2008 IEEE International Conference on
Robotics and Automation
Pasadena, CA, USA, May 19-23, 2008

Stability of Double-Integrator Plants Controlled using Real-Time SLAM Maps

Franz Hover

Abstract— **The position control of agents having double-integrator dynamic response, with explicit dependence on an estimator that simultaneously reconstructs agent states and refines a feature map, is considered. We show that under broad conditions the phase and gain margins of an equivalent control system based on absolute navigation sensors can be significantly eroded when the dynamic map is used. The closed-loop bandwidth achieved depends strongly on the bandwidth of the filter, which is in turn dictated by the ratio of process to sensor noise, on the control gains, and on the arrangement of features bing tracked. The findings are relevant to the integration of SLAM algorithms in high-performance positioning of mobile agents. A specific example is the case undersea robots working near large structures, where traditional navigation systems may not be practical.**

I. INTRODUCTION

The objective of a simultaneous localization and mapping (CML or SLAM) algorithm is to achieve at the same time localization of an agent within an environment, and continual refinement of a map of environmental features. The feature map that is constructed augments or may completely replace a conventional set of positioning sensors. For instance, a mobile robot may have odometry sensors (subject to drift), but no heading or absolute position data. In these cases, sensors that measure proximity and orientation relative to walls, openings, or other structures are used as primary inputs for a SLAM algorithm.

Early work on this topic includes Smith *et al.* [1], Moutarlier and Chatila [2], and Leonard & Durrant-Whyte [3], where an Extended Kalman Filter was applied to track a large state vector containing conventional vehicle states expressed in absolute coordinates and map feature locations. More recent work has focussed on dealing with complexity in large maps. The EKF has been used extensively, and was assessed critically by Dissanayake *et al.* [4], who showed that the maps converge and in fact that an agent starting with no *a priori* information can localize itself. The probabilistic scheme by Thrun *et al.* [5] is a high-dimension maximum-likelihood estimator, with enhanced robustness properties that are not shared by the EKF.

The SLAM algorithm was studied at a more fundamental level by Gibbens *et al.* [6], considering a single agent whose velocity is directly driven by a control and a random process ($\dot{x} = u + q$). The state vector is augmented with feature positions. The objective of this canonical problem is to estimate the relative positions of the features from a vehicle velocity measurement and a number of (identically noisy) range measurements; the paper illustrates that the Fisher information divided by the feature process noise is a critical parameter governing the rate of convergence in the feature ranges: $\sqrt{qn/r}$ is the characteristic time constant, where q is the process noise, r the sensor noise, and n the number of range sensors. The basic structure of the observation operator for the SLAM problem is also brought out in the example from this paper; measurements are regular combinations of elements in the aggregate state vector, resulting naturally from range and angular measurements. Beyond the single degree-of-freedom case, the observation operator depends strongly on the geometry of the map. Mourikis and Roumeliotis [7] show for a robot moving at nonzero speed within a plane, the maximum positioning error covariance scales with $\sqrt{qr/n}$, a result also seen in the case treated by Gibbens *et al.*

Complex or simple, the great majority of work with SLAM has been performed in the context of a mobile agent that is either not controlling its own motions using the map, or is controlling them at a bandwidth that is safely below the that of the mapping filter. Considering the case of a truly dynamic vehicle (e.g., an underwater vehicle or unmanned aircraft that can drift), and low-level flight control using localization in a map, clearly the mapping and the feedback processes have to be integrated to achieve the best performance and robustness. Indeed such a consideration is unavoidable when the agent has to rely on the map for its own dynamic stability, and the features are the only navigational aid. One real-world scenario is autonomous underwater vehicle inspection of a ship hull, where a compass and long-baseline acoustic navigation may not be available. In such cases it is also reasonable to ask why is high-fidelity-control required while building a map? One answer is that an agent whose main task is to develop a good map in minimum time, and with limited observations, benefits from controlled motions, such as driving in a straight line while rejecting disturbances (see, e.g., [8], [9]). Once again, the integration of the map-building and the control processes is critical.

Here we study this fundamental interaction, fixing the physical plant to be a double integrator, or a coupled set of double integrators operating in a plane. A specific control system is explored that seeks to position the agent relative to the estimated features, making direct use of basic SLAM algorithm results. Our scope is to understand the limits of robustness and performance in this closed-loop system, under the assumptions of linearization, and steady-state estimator

Research is supported by the Office of Naval Research Grant N00014-06-10043, monitored by Dr. T.F. Swean.

