v(t) = \left. \frac{\partial \mathbf{f}(\mathbf{x}(t))}{\partial \mathbf{x}_v(t)} \right|_{\mathbf{x}_v(t)} \quad (3.10)$$

is the Jacobian of the vehicle model function evaluated at the current vehicle state. The structure of the covariance equation indicates that the prediction step will update the vehicle covariance and the cross covariances between the current vehicle state and the delayed states.

The state prediction using (3.9) is carried out numerically using a Runge-Kutta approximation. The integration produces the mean vehicle state $\bar{\mathbf{x}}_v^-$ and covariance \mathbf{P}_{aug}^- at time t_k . The update for the new state to incorporate the measurement at t_k is then accomplished

using the standard EKF discrete time update equations

$$W = P_{aug}^- H^\top \left[H P_{aug}^- H^\top + R \right]^{-1} \quad (3.11a)$$

$$\bar{\mathbf{x}}_{aug}[t_k] = \bar{\mathbf{x}}_{aug}^- + W[\mathbf{z}[t_k] - \mathbf{h}_n(\mathbf{x}_v[t_k], \mathbf{x}_{v(sensor)})] \quad (3.11b)$$

$$P_{aug}[t_k] = [\mathbf{I} - WH]P_{aug}^-[\mathbf{I} - WH]^\top + WRW^\top. \quad (3.11c)$$

Here the Jacobian H of the measurement equation $\mathbf{h}_n(\mathbf{x}_v[t_k], \mathbf{x}_{v(sensor)})$ taken over the entire state vector is needed. Similar to the vehicle model Jacobian, the required matrix contains zeros over the delayed state portion of the state vector. For a navigation sensor updating m states the Jacobian takes the form $H = [H_n, 0_{[m \times 6N]}]$, where

$$H_n = \left. \frac{\partial \mathbf{h}_n(\mathbf{x}[t_k])}{\partial \mathbf{x}_v[t_k]} \right|_{\mathbf{x}_v[t_k]}. \quad (3.12)$$

In accordance with the mixed-coordinate implementation the matrix R contains the appropriate measurement covariance for the navigation sensor expressed in that sensor's measurement frame. Typical values for the measurement covariances are show in Table 5.2 in Chapter 5.

3.4 Sub-map creation

The bathymetric sub-maps are created online as the navigation data filtering progresses. Each map contains points that are defined with respect to a local sub-map origin contained in delayed portion of the state vector. The delayed state poses consist of 6 DOF pose defined by

$$\mathbf{x}_{s_i} = [x, y, z, \phi, \theta, \psi]^\top. \quad (3.13)$$

The first sub-map origin is taken as the initial vehicle pose at the start of the filtering. Successive map origins are defined when the currently active sub-map meets a closure condition based on the structure of the terrain within the map or a limit on the navigation uncertainty. These conditions are described in Section 3.4.1. The state vector augmentation is completed as

$$\mathbf{x}_{aug} = [\mathbf{x}_v^\top, \mathbf{x}_{s_1}^\top, \dots, \mathbf{x}_{s_N}^\top]^\top \xrightarrow{\text{new map}} \mathbf{x}_{aug} = [\mathbf{x}_v^\top, \mathbf{x}_{s_1}^\top, \dots, \mathbf{x}_{s_N}^\top, \mathbf{x}_v^\top]^\top. \quad (3.14)$$

When new sub-maps are created additional rows and columns are also added to the covariance matrix. These new elements will be non-zero because the current state is correlated with all currently held delayed states. The growth of the covariance matrix is written in a block form as

$$P_{aug} = \left[\begin{array}{c|c} P_{\mathbf{x}_v \mathbf{x}_v} & P_{\mathbf{x}_v \mathbf{x}_s} \\ \hline P_{\mathbf{x}_s \mathbf{x}_v} & P_{\mathbf{x}_s \mathbf{x}_s} \end{array} \right] \xrightarrow{\text{new map}} P_{aug} = \left[\begin{array}{c|c|c} P_{\mathbf{x}_v \mathbf{x}_v} & P_{\mathbf{x}_v \mathbf{x}_s} & P_{\mathbf{x}_v \mathbf{x}_v} \\ \hline P_{\mathbf{x}_s \mathbf{x}_v} & P_{\mathbf{x}_s \mathbf{x}_s} & P_{\mathbf{x}_s \mathbf{x}_v} \\ \hline P_{\mathbf{x}_v \mathbf{x}_v} & P_{\mathbf{x}_v \mathbf{x}_s} & P_{\mathbf{x}_v \mathbf{x}_v} \end{array} \right]. \quad (3.15)$$

This covariance matrix augmentation allows the new sub-map origin to inherit the correct uncertainty of its position estimate and correlation with the other delayed states.

The raw data within each sub-map consists of a set of 3D points

$$\mathcal{M}_i = \{\mathbf{m}_i[1], \mathbf{m}_i[2], \dots, \mathbf{m}_i[n]\}, \quad (3.16)$$

where $\mathbf{m}_i[1 \dots n] = [x, y, z]^\top$. The points are created from the beam ranges using the position and attitude from the state vector at the ping time t_p and the composition sequence

$$\mathbf{x}_{s_i r_k} = (\ominus \mathbf{x}_{s_i} \oplus \mathbf{x}_{v_p}(t_p)) \oplus \mathbf{x}_{v_s} \oplus \mathbf{x}_{sb_k} \oplus \mathbf{x}_{b_k r_k} \quad (3.17a)$$

$$= \mathbf{m}_i[k]. \quad (3.17b)$$

The various pose vectors in this sequence are shown in Fig(3-2). The vehicle pose $\mathbf{x}_{v_p}(t_p)$ is extracted from the state vector once $\bar{\mathbf{x}}_v[t_p]^-$ is created using the prediction equations (3.9). The sonar transforms \mathbf{x}_{sb_k} and $\mathbf{x}_{b_k r_k}$ are taken from the beam-formed and processed data described in Chapter 2. The use of local map origins allows for easy manipulation of the sub-maps. Once the sub-maps are created they are considered rigid and any motion of a sub-map will be accomplished by updating the sub-map origin \mathbf{x}_{s_i} . The EKF algorithm which accommodates vehicle trajectory integration and sub-map creation is summarized in Algorithm 1.

Algorithm 1 EKF Loop The main *continuous-discrete* EKF loop alternates between handling navigation sensor data and sub-map creation. Algorithm 2 continues with the details of making a relative pose measurement between delayed states using the available sub-maps.

```

while Running the filter,  $t < t_{end}$  do
  Get times  $[t_{sonar}, t_{navigation}]$  to the next sonar and navigation measurement.
   $t^* = \min[t_{sonar}, t_{navigation}]$ .
  Predict the state for time  $t^*$ ,  $\bar{\mathbf{x}}_{aug}(t^*)^-$ ,  $P_{aug}(t^*)^-$ , using (3.9).
  if  $t^*$  from sonar then
    Extract  $\mathbf{x}_{v_p}$  and add the ping to currently open sub-map using (3.17).
    Call Algorithm 2 Sub-map handling.
  else
    Get the navigation measurement  $\mathbf{z}[t^*]$ 
    Predict the navigation measurement,  $\hat{\mathbf{z}}[t^*]$ , using (3.6)
    Update the state vector:  $\mathbf{x}_{aug}[t^*]^- \rightarrow \mathbf{x}_{aug}[t^*]$ ,  $P_{aug}[t^*]^- \rightarrow P_{aug}[t^*]$  using (3.11).
  end if
   $t = t^*$ 
end while

```

3.4.1 Dynamic map sizing

The primary assumption supporting the algorithmic generation of sub-maps is that the short time scale DR navigation produced by the filter integration is sufficiently accurate to

create sub-maps that represent the true world. Without external ground referenced position measurements the dead reckoning error will grow without bound and a sub-map will become arbitrarily distorted if it is not closed. This eventuality suggests there is an optimal size at which to break a terrain map and begin another. The selection of this break point involves the following trade offs.

- A sub-map should be small enough that it does not contain a significant amount of internal error or distortion caused by accumulated dead reckoning error. Since the maps are considered rigid once formed any internal distortions will only degrade the ability of match sections of terrain.
- A sub-map should be large enough to contain sufficient 3D information that it can be registered unambiguously to another map. Small maps will contain less internal distortion but be more difficult to register.

Given these criteria there are a few obvious limitations in applying this technique. First, the DR navigation must be reasonable enough that sub-maps can be made at all. Second, the sea floor can not be flat and featureless to the point where terrain matching is not possible. Fortunately, there is little interest in mapping such areas.

In an effort to algorithmically break and initiate sub-maps the characteristics of the sub-maps are monitored as sonar pings are added and the map size increases. To monitor the amount of 3D spatial information in a map there are several possible options. Ideally, a single statistic, that is not computationally expensive to compute, would indicate the potential for any sub-map to be registered correctly. The following list presents some possible measures.

Normal space occupancy

As an improvement to the performance of iterative closest point registration algorithms Rusinkiewicz [116] has proposed a normal based sampling technique where the input point cloud is down sampled by selecting points over the space of surface normal as uniformly as possible. The idea is to help the point matching solution by using all of the available constraining geometry. To convert this into a “registerability” test, the increase in normal space occupancy can be monitored as sonar ranges are added to a sub-map and the map’s geometry changes. Surface normals can be estimated using the principal component analysis (PCA) technique described in Appendix B.3. This test can be performed using a threshold for the number of occupants needed to consider a normal space bin occupied and a threshold for the total number of occupied bins that would suggest a good registration. A non-zero threshold on the number of occupants per bin helps suppress errant surface normals, caused by poor surface sampling and noisy data, from falsely populating the normal space. The image in Fig(3-4(a)) shows how the normal space occupancy changes as terrain is accrued into sample sub-maps.

Principal components

An inexpensive test for 3D structure can be made using the condition number of the covariance matrix

$$C = \sum_{i=1}^N [(\mathbf{m}_i - \bar{\mathbf{m}})(\mathbf{m}_i - \bar{\mathbf{m}})^\top] \quad (3.18)$$

where, $\mathbf{m}_i = [x, y, z]^\top$ is one of N points in the map \mathcal{M}_i and $\bar{\mathbf{m}}$ is the centroid of the point cloud. If this matrix is poorly conditioned the principal components of the sub-map describe an approximately planar surface, or the aspect ratio of the map is far from one to one. In general the condition number will be large when the map is started, decrease as the map gathers terrain and then increase again once the along track distance of the maps significantly exceeds the cross track width of the map. Although it is impossible to distinguish between these two cases using the condition number at one instant in time, a large condition number suggests a map with poor registration characteristics. The graph in figure Fig(3-4(b)) shows how the condition number will change for a selection of sample sub-maps.

Auto correlation

During the map creation the shape of the auto correlation surface produced by correlating a gridded version of a sub-map with itself can be used as a map breaking test. Ideally, the terrain within the map will contain enough relief that the sub-map will only correlate with itself for small displacements. This would suggest that additional maps covering the same area will have the same desirable property for registration. The correlation can be calculated with a gridded version M_i of map \mathcal{M}_i . The gridded surface should represent a height map of the form $z = f(x, y)$ on a regularly sampled mesh. The correlation surface \mathcal{C} is defined as

$$C(x, y) = \frac{1}{N_{x,y}} \sum_{\mathcal{I}_{x,y}} (M_i(x, y) - M_i(0, 0))^2 \quad (3.19)$$

where, $\mathcal{I}_{x,y}$ represents the set of $N_{x,y}$ the common bins between the map and the shifted map that overlap.

The correlation surface can be approximated using a quadratic surface fit of the form

$$C(x, y) \sim c + \begin{bmatrix} a \\ b \end{bmatrix}^\top \begin{bmatrix} x \\ y \end{bmatrix} + \begin{bmatrix} x \\ y \end{bmatrix}^\top \mathbf{H} \begin{bmatrix} x \\ y \end{bmatrix}, \quad \text{where } \mathbf{H} = \begin{bmatrix} d & e \\ e & f \end{bmatrix}. \quad (3.20)$$

The curvature of the correlation surface is described by the matrix \mathbf{H} . This matrix should be positive definite, well conditioned and have Eigen values that both exceed a threshold. When these conditions are satisfied the sub-map can be broken. Unfortunately, the gridding and correlation can be expensive to use as a continuous map monitoring test and is only used in a batch sense after increments of terrain are added to the sub-maps.

Growing navigation error

As an attempt to limit the mapping error internal to a sub-map, a map breaking test can be made to compare the vehicle navigation error with respect to the current sub-map origin and the placement error of the sonar range points relative to the vehicle. The purpose of this test is to break a sub-map when the vehicle positioning error grows larger than the error associated with the mapping sensor itself. This would suggest that the limiting factor in overall mapping accuracy is becoming the vehicle's lack of navigation rather than the sonar sensor itself.

To develop this test the placement of a ping into a sub-map relative to the sub-map origin using (3.17) can be rewritten as

$$\mathbf{x}_{s_i r_k} = (\ominus \mathbf{x}_{s_i} \oplus \mathbf{x}_{v_p}(t_p)) \oplus (\mathbf{x}_{v_s} \oplus \mathbf{x}_{s_b k} \oplus \mathbf{x}_{b_k r_k}) \quad (3.21a)$$

$$= \mathbf{x}_{s_i v} \oplus \mathbf{x}_{v r_k}. \quad (3.21b)$$

Here $\mathbf{x}_{s_i v}$ represents the vehicle pose relative to the sub-map origin and $\mathbf{x}_{v r_k}$ represents the placement of the range point relative to the vehicle. The combined uncertainty for point placement into the current sub-map can then be written as,

$$\mathbf{P}_{\mathbf{x}_{s_i r_k} \mathbf{x}_{s_i r_k}} = \mathbf{J}_{s_i r_k \oplus} \begin{bmatrix} \mathbf{P}_{\mathbf{x}_{s_i v} \mathbf{x}_{s_i v}} & \mathbf{0}_{6 \times 6} \\ \mathbf{0}_{6 \times 6} & \mathbf{P}_{\mathbf{x}_{v r_k} \mathbf{x}_{v r_k}} \end{bmatrix} \mathbf{J}_{s_i r_k \oplus}^\top \quad (3.22a)$$

$$= \mathbf{J}_{s_i r_k \oplus 1} \mathbf{P}_{\mathbf{x}_{s_i v} \mathbf{x}_{s_i v}} \mathbf{J}_{s_i r_k \oplus 1}^\top + \mathbf{J}_{s_i r_k \oplus 2} \mathbf{P}_{\mathbf{x}_{v r_k} \mathbf{x}_{v r_k}} \mathbf{J}_{s_i r_k \oplus 2}^\top \quad (3.22b)$$

where, $\mathbf{P}_{\mathbf{x}_{s_i v} \mathbf{x}_{s_i v}}$ captures the sub-map-to-vehicle uncertainty, $\mathbf{P}_{\mathbf{x}_{v r_k} \mathbf{x}_{v r_k}}$ captures the vehicle-to-point uncertainty and the cross correlations are zero. The Jacobian $\mathbf{J}_{s_i r_k \oplus} = [\mathbf{J}_{s_i r_k \oplus 1}, \mathbf{J}_{s_i r_k \oplus 2}]$ comes from the composition operation $\mathbf{x}_{s_i v} \oplus \mathbf{x}_{v r_k}$, see Appendix A.2.1. The ping placement error has been calculated and plotted for the individual ranges in two consecutive sub-maps in Fig(3-3). This figure shows that the error internal to a sub-map will grow away as the vehicle moves away from the sub-map origins.

Equation (3.22) indicates that the sub-map-to-vehicle and vehicle-to-point errors are additive. Although strictly speaking any vehicle pose error will combine with the range point placement error, it is more realistic to consider the situation where the vehicle position error begins to dominate the range placement error. As a test we can compute the vehicle error online during the filtering process and pre-compute a threshold value for the vehicle-to-point error using some typical error values for a sonar measurement.

For online computation $\mathbf{P}_{\mathbf{x}_{s_i v} \mathbf{x}_{s_i v}}$ can be created using the Jacobian $\ominus \mathbf{J}_{s_i v \oplus}$ associated with the relative pose operation $\ominus \mathbf{x}_{s_i} \oplus \mathbf{x}_{v_p}$ between the current vehicle position and the sub-map origin. This uncertainty can be written as

$$\mathbf{P}_{\mathbf{x}_{s_i v} \mathbf{x}_{s_i v}} = \ominus \mathbf{J}_{s_i v \oplus} \mathbf{P}_{aug \ominus} \mathbf{J}_{s_i v \oplus}^\top. \quad (3.23)$$

Within this covariance matrix the most important components are those associated with the vehicle position in $[x, y, z]$. The attitude errors can be lumped in with the sonar errors as they will directly contribute to the error volume the sonar range points are placed in.

The attitude errors are also related to measurements from obtainable references like North and the gravity vector, and not subject to the large scale error growth of the vehicle position estimates.

Pre-computing a sonar mapping error threshold requires an error statistic for the sonar range accuracy, an error measure related to the sonar beam pattern, as well as an assumed beam range. These statistics will affect the transformations \mathbf{x}_{vs} , \mathbf{x}_{sb} and \mathbf{x}_{br} between the vehicle and the range point placement. If the errors are approximated as Gaussian for computational convenience the respective covariance matrices take the form

$$\mathbf{P}_{\mathbf{x}_{vs}\mathbf{x}_{vs}} = \text{diag}([0, 0, 0, \sigma_{vr}^2, \sigma_{vp}^2, \sigma_{vh}^2]) \quad (3.24a)$$

$$\mathbf{P}_{\mathbf{x}_{sb_k}\mathbf{x}_{sb_k}} = \text{diag}([0, 0, 0, \sigma_a^2, \sigma_w^2, 0]) \quad (3.24b)$$

$$\mathbf{P}_{\mathbf{x}_{b_k r_k}\mathbf{x}_{b_k r_k}} = \text{diag}([0, 0, \sigma_r^2, 0, 0, 0]). \quad (3.24c)$$

where, σ_{vr}^2 , σ_{vp}^2 and σ_{vh}^2 represent the vehicle frame attitude errors. Typical values for these are given in Table 5.2. For the sonar σ_a^2 , σ_w^2 and σ_r^2 are conservative estimates for the standard deviations of an individual beam angle, the for-aft beam width and the range error respectively. Using (3.24) the vehicle-to-point uncertainty $\mathbf{P}_{\mathbf{x}_{vr_k}\mathbf{x}_{vr_k}}$ is written as

$$\mathbf{x}_{vb_k} = \mathbf{x}_{vs} \oplus \mathbf{x}_{sb_k} \quad (3.25a)$$

$$\mathbf{P}_{\mathbf{x}_{vb_k}\mathbf{x}_{vb_k}} = \mathbf{J}_{vb\oplus} \begin{bmatrix} \mathbf{P}_{\mathbf{x}_{vs}\mathbf{x}_{vs}} & \mathbf{0}_{6 \times 6} \\ \mathbf{0}_{6 \times 6} & \mathbf{P}_{\mathbf{x}_{sb_k}\mathbf{x}_{sb_k}} \end{bmatrix} \mathbf{J}_{vb\oplus}^\top, \quad (3.25b)$$

followed by

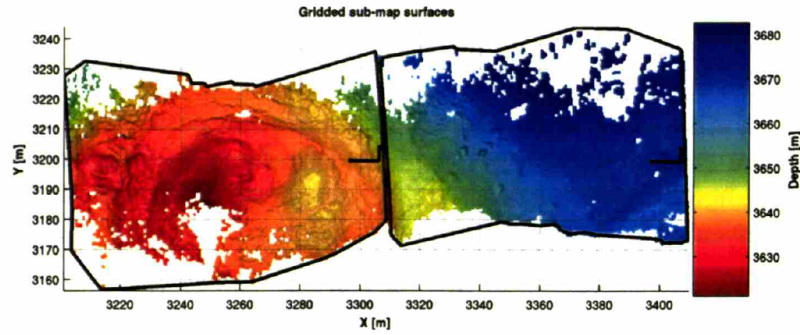
$$\mathbf{x}_{vr_k} = \mathbf{x}_{vb_k} \oplus \mathbf{x}_{b_k r_k} \quad (3.26a)$$

$$\mathbf{P}_{\mathbf{x}_{vr_k}\mathbf{x}_{vr_k}} = \mathbf{J}_{vr\oplus} \begin{bmatrix} \mathbf{P}_{\mathbf{x}_{vb_k}\mathbf{x}_{vb_k}} & \mathbf{0}_{6 \times 6} \\ \mathbf{0}_{6 \times 6} & \mathbf{P}_{\mathbf{x}_{b_k r_k}\mathbf{x}_{b_k r_k}} \end{bmatrix} \mathbf{J}_{vr\oplus}^\top. \quad (3.26b)$$

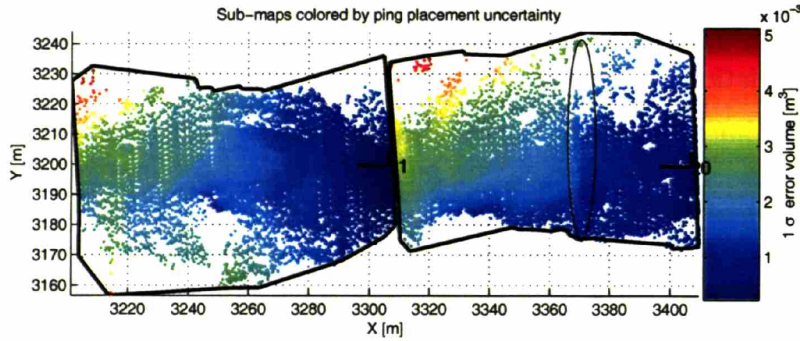
An algorithmic test to close a sub-map can be implemented by comparing a determinants of the of upper left 3×3 blocks extracted from $\mathbf{P}_{\mathbf{x}_{s_i v}\mathbf{x}_{s_i v}}$ and $\mathbf{P}_{\mathbf{x}_{vr_k}\mathbf{x}_{vr_k}}$. When $\det(\mathbf{P}_{\mathbf{x}_{s_i v}\mathbf{x}_{s_i v}})$ exceeds $\det(\mathbf{P}_{\mathbf{x}_{vr_k}\mathbf{x}_{vr_k}})$ the navigation error is dominating the ping placement error and the sub-map can be broken and a new map started. Fig(3-4(c)) show examples of this test for several example sub-maps.

Additional map checks

Aside from the map breaking conditions described above other simple checks are also employed during the sub-map generation. The minimum and maximum map size are limited and tested by an online by a calculation of map area based on a bounding border of the sub-map point cloud. The maximum map size limit serves to bound potential errors caused by angular error in the registration process. Although smaller internal sections of the sub-maps will overlap with other sub-maps during registration, the resulting transform is applied to the entire map as a rigid body. As a result the end points of large maps with high aspect ratios can be displaced significantly due to a lever arm effect. Thus, it is desirable to limit the total sub-maps size independent of the other end conditions. To support a generaliza-



(a) Gridded terrain



(b) Pings colored by uncertainty

Figure 3-3: Sonar measurement error growth internal to a sub-map. (a) Two consecutive sub-maps created by the vehicle moving right to left. (b) The growth of total mapping error internal to the sub-maps. Note that the uncertainty “resets” when the second map is started and $P_{\mathbf{x}_{s_i u} \mathbf{x}_{s_i v}}$ “starts over” from the new map origin. The ellipse in (b) indicates a step change where the vehicle’s forward progress stopped momentarily and position uncertainty continued to grow. The mapping error for each ping is calculated as a 1σ uncertainty volume using the square root of the determinant of the upper left 3×3 block of $P_{\mathbf{x}_{s_i r_k} \mathbf{x}_{s_i r_k}}$. The values $\sigma_a = .1^\circ$, $\sigma_w = .3^\circ$, $\sigma_r = .1M$ were used in the calculation of $P_{\mathbf{x}_{s_i r_k} \mathbf{x}_{s_i r_k}}$.

tion to 3D mapping the map area can be calculated after orthographically projecting the map point cloud onto a plane described by the normal vector from a map-wide PCA, see Appendix B.2. Additionally, the map aspect ratio can be monitored to serve as a map breaking condition after the minimum map area condition is satisfied.

Examples of online calculation of the sub-map properties described above are shown in Fig(3-4).

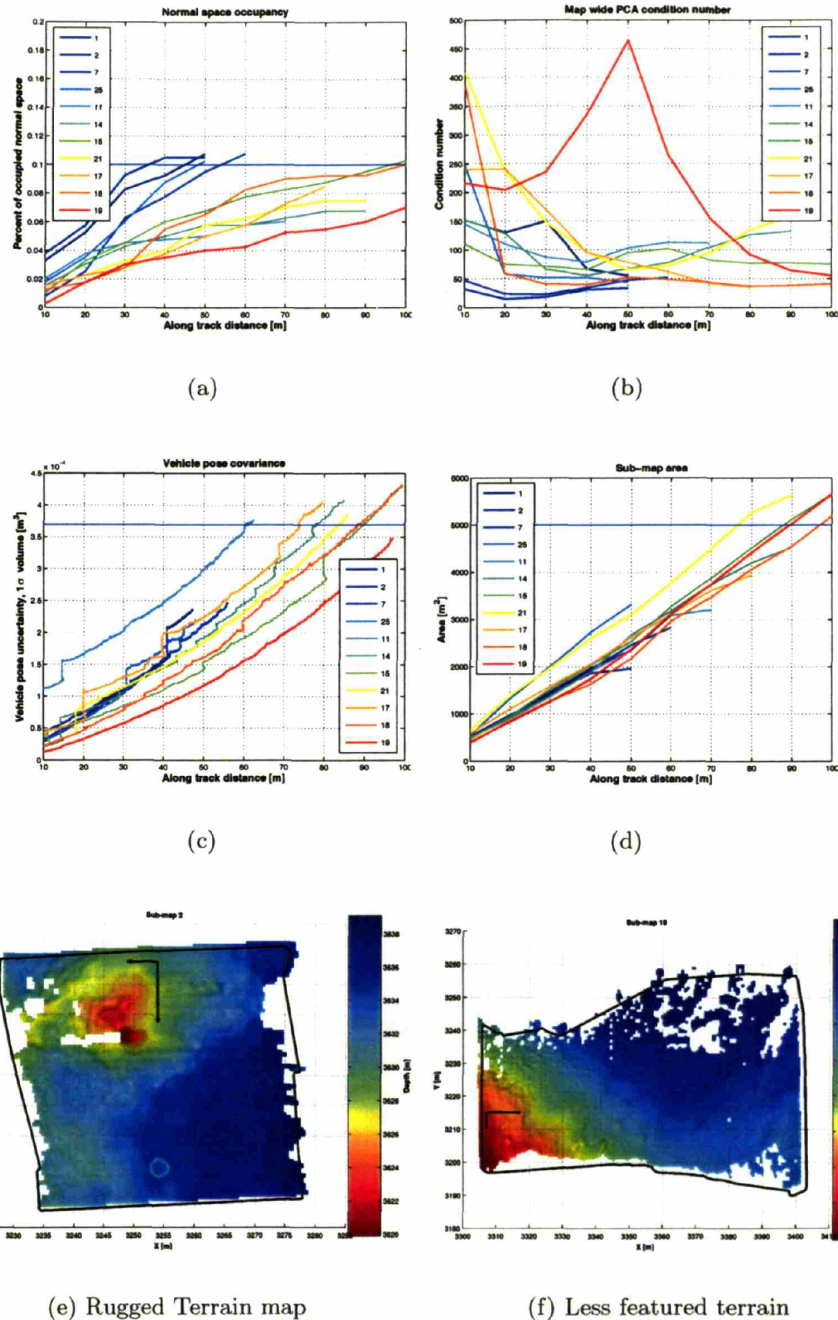


Figure 3-4: Examples of the changing sub-map properties. The changing properties in figures (a), (b), (c) and (d) are plotted as a function along track mapping distance. Two example maps are shown in (e) and (f). The along track direction is indicated by the long axis arrow. (a) The normal space occupancy condition was set to .1 and several well featured maps were broken when it was exceeded. (b) The principal component condition number shows considerable variability but does highlight approximately planar terrain. The first 50 meters of map 19 (e) are very planar. (c) The vehicle position uncertainty as calculated by (3.23), increases until it exceeds the sonar mapping uncertainty threshold. Vertical sections in the uncertainty lines indicate that the ROV held position momentarily and continued to accumulate position uncertainty. (d) The increasing map area is also monitored.

