

MIT Open Access Articles

*Stability and robustness analysis tools for
marine robot localization and SLAM applications*

The MIT Faculty has made this article openly available. **Please share**
how this access benefits you. Your story matters.

Citation: Englot, B., and F. Hover. "Stability and robustness analysis tools for marine robot localization and SLAM applications." Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on. 2009. 4426-4432. © Copyright 2009 IEEE

As Published: <http://dx.doi.org/10.1109/IROS.2009.5354171>

Publisher: Institute of Electrical and Electronics Engineers

Persistent URL: <http://hdl.handle.net/1721.1/59291>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Stability and Robustness Analysis Tools for Marine Robot Localization and SLAM Applications

Brendan Englot and Franz Hover, *Member, IEEE*

Abstract—Our aim is to explore the fundamental stability issues of a robotic vehicle carrying out localization, mapping, and feedback control in a perturbation-filled environment. Motivated by the application of an ocean vehicle performing an autonomous ship hull inspection, our planar vehicle model performs localization using point features from a given map. Cases in which the agent must update the map are also considered. 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. These theorems yield a stability test for the feedback controller. Perturbations are then considered using a theorem on the convergence on the perturbed system transition matrix, yielding a robustness test for the estimator. Together, these tests form a set of tools which can be used in planning and evaluating the robustness of marine vehicle survey trajectories, which is demonstrated through experiment.

I. INTRODUCTION

Simultaneous localization and mapping (SLAM) algorithms permit an agent to build and refine a map of its environment while at the same time using the map to estimate its position and orientation within the environment [1]. In many applications, an estimation-theoretic approach is used and the extended Kalman filter can be employed to carry out the algorithm [2], [3]. Despite the widespread use of SLAM, a limited body of work exists on the stability of the integrated localization, mapping, and dynamic control process. Stability of the linear Kalman filter in the specialized case of a one degree-of-freedom monobot was assessed by Vidal-Calleja, Andrade-Cetto, and Sanfeliu [4], and its observability by Andrade-Cetto and Sanfeliu [5], leading to the conclusion that the partial observability of the filter yields marginally stable estimation error dynamics. Hover [6] analyzed the stability margins of a localization estimator with closed-loop control for 1-DOF and planar vehicles with double integrator plants. This analysis considered the regulation problem, using a constant-gain controller and estimator only. In the present work, we will extend the focus to a planar vehicle with time-varying controller and estimator gains, allowing travel anywhere in the 2-D plane.

Integrated localization, map-building and control, and the related problem of localization and control based on a given map, are 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 a prior map of the hull is often, but not always, available, and underwater operation near a large steel structure prevents use of compass, GPS, or long baseline acoustic tracking as consistent sensors. Random disturbances in the ocean environment necessitate the use of high-fidelity feedback control, and hence a better understanding of integrated localization, mapping, and dynamic control is desirable. Figure 1 illustrates a recent application of EKF-SLAM to the task of ship hull inspection, in which the MIT-Bluefin Hovering Autonomous Underwater Vehicle (HAUV) [7] constructed a map in real-time of six mine-shaped features planted on the hull of the *USS Saratoga* in May 2008. The use of imaging sonar allows the vehicle-relative range and bearing of features to be detected [8]. Completion of a feature-based estimation task in real-time has led to the goal of implementing integrated mapping and control with the HAUV.

In our approach to integrated mapping and control, it is important to note that mapping is considered in the context of use and *refinement* of a map that is given *a priori*, and not the building of an entirely new map. Refinement of the map entails updating or correcting the features already present on the map, but does not include the addition of new features. We also consider the case of *map exploitation*, in which the *a priori* map of known features is used for the sole purpose of localization, and is not updated or corrected. Although localization by exploiting an *a priori* map can also be achieved using iterative closest point (ICP) algorithms [9] and a variety of other iterative algorithms prevalent in visual servoing research [10], our discussion of localization will be limited to an estimation-theoretic SLAM framework, as this framework enables a series of stability guarantees relevant to the problem of autonomous ship hull inspection.