F. Hover is with the Department of Mechanical Engineering, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139 USA hover@mit.edu

and control gains. Although the scenarios we describe may not be useful in today's common applications, we believe that the insight obtained from a straightforward, classical analysis will inform the broader problem.

II. SINGLE DEGREE-OF-FREEDOM DOUBLE INTEGRATOR

A. Statement of Problem

Expanding the notation from Gibbens *et al.* [6], the formulation is as follows:

$$\begin{Bmatrix} \ddot{x} \\ \dot{x} \\ \dot{p}_1 \\ \vdots \\ \dot{p}_n \end{Bmatrix} = \begin{bmatrix} 0 & 0 & 0 & \cdots \\ 1 & 0 & 0 & \cdots \\ 0 & 0 & 0 & \cdots \\ \vdots & & & \\ \vdots & & & \end{bmatrix} \begin{Bmatrix} \dot{x} \\ x \\ p_1 \\ \cdots \\ p_n \end{Bmatrix} + \begin{bmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} X + \begin{bmatrix} 1 & 0 & 0 & \cdots & \\ 0 & 0 & 0 & \cdots & \\ 0 & 1 & 0 & \cdots & \\ \vdots & & & 1 & 0 \\ \vdots & & & 0 & 1 \end{bmatrix} \begin{Bmatrix} w_1 \\ w_2 \\ \cdots \\ w_{n+1} \end{Bmatrix}$$

$$\underline{\dot{x}} = A\underline{x} + Bu + G\underline{w} \tag{1}$$

where the agent position and velocity are x and $\dot{x}$, respectively, the n feature locations are p_i, the control input is X, and the process noise $\underline{w}$ is Gaussian white noise, with diagonal covariance matrix Q. The observation set includes an agent velocity, absolute position, and the set of n ranges:

$$\underline{y} = \begin{bmatrix} 1 & 0 & \cdots & \\ 0 & 1 & 0 & \cdots \\ 0 & 1 & -1 & \\ \vdots & & & \\ 0 & 1 & \cdots & -1 \end{bmatrix} \begin{Bmatrix} \dot{x} \\ x \\ p_1 \\ \cdots \\ p_n \end{Bmatrix} + \begin{Bmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \\ \cdots \\ \nu_{n+2} \end{Bmatrix}$$

$$= C\underline{x} + \underline{\nu} \tag{2}$$

The sensor noise is Gaussian white noise with diagonal covariance matrix R. As is well-known in filter design, the relationship between Q and R determines the bandwidth of the filter. Since our main objective in this paper is to understand the dynamic interaction between a SLAM filter and a controller, we have allowed for the feature process noise to be nonzero. This contradicts the usual SLAM procedure of allowing no process noise in the feature locations, consistent with the construction of a static map. On the other hand, it is well-known in the case of parameter estimation using the EKF that process noise is required if the correct parameter values are to be recovered [10]. In the present paper, our emphasis is on the control loop stability, not bias in the filter. In our results below, we do consider zero process noise as a limiting case.

The observation equations as written contain absolute velocity and position measurements, which can be effectively removed from the filter design by making their noise levels very large. Many robots have velocity measurements available, for example from a Doppler velocimetry logger, whereas the absolute position is maintained here only to keep the system formally observable.

The regulating control law is written as: $X = [-k_d \; -k_p \; -k_p/n \; \cdots \; -k_p/n]\,\underline{\hat{x}} = -K_c \; \underline{\hat{x}}$. The double integrator plant with this controller gives the simple characteristic equation $s^2 + k_d s + k_p = 0$ for the closed-loop system without the estimator dynamics. If the gains are chosen for damping ratio $\sqrt{2}/2$, this achieves sixty degrees phase margin and infinite upward gain margin, the same margins afforded by a steady-state Kalman filter loop. The steady-state estimator evolves according to the usual linear filter equation $\dot{\underline{\hat{x}}} = A\underline{\hat{x}} + K_f\left(\underline{y} - C\underline{\hat{x}}\right)$.

As we describe below, the most interesting features of the SLAM control problem are evident when neither the absolute position nor the velocity is measured, although certainly each variation deserves to be considered in turn. This limiting condition is a range-only problem in one dimension, or a range-and-bearing-only problem in two dimensions (next section).