3.5 Relative pose measurements

Terrain matching and registering sub-maps will create relative pose constraints between the delayed states that reduce the growing uncertainty created by the DR process model. At a map closure, links to previously defined maps can be hypothesized in a straight forward manner. Since the delayed states represent the sub-map origins in a single coordinate frame, checks can be made for overlap using the intersection of the sub-map borders. This can be done in an all-against-all manner to check for all possible links, including new links between previously delayed states that may now be possible because of the trajectory refinement, at $O(N^2)$ cost. Experimentally, a simpler test for the current map against all prior maps, $O(N)$ computation, has worked well. This is due to the large sub-map size relative to magnitude of pose uncertainty. Very few links should need to be re-established if the survey pattern allows for continual small scale adjustment of the map origins. Simple checks on the size and shape of the intersection region between sub-map borders are also used to avoid the risk of ill-conditioned matches being made, such as maps which overlap along a long thin strip rather than an approximately square area. As a ad-hoc method of incorporation position uncertainty into the link proposal, the borders of the sub-maps can be “dilated” in proportion to the $[x, y]$ pose uncertainty of the map origins. The similar problem of uncertainty based link proposal has been addressed in visual based systems [36, 110] where problem is inherently harder due to the more equal relation between camera foot print size on the sea floor and the vehicle pose uncertainty.

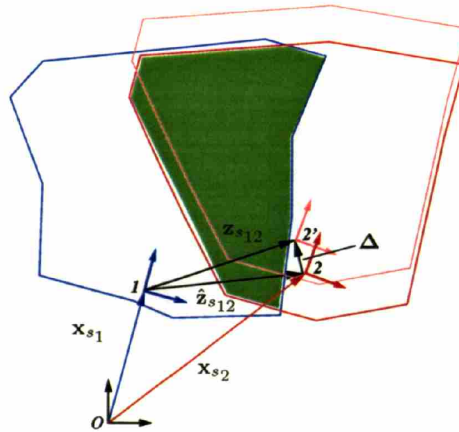


Figure 3-5: Vector diagram showing a sub-map registration measurement. Two overlapping sub-maps (red) and (blue) are shown with their locations in the local level frame indicated by x_{s1} and x_{s2} . The vector \hat{z}_{s12} indicates the relative pose between the origins as predicted by the EKF. The addition correction vector, Δ , is determined by the map registration and used to create the actual terrain relative measurement z_{s12} .

The vector diagram in Fig(3-5) illustrates the measurement between two overlapping sub-maps. The measurement model to predict a relative link given the augmented state vector is the *tail-to-tail* operation defined in Appendix A.2.3. This operation is written in

the form of a non-linear measurement model

$$\hat{\mathbf{z}}_{s_{ij}} = \mathbf{h}_s(\mathbf{x}_{s_i}, \mathbf{x}_{s_j}) \quad (3.27a)$$

$$= \ominus \mathbf{x}_{s_i} \oplus \mathbf{x}_{s_j}, \quad (3.27b)$$

whose arguments are the sub-map origins. The accompanying Jacobian for the measurement with respect to the entire state vector will be a sparse matrix

$$\mathbf{H}_{s_{ij}} = \left[0, \frac{\partial \mathbf{h}_s(\mathbf{x}_{s_i}, \mathbf{x}_{s_j})}{\partial \mathbf{x}_{s_i}}, 0, \frac{\partial \mathbf{h}_s(\mathbf{x}_{s_i}, \mathbf{x}_{s_j})}{\partial \mathbf{x}_{s_j}}, 0 \right]. \quad (3.28)$$

This Jacobian can be used to predict the uncertainty estimate of the relative pose measurement as

$$\mathbf{P}_{\mathbf{x}_{s_{ij}} \mathbf{x}_{s_{ij}}} = \mathbf{H}_{s_{ij}} \mathbf{P}_{aug} \mathbf{H}_{s_{ij}}^\top. \quad (3.29)$$

If a terrain measurement is made by the registration process it can be incorporated into the filter in a similar manner as a navigation measurement. The actual measurement $\mathbf{z}_{s_{ij}}$ will also have an accompanying measurement covariance estimate $\mathbf{R}_{\mathbf{z}_{s_{ij}} \mathbf{z}_{s_{ij}}}$. The prediction equations (3.9) are used to obtain $\bar{\mathbf{x}}_{aug}^-[t_p]$ for the time when the last ping is placed in the sub-map and the map is closed. The update equations 3.11 are then used to update the entire state vector after replacing $\mathbf{h}_n(\cdot)$ with $\mathbf{h}_s(\mathbf{x}_{s_i}, \mathbf{x}_{s_j})$, \mathbf{H} with $\mathbf{H}_{s_{ij}}$ and \mathbf{R} with $\mathbf{R}_{\mathbf{z}_{s_{ij}} \mathbf{z}_{s_{ij}}}$. The on-line sub-map registration is summarized in Algorithm 2.

3.6 Summary

This chapter has described an EKF framework to both estimate the vehicle trajectory using navigation data and incorporate terrain relative measurements between previously visited vehicle states. The details for constructing point cloud terrain maps on-line were discussed and several tests to evaluate the potential for the terrain within a map to be registered were given. These tests serve as break points for the completion of a sub-map and the start of another. It was also shown that an estimate of the point placement error internal to the sub-map can be calculated from the filter covariances. This error can also be used as a test to determine when the vehicle position error begins to dominate the error in point placement due to the sonar inaccuracies alone. Although map breaking tests involve tunable thresholds that have to be set in accordance with the data set at hand, the tests themselves indicate that map breaking conditions can be formulated to use the mapping data incrementally in an automatic fashion. Lastly, the measurement model for the relative pose measurements between delayed states was given. These measurements will be produced by the registration procedure and used to constrain the potentially large scale growth in due to the DR vehicle model.

Algorithm 2 Sub-map handling The sub-map handling process monitors the sub-map properties as the maps are created, closes maps and handles registration to previous maps. Algorithm 3 will continue when overlapping sub-maps are determined and may return with the delayed state measurement $\mathbf{z}_{s_{ij}}$. The measurement acceptance test mentioned in this algorithm will be described in Section 5.3.2.

```

Determine terrain properties of  $\mathcal{M}_i$ , Section 3.4.1.
if map  $\mathcal{M}_i$  gets closed at time  $t^*$  then
  Create delayed state  $\mathbf{x}_{s_{i+1}}$  for the next map  $\mathcal{M}_{i+1}$ 
  Find the border contour  $\mathcal{C}_i$  of  $\mathcal{M}_i$ 
  for  $j = 1 \dots i - 1$  do
    if  $\mathcal{C}_i \cap \mathcal{C}_j$  passes intersection tests then
      Try and register  $\mathcal{M}_i$  and  $\mathcal{M}_j$ , Call Algorithm 3 Map Registration
      if Algorithm 3 returns a measurement  $\mathbf{z}_{s_{ij}}$  then
        Compute  $\hat{\mathbf{z}}_{s_{ij}}$  and  $\mathbf{H}_{s_{ij}}$  with (3.27) and (3.28)
        Update,  $\mathbf{x}_{aug}[t^*]^- \rightarrow \mathbf{x}_{aug}[t^*]$ ,  $\mathbf{P}_{aug}[t^*]^- \rightarrow \mathbf{P}_{aug}[t^*]$ , using (3.11).
        Do measurement acceptance test, (See Section 5.3.2)
        if Passes test then
          Keep update state vector
        else
          Revert the state and covariance back to  $\mathbf{x}_{aug}[t^*]^-$  and  $\mathbf{P}_{aug}[t^*]^-$ .
          Flag this map pair as a bad match.
        end if
      end if
    end if
  end for
end if

```

Chapter 4

Terrain registration

4.1 Introduction

This chapter outlines the techniques developed to pairwise register bathymetric sub-maps. The registration is performed to obtain the 6 DOF transformation that aligns two small sections of bathymetry created using short term vehicle navigation. The alignment process is divided into two steps. The first step uses a 2D correlation to find the $[x, y]$ translation which best aligns the sub-maps. The second step uses an iterative closest point algorithm to determine a final 6 DOF transformation. It is shown that incorporating sonar return attributes into the ICP point selection step improves the matching convergence. The results of the registration are evaluated using an error metric based on the pairwise error between corresponding points from each map. Lastly, the registration error estimates required to incorporate the relative transform as a measurement for the delayed state EKF are discussed.

4.2 Relative position measurements

4.2.1 Selection of methods

The proposed map matching is accomplished using correlation and ICP methods that attempt to minimize surface wide errors in map registration rather than the errors between specific features that have been extracted from the surfaces themselves. Experimentation has indicated that identifying and utilizing geometric features extracted from acoustically mapped terrain data is problematic. The view point dependent nature of acoustic scattering will cause the same feature imaged from multiple vantage points to appear differently and have different error statistics. (See Chapter 2). This fact violates the primary assumption that feature-based registration methods make regarding features that are invariant to view-point [64, 65]. An additional motivation for choosing featureless methods is the desire to register the individual maps into a single consistent point cloud. Ideally, the registration will recover the transform i_jT that allows \mathcal{M}_j to be combined with \mathcal{M}_i and describe a single surface \mathcal{C} in the overlapping area

$$\mathcal{C} = \mathcal{M}_i \cup {}^i_jT\mathcal{M}_j. \quad (4.1)$$

Methods that utilize a set of surface features during registration are susceptible to proposing transforms that best match the feature points at the expense of errors distributed over the entire surface. Although feature based methods have been proposed for underwater navigation [86, 101, 113, 123, 124], their utility has not been demonstrated for creating a consistent terrain representation.

The map alignment used here is broken down into a two step process which offers robustness and minimizes the chance for obtaining a false match. The 2D correlation measurements are used to determine the large scale shifts that are possible given the potential for large $[x, y]$ positioning errors. Compared to $[x, y]$ translational errors, the maps relative orientation and depth are very well known by direct measurement using navigation sensors. This fact makes 2D correlation an effective alignment method. The ICP step serves to refine the 6 DOF transform between the maps and uses the correlation result as an initial guess. This refinement is meant to better the correlation measurement in $[x, y]$ and also correct for small changes in depth and attitude. Typical errors related to marine attitude sensors suggest that the angular refinements will be on the order of a few degrees. The quality of this initial guess, as aided by correlation and navigation sensors, provides the ICP algorithm with a good starting point. Of greater concern to the ICP algorithm is the quality of the point cloud data. In comparison to laser scanner data, which is known to register well with ICP methods, acoustically mapped terrain will have a higher level of noise relative to the feature scales within the point clouds. It should also be suspected that the acoustic maps have biases in them related to the finite beam width of the sensor.

4.2.2 Sub-mapping specifics

The map matching process has been developed to perform pairwise registrations using the common area between two overlapping sub-maps created with the EKF filtering algorithm. The inputs to the registration process are

- the point clouds sets \mathcal{M}_i and \mathcal{M}_j , each described in their own local reference frames,
- the initial guess for the relative transform between the reference frame, $\mathbf{x}_{s_{ij}}$, created with equation (3.27),
- and an uncertainty estimate for the transform, $P_{s_{ij}} = H_s P_{aug} H_s^\top$, obtained from the filter covariance and the Jacobian of the measurement function (3.27).

For generality the registration can be performed in either of the input map reference frames, and the fact that the EKF state vector contains the sub-map pose locations in a common coordinate frame can be ignored. Thus, as an initial step, the predicted relative pose $\mathbf{x}_{s_{ij}}$ is used to transform the point cloud in map \mathcal{M}_j into reference frame i . This is expressed using the transform operator ${}^i_j T$ (described in Appendix A) as ${}^i \mathcal{M}_j = {}^i_j T \mathcal{M}_j$. Two example sub-maps that have been transformed into a single coordinate frame are shown in Fig(4-1). The sub-map registration will then produce the residual vector Δ , shown in Fig(3-5), that corrects the EKF prediction of the sub-map relative pose. The relative pose measurement returned to the filter is constructed by

$$\mathbf{z}_{s_{ij}} = \Delta \oplus \mathbf{x}_{s_{ij}}. \quad (4.2)$$

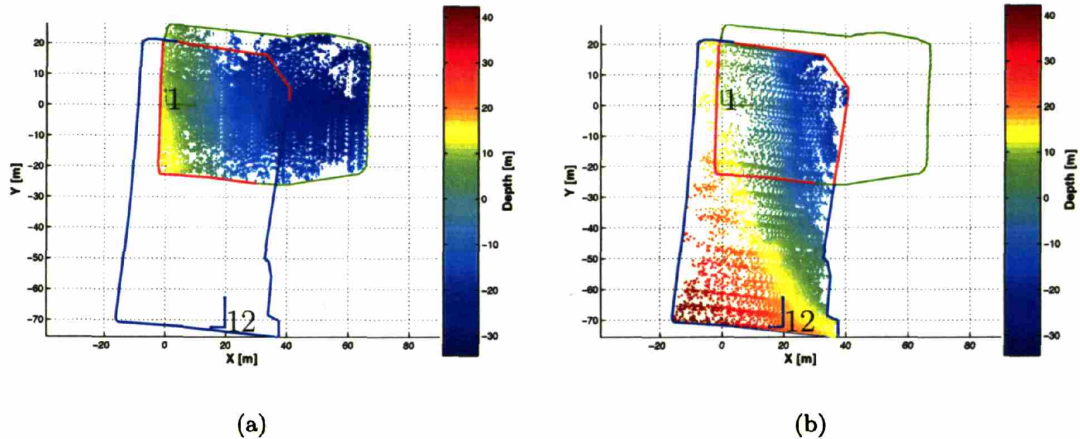


Figure 4-1: Sub-maps can be projected into a common coordinate frame prior to registration. Here map 12 is projected into frame 1. The red bordered region will be used for registration and the $[x, y, z]$ points are color coded by height in z direction. The number of displayed points has been down sampled from the actual number by 5.

4.3 Methods

4.3.1 Correlation

The initial sub map alignment step uses a 2D correlation to determine the $[x, y]$ translation that best aligns the two sub-maps. Correlation has proven to be a robust alignment method with a low probability of generating false matches due to local minima. The correlation measurement is generated by first gridding the sub-map point clouds to create a depth image sampled on a regular grid. The gridding approach used for this application is very general and described in Appendix B. The sub-map \mathcal{M}_i will produce the gridded surface M_i , and a single grid node on M_i is denoted by $M_i[k]$.

The correlation between the depth images is calculated by displacing the gridded surface iM_j with respect to the fixed grid M_i and computing the normalized sum of the squared node-wise depth disparities over the set of common grid nodes $\mathcal{I}_{x,y}$,

$$C(x, y) = \frac{1}{N_{x,y}} \sum_{k \in \mathcal{I}_{x,y}} \frac{1}{\sigma_{x,y}^2[k]} (M_i[k] - {}^iM_j(x, y)[k])^2. \quad (4.3)$$

The set $\mathcal{I}_{x,y}$, which indexes the common grid cells, is recalculated for each $[x, y]$ shift to account for irregularly shaped borders and missing data in the overlapping region. The number of common grid cells at each shift is $N_{x,y}$. The uncertainty measure $\sigma_{x,y}^2[k]$ represents the uncertainty of the depth differences assuming the node depths are random variables. A conservative estimate for this value can be made by ignoring the cross correlation of the sub-map origins and using

$$\sigma_{x,y}^2[k] = {}^i\sigma^2[k] + {}^j\sigma_{x,y}^2[k], \quad (4.4)$$

where ${}^i\sigma^2[k]$ and ${}^j\sigma_{x,y}^2[k]$ are values for the depth uncertainty in each map taken from the

range points nearest to the grid nodes. The point depth uncertainties are calculated using (3.22). Incorporating this uncertainty measure into (4.3) serves to capture the growing nature of the mapping uncertainty internal to each sub-map.

The correlation measurement for the alignment transform is then

$$\Delta_c = \arg \min_{[x,y]} C(x, y). \quad (4.5)$$

To check that Δ_c does in fact correspond to a local minima, a quadratic surface of the form shown in (3.20) can be fit to $C(x, y)$ in the neighborhood of the calculated minimum. A minimum is verified by a positive definite Hessian matrix H .

The size of the window over which x and y are varied for the correlation can be set using the uncertainty estimate $P_{\mathbf{x}_{s_{ij}} \mathbf{x}_{s_{ij}}}$ for the initial estimate of the relative transform $\mathbf{x}_{s_{ij}}$ provided by the EKF. An $[x, y]$ 99% χ^2 confidence ellipse can be calculated from the upper left 2 block of $P_{\mathbf{x}_{s_{ij}} \mathbf{x}_{s_{ij}}}$. A sample correlation measurement is summarized in Fig(4-2).

Uncertainty estimation

To use the correlation measurement in the EKF a first order estimate of its uncertainty is needed. To develop this the correlation measurement, (4.3), can be considered proportional to a log likelihood expression for the aligning $[x, y]$ shift. This likelihood would consider the individual grid cell depth disparities as zero mean independent Gaussian measurements and be written as

$$L(\mathbf{y}|\mathbf{x}(x, y)) = \frac{1}{\sqrt{(2\pi)^N \det(\Lambda)}} \exp^{-\frac{1}{2}(\mathbf{y}-\mathbf{x}(x, y))^\top \Lambda^{-1}(\mathbf{y}-\mathbf{x}(x, y))}, \quad (4.6)$$

where the vector \mathbf{y} represents depths in the fixed map \mathcal{M}_i and $\mathbf{x}(x, y)$ represents a measurement of these depths as a function of the translation $[x, y]$. The matrix Λ is the diagonal matrix of node wise depth uncertainties, $\Lambda = \text{diag}[\sigma_{x,y}^2[1], \dots, \sigma_{x,y}^2[N]]$. In the neighborhood of maximum, this likelihood can be approximated by a Gaussian over the x and y translations directly. This will generate an estimate of the measurement uncertainty that is related to the curvature of the likelihood. The approximation can be written as

$$C(x, y) = -\ln L(\mathbf{y}|\mathbf{x}(x, y)) \quad (4.7a)$$

$$\sim -\ln \left[\frac{1}{c} \exp^{-\frac{1}{2\sigma_x^2\sigma_y^2(1-\rho^2)} \left[\left(\frac{x-m_x}{\sigma_x}\right)^2 + 2\rho\left(\frac{x-m_x}{\sigma_x}\right)\left(\frac{y-m_y}{\sigma_y}\right) + \left(\frac{y-m_y}{\sigma_y}\right)^2 \right]} \right] \quad (4.7b)$$

where, the minimum occurs as $\Delta_c = [m_x, m_y]^\top$ and $c = 2\pi\sigma_x\sigma_y\sqrt{1-\rho^2}$. The parameters σ_x , σ_y , m_x , m_y and ρ can be solved for using a least squares fit with the values of $C(x, y)$ around the minimum. The covariance matrix for the correlation measurement is then written as

$$\mathbf{R}_c = \begin{bmatrix} \sigma_x^2 & \rho\sigma_x\sigma_y \\ \rho\sigma_x\sigma_y & \sigma_y^2 \end{bmatrix}. \quad (4.8)$$

The initial guess for the uncertainty parameters in the least square solution can be come for a simpler surface fit around the minimum of the form $C(x, y) \sim [x, y]^\top \mathbf{H}[x, y]^\top$ and

solving $R_c \sim H^{-1}$. Similar derivations for terrain correlation uncertainty can be found in [104–106]. In practice the correlation surface is usually represented well by a quadratic and this approximation seems reasonable. Solving for m_x and m_y is also used to generate a sub grid cell estimate of the minimum location. Most importantly, the Eigen vectors of P_Δ are able to capture the orientation of the uncertainty, Fig(4-2(f)).

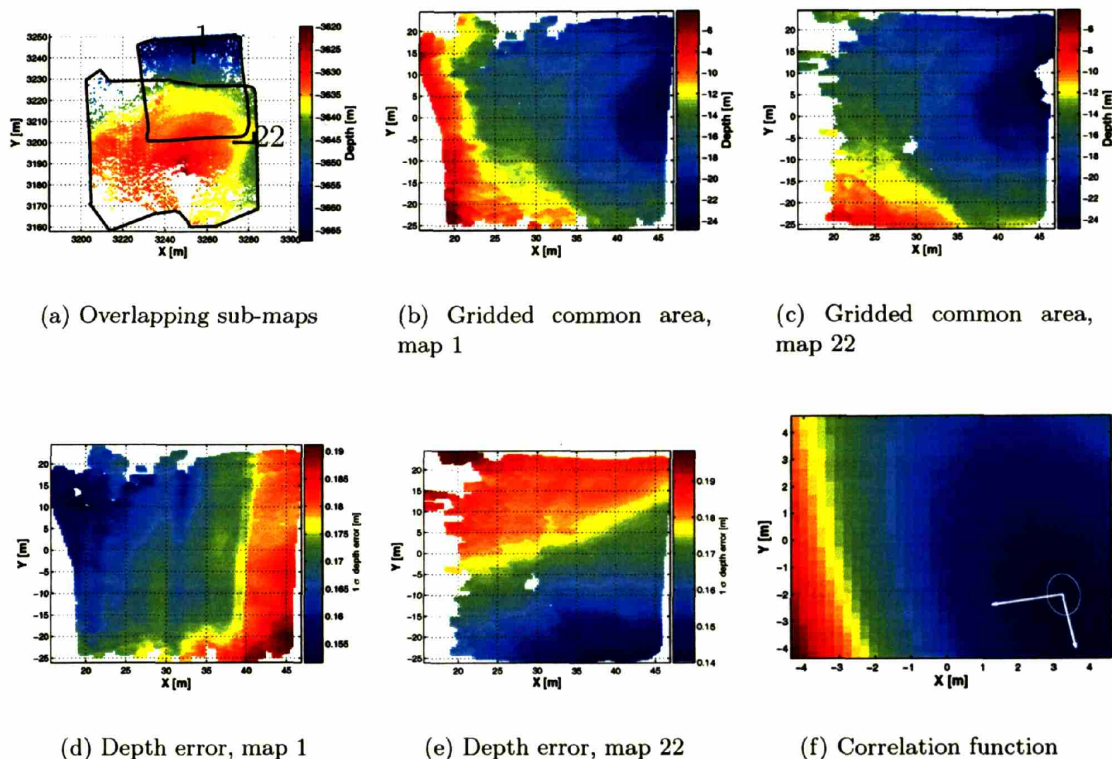


Figure 4-2: Sample correlation measurement. (a) Two overlapping sub-maps. (b,c) The gridded depth images of the overlapping region for each map, referenced to frame 1. (d,e) Depth uncertainty of each map. (f) The correlation result $C(x,y)$ shown with the principal axis of the uncertainty matrix P_Δ indicated by the vectors located as the minimum correlation score.

4.3.2 Iterative closest point matching

The ICP registration step is used to refine the 2D registration. This step should reduce the total error between the sub-map surfaces and offer robustness to errors in heading, pitch, and roll that affect the sub-map origins. The point-to-point and point-to-plane ICP algorithms were evaluated to see which performed better for acoustically created maps. Both methods attempt to minimize a distance measure calculated for a subset of nearest neighbor points between selected from the overlapping region of the two point clouds. The application of these methods for bathymetry has not been well tested, but preliminary results suggest reliable registration can be performed when using sonar data [25, 114, 134]. Unfortunately, generalizations about the applicability to sonar mapping are difficult to make due to the strong dependence on data quality, point selection, and error metrics [116]. The performance

evaluation presented here used repeated trials from randomized starting locations around a nominal solution for many sub-map pairs. The ability to repeatedly register a sub-map pair to a single pose solution from many different starting locations indicates the robustness and consistency of the registration.

The following, somewhat standard, ICP augmentations were used for both the point-to-point and point-to-plane implementations.

- Points positioned on or near the boundaries between maps were not used as point pairs.
- After point pair selection, those pairings with link lengths greater than two standard deviations from the mean link length were rejected.
- The point selection used between 500 and 2000 points from the common regions.
- Paired points with surface normals differing by more than 45° were rejected.

The initial location of map \mathcal{M}_j at the start of the ICP registration is produced from

$${}^i\mathcal{M}_j = {}^i T_c \mathcal{M}_j, \quad (4.9)$$

where the transform ${}^i T_c$ relates the composition of the correlation measurement and the initial relative pose guess from the EKF, ${}^i T_c \rightleftharpoons \Delta_c \oplus \mathbf{x}_{s_{ij}}$.

Point-to-point

The cost function for the point-to-point method [10] penalizes the sum of squared distances for the selected point pairings and can be written as

$${}^i T = \arg \min_T \sum_{i=1}^k w_k \| T \mathbf{m}_j[k] - \mathbf{m}_i[k] \|^2, \quad (4.10)$$

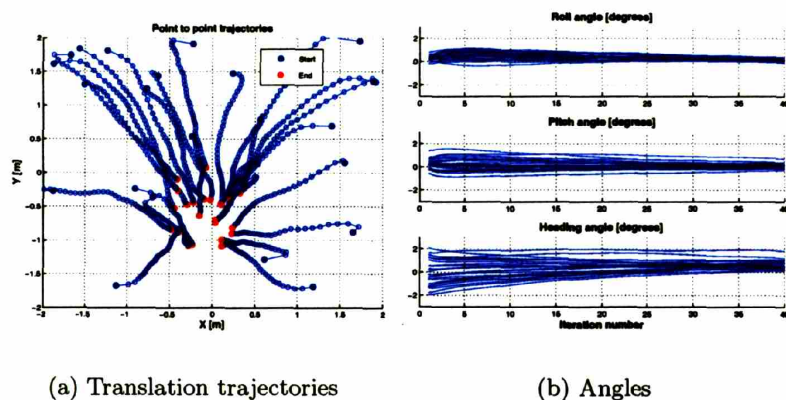
where, $\mathbf{m}_j[k]$ and $\mathbf{m}_i[k]$ make a nearest neighbor pair, and w_k is used to weight the contribution of individual point pairs to the cost function. For a given set the point pairings, Horn's closed form least square quaternion solution [57] can be used to determine the transform T that brings the pairs into alignment. The basic algorithm repeats the following steps,

- select point pairs,
- solve for T using Horn's algorithm,
- apply T to the points $\mathbf{m}_j[\cdot]$
- repeat,

until the alignment transform T between steps approaches unity, indicating the point set alignment has stabilized.