Motivated by the example of ship hull inspection, this analysis considers a planar 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. Section II introduces the vehicle model and the control and estimation strategy. In Section III it will be shown that integrated map refinement and control can achieve uniform stability in the sense of Lyapunov using a theorem on the transition matrix Euclidean norm. A complementary theorem will be used to show that integrated map exploitation and control can achieve uniform asymptotic stability. Section IV considers the effect of perturbations on stability and introduces a robustness performance metric. This metric can be used to evaluate the conditioning of the estimator, and in particular to evaluate the variation in robustness that results

This work was supported by the Office of Naval Research under Grant N00014-06-10043, monitored by Dr. T.F. Swain

B. Englot and F. Hover are with the Department of Mechanical Engineering, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge MA 02139 USA benglot@mit.edu hover@mit.edu




Fig. 1. 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 photo of the HAUV and a sonar mosaic image of the targets placed on the ship hull (mosaic image provided courtesy of AcousticView www.acousticview.com).

from a variation in the geometric pattern of features on the map. Section V introduces experimental results which validate these predictions of geometry-dependent robustness. We then discuss how this metric can be used in guiding the choice of vehicle survey trajectories, using the problem of selecting an HAUV ship hull inspection trajectory as an example.

II. A SIMPLE MARINE VEHICLE MODEL

The dynamics of a holonomic marine vehicle operating in a 2-D plane may be described in discrete time as follows:

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

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

The body-referenced forward velocity, sway velocity, and yaw rate are described by u , v , and $\dot{\phi}$, 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 , 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. Throughout the analysis and experiments to follow, this model specifically describes the holonomic platform pictured in Figure 2, whose time constant of linear motion is approximately two seconds, and whose time constant of angular motion is approximately one second.

A convenient way to permit KF-based SLAM to reconstruct the vehicle state using the measurement of point features is the addition of quasi-states representing the horizontal and vertical positions of the features in the plane. We will assume that the quasi-states are permanently fixed in space and have no dynamics. The aggregate state vector appears as follows:

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

The vehicle states at time k are contained within $\underline{x}_{v|k}$. The contents of this state vector permit the vehicle-relative range and bearing measurements of each feature to be assembled.

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:

$$\begin{aligned} \underline{x}_{k+1} &= f(\underline{x}_k) + B(\underline{u}_{OLk} - G_k N \delta \hat{\underline{x}}_k) + \Gamma \underline{w}_k \\ \underline{z}_{k+1} &= h(\underline{x}_{k+1}) + \underline{v}_{k+1} \end{aligned} \quad (3)$$

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 (1) 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 $2n$ quasi-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 map refinement 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 :

$$\begin{aligned} \delta \hat{\underline{x}}_{k+1} &= \delta \hat{\underline{x}}_{k+1|k} + K_{k+1} [\delta \underline{z}_{k+1} - H(\bar{x}_{k+1}) \delta \hat{\underline{x}}_{k+1|k}] \\ \text{where } \delta \hat{\underline{x}}_{k+1|k} &= F(\bar{x}_k) \delta \hat{\underline{x}}_k - B G_k N \delta \hat{\underline{x}}_k \end{aligned} \quad (4)$$




Fig. 2. The model used in this stability analysis describes a holonomic floating platform with three degrees of freedom, a scaled-down mockup of a ship hull inspection vehicle. The platform is equipped with a Hokuyo-URG laser and four bilge pumps mounted on foam pontoons which act as water jet thrusters.