The system model above is also representative of a cruising vehicle, wherein the ranges are employed to control heading:

$$\begin{Bmatrix} \dot{\phi} \\ \dot{x} \end{Bmatrix} = \begin{bmatrix} 0 & 0 \\ U & 0 \end{bmatrix} \begin{Bmatrix} \phi \\ x \end{Bmatrix} + \begin{bmatrix} cU \\ 0 \end{bmatrix} \delta + \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{Bmatrix} w_1 \\ w_2 \end{Bmatrix}. \tag{3}$$

Here the heading is ϕ and the cross-track position is x; the steady speed is U, and c is a constant. The cruising problem is nearly identical to the one-dimensional hovering problem above, with the exception that physical process noise can be accepted in both of the states. Namely, the vehicle is perturbed in the heading and the cross-track directions independently. These are imposed as kinematic not inertial effects, however.

B. Restrictions

The controller and filter gains are taken here to be time-invariant. This point emphasizes that our analysis is focussing on stability and bandwidth properties at a nominal condition. Clearly when the agent moves a large distance relative to the features, then the steady-state assumption should be reevaluated, and the time-varying Riccati equation can be used.

We assume that the environment is always observable, that is, that no particular motion on the part of the agent is required to resolve the feature locations. The problem of choosing trajectories so as to maximize the quality of the feature estimates has been addressed by Feder *et al.* [11], among others.

We do not consider here the changes in system dynamic behavior that arise from the arrival or loss of feature ranges. Similarly, there is no account in our analysis of feature association errors; in realistic problems, such errors will have a major impact on performance of the system.

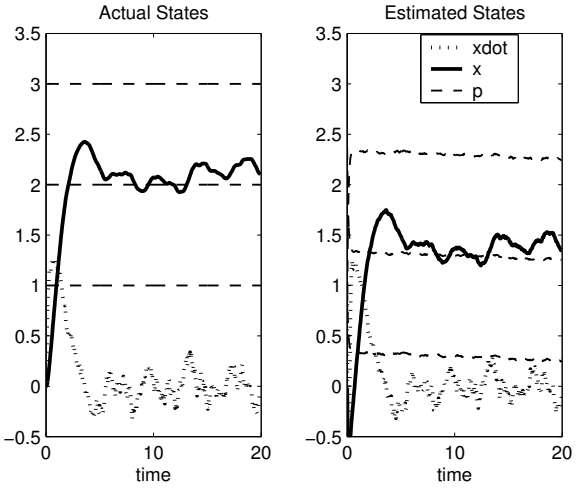

Fig. 1. Convergence of the vehicle position to the centroid of the features occurs at the bandwidth of the control; the slower dynamic response of the filter is not implicated because its mean is unchanged from the initial conditions, and evolves autonomously. This figure shows random initial conditions for the vehicle state and the feature estimates, with the three true feature locations [1,2,3]. Parameters are: n = 3, $r_{11} = r_{22} = 1e10$, $r_i = 1e{-}4$, $q_{11} = 1.0$, $q_{22} = 0$, $q_i = 0.01$, and $\omega_c = 1$. The continuous-time system was discretized in a manner consistent with the Ito calculus.

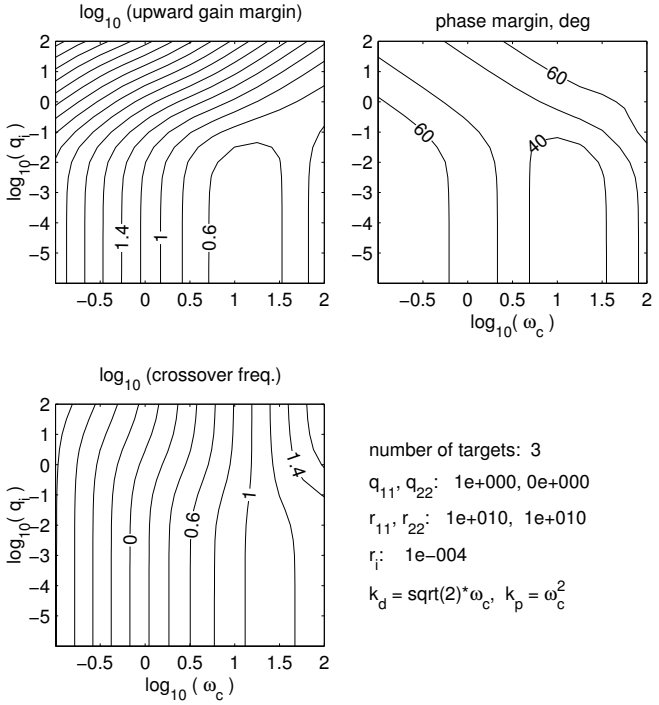

Fig. 2. For positioning in one dimension, gain and phase margins along with achievable bandwidth of the closed-loop system are given, as a function of control bandwidth ω_c and feature process noise q_i.