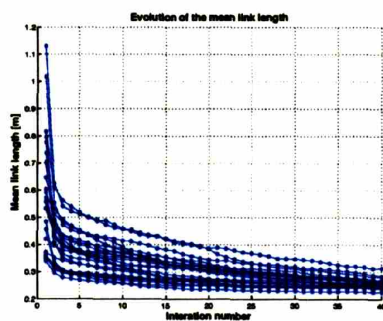
Application of this method to the bathymetric sub-maps has shown poor convergence properties for repeated randomized trials over many different map pairs. An example is shown in Fig(4-3). The trajectories of the transform parameters do not exhibit strong convergence to single solution. This is not unexpected, as the method is known to produce slow convergence when applied to noisy surface data [43, 111]. Examination of the point pair links between the sub-maps shows it is common to generate a large number of links that

do not suggest a consistent direction of motion for improved alignment. This will occur where the surfaces parallel each other and is likely to happen when the terrain is relatively flat. Additionally, the point-to-point approach is prone to aligning artifacts in the scanning pattern when the underlying terrain shows little structure of its own. The striped pattern created from the individual multibeam pings, seen in Fig(4-1), can cause the alignment to find an incorrect match when the maps are created by the vehicle flying on reciprocal headings. The weighting of individual point pair distances proportional to the point placement error using w_k was not found to affect the convergence behavior significantly.



(a) Translation trajectories

(b) Angles



(c) Pairwise error reduction

Figure 4-3: The typical convergence of the point-to-point ICP algorithm for the terrain sub-maps using initial guesses that were randomized around an assumed map alignment. (a) Translation trajectories in $[x, y]$. (b) The change in orientation parameters. (c) The reduction of the point pair distance during convergence. Translation was randomized over a 2m radius, heading over 3° and, pitch and roll over 2° .

Point-to-plane

The point-to-plane ICP method [27] seeks to minimize the sum of squared distances between the closest point pairings along a local surface normal. This distance measure only penalizes surface separation along the surface normal direction and allows contacting surfaces to slide

tangentially without penalty. Although this method can be more prone to local minima than the point-to-point method, it performs significantly better in practice.

The minimization cost function for the normal projected distance is written as

$${}^i_j T = \arg \min_T \sum_{k=1}^n \|(T \mathbf{m}_j[k] - \mathbf{m}_i[k]) \cdot \mathbf{n}_i[k]\|, \quad (4.11)$$

where $\mathbf{n}_i[k]$ is an estimate of the surface normal at point $\mathbf{m}_i[k]$. Since a closed form solution for ${}^i_j T$ does not exist, the cost needs to be reduced with an iterative process. The error $d[k]$ for an individual point pairing can be rewritten using the components of T as

$$d[k] = (\mathbf{R} \mathbf{m}_j[k] + \mathbf{t} - \mathbf{m}_i[k]) \cdot \mathbf{n}_i[k] \quad (4.12a)$$

$$\sim (\mathbf{m}_j[k] - \mathbf{m}_i[k]) \cdot \mathbf{n}_i[k] + \mathbf{r} \cdot (\mathbf{m}_j[k] \times \mathbf{n}_i[k]) + \mathbf{t} \cdot \mathbf{n}_i[k], \quad (4.12b)$$

where $\mathbf{r} = [r_x, r_y, r_z]$ simplifies the rotation \mathbf{R} by assuming the cost will be minimized over small angular displacements. Using this substitution, the minimization can be rewritten with respect to the parameter vector $\mathbf{x} = [\mathbf{r}^\top \mathbf{t}^\top]^\top$ as

$$\mathbf{x}^* = \arg \min_{[\mathbf{r}^\top \mathbf{t}^\top]^\top} \sum_{k=1}^n ((\mathbf{m}_j[k] - \mathbf{m}_i[k]) \cdot \mathbf{n}_i[k] + \mathbf{r} \cdot (\mathbf{m}_j[k] \times \mathbf{n}_i[k]) + \mathbf{t} \cdot \mathbf{n}_i[k])^2. \quad (4.13)$$

Taking partial derivatives of this cost function and setting them to zero results in a 6×6 matrix equation of the form

$$\mathbf{F}^\top \mathbf{F} \mathbf{x} = \mathbf{F}^\top \mathbf{b}, \quad (4.14)$$

where \mathbf{F} is a matrix of Jacobians relating the change in parameters $[\mathbf{r}^\top \mathbf{t}^\top]^\top$ to the change in each point pairwise distance $d[k]$

$$\mathbf{F}_{(6 \times n)} = \begin{bmatrix} \mathbf{m}_j[1] \times \mathbf{n}_i[1] & \cdots & \mathbf{m}_j[n] \times \mathbf{n}_i[n] \\ \mathbf{n}_i[1] & \cdots & \mathbf{n}_i[n] \end{bmatrix}^\top, \quad (4.15)$$

and \mathbf{b} is the residual vector

$$\mathbf{b} = [(\mathbf{m}_j[1] - \mathbf{m}_i[1]) \cdot \mathbf{n}_i[1] \quad \cdots \quad (\mathbf{m}_j[n] - \mathbf{m}_i[n]) \cdot \mathbf{n}_i[n]]_{(n \times 1)}^\top. \quad (4.16)$$

The basic algorithm proceeds by repeating the following steps until the newly determined transform approaches identity.

- Select point pairs
- Calculate \mathbf{F} and \mathbf{b}
- Compute $\mathbf{x} = (\mathbf{F}^\top \mathbf{F})^{-1} \mathbf{F}^\top \mathbf{b}$
- Apply an updated transform $T \Leftarrow \mathbf{x}$ to the points $\mathbf{m}_j[\cdot]$
- Repeat

The final values for \mathbf{R} and \mathbf{t} are then related to the ICP transform Δ_{icp} .

The point-to-plane method requires the calculation of surface normals for the selected points. A robust normal estimation can be made using local principal component analyzes

on groups surface of points [91, 108]. This technique determines the direction of minimal variance that indicates the normal to a tangent plane (see Appendix B.2). The normals calculated in this manner have been accurate enough to successfully apply the point-to-plane ICP to bathymetric data. The actual accuracy of the normal estimation however is difficult to quantify given the lack of ground truth. A detailed description covering the effects of point cloud noise on the normal calculations has been presented by Mitra [91]. Beyond the basic PCA implementation, an additional step to reject surface normals for sonar returns with long return pulse duration can be used. Also, varying the size of the region the local points are collected from in proportion to the point sample density helps the normal estimation consistency.

Point selection

As noted by Gelfand [43], the stability of the ICP solution is related to structure of the of the matrix $F^T F$ which encodes all of the point cloud shape and normal information. If this matrix is poorly conditioned the solution will be unconstrained in one or more directions. This is seen by writing the change in total cost (4.13) as

$$\Delta d^2 = [\mathbf{r}^T \mathbf{t}^T] F^T F [\mathbf{r}^T \mathbf{t}^T]^T. \quad (4.17)$$

When $F^T F$, is poorly conditioned there will exist motions in $[\mathbf{r}^T \mathbf{t}^T]^T$ that produce little variation in cost. To combat this problem methods have been proposed to select the set of matching points from the point clouds that better constrain the solution by minimizing the condition number of $F^T F$. Gelfand [43] presents a method to select points based on constraining the Eigen vectors of $F^T F$. In practice this method has proven to be sensitive to the aspect ratio of the common area between maps and tends to select points near the region borders. This behavior is undesirable in the context of bathymetric sub-mapping for two reasons. First, the aspect ratio of the sub-maps is highly variable due to their incremental assembly. Second, the map edge points will typically be imaged off normal incidence with the sea floor and be more prone to error, Fig(4-4(c)).

As an alternative approach to addressing the solution stability, a normal space sampling method is used. In this method the space spanned by the region wide set of surface normals is gridded and points are chosen as uniformly as possible from the populated grid cells [116]. The idea is to utilize as much constraining geometry as possible. To tailor this to sonar mapping the selection of the points from each normal bin is made according to the shortest returned pulse duration. The objective is to cover the normal space as completely as possible while choosing points that were generated with the potentially most accurate sonar range measurements. The obvious caveat to this approach is that points imaged at favorable near normal incidence from one vantage point will not be imaged favorably from a different vantage point. To account for this the number of points selected from the first surface is increased. After the nearest neighbor points on the second surface are found, the number of links is reduced to the desired number by rejecting links to points on the second surface which have long returned pulse durations. This double sorted duration based sampling has improved the convergence properties of the point-to-plane ICP. Fig(4-4) shows an example of the point selections for both standard and duration preferred normal sampling. As shown in this case, the duration preferred sampling will tend to condense

the selected points into areas that have been favorably imaged. In some sense, this behavior is suggestive of a feature based registration approach, but still maintains a large number of corresponding points across the surface. The convergence behavior for randomized down sampling, standard normal space sampling, and duration preferenced sampling are shown in Fig(4-5). An outline of the sampling algorithm is given in Appendix B.3.

It is worth noting that surface shapes do exist which prevent a normal space sampling strategy from constraining the matrix $F^T F$. For example, if the mapped terrain describes a $\frac{1}{2}$ sphere the normals space will be fully populated yet the solution is completely unconstrained in three rotations [43]. In the context of bathymetric mapping such cases can be considered “pathological” and present a more significant problem to normal space sampling methods applied to 3D object modeling. Also, simple checks can be made to ensure that the selected points have a minimal coverage over the common region and are not overly concentrated in a single area.

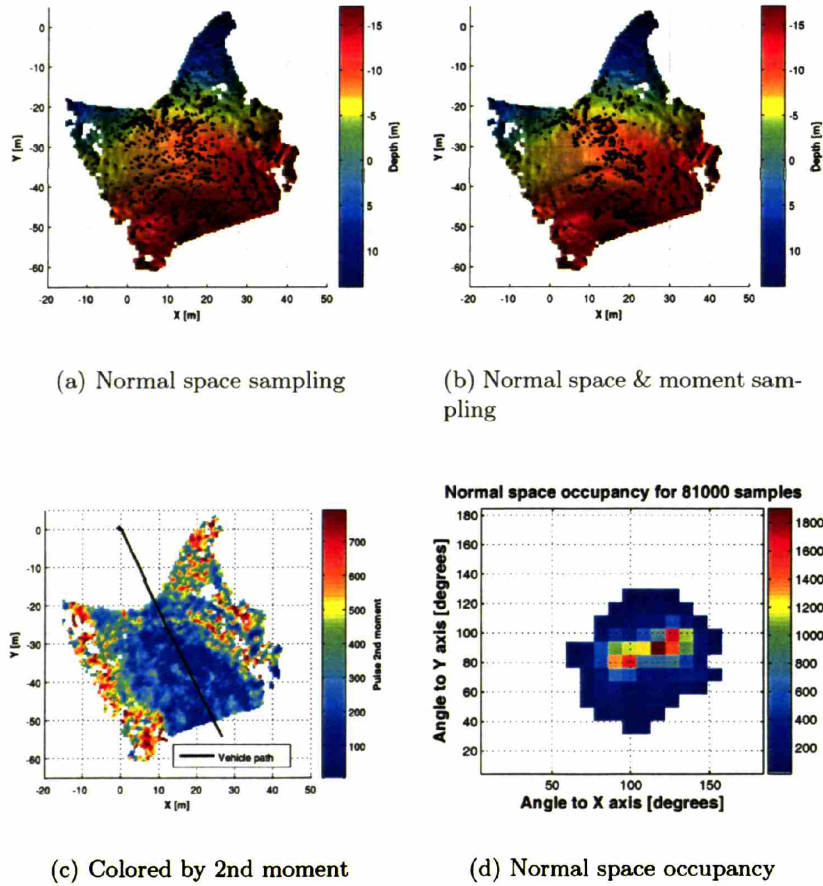


Figure 4-4: Comparison of standard and return duration preferenced normal space sampling. (a) Map points that would be selected using standard normal based sampling. (b) Points selected for pulse duration normal sampling. (c) Terrain color coded by received pulse duration. The black line indicates the vehicle path. Note the longer pulse duration returns closer to the edges of the map. (d) The normal space occupancy for this section of terrain.

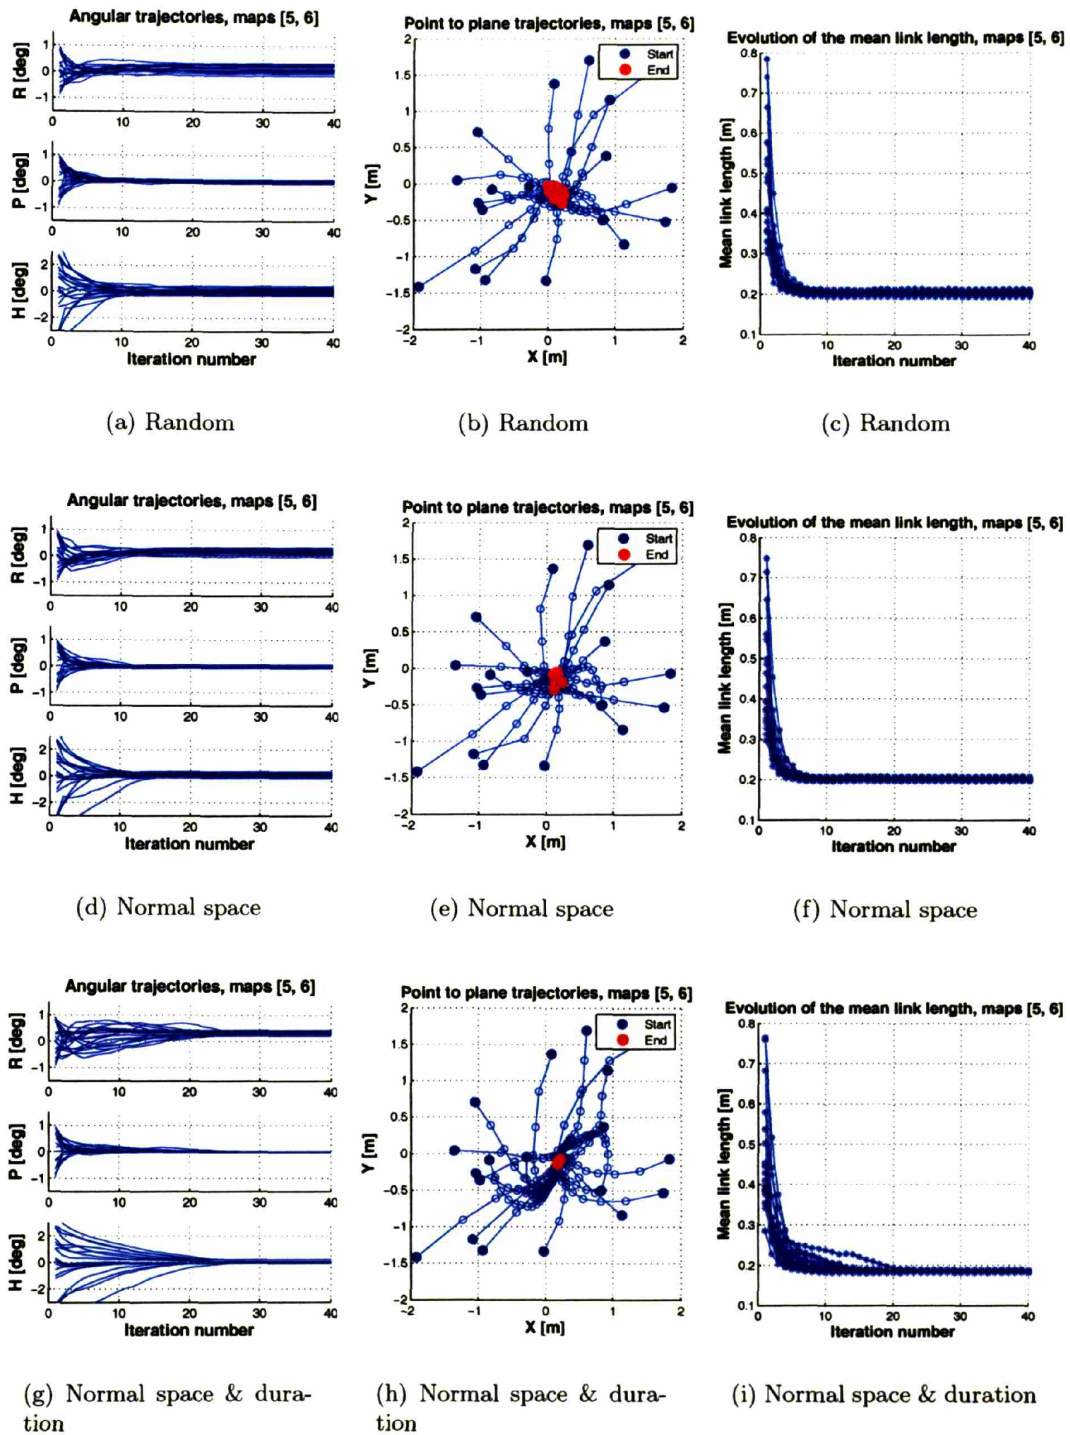


Figure 4-5: Sample convergence behavior for the three different point selection methods. (a,b,c) Random down sampling. (d,e,f) Normal space sampling. (g,h,i) Pulse duration preferred sampling. Randomized starting points were used and duration based sampling produced the tightest convergence behavior in both translation and angular motion, and resulted in the smallest pairwise links lengths. In all cases the convergence was superior to the point-to-point method. Translation was randomized over a 2m radius, heading over 3° and, pitch and roll over 2° .

ICP error estimation

An uncertainty estimate for the point-to-plane estimate can be made once the iteration has converged. Using the value of \mathbf{F} at the solution, the covariance of the transform parameters can be related to the nearest neighbor point distances. The relationship follows as

$$\mathbf{R}_{icp} = E[\mathbf{xx}^T] \quad (4.18a)$$

$$= E[(\mathbf{F}^T \mathbf{F})^{-1} \mathbf{F}^T d d^T \mathbf{F} (\mathbf{F}^T \mathbf{F})^{-1}] \quad (4.18b)$$

$$= (\mathbf{F}^T \mathbf{F})^{-1} \mathbf{F}^T E[dd^T] \mathbf{F} (\mathbf{F}^T \mathbf{F})^{-1} \quad (4.18c)$$

$$= (\mathbf{F}^T \mathbf{F})^{-1} \mathbf{F}^T \Lambda \mathbf{F} (\mathbf{F}^T \mathbf{F})^{-1}, \quad (4.18d)$$

where $\Lambda = \text{diag}[\sigma_1^2 \cdots \sigma_n^2]$ is a diagonal matrix containing a scalar variance statistic for each point pair link length. These individual pair wise variances can be computed from the uncertainty related to each point $\mathbf{P}_{\mathbf{x}_{s_i} r_k \mathbf{x}_{s_i} r_k}$. Alternatively, if all the point pair variances are assumed equal this simplifies to $\mathbf{R}_{icp} = \sigma^2 (\mathbf{F}^T \mathbf{F})^{-1}$. In this case it is clear that the poorly constrained directions will map directly to large error covariances. An estimate for σ^2 can be generated from the pairwise link lengths as $\sigma^2 = \text{var}(d[1 \cdots n])$. In practice this estimate of uncertainty has proven to be over confident in magnitude, but correct in the orientation when $\mathbf{F}^T \mathbf{F}$ is poorly conditioned. This was checked using the stopping points of the randomized trials, Fig(4-5), and a χ^2 bound on the calculated R matrix, Fig(4-6). To compensate for this the calculated covariance $\mathbf{R} = \sigma^2 (\mathbf{F}^T \mathbf{F})^{-1}$ can be scaled in the algorithm.

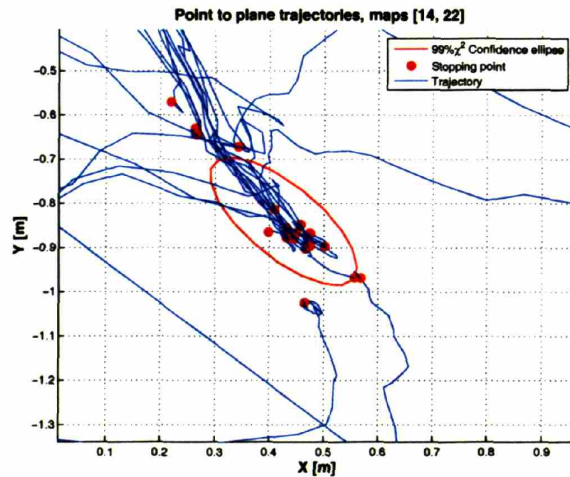


Figure 4-6: Example confidence ellipse for the ICP registration error. The confidence ellipse is drawn at the convergence centroid of several random trials, similar to those shown in Fig(4-5). This particular case shows the error covariance to be over confident. The shape of the convergence paths indicates the less constrained directions.

Limitations to ICP and pairwise measurements

The above uncertainty measure does not capture the possibility of the ICP algorithm converging to a local minima and instead of the actual solution. To guard against local minima a terrain consistency check will be performed using a surface error evaluation with other nearby maps. This will be discussed in section Section 5.3.2.

Also, the EKF formulation has assumed that different pairwise registration measurements involving common sub-maps are independent. This assumption is commonly made in similar constraint based SLAM algorithms [15, 42, 49, 83] due to the difficulty in calculating the measurement cross correlations explicitly. If the individual pairwise matches do not share any common points this assumption can be justified. However, due to the pulse duration sampling strategy it is likely that the same points will be used for different measurements involving the same areas on a sub-map. Unfortunately, the difficulty in accurately predicting even the direct pairwise covariance suggests the required cross correlations to remove the independence assumption would not be easily calculated. As such, the independence assumption between the measurements is used without complete justification.

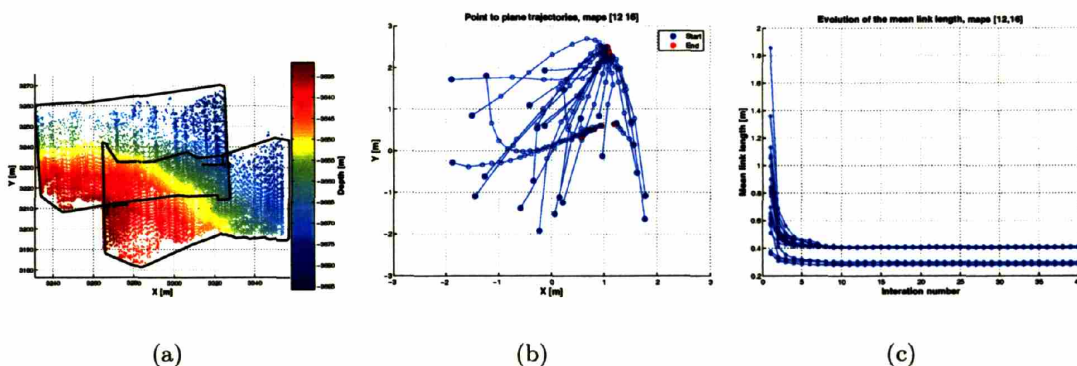


Figure 4-7: An example of the ICP solution finding a local minima . (a) Two sub-maps with a low level of constraining terrain. (b) Two attraction regions for the translation solution. (c) Two different values for the final links lengths suggesting one of the attraction regions is a better match than the other.

4.4 Measurement evaluation

4.4.1 Surface error

To further evaluate the performance of the terrain matching it is necessary to define an error metric which penalizes sub-map mis-registration and highlights inconsistencies when maps are not aligned correctly. Ideally, two correctly registered point cloud sets \mathcal{M}_i and \mathcal{M}_j will combine to produce a composite point cloud set \mathcal{C} which describes a single surface. The surface error statistics of \mathcal{C} should approach those of the noisier surface, \mathcal{M}_i or \mathcal{M}_j , in the common region. Any additional error in the overlapping region should be attributed to mis-registration.

Within the bathymetric mapping community no standard method for reporting mapping errors exists. At the larger scales for ship based surveys, depth errors can be calculated by binning the depth soundings into grid cells and calculating the variance in depth per cell [61]. A more sophisticated method proposed by Calder [20] [21] generates a depth variance statistic as part of its depth estimation at grid points. These approaches assume that the seafloor can be modeled as a height map of the form $z = f(x, y)$. At a large scale the assumption is valid, however for vehicle based mapping over rugged terrain the height map assumption is less justified. The depth variance calculated for points contained in any $[x, y]$ grid cell will over predict errors in regions of sloping and featured terrain, and is not applicable general 3D mapping.

To develop a better error metric that can be applied to point cloud data the following criteria are required.

- The error measure should remain as independent as possible from bin size.
- The measure should be applicable to poorly registered maps. This would include composite point clouds containing “air-gaps” in regions where the surfaces do to touch.
- The measure should utilize the fact that bathymetry surveys can be broken down into sub-maps.

A measure that satisfies these conditions will be broadly applicable to 3D mapping in more complex environments and be able to highlight many of the artifacts commonly present in bathymetric maps.

4.4.2 Principal component analysis (PCA)

The 3D graphical modeling community has addressed surface error measurements using localized PCA. Composite graphical models of complex objects created from multiple range scans are often stored in point cloud form. Error statistics for the model surface can be generated by grouping points locally around the surface and performing individual principal component analysis on the groups [66, 108]. A surface variance calculated in this manner represents the orthogonal projection error of the points onto a locally fit plane. Practically this is accomplished by binning the point data into grid cells or voxels and then obtaining the PCA normal direction and variance estimates described in Appendix B.2. Generating reliable results using this idea requires the noise and mis-registration errors within the composite point cloud to be smaller than the feature scale of the shape itself. This condition generally holds true in the realm of graphical modeling using laser scan data. Experimentation, however, suggests that point clouds generated from acoustically mapped terrain do not generally satisfy this condition. Maps created acoustically will have a higher ratio of noise to surface feature size and are prone to greater registration error. Fig(4-8) shows how the PCA based error estimate will break down as registration errors grow. In the limit the surface error calculated using this method is bounded by the choice of bin cell size. This makes applying the method difficult. Choosing a bin size too small will cap the error value causing it to under predict. Choosing a large cell size allows surface features on the same length scale as the bin size to bias the error estimate even if the underlying point cloud is perfectly registered.




Figure 4-8: Lower and upper bounds of the PCA surface error metric when applied to a composite 2D point cloud. (a) Sketch showing three well registered point clouds (color coded points) where the PCA correctly predicts the surface variance. Arrows indicate the normal to the fitted line within each bin (vertical black lines). (b) Three mis-registered point clouds where the PCA method fails. The right most bin shows the fitted line becoming vertical as the mis-registration error becomes larger than the bin size. (c) The upper and lower bounds for the calculated variance. The upper bound is limited by the bin size directly.