The nonlinear state transition function $f(x)$ and the measurement function $h(x)$ are now replaced by $F(\bar{x}_k)$ and $H(\bar{x}_k)$, the corresponding Jacobians, which are linearized about the nominal trajectory at each time step. The term $\delta \tilde{z}_k$ 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 [11]). 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 the 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 (3) and (4) yields the following compact formulation:

$$\begin{bmatrix} \delta \tilde{x}_{k+1} \\ \delta \tilde{z}_{k+1} \end{bmatrix} = \begin{bmatrix} F(\bar{x}_k) - BG_k N & BG_k N \\ 0 & F(\bar{x}_k) - E_k \end{bmatrix} \begin{bmatrix} \delta \tilde{x}_k \\ \delta \tilde{z}_k \end{bmatrix} \quad (5)$$

where $E_k = K_{k+1} H(\bar{x}_{k+1}) F(\bar{x}_k)$

$$\delta \tilde{z}_k = \delta \tilde{x}_k - \delta \hat{x}_k$$

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 (5) 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.

III. UNPERTURBED STABILITY ANALYSIS

A. Stability Theorems

Before we assess the stability of (5), two stability theorems are laid out. These theorems are defined for use

with discrete systems by Willems [12].

Theorem 1. The null solution of (5) 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_1 in which we wish the system to stay, we can identify a region R_2 in which the system must start, independent of initial time k_0 .

Theorem 2. The null solution of (5) 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 (5) 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 (5), the transition matrix from any k_0 to any time k along the nominal trajectory can be computed.

Although the framework described in Section II is designed for a nominal trajectory that connects two waypoints, this scenario can be augmented for analysis at infinite time. In the analysis to follow a periodic path between the waypoints depicted in Figure 3 will be executed, in simulation and in experiment, by the holonomic platform displayed in Figure 2. One period consists of a forward trip along the path, followed by a trip in reverse back to the origin. It is also assumed that the vehicle uses the same three features throughout.

B. Stability Results for Map Refinement

First we consider the case in which the vehicle uses the features for localization and can simultaneously adjust its estimates of the feature locations. The first plot of Figure 4 displays the values of the transition matrix norm for the $\Phi(k, k_0)$ that transit from every initial time k_0 to every final time k along two cycles of the nominal trajectory. The convergence of the norm to a periodic surface independent of k_0 permits the conclusion that the transition matrix norm can be assigned a bound M which it will not surpass, and that this bound can be taken independent of k_0 . Hence, Theorem 1 is satisfied and we may conclude the nominal trajectory investigated is uniformly stable in the sense of Lyapunov.




Fig. 3. The periodic vehicle trajectory analyzed in this study is depicted, along with the configuration of three features used in the results to follow. The nominal changes in heading commanded across this trajectory are indicated by the body-fixed coordinates u and v and inertial angular coordinate ϕ . Parameters that are used in our vehicle simulation and experiment are listed.

This is the strongest stability guarantee we can obtain for a system that includes the feature quasi-states, which allow map refinement to occur as the features are measured. It also confirms the intuition that a system cannot drive both the vehicle and feature state estimates to zero if the map used for pre-computing the LKF filter gains and Jacobians is in need of refinement.

C. Stability Results for Map Exploitation

In search of a stronger stability guarantee we consider map exploitation, in which the *a priori* map is assumed to contain the correct feature locations and only the vehicle states are estimated (i.e. localization rather than full SLAM). The upper right plot of Figure 4 shows the transition matrix norm for the $\Phi(k, k_0)$ that transit from every initial time k_0 to every final time k along two cycles of the nominal trajectory. Unlike before, the norm converges to a value of zero, and it does so irrespective of initial time k_0 . This means that not only may we identify a bound M for the norm according to Theorem 1, but we may also apply Theorem 2 since the norm of $\Phi(k, k_0)$ approaches zero as k approaches infinity, independent of our choice of k_0 . Thus, map exploitation is capable of achieving uniform asymptotic stability, since the vehicle and its state estimates can be driven to zero if the map is accurate.