C. Estimator Properties

The C-matrix has a specific difference structure, and as noted, measurements of $\dot{x}$ and x, vehicle velocity and position, are included here mainly so as to make the system fully observable for computations. When considering the SLAM problem, r_x and $r_{\dot{x}}$ can be set very large so that these measurements are not used by the filter. There is no conceptual difference between this latter problem and reformulating the system with position states $x - p_i$.

The feature estimated locations are allowed to drift, with their mobility scaled by the process noise in each channel and governed also by the range measurement noise, as is the norm in filter design. Relative to the base case of only a velocity and an absolute position measurement (that is, no features), with position sensor noise covariance equal to that of the ranges, the SLAM control problem is differentiated primarily by the choice of process noise level on the feature locations, q_i. Clearly, setting these to large values will enable the estimates to converge more quickly. When they are small, then the measurements have the same low-frequency stability as does an absolute measurement, and the characteristics of the base case are recovered exactly, with the necessary addition of a pole at the origin.

The estimator poles for this case are characterized as follows. With $q_{22} = 0$, for any number of features greater than zero, there exists one pole at the origin, $n-1$ real poles at $-\sqrt{q_i/r_i}$, and two poles at the roots of

$$s^2 + s\sqrt{\frac{2\sqrt{q_{11}r_i n} + q_i}{r_i}} + \sqrt{\frac{q_{11}n}{r_i}} = 0. \tag{4}$$

This implies that the basic estimator bandwidth is $\omega_e = (q_{11}n/r_i)^{1/4}$, and the damping ratio is

$$\zeta_e = \frac{1}{2}\sqrt{2 + \frac{q_i}{\sqrt{r_i q_{11} n}}}. \tag{5}$$

These two poles become faster with increased n. The damping ratio is nominally $\sqrt{2}/2$, but increases with the mobility of the features q_i.

An interesting property of the system arises when the vehicle has no absolute position measurement: the estimator pole at the origin means that the drift is arbitrary, but the mean of the feature estimates changes very little. Because the controller in fact tracks the mean, we find that the vehicle can position itself well before the feature estimates have stabilized. This is illustrated in Figure 1. The controller as written will drive the system to the "middle" of the feature estimates. A second useful fact emerges on writing the total system dynamic equation:

$$\frac{d}{dt}\left\{\begin{array}{c}\underline{x}\\ \underline{\hat{x}}\end{array}\right\} = \left[\begin{array}{cc} A & -BK_c \\ K_f C & A - BK_c - K_f C\end{array}\right]\left\{\begin{array}{c}\underline{x}\\ \underline{\hat{x}}\end{array}\right\} \tag{6}$$

with the assumption of no modeling errors. Because of the structure in C, the filter gain matrix K_f also has a specific form; the columns [1,2,2+n+1,2+n+2] of $[K_f C \;\; A - BK_c - K_f C]$ are zero below the first two rows, so that the feature estimates do not depend at all on the actual vehicle states or their estimates. Hence the feature estimate dynamic response is autonomous, and this is a defining feature of the single degree-of-freedom SLAM-control problem.

D. Characteristics of the Closed-Loop System

The concatenation of an estimator and controller is generally referred to as the linear quadratic gaussian (LQG) regulator when the controller gains derive from a steady-state LQR formulation, and the filter employs a steady-state Kalman filter gain. While the robustness and performance properties of the LQR and the Kalman filter are well-known, the LQG inherits neither, and cases of extreme sensitivity exist [12], [13]. On the other hand, classical measures of robustness including the Nyquist and Bode plots remain useful in designing gain and phase margins to achieve maximum closed-loop bandwidth.

Free parameters for this problem include the process noise on the agent itself (intensity q_{11} and q_{22}), process noise on each of the features (q_i), assumed to be uniform across all features, the absolute sensor noise (r_{11} and r_{22}), the noise on each range measurement (r_i), the number of sensors n, and the design bandwidth of the basic control loop ω_c. Presumably $q_{\dot{x}}$, $r_{\dot{x}}$ and r_x, r_i can be chosen based on knowledge of the environment and the sensors. As noted by Gibbens *et al.* [6], any number of features for this problem is equivalent to the case of one feature with equivalent Fisher information. We have performed analysis of this system varying ω_c and q_i, with a fixed number of features.