4.4.3 Point based errors

Point-to-point distance measurements have proven to be more robust than surface variance calculations in capturing the potential registration error between bathymetric sub-maps. Point-to-point distances can be used to evaluate both pairwise map registration and the error in a composite surface created from more than two maps. The pairwise case is discussed here and the multiple map case is described in Section 5.3.1.

To develop with measure the distance between a specific point \mathbf{p}_1 in map \mathcal{M}_1 to the closest point in map \mathcal{M}_2 can be formalized using the euclidean norm $\|\cdot\|$ as

$$d(\mathbf{p}_1, \mathcal{M}_2) = \min_{\mathbf{p}' \in \mathcal{M}_2} \|\mathbf{p}' - \mathbf{p}_1\|. \quad (4.19)$$

Calculating $d(\mathbf{p}_1[k], \mathcal{M}_2)$ for every point k in the common region between maps \mathcal{M}_1 and \mathcal{M}_2 directly indicates the registration error between the point sets. An example is shown in Fig(4-9(b)). Unfortunately, calculation of the nearest point distance for each point in the common area is computationally expensive for maps containing $O(100,000)$ points. Layered data structures such as k-d trees [8, 45] can reduce this cost to $O(\log(N))$. To reduce computation further the intersecting region on \mathcal{M}_1 can be gridded and the point-to-point distance for n points in each grid cell j can be averaged as

$$D_j(\mathcal{M}_1, \mathcal{M}_2) = \frac{1}{n} \sum_{k=1}^n d(\mathbf{p}_1[k], \mathcal{M}_2). \quad (4.20)$$

This bin-wise average shown in figure Fig(4-9(c)) approximates the dense point map quite well. Computing a histogram for the error based on all points or the binned error shows how the registration errors are distributed, Fig(4-9(f)). Correctly registering two sub-maps

should move the mass of the histograms toward zero. The lower bound on the point-to-point error metric for a perfect registration is proportional to the sample density of the points on the terrain surface. To quantify the registration error with a simpler statistic the mean, median, and variance of the point error distribution can be calculated. Tests for the reduction in these three statistics are used in assessing whether a registration was successful and if the determined relative pose transform should be returned to the EKF algorithm as a measurement. Fig(4-11(c)) shows the error histograms for a map pair at different stages of the registration process. The registration algorithm will first attempt a correlation measurement and check if the surface error is reduced from the initial EKF proposed error. If the error is reduced the ICP matching is performed. If the ICP measurement is able to improve the surface error it is returned for the EKF update. If the error is not reduced the correlation measurement is returned. An outline of the complete registration process showing these error reduction checks is given in Algorithm 3.

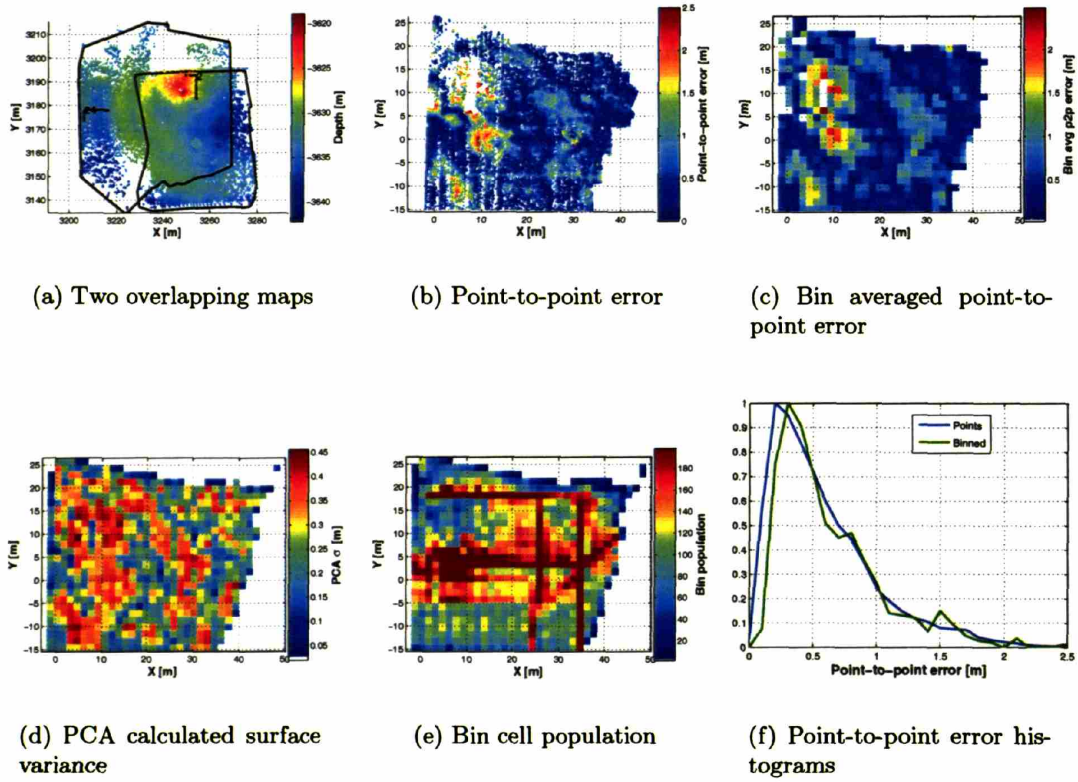


Figure 4-9: Point based error metric example. (a) Two sub-maps are shown in the world coordinate frame. (b) The point based error for all closest pairwise matches (shown in frame 2). Several areas show significant error. (c) The binned point based error measure using 4 points per bin approximates (b). (d) The surface variance calculated using the PCA method shows the surface errors less clearly. (e) The number of points in each bin. (f) The histograms for the point-based error and the binned point-based error. The binned error measure approximates the all point error reasonably.

Algorithm 3 Map Registration The registration process to return a relative pose measurement to the sub-mapping EKF.

```

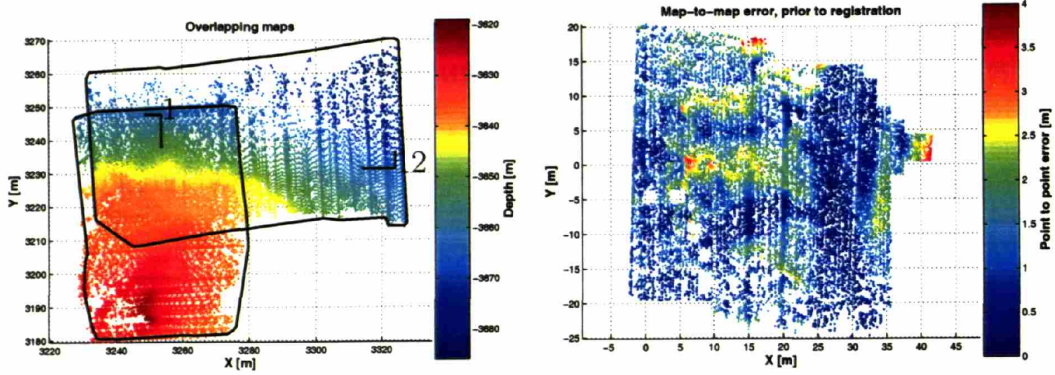
Transform point cloud  $\mathcal{M}_j$  into frame to make  ${}^i\mathcal{M}_j$ 
Grid point maps to make  $M_i$  and  $M_j$ .
Set the size of correlation window using  $P_{\mathbf{x}_{s_{ij}}\mathbf{x}_{s_{ij}}}$ 
Attempt 2D correlation to obtain  $\Delta_c$  and  $R_c$ .
if Correlation H positive definite & surface error reduced then
  Attempt ICP between  $\mathcal{M}_i$  and  ${}^i\mathcal{M}_j$ 
  if ICP convergent & ICP surface error < Correlation surface error then
    Return measurement  $\mathbf{z}_{s_{ij}} = \Delta_{icp} \oplus \Delta_c \oplus \mathbf{x}_{s_{ij}}$  and  $R_{\mathbf{z}_{s_{ij}}\mathbf{z}_{s_{ij}}} = R_{icp}$  to Algorithm 2.
  else
    Return measurement  $\mathbf{z}_{s_{ij}} = \Delta_c \oplus \mathbf{x}_{s_{ij}}$  and  $R_{\mathbf{z}_{s_{ij}}\mathbf{z}_{s_{ij}}} = R_c$  to Algorithm 2.
  end if
end if

```

4.5 Summary

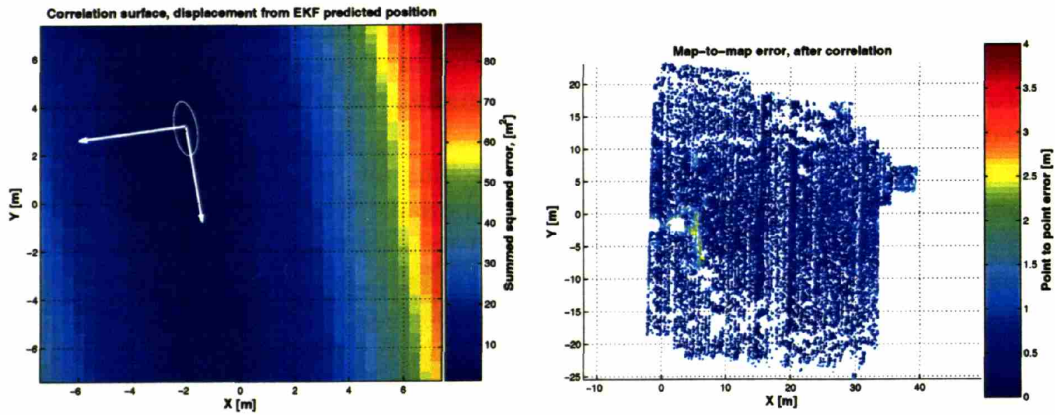
This chapter has outlined the pairwise registration process used to create relative pose links between sub-map origins. A complete example of the registration process is shown in Fig(4-10) and Fig(4-11). This example illustrates the reduction of surface error during the registration steps and shows the surface error that results after the relative pose measurement is incorporated into the EKF filter. Due to the filter's own estimate of $\mathbf{x}_{s_{ij}}$ the surface error after the incorporation into the filter is generally greater than what the correlation or ICP predict independently.

The sequential application of correlation matching and an ICP algorithm is robust to local minima and improves the registration of sub-maps. The point-to-plane ICP method has superior convergence properties over the point-to-point method. This is most likely due to the "sliding" the point-to-plane cost function allows and the predominately low relief maps that slow convergence for the point-to-point method. In addition the point selection preferenced on the returned sonar pulse duration improves the convergence behavior of the point-to-plane algorithm. Finally, it was shown that a point-to-point error metric more accurately captures the registrations errors than the PCA based method typically used by the graphical modeling community.



(a) Overlapping maps

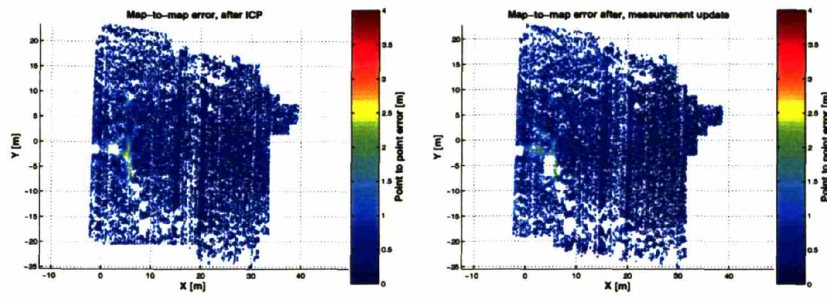
(b) Initial map-to-map error



(c) Correlation surface

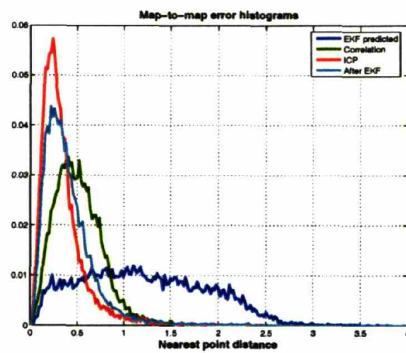
(d) Map-to-map error after correlation

Figure 4-10: Summary of the steps in the sub-map registration sequence. (a) Two overlapping sub-maps. (b) The surface error over the common area prior to registration. This alignment is provided by the EKF predicted relative pose $\mathbf{x}_{s_{i,j}}$. (c) The correlation surface and uncertainty orientation. (d) The surface error after shifting 12 using correlation solution. Note the error is significantly reduced.




(a) Map-to-map error after ICP


(b) Map-to-map error after EKF update



(c) Histograms of map-to-map surface errors



(d) Gridded surface prior to registration



(e) Gridded surface after registration

Figure 4-11: Registration example continued. (a) The surface error after applying the ICP solution. (b) The surface error using the sub-map positions obtained after the ICP measurement is incorporated by the state update equations. (c) Histogram showing the reduction of error between the initial alignment, the correlation and ICP measurements. The “After EKF” line shows the error using the updated map origins. (d,e) Gridded versions of the maps before and after registration. The similarity of these images indicates how mis-registration, as show in Fig(4-10(b)) is not easily discernible from a gridded map alone.

Chapter 5

Experimental results and validation

5.1 Introduction

This chapter presents a validation of the proposed sub-mapping algorithm using results from a real world large scale mapping experiment. It is shown that a better terrain map can be produced using a sub-map framework than using more standard navigation filtering techniques. To compare and evaluate maps a point-based error metric is developed to indicate the total amount on mis-registration in a composite point cloud that describes the entire map. The sub-mapping method is able to reduce this surface error significantly. Tests are also done to show the robustness the of sub-mapping method to some common error sources present in robotic surveys. Some additional details of the algorithm are also discussed and a post processing pose refinement step is presented.

5.2 Survey description

The bathymetric surveys presented here were specifically designed to test the proposed sub-mapping algorithm, Fig(5-1). By general underwater surveying standards, these surveys contain an extreme amount of bottom coverage and numerous crossing tracklines that would normally increase the potential for registration errors caused by uncertain navigation. Ideally however, the sub-mapping algorithm will take advantage of the crossing tracklines lines and limit the surface errors in these regions. The survey patterns are consistent with the previously stated assumption that underwater surveys can be designed to avoid the large loop closures known to cause difficulty for SLAM algorithms. The specific details for the surveys are given in Table 5.1. Although the surveys were completed with an ROV the survey design is consistent with AUV mapping capabilities.

The vehicle platform used for this work was the JASON ROV, which is part of the US National Deep Submergence Facility, Fig(5-2). The ROV contains on board navigation sensors for three axis attitude, three axis bottom relative velocity and, surface relative depth. Table 5.2 shows the specific characteristics of the navigation sensors. The individual sensor measurements are recorded asynchronously at rates varying between 5 and 10 Hz. Acous-

tic long baseline navigation fixes from external beacons are also obtained at a 10 second interval using a vehicle mounted transponder. The JASON system is position controlled in real time using dead reckoning navigation based on the integration of Doppler velocity log (DVL) velocities and measured attitude [72]. During the surveys the LBL fixes are used to periodically “reset” the DR navigation and keep the vehicle close to intended survey path. The LBL fixes are not, however, used in real time in a continuous filtering sense.

Table 5.1: *Summary of survey details*

Detail	Description
Vehicle speed	~.25 m/s
Vehicle altitude	Survey 1, 15 - 20 m Survey 2, 25 - 30 m
Path length & duration	Survey 1, ~5.1 km ~13 hours Survey 2, ~1.8 km ~4.5 hours
Sonar frequency	200 kHz
Outgoing pulse length	50 - 75 μ s
Range resolution	~4 cm per sample
Ping rate	~1 ping per second
Sonar transmitting beams angles	120° athwart ships, 3° fore-aft
Beamforming	128 beams uniform across 120°
Total number of pings	~30,000 each survey

- Calibration details for the SM2000 can be found in [28, 29, 60].

Table 5.2: *Navigation sensors*

Measurement	Sensor	Precision
Heading (north seeking)	FOG ¹	$\pm 1^\circ$
Pitch/Roll	Tilt sensors	$\pm 0.1^\circ$
Depth (surface relative)	Pressure sensor	± 0.01 m
Vehicle velocity (bottom relative)	Acoustic Doppler (DVL)	± 0.01 m/s
Position (x,y)	Long Base Line	$O(1$ m)

¹ The Fiber Optic Gyro (FOG) also has the desirable property of zero heading dependent deviation.

The multibeam sonar used for this work was the SM2000 sonar (Kongsberg-Mesotech Ltd). The details specific to the sonar are given in Table 5.1. The sonar ping rate was chosen to ensure a dense bottom coverage for the nominal altitude and fore-aft beam width. A procedure to determine the sonar pose offset, \mathbf{x}_{vs} , with respect to the vehicle body frame is given in Appendix C. It is important this be done prior to running the sub-mapping algorithm as the errors produced by this offset being incorrect will translate directly into sub-map motion during the registration process. The sonar data was batch processed using

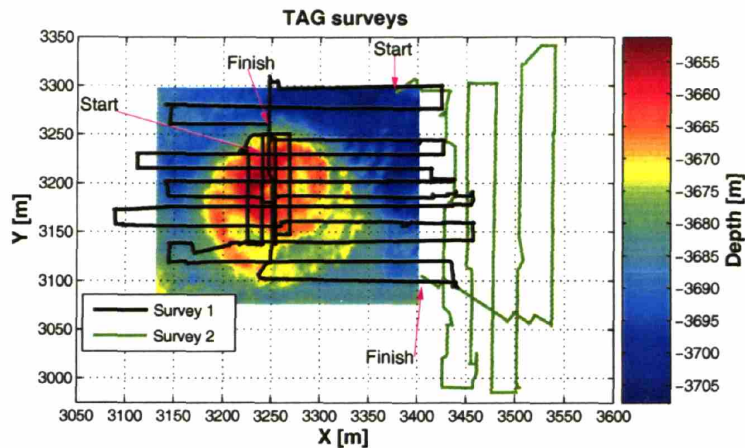



Figure 5-1: TAG surveys. The two survey paths are shown over a low resolution map of the TAG mound. The black survey (survey 1) contains multiple crossings over the main hydrothermal vent. The sloping sides of the mound are at an approximately 45° angle. The smaller survey (survey 2) was completed with a single crossing line and slightly wider trackline spacing.



(a) JASON

(b) The JASON surface navigation display

Figure 5-2: The JASON ROV shown prior to launch over the TAG hydrothermal mound. (a) SM2000 sonar (receiving head highlighted) mounted in a down looking configuration. The transmit array is hidden by the DVL. The ROV is controlled actively in translation, yaw and depth. Pitch and roll rely on passive stability. (b) A screen shot of the DVLNAV topside navigation system [72] available to the navigator on watch. The position of JASON, the clump weight MEDEA, the ship, and the LBL position fixes (+ symbols) are shown in real time. The trackline of the vehicle is shown as a bread crumb trail (green dotted line). During operation JASON uses Doppler based DR navigation for closed loop speed and position control. The navigator is able to monitor the LBL fixes in real time and resets the position estimate to an LBL fix when the difference between the DR and LBL position grows.

the methods detailed in Chapter 2.

The remaining sections of this chapter present data and results from survey 1, the larger of the two surveys shown in Fig(5-1). Similar plots for survey 2 are shown and described in Appendix D. The data from survey 2 has been processed with an identical set of algorithm parameters and shows similar results as survey 1. The trackline pattern in survey 2 resembles a more standard survey with several parallel tracklines and a crossing line. The second survey was performed with slightly wider trackline spacing and a higher flying altitude above the bottom.

The detailed plot of the survey 1 tracklines in Fig(5-3) shows the vehicle trajectory generated by DR navigation using the DVL velocity and attitude measurements only. The tracklines diverge from the LBL fixes due to the integration of velocity and heading noise, and error in the knowledge of the offsets \mathbf{x}_{va} and \mathbf{x}_{vv} between the vehicle frame, and the attitude and velocity sensors respectively. Since the DR navigation is completely relative, the start of the DR track is shifted to coincide with an LBL position fix at a similar time. The covariance ellipses spaced periodically along the track show the growing uncertainty in the position estimate with time.




Figure 5-3: Survey 1 tracklines in detail. The DR tracklines for survey 1 are shown with the LBL navigation fixes and $99\% \chi^2$ uncertainty ellipses. The ellipses show a bound for the DR position estimate in $[x, y]$ and grow steadily over the course of the survey. Note the lack of LBL data for a section of the map (middle right). This “shadow zone” was most likely caused by terrain interfering with the direct acoustic path. The three LBL beacons were located at the following $[x, y]$ locations, $[1572.5, 3151.3], [3780.9, 4746.3], [4264.3, 2275.2]$.

The total sonar sounding density is shown in Fig(5-4(a)). The width of the sonar swath on the bottom for a single trackline is shown in Fig(5-4(b)). Along the highlighted trackline (red) four sub-maps were created. The swath width extends to the neighboring tracklines

for creating a significant amount of redundant coverage. The actual swath width on the bottom will vary according to the terrain slope and the closely spaced tracklines ensure coverage in rugged terrain.

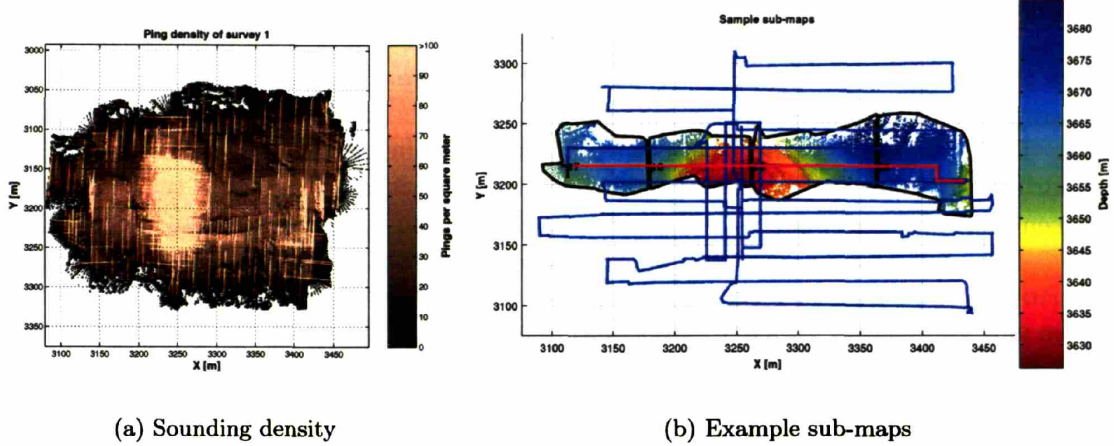


Figure 5-4: Sonar sounding density for survey 1. (a) The sounding density on the bottom. The “hashed” features indicate where the vehicle’s forward motion stopped and the sonar continued to ping. (b) Example sub-maps created along the highlighted (red) trackline. The track line spacing was set to obtain approximately 200% bottom coverage, ie. complete coverage to the adjacent line. The variability in track width suggests that closely spaced lines are needed to ensure complete bottom coverage.

5.3 Complete maps

A single composite terrain map is created from the union of the individual sub-map point clouds once they are independently transformed to the common vehicle local level coordinate frame. The composite point cloud \mathcal{C} is written as

$$\mathcal{C} = \{ {}^1T\mathcal{M}_1 \cup {}^2T\mathcal{M}_2 \cup \dots \cup {}^nT\mathcal{M}_n \}, \quad (5.1)$$

where the transform parameters, ${}^i T \Leftarrow \mathbf{x}_{s_i}$, for each map \mathcal{M}_i are determined from the final estimate of the delayed state vector \mathbf{x}_{aug} . Ideally, given perfect sensor measurements and exact map registrations, this composite point cloud would describe a zero thickness point sampled surface. More likely though, the point sampled surface will have a “thickness” related to errors in the individual maps themselves and registration errors where the maps are mis-aligned with each other. To evaluate the registration error within the composite point cloud knowledge of the origin sub-map for each point should be retained. If the origin map for each point is known, a point based error metric can be constructed to show how the worst case sub-map alignment error is distributed across the surface. For similar reasons to those discussed in Section 4.4, the vertical variance and a surface variance relative to a fitted plane do not accurately represent the error in the composite point cloud. Therefore,

the following point based error method is used to evaluate the mapping error.

5.3.1 Composite surface errors

To extend the pairwise point-to-point error metric of Section 4.4.3 to surfaces composed of more than two maps it is necessary to consider that there will be a set of maps contributing points to any patch on the surface. Thus, the measure of total mapping accuracy for a surface patch should calculate the largest amount of mis-registration present among these maps. To develop this the nearest neighbor point-to-point distance used for pairwise error measuring, (4.19), needs to be extended. The first step is to determine a set wise distance between one point and many other sub-maps. The second step is to find the largest set wise distance amongst the entire set of maps contributing points to a surface patch. To evaluate the error across the surface, patches can be created by binning the surface in $[x, y]$ if the terrain is relatively flat, or sorted the composite point cloud into 3D voxels in areas of high relief.

Once the surface is binned consider a point \mathbf{p}_i from map \mathcal{M}_i located in surface bin j . Also, consider the set of maps \mathcal{T}_j that have contributed points to bin j AND all the bins that surround bin j . Set \mathcal{T}_j , defines the set of maps that points in map \mathcal{M}_i are measure to. By requiring all members of \mathcal{T}_j to contribute points to the bins surrounding bin j biasing the error near the sub-map boundaries is avoided. From the point \mathbf{p}_i a maximal mis-registration distance can then be defined as

$$m_j(\mathbf{p}_i, \mathcal{T}_j) = \max_{\mathcal{M}_k \in \mathcal{T}_j} d(\mathbf{p}_i, \mathcal{M}_k), \quad (5.2)$$

where $d(\mathbf{p}_i, \mathcal{M}_k)$ is the distance from \mathbf{p}_i to the nearest point in sub-map \mathcal{M}_k . This measure can be used to determine how well a given map is registered to a region covered by other maps. This is similar in form to the Hausdorff distance [30], except that this considers the distance from one point to multiple point sets rather than the distance between two sets with multiple points each.

To estimate the total surface error within bin j , the maximum value of the single map error needs to be calculated over all maps that have contributed points to j . To do so, the set of maps contributing points to bin j can be defined as \mathcal{F}_j . Note that \mathcal{F}_j can be larger than set \mathcal{T}_j since it only requires contribution of points to bin j and not all of j 's neighbors. If a set a points \mathcal{P} is chosen, one at random for each map in \mathcal{F}_j , one instance of the total map error for bin j can be written as

$$M_j(\mathcal{F}_j, \mathcal{T}_j) = \max_{\mathbf{p}_k \in \mathcal{P}} m_j(\mathbf{p}_k, \mathcal{T}_j). \quad (5.3)$$

Several instances of this measurement can be averaged to reduce the variance associated with picking the random points from each map. The motivation for the surface binning and random point selection is to reduce the computation while still generating an error estimate that shows the surface errors clearly. The diagram in Fig(5-5) gives an example of this map-to-map error measurement for a 2D slice showing the “thickness” of a composite point cloud for several maps that are not correctly registered. Several examples of this measurement calculated for the real terrain are shown in Fig(5-6).




