

PERCEPTION, STABILITY ANALYSIS, AND MOTION PLANNING FOR AUTONOMOUS SHIP HULL INSPECTION

Brendan Englot, Hordur Johannsson, and Franz Hover
Massachusetts Institute of Technology

77 Massachusetts Avenue, Cambridge MA 02139 (617)-324-3994, benglot@mit.edu, hordurj@mit.edu, hover@mit.edu

Abstract— Recent progress in the use of feature-based navigation for ship hull inspection performed by an autonomous underwater vehicle is discussed. Localization and map-building results using both imaging and bathymetry sonar are presented. We then assess the dynamic stability of the closed-loop localization and feedback control process, and formulate a robustness-optimal motion planning framework designed to maximize a marine vehicle’s stability margins against ocean disturbances.

I. INTRODUCTION

Feature-based navigation methods involving localization, map-building, and simultaneous localization and mapping (SLAM) have become ubiquitous as a means of perception for autonomous mobile robots [1], [2], [3]. Using the natural features of the surrounding environment for navigation is especially important in applications where traditional odometry and direction sensors are unavailable. One such example is ship hull inspection performed by an autonomous underwater vehicle, in which sonar imaging and range-sensing present cost-effective alternatives to high precision inertial navigation, and underwater operation near a large steel structure prevents use of compass, GPS, or long baseline acoustic tracking as consistent sensors. In addition, random disturbances in the ocean environment necessitate the use of high-fidelity feedback control, and hence we not only wish to utilize feature-based methods for estimation, but we seek a better understanding of integrated localization, mapping, and dynamic control.

In our approach to localization and mapping for hull inspection, we first present the results of an extended Kalman filter SLAM framework that was implemented successfully during hull inspection exercises at AUVFest 2008. The MIT-Bluefin HAUV [4], pictured in Figure 1, identified a series of mine-shaped training targets using DIDSON imaging sonar [5] and mapped these targets in real-time. We also present the results of mapping using BlueView bathymetry sonar, which provides an informative view of running gear and other complex hull structures.




Figure 1: The MIT-Bluefin Hovering Autonomous Underwater Vehicle, designed to perform autonomous ship hull inspections.

In our consideration of dynamic control, we then expand on previous analyses of the integrated localization, mapping, and control process [6], [7], [8], which have considered single and three degree-of-freedom vehicles governed by constant-gain controllers and estimators. In the present work, motivated by the application of ship hull inspection, we consider a planar marine vehicle using range and bearing measurements of a set of point features to traverse a path expansive enough to require time-varying controller and estimator gains. The stability of the controller and estimator duo is investigated using a pair of theorems requiring boundedness and convergence of the transition matrix Euclidean norm. Perturbations are then considered using a theorem on the convergence of the perturbed system transition matrix, yielding a robustness test for the closed-loop system. Specifically, we show that the robustness of a marine vehicle varies as a function of the sequence by which the vehicle observes its surrounding features. Together, these stability and robustness tests form a set of tools which can be used in planning and evaluating the robustness and performance of marine vehicle survey trajectories.

Finally, the stability analysis is utilized in the development of a path-planner which uses an augmented version of the A* algorithm [9] to plan survey trajectories which deviate from the shortest path to improve the marine vehicle’s robustness to perturbations.

II. REAL-TIME FEATURE-BASED MAP-BUILDING

At AUVfest 2008, an event organized for collective experiments and demonstrations of Navy-sponsored AUV technology, the MIT-Bluefin HAUV performed a hull inspection exercise aimed at the identification of mine-shaped training targets planted on a ship hull. The task was performed on the hull of the *USS Saratoga*, a retired naval aircraft carrier. The HAUV performed surveys along a portion of the hull very similar in curvature and orientation to a flat vertical wall. The HAUV used its IMU, DVL, and depth sensor to localize along with layered proportional-integral-derivative (PID) control to correct perceived errors as it carried out the hull survey. While these standard hull survey processes were performed using the vehicle’s internal computer, the feature extraction and map-building processes were carried out on a topside computer. The computer intercepted real-time image data from the DIDSON, odometry data from the DVL, and used these to produce real-time localization and map estimates. Two algorithms were run in series to produce these estimates, a feature extraction algorithm which identified the vehicle-relative range and bearing of mine-shaped training targets observed in DIDSON imagery, and an Extended Kalman Filter algorithm which recursively estimated the vehicle pose and velocity, and the location of each target on the hull.

A. Feature Extraction Algorithm

The goal of the feature extraction algorithm is to identify mine-shaped training targets and approximate their location by designating a single entry of a DIDSON frame’s image matrix to represent a sighted target. Each entry of the matrix, which consists of an intensity value, corresponds to a vehicle-relative range and bearing identified by the row and column of the entry. If a color spectrum is assigned to the intensity values in the image matrix, the image will appear similar to the example in Figure 2, which contains one of the mine-shaped training targets used on the *USS Saratoga*. To identify the targets, knowledge of their intensity signature in DIDSON imagery is used, specifically the high intensity returned by the target itself and the low intensity returned by the target’s “shadow”. The significant contrast between the shadow and the surrounding image data, especially the contrast between the shadow and the target itself, can be exploited to produce an algorithm which runs fast enough to allow mapping to be executed in real-time.

Sections of the image matrix are evaluated iteratively to obtain a mean difference in intensity between a small rectangular portion of the matrix (designed to capture the low intensity of the shadow) and the ring that immediately surrounds it (designed to capture the high intensity just

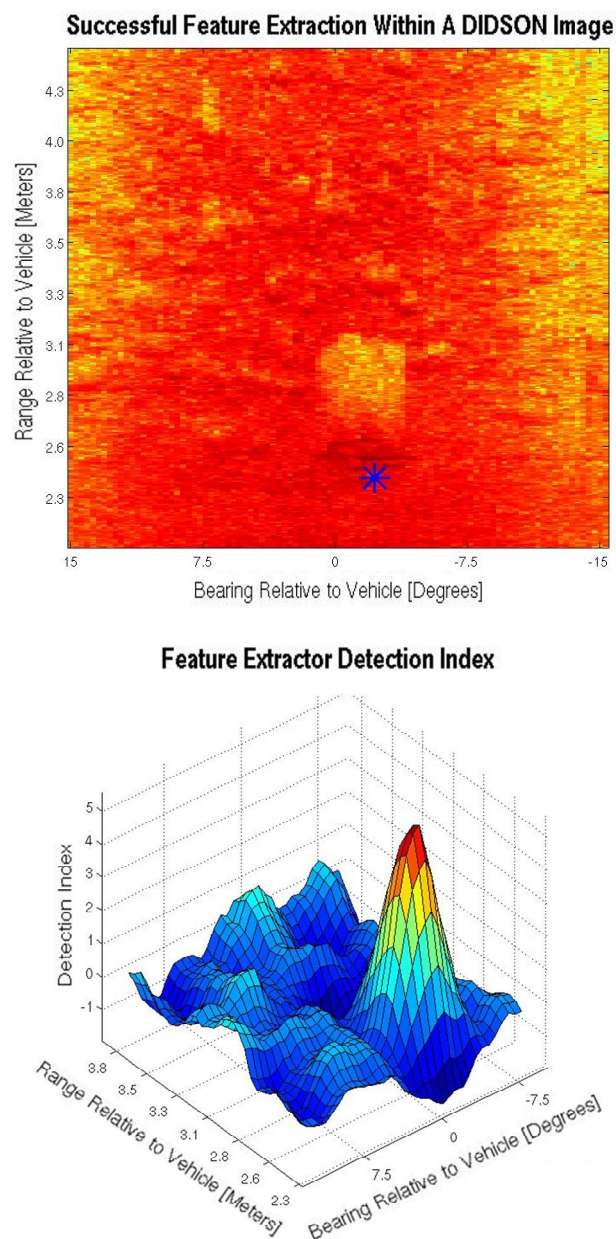


Figure 2: Performance of real-time feature extractor is demonstrated using two example DIDSON frames. The raw data is presented at left, and the feature extractor detection index for each rectangular quadrant of the image that was processed is displayed at right. Areas where features were identified (indicated by the blue asterisk) correspond to high peaks in the feature detection index.