In the general case, the performance and robustness of any such system critically depend on the interplay of the controller and the filter designs. Gain and phase margins, as well as crossover frequency for the present system are obtained by breaking the loop at the controller output, that is, at the signal u. Hence gain and phase margins are relative to the plant input; this is natural for a double-integrator plant, where the major uncertainties may be in the actuators themselves. Figure 2 illustrates a case with no velocity or absolute position information available. When the features' estimates are constrained by small q_i, slowing down the filter poles, interference with the controller leads to a significant degradation of gain and phase margins as expected: the base controller has infinite upward gain margin and sixty-degree phase margin. In these cases also, the crossover frequency varies with respect to the control bandwidth ω_c. The findings suggest that the designer should *not* set the feature drift q_i to zero, as in Dissanayake *et al.* [4], when robustness in the closed-loop system is a concern. If q_i is large, then we see the margins are improved toward the control loop robustness levels, and the overall bandwidth can be increased arbitrarily. Improved controller robustness is obtained at the cost of the feature position estimates moving around more freely, as in Figure 1.

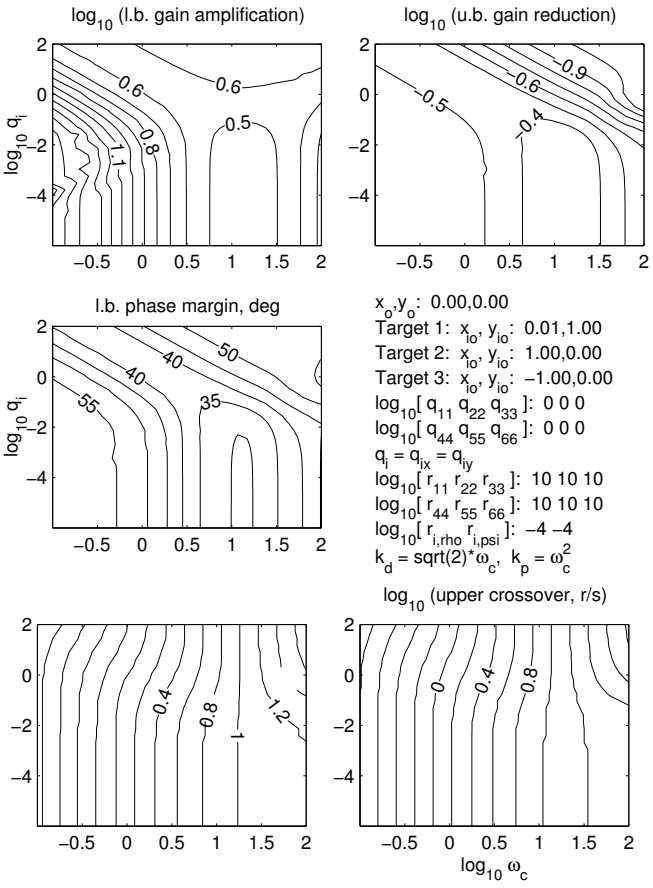

Fig. 3. For positioning in the horizontal plane, gain and phase margin bounds along with upper and lower bandwidths of the closed-loop system are given, as functions of control bandwidth ω_c and feature process noise q_i. u.b. and l.b. indicate upper and lower bounds, respectively.

III. THREE DOUBLE INTEGRATORS

A. Statement of Problem

A related holonomic two-dimensional problem, where the plant contains two translational and one rotational degrees of freedom, is as follows:

$$\begin{Bmatrix} \dot{u} \\ \dot{v} \\ \dot{r} \\ \dot{x} \\ \dot{y} \\ \dot{\phi} \end{Bmatrix} = \begin{bmatrix} 0_{3\times3} & 0_{3\times3} \\ I_{3\times3} & 0_{3\times3} \end{bmatrix} \begin{Bmatrix} u \\ v \\ r \\ x \\ y \\ \phi \end{Bmatrix} + \begin{bmatrix} I_{3\times3} \\ 0_{3\times3} \end{bmatrix} \begin{Bmatrix} X \\ Y \\ N \end{Bmatrix} + \begin{bmatrix} I_{3\times3} & 0_{3\times3} \\ 0_{3\times3} & 0_{3\times3} \end{bmatrix} \underline{w}_{3\times1}. \tag{7}$$

The forward velocity, sway velocity, and yaw rate are u, v, and ϕ, respectively, whereas x, y, and ϕ are their integrals. There are three control channels, X, Y, and N. The process noise is applied to the acceleration channels, with covariance $diag[q_{11}\ q_{22}\ q_{33}]$. As in the single integrator case, the state is augmented with the Cartesian feature locations.