Figure 5-5: *Map-to-map error example sketch. This illustrates the calculation of the map-to-map error in equation 5.3. The colored points represent members of individual maps and the vertical divisions represent bins. Within each bin, a point is chosen at random from each map. The thin arrows indicate the closest pairs of all points from the other maps. The bolded blue arrows indicate the maximum mis-registration error within each bin. Also note that when determining the closest points allowing the search outside of the immediate bin will avoid bin size related artifacts. The magenta arrow indicates how a nearest green→blue point would incorrectly be used if searching was only allowed inside a given bin. Finally, note that the right most bin with pairings does not show any Map 3 (green) pairings. This is because there are no Map 3 points in both surrounding bins.*

To convert the bin-wise error measurements into a scalar value for the whole map, the mean or median of the error over all bins can be calculated. The lower bound for the measure is related to the surface sampling density, as even perfectly aligned maps will have a sample-to-sample distance. To try and remove this lower bound, and have an error measure that can approach zero, it is tempting to use the point-to-plane distance as defined in (4.12) instead of the point-to-point distance. This would project the point-to-point distance onto a local surface normal and allow the error to go to zero where the surfaces are perfectly aligned. Experience, however, suggests that the difficulty in estimating surface normals for any given point will cause this to be a noisier measurement of the error. The improved performance of the ICP registration (Section 4.3.2) based on the preferenced sampling for points that have potentially better range accuracy, and consequently better surface normal estimates, is consistent with this observation. As such the point-to-point error is used instead of the point-to-plane and the lower bound is noted.

Using this measure of error the incremental change in total surface mis-registration can be monitored as sub-maps are created and relative pose measurements between the sub-maps are made. The surface error will increase as maps are added using the filtered EKF state estimate for the initial map positions. When a relative pose measurement is made between two maps, the base positions of all the maps are adjusted by the update equations (3.11). If the measurement is correct, not a local minimum, the surface error should be reduced over all maps. An example of this reduction in error is shown between figures 5-6(c) and 5-6(d).

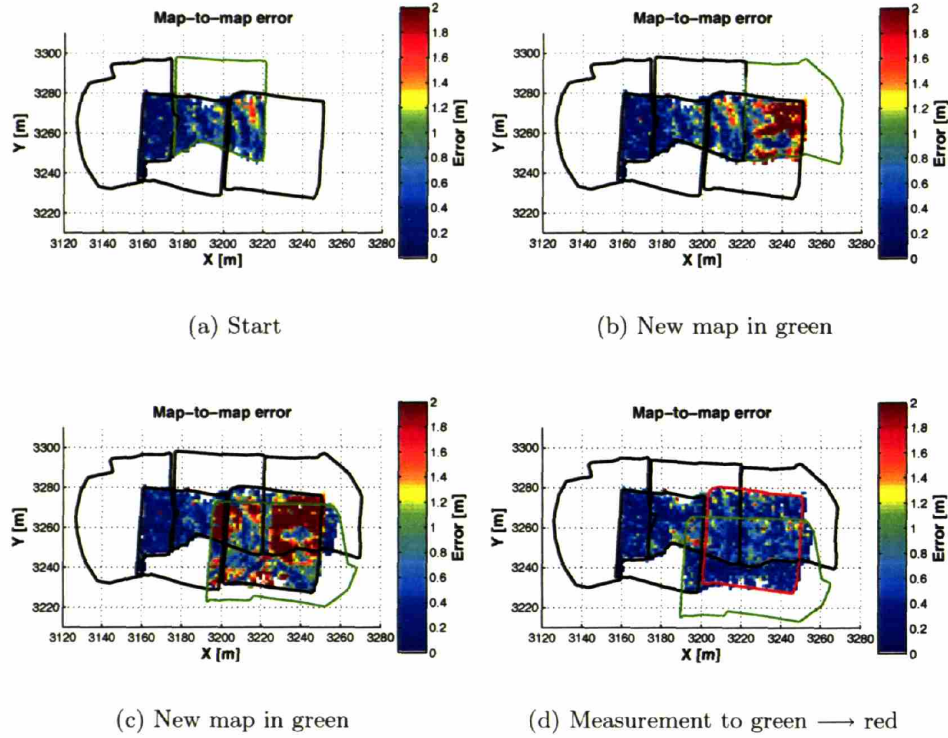


Figure 5-6: Incremental changes in the total map error. The total registration error is shown for the addition of two new sub-maps overlapping a previously mapped area. (a) The existing sub-map borders (black) and total mapping error. Note that the map-to-map error is only calculated where the maps overlap. (b,c) show the error after two additional maps are added. (d) The reduced error after the newest map is registered to an overlapping map (red). 1.5 meter binning was used for the surface error calculation. Also, note that sharp changes in the error can occur at the map boundaries and highlight specific maps that are not well registered.

5.3.2 Terrain consistency checking

Recursive Kalman estimators are known to diverge from a correct state estimate when biased or unmodeled measurements are incorporated [5]. For sub-mapping this can arise when the pairwise map registration returns a measurement corresponding to a local minima instead of the ideal terrain match. Divergence is generally undetectable from examination of the filter covariances directly. One possible check is to monitor whether the normalized innovation

$$\epsilon = (\mathbf{z}_{s_{ij}} - \hat{\mathbf{z}}_{s_{ij}})^\top [\mathbf{H}\mathbf{P}_{aug}^-\mathbf{H}^\top + \mathbf{R}]^{-1} (\mathbf{z}_{s_{ij}} - \hat{\mathbf{z}}_{s_{ij}}), \quad (5.4)$$

for a particular relative pose measurement measurement stays within a 6 DOF χ^2 bound. Violating this bound could indicate that the measurement corresponds to an unlikely terrain match outside of reasonable state uncertainty bounds. However, satisfying the bound does not imply a correct match because the measurement $\mathbf{z}_{s_{ij}}$ could come from a nearby local minima and not the ideal match.

As an alternative, a consistency check can be made using the map-to-map error to verify an improvement after a relative pose measurement is incorporated into the filter. To do this, the median of the binned map-to-map error is calculated for a composite point cloud generated from both maps involved in the pairwise measurement AND all of their overlapping neighbors. The Kalman update, equation (3.11), for the delayed state measurement is then performed and the median error for the point cloud generated using the updated sub-map poses is calculated. If error increases the measurement is rejected and the state is reverted to a copy stored prior to the measurement update. The increase in error would suggest that the measurement is incorrect, and using it requires the sub-maps to be moved in an inconsistent way. This simple check prevents the effects of a single mis-registration from propagating through the entire map. The pairwise measurement shown in Fig(5.3) would pass this test, as the surface error over all overlapping maps is reduced after the measurement.

In general the Kalman filter equations offer no guarantees with regard to the actual surface error. The filter covariances, which indicate the uncertainty in the sub-map locations, are only affected in two ways. Process noise in the vehicle model will increase the state covariance and any measurement, even an erroneous one, will decrease the state covariance. Thus, a decreasing covariance for the delayed state locations is not a sufficient condition for improving map accuracy. The consistency check described here is based on the only available information that is external to the filter itself. The drawbacks of this check are that it is expensive to compute and that it can be too permissive if the underlying terrain has a low level of features which generate error when mis-registered.

It should be noted that the amount of surface error caused by a terrain mis-match is a function of the terrain itself. In very featured terrain registration errors a noticeable because any error in sub-map placement generates inconsistency. If the terrain has a low level of relief mis-registrations can be undetected. In the limiting case, sub-maps from a perfectly flat bottom can be arbitrarily placed without error. This terrain dependence is the primary reason all of the error comparisons used in the algorithm are relative. Defining fixed thresholds on the acceptable amounts of surface error is not generally feasible.

5.3.3 Preliminary maps

Two “standard” mapping methods to compare the proposed sub-mapping algorithm against are, mapping using DR navigation and mapping using a Kalman filtered combination of vehicle velocity, attitude, depth and LBL measurements. The DR navigation alone is not ideal, due to the lack of ground referenced measurements, but is representative of what a robotic vehicle using today’s most accurate commercially available navigation instruments is capable doing on its own. The images in Fig(5-7) show the sub-map layout and surface error for a survey using DR navigation only. The sub-maps were created so the map-to-map error metric could be used, but no relative pose measurements between the sub-maps were made. The overlap map, Fig(5-7(a)), shows the redundancy at the center of the survey where some sections of the bottom were imaged as many as ten times. The tracklines for this survey are those shown in Fig(5-3).

The mapping error produced when LBL fixes are incorporated in the navigation estimation is shown in Fig(5-8(a)). This error map shows a clear improvement in comparison

to the DR navigation and the error is primarily located around the sloping outskirts of the mount where the potential for large mis-registration errors is the highest. The tracklines shown in Fig(5-8(b)) are consistent with the LBL positions. Prior to filtering numerous LBL outliers were removed by hand to prevent obviously erroneous fixes from being incorporated into the filter. The LBL fixes were assigned covariances of $\sigma_x^2 = \sigma_y^2 = 1\text{m}$. Similar surface error results are obtained using a causal filter and non-causal smoother [63] on the same data. The gridded version of the terrain created with this method is shown in Fig(5-11(b)). A vertical slice through the terrain, which shows the disparity between the individual sub-maps, is shown in Fig(5-13(a)).

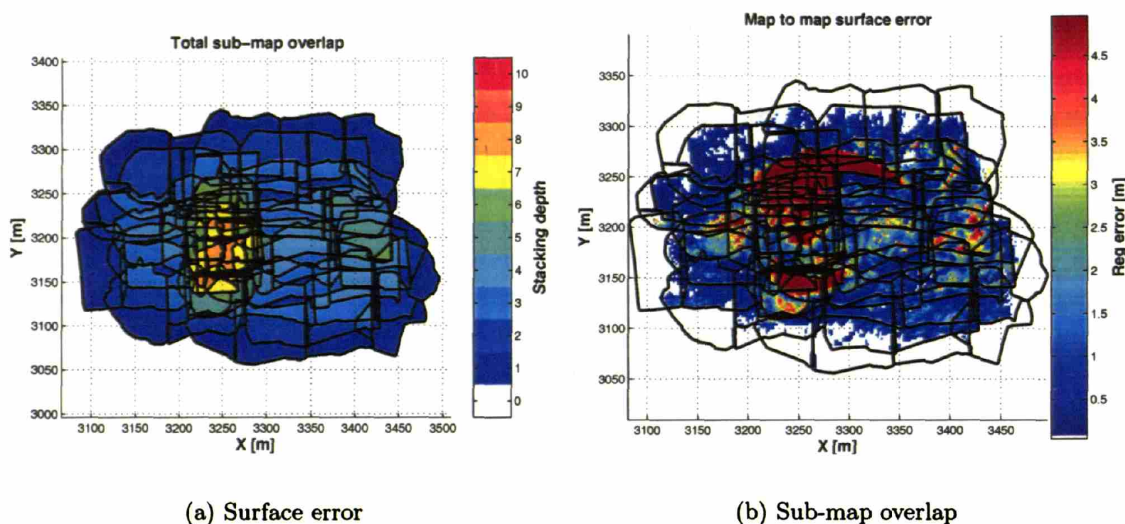


Figure 5-7: DR navigation results. (a) The color coded stacking depth of the sub-maps. (b) The map-to-map surface error for the composite map created using DR navigation. Note that the errors are large where the stacking depth is also large. This indicates that more overlapping coverage is leading directly to more inconsistent mapping.

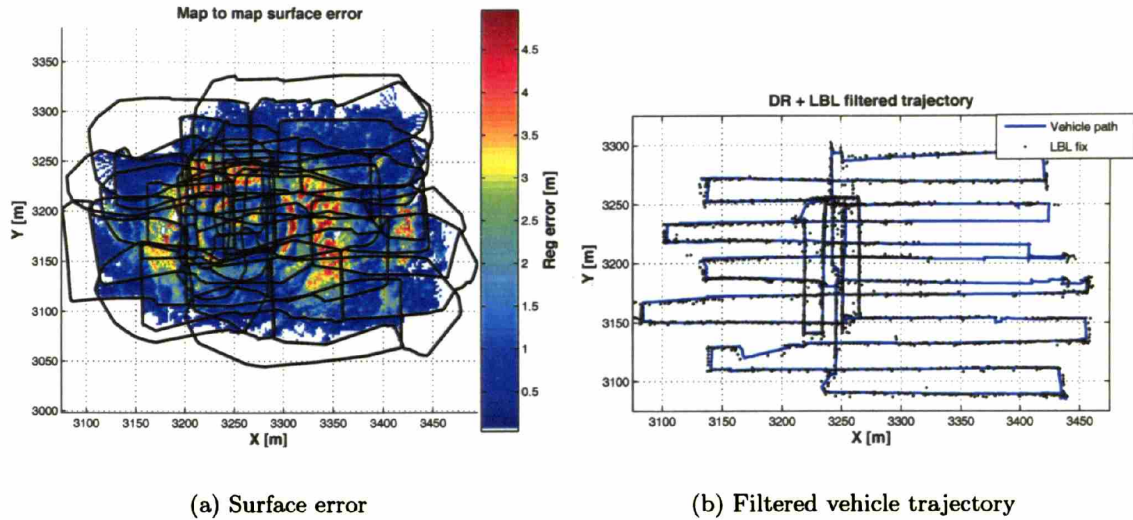


Figure 5-8: Results for LBL filtered navigation. (a) The resulting map-to-map error when LBL navigation is used to generate a map. Note that the error is less than that shown in Fig(5-7(b)) for the DR mapping. The distribution of error is related to the generally circular shape of the TAG mound and located on the steeper terrain slopes (refer to Fig(5-1)). (b) The estimated vehicle path produced by the causal Kalman filter. The filtered trajectory is aligned with the LBL fixes suggesting the drift associated with the DR navigation only has been removed.

5.4 Sub-mapping results

The proposed sub-mapping algorithm will generate a network of constraints between the sub-map origins when applied to the same data as above, Fig(5-9). The links are proposed based on the intersection of sub-map borders (Section 3.5) and relative pose measurements are attempted according to the steps mentioned in Algorithm 3 in Section 3.5. As a result of these pairwise measurements the sub-map origin position uncertainty no longer grows steady along the survey path and is instead related to the link topology. Poses topologically farther from the start of the survey and less connected tend to show larger position uncertainty. This is not a definite statement because the calculated covariances of the individual pairwise measurements used during the filter updates vary from one pair to the next. The failed links shown in Fig(5-9) correspond to links that were proposed due to overlapping map borders but were not established because they failed the error reduction tests mentioned in Algorithm 3 or the incorporation of the measurement failed the consistency check in described in Section 5.3.2.

The map-to-map surface error for the composite point cloud is shown in Fig(5-10). The error has been significantly reduced from the filtered LBL map. Most importantly, error is distributed relatively uniformly across the surface and in general not proportional to the number of overlapping sub-maps as seen in Fig(5-7(a)). Surface error growth in regions of high overlap would indicate repeatedly poor registration. The two areas of the error remaining in the map can be related to two specific sub-maps within which the ROV




Figure 5-9: *Sub-mapping pose network. This pose network was established by the sub-mapping algorithm. Nodes indicate the location of the sub-map origins. Blue links indicate consecutive poses in time. Green links indicate where relative pose measurements were made. Magenta links indicate links that were tried but not established. The uncertainty ellipses have been scaled in size by 8 times for visibility. Note that the poses fall into alignment with the LBL fix locations even though this algorithm did not utilize LBL measurements. This survey consisted of 62 sub-maps and 92 established links.*

was “yanked” by the tether and the constant velocity assumption was intensely violated. The terrain within these maps is distorted in an unmodeled way and will always present a registration problem. The important thing to note however, is that the error caused by these maps remains localized and does not propagate through the entire composite map.

A gridded version of the terrain created using the sub-mapping approach is shown in Fig(5-11(b)). This terrain maps shows considerably more detail than the LBL constructed map. As an example of the detail, a close up view of an Ocean Drilling Program re-entry cone is shown in Fig(5-12). This feature can be clearly seen in the sub-mapped terrain and is completely obscured in the LBL map. The slices through the terrain shown in Fig(5-13) also highlight the more consistent nature of the sub-mapped terrain.

Relation to LBL errors

A closer examination of Fig(5-9) reveals that the sub-map origins align with nearby LBL fixes when the pose network is shifted to align with a single LBL fix at the start of the survey. It is worth noting however, that the difference in map accuracy between the map created using LBL measurements and the sub-mapping map suggests that proximity to the LBL fixes, as in Fig(5-8(a)), does not guarantee surface consistency. This observation can be discussed with regard to both composite map consistency and the accuracy of the sub-map pose network locations.

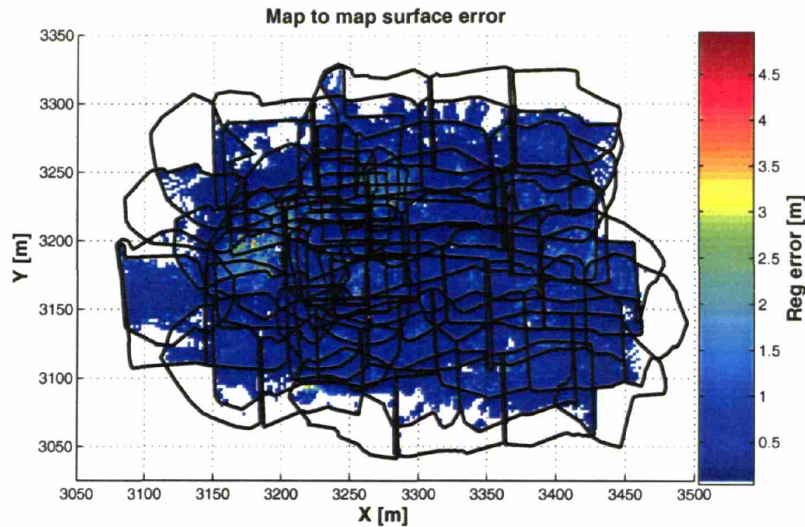



Figure 5-10: *The map-to-map surface error for the sub-mapped terrain is significantly reduced compared to the DR and LBL filtered map errors. The two small regions of error that exists can be attributed to internal distortions in two maps that were affected by a “yank” of the ROV by its tether. Maps with internal distortion will always mis-register with all or part of the surrounding maps.*

The improved surface accuracy with sub-mapping suggests that the simplistic Gaussian modeling of LBL position fix errors, used in the Kalman filter to incorporate the LBL data, is not sufficient. The non-Gaussian nature of the LBL errors [12,63,145] more likely requires a richer error treatment if the fixes are to be used in an automated filter. The improvement in map quality should be more precisely stated as an improvement to LBL mapping, when LBL measurements are handled with Gaussian assumptions. Used in this manner the LBL measurements are not helping the terrain consistency and are instead causing surface errors.


With regard to map positioning, the improvement in map quality does not directly imply that the sub-mapping method will produce better position estimates of the sub-map origins than the LBL measurements. The errors associated with LBL navigation will vary across a survey area, but remain bounded. The sub-mapping pose network is relative, and position uncertainty for the sub-map origins will grow as an unbounded function of topological distance away from any point in the network assumed to be known. These two types of error are different and make a direct comparison between the accuracy of the sub-map origins and the LBL position fixes difficult. For a pose network large enough there will always exist some distance across the network where the positioning errors of the sub-maps origins relative to each other will exceed the LBL positioning error.

Survey density

The tracklines used for survey 1 contain a significant amount of overlap and allow a dense sub-map pose network to be created. To test the sub-mapping concept for a more typical survey pattern with less overlap the survey can be reprocessed with some of the sub-maps



(a) Constructed using EKF sub-mapping



(b) Constructed using sub-maps

Figure 5-11: Comparison between the sub-map created terrain and the LBL created terrain. (a) Terrain created with LBL filtered nav. (b) Terrain produced by sub-mapping. The sub-mapped terrain shows significantly more detail and less scan pattern artifact. Note that the two areas in (b) that do show some pattering (circled), correspond to areas that show error in Fig(5-10). The arrow in (b) indicates the location of an ODP re-entry cone.




Figure 5-12: Close up of mapping detail. (a) Terrain map showing the bathymetry created for the 4 m diameter Ocean Drilling Program re-entry cone. (b) Re-entry cone photographed by JASON's still camera. Note that the concave shape of the cone is captured in the bathymetry.





Figure 5-13: Slices in the XZ plane to illustrate mis-registration. (a) Slice through the LBL created terrain map, with the points color coded by sub-map number. Note the “gaps” between maps indicating registration error. (b) Slice through the sub-map created terrain. The maps more clearly define a single surface and the main peak on the TAG mound is more clearly represented.

prevented from accepting links. The plots in Fig(5-14) show the results when the sub-maps from every other line are forbidden from acquiring links. The resulting surface error plots still show a significant improvement over the DR and LBL mapping cases. It is important to note that since fewer maps are used to show the error in Fig(5-14(c)) there are less opportunities to create surface error. To compensate for this all of the sub-maps can be included in the map and the surface error can be directly compared to the previous maps with the same amount of data. This error is shown in Fig(5-14(d)). The similar amount of error between the full survey and the reduced survey suggests that a small number of links will help reduce the DR related errors and produce an improved map. The data for survey 2, presented in Appendix D, shows another survey topology that is also more typical of a standard vehicle-based survey.

5.5 Robustness to common errors

Automated processing of vehicle navigation data is often complicated by unknown offsets between the vehicle body frame and the navigation sensors, and unmodeled sensor biases. This section describes several typical sensor offsets that cause problems when processing navigation data and shows how they will affect the sub-mapping algorithm. Although the sensor offsets can be measured approximately, errors in their knowledge will create differences between the survey pattern the vehicle actually flew under closed loop control and the pattern recreated in post processing by examination of the logged navigation data. For this discussion a survey can be considered as a sequence of tracklines each specified by a heading, speed, and starting position. If external ground referenced navigation, such as LBL, is available *position* will have a definite origin and orientation. If ground-based measurements are not, DR navigation is the only option and position needs to be defined relative to an arbitrarily chosen origin. Without loss of generality, the position and orientation of this origin can be made coincident with the vehicle body frame pose at some time, as measured using the onboard navigation sensors. The orientation of this frame will be determined with the heading measurement provided by the heading sensor. It can also be assumed that the ground relative velocity measurements, from the DVL, will correspond to the velocities measured along the vehicle body frame axes. Fig(5-15(a)) shows a sketch of a vehicle performing a four leg survey with the heading and velocity sensors correctly oriented to the vehicle and no unknown biases.

When an unknown static offset affects the heading measurement, Fig(5-15(b)), the actual vehicle path over the bottom will differ from what the navigation sensors would indicate. The recorded data would suggest the vehicle also flew the path shown in Fig(5-15(a)), the measured heading was identical for each leg and forward motion occurred purely in surge for both cases. The actual pattern is entirely self consistent, but the vehicle has surveyed a different part of the seafloor. Because the sub-mapping algorithm is entirely relative, this error between the desired survey and the actual survey is unobservable.

As shown in Fig(5-15(c)) a static DVL offset will shear the mapped bathymetry swath but still created a square trackline crossing. As such the error between reality and the recorded trajectory is not observable from the crossing location and only observable in the terrain distortion occurring when the sonar data is mapped using the navigation data in post processing, Fig(5-16). This type of distortion will potentially affect the registration

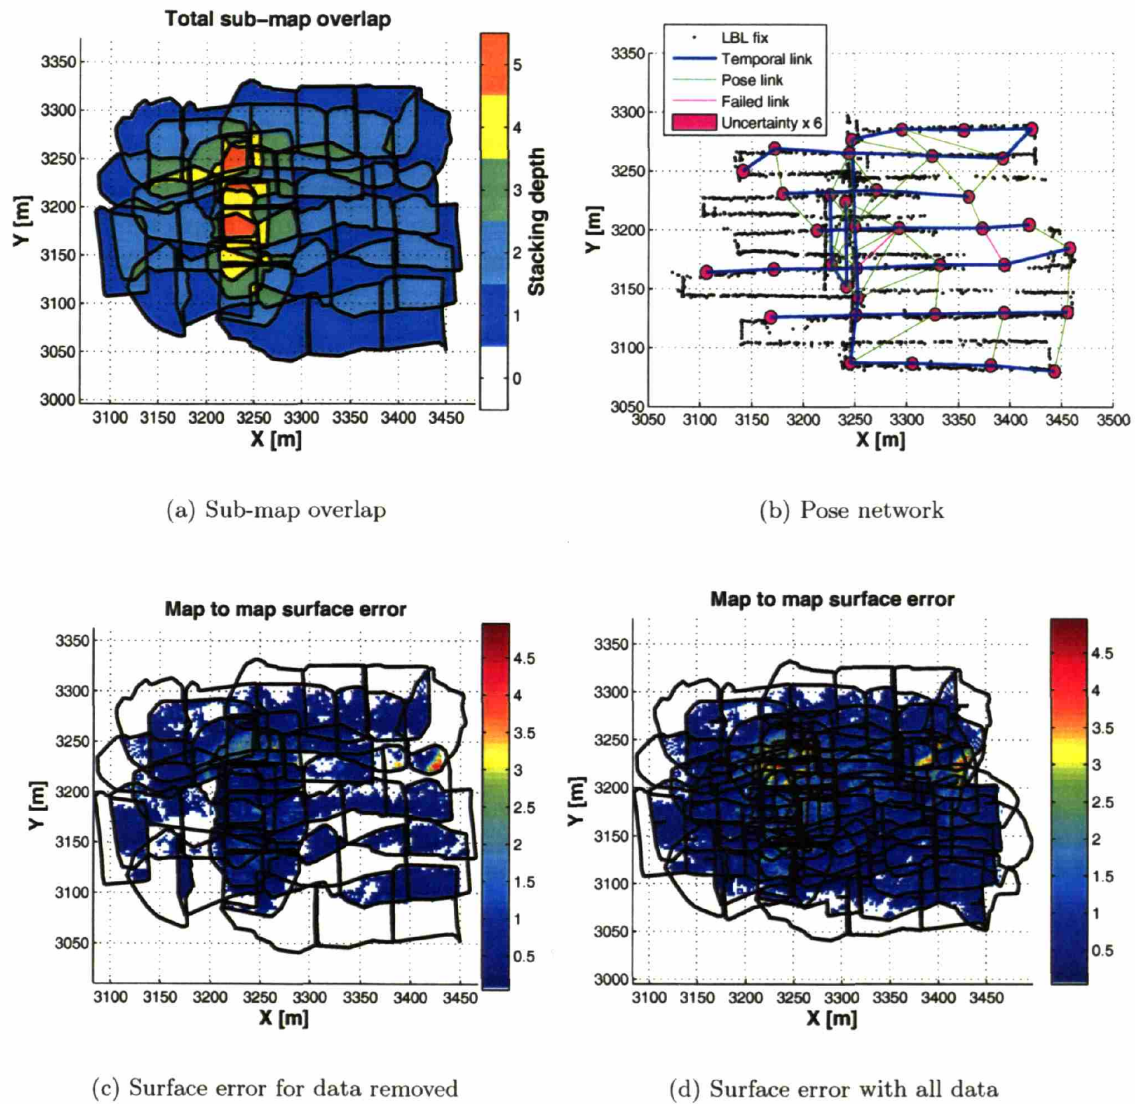


Figure 5-14: Results for survey 1 with the sub-maps from every other line prevented from linking to the pose network. (a) The overlap plot for the sub-maps allowed to link. The amount of overlap is reduced compared to the original survey, Fig(5-7(a)). (b) The pose network created by the sub-mapping algorithm has fewer links and slightly larger error covariance for sub-map locations. (c) The map-to-map surface error for the reduced number of sub-maps. (d) The mapping error when the links and navigation in (b) are used with all of the mapping data. This error is still reduced significantly from the LBL mapping case, Fig(5-8(a))

of sub-maps, as once created the sub-maps are considered rigid. Although some work on non-rigid registration has been done [19,59], attempting to parameterize this distortion and account for it has not been attempted here.