outside of the shadow). This heuristically-designed feature extraction technique yielded selective and consistent identification of training targets, although more advanced and widely applicable methods are available and in development by our collaborators as our mapping exercises grow in complexity [10].

B. Localization and Mapping Algorithm

An EKF is employed to estimate the pose and velocity of the HAUV and to construct the map of hull features. The EKF approach was selected to enable the probabilistic

mapping of the training targets while conforming to the framework of our feedback design approach. Although the HAUV is a six degree-of-freedom vehicle subjected to hydrodynamic drag, a number of approximations are made to keep the vehicle model as simple as possible. First, a planar three degree-of-freedom vehicle model is used since the HAUV is surveying an approximately vertical wall at fixed range, fixed heading, and zero pitch. As a result, only the roll angle and x-y position are necessary to locate the vehicle and determine its orientation relative to the hull. In addition, the hull survey run on the Saratoga consists of horizontal and vertical straight-line trajectories only, with no planned variation in roll angle. For this reason, the vehicle's orientation in roll will be decoupled from its dynamics in x and y. In reality the vehicle will be perturbed in range, heading, pitch, and roll, but because the HAUV layered control acts to correct any errors in these degrees of freedom, they are assumed to be of constant magnitude as described above (although we still estimate the roll angle for the purposes of monitoring angular perturbations). One final simplification is the approximation that the vehicle dynamics consists of double integrators in each of its three planar, uncoupled degrees of freedom (i.e., hydrodynamic drag is neglected). Since this model will not be used to compute a control action for the vehicle and is being used for estimation only, this assumption will not hinder performance of the algorithm. Given these assumptions, the HAUV dynamic equations, which use an Euler discretization, appear as follows:

$$\begin{bmatrix} u_{k+1} \\ v_{k+1} \\ \dot{\varphi}_{k+1} \\ x_{k+1} \\ y_{k+1} \\ \varphi_{k+1} \end{bmatrix} = \begin{bmatrix} I_{3 \times 3} & 0_{3 \times 3} \\ I_{3 \times 3} \Delta T & I_{3 \times 3} \end{bmatrix} \begin{bmatrix} u_k \\ v_k \\ \dot{\varphi}_k \\ x_k \\ y_k \\ \varphi_k \end{bmatrix} + \begin{bmatrix} I_{3 \times 3} \\ 0_{3 \times 3} \end{bmatrix} \begin{bmatrix} w_1 \\ w_2 \\ w_3 \end{bmatrix} \quad (1)$$

$$\underline{x}_{v|k+1} = F \underline{x}_{v|k} + \Gamma \underline{w}_k$$

The body-referenced sway velocity, heave velocity, and roll rate are described by u , v , and φ , respectively, and x , y , and φ represent the horizontal, vertical, and angular position of the vehicle relative to the ship hull. Process noise w_i , which is zero mean Gaussian white noise with diagonal covariance matrix Q , is applied in each degree of freedom. The notation \underline{x}_v refers to a column vector which contains all six vehicle states.

Features observed on the hull are approximated as point features defined by an x - y position. Because the EKF is used to estimate both the vehicle states and the feature locations, the feature locations must be included in the state vector. The features are assumed to be permanently fixed to the ship hull and have no dynamics. The aggregate state vector containing both vehicle and feature

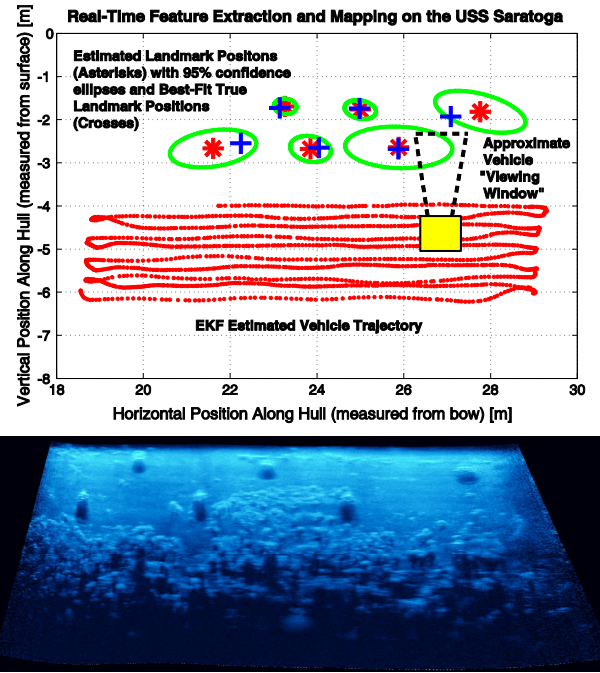


Figure 3: Real-time map and vehicle localization data obtained from a survey of the USS Saratoga in May 2008 using an Extended Kalman Filter. At bottom, a sonar mosaic image of the targets placed on the ship hull (courtesy of AcousticView <http://www.acousticview.com>).

states has the following structure:

$$\underline{x}_k = [\underline{x}_{v|k} \quad x_1 \quad y_1 \quad x_2 \quad y_2 \quad \dots \quad x_n \quad y_n]^T \quad (2)$$

The size of the system transition matrix F must be increased to accommodate these new states. Because the features are static, two rows and columns of zeros must be added to F for each new feature, with a single entry of magnitude one along the diagonal to propagate the constant value of each static feature state. Although the above state vector is depicted with n features, the number of features in the state vector changes over the course of the algorithm. The algorithm is initialized with zero features in the state vector, and new features are added as they are observed by the vehicle. Features are never removed from the state vector, regardless of the time passed since they were last sighted.

The measurement process returns DVL odometry and feature sightings extracted from DIDSON data. The feature measurements consist of range and bearing from the vehicle to the training targets, which are nonlinear functions of the system states in (2) and require the use of the extended rather than linear Kalman filter. The EKF algorithm produces an estimate of the state vector given by (2) and the associated estimation error covariance matrix.