The linearized range measurements to n features are given by:

$$\begin{aligned} \rho_{i,o} + \rho_i &= \sqrt{\tilde{x}^2 + \tilde{y}^2} + \nu_{i,\rho} \longrightarrow \\ \rho_i &= \frac{x_{i,o} - x_o}{\rho_{i,o}}(x_i - x) + \frac{y_{i,o} - y_o}{\rho_{i,o}}(y_i - y) + \nu_{i,\rho}, \end{aligned} \tag{8}$$

where $\rho_{i,o}$ is the nominal range to the i'th feature, and ρ_i its perturbation. We use the definitions $\tilde{x} = x_{i,o} + x_i - x_o - x$ and $\tilde{y} = y_{i,o} + y_i - y_o - y$, where x_o and y_o are the nominal vehicle positions, and x and y are their perturbations; $x_{i,o}$ and $y_{i,o}$ are the n nominal feature locations, and x_i and y_i are their perturbations. .

Additionally, an angular measurement to the i'th feature is available as

$$\psi_{i,o} + \psi_i \quad = \arctan\left(\frac{y_{i,o}+y_i-y_o-y}{x_{i,o}+x_i-x_o-x}\right) - \phi_o - \phi + \nu_{i,\psi}. \tag{9}$$

Here ϕ_o is the nominal vehicle heading and ϕ is its perturbation. We assume there are no quadrant changes, so that the $arctan(y/x)$ function replaces the usual $arctan2(y,x)$. This has the linearization

$$\begin{aligned} \psi_i &= -\frac{1}{1+g^2}\left[\frac{g}{x_{i,o}-x_o}\right](x_i - x) + \\ &\quad \frac{1}{1+g^2}\left[\frac{1}{x_{i,o}-x_o}\right](y_i - y) - \phi, \text{ where} \\ g &= \frac{y_{i,o}-y_o}{x_{i,o}-x_o}. \end{aligned} \tag{10}$$

The objective of the controller is to maintain zero position relative to the centroid of the features, and zero average bearing to the features. In the examples we show below, there is no compass or yaw rate measurement used.

B. Restrictions

As in the single-DOF case, this analysis assumes the estimator and control gains are constant. Motions from the nominal position are considered to be small, and all features are observed at all times.

C. Estimator Properties

The general properties of the filter loop, with no absolute sensors available, are as follows. Clearly there are three zero eigenvalues, corresponding with the drift of $[x, y, \phi]$. When only one feature is available, a fourth eigenvalue at zero occurs because the vehicle can pinwheel around the feature, while having the same range and bearing data. With two features, we see an increase in the bandwidth of the filter, consistent with the results above for a single degree of freedom, and only three zero poles; there are three complex pairs and a single (stable) real pole. For more than two features, real or complex pole pairs are added, while the overall system response becomes increasingly fast. This system does not share the autonomous feature estimate of the one-dimensional case, due to the richer structure of C; inspection of $K_f C$ will show that the evolution of $[x, y, \phi]$ now bears directly on the feature estimates.

D. Characteristics of the Closed-Loop System

Gain and phase margin estimates of the multivariable system are computed according to standard techniques [14]

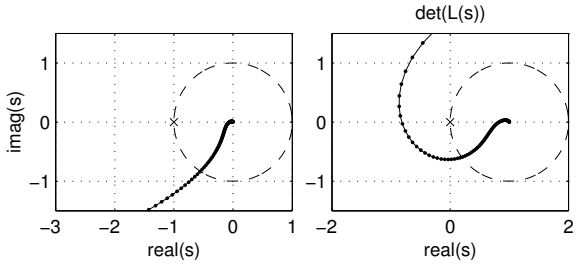

Fig. 4. Single-DOF (top) and multivariable (bottom) Nyquist plots show a basic difference that limits the gain reduction margin in the multivariable case. Shown is the case of $\omega_c = 1$, $q_i = 0.01$, with all other parameters the same as in Figures 2 and 3. $L(s)$ is the loop transfer matrix, with the loop broken at the plant input.