When the heading source has deviation, or heading dependent bias, the crossing points of the tracklines will change Fig(5-15(d)). In this case also, the recorded navigation data will suggest the vehicle performed the survey as in Fig(5-15(a)). Since the actual crossing point over the previous trackline will show inconsistency, this error is observable in the sub-mapping context and can be corrected for. This is illustrated by the test shown in Fig(5-17). The heading data for survey 1 was corrupted intentionally prior to the running the algorithm. The corrupted tracklines show significant error with respect to the LBL position fixes. As shown in Fig(5-17(b)) the sub-mapping was able to compensate for the error and create a network consistent with the LBL. The mapping error, Fig(5-17(c)), is still comparable to the unbiased sub-mapping case and only shows one high error area related to the final maps which were not linked back to earlier sub-maps. It should be noted however that although the effect of the bias has been largely removed, it was not modeled and is in some sense a pleasant but undeserving result. The amount of unmodeled bias that can be removed by the sub-mapping EKF algorithm is a function of the process noise and measurement noise used in the model. A large vehicle process noise will result in the filter utilizing the terrain matches heavily at the expense of filtering the navigation sensor noise.

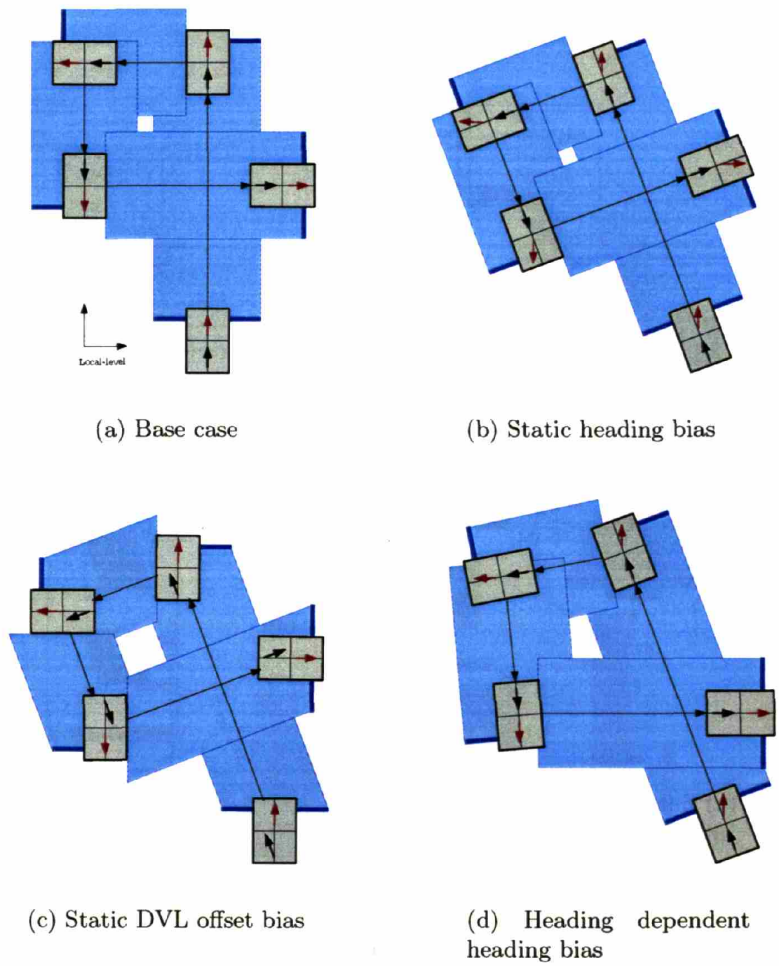
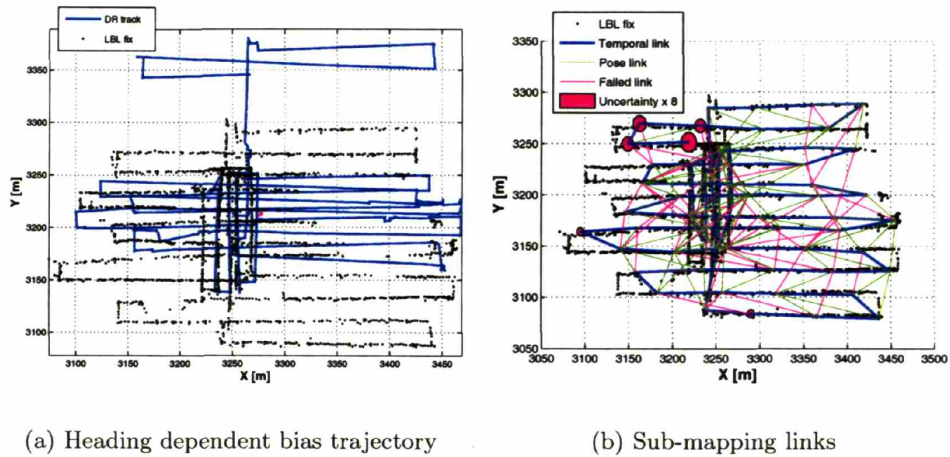



Figure 5-15: Example vehicle trajectories affected by heading and DVL offset error. (a) A sample four leg path driven by a vehicle (gray box) with a heading sensor (red arrow) and DVL (black arrow) mounted correctly on the vehicle. The overlapping sonar swath is blue. The local-level orientation is coincident with the heading sensor at the start. The direction of motion is determined by the direction of the DVL arrow, as would be for an ROV or AUV in closed loop control with a commanded surge velocity and zero commanded sway. The vehicle orientation is defined by the heading sensor arrow. (b) The actual bottom track the vehicle would produce with a static heading offset and the identical commanded path as (a). Note that measured vehicle navigation data would suggest the vehicle flew an identical path as (a). (c) Actual vehicle path for a static DVL offset. The mapped swath is now distorted by a shear. (d) The vehicle path for a heading dependent heading bias. The swath remains square to the vehicle, but the crossing point is in a different location.



(a) Heading dependent bias trajectory

(b) Sub-mapping links



(c) Map error

Figure 5-17: Results for simulated heading dependent bias. (a) DR tracklines when a heading dependent bias of $2^\circ \cos(\psi - \frac{\pi}{4})$ is added to the actual heading. (b) The pose network developed by the sub-mapping algorithm. The poses align with the LBL fixes again and the effect of the heading bias has been removed. (c) The resulting map error. The surface error and sub-map locations are comparable to the un-biased case. The one selection of high error is the result of the last few maps (large ellipses in (b)) not being linked back to the previous maps.

5.6 Additional map refinements

To achieve an additional reduction in total map error the final sub-map poses extracted from the state vector of the delayed state filter can be used as an initial guess for a final pose optimization. This step can be used to counter the effect of unmodeled biases and linearization errors that have been incorporated into the EKF solution. Many SLAM techniques [15, 41, 49, 83] use the formulation of a constraint network to estimate global poses from strictly relative measurements. Although in implementation the solutions vary, the problem is commonly posed as an optimization of a vector valued cost function which relates the individual pose locations to a measure of disparity across all of the available constraints. A source of difficulty for these methods is the creation of a good initial guess for the pose locations to start the optimization. For the problem at hand an initial guess is provided directly from the final delayed state vector \mathbf{x}_{aug} and the constraints are the pairwise terrain registration measurements $\mathbf{z}_{s_{ij}}$ already created for the proposed links.

For each link a disparity transform can be written using the composition sequence

$$\mathbf{e}_{s_{ij}} = \ominus \mathbf{x}_{s_i} \oplus \mathbf{x}_{s_j} \oplus \mathbf{z}_{s_{ij}}, \quad (5.5)$$

that loops from the local level origin, through the relative pose constraint and back to the origin. If all the relative pose constraints $\mathbf{z}_{s_{ij}}$ are satisfied exactly by the location of map origins i and j in the local level frame $\mathbf{e}_{s_{ij}}$ will be the identity transform. If not $\mathbf{e}_{s_{ij}}$ represents a small displacement required to close the pose loop. The pose optimization problem to minimize the size of the disparity transforms over all sub-map locations is formulated for M links as

$$\mathcal{T}^* = \arg \min_{\mathcal{T}} \sum^M \mathbf{e}_{s_{ij}}^\top \mathbf{R}_{\mathbf{z}_{s_{ij}} \mathbf{z}_{s_{ij}}}^{-1} \mathbf{e}_{s_{ij}} \quad (5.6)$$

where $\mathcal{T} = \{\mathbf{x}_{s_1}, \dots, \mathbf{x}_{s_N}\}$ is the set of N sub-map origin positions and $\mathbf{R}_{\mathbf{z}_{s_{ij}} \mathbf{z}_{s_{ij}}}$ is the uncertainty associated with each relative pose measurement. The particular representation of the disparity transform will affect the solution of this optimization problem. Pennek [109] has suggested the axis angle representation of the error transform. Standard optimization packages, such as Matlab's optimization toolbox, can readily handle equations in the form of (5.6).

The solution obtained from (5.6) can be used to reconstruct a refined composite terrain map which shows reduced surface error, Fig(5-18). It should be noted however that, this solution does not penalize the actual surface error and can potentially result in pose refinement at the expense of increased surface error. To combat this (5.6) can be augmented with addition terms to penalize deviations between the pose variables and measurements made directly with navigation sensors at times closest to when the sub-map origins were originally defined. Of the 6 degrees of freedom associated with each origin, x and y translation will be the most uncertain and can be left unpenalized. Pitch, roll, heading and depth however are all measured with respect to stable references and can be more heavily penalized. The

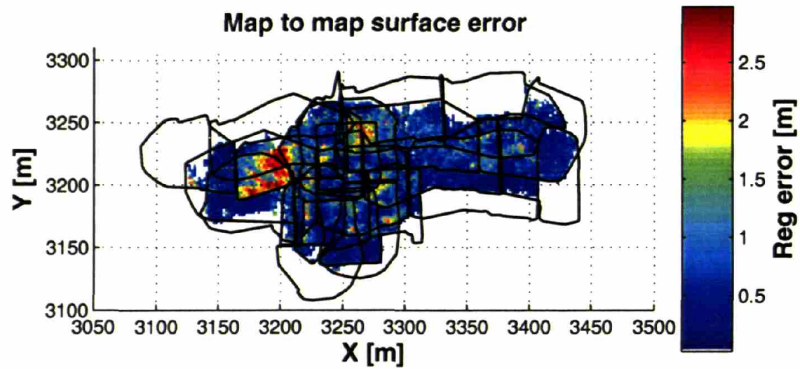
new cost function takes the form

$$\mathcal{T}^* = \arg \min_{\mathcal{T}} \left(\sum^M \mathbf{e}_{s_{ij}}^\top \mathbf{R}_{\mathbf{z}_{s_{ij}}}^{-1} \mathbf{e}_{s_{ij}} + \sum^N (\mathbf{x}_{s_i} - \mathbf{z}_{s_i})^\top \mathbf{W}_i (\mathbf{x}_{s_i} - \mathbf{z}_{s_i}) \right), \quad (5.7)$$


where \mathbf{W}_i is the navigation weighting factor and \mathbf{z}_{s_i} is the collection of navigation measurements associated with pose \mathbf{x}_{s_i} . The weighting can be chosen similar to the measurement covariances of the navigation sensors themselves. It should be noted however, that in using (5.7) the solution can be re-constrained to navigation sensor readings that were biased and the sub-mapping algorithm compensated for. Obtaining improvement with either of these cost functions has required repeated iteration on the weightings and in general should not be considered guaranteed because the surface error is not directly penalized.

5.7 Summary

The results from a real world data set presented in this chapter show the fully automated sub-mapping method can significantly improve terrain mapping consistency when compared to more standard mapping methods using DR and LBL navigation. To show this a point based surface error metric was defined to indicate the total amount of mis-registration within the complete terrain map created by the union of individual sub-maps. The reduction of this error indicates consistency between the mapping data and the navigation data. The sub-mapping algorithm was able to reduce the surface error when applied to both test surveys, and showed robustness to a common heading deviation error. A mapping consistency check based on surface error was also defined. This check adds robustness to the algorithm and prevents errors associated with local minima in the terrain registration step from degrading the entire map. Lastly, two pose refinement steps were presented to adjust the final sub-map positions and potentially reduce the surface error further.



(a) Maps error using EKF poses



(b) Maps error using optimized poses

Figure 5-18: Comparison of surface error before and after the non linear least squares pose optimization. (a) Error map for a section of survey 1. (b) The error map after pose optimization. The overall surface error is generally reduced by the pose optimization. However, there are regions where the error increases. This should be expected as the pose refinement does not directly penalize surface error.

Chapter 6

Conclusions

6.1 Introduction

This thesis has presented a methodology for reducing the navigation related errors that currently limit the accuracy of the vehicle-based bathymetric mapping. The proposed approach has been motivated by the simple observation that, the range accuracy of a sonar measurement relative to the vehicle, will in general be better than the accuracy of the vehicle's own position estimate. With this in mind, the sub-mapping algorithm was designed to utilize the accurate short term navigation provided by high quality navigation sensors and break the entire mapping problem into smaller sections with limited individual error. The sub-mapping concept has proven to be an effective way of creating addition constraints which reduce the corrupting affects of large scale vehicle positioning errors. The sub-maps also allow for the construction and evaluation of an entire bathymetric map.

6.2 Summary

The individual aspects of this thesis can be summarized as follows.

- **Sonar Processing** The sonar processing as described in Chapter 2 was designed to automatically process multibeam data into individual beam ranges with an accompanying "pulse duration" measurement. The simple second moment measurement of returned pulse duration was shown to correlate well with the beam angle of incidence to the seafloor, and is used as an indicator of range measurement accuracy.
- **Delayed state filter** A delayed state extended Kalman filter was used to both filter vehicle navigation data and archive previously visited vehicle positions. The delayed state vector allows relative position measurements based on registered terrain maps to be incorporated into the vehicle position estimation. This allows the navigation to be constrained by the mapping data itself.
- **Sub-map creation** Small bathymetric sub-maps were created using short term DR navigation. It was shown that filter covariances can be used to estimate the uncertainty of the mapping data within the sub-maps. Tests were presented to monitor the geometric properties of the sub-maps as they are created.

- **Sub-map registration** A procedure was presented to pairwise register bathymetric sub-maps using a two dimensional correlation and a six DOF point cloud registration. It was shown that a point-to-plane ICP will provide better convergence properties than the point-to-point method when applied to bathymetric data. A point selection algorithm based on uniform normal space sampling and returned acoustic pulse duration was shown to improve convergence of the point-to-plane method.
- **Complete map evaluation** A point based error metric was developed to evaluate to distribution of registration error in the composite map created from the union of individual sub-maps. Using this measure of surface error a consistency test was developed to indicate incorrect sub-map registrations during the filtering process.
- **Experimental results** The completely automatic processing of a deep water data set was presented. The sub-mapping method was able to produce more accurate maps than can be created using DR navigation alone or LBL filtered navigation. The proposed sub-mapping algorithm was also shown to compensate for heading dependent heading sensor bias.

6.3 Limitations & Future Work

6.3.1 Ground truth

The results presented in this thesis have been judged on the basis of self consistency. The point based multiple map error metric is able to highlight inconsistencies, but an overall ground truth is still missing. Additional experimental work will be needed to resolve some remaining issues regarding the true accuracy of the individual maps and the accuracy of the pairwise registration.

6.3.2 Navigation

SLAM framework

The implementation of the delayed state EKF has proved convenient for navigation filtering and easy manipulation of uncertainty estimates. This solution however does not scale well due to the $O(n^2)$ update computation. The adoption of another SLAM solution methodology with more desirable computational properties is an necessary extension. Potential avenues would include information form solutions [37] and constraint based approaches [15] [41]. For a constraint based approach, a fixed state size Kalman estimator could be used to create the individual sub-maps from the vehicle navigation data. A more computation attractive framework would allow for a real-time extension of the sub-mapping method. The current implementation is strictly casual, but limited to less than 100 sub-maps. More amiable computation would also allow for a multi-scale implementation of the sub-map algorithm where maps are broken into smaller pieces with new local origins as the estimate of there position becomes more confident.

LBL characterization

Characterizing the errors associated with LBL navigation is difficult, as LBL is typically the only ground referenced measurement available in typical AUV and ROV deployment scenarios. The presented terrain registration should provide the ability to ground reference the vehicle position over the course of a survey and allow meaningful LBL measurement residuals to be calculated. This could shed light on some of the persistent and time dependent (tide related) errors in LBL navigation.

6.3.3 Terrain registration

Section 5.6 presented a method for potentially reducing the surface error further by applying a non-linear optimization over the pairwise constraints developed by the sub-map registration. The word potentially is used because this optimization does not penalize surface error directly. In fact, this refinement can be framed as the more general and unsolved problem of multi-view point cloud alignment. Although many solutions have been proposed [7, 9, 14, 112, 119, 136], none have developed a computationally efficient method for directly penalizing surface error. Unlike the similar problem of bundle adjustment in computer vision [52], where a distance measure between specific features can be defined, the distance between point clouds is not easily obtained. The proposed sub-mapping algorithm has converted bathymetric mapping to multi-view registration. However, for sub-maps created in this manner the problem is further complicated by errors internal to the sub-maps themselves. This is not generally addressed in multi-view registration and a straight forward way to characterize this error is not immediately apparent.

6.3.4 Acoustic modeling

The acoustic modeling used in this thesis has been intentionally simple. The returned pulse duration measure was created as a proxy for an individual return's range accuracy. The Gaussian assumptions for beam width were made for computational simplicity. A more detailed investigation into the affects of beam width and rough surface scattering should produce more accurate estimates of range uncertainty. The use of "point clouds" is also an approximation for a finite beam width sensor. The point cloud approximation and the more sophisticated error model may be better handled using a particle sampled representation [77] to generate a denser point cloud with statistics consistent with the true nature of the errors.

Appendix A

Relative pose transformations

The following sections summarize the notation used to represent the coordinate frame relationships used in this thesis.

A.1 Basic definitions

Vectors written in the form

$$\mathbf{x}_{ij} = [x, y, z, \theta, \phi, \psi]^T \quad (\text{A.1})$$

describe the spatial relationship of reference frame j with respect to frame i , Fig(A-1). The parameters $[x, y, z]$ determine the vector ${}^i\mathbf{t}_{ij} = [x, y, z]^T$ that points from the origin of frame i to the origin of frame j as expressed in coordinate frame i . The angular parameters $[\theta, \phi, \psi]$ represent the sequence of rotations about the z axis, then y' axis and finally x'' axis that take the orientation of frame i to the orientation of frame j . Although this notation follows that used by Smith [126], the rotation sequence differs and follows the convention used by Fossen [40]. As a result the direct application of Smith's detailed equations requires a re-ordering of the angular sequence.




Figure A-1: Basic sketch of the coordinate frames.

These parameters can be written as a transformation to express any point ${}^j\mathbf{p}$, originally expressed in frame j , as the point ${}^i\mathbf{p}$ expressed in frame i . The transformation operator i_jT

is applied as

$${}^i\mathbf{p} = {}^i T^j \mathbf{p} \quad (\text{A.2a})$$

$$= {}^i\mathbf{t}_{ij} + {}^i\mathbf{R}^j \mathbf{p}, \quad (\text{A.2b})$$

and requires the rotation matrix ${}^i\mathbf{R}$. This matrix is written in terms of the individual rotations as,

$${}^i\mathbf{R} = \mathbf{R}_z(\psi)^\top \mathbf{R}_y(\phi)^\top \mathbf{R}_x(\theta)^\top \quad (\text{A.3a})$$

$$= \begin{bmatrix} \cos(\psi) & \sin(\psi) & 0 \\ -\sin(\psi) & \cos(\psi) & 0 \\ 0 & 0 & 1 \end{bmatrix}^\top \begin{bmatrix} \cos(\phi) & 0 & -\sin(\phi) \\ 0 & 1 & 0 \\ \sin(\phi) & 0 & \cos(\phi) \end{bmatrix}^\top \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\theta) & \sin(\theta) \\ 0 & -\sin(\theta) & \cos(\theta) \end{bmatrix}^\top. \quad (\text{A.3b})$$

There is a common relationship between with pose vector components, the translation vector and rotation matrix, and transform operator

$$\mathbf{x}_{ij} \Leftrightarrow \{{}^i\mathbf{t}_{ij}, {}^i\mathbf{R}\} \Leftrightarrow {}^i T, \quad (\text{A.4})$$

where any one can be determined from the others.

To accommodate sub-maps the sets of points $\{\mathbf{m}_i[1], \mathbf{m}_i[2], \dots, \mathbf{m}_i[n]\}$ contained in map \mathcal{M}_i are expressed in the base reference frame for map \mathcal{M}_i . To move these points to a new reference frame the transform ${}^j T$ can be applied to map \mathcal{M}_i as

$${}^j \mathcal{M}_i = {}^j T \mathcal{M}_i. \quad (\text{A.5})$$

After the transform operation the individual points can be written as $\{{}^j \mathbf{m}_i[k], {}^j \mathbf{m}_i[2], \dots, {}^j \mathbf{m}_i[n]\}$.

A.2 Additional relations

A.2.1 Head-to-tail

The *composition*, or sequential linking, of two reference frame relations will produce a single relation. To specify this relation the *head-to-tail* operation

$$\mathbf{x}_{ik} \triangleq \mathbf{x}_{ij} \oplus \mathbf{x}_{jk} \quad (\text{A.6})$$

is used. The composition has removed the intermediary frame j . The calculation of the parameters for \mathbf{x}_{ik} is accomplished by

$${}^i\mathbf{t}_{ik} = {}^i\mathbf{t}_{ij} + {}^i\mathbf{R}^j \mathbf{t}_{jk} = {}^i T^j \mathbf{t}_{jk} \quad (\text{A.7a})$$

$${}^i\mathbf{R} = ({}^i\mathbf{R})({}^j\mathbf{R}). \quad (\text{A.7b})$$

The individual parameters for roll, pitch and heading $[\theta, \phi, \psi]_{ik}$ can be solved for using the elements of ${}^i\mathbf{R}$ [126].

The Jacobian of this relationship with respect to the individual parameters is calculated as

$$\mathbf{J}_{\oplus} \triangleq \left. \frac{\partial \mathbf{x}_{ik}}{\partial (\mathbf{x}_{ij}, \mathbf{x}_{jk})} \right|_{6 \times 12} \quad (\text{A.8a})$$

$$= [\mathbf{J}_{\oplus 1}, \mathbf{J}_{\oplus 2}]. \quad (\text{A.8b})$$

The Jacobian is used to propagate a first order estimate of the relationship covariance when the individual parameters are considered random variables with their own covariance and cross covariance estimates.

$$\mathbf{P}_{x_{ik}x_{ik}} = \mathbf{J}_{\oplus} \begin{bmatrix} \mathbf{P}_{x_{ij}x_{ij}} & \mathbf{P}_{x_{ij}x_{jk}} \\ \mathbf{P}_{x_{ij}x_{jk}} & \mathbf{P}_{x_{jk}x_{jk}} \end{bmatrix} \mathbf{J}_{\oplus}^{\top} \quad (\text{A.9})$$

A.2.2 Inverse

The *inverse* operation can be used to change the direction of a pose relation. This is defined as

$$\mathbf{x}_{ji} \triangleq \ominus \mathbf{x}_{ij}. \quad (\text{A.10})$$

The individual parameters are calculated from

$${}^j \mathbf{t}_{ji} = -{}^i \mathbf{R}^{\top} {}^i \mathbf{t}_{ij} \quad (\text{A.11a})$$

$${}^j \mathbf{R} = {}^i \mathbf{R}^{\top}. \quad (\text{A.11b})$$

The Jacobian relation

$$\mathbf{J}_{\ominus} = \left. \frac{\partial \mathbf{x}_{ji}}{\partial \mathbf{x}_{ij}} \right|_{6 \times 6}, \quad (\text{A.12})$$

can then be used to convert the covariance estimate accordingly.

$$\mathbf{P}_{x_{ji}x_{ji}} = \mathbf{J}_{\ominus} \mathbf{P}_{x_{ij}x_{ij}} \mathbf{J}_{\ominus}^{\top} \quad (\text{A.13})$$

A.2.3 Tail-to-tail

Lastly, the *tail-to-tail* relation can be used to derive the intermediary relation between two poses relations

$$\mathbf{x}_{jk} = \mathbf{x}_{ji} \oplus \mathbf{x}_{ik} \quad (\text{A.14a})$$

$$= \ominus \mathbf{x}_{ij} \oplus \mathbf{x}_{ik}. \quad (\text{A.14b})$$

The parameters for \mathbf{x}_{jk} can be calculated from

$${}^j \mathbf{t}_{jk} = -{}^i \mathbf{R}^{\top} {}^i \mathbf{t}_{ij} + {}^i \mathbf{R}^{\top} {}^i \mathbf{t}_{ik} \quad (\text{A.15a})$$

$${}^j \mathbf{R} = ({}^i \mathbf{R})^{\top} ({}^i \mathbf{R}). \quad (\text{A.15b})$$

The Jacobian for this relation is calculated as

$${}_{\ominus}\mathbf{J}_{\oplus} = \frac{\partial \mathbf{x}_{jk}}{\partial (\mathbf{x}_{ij}, \mathbf{x}_{ik})} \Big|_{6 \times 12} \quad (\text{A.16a})$$

$$= \frac{\partial \mathbf{x}_{jk}}{\partial (\mathbf{x}_{ji}, \mathbf{x}_{ik})} \cdot \frac{\partial (\mathbf{x}_{ji}, \mathbf{x}_{ik})}{\partial (\mathbf{x}_{ij}, \mathbf{x}_{ik})} \quad (\text{A.16b})$$

$$= \mathbf{J}_{\oplus} \cdot \begin{bmatrix} \mathbf{J}_{\ominus} & \mathbf{0}_{6 \times 6} \\ \mathbf{0}_{6 \times} & \mathbf{I}_{6 \times 6} \end{bmatrix} \quad (\text{A.16c})$$

$$= [\mathbf{J}_{\oplus 1} \mathbf{J}_{\ominus}, \mathbf{J}_{\oplus 2}] . \quad (\text{A.16d})$$

Appendix B

3D Point set matching

B.1 Surface gridding

Surface gridding can be considered broadly as any process which takes a point cloud of non-uniformly spaced surface samples and generates samples at the nodal points of a pre-determined grid. For the figures presented in this thesis a Gaussian weighted gridding is used that assumes the terrain can be described as a height map. This method weights the contributions of the point samples to the depth at the grid nodes based on a radially symmetric Gaussian function, Fig(B-1). The parameters of the Gaussian dictate how far the influence of a single point will spread. Typical values for these parameters create a Gaussian kernel with a standard deviation proportional to the sonar foot print size on the bottom. Although more sophisticated surface gridding methods exist this simple method has proven sufficient to create grids for individual sub-maps used in the registration process and for displaying the complete composite maps. The extension to true 3D gridding can be made using mesh generation methods that consider the direction the surface is imaged from [31,62]. Notationally, the gridded version of a sub-map point cloud \mathcal{M}_i is represented by M_i .

