LINEARIZED HOVERING CONTROL WITH ONE
OR MORE AZIMUTHING THRUSTERS

F.S. Hover, Massachusetts Institute of Technology, United States

SUMMARY: We propose a simple method of control system design for marine vehicles with one or more azimuthing propulsors, and specifically for the case where the speed of the actuator is on the same time scale as the plant dynamic response, thus making the assumption of a separation of time scales invalid. By setting a fixed, regular azimuth trajectory, the control problem is simplified sufficiently to allow a fully linear design approach, for which bandwidth achieved, robustness, and disturbance and noise rejection, will be more tangible than in the nonlinear cases. Several simulation examples are presented for a new vehicle that is in development; the approach would apply directly to the cases of multiple propulsors and dynamic positioning as well.

1. INTRODUCTION

The use of azimuthing propulsors in both ships and floating structures is well-established and increasing. The recent development of high-power electric motor components, which can be mounted and azimuthed outside the hull, both frees up layout space inside the vessel, and enhances maneuverability. The thrusters can be used for main propulsion and as steering devices while underway, and as directional jets for station-keeping.

Large dynamically-positioned craft usually rely on comparatively high-speed actuators in the control system design: the dynamic response of the thruster (both the propeller and the azimuth drive) is assumed to be much faster than the vessel response. This allows a nearly quasi-static view of the thrusters, in the sense of their directional force production. They may, however, be subject to slew rate and azimuth position limitations. Overall, dynamic positioning today is predicated on *de facto* separation of a vessel controller - providing forces and moments in the frame of the vessel - and a separate process which drives the thrusters and their azimuth angles, so as to best reconstruct the command from the vessel controller. To consider a few examples illustrating the quasi-static assumption, the controller described by Sorensen *et al.* achieves a notch at the resonant surge frequency of a semi-submersible, with an effective bandwidth in the neighborhood of $0.02Hz$ (1). Another paper by the same author uses a pseudo-inverse solution to achieve control of a $4200t$ vessel; the data show that the vessel can follow trajectories with a time constant of better than about one hundred seconds (2). In these cases and others, the quasistatic assumption seems valid, works well, and can be developed more fully for enhanced characteristics, e.g., (3,4,5,6).

For certain marine systems, however, this separation of vessel and thruster controllers is inadequate. As with most inner-outer control designs, the critical condition is that the inner loop (thruster controller) achieves a much faster closed-loop bandwidth than does the outer loop; the factor in time scales is usually at least three to five. When this assumption fails, for example if the thrusters can provide enough authority to drive the plant at a concomitant bandwidth, a more integrated approach has to be taken. Sordalen, for example, poses a controller design scheme that does limit the rate and bandwidth of the azimuth action, but only through tuned filters and nonlinear components; hence an explicit estimate of specific robustness or performance levels is unavailable (7). More broadly, it is uncertain from these uncoupled, or inner-outer, loop designs whether systematic robustness against modelling errors is achieved. Similarly, such fundamental properties as closed-loop bandwidth and the impact of sensor noise and disturbances, cannot be easily determined.

The current work studies the application of azimuthing thrusters in a case where the time scale of azimuthing action is close to that of the vehicle's dynamic response. We consider in particular a vehicle that is being developed at the Massachusetts Institute of Technology Sea Grant Autonomous Vehicles Laboratory. As shown in Figure 1, this is a streamlined craft with an unusual arrangement of thrusters. This configuration derives from the mission it is designed for:

1. a very fast descent from the surface to a maximum depth of $3000m$. The descent is largely unpowered, and employs a large drop weight.

2. a short survey or sampling mission, which requires low-speed (0.5-1 m/s) flight or hovering

in place;

3. a very fast ascent to the surface, having released a second drop weight.

The vehicle is intended for fast exploration in a sparse spatial domain. It takes advantage of the fact that surface vessels - that is, oceanographic ships - travel much faster than most submerged AUV's. Other attributes of its mission design are detailed elsewhere (8). The vehicle as designed has four thrusters; two cross-body units and two more mounted on a servo-controlled pitching assembly with a slip ring. Hence, the thruster pair can be rotated to any pitch angle without a limit to the accrued number of turns. The pitching thrusters were chosen for two reasons. First, the fast ascent and descent require a low-drag vehicle, and any additional thrusters would add significant drag. Second, the thrusters are quite expensive relative to other components on the vehicle, even to the rotating assembly.