as:

$$gm \geq \max\left[1 + \frac{1}{||T||_\infty}\ ,\ 1 + \frac{1}{1 - 1/||S||_\infty}\right] \quad \text{(amplification)} \tag{11}$$

$$gm \leq \min\left[1 - \frac{1}{||T||_\infty}\ ,\ \frac{1}{1 + 1/||S||_\infty}\right] \quad \text{(reduction)}$$

$$pm \geq \max\left[2\arcsin\left(\frac{1}{2||T||_\infty}\right)\ ,\ 2\arcsin\left(\frac{1}{2||S||_\infty}\right)\right] \quad \text{(phase)}$$

where $S(j\omega) = (I + C(j\omega)P(j\omega))^{-1}$ and $T(j\omega) = I - S(j\omega)$; S is the sensitivity function and T is the complementary sensitivity. As in our one-dimensional case, the loop is broken at the controller output so that the gain and phase margins are referenced here. We note these are bounds only, that may be conservative. Also, the above formulas lead to a small amount of jitter and non-uniformity in the figures.

Results are shown in Figure 3, with a "good" - or distributed - geometry of three features. The parameters chosen for the plot are very similar to those for Figure 2, and hence the effects of the additional degrees of freedom are apparent. The upward gain margin of the planar problem does not go to high levels as it does in the single-DOF case, when q_i is large. Instead, values of about three to four are achieved over most of the range of q_i. Instability due to gain reduction does not occur for the single-DOF case, because the Nyquist plot has no encirclements of the critical point; in the multivariable condition, however, Figure 4 shows that the determinant of the open-loop transfer matrix $L(j\omega)$ can go to zero with a gain reduction of about one-half. The phase margin bounds in the planar positioning problem are slightly lower than the margin of the one-DOF case, but overall they are still acceptable, and in fact may be preferred from the point of view of overshoot and settling time. For the multivariable case, the margins given are guaranteed bounds, and may be conservative. Figure 4 suggests that for this particular example at least, the bounds are reasonably accurate.

The fastest and slowest crossover frequencies (from singular values) are quite close together, and this indicates that bandwidths among the surge, sway, and rotational degrees of freedom move together, in accordance with the use of a

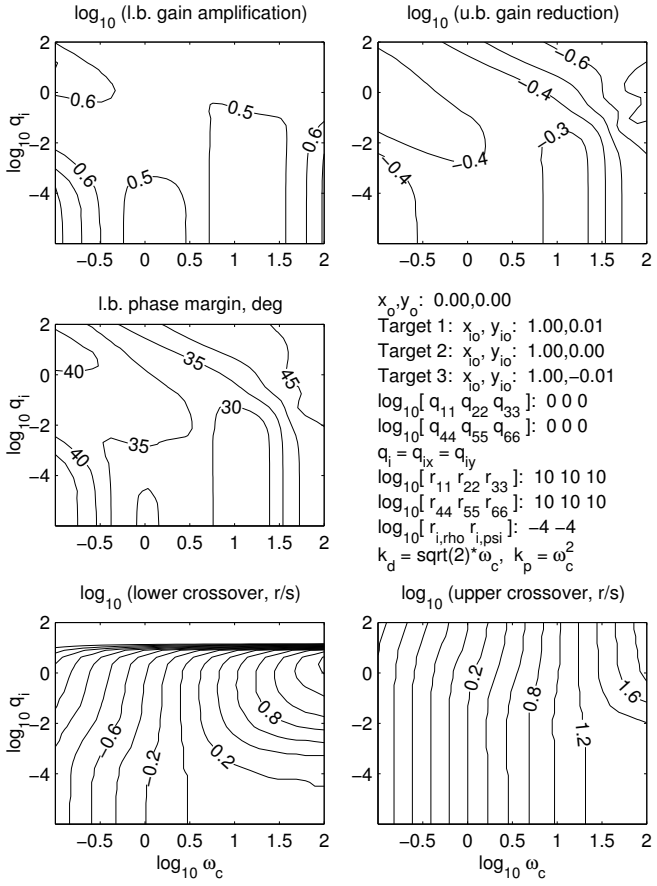

Fig. 5. Clustered feature locations affect the stability margins of the closed-loop system; except for the feature locations, this case is identical to that shown in Figure 3.

single derivative and proportional gain choice for all three.