C. Experimental Results

The HAUV survey covered a section of the hull approximately 12m by 2m in size on which nine mine-




Figure 4: BlueView MB2250 bathymetry sonar data collected from an HAUV survey of a 65' retired Coast Guard Inland Buoy Tender, featuring 3D visualization of the running gear and keel cooling pipes. For comparison, an imaging sonar mosaic of the same vessel (provided courtesy of AcousticView).

shaped training targets were mounted; a sonar mosaic displaying the layout of the targets underwater is given by Figure 3. Biofouling on the hull hindered the visibility of three targets, and so only six were used in the real-time localization and mapping exercise. Although localization was performed concurrently with the mapping process during the hull surveys, the sparsity of feature sightings required the localization process to depend heavily on vehicle odometry from the DVL. Results from a representative hull survey are displayed in Figure 3. The EKF-estimated vehicle trajectory is plotted alongside the map of training targets, with red asterisks marking the estimated feature locations. Ninety-five percent confidence ellipses derived from the error covariance matrix are plotted around each feature on the map. In the cases of features sighted thirty to forty times, the filter's confidence in the feature estimate yielded an ellipse of about 0.5m in diameter. Features with the largest ellipses were sighted only a handful of times, with confidence ellipses as large as 2m in size.

III. 3D MAPPING WITH BATHYMETRY SONAR

The HAUV has also been deployed with bathymetry sonar to produce 3D rather than 2D hull visualizations, using both the BlueView 3D MicroBathymetry (MB2250) sonar and the DIDSON with a one-degree concentrator lens.

A. Profile extraction

When running in bathymetry mode the sonar is pointed directly up towards the ship hull, giving the best resolution for measuring range to the hull. The sonar returns a series of range values (of varying intensity) corresponding to each of several vehicle-relative bearing values, and the range value with the largest intensity is selected to represent its respective bearing value. Using the coordinate transformation between the sonar (which is pitched relative to the vehicle) and the HAUV (whose angular position is estimated using integrated gyro data), the range data is mapped in 3D space. Our most recent bathymetry-enabled mapping results, displayed in Figure 4, pair the range data with vehicle positions obtained from integrated DVL data. Use of the bathymetry sonar allows 3D visualization of the running gear and keel cooling pipes on the hull of a 65-foot buoy tender.

B. Submapping and Registration

Over a short period of time the HAUV can accurately report its location using its ring laser gyro and DVL, but the accuracy of its position estimate gradually drifts when relying solely upon dead reckoning (drift of the current IMU is approximately 10 degrees/hr). Current work in progress to improve the accuracy of HAUV bathymetric mapping includes use of a submapping algorithm, based on work by Roman et al. [11], which uses sonar

bathymetry data to bound long-term navigation drift. The algorithm creates submaps by combining a sequence of sonar scans into a single point cloud. To improve localization accuracy the mission must contain overlap between submaps, and when overlap occurs we can register two adjacent submaps together to obtain a constraint between two points in the vehicle trajectory. These constraints can then be used to optimize the estimate of the vehicle trajectory and bound the vehicle's position error.

A method commonly used to register two point clouds is the iterative closest point algorithm (ICP) [12], [13]. In its simplest form it consists of finding the correspondences among two point clouds, solving for the coordinate transformation that minimizes the cost function, applying the transformation, and repeating the process until the change in error falls below a designated threshold. Using the information from the vehicle navigation system and the registration of overlapping submaps, we can then form a pose graph representing the vehicle trajectory over the entire hull survey, accompanied by a more accurate map of the hull.

IV. DYNAMIC STABILITY OF AUV SURVEY TRAJECTORIES

Turning our attention to dynamic control, we proceed with an analysis aimed at evaluating the stability of an HAUV survey trajectory which utilizes a feature-based estimation process similar to that of Section II.

A. Marine Vehicle Model

Seeking a more descriptive dynamic model for the purposes of stability analysis, we now present the holonomic marine vehicle dynamics using the following discrete time equations:

$$\begin{bmatrix} u_{k+1} \\ v_{k+1} \\ \dot{\varphi}_{k+1} \\ x_{k+1} \\ y_{k+1} \\ \varphi_{k+1} \end{bmatrix} = \begin{bmatrix} u_k + \Delta T b_u(u_k)/m \\ v_k + \Delta T b_v(v_k)/m \\ \dot{\varphi}_k + \Delta T b_\varphi(\dot{\varphi}_k)/J \\ x_k + \Delta T(u_k \cos \varphi_k - v_k \sin \varphi_k) \\ y_k + \Delta T(u_k \sin \varphi_k + v_k \cos \varphi_k) \\ \varphi_k + \Delta T \dot{\varphi}_k \end{bmatrix} + \begin{bmatrix} I_{3 \times 3} \\ 0_{3 \times 3} \end{bmatrix} \begin{bmatrix} U_1 \\ U_2 \\ U_3 \\ J \end{bmatrix} + \begin{bmatrix} I_{3 \times 3} \\ 0_{3 \times 3} \end{bmatrix} \begin{bmatrix} w_1 \\ w_2 \\ w_3 \end{bmatrix}$$

$$\underline{x}_{v|k+1} = f(\underline{x}_{v|k}) + B \underline{u}_k + \Gamma \underline{w}_k \quad (3)$$

The body-referenced forward velocity, sway velocity, and yaw rate are described by u , v , and $\dot{\varphi}$, respectively, and x , y , and φ represent the horizontal, vertical, and angular position of the vehicle in the inertial plane. Hydrodynamic drag b is expressed as a function of velocity in each degree of freedom, and vehicle mass and rotational inertia are described by m and J . Surge, sway, and yaw commands are applied to the channels U_1 , U_2 , and U_3 , respectively, and process noise w_i , which is zero mean Gaussian white noise with diagonal covariance

matrix Q , is also applied to each channel. As in Section II, the six vehicle states will be combined with $2n$ feature states in the aggregate state vector expressed in (2).

A nominal trajectory is generated for the vehicle to send it to a desired waypoint from its starting position in the plane. An open-loop input trajectory delivers a nominal command at each time step, and a closed-loop control correction is used to counteract disturbances. The closed-loop system and measurement dynamics is given by:

$$\begin{aligned} \underline{x}_{k+1} &= f(\underline{x}_k) + B(\underline{u}_{OLk} - G_k N \delta \hat{\underline{x}}_k) + \Gamma \underline{w}_k \\ \underline{z}_{k+1} &= h(\underline{x}_{k+1}) + \underline{v}_{k+1} \end{aligned} \quad (4)$$

The term $\delta \hat{\underline{x}}(k)$ represents the deviation of the state estimate from the nominal state trajectory, used as an error signal for the controller. The nonlinear functions $f(x)$ and $h(x)$ are used to represent, respectively, the state transition relationships in (3) and the nonlinear measurement of range and bearing relative to each of the features. The sensor noise term \underline{v}_k represents zero mean Gaussian white noise with diagonal covariance matrix R . G_k is a time-varying matrix of controller gains, which we compute optimally using the discrete-time matrix Riccati equation. N is a 6 by $6+2n$ stripping matrix needed to extract the vehicle states from the state vector, discarding the $2n$ feature states for the purposes of control. Because the feature states have no dynamics, the lower part of $f(x)$ is the identity, the lower part of B is zero, and the lower part of Γ is zero.