While azimuthing actuation is inherently a nonlinear control problem, our aim here is to develop some capability for the hovering problem with linear control. We believe that the significant advantages of linear control make it worthwhile to consider, namely ease of design and programming, known robustness, and allowance for state reconstruction. Indeed, there are ample cases documented where simple controllers, such as the linear quadratic regulator we consider below, achieve good performance and robustness for strongly nonlinear plants, even relative to some very elegant nonlinear control schemes, e.g., (9).

It should be noted that hovering control is fundamentally different than the cruising problem, where the thrust vectoring can be considered a perturbation to the steady powering condition; in hovering, both axes of the azimuthing thrusters have to be used in an unpredictable way, to reject disturbances and sensor noise. Because the desired trajectory is null, no feedforward terms, of the type seen in feedback linearization approaches, are applied.

2. NONLINEAR THRUSTER CONTROL

The basic problem in the case of controllable angle (e.g., pitch) u_p and thrust u_t is illustrated by the simplified equations governing the new Sea Grant vehicle:

$$
\begin{aligned}
m\dot{U} &= f(U, V, \dot{\theta}, \theta) + cos(u_p)u_t & (1) \\
m\dot{V} &= g(U, V, \dot{\theta}, \theta) + sin(u_p)u_t \\
J\ddot{\theta} &= h(U, V, \dot{\theta}, \theta) + r_y cos(u_p)u_t + r_x sin(u_p)u_t,
\end{aligned}
$$

where U and V are the surge and heave body-referenced velocities, θ is the pitch angle, m and J are representative linear and rotational mass, respectively, and functions f, g, and h capture complex fluid

dynamic effects, which we will approximate and linearize for the example cases below. The pitch angle of the thruster is u_p, and the thrust level is u_t.

This system, not dissimilar to the dynamic positioning problem with azimuthing propulsors, is nonlinear in the inputs. Controller design quickly becomes a very difficult if there are significant rate limits, $|\dot{u}_p| \leq \overline{\dot{u}}_p$ or, even more paralyzing, an angular limit: $\underline{u}_p \leq u_p \leq \overline{u}_p$. Moreover, the azimuthing thrust arrangement fails basic tests (Lie brackets) for controllability, in sympathy with other types of non-holonomic physical systems (e.g., automobiles), or systems with nonholonomic actuation, such as ours (10). The steering sinusoids approach is the essential strategy for the general problem of this type, and it has been successfully applied to marine vehicles (11,12,13).

A simple quasi-static approach, employing a separation of controllers, proceeds naturally as follows:

- Develop a control scheme in the body degrees of freedom, e.g., surge, sway, yaw. We call this a *vessel control* in the sense that all actuator dynamics are put aside; we say for the case of full-state feedback that $u = -Kx$, where the control vector u comprises commands in the vehicle U and V directions, and x is the vehicle (rigid body) state vector.

- Create a mapping that transforms these body-referenced actions into thruster coordinates, e.g., azimuth angle and thrust level. We call this the *nonlinear thruster control*. One simple strategy we consider below, assuming a hard velocity constraint but no position constraints, is as follows:

$$
\begin{aligned}
u_{p,ref} &= atan2(u_U, u_V) & (2) \\
& \quad \text{(reference value for angle)} \\
u_t &= \sqrt{u_U^2 + u_V^2} \\
\dot{u}_p &= \overline{\dot{u}}_p \, sgn(u_{p,ref} - u_p)
\end{aligned}
$$

In the last equation, the signum function may be smoothed around the origin so as to eliminate chatter. Additionally, its argument should be clipped to the interval $(-\pi, \pi]$ so that the instantaneous preference will always be for the device to spin the shortest distance. These equations resist linearization about any given point because rejecting disturbances physically requires that both of the U and V degrees of freedom can be actuated. Clearly, however, if $\dot{u}_{p,max}, \dot{u}_{p,min}$ are large enough to keep up with the vessel control signal, then this is a reasonable approach. When redundant thrusters are available, this problem can be posed as a quasi-static optimization. Note that there are no free design variables in our thruster controller, other than the smoothing of the signum function. Nonetheless, the system behavior - including

performance and robustness - cannot be analyzed except through simulations.