To demonstrate that computation of $\|\Phi(k, k_0)\|$ can be used as a necessary and sufficient indicator of system stability, the bottom of Figure 4 contains a map exploitation case in which the controller gains used previously were amplified by factors of five and six, respectively. The system remains stable using a factor of five, but a factor of six is sufficient to render the system unstable, as $\|\Phi(k, k_0)\|$ is unbounded and grows rapidly with increasing k . In this manner, we can gauge the upward gain margin of the system by analyzing $\|\Phi(k, k_0)\|$. Although this method allows detection of an ill-conditioned controller, there are aspects of the localization estimator which can pass through unnoticed when ill-conditioned (such as a map whose geometry will not allow successful localization). For this reason, we must also consider the effect of perturbations on the system, which are needed to bring

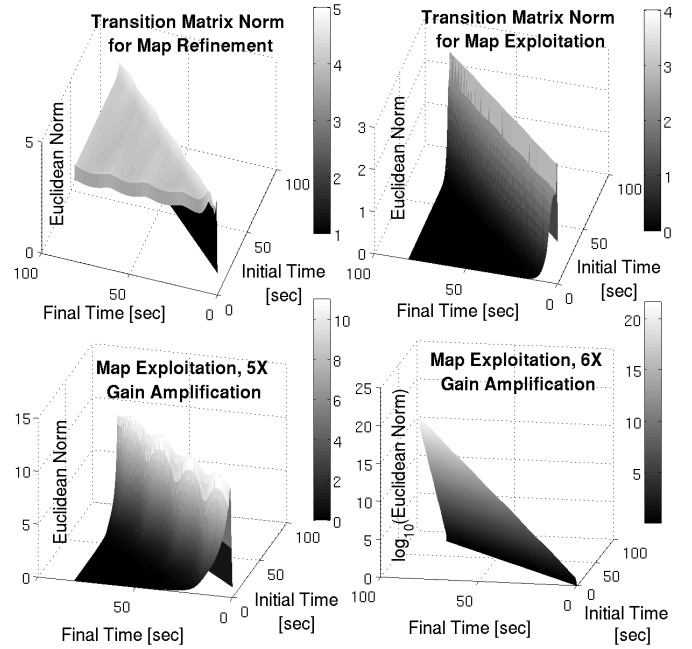


Fig. 4. The transition matrix norm is plotted for the trajectory and map depicted in Figure 3, using the framework presented in Section II. The plots encompass two complete cycles of the nominal trajectory, using the parameters identified in Figure 3.

about the failure of the filter due to certain aspects of filter conditioning.

IV. PERTURBED STABILITY ANALYSIS

A. 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:

$$\begin{aligned}\delta \underline{x}_{k+1} &= (F_k + \Delta F_k) \delta \underline{x}_k - BG_k N \delta \hat{x}_k + \Gamma w \\ \delta \hat{x}_{k+1} &= \delta \hat{x}_{k+1|k} + K_{k+1} [\delta z_{k+1} - H_{k+1} \delta \hat{x}_{k+1|k}] \\ \delta \hat{x}_{k+1|k} &= (F_k - BG_k N) \delta \hat{x}_k \\ \delta z_{k+1} &= (H_{k+1} + \Delta H_{k+1}) \delta x_{k+1} + v_{k+1}\end{aligned}\quad (6)$$

These equations can be manipulated in a manner similar to (5), and all terms containing correction matrices ΔF and ΔH can be collected in an aggregate perturbation matrix ΔA , where ΔA is defined as follows:

$$\begin{aligned}\Delta A_k &= \begin{bmatrix} \Delta F_k & 0 \\ M_k & -K_{k+1} \Delta H_{k+1} B G_k N \end{bmatrix} \\ M_k &= \Delta F_k - K_{k+1} (\Delta H_{k+1} (F_k + \Delta F_k + B G_k N) + H_{k+1} \Delta F_k)\end{aligned}\quad (7)$$