The use of a nominal vehicle trajectory permits a linearized Kalman filter to serve as the estimator for vehicle localization. This strategy allows the vehicle to move to any desired location in the plane, as long as an approximate layout of features is known in advance. It is also assumed that feature association can be performed successfully, and that ΔT between measurements is constant. These assumptions will allow vehicle pose estimation and refinement of a feature-based map to occur using a precomputed set of gains and Jacobians. The estimation equation is written in terms of deviation from the nominal trajectory, \bar{x}_k :

$$\delta \hat{\underline{x}}_{k+1} = \delta \hat{\underline{x}}_{k+1|k} + K_{k+1}[\delta \underline{z}_{k+1} - H(\bar{x}_{k+1})\delta \hat{\underline{x}}_{k+1|k}]$$

$$\text{where } \delta \hat{\underline{x}}_{k+1|k} = F(\bar{x}_k)\delta \hat{\underline{x}}_k - B G_k N \delta \hat{\underline{x}}_k \quad (5)$$

The nonlinear state transition function $f(x)$ and the measurement function $h(x)$ are now replaced by their corresponding Jacobians $F(\bar{x}_k)$ and $H(\bar{x}_k)$, which are linearized about the nominal trajectory at each time step. The term $\delta \underline{z}_k$ represents the deviation of the measurement from the deterministic measurement along the nominal trajectory. K_k is the time-varying Kalman gain, which, as per the linearized Kalman filter framework, is computed in advance along each step of

the nominal trajectory (see Gelb [14]). Although an EKF is likely to yield better estimation in the presence of perturbations, its nonlinearity and dependence on the vehicle's noise-influenced trajectory do not allow linear matrix computation or computation in advance of the vehicle's deployment. Thus the LKF will serve as our estimator at the risk of inaccuracy in the presence of large perturbations and with the benefit of enabling a more descriptive stability analysis.

Thus far the only linearization approximations have been those which are called for specifically by the linearized Kalman filter. To enable our stability analysis, it will be further assumed that dynamics of the true physical plant are well approximated by the state transition Jacobian, and that the true measurement process is also well approximated by the measurement Jacobian. Simplification of (4) and (5) yields the following compact formulation:

$$\begin{bmatrix} \delta \underline{x}_{k+1} \\ \delta \tilde{\underline{x}}_{k+1} \end{bmatrix} = \begin{bmatrix} F(\bar{x}_k) - BG_k N & -BG_k N \\ 0 & F(\bar{x}_k) - E_k \end{bmatrix} \begin{bmatrix} \delta \underline{x}_k \\ \delta \tilde{\underline{x}}_k \end{bmatrix}$$

$$\begin{aligned} \text{where } E_k &= K_{k+1} H(\bar{x}_{k+1}) F(\bar{x}_k) \\ \delta \tilde{\underline{x}}_k &= \delta \underline{x}_k - \delta \hat{\underline{x}}_k \end{aligned} \quad (6)$$

The upper half of the state vector contains deviations from the nominal trajectory, rather than the full states, and the lower half of the state vector contains the estimation error. Equation (6) provides us with an equilibrium point of zero throughout the system's operation, and it also offers an expression of the closed-loop system dynamics which lends itself to Lyapunov stability analysis.

B. Stability Theorems

Before we assess the stability of (6), two stability theorems are laid out. These theorems are defined for use with discrete systems by Willems [15].

Theorem 1. The null solution of (6) is stable in the sense of Lyapunov if and only if there exists a bound M , for any k_0 , such that the following inequality holds for all $k \geq k_0$:

$$\|\Phi(k, k_0)\| \leq M$$

$\Phi(k, k_0)$ is the transition matrix which propagates the system to time k from time k_0 . If M can be taken independently of k_0 , then the solution is uniformly stable in the sense of Lyapunov. In other words, for any region R in which we wish the system to stay, we can identify a region r in which the system must start, independent of initial time k_0 .

Theorem 2. The null solution of (6) is asymptotically stable if and only if the conditions of Theorem 1 for stability in the sense of Lyapunov are satisfied, and:

$$\lim_{k \rightarrow \infty} \|\Phi(k, k_0)\| = 0$$

The solution is uniformly asymptotically stable if the above is satisfied and the bound M of Theorem 1 can be taken independently of k_0 . If perturbed, a uniformly asymptotically stable system will return to a state of equilibrium and will do so independently of initial time k_0 .

The norm used in Theorems 1 and 2 is the Euclidean or spectral norm, equivalent to the largest singular value of the transition matrix. The state transition matrix of (6) is equivalent to $\Phi(k+1, k)$, the transition matrix of Theorems 1 and 2 when a transition from k to $k+1$ is made. By multiplying the successive state transition matrices of (6), the transition matrix from any k_0 to any time k along the nominal trajectory can be computed.

Although the control and estimation framework described above is designed for a nominal trajectory that connects two waypoints, this scenario can be augmented for analysis at infinite time. We will hereafter assume that our ship hull survey vehicle travels between a series of waypoints periodically and for infinite time, allowing assessment of stability using Theorems 1 and 2. Our previous findings have demonstrated that a vehicle which can adjust its estimates of both the vehicle and feature states can satisfy Theorem 1 but not 2, and a vehicle which is allowed only to adjust its vehicle state estimates (i.e., localization only, rather than full SLAM), can satisfy both theorems [16].

C. Perturbed Stability Analysis

In addition to guarantees of stability, we also seek a better understanding of how the closed-loop hull inspection vehicle will reject ocean disturbances. Although stability margins are difficult to quantify for a linear time-varying system, analysis of the transition matrix of the perturbed system will provide some insight.

1) Problem Formulation

To understand the system's behavior in the presence of perturbations we must consider how the governing equations change when the vehicle is displaced from the nominal trajectory. Because a displacement from the nominal trajectory renders the linearization of the nominal state transition Jacobian F and measurement Jacobian H incorrect, correction terms ΔF and ΔH are needed to express the true location of the vehicle. Despite this, the need for correction is unknown to the estimator. Using these correction terms the propagation of the state and the estimate appears as follows:

$$\begin{aligned} \delta \underline{x}_{k+1} &= (F_k + \Delta F_k) \delta \underline{x}_k - BG_k N \delta \hat{\underline{x}}_k + \Gamma \underline{w} \\ \delta \hat{\underline{x}}_{k+1} &= \delta \hat{\underline{x}}_{k+1|k} + K_{k+1} [\delta \underline{z}_{k+1} - H_{k+1} \delta \hat{\underline{x}}_{k+1|k}] \end{aligned}$$

$$\delta \hat{\underline{x}}_{k+1|k} = (F_k - BG_k N) \delta \hat{\underline{x}}_k \quad (7)$$

$$\delta \underline{z}_{k+1} = (H_{k+1} + \Delta H_{k+1}) \delta \underline{x}_{k+1} + \underline{v}_{k+1}$$

These equations can be manipulated in a manner similar to equation (6), and all terms containing the correction matrices ΔF and ΔH can be collected in an aggregate perturbation matrix ΔA , where ΔA is defined as follows:

$$\begin{bmatrix} \Delta F_k & 0 \\ [\Delta F_k - K_{k+1}(\Delta H_{k+1}(F_k + \Delta F_k + BG_k N) + H_{k+1}\Delta F_k)] & -K_{k+1}\Delta H_{k+1}BG_k N \end{bmatrix} \quad (8)$$