B.2 PCA surface normal estimation

Normal estimation

Given a set of the 3D points that describe a surface, a robust surface normal estimation can be achieved using a principal component analysis over localized groupings of the points. This method is used as a standard in many surface registration and representation techniques [56,66,91,108]. For a given sample point \mathbf{p}^* in a map, a set of points $\mathcal{P} = \{\mathbf{p}_1, \dots, \mathbf{p}_n\}$ located within a predefined spherical radius of \mathbf{p}^* is created. A local covariance matrix can be calculated as:

$$\mathbf{C} = [\mathbf{p}_1 - \bar{\mathbf{p}}, \dots, \mathbf{p}_n - \bar{\mathbf{p}}] [\mathbf{p}_1 - \bar{\mathbf{p}}, \dots, \mathbf{p}_n - \bar{\mathbf{p}}]^\top, \quad (\text{B.1})$$

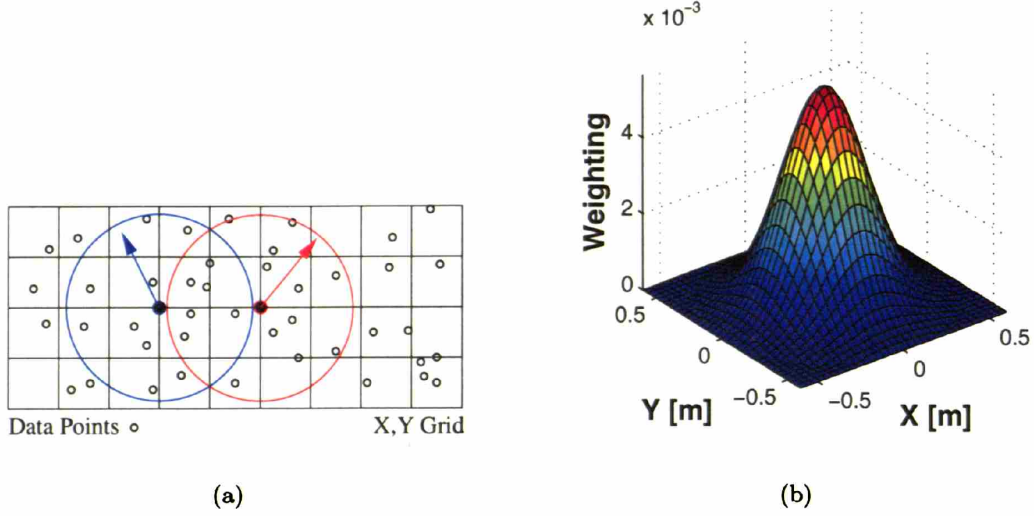


Figure B-1: Surface gridding sketches. (a) Points contributing to a grid node are selected to be inside a specified radius from the node. (b) A Gaussian kernel weights the individual point contributions to the node depth.

where, the centroid of the point set $\bar{\mathbf{p}}$ is determined by

$$\bar{\mathbf{p}} = \frac{1}{n} \sum_{i=1}^n \mathbf{p}_i. \quad (\text{B.2})$$

Let $\mathbf{E} = [\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3]$ be the matrix of Eigen vectors and $\{\lambda_1, \lambda_2, \lambda_3\}$ the Eigen values of \mathbf{C} . If $\lambda_1 < \lambda_2 < \lambda_3$ are the Eigen values of \mathbf{C} and span \mathbb{R}^3 , \mathbf{v}_1 is normal to the surface tangent plane spanned by \mathbf{v}_2 and \mathbf{v}_3 at \mathbf{p}^* . This defines the surface normal estimate $\mathbf{n}^* = \mathbf{v}_1$ and indicates the direction of minimal projected residual variance to the tangent plane.

When an entire map is considered, additional tests can be used to avoid spurious normals at the map edges or in regions of low sample density. In this application checks are made for a minimum number of points contained within the sphere. If too few points surround \mathbf{p}^* the surface normal is not calculated and \mathbf{p}^* is left out of any subsequent operations that require a surface to be defined. A check can also be made on the condition number of \mathbf{C} . As the ratios $\frac{\lambda_3}{\lambda_1}$ and $\frac{\lambda_2}{\lambda_1}$ approach 1, the point set \mathcal{P} more likely describes a spherical or cylindrical collection of points rather than a planar surface. In the cases where either ratio is below a preset threshold, typically $O(10)$, the normal at \mathbf{p}^* is also left undefined. A more detailed analysis on the effect of sample density and surface curvature has been shown by Mitra [91]. Due to the directional ambiguity associated with \mathbf{v}_1 , additional steps can be taken to ensure normal direction consistency among neighbors [56].

Surface error

The local surface variance can also be estimated using the principal components [108]. The Eigen values of C equal the sum of the squared projections along the principal directions, ie

$$\lambda_j = \sum_{i=1}^n (\mathbf{v}_j \cdot (\mathbf{p}_i - \bar{\mathbf{p}}))^2. \quad (\text{B.3})$$

Thus the variance in the j_{th} principal directions is

$$\sigma_j^2 = \frac{1}{n-1} \lambda_j, \quad (\text{B.4})$$

and can be used as a measure of the surface error in a region surrounding point \mathbf{p}_i .

B.3 Point sampling methods

Rusinkiewicz [116] suggests that a surface normal based sampling approach can be used to down sample a point cloud prior registration. This is accomplished by discretizing the space of surface normal directions into bins and sampling uniformly amongst the bins to obtain the down sampled point set. As a robustness measure for sonar data registration, a step to sort the points in each bin by returned acoustic pulse duration can be added to the procedure detailed in Algorithm 4.

Algorithm 4 Normal based sampling This can be done to select M points from a set size N .

- 1: For each point in \mathcal{P} determine a surface normal using the PCA method, Appendix B.2.
 - 2: Define normal space bins, \mathcal{B}_{xy} , over the ranges $[-\pi \leq \angle_x \leq \pi]$ & $[-\pi \leq \angle_y \leq \pi]$.
 - 3: For each point determine the angles between the normal vector and the x and y axes.
 - 4: Sort all points with defined normals into the bins.
 - 5: **if** Using pulse duration sampling **then**
 - 6: Order points in all bins by increasing returned duration.
 - 7: **else**
 - 8: Randomize the ordering of points within each bin.
 - 9: **end if**
 - 10: Let M be the number of remaining points to be selected and n be the number of populated bins.
 - 11: **while** $M > n$ **do**
 - 12: Select the first point from each populated bin.
 - 13: Set $M = M - n$
 - 14: Find the new n value
 - 15: **end while**
 - 16: Select M points, by choosing the first point from the remaining n bins randomly.
-

Appendix C

Sonar sensor offset refinement

Placing the sonar ranges in space to build a map requires knowledge of the vehicle-to-sonar offset \mathbf{x}_{vs} shown in Fig(3-2). The sensor offset can be physically measured with limited precision, but will usually require an additional correction to be determined by an investigation of the mapping data. Errors in the offset will causes range points to be mapped inconsistently even when the vehicle position is known precisely. The basic idea is to find the offset vector that minimizes a measurement of the surface error for a region mapped from several vantage points where the vehicle navigation is well known. A methodology for doing this is given by Singh [121, 122]. Here, the map-to-map error metric developed in Section 5.3.1 is used in a similar manner to refine the hand measured estimate of \mathbf{x}_{vs} .

A short length of vehicle trajectory which contains the vehicle flying a U-turn is taken from the TAG data set, as described in Chapter 5. Over this short section of trackline the vehicle navigation is assumed to be exact and all the error in the mapped surface is attributed to error in the sensor offset. Without precise position measurements, as used by Singh, the next most reasonable step is to select a short section of DR navigated trackline that has some small amount of error, Fig(C-1). For this section of trackline the terrain on the interior of the U-turn is imaged by the sonar three different times.

The optimal sensor offset will minimize the binned composite surface error over the multiply mapped region. Assuming that \mathbf{x}_{vs} is parameterized with three translations and three rotations, $\mathbf{x}_{vs} = [t_x, t_y, t_z, \theta, \phi, \psi]$, the minimization is written as

$$\mathbf{x}_{vs}^* = \min_{t_x, t_y, t_z, \theta, \phi, \psi} \sum_{j=1}^{N_{bins}} M_j(\mathcal{F}_j, \mathcal{T}_j) \quad (\text{C.1})$$

where, $\mathcal{T}_j = \mathcal{F}_j = \{\mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3\}$ are the overlapping maps and $M_j(\mathcal{F}_j, \mathcal{T}_j)$ is the map-to-map error for a single bin in the common area described in Section 5.3.1. Fig(C-2) shows a comparison of the map-to-map error over the common region for two different values of the roll offset. Of the angular offsets, roll will affect the surface error the most significantly [122]. Equation C.1 can be solved numerically to yield the final set of offset parameters.

There a multiple choices for the surface error metric used in (C.1). The sum over the bin-wise variances to fitted planes, described in Appendix B.2, and the sum of bin-wise variance in the z direction can also be used. These would both be computed after the




Figure C-1: Three sub-maps are created from the U-turn data. The map borders are colored to match the vehicle position track in $[x, y]$ for the three sections of the U.

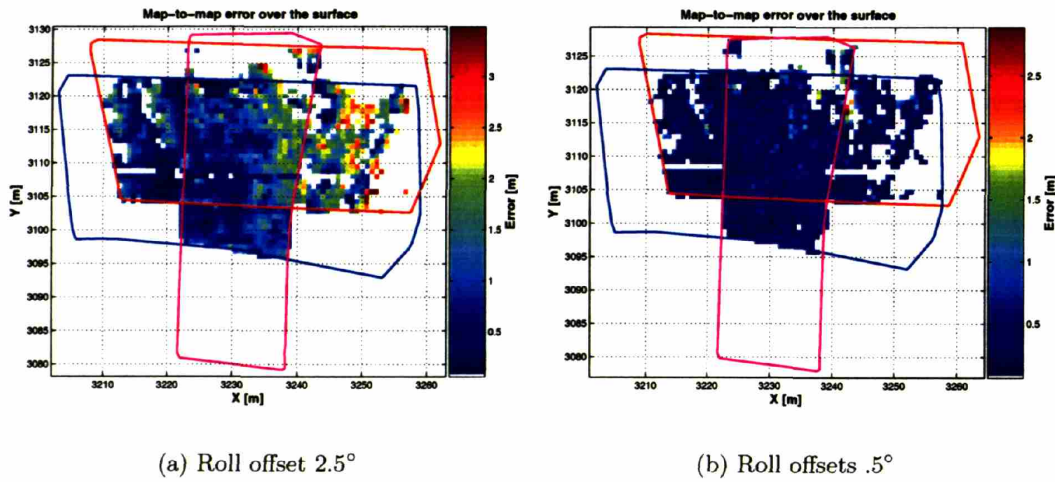


Figure C-2: Different roll offsets will change the amount of surface error in the overlapping region. These plots were made using 1.5m binning.

maps $\{\mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3\}$ have been merged into a single composite point cloud. The plots in Fig(C-3) show the change in all three surface error metrics as a function of the roll, pitch and heading offsets. The roll and pitch offsets show a clear minimum in error. The heading offset, Fig(C-3(c)), is the most difficult to estimate since the surface variance does not change significantly when the offset is varied. For all three offset angles the map-to-map error metric indicates the change in surface error as well or better than the fitted plane error or z variance error.

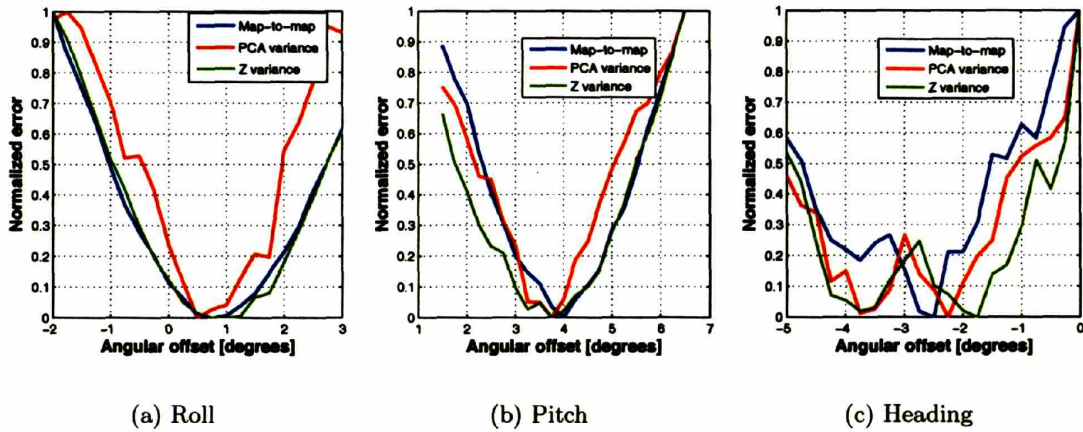


Figure C-3: Variations in surface error for the vehicle-to-sonar roll, pitch and heading offsets. These images show how the different surface error calculation methods are able to capture the change in surface error. Roll (a) and pitch (b) show clear minimums at particular offset values while heading (c) is less clear. The different error measurement have be normalized for comparison.

Appendix D

TAG survey 2

The results for the second TAG survey are shown here. The second survey was completed with the vehicle flying at a higher altitude and with wider spaced tracklines. The algorithm parameters used to process this data set were identical to those used to process survey 1. Qualitatively similar results were obtained. The sub-mapping algorithm was able to create a pose network and produce a terrain map with less surface error than both DR navigation alone and LBL filtered navigation. Fig(D-1) shows the tracklines and the growing uncertainty ellipses for the DR navigation.




Figure D-1: Dead reckoning only tracklines for the second survey shown with LBL fixes and growing $99\% \chi^2$ uncertainty ellipses. Note that the northern ends of the lines align with the LBL fixes while the southern ends of the lines do not. This suggests there could be a location dependent bias in the LBL fixes.

The surface error for the DR navigated map is shown in Fig(D-2(a)). The overall level of error is less than that for survey 1. A difference in the surface error should be expected. The amount of detectable surface error is a function of both the navigation errors and the underlying terrain itself. The actual terrain for survey 2 is shown in Fig(D-4). The sub-map overlap plot in Fig(D-2(b)) shows that the surface sampling for this survey is less redundant than survey 1.

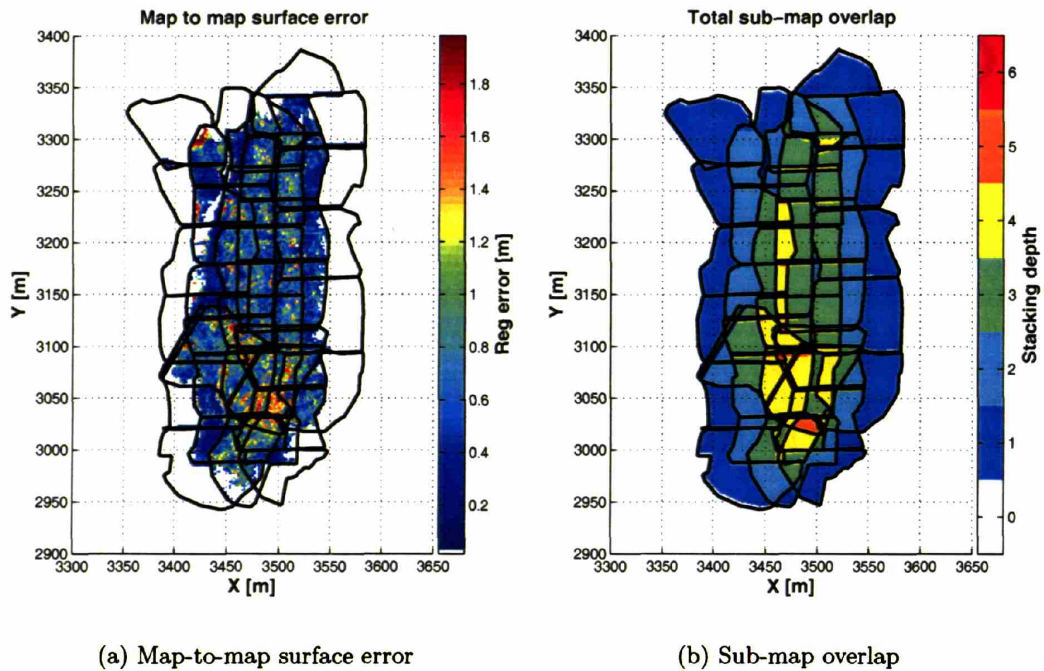


Figure D-2: DR navigation results for survey 2. (a) The map-to-map surface error using the DR navigation shown in Fig(D-1). (b) The total sub-map overlap.

The surface error from the sub-mapped terrain and the pose network for this survey are shown in Fig(D-3). The surface error is reduced from that shown in Fig(D-2(a)). Also, the surface error is more uniformly distributed across the terrain and does not show an increase in the neighborhood of the crossing trackline as visible in Fig(D-2(a)). The pose network shows a network dependent error growth instead of a time dependent growth. It can be noticed that even after sub-mapping the southern sub-map origins still do not align with the LBL fixes. This further suggests that the LBL is biased in this region of the survey area.

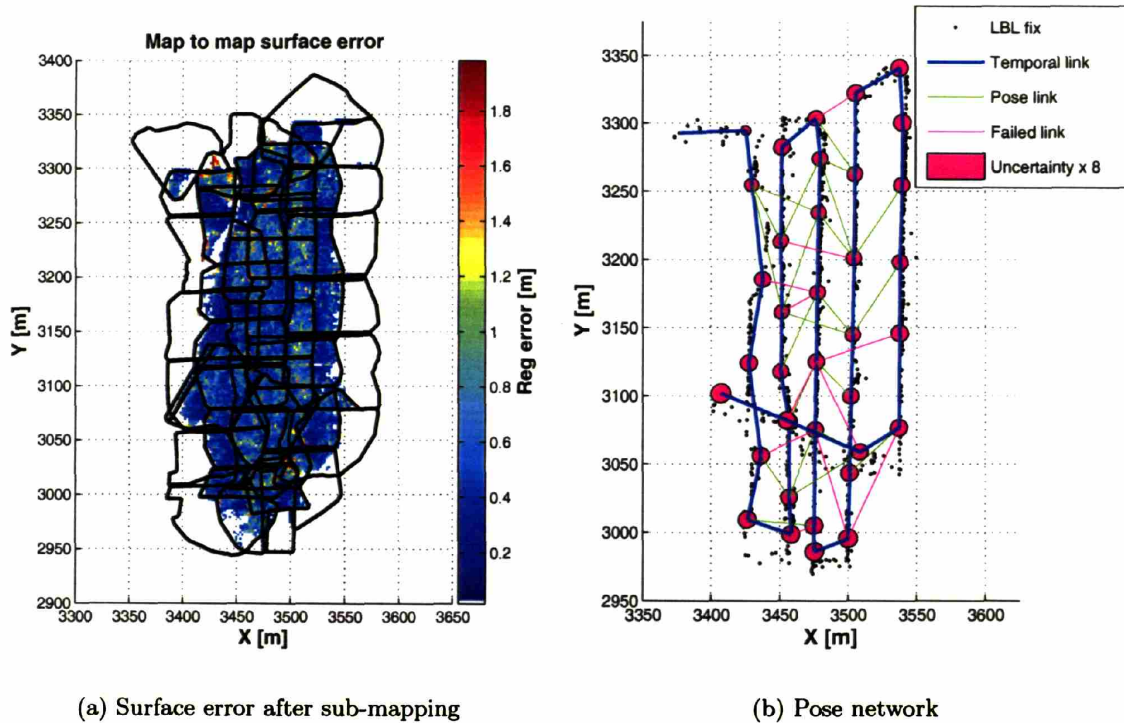
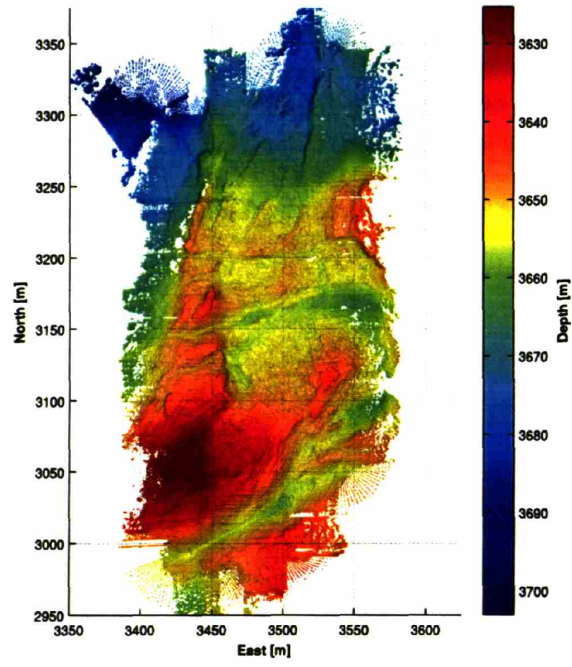
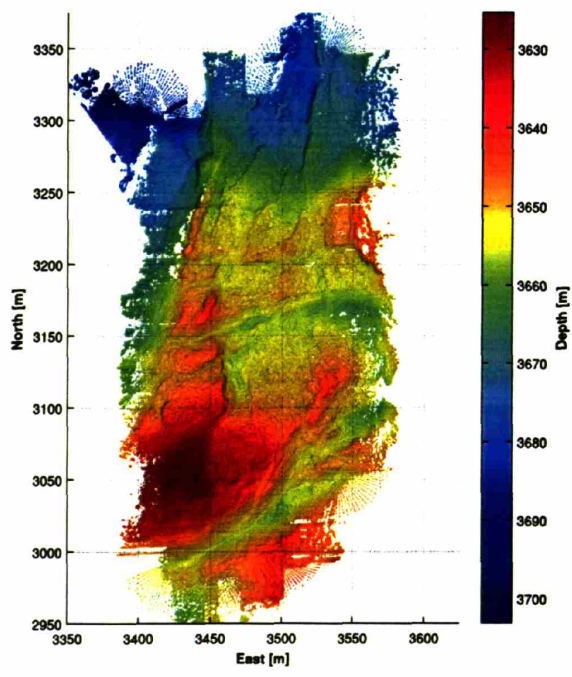


Figure D-3: Survey 2 surface error after sub-mapping. (a) The surface error for the composite map created with the sub-mapping algorithm. (b) The pose network and final sub-map origin covariances. The time dependent growth of the pose uncertainties has been eliminated and the error is now network dependent.

The terrain maps for survey 2 are shown in Fig(D-4). For this survey the difference between the terrain maps is difficult to detect. This difficulty in apparent accuracy speaks to the utility of the map-to-map error for indicating the map regions with errors that would otherwise be unknown.



(a) DR terrain



(b) Sub-mapping terrain

Figure D-4: *The gridded terrain for survey 2. (a) Terrain created from the DR navigation. (b) Terrain created from the sub-mapping algorithm.*

Bibliography

- [1] D. Alexandrou. Sea Beam Sidelobe Interference Cancellation. In *OCEANS*, volume 17, pages 472 – 476, 1985.
- [2] P. K. Allen, I. Stamos, A. Troccoli, B. Smith, Leordeanu M., and Hsu Y. C. 3d modeling of historic sites using range and image data. In *IEEE Conference on Robotics and Automation*, pages 145–150, Tapei, 2003.
- [3] A. Bachmann and S. B. Williams. Terrain Aided Underwater Navigation - A deeper insight into generic Monte Carlo localization. *Australasian Conference on Robotics and Automation*, 2003.
- [4] R.D. Ballard, A.M. McCann, and D.R. a.e.a. Yoerger. The Discovery of Ancient History in the Deep Sea Using Advanced Deep Submergence Technology. *Deep Sea Research*, 1(47):1591–1620, 2000.
- [5] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan. *Estimation with Applications to Tracking and Navigation*. John Wiley & Sons, Inc, 2001.
- [6] A. Basu and N. K. Saxena. A Review of Shallow-Water Mapping Systems. *Marine Geodesy*, 22(4):249–257, 1999.
- [7] R. Benjemaa and F. Schmitt. A Solution for the Registration of Multiple 3D Point Sets Using Unit Quaternions. In *European Conference on Computer Vision*, pages 34–50, 1998.
- [8] J. L. Bentley. Multidimensional Binary Search Trees Used for Associative Searching. *Communication of the ACM*, 18(9):509–517, 1975.
- [9] R. Bergevin, M. Soucy, H. Gagnon, and D Laurendeau. Towards a General Multi-view Registration Technique. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(5):540–547, 1996.
- [10] P. J. Besl and N. D. McKay. A method for registration of 3-d shapes. *IEEE PAMI*, 14:239–256, 1992.
- [11] B. S. Bingham, D. A. Mindell, D. R. Yoerger, B. Foley, and W. Seering. Acoustic Multipath Identification with Expectation-Maximization. In *Proc. IEEE Oceans*, pages 2388–2395, San Diego, Sept 2003.