The performance of a control system of this type is given in Figure 2 for a representative vehicle model. The model includes the pitch/heave coupling induced by the large horizontal fin area, as well as the righting moment due to buoyancy; it has a lightly damped mode at 1.53 radians per second and no zeros. The vehicle controller is an LQR design (the full state is assumed to be known), such that the slow and fast closed-loop poles are at about one and 3.6 radians per second, respectively. As is well-known, the LQR implemented perfectly carries a sixty degree phase margin, infinite gain reduction margin, and a factor of two gain amplification margin. While these are generally considered to be overly conservative relative to other optimal control techniques, the LQR remains a viable basic control strategy, and will be the focus of our work here.

The above nonlinear thruster control algorithm is applied to the LQR-generated control vector, with a cycle time of one second. This is reasonable, if a bit fast, from the point of view of the pitch drive system capability, and the ability of the propulsor to modulate thrust levels during one cycle. We see that a clear degradation of the closed-loop response occurs when the nonlinear thrust control is used, most notably in the surge DOF, where several oscillations appear. This controller also uses up to four times the thrust command seen in the direct LQR, and has significantly more azimuth action. These negative effects are the solely the result of the limitation on thruster slew rate; they are reduced for a faster cycle time in the azimuth, and exacerbated for a slower cycle time. In fact, with $\Delta t = 1.1$ seconds (not shown), we find the nonlinear thruster control is unacceptable, so it certainly cannot be considered a robust design relative to cycle time. The response also depends unpredictably on both the initial condition of the plant and on the initial angle of the thruster.

3. SEQUENTIAL LINEARIZATION

We now introduce the major approach of the paper, which we will show provides a more easily designed and better performing closed-loop system than the nonlinear thruster control above. The key concept is that of a constant azimuth rate, or alternatively, a cyclic azimuth trajectory, which reduces the nonlinear input problem to the linear time-varying form

$$\dot{x} = Ax + B(t)u, \qquad (3)$$

where the matrix $B(t)$ captures the variations in thrust direction with terms of the form $cos(2\pi t/\Delta t)$ and so on. It is not dissimilar to the steering sinusoids mentioned previously; this is a case of so-called extended linearization, enabled by the explicit elimination of the angle control variable.

Typical feedback approaches for dealing with such time-varying systems derive from gain-scheduling, wherein a family of control gains is generated as if each point of operation were the only one. This can work quite well in practice, although without further analysis it is unclear what are the limits of the approach. Clearly, any set of static optimal controllers based on piecewise linearization will suffer degraded performance and robustness properties, because a given steady-state Riccati equation minimizes an integral cost, as opposed to an instantaneous one. At the next linearization point, the design from the previous one has become suboptimal. This deleterious fact is illustrated in Figure 3, for a LQR design with an undamped oscillator plant having natural frequency of one radian per second. The top row shows the performance of the discrete-time LQR with a full-width zero-order hold assumption: that is, the control is assumed constant across Δt, and this is how it is applied. The second row shows the performance of an equivalent controller design with the discretization and design time step set to $\Delta t/2$. The control is only executed, however, on alternate steps, that is, once per Δt. We observe that for smaller Δt, this approach is acceptable, but that serious degradation occurs as the Nyquist rate $(\pi/\Delta t)$ nears the natural mode.

This uncertainty can be completely eliminated by translation of the continuous plant into discrete-time, employing the idea of *sequential linearization* to expand the control vector, and then designing controllers entirely in the discrete domain. While the discretization involves no significant approximations, the implementation of any discrete-time controller still has a strong dependency on sampling rate. Considering the field of discrete linear systems, several basic methods are available for assessing their stability, under structured and unstructured uncertainty models (14, 15). Guaranteed robustness margins of the discrete-time linear quadratic regulator depend on penalty matrices Q and R, unlike the continuous-time case, and are not nearly as good as are found in the continuous-time case. The margins do converge in the cases of large R (expensive control) and of course when the time step is very fast compared to the system dynamic response (16). For this reason, the question of robust performance in controller design was addressed by several authors using a special cost function comprising a robustness part and a separate LQR performance part. The robustness part typically makes use of the generic methods for stability analysis mentioned above (17). Our expansion of the control vector into sequential components is notably similar to the expansion used for discrete-time control with computation delays, for which several papers indicate a strong variation in the robustness properties

as control cost R is changed (18, 19).