The aggregate perturbation matrix is combined with the state transition matrix of equation (6) as well as the process and measurement noise to express the system equations as follows:

$$\begin{bmatrix} \delta \underline{x}_{k+1} \\ \delta \tilde{\underline{x}}_{k+1} \end{bmatrix} = \begin{bmatrix} A_k \\ \Gamma \end{bmatrix} \begin{bmatrix} \delta \underline{x}_k \\ \delta \tilde{\underline{x}}_k \end{bmatrix} + \begin{bmatrix} \Delta A_k \\ 0 \end{bmatrix} \begin{bmatrix} \delta \underline{x}_k \\ \delta \tilde{\underline{x}}_k \end{bmatrix} + \begin{bmatrix} \Gamma \\ K_{k+1} \end{bmatrix} \begin{bmatrix} \underline{w}_k \\ \underline{v}_k \end{bmatrix} \quad (9)$$

The state transition matrix of equation (5) is represented here by A_k . By taking the Euclidean norm of both sides of this equation and applying the Bellman-Gronwall lemma, Chen and Dong [17] first described a sufficiency condition for asymptotic stability of a linear system subjected to a perturbation ΔA , which is now a well-documented result [18], [19].

Theorem 3. The null solution of equation (9) is uniformly asymptotically stable if two conditions are satisfied. First, the system must be uniformly asymptotically stable in the absence of perturbations, indicated by the following:

$$\|\Phi(k, k_0)\| \leq mr^k$$

This means that the Euclidean norm of the state transition matrix must be bounded by a discrete exponential with parameters m and r , and that this bound holds independently of k_0 . Second, for a series of perturbation matrices ΔA_k , the following must also hold for all k :

$$\|\Delta A_k\| \leq \frac{1-r}{m}$$

If both conditions are satisfied, then the system will remain uniformly asymptotically stable in the presence of plant perturbations ΔA_k .

2) Evaluating System Robustness

The conservative nature of Theorem 3 requires that the unperturbed system under consideration must be uniformly asymptotically stable, which means any system that fails to satisfy Theorem 2 cannot satisfy Theorem 3. Although this means we can only guarantee perturbed system stability for closed-loop vehicles using features for

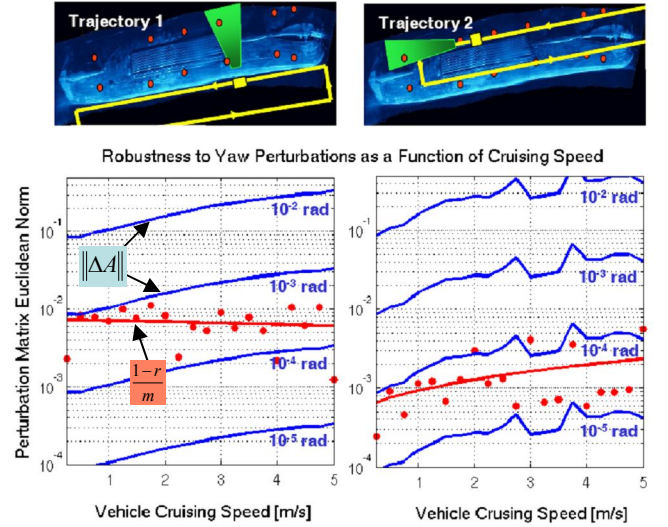


Figure 5: The relative robustness of two HAUV candidate trajectories is evaluated. The perturbation matrix norm $\|\Delta A\|$ is plotted in the angular degree of freedom as a function of vehicle cruising speed for angular disturbances of incrementally varying magnitude (the solid blue lines). Alongside each set of perturbation matrix curves is plotted the exponential bounding curve performance metric, indicating the maximum perturbation size for which each feature layout is asymptotically stable (expressed as a series of red points with a best-fit line).

localization rather than full SLAM, it is conjectured that aspects of filter conditioning (such as the choice of nominal trajectory) which are best-suited for robustness in cases of localization will also be best-suited for robustness in cases of localization and mapping.

One challenge in implementing Theorem 3 is choosing an appropriate perturbation matrix ΔA . For this reason, we evaluate the theorem at perturbations of incrementally varying magnitude, and separately in each degree of freedom. In simulation, the vehicle is displaced from its nominal trajectory by a given distance at each point along the trajectory (in a single degree of freedom), and the worst-case perturbation norm over the entire trajectory is chosen as $\|\Delta A\|$ to represent the magnitude of the displacement. The effect of combining perturbations in multiple degrees of freedom is not considered.

Next we obtain the parameters r and m , which, according to Theorem 3, must comprise an exponential function that bounds the transition matrix norm for any choice of k and k_0 along the nominal trajectory. To achieve this, we find the k_0 which yields the worst-behaved plot of $\|\Phi(k, k_0)\|$ and produce an exponential function which bounds all values on this curve efficiently. By comparing the parameters of the exponential bounding curve with the norm of ΔA , it is apparent when the system is guaranteed asymptotically stable. In addition, by comparing system conditions which achieve a stability guarantee for differing perturbation magnitudes, the relative robustness of two trajectories, maps, or otherwise can be compared. In this manner the second inequality of

Theorem 3 is used as a performance metric.

D. Robustness of HAUV Survey Trajectories

As an illustrative example of how the above stability theorems may be used to analyze a closed-loop vehicle survey, we consider the problem of choosing a robust trajectory for a HAUV hull inspection. The two survey trajectories depicted in Figure 5 were selected heuristically based on simplicity and apparent hull coverage, and both have been used frequently in vehicle experiments performed on the hull that is pictured (the same hull depicted in Figure 4). The simulated survey will use DVL odometry and range-and-bearing measurement of the hull's zinc anodes (marked in red in Figure 5) for navigation. The robustness of these trajectories is challenging to predict since the viewing window of the DIDSON imaging sonar, pictured in Figure 5, is limited in size, and only a subset of the zincs will be observed at each time step. If the cruising speed of the vehicle is varied, the frequency with which the features will be observed will vary in turn.

For a variety of cruising speeds, the maximum perturbation magnitude for which stability is guaranteed is computed for each of the two candidate trajectories. At the bottom of Figure 5, $\|\Delta A\|$ is plotted for varying perturbation magnitudes in the angular degree of freedom. This perturbation norm is compared with the r and m parameters of the exponential fit to $\Phi(k, k_0)$ for each of the cruising speeds examined. To offer an example of what we may conclude from these plots, the bottom right plot of Figure 5 shows that for Trajectory 2, the vehicle is guaranteed stable against angular perturbations of order 10^{-4} radians for a cruising speed of 3 m/s, but this guarantee cannot be made for a cruising speed of 4 m/s. Hence, a vehicle executing trajectory 2 is more robust traveling at 3 m/s than at 4 m/s, since it is guaranteed stable against larger-sized perturbations. In general, Figure 5 permits the conclusion that a vehicle executing Trajectory 1 is guaranteed stable against larger-sized perturbations than a vehicle executing Trajectory 2. Hence, our performance metric would indicate that Trajectory 1 is the robustness-optimal choice from among the two candidates.

This method for comparing the relative robustness of two nominal trajectories is not limited to this particular choice of dynamic model or suite of sensors, but can be applied to any closed-loop system for which asymptotic stability is attainable and whose dynamics may be expressed using a linear time-varying state space framework.

