

# Towards Autonomous Ship Hull Inspection using the Bluefin HAUV

Michael Kaess, Hordur Johannsson, Brendan Englot,  
Franz Hover, John Leonard  
Massachusetts Institute of Technology  
Cambridge, MA 02139, USA  
{kaess, hordurj, benglot, hover, jleonard}@mit.edu

In this paper we describe our effort to automate ship hull inspection for security applications. Our main contribution is a system that is capable of drift-free self-localization on a ship hull for extended periods of time. Maintaining accurate localization for the duration of a mission is important for navigation and for ensuring full coverage of the area to be inspected. We exclusively use onboard sensors including an imaging sonar to correct for drift in the vehicle's navigation sensors. We present preliminary results from online experiments on a ship hull. We further describe ongoing work including adding capabilities for change detection by aligning vehicle trajectories of different missions based on a technique recently developed in our lab.

## 1 Introduction

Bluefin Robotics and MIT have built a ship hull inspection vehicle, called the hovering autonomous underwater vehicle (HAUV) [11], shown in Figure 1. The HAUV is equipped with a Doppler velocity log (DVL) to measure velocity relative to a surface, a ring laser gyro for attitude measurements and a dual frequency identification sonar (DIDSON) [2] for imaging the structures being inspected. The vehicle gets constantly improved [12], but the sensor suite used in this work remains the same.

A drift-free position estimate is needed to ensure full coverage of the area being inspected and for reporting the exact location of detected targets. However, the HAUV only uses its gyro and DVL sensors to estimate its position and heading, resulting in drift over time. A magnetic compass is not being used, because they perform poorly near the large metal structures of ship hulls. The vehicle estimates




Figure 1: Top view of the Bluefin-MIT hovering autonomous underwater vehicle (HAUV). The vehicle is equipped with a Doppler velocity log (DVL), an imaging sonar, a ring laser gyro (internal), an optical camera and a light strobe. The sonar and DVL can be actuated independently to optimally align the sensors to the surface being inspected.

its heading by integrating the rotation rates from the gyro. The position is then estimated by dead reckoning using its heading estimate together with the velocities reported by the DVL. Even though the sensors locally provide very accurate measurements, small errors will accumulate over time and the position estimate will drift without bounds.

One solution to eliminating drift would be the use of external infrastructure for global positioning. However, that is difficult underwater and in particular in harbor environments. GPS signals do not propagate through the water, instead requiring deployment of acoustic beacons. While acoustic beacons are frequently being used for accurate localization in deep water, it is challenging to do so in shallow and cluttered harbor environments. To avoid these complications and simplify deployment of the vehicle, we focus on using only onboard sensors for navigation.

Using only onboard sensors for navigation, we are left with using the environment structure itself, in this case the ship hull, for global positioning. Using the imaging sonar we can observe sparse features on the ship hull caused by protruding objects such as anodes and welding lines. Re-observing these features by revisiting previously observed parts of the ship hull, we can correct for drift accumulated in the navigation estimate. However, initially we do not know the location of the




Figure 2: Two imaging sonar frames from the King Triton vessel with extracted features marked in red and green, respectively. On the right hand side the second frame is shown again with features aligned from both frames.

features. The problem of creating a map of those features while using them for localization is called simultaneous localization and mapping (SLAM) and has been extensively studied in the robotics literature, see [5] for a survey.

In this paper we present our work on imaging sonar registration combined with advanced SLAM techniques to achieve drift-free navigation. In previous work, Walter et al. [13] used manually extracted landmarks for navigation. An automatic feature detector and a landmark formulation using an EKF filter was used in [6]. Our main contribution is an automated dense feature extraction and matching that allows for effective registration and loop closures. As we use dense feature matching, we do not explicitly model landmarks, requiring a different solution to the SLAM estimation problem based on a smoothing approach.

## 2 Drift-Free Navigation

Our approach to drift-free navigation consists of two main components: Sonar scan matching to obtain geometric alignment between sonar frames, and the actual pose estimation that solves the resulting SLAM problem.

### 2.1 Sonar Scan Matching

Our system allows for drift-free navigation along the flat parts of the hull without depending on any external infrastructure. In this paper we assume that the areas

being inspected are roughly horizontal, restricting us to locally flat surfaces, which cover large parts of large ship hulls. The dense feature extraction allows aligning of frames based on existing laser scan matching techniques for land-based vehicles, and is particularly useful to disambiguate between places when finding large loop closures. An example of extracted and aligned features is shown in Figure 2.

We extract reasonably stable features from a sonar image based on strong gradients. The image is first smoothed using a median filter to reduce noise. Next the gradient is calculated and all points above a threshold accepted. Finally the points are clustered and small clusters are discarded. The remaining clusters are used for registration with another frame.