Figure 3 shows results with three features spaced one-hundred-eighty degrees apart relative to the vehicle, that is at $[0\ 90\ 180]$ degrees. In Figure 5 we illustrate the result when the three features are close together, separated by only a few degrees. Intuition points to a far less robust system, and this is borne out by the fact that the gain amplification margin now does not necessarily increase at lower controller bandwidth, as seen when the features are spread out. Similarly, the reduction margin is somewhat poorer, and phase margin is degraded by five to ten degrees overall. Most notably, however, the lower crossover frequency has an entirely new characteristic: at low q_i, the highest bandwidth that can be achieved is less than 1.5, even with very high controller gains. However, this lower frequency moves to extremely low values with $q_i > 10$, the major slowdown beginning at $q_i \simeq 1$. Needless to say, either of these cases would constitute a poor operating condition.

IV. SUMMARY

Our goal has been to apply simple robustness analysis to several nominal cases in SLAM-based dynamic positioning, so as to understand the basic tradeoffs facing the designer of SLAM-capable agents. The one-DOF case is characterized by autonomous feature estimation due to a special structure in the closed-loop system matrix. At the same time, the control loop discussed operates relative to the mean of feature estimated positions, achieving a dynamic response that can be faster than that of the estimated features. In the multiple-dimension case, conventional bounds show that the geometry of the features can significantly influence robustness, with examples indicating a serious but not catastrophic degradation as the features cluster. The intensity of process noise on the feature states is an important factor, and significant robustness loss can occur if this noise level is set to zero.

Overall, behavior of the combined localization and control loop is affected by all of the governing parameters and conditions, and designers must account for the various interactions to create systems that operate with acceptable robustness.

REFERENCES

[1] R. Smith, M. Self, and P. Cheeseman, "A Stochastic Map For Uncertain Spatial Relationships," *4th Int. Symp. on Robotics Research*, MIT Press,1987.

[2] P. Moutarlier and R. Chatila, "Stochastic Multisensory Data Fusion for Mobile Robot Location and Environment Modeling," *5th Int. Symposium on Robotics Research*,Tokyo,1989, pp. 207-216.

[3] 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.

[4] G.M.W.M. Dissanayake, P. Newman, S. Clark, H.F. Durrant-Whyte, and M. Csorba, A Solution to the Simultaneous Localiziation and Map Building (SLAM) Problem, *IEEE Trans. on Robotics and Automation*, vol. 17(3), 2001, pp. 229-241.

[5] S. Thrun, W. Burgard, and D. Fox, A Probabilistic Approach to Concurrent Mapping and Localization for Mobile Robots, *Machine Learning*, vol. 31, 1998, pp. 29-53.

[6] P.W. Gibbens, G.M.W.M. Dissanayake, and H.F. Durrant-Whyte, "A Closed-Form Solution to the Single-Degree of Freedom Simultaneous Localisation and Map Building (SLAM) Problem," *Proc. 39th IEEE Conf. on Decision and Control*, Sydney, 2000, pp. 191-196.

[7] A.I. Mourikis and S.I. Roumeliotis, "Analysis of Positioning Uncertainty in Simultaneous Localization and Mapping (SLAM)," *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Sendai, Japan, 2004, pp. 13-20.

[8] M. Bryson and S. Salah, "Active Airborne Localisation and Exploration in Unknown Environments using Inertial SLAM," IEEE Aerospace Conf. Proc., Big Sky, MT, USA, 2006.

[9] C. Leung, S.D. Huang, N. Kwok, and G. Dissanayake, Planning Under Uncertainty Using Model Predictive Control for Information Gathering, *Robotics and Autonomous Systems*, vol. 54(11), 2006, pp. 898-910.

[10] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, Estimation with Applications to Tracking and Navigation. New York: John Wiley & Sons, 2001.

[11] H.J.S. Feder, J.J. Leonard, and C.M. Smith, Adaptive Mobile Robot Navigation and Mapping, *Int. J. Robotics Research*, vol. 18(7), 1999, pp. 650-668.

[12] J.C. Doyle, Guaranteed Margins for LQG regulators. *IEEE Trans. on Automatic Control*, vol. 23(4), 1979, pp. 756-757.

[13] J.C. Doyle and G. Stein, Robustness with Observers, *IEEE Trans. Automatic Control*, vol. 24(4), 1979, pp. 607-611.

[14] N.A. Lehtomaki, N.R. Sandell, and M. Athans, Robustness Results in Linear-Quadratic Gaussian Based Multivariable Control Designs, *IEEE Trans. Automatic Control*, vol. 26(1), 1981, pp. 75-92.