The aggregate perturbation matrix is combined with the state transition matrix of (5) as well as the process and measurement noise to express the system equations as follows:

$$\begin{bmatrix} \delta \underline{x}_{k+1} \\ \delta \underline{\tilde{x}}_{k+1} \end{bmatrix} = [A_k] \begin{bmatrix} \delta \underline{x}_k \\ \delta \underline{\tilde{x}}_k \end{bmatrix} + [\Delta A_k] \begin{bmatrix} \delta \underline{x}_k \\ \delta \underline{\tilde{x}}_k \end{bmatrix} + \begin{bmatrix} \Gamma & 0 \\ [I - K_{k+1}(H_{k+1} + \Delta H_{k+1})]\Gamma & K_{k+1} \end{bmatrix} \begin{bmatrix} \underline{w}_k \\ \underline{v}_k \end{bmatrix} \quad (8)$$

The state transition matrix of (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 [12] first described a sufficiency condition for asymptotic stability of a linear system subjected to a perturbation ΔA , which is now a well-documented result [13], [14].

Theorem 3. The null solution of (8) 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 .

B. Evaluating System Robustness

The conservative nature of Theorem 3 requires that the unperturbed system under consideration must be uniformly asymptotically stable, which excludes cases of map refinement from evaluation since these cases achieve Lyapunov stability at best. Despite this, it is conjectured that aspects of filter conditioning (such as the choice of nominal trajectory) which are best-suited for robustness in cases of map exploitation will also be best-suited for robustness in cases of map refinement.

One challenge in implementing Theorem 3 is choosing an appropriate perturbation matrix ΔA . For this reason, we have evaluated the theorem at perturbations of incrementally varying magnitude in each degree of freedom. In simulation, the vehicle is displaced from the 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 ΔA to represent the magnitude of the displacement. The effect of combining perturbations in multiple degrees of freedom is not considered.

We obtain r and m by bounding the surface plot of $\|\Phi(k, k_0)\|$ with an exponential. Specifically, we find the k_0 which yields the worst-behaved curve and ensure that the exponential 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.

As an illustrative example, the effect of feature spacing on vehicle robustness can be investigated using this performance metric. Starting with three features condensed to a single point and gradually separating them until they achieve the configuration of Figure 3, the maximum perturbation magnitude for which stability is guaranteed can be computed for each feature configuration. This is demonstrated in Figure 5, in which $\|\Delta A_k\|$ is plotted for varying configurations of features, for varying disturbance magnitudes, and separately in each degree of freedom. The perturbation norm is compared with the r and m parameters of the exponential fit to $\|\Phi(k, k_0)\|$ for each configuration. To offer an example of what we may conclude from these plots, the bottom right plot of Figure 5 shows that the vehicle is guaranteed stable against angular perturbations of order 10^{-3} radians for a feature spacing of 0.1 meters, but this guarantee cannot be made for a feature spacing of 0.01 meters. Hence, a vehicle observing features with a spacing of 0.1 meters is more robust than if it were observing features with a spacing of 0.01 meters, since it is guaranteed stable against larger-sized perturbations. In general, Figure 5 permits the conclusion that maps with a wider feature spacing are guaranteed stable against larger-sized perturbations, although the performance metric exhibits asymptotic behavior and there is a point beyond which no significant gains in robustness are achieved by spreading the features further.

Although the perturbations for which stability is guaranteed are often small in size (it is almost certain that perturbations larger than 10^{-3} radians will be encountered during vehicle tests), we aim to use Theorem 3 as a means of comparing system configurations. The system configurations which yield the maximum stability guarantee, despite the magnitude of the guarantee itself, will be best-equipped for operation in a noise-filled environment.