Registration of the features extracted from two frames is performed using the normal distribution transform (NDT) algorithm by Biber and Strasser [3]. The NDT algorithm uses a grid to discretize space. For each cell the mean and variance of the points that lie inside the cell are calculated. Four overlapping grids are used that are shifted to alleviate the effect of discontinuities resulting from discretization. The NDT provides a compact representation suitable for optimization and does not require known correspondences between the features. Successful registrations are determined based on a score obtained by evaluating the corresponding Gaussian at each point. The registration provides the vehicle transformation between the times when the two images were taken.

## 2.2 Pose Estimation

Successful registrations between sequential pairs of poses as well as arbitrary pairs for loop closing provide constraints that are combined into an estimation problem to correct for drift. Simply accumulating sequential measurements leads to drift. Loop closing constraints are obtained by scan matching the current frame to frames encountered when previously visiting the same area of the hull. These constraints can be combined into a least-squares formulation that takes all measurements into account simultaneously.

To allow for real-time operation, we use an efficient formulation of the underlying optimization problem. Note that traditional filtering-based methods cannot be applied here, as that formulation does not allow to include constraints to previously estimated poses. Instead of earlier EKF formulations, we use a smoothing method that estimates the complete trajectory instead of only the current pose. This allows for a sparse formulation that actually simplifies the problem, while avoiding approximations introduced by the EKF. In particular, we use an online algorithm called incremental smoothing and mapping (iSAM) [7].

iSAM maintains the square root information matrix, which can be obtained by matrix factorization of the information matrix or the underlying measurement Jaco-

bian. The key to an efficient solution is to find a good variable ordering. Depending on the variable order, the square root matrix can have a smaller or larger number of entries. Finding the best ordering is prohibitively expensive, but good heuristics are available such as the column approximate minimum degree ordering (COLAMD) algorithm [4]. Instead of factoring the matrix each time a new measurement is taken, iSAM updates the existing matrix factorization using Givens rotation, providing a real-time solution even for larger problems than will be encountered during ship hull inspection.

### 3 Experiments and Results

We present results from experimental evaluation on the King Triton vessel in Boston harbor. The HAUV autonomously navigated for 45 minutes in a rectangular pattern of 1 meter by 2 meter underneath the ship. The results in Figure 3 show that drift was successfully corrected. In red color the trajectory according to the vehicle's dead reckoning estimate is shown as well as two manually selected features for reference. Blue color shows the data obtained online using our drift correction.

While the slow long-term drift is introduced by the gyro and DVL combination, our approach also compensated for the motion of the vessel itself. The vessel moved visibly in the swell even though it was tied down. Our approach allows correcting for this motion without requiring the DVL to point to the ship hull. We have generally found it advantageous to point the DVL down to avoid interruptions when some or all of the DVL's sonar beams move outside the boundaries of the hull.

We have recently performed further experiments, one set of which was run in a tank in our lab. The small tank environment provides a difficult setting for the DVL. Multipath leads to significantly more drift, which was successfully corrected by our approach. We also surveyed a 10 meter by 20 meter area at the bottom of the Charles river for a duration of two hours, while drift was successfully corrected.

During our experiments, for simplicity we currently use the onboard processor only for the actual control, while all processing is performed on shore. The vehicle is tethered using an optic fiber, and laptops are used for processing. This allows us to focus on algorithms rather than optimization and allows for rapid development. It should be noted that for real deployment of the vehicle the algorithms could be run on a second onboard computer. Instead of sending back all sensor data, select information such as the map and individual scans showing potential targets can be sent back using acoustic modems. Because of the shallow and cluttered workspace, special efforts are needed for the acoustic communication [1] that are also funded under the same project. Target detection is performed using software from Seebyte [10], also funded under this project.