V. ROBUSTNESS-OPTIMAL MOTION PLANNING

The process of finding the most robust closed-loop

$u = \begin{bmatrix} -f_u \\ f_v \\ 0 \end{bmatrix}$ cost = $\sqrt{2} C_0$	$u = \begin{bmatrix} 0 \\ f_v \\ 0 \end{bmatrix}$ cost = C_0	$u = \begin{bmatrix} f_u \\ f_v \\ 0 \end{bmatrix}$ cost = $\sqrt{2} C_0$
$u = \begin{bmatrix} -f_u \\ 0 \\ 0 \end{bmatrix}$ cost = C_0	$u = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$ cost = $\frac{1}{2} C_0$	$u = \begin{bmatrix} f_u \\ 0 \\ 0 \end{bmatrix}$ cost = C_0
$u = \begin{bmatrix} -f_u \\ -f_v \\ 0 \end{bmatrix}$ cost = $\sqrt{2} C_0$	$u = \begin{bmatrix} 0 \\ -f_v \\ 0 \end{bmatrix}$ cost = C_0	$u = \begin{bmatrix} f_u \\ -f_v \\ 0 \end{bmatrix}$ cost = $\sqrt{2} C_0$

Figure 6: The discrete input choices used for implementation with A* are displayed in a spatial grid. The distance of each choice from the center corresponds to the relative energy required to enact its input thrust command. Input choices with equivalent cost are assigned the same color.

system configuration for a holonomic marine vehicle is next combined with a globally optimal kinodynamic motion planning algorithm. As introduced above, a system configuration is designated the most robust configuration if it achieves a guarantee of asymptotic stability against a larger sized-perturbation than other competing system configurations. Rather than heuristically choosing a handful of system configurations and evaluating their relative closed-loop robustness, a modified A* algorithm is used to select the most robust trajectory from among all trajectories in the vehicle state space that connect a given start waypoint and a given goal waypoint. The robustness-optimal trajectories selected by the kinodynamic planning algorithm are compared with A* trajectories that do not consider the robustness performance metric.

A. Kinodynamic Planning in Control Input Space

Figure 6 displays the nine discrete control input choices used by the planning algorithm and their associated costs tabulated in a spatial grid. In this grid, a positive surge thrust is located one unit to the right from the center, and a negative surge thrust is located one unit to the left from the center. A positive sway thrust is located one unit up from center, and a negative sway thrust is one unit down. Applying positive thrust in surge and sway corresponds to the grid's upper right corner, and so on. The location of these input choices on the nine-by-nine spatial grid corresponds approximately to the energy required to implement them. Costs are assigned to the eight outermost input choices based on the Euclidean distance of each grid entry from the center. The center square is assigned a non-zero cost so the vehicle does not find sitting idle to be the cheapest action. In addition, this cost is less than the cost of single-actuator thrust so the vehicle

will opt to coast along without using additional thrust once it is close to the goal, instead of hitting the goal at top speed.

B. Cost-to-Go Heuristic

To ensure that the A* algorithm converges quickly to the goal, an admissible heuristic must be employed which accurately approximates the cost required to reach the goal without over-estimating the true minimum cost of travel to the goal node from the node under consideration. To achieve this, three simplifying assumptions are made. First, all hydrodynamic drag forces are neglected. It is assumed that the marine vehicle behaves as a double integrator in each degree of freedom. Less input thrust is required to move a double integrator from one location to another than to move an equivalent mass with any amount of hydrodynamic drag, and so this assumption will aid in producing an underestimate. The second simplifying assumption is that regardless of the vehicle pose at the node being evaluated, it is assumed that a thruster is pointing directly at the goal so that the double-thruster cost will never be applied in generating the cost-to-go. The only input cost that will be used is the cost associated with running a single thruster. The third assumption projects the vehicle velocity onto a vector connecting the vehicle's current position and the goal, neglecting the velocity component normal to this vector. The only thrust required by the heuristic will be thrust directed at the goal. The resulting procedure for computing the cost-to-go is given by (10):

$$t_{goal} = \frac{-V_0 + \sqrt{V_0^2 + 2 \frac{F_0}{m} d}}{\frac{F_0}{m}} \quad (10)$$

$$C_{go} = \text{floor}(t_{goal}/T_{plan}) * C_0$$

This procedure directly computes the time required to reach the goal from the current vehicle position, t_{goal} . This closed-form solution is possible because of the double integrator assumption. V_0 represents the net velocity in the direction of the goal, which is positive when the vehicle is moving toward the goal and negative when the vehicle is moving away from the goal. The distance between the current position and the goal is d , and F_0 represents the fixed amount of thrust that may be applied by a single thruster. The mass of the vehicle is m , T_{plan} is the fixed time interval over which each input command is applied, the "floor" function rounds its argument down to the nearest integer, and C_0 is the baseline unit of cost introduced in Figure 6. It has not been proven that this procedure yields a cost-to-go heuristic that is always admissible, but it has thus far yielded satisfactory planning results and fast algorithm convergence. The only possibility left to chance by this heuristic is that

thrusting, followed by coasting, may sometimes be the most cost-efficient behavior. Since computing the exact time to turn off the thrusters and begin coasting in the computation of C_{go} is yet another optimization problem, the double integrator and velocity vector assumptions are used instead to offset any possibility of getting to the goal at a lower cost by coasting.

C. Robustness-Augmented Cost Function

The final and most unique feature of this kinodynamic A* algorithm is the consideration of robustness against perturbations to the marine vehicle. In addition to the cost-to-come and the cost-to-go, the cost function used in this analysis includes a term which evaluates the robustness of the path from the start node to the current node by applying the performance metric introduced in Section IV. This new cost is concerned only with the specific path traveled so far, and so it is an additional cost-to-come. The complete cost function is given by (11):

$$C_{total} = C_{come} + C_{go} + C_{robust, come} \quad (11)$$

$$C_{robust, come} = W * \log_{10} \left(\|\Delta A\|_{\max} - \frac{1-r}{m} \right)$$

The robustness component of the cost function consists of a tunable weight W multiplying the base ten logarithm of the difference between a perturbation matrix norm and an exponential performance metric, the two components of the inequality of Theorem 3. The robustness parameters are obtained by designating the trajectory under consideration between the start node and the current node to be the vehicle's nominal trajectory. A series of LKF and feedback control gains are computed for this nominal trajectory so that a closed loop transition matrix may be formulated for the system. The perturbation matrix is computed according to (8) for a benchmark perturbation size selected for the entire duration of the algorithm. The perturbation is applied to the vehicle in positive and negative x and y at each sampling instant of the trajectory, and the ΔA matrix with the largest Euclidean norm over the entire nominal trajectory, over all perturbations considered, is selected for the robustness cost function. A large perturbation size is selected so that ΔA is always greater than the robustness performance metric and the argument of the logarithm in (11) is always positive. The exponential parameters r and m , which were first introduced in Theorem 3, are obtained by bounding the marine vehicle's transition matrix norm along the A*-formulated nominal trajectory with an exponential function. In this procedure, the norm of $\Phi(k, k_0)$ from every instant k_0 to every instant k along the entire nominal trajectory is no longer considered. Instead r and m are obtained by fitting an exponential function to