V. EXPERIMENTAL RESULTS

We verified these predictions experimentally using the holonomic floating platform pictured in Figure 2. This vehicle carried out map exploitation and attempted to execute the trajectory of Figure 3 using feature spacings of 0.3m, 0.05m, and 0m. These spacings were chosen since 0.3m is near the asymptotic upper limit of the performance metric in Figure 5, the performance metric for 0.05m is 50% smaller, and 0m has a performance metric of zero, since $\|\Phi(k, k_0)\|$ does not converge to zero when all three features share the same location (and so this case is never guaranteed asymptotically stable). The top of Figure 6 demonstrates that the 0.3m spacing yielded the most effective closed-loop vehicle, as the platform adhered closely to the nominal

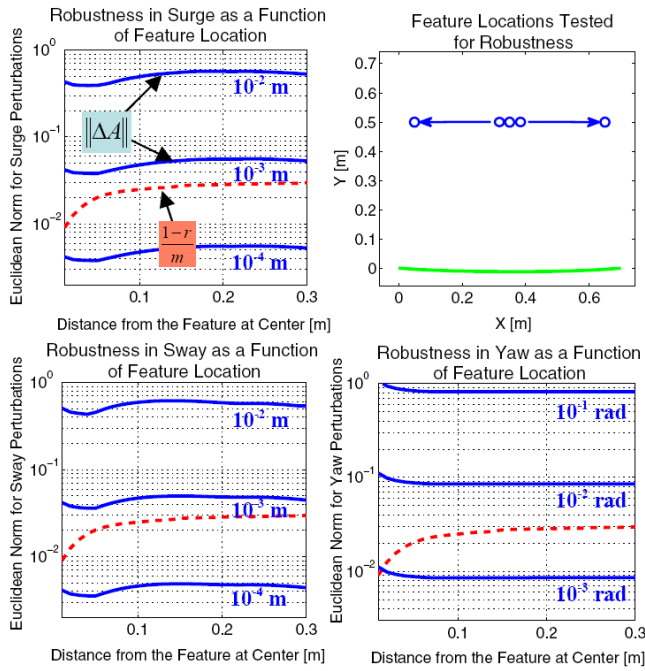


Fig. 5. The perturbation matrix norm is plotted in each degree of freedom as a function of feature spacing for disturbances of incrementally varying magnitude (captured on the solid blue lines). Alongside each set of perturbation matrix curves is plotted the exponential bounding curve performance metric (the dashed red line), indicating the maximum perturbation size for which each feature layout is asymptotically stable. The smallest feature spacing considered in this plot is 0.01m, since a spacing of 0m yields a performance metric of zero, which cannot be expressed on the above logarithmic scale.

trajectory and the estimator adhered closely to the vehicle's true location, although there were small errors in following the nominal trajectory due to perturbations in the water tank testing environment and interference from the vehicle tether. The 0.05m spacing caused errors in the estimator, which at times remained closer to the nominal trajectory than to the vehicle's true location and thus the vehicle failed to execute the entirety of the nominal trajectory. In the 0m spacing case (which effectively consists of a single feature on the map) the estimator failed completely and the vehicle was driven off of the nominal trajectory. The bottom of Figure 6 isolates performance in the angular degree of freedom, and compares filter estimation error with true displacement of the platform from the nominal trajectory. Here it becomes clear that in the 0.05m and 0m cases, instability results, as the estimator failed to converge to the true location of the vehicle. Although the robustness performance metric is only a sufficiency condition and doesn't explicitly predict vehicle instability, in this case it called for selection of the map best-suited for localization and closed-loop control, avoiding less robust feature configurations which ultimately brought about system instability.

Looking ahead to implementation on the HAUV itself, the robustness metric was also used to investigate which of two candidate trajectories is a more robust ship hull survey trajectory. The HAUV hull survey trajectories depicted in Figure 7 were selected heuristically based on simplicity