- [12] Brian Bingham. *Precision Autonomous Underwater Navigation*. PhD thesis, Massachusetts Institute of Technology, 2003.
- [13] J.S. Bird and P Kraeutner. Cramer-Rao Bound Investigation of Swath Bathymetry Accuracy. In *OCEANS, 2001. MTS/IEEE Conference and Exhibition*, volume 3, pages 1640 – 1647, Nov 2001.
- [14] G. Blais and M.D. Levine. Registering Multiview Range Data to Create 3D Computer Objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(8):820 – 824, 1995.
- [15] M. Bosse. *ATLAS, A Framework for Large Scale Automated Mapping and Localization*. PhD thesis, Massachusetts Institute of Technology, 2004.
- [16] M. Bosse, J. Newman, P. and Leonard, and S. Teller. An ATLAS Framework for Scalable Mapping. Technical report, Massachusetts Institute of Technology, October 2002.
- [17] O. Bozma and R. Kuc. A Physical Model-Based Analysis of Heterogenous Environments Using Sonar-Endura Method. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(5):497–506, May 1994.
- [18] N.A Brokloff. Matrix Algorithm for Doppler Sonar Navigation. In *IEEE OCEANS 94'*, pages 378–383, Brest, France, 1994.
- [19] B. Brown and S. Rusinkiewicz. Non-Rigid Range-Scan Alignment Using Thin-Plate Splines. In *Symposium on 3D Data Processing, Visualization, and Transmission*, 2004.
- [20] B. R. Calder and L. Mayer. Robust Automatic Multibeam Bathymetric Processing. In *Proc. US Hydrographic Conference*, 2001.
- [21] L. A Calder, B. R.; Mayer. Automatic processing of high-rate, high-density multibeam echosounder data. *Geochemistry Geophysics Geosystem*, 4(6):1–23, 2003.
- [22] G. Canepa, O. Bergem, and N. Pace. A New Algorithm for Automatic Processing of Bathymetric Data. *IEEE Journal of Oceanic Engineering*, 28(1):62–76, 2003.
- [23] S. M. Carbotte, W. B. F. Ryan, W. Jin, M. Cormier, E. Bergmanis, J. Sinton, and S. White. Magmatic subsidence of the East Pacific Rise (EPR) at 18° 14'S Revealed Through Fault Restoration of the Ridge Rrest Bathymetry. *Geochemistry Geophysics Geosystem*, 4(1), January 2003.
- [24] R. N. Carpenter and M. R. Medeiros. Concurrent mapping and localization and map matching on autonomous underwater vehicles. In *In Proc. IEEE Oceans*, pages 380–390, 2001.
- [25] U. Castellani, A. Fusiello, and V Murino. Registration of Multiple Acoustic Range Views for Underwater Scene Reconstruction. *Computer Vision and Image Understanding*, 87:78–89, 2002.

- [26] J. A. Castellanos, J. Neira, and J. D. Tards. Limits to the Consistency of EKF-based SLAM. In *5th IFAC Symp. on Intelligent Autonomous Vehicles, IAV'04*, Lisbon, Portugal, July 2004, July 2004.
- [27] Y. Chen and G. Medioni. Object Modelling by Registration of Multiple Range Images. *Image and Vision Computing*, 10:145–155, April 1992.
- [28] D. Chu, K. C. Baldwin, K. G. Foote, Y. Li, L. Mayer, and G. D. Melvin. Multibeam Sonar Calibration: Target Localization in Azimuth. In *IEEE Oceans*, pages 2506–2510, Honolulu, HI, 2001.
- [29] D. Chu, K. G. Foote, and L. C. Hufnagle. Measurements of Multibeam Sonar Directivity Patterns. In *IEEE Oceans*, pages 1411–1414, 2002.
- [30] P. Cignoni, C. Rocchini, and R. Scopigno. Metro: Measuring Error on Simplified Surfaces. *Computer Graphics Forum*, 17(2):167–174, 1988.
- [31] B. Curless and M. Levoy. A Volumetric Method for Building Complex Models from Range Images. *Computer Graphics (SIGGRAPH 96)*, 1996.
- [32] C de Moustier. Beyond Bathymetry: Mapping acoustic backscattering from the deep seafloor with Sea Beam. *Journal of the Acoustical Society of America*, 79(2):316–331, Feb. 1986.
- [33] C. de Moustier and M. C. Klienrock. Bathymetric Artifacts in Sea Beam Data: How to Recognize Them and What Causes Them. *Journal of Geophysical Research*, 19(B3):3407–3424, March 1986.
- [34] D di Massa. *Terrain-Relative Navigation for Autonomous Underwater Vehicles*. Phd, Woods Hole Oceanographic Institution, 1997.
- [35] A. Elfes. Using Occupancy Grids for Mobile Robot Perception and Navigation. *Computer*, 22(6):46–57, 1989.
- [36] R. Eustice, O. Pizarro, and H. Singh. Visually augmented navigation in an unstructured environment using a delayed state history. In *IEEE ICRA 2004*, April 2004.
- [37] R. Eustice, H. Singh, and J. Leonard. Exactly Sparse Delayed-State Filters. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, page To Appear, Barcelona, SPAIN, April 2005.
- [38] H. K Farr. Multibeam Bathymetric Sonars: Sea Beam and Hydrochart. *Marine Geology*, 4:77–93, 1980.
- [39] S. Fleischer. *Bounded-Error Vision Based Navigation of Autonomous Underwater Vehicles*. PhD thesis, Stanford University, 2000.
- [40] Thor Fossen. *Guidance and control of ocean vehicles*. Chichester, New York, 1994.

- [41] U. Frese and T. Duckett. A multigrid approach for accelerated relaxation-based SLAM. In *In Proc. IJCAI Workshop on Reasoning with Uncertainty in Robotics*, pages 39–46, 2003.
- [42] U. Frese and G. Hirzinger. Simultaneous localization and mapping - a discussion. In *Proc. IJCAI Workshop on Reasoning with Uncertainty in Robotics*, August 2001.
- [43] N. Gelfand, L. Ikemoto, S. Rusinkiewicz, and M. Levoy. Geometrically stable sampling for the icp algorithm. In *Proc. the International Conference on Recent Advances in 3-D Digital Imaging and Modeling*, October 2003.
- [44] G. Godin, D. Laurendeau, and R. Bergevin. A Method of Registration of Attributed Range Image. In *Proc. the International Conference on Recent Advances in 3-D Digital Imaging and Modeling*, pages 179–186, Quebec, Canada, 2001.
- [45] M. Greenspan and M. Yurick. Approximate K-D Tree Search for Efficient ICP. In *Proc. the International Conference on Recent Advances in 3-D Digital Imaging and Modeling*, Canada, 2003.
- [46] A. Gruen and D. Akca. Least Squares 3D Surface Matching. In *ISPRS Working Group V/1 Panoramic Photogrammetry Workshop, Remote Sensing and Spatial Information Science*, volume 34, Dresden, 2004.
- [47] J. Guivant and E. Nebot. Optimization of the simultaneous localization and map building algorithm for real time implementation. *IEEE Transactions on Robotics and Automation*, 17(3):242–257, June 2001.
- [48] F. Gustafsson, F. Gunnarsson, N. Bergman, U. Forssell, J. Jansson, R. Karlsson, and P.-J. Nordlund. Particle Filters For Positioning, Navigation, and Tracking. In *IEEE Transactions on Signal Processing*, volume 50, pages 425–437, Feb 2002.
- [49] J. Gutmann and K. Konolige. Incremental mapping of large cyclic environments. In *In Proc. IEEE Int. Symposium on Computational Intelligence in Robotics and Automation - CIRA-99*, November 1999.
- [50] E. Hammerstad. EM Technical Note: Multibeam Echo Sounder Accuray. Technical report, Kongsberg-Simrad, 2001.
- [51] R. Hare, A. Godin, and L. A. Mayer. Accuracy Estimation of the Canadian swath (Multibeam) and Sweep (MultiTransducer) Souding Systems. Technical report, Canadian Hydrographic Service, 1994.
- [52] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521540518, second edition, 2004.
- [53] A. J. Healey, P. E. An, and D. B. Marco. On Line Compensation of Heading Sensor Bias for Low Cost AUVs. In *Proceedings of the IEEE Symposium on Autonomous Underwater Vehicle Technology*, Cambridge, MA, Aug 1998.

- [54] Natasha Hennis. Automatic Outlier Detection in Multibeam Data. Master's thesis, Delft University of Technology, 2003.
- [55] D. R. Herlihy, T. N. Stepka, and T. D. Rulon. Filtering erroneous soundings from multibeam survey data. *International Hydrography Review*, 2:67–76, Sept 1992.
- [56] H. Hoppe, T. DeRose, and T. Duchamp. Surface Reconstruction from Unorganized Points. *Computer Graphics*, 26(2):71–76, July 1992.
- [57] B.K.P. Horn. Closed-Form Solution of Absolute Orientation Using Unit Quaternions. *Journal of the Optical Society of America*, 4(4):629–642, April 1987.
- [58] Daniel Huber. Automatic 3d modeling using range images obtained from unknown viewpoints. In *Proc. the International Conference on Recent Advances in 3-D Digital Imaging and Modeling*, pages 153–160. IEEE Computer Society, May 2001.
- [59] L. Ikemoto, N. Gelfand, and M. Levoy. A Hierarchical Method for Aligning Warped Meshes. In *Proc. the International Conference on Recent Advances in 3-D Digital Imaging and Modeling*, pages 434 – 441, 2003.
- [60] D. Jackson and C. Jone. Report of Calibration fo WHOI-DSL SM2000. Technical report, Applied Physics Laboratory, 2000.
- [61] M. Jakobsson, B. Calder, and L. Mayer. On the Effect of Random Errors in Gridded Bathymetric Complilations. *Journal of Geophysical Research*, 107(B12), 2002.
- [62] M. Jakuba and D. Yoerger. Personal communication regarding map making at the lost city hydrothermal vent cite using the abe auv.
- [63] M. V. Jakuba and D. R. Yoerger. High-Resolution Multibeam Sonar Mapping With The Autonomous Benthic Explorer (ABE). In *UUST 2003*, Durham NH, August 2003.
- [64] A Johnson. *Spin-Images: A Representation for 3-D Surface Matching*. PhD thesis, Carnegie Mellon University, 1997.
- [65] A. E. Johnson. Surface Landmark Selectionand Matching in Natural Terrain. *Computer Vision and Pattern Recognition*, 2:413–420, 2000.
- [66] A. Kalaiiah and A. Varshney. Statistical Point Geometry. In *Eurographics Symposium on Geometry Processing*, pages 107–115, 2003.
- [67] B. Kamgar-Parsi. Registratation algorithms for geophysical maps. In *IEEE OCEANS*, volume 2, pages 974 – 980, 1997.
- [68] B. Kamgar-Parsi, L.J. Rosenblum, F.J. Pipitone, and J.L. Davis, L.S.and Jones. Toward an automated system for a correctly registered bathymetric chart. *IEEE Journal of Oceanic Engineering*, 14(4):314 – 325, 1989.

- [69] J.L. Kamgar-Parsi, B. Jones and A. Rosenfeld. Registration of Multiple Overlapping Range Images: Scenes Without Distinctive Features. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(9):857–871, 1991.
- [70] R. Karlsson and F. Gustafsson. Particle filter for underwater terrain navigation. In *Statistical Signal Processing, 2003 IEEE Workshop on*, pages 526 – 529, Sept 2003.
- [71] J. Kinsey and L. Whitcomb. Towards in-situ Calibration of Gyro and Doppler Navigation Sensors for Precision Underwater Vehicle Navigation. In *Proc. International conference on robotics and automation, ICRA 2002*, pages 4016–4023, 2002.
- [72] J. C. Kinsey and L. L. Whitcomb. Preliminary Field Experience With the DVLNAV Integrated Navigation System for Manned and Unmanned Submersibles. In *In Proceedings of the 1st IFAC Workshop on Guidance and Control of Underwater Vehicles*, pages 83–88, April 2003.
- [73] James C. Kinsey and Louis L. Whitcomb. Adaptive Identification on the Group of Rigid Body Rotations. In *IEEE Conference on Robotics and Automation*, April 2005. Accepted, to appear. Preprint available at robotics.me.jhu.edu/~jkinsey/publications.html.
- [74] B.J. Kraft and C.P. de Moustier. Variable Bandwidth Filter for Multibeam Echo-sounding Bottom Detection. In *OCEANS '04. MTS/IEEE TECHNO-OCEAN*, volume 2, pages 1154–1158, 2004.
- [75] T Kweon, S. Kanade. High-resolution terrain map from multiple sensor data. *IEEE Transactions on Patterns Analysis and Machine Intelligence*, 14(2):278–292, February 1992.
- [76] J. Leal, S. Scheduling, and G. Dissanayake. Stochastic Simulation in Surface Reconstruction and Application to 3d Mapping. In *IEEE 2002 International Conference on Robotics and Automation*, pages 1765–1770, Washington DC, May 2002.
- [77] Jeffery Leal. *Stochastic Environment Representation*. PhD thesis, The University of Sydney, ACFR, 2003.
- [78] S. Lee. Constraining Navigation by Matching Swath Bathymetry and Gravity Measurements at Ship Track Crossovers. *Marine Geodesy*, 23(1):31–53, 2000.
- [79] J. Leonard, A. Bennet, C. Smith, and J. Feder. Autonomous Underwater Vehicle Navigation. Technical report, MIT Marine Robotics Laboratory Technical Memorandum 98-1, 98.
- [80] J. Leonard and P. Newman. Consistent, Convergent, and Constant-time SLAM. In *Proc. 18th Int. Joint Conf. on Artificial Intelligence (IJCAI-03)*, pages 1143–1150, 2003.
- [81] J. Leonard, R. Rikoski, P. Newman, and M. Bosse. Mapping partially observable features from multiple uncertain vantage points. *The International Journal of Robotics Research*, pages 943–975, October–November 2002.

- [82] C.B. Lirakis and K.P. Bongiovanni. Automated multibeam data cleaning and target detection. In *OCEANS 2000 MTS/IEEE Conference and Exhibition*, volume 1, pages 719 – 723, 2000.
- [83] F. Lu and E. Milios. Globally consistent range scan alignment for environment mapping. *Autonomous Robots*, 4:333–349, 1997.
- [84] X Lurton. Swath Bathymetry Using Phase Difference: Theoretical Analysis of Acoustical Measurement Precision. *IEEE Journal of Oceanic Engineering*, 25(3):351–363, July 2000.
- [85] S. Majumder, J. Rosenblatt, S. Schedling, and H. Durrant-Whyte. Map Building and Localization for Underwater Navigation. In *International Symposium On Experimental Robotics (ISER 2000)*, Honolulu, Hawaii, 2000.
- [86] Somajyoti Majumder. *Sensor fusion and feature based navigation for sub-sea robots*. PhD thesis, The University of Sydney, ACFR, 2001.
- [87] M. Mandt, K. Gade, and B. Jalving. Integrating DGPS-USBL Position Measurements With Inertial Navigation in the HUGIN 3000 AUV. In *Proc. 8th Saint Petersburg International Conference on Integrated Navigation System*, Saint Petersburg, Russia, 2001.
- [88] L. A. Mayer, B. R. Calder, J. Schmidt, and Malzone C. Providing the third dimension: High resolution Multibeam sonar as a tool for archaeological investigations an example from the D-Day beaches of Normandy. In *Conference Proceedings: U.S. Hydrographic Conference*, Biloxi MS, 2003.
- [89] H. Medwin and C. Clay. *Fundamentals of Acoustical Oceanography*. Academic Press Inc., London, Date 1998.
- [90] P. H. Milne. *Underwater Acoustic Positioning Systems*. Gulf Publishing Company, Houston, 1983.
- [91] N. Mitra and A. Nguyen. Estimating Surface Normals in Noisy Point Cloud Data. In *Symposium on Computational Geometry 2003*, pages 322–328, San Diego, June 2003.
- [92] N. J. Mitra, Gelfand N., Pottmann H., and Guibas L. Registration of point cloud data from a geometric optimization perspective. In *Symposium on Geometry Processing*, 2004.
- [93] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit. FastSLAM: A Factored Solution to the Simultaneous Localization and Mapping. In *Proc. AAAI-02*, 2002.
- [94] H. Moravec. Robot Spatial Perception by Stereoscopic Vision and 3D Evidence Grids. Technical Report CMU-RI-TR-96-34, Carnegie Mellon University, 1996.
- [95] H. Moravec and A. Elfes. High Resolution Maps from Wide-Angle Sonar. In *IEEE Conference on Robotics and Automation*, March 1985.

- [96] V. Murino and A. Trucco. Three Dimensional Image Generation and Processing in Underwater Acoustic Vision. *Proceedings of the IEEE*, 88(12):1903–1947, December 2000.
- [97] V. Murino, A. Trucco, and C.S. Regazzoni. A Probabilistic Approach to the Couples Reconstruction and Restoration of Underwater Acoustic Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(1):9–21, January 1998.
- [98] F. Napolitano, T. GaiFFE, Y. Cottreau, and T. Loret. Phins: The First High Performance Inertial Navigation System Based On Fibre Optic Gyroscopes. From the PHINS product web page: <http://www.ixseea.com>.
- [99] J. Neira and J. D. Tardos. Data Association in Stochastic Mapping Using the Joint Compatibility Test. *IEEE Transactions on Robotics and Automation*, 17(6):890–897, Dec 2001.
- [100] P. Newman and J. Leonard. Pure Range-only Sub-sea SLAM. In *IEEE Conference on Robotics and Automation*, volume 2, pages 1921 – 1926, Sept 2003.
- [101] P. M. Newman, J. J. Leonard, and R. J. Rikoski. Towards Constant-Time SLAM on an Autonomous Underwater Vehicle Using Synthetic Aperture Sonar. In *Proceedings of the Eleventh International Symposium on Robotics Research*, Sienna, Italy, Oct 2003.
- [102] C.E Nishimura and D. W. Forsyth. Improvements in Navigation using Sea Beam Crossing Errors. *Marine Geophysical Researches*, 9:333–352, 1988.
- [103] A. Nütcher, H. Surmann, K. Lingemann, and S. Thrun. 6D SLAM with an application in Autonomous Mine Mapping. In *In Proc. International Conference on Robotics and Automation*, volume 2, pages 1998–2003, April 2004.
- [104] I. Nygren and M. Jansson. Terrain Navigation for Underwater Vehicles Using the Correlator Method. *IEEE Journal of Oceanic Engineering*, 29(3):906–915, July 2004.
- [105] C.F. Olson. Landmark Selection for Terrain Matching. In *Computer Vision and Pattern Recognition*, volume 2, pages 52 – 57, 2000.
- [106] C.F. Olson. Maximum-likelihood template matching. In *Robotics and Automation*, volume 2, pages 1447 – 1452, 2000.
- [107] M. A. Paskin. Thin junction tree filters for simultaneous localization and mapping. In *In Proc. of the Eighteenth International Joint Conference on Artificial Intelligence*, pages 1157–1164, 2003.
- [108] M. Pauly, M. Gross, and L. Kobbert. Efficient simplification of point-sampled surfaces. In *Proceedings of the conference on Visualization*, pages 163–170, 2002.
- [109] X. Pennec and J. Thirion. A framework for uncertainty and validation of 3-d registration methods based on points and frames. *International Journal of Computer Vision*, 25(3):203–229, 1997.

- [110] O. Pizarro. *Large Scale Structure from Motion for Autonomous Underwater Vehicle Surveys*. PhD thesis, MIT & WHOI Joint Program in Oceanography and Oceanographic Engineering, 2004.
- [111] H. Pottmann, Q. X. Huang, Yang Y. L., and S. M. Hu. Geometry and convergence analysis of algorithms for registration of 3D shape. Technical Report Tech.Rep. 117, Geometry Preprint Series, Vienna University of Technology, 2004.
- [112] K. Pulli. Multiview Registration for Large Data Sets. In *Second International Conference on 3D Digital Imaging and Modeling*, pages 160–168, 1999.
- [113] J.J Rikoski, R.J. Leonard. Trajectory sonar perception. In *IEEE Conference on Robotics and Automation*, volume 1, pages 963– 970, 2003.
- [114] C. Roman and H. Singh. Micro-bathymetric Mapping Using Acoustic Range Images. In *IEEE/MTS Oceans'04*, pages 1574–1579, Japan, Nov 2004.
- [115] L. Rosenblum and B. Kamgar-Parsi. Reconstruction of Small Underwater Objects Using High-resolution Sonar Data. In *Autonomous Underwater Vehicle Technology*, pages 228 – 235, June 1992.
- [116] S. Rusinkiewicz and M Levoy. Efficient variants of the ICP algorithm. In *Proc. the International Conference on Recent Advances in 3-D Digital Imaging and Modeling*, pages 145–152, Canada, 2001.
- [117] C. Schlegel and T. Kampke. Filter design for simultaneous localization and map building (SLAM). In *In Proc. IEEE International Conference on Robotics & Automation*, May 2002.
- [118] T. Shank, D. Fornari, D. Yoerger, S. Humphris, A. Bradley, S. Hammond, and *et al.* Deep Submergence Synergy: Alvin and ABE Explore the Galpagos Rift at 86. *EOS*, Oct, 7 2003.
- [119] G.C. Sharp and D.K. Lee, S.W.and Wehe. Multiview Registration of 3D Scenes by Minimizing Error Between Coordinate Frames. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(8):1037–1050, 2004.
- [120] D.A. Simon, M. Hebert, and T. Kanade. Real-time 3-D Pose Estimation Using a High-speed Range Sensor. In *IEEE Conference on Robotics and Automation*, pages 2235–2241, May 1994.
- [121] H. Singh, O. Pizarro, L.L. Whitcomb, and D.R. Yoerger. In-Situ Attitude Calibration for High Resolution Bathymetric Surveys with Underwater Robotic Vehicles. In *In Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1767–1774, April 2000.
- [122] H. Singh, L.L. Whitcomb, D.R. Yoerger, and O. Pizarro. Microbathymetric Mapping from Underwater Vehicles in the Deep Ocean. *Computer Vision and Image Understanding*, 79(1):143–161, July 2000.

- [123] M. Sistiage, J. Opderbecke, and M. J. Aldon. Depth Image Matching for Underwater Vehicle Navigation. In *Image Analysis and Processing, 1999. Proceedings. International Conference on*, pages 624 – 629, Sept 1999.
- [124] M. Sistiage, J. Opderbecke, M. J. Aldon, and V. Riguard. Map Based Underwater Navigation Using a Multibeam Echosounder. In *OCEANS '98 Conference Proceeding*, volume 2, pages 747 – 751, Sept 1998.
- [125] C.M. Smith, J.J. Leonard, A.A. Bennett, and C. Shaw. Feature-Based Concurrent Mapping and Localization for Auvs. In *MTS/IEEE Conference Proceedings OCEANS '97*, volume 2, pages 896–901, October 1997.
- [126] R. Smith, M. Self, and P. Cheeseman. *Estimating uncertain spatial relationships in robotics*, pages 167–193. Autonomous Robot Vehicles, Springer-Verlag, 1990.
- [127] I. Stamos and P Allen. Geometry and Texture Recovery of Scenes of Large Scale. *Computer Vision and Image Understanding*, 88(2):94–118, 2002.
- [128] K. Stewart. *Multisensor Modeling Underwater with Uncertain Information*. Phd, Woods Hole Oceanographic Institution, 1988.
- [129] K. Stewart. A Model-Based Approach to 3-D Imaging and Mapping Underwater. *Journal of Offshore Mechanics and Arctic Engineering*, 112, 1990.
- [130] I. Tena Ruiz, de Raucourt S., Y. Petillot, and D.M Lane. Concurrent Mapping and Localization Using Sidescan Sonar. In *IEEE Journal of Oceanic Engineering*, volume 29, pages 442–456, April 2004.
- [131] S. Thrun. Learning Occupancy Grids with Forward Models. In *2001 IEEE,RSJ International Conference on Intelligent Robots and Systems*, pages 1676–1681, Maui, Hawaii, November 2001.
- [132] S. Thrun. Using Em to Learn 3D Models of Indoor Environments with Mobile Robots. In *Eighteenth International Conference on Machine Learning*, Williams College, July 2001.
- [133] S. Thrun, Y. Liu, Z. Koller, H. Ghahramani, H. Durrant-Whyte, and A. G. Ng. Simultaneous mapping and localization with sparse extended information filters. *International Journal of Robotics Research*, 23:693–716, 2004.
- [134] U. U. Castellani, A. Fusiello, V. Murino, L. Papaleo, E. Puppo, S. Repetto, and M. Pittore. Efficient On-line Mosaicing from 3D Acoustical Image. In *Int. Conf. Oceans '04 MTS/IEEE*, pages 865–872, Kobe, Japan, 2004.
- [135] R.J. Urick. *Principles of Underwater Sound*. McGraw-Hill, New York, 1975.
- [136] Eggert D. W., A. W. Fitzgibbon, and R. B. Fisher. Simultaneous Registration of Multiple Range Views for Use in Reverse Engineering of CAD Models. *Computer Vision and Image Understanding*, 69(3):253–272, 1998.

- [137] R. Whitaker and E.L. Juarez-Valdes. On the Reconstruction of Height Functions and Terrain Maps from Dense Range Data. *IEEE Transactions on Image Processing*, 11(7), July 2002.
- [138] L.L. Whitcomb, D.R. Yoerger, and H. Singh. Towards Precision Robotic Manuvering, Survey and Manipulation in Unstructured Undersea Environments. In *Robotics Research - The Eighth International Symposium*, pages 45–54, Springer Verlag, London, 1998.
- [139] L.L. Whitcomb, D.R. Yoerger, and H. Singh. Advances in Doppler-Based Navigation of Underwater Robotic Vehicles. In *Proceedings of the 1999 International Conference on Robotics and Automation*, volume 1, pages 399–406, 1999.
- [140] L.L. Whitcomb, D.R. Yoerger, and H. Singh. Combined Doppler/Lbl Based Navigation of Underwater Vehicles. In *Proceedings of the 11th International Symposium on Unmanned Unthethered Submersible Technology*, Durham, New Hampshire, May 1999.
- [141] S. B. Williams. A Terrain-aided Tracking Algorithm for Marine Systems. In *The 4th International Conference on Field and Service Robotics*, July 2003.
- [142] S. B. Williams, G. Dissanayake, and H. Durrant-Whyte. Towards Terrain-Aided Navigation for Underwater Robotics. *Advanced Robotics*, 15(5):533–550, 2001.
- [143] L. Yang, T. Taxt, and F. Albrechtsen. Bottom Detection for Multibeam Sonars with Active Contours. In *OCEANS '97. MTS/IEEE Conference Proceedings*, volume 2, pages 943–950, Oct. 1997.
- [144] D. R. Yoerger, D. S. Kelley, and J. R. Delaney. Fine-Scale Three-Dimensional Mapping of a Deep-Sea Hydrothermal Vent Site Using the Jason ROV System. *The International Journal of Robotics Research*, 19:1000–1014, 2000.
- [145] D.R. Yoerger, A. Bradley, M.-H. Cormier, and W.B.F. Ryan. Fine-Scale Seafloor Survey in Rugged Deep-Ocean Terrain with an Autonomous Robot. In *IEEE ICRA*, 2000.
- [146] D.R. Yoerger, A. Bradley, M.-H. Cormier, W.B.F. Ryan, and B.B. Walden. High Resolution Mapping of a Fast Spreading Mid-Ocean Ridge with the Autonomous Benthic Explorer. In *UUST99*, 1999.



Room 14-0551
77 Massachusetts Avenue
Cambridge, MA 02139
Ph: 617.253.5668 Fax: 617.253.1690
Email: docs@mit.edu
<http://libraries.mit.edu/docs>

DISCLAIMER OF QUALITY

Due to the condition of the original material, there are unavoidable flaws in this reproduction. We have made every effort possible to provide you with the best copy available. If you are dissatisfied with this product and find it unusable, please contact Document Services as soon as possible.

Thank you.

Some pages in the original document contain color pictures or graphics that will not scan or reproduce well.