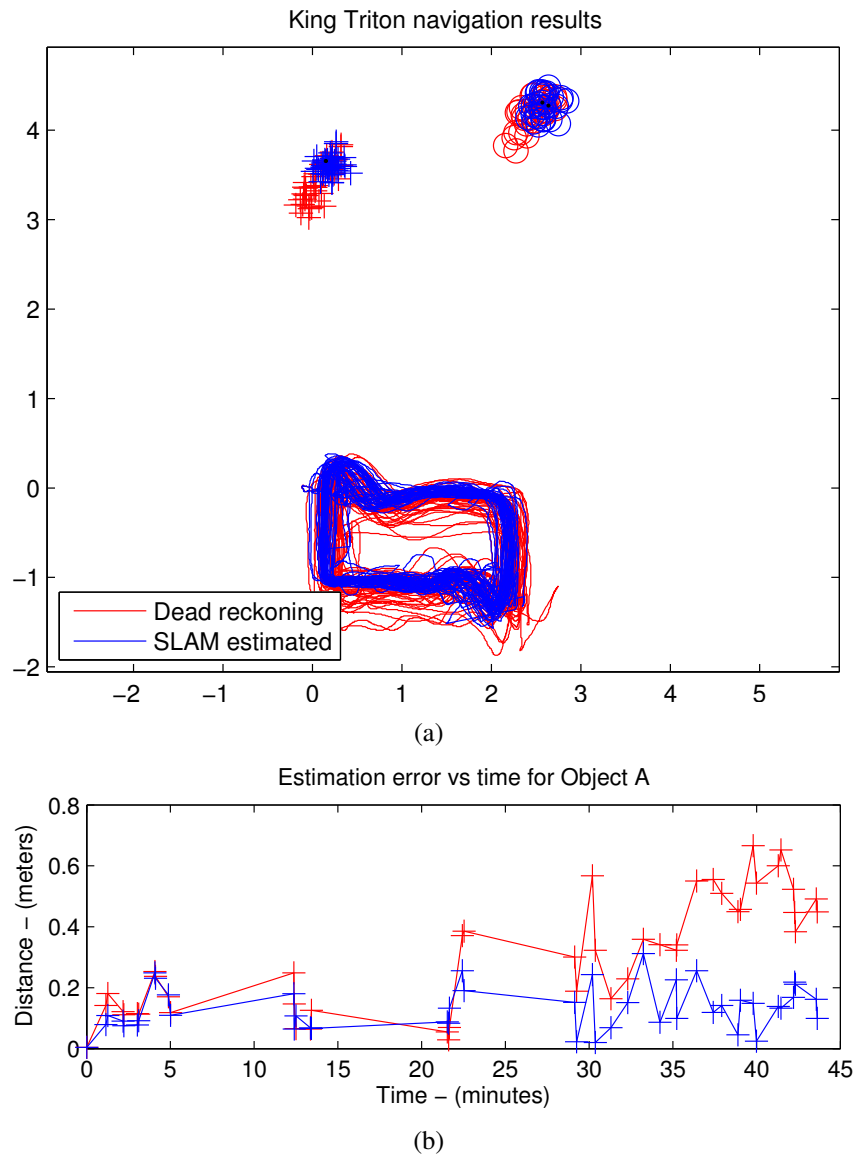


Figure 3: Navigation result from the King Triton vessel in Boston harbor. (a) The rectangular trajectory of the vehicle before (red) and after (blue) correction is shown, as well as the location of two reprojected features. As ground truth position is difficult to obtain underneath the ship, we have manually labeled two features in the sonar images and show their reprojections from many frames. The data obtained from our system (blue) is fairly stable, while the result based on the internal estimate (red) shows significant drift. (b) The drift of one of the features is shown over time.




Figure 4: Complex parts of the ship hull require a different approach using profiling sonar. (a) The rudder and propeller of a large ship. (b) Uncorrected point cloud of the complex area of the Oceanus vessel at WHOI in Woods Hole, MA.

## 4 Ongoing Work

There are multiple ongoing efforts to extend our current system to fully autonomous ship hull inspection of the complete hull. We are also interested in combining sonar and visual data for navigation, and in performing change detection.

### 4.1 Change Detection

When a ship hull has already been mapped in a previous session, the recorded data can be used to detect changes on the hull, making it easier to identify potential targets. The same feature extraction and matching process is used to align the current vehicle trajectory with previously recorded data. We propose to use the concept of anchor nodes [9] to align the trajectory of an ongoing run online with recorded data. After successful alignment, the sonar data can be compared to detect changes.

### 4.2 Complex Areas

Complex areas of the ship hull require a different strategy using profiling sonar instead of imaging. Our current approach requires nearly flat surfaces because of the ambiguity inherent in the imaging sonar data, i.e. a point in the sonar image does not correspond to a single point in space. Profiling sonar on the other hand does not have this ambiguity, as the opening angle of the sonar is very small, close to one degree. Switching the DIDSON sonar sensor between imaging and profiling mode only requires replacing its acoustic lens.




Figure 5: Images from the HAUV camera providing additional details useful for both navigation and inspection. We plan to combine sonar and vision data for navigation. The images can additionally help an operator to identify a detected object.

With profiling, the individual scans do not overlap, and multiple scans have to be accumulated into a small submap that can then be used for matching in order to correct drift over time. The small submaps are very accurate because of the high quality of the navigation sensors. Instead of scan matching, we align the 3D points of the submaps. Drift correction is then performed using the same estimation technique as before. An uncorrected point cloud of the complex parts of a ship hull is shown in Figure 4.