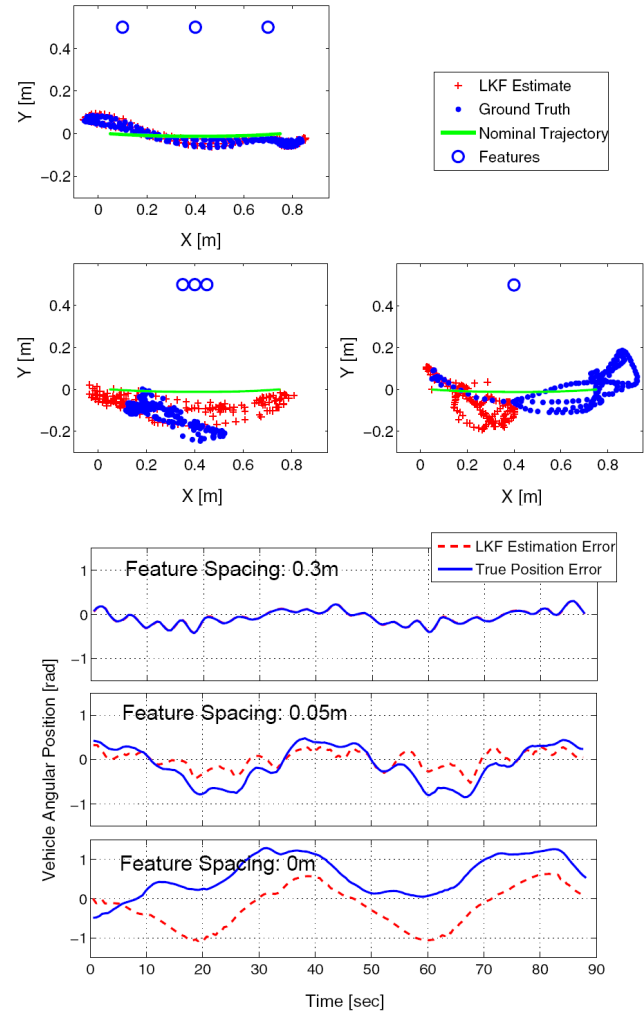


Fig. 6. The floating platform of Figure 2 was driven along the nominal trajectory of Figure 3 for three different configurations of three point features. In the upper half of the figure, the feature locations, nominal trajectory, vehicle position estimates, and vehicle position ground truth are plotted for each case. In the lower half, the angular degree of freedom is isolated and estimation error is plotted alongside vehicle displacement from the nominal trajectory. All plots display data from two complete cycles of the nominal trajectory.

and apparent hull coverage, and both have been used frequently in vehicle experiments performed on the hull that is pictured, a flat-bottomed boat approximately 18 m in length. Trajectory 1 orients the imaging sonar footprint along the width of the hull, and Trajectory 2 orients the sonar footprint along the length of the hull. The robustness of these trajectories is harder to predict than the trajectory of Figure 3 since the viewing window of the HAUV imaging sonar is limited in size, and only a subset of the features will be observed at each time step. If the cruising speed of the vehicle is varied, the frequency with which features will be observed will vary in turn. The hull survey was treated as a case of map exploitation, and locations of the hull's zinc anodes were used as features on an accurate given map. Odometry in each degree of freedom was added among the measurement capabilities since there are short

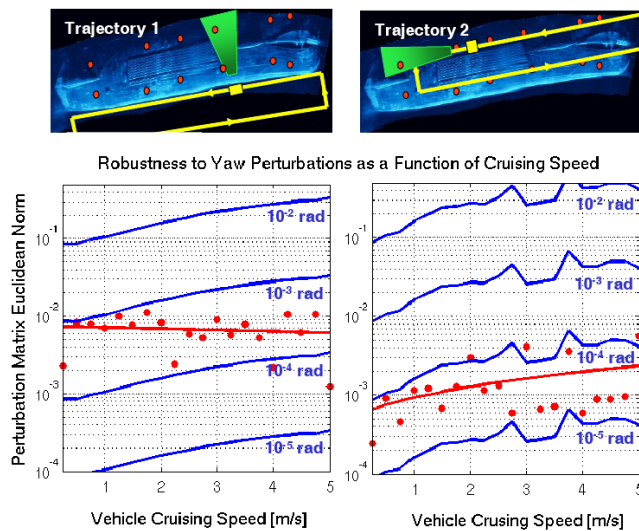


Fig. 7. The robustness performance metric is used to evaluate two HAUV candidate trajectories, for which the vehicle path, sonar viewing window, and feature locations are depicted. The perturbation matrix norm is plotted in the angular degree of freedom as a function of vehicle cruising speed for angular 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).