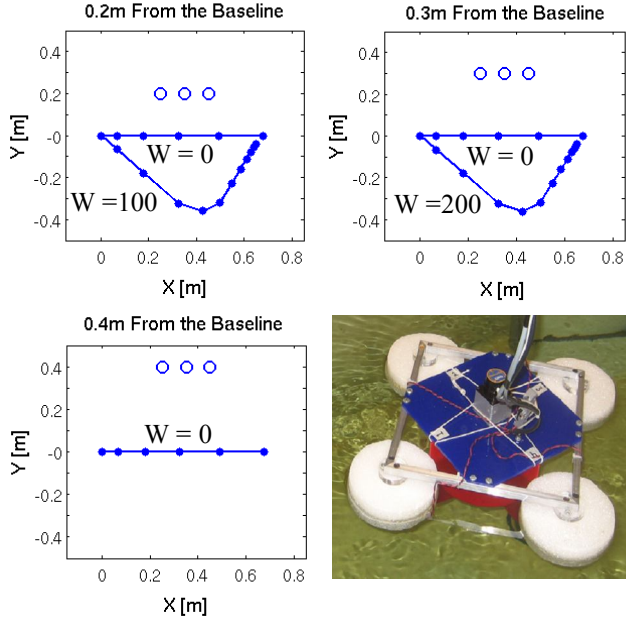


Figure 7: Results from regular and robustness-augmented A* when the start and goal waypoints are separated by a horizontal distance of 0.7m. Three cases with different vertical distances from the three-feature baseline are considered. A robustness weight of 100 is required to produce the curved trajectory for the 0.2m case, a weight of 200 is required to produce the curved trajectory for the 0.3m case, and no weight is capable of producing a curved trajectory in the 0.4m case. The points plotted represent the location of the vehicle at each sampling instant along the A*-planned path. A photo of the laboratory test platform is also provided.

$|\Phi(k, k_0)|$ for the single k_0 which corresponds to the start node. This simplification is made to improve computational efficiency of the algorithm, since many hundreds of nominal trajectories might be evaluated over the course of the A* algorithm, and computing the norm of $\Phi(k, k_0)$ for every possible k and k_0 is an expensive task. The only consequence is that the stability guarantees of Theorems 2 and 3 now pertain to asymptotic rather than uniform asymptotic stability.

D. Performance of Algorithm

The motion planning results to follow discuss trajectories which have been planned for a scaled-down holonomic laboratory platform (pictured in Figure 7) rather than the HAUV itself. Testing the algorithm on this platform will allow simple plans using small numbers of features to be executed before scaling up to the HAUV. To evaluate the A* algorithm's performance, a set of three collinear features is selected and (0,0) and (0,0.7) are chosen as start and goal waypoints in the x - y plane. Figure 7 displays the three-feature map and the trajectories produced by A* with and without a robustness cost for three different separation distances between the features and the waypoints.

In the absence of a robustness cost (achieved by setting the weight W from (11) to zero), A* yielded a horizontal

straight-line trajectory between the two waypoints. This makes intuitive sense as it achieves the shortest spatial distance between the two points and it uses the cheapest and most basic application of thrust. In each case where W was set to zero the algorithm used two consecutive applications of the surge thruster to reach the goal. By setting the robustness weight to larger values, a curved trajectory eventually emerged which departed from the energy-efficient path to stay in more robust locations relative to the features. In all locations where the curved path emerged, it consisted of a simultaneous application of thrust in positive surge and negative sway, followed by an application of positive thrust in sway, after which the vehicle refrains from using thrusters and coasts to the goal for the remainder of the trajectory. Relative to the 0.2m vertical separation between the waypoints and the three-feature baseline, the 0.3m separation required twice the weighting factor to force the vehicle away from the regular A* path, and no weight was found which could force the trajectory with 0.4m separation to curve its path. These results suggest that as the three-feature baseline is moved closer to the designated waypoints, the regular A* trajectory becomes less and less robust.

E. Jacobian Sensitivities

As numerical computation revealed, the measurement Jacobian is significantly more sensitive to perturbations along the horizontal straight-line trajectories than to perturbations along the robustness-optimal curved trajectories for the 0.2m and 0.3m separation distances. The part of the trajectory with the greatest sensitivity was the portion near the center of the three-feature baseline. An examination of the perturbation norm $\|\Delta A\|_{max}$ and the ratio of exponential parameters $(1-r)$ to m revealed that the greatest sensitivity is present in the perturbation norm, and it is contributed largely by the error ΔH in the nominal measurement Jacobian when the vehicle is perturbed. ΔH represents, as it does in (7) and (8), the difference between the nominal measurement Jacobian used by the LKF for localization and a Jacobian linearized about the vehicle's perturbed location. This error in the measurement Jacobian is displayed graphically in Figure 8, which illustrates the size of ΔH for perturbations of 0.1m in x and in y , respectively, for each of the A*-computed nominal trajectories that is plotted in Figure 7. The size of the ΔH matrix is represented by computing the Euclidean norm. The location of each bar in Figure 8 corresponds approximately to the location of the respective point in x along the A*-computed nominal trajectory.

Figure 8 demonstrates that the measurement Jacobian is especially sensitive to perturbations from the 0.2m horizontal straight-line trajectory, which was computed

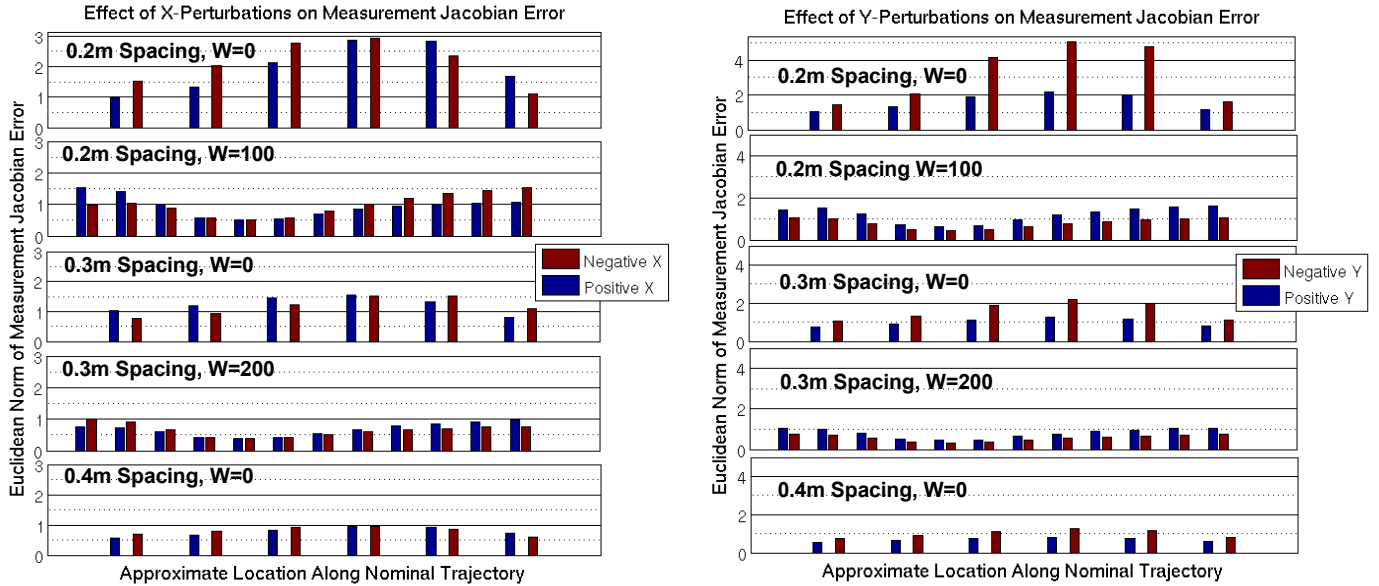


Figure 8: The Euclidean norm of the Jacobian error ΔH induced by perturbations of positive and negative 0.1m is plotted separately in x and in y for each of the five A*-computed trajectories displayed in Figure 7. The worst Jacobian errors occur along the horizontal straight-line trajectories that are close to the three-feature baseline.