### 4.3 Combining with Camera

When the water quality allows, we can additionally make use of the camera carried by the HAUV. A flash bulb on top of the HAUV provides the necessary lighting underneath a ship. Ryan Eustice at the University of Michigan provides the visual SLAM capabilities [8]. We are planning to combine data from both camera and imaging sonar into one common framework. The advantage of combining both sensors are twofold: First we can expect higher accuracy, as multiple independent measurements become available. Second, the acoustic and visual sensors complement each other in that they observe different features: A part of the ship hull that is featureless for the sonar might still contain visual features, such as differences in paint, and the same is true the other way around. Camera images as shown in Figure 5 can also be useful for an operator to identify a detected object.

### 4.4 Planning for Coverage

We are actively developing algorithms to achieve complete sensor coverage of complex, three-dimensional structures surveyed by an autonomous agent with multiple degrees of freedom. Input to the algorithm is a closed triangular mesh that can either






Figure 6: Coverage planning for the complex parts of the ship hull. (a) A model of the propeller and rudder. The ship hull is shown in green and used for collision avoidance. (b) A vehicle trajectory that ensures full coverage obtained by ongoing work in our lab.

be obtained from blueprints of the vessel, or by a separate run using the profiling sonar, starting from an approximate shape of the ship hull. An example trajectory calculated by our approach is shown in Figure 6.

## Acknowledgments

This work is supported by the Office of Naval Research under grant N00014-06-10043, monitored by Dr. T.F. Swean, M. Zalesak and V. Steward. We would like to thank J. Vaganay, K. Shurn and M. Elkins from Bluefin Robotics for their great support during the course of the experiments.

## References

- [1] P.-P.J. Beaujean, E.A. Carlson, J. Spruance, and D. Kriel. HERMES - a high-speed acoustic modem for real-time transmission of uncompressed image and status transmission in port environment and very shallow water. In *Proceedings of the IEEE/MTS OCEANS Conference and Exhibition*, pages 1–9, 2008.
- [2] E. Belcher, W. Hanot, and J. Burch. Dual-frequency identification sonar (DID-SON). In *Underwater Technology, 2002. Proceedings of the 2002 International Symposium on*, pages 187–192, 2002.
- [3] P. Biber and W. Strasser. The normal distributions transform: a new approach to laser scan matching. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, volume 3, pages 2743–2748, Oct 2003.

- [4] T.A. Davis, J.R. Gilbert, S.I. Larimore, and E.G. Ng. A column approximate minimum degree ordering algorithm. *ACM Trans. Math. Softw.*, 30(3):353–376, 2004.
- [5] H.F. Durrant-Whyte and T. Bailey. Simultaneous localisation and mapping (SLAM): Part I. *Robotics & Automation Magazine*, 13(2):99–110, Jun 2006.
- [6] B. Englot, H. Johannsson, and F. Hover. Perception, stability analysis, and motion planning for autonomous ship hull inspection. In *Proceedings of the International Symposium on Unmanned Untethered Submersible Technology (UUST)*, 2009.
- [7] M. Kaess, A. Ranganathan, and F. Dellaert. iSAM: Incremental smoothing and mapping. *IEEE Trans. Robotics*, 24(6):1365–1378, Dec 2008.
- [8] A. Kim and R. Eustice. Pose-graph visual SLAM with geometric model selection for autonomous underwater ship hull inspection. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, Oct 2009.
- [9] B. Kim, M. Kaess, L. Fletcher, J. Leonard, A. Bachrach, N. Roy, and S. Teller. Multiple relative pose graphs for robust cooperative mapping. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 3185–3192, Anchorage, Alaska, May 2010.
- [10] P.Y. Mignotte, S. Reed, A. Cormack, and Y. Petillot. Automatic ship hull inspection - the detection of mine-like targets in sonar data using multi-CAD fusion and tracking technologies. In *Proceedings of Institute of Acoustics International Conference on the Detection and Classification of Underwater Targets*, Sep 2007.
- [11] J. Vaganay, M. Elkins, D. Esposito, W. O’Halloran, F. Hover, and M. Kokko. Ship hull inspection with the HAUV: US Navy and NATO demonstrations results. In *Proceedings of the IEEE/MTS OCEANS Conference and Exhibition*, volume 1, pages 761–766, 2007.
- [12] J. Vaganay, L. Gurfinkel, M. Elkins, D. Jankins, and K. Shurn. Hovering autonomous underwater vehicle - system design improvements and performance evaluation results. In *Proceedings of the International Symposium on Unmanned Untethered Submersible Technology (UUST)*, 2009.
- [13] M. Walter, F. Hover, and J. Leonard. SLAM for ship hull inspection using exactly sparse extended information filters. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 1463–1470, 2008.