instances along the trajectories where no features are visible. The addition of odometry still permits the expression of the closed-loop vehicle as a linear time-varying system, and so the robustness analysis techniques continue to apply. A simulation of the HAUV's vehicle dynamics was used to apply Theorem 3 to each candidate trajectory for a variety of cruising speeds, and the results demonstrated that Trajectory 1 is guaranteed to tolerate perturbations approximately an order of magnitude greater than those for which Trajectory 2 is equipped. Figure 7 displays these results in the angular degree of freedom. Hence, our performance metric would indicate that Trajectory 1 is the robustness-optimal choice from among the two candidates.

VI. CONCLUSION

First, it has been shown that $\Phi(k, k_0)$, a holonomic marine vehicle's linearized time-varying transition matrix, can be used as a necessary and sufficient indicator of stability in inspecting the vehicle's controller. Analysis of $\Phi(k, k_0)$ also reveals that map refinement (i.e., localization and mapping) can achieve Lyapunov stability at best, while map exploitation, which is strictly localization, can achieve asymptotic stability.

Second, it has been shown that by considering perturbations to the system, a performance metric can be derived which serves as a sufficient condition for asymptotic stability in the presence of perturbations. This metric also serves a more useful purpose of allowing a comparison of robustness between different system configurations. The performance metric allows us to discern the impact of subtle aspects of

filter conditioning, such as the geometric configuration of map features, on the robustness of the marine vehicle.

Although the stability and robustness analysis tools presented here were developed for the analysis of vehicles employing feature-based navigation processes, they can also be applied to any system whose closed-loop dynamics may be expressed using a linear time-varying state space model. There are no additional limitations which restrict the quantity or type of sensors that may be blended to produce the state estimate, allowing systems like the HAUV, which integrate a large suite of navigation sensors, to be accommodated.

In addition to evaluating feature configurations, computing the robustness performance metric yields a procedure for discerning which of several candidate trajectories is best-equipped to tolerate perturbations, which we aim to develop into a robustness-optimal motion planning algorithm. By augmenting A* to consider robustness in addition to path length in its cost-to-come, we hope to enable hull surveys to divert from the shortest survey path to gain robustness against perturbations.

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] 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.
- [5] 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.
- [6] 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.
- [7] 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.
- [8] E. Belcher, B. Matsuyama, and G. Trimble, "Object Identification with Acoustic Lenses," *Proc. IEEE OCEANS Conf.*, Honolulu, 2001, pp. 6-11.
- [9] R. Madhavan, G. Dissanayake, and H.F. Durrant-Whyte, "Map Building and Map-Based Localization in an Underground Mine by Statistical Pattern Matching," *Proc. IEEE Int. Conf. on Pattern Recognition*, vol. 2, 1998, pp. 1744-1746.
- [10] P. Corke, "Mobile Robot Navigation as a Planar Visual Servoing Problem," R. Jarvis and A. Zelinsky, eds., *Robotics Research: The Tenth International Symposium*, Heidelberg: Springer-Verlag, 2003, pp. 361-371.
- [11] A. Gelb, ed., *Applied Optimal Estimation*, Cambridge, MA: The MIT Press, 1984.
- [12] J.C. Willems, *Stability Theory of Dynamical Systems*, London: Thomas Nelson and Sons Ltd., 1970.
- [13] B. Chen and T. Dong, "Robust Stability Analysis of Kalman-Bucy Filter Under Parametric and Noise Uncertainties," *Int. J. Control*, vol. 48(6), 1988, pp. 2189-2199.
- [14] A. Weinmann, *Uncertain Models and Robust Control*, New York: Springer-Verlag/Wien, 1991.
- [15] G. Chen, ed., *Approximate Kalman Filtering*, Singapore: World Scientific, 1993.