with a robustness weight W of zero. It was second-most sensitive to perturbations from the horizontal trajectory spaced 0.3m from the baseline. Both robustness-augmented curved trajectories yield significant improvements in the sensitivity of the measurement Jacobian, and also in the overall robustness cost assigned to the trajectory. In particular, for the 0.3m spacing case the robustness-augmented trajectory accumulates a robustness cost-to-come of 1, while the regular A* trajectory accumulates a robustness cost of 47. Changes in nominal trajectory had almost no effect on the ratio produced using exponential parameters m and r , which are nearly identical in value for all trajectories. All of the sensitivity to perturbations was exhibited in the perturbation norm $\|\Delta A\|_{max}$, and in particular this sensitivity can be traced to the measurement Jacobian.

VI. CONCLUSION

We have reported recent progress in the effort to achieve autonomous ship hull inspection in the areas of perception, stability analysis, and motion planning. First, mine-shaped training targets were identified and mapped by the HAUV in real time, and bathymetry sonar is used to construct a detailed 3D representation of a ship's running gear.

In addition, a stability analysis has been applied to a planar model of a ship hull inspection vehicle, and it has been shown that $\Phi(k, k_0)$, the vehicle's linearized time-varying transition matrix, can be computed to evaluate the stability of the unperturbed system. It has also been shown that by considering perturbations to the system, a performance metric can be derived which allows a

comparison of robustness among different system configurations. This performance metric allows us to discern the impact of subtle aspects of filter conditioning, such as the sequence in which features are observed, on the robustness of the closed-loop vehicle. The use of this technique yields a procedure for evaluating which of several candidate trajectories is best-equipped to tolerate perturbations.

Finally, a kinodynamic planning algorithm was introduced which deviates from the shortest path between waypoints to achieve a stronger guarantee of asymptotic stability in the presence of perturbations. We hope that continued consideration of the dynamic stability of an integrated localization, mapping, and control process will allow for the implementation of more robust and aggressive motion planning and control algorithms on the HAUV as we enhance its reliance on feature-based navigation.

VII. ACKNOWLEDGEMENTS

Thanks to Jerome Vaganay of Bluefin Robotics for providing HAUV field testing support. This work is supported by the Office of Naval Research under Grant N00014-06-10043, monitored by Dr. T.F. Swain.

REFERENCES

- [1] R.C. Smith and P. Cheeseman, "On the representation and estimation of spatial uncertainty," *Int. J. Robotics Research*, vol. 5(4), 1986, pp. 56-68.
- [2] J.J. Leonard and H.F. Durrant-Whyte, "Mobile Robot Localization by Tracking Geometric Beacons," *IEEE Trans. on Robotics and Automation*, vol. 7(3), 1991, pp. 376-382.
- [3] G. Dissanayake, P. Newman, S. Clark, H.F. Durrant-Whyte, and M.Csorba, "A Solution to the Simultaneous Localization and Map

- Building (SLAM) Problem,” *IEEE Trans. on Robotics and Automation*, vol. 17(3), 2001, pp. 229-241
- [4] F. Hover, et al., “A Vehicle System for Autonomous Relative Survey of In-Water Ships,” *Marine Technology Society Journal*, vol. 41(2), 2007, pp. 44-55.
- [5] E. Belcher, H. Dinh, D. Lynn, and T. Laughlin, “Beamforming and Imaging with Acoustic Lenses in Small, High Frequency Sonars,” *Proc. IEEE OCEANS Conf.*, Seattle, 1999, pp. 1495-1499.
- [6] T. Vidal-Calleja, J. Andrade-Cetto, and A. Sanfeliu, “Estimator Stability Analysis in SLAM,” *Proc. 5th IFAC/EURON Symp. on Intelligent Autonomous Vehicles*, Lisbon, July 2004.
- [7] J. Andrade-Cetto and A. Sanfeliu, “The Effects of Partial Observability in SLAM,” *Proc. IEEE Int. Conf. on Robotics and Automation*, New Orleans, 2004, pp. 397-402.
- [8] F. Hover, “Stability of Double-Integrator Plants Controlled using Real-Time SLAM Maps,” *Proc. IEEE Int. Conf. on Robotics and Automation.*, Pasadena, 2008, pp. 637-642.
- [9] P. Hart, N. Nilsson, and B. Raphael, “A Formal Basis for the Heuristic Determination of Minimum Cost Paths,” *IEEE Trans. On Systems Science and Cybernetics*, vol. 4(2), 1968, pp. 100-107.
- [10] C. Roman and H. Singh, “Consistency-based Error Evaluation for Deep Sea Bathymetric Mapping with Robotic Vehicles,” *Proc. IEEE Int. Conf. on Robotics and Automation*, Orlando, 2006, pp. 3568-3574.
- [11] S. Reed, Y. Petillot, and J. Bell, “An Automatic Approach to the Detection and Extraction of Mine Features in Sidescan Sonar,” *IEEE Journal of Oceanographic Engineering*, vol. 28(1), 2003, pp. 90-105.
- [12] P.J. Besl and N.D. McKay, “A method for registration of 3-d shapes,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 14(2), 1992, pp. 239-256.
- [13] S. Rusinkiewicz and M. Levoy, “Efficient Variants of the ICP Algorithm,” *Proc. Third Int. Conf. on 3-D Imaging and Modeling*, Quebec, 2001, pp. 145-152.
- [14] A. Gelb, ed., *Applied Optimal Estimation*, Cambridge, MA: The MIT Press, 1984.
- [15] J.C. Willems, *Stability Theory of Dynamical Systems*, London: Thomas Nelson and Sons ltd., 1970.
- [16] B. Englot, *Stability and Robustness Analysis Tools for Marine Robot Localization and Mapping Applications*, Master’s Thesis, Massachusetts Institute of Technology, 2009.
- [17] B. Chen and T. Dong, “Robust stability analysis of Kalman-Bucy Filter under parametric and noise uncertainties”, *Int. J. Control*, 48, 1988, pp. 2189-2199.
- [18] A. Weinmann, *Uncertain Models and Robust Control*, New York: Springer-Verlag/Wien, 1991.
- [19] G. Chen, ed., *Approximate Kalman Filtering*, Singapore: World Scientific, 1993.
-