We now describe the implementation of the sequential linearization. We consider here a single actuator, although every aspect of the approach is easily extended to the case of multiple actuators. Given a basic linear time-invariant system description $\dot{x} = Ax + Bu$, the zero-hold model for an actuator requires the level of the actuator output to be constant over the sampling interval. This is expressed as:

$$x(t_{k+1}) = e^{A\Delta t}x(t_k) + \int_{t_k}^{t_{k+1}} e^{A(t_{k+1}-s)}Bu(s)ds, \quad (4)$$

and is recognized as a convolution over the sampling interval (20). Next we invoke the notion of a steady physical cycle time; in the case of a thruster azimuthing at constant rate, this cycle time corresponds with one turn. This is the same as $\Delta t = t_{k+1} - t_k$. The idea now is that one physical cycle can be split into n non-overlapping time intervals (but not necessarily evenly spaced), wherein the thrust action u_i is in force, and constant over the interval. The equivalent n-input system, with sampling coincident with the beginning of the first interval per cycle is

$$
\begin{aligned}
x(t_{k+1}) &= e^{A\Delta t}x(t_k) + \qquad\qquad (5)\\
&\quad \int_{t_k}^{t_k+\Delta t} e^{A(t_{k+1}-s)}B_1 u_1(s)ds +\\
&\quad \int_{t_k+\Delta t/n}^{t_k+2\Delta t/n} e^{A(t_{k+1}-s)}B_2 u_2(s)ds + \cdots\\
&= e^{A\Delta t}x(t_k) +\\
&\quad \int_{t_{k+1}-\Delta t/n}^{t_{k+1}} e^{A(t_{k+1}-s)}ds \times\\
&\quad \sum_{i=1}^{n} e^{A\Delta t(n-i)/n}B_i u_i(t_k)\\
&= \Phi x(t_k) + \Gamma' \times\\
&\quad \left[e^{A(n-1)\Delta t}B_1 \; e^{A(n-2)\Delta t}B_2 \; \cdots \right.\\
&\qquad \left. e^{A\Delta t}B_{n-1} \; B_n \right] u(t_k)\\
&= \Phi x(t_k) + \Gamma u(t_k).
\end{aligned}
$$

Here, B_i denotes the i'th column of the input matrix B. Note that this analysis includes no approximations, except for the zero-hold behavior of the input; the mapping could be smoothed with more linearization intervals.

For the oscillator example of Figure 3, wherein a cycle is split into two subcycles, the third row of plots indicates that the half-width control action is comparable with, but not exactly the same as, the full ZOH approach. In this case, the system model which creates the LQR design has built into it the fact that the control is shut off for half of every Δt, and so maintains the basic capabilities and limitations of the full-width design. Concerning robustness, the guaranteed LQR bounds for this particular design

are inconsequential because of the specific design parameters used, but the resulting closed loop matrix $\Phi - \Gamma K$ can be assessed via Martin's and Hewer's procedure for unstructured perturbations of the form $eig(G) = eig(G_{design} + \Delta G)$. Stability is maintained for maximum singular values of unstructured variations in $\Phi - \Gamma K$ less than 0.42, 0.45, and 0.30, for the three time steps of [1.5, 2.0, 2.5] seconds, respectively. These are substantial margins, covering gain uncertainties, phase uncertainties, and real paramater variations in Φ.

We reiterate that any discrete-time control system is subject to the sampling theorem. This constraint is manifested by a loss of control of modes above the Nyquist rate, and reduced gain and phase margins if the Nyquist rate is not high enough. In the present case, the time steps are extremely large for the problem, as indicated by the discretization shown in Figure 3.

4. APPLICATION

Returning to the three-degree of freedom vehicle model, the thrust actuation can be decomposed into four separate portions we denote u_1, u_2, u_3, u_4; the first and third access the same DOF, as do the second and the fourth:

$$
\begin{aligned}
U,\theta: &\quad u_1, &\quad t_k < t < t_k + \Delta t/4 &\qquad (6)\\
V,\theta: &\quad u_2, &\quad t_k + \Delta t/4 < t < t_k + 2\Delta t/4\\
U,\theta: &\quad -u_3, &\quad t_k + 2\Delta t/4 < t < t_k + 3\Delta t/4\\
V,\theta: &\quad -u_4, &\quad t_k + 3\Delta t/4 < t < t_k + \Delta t.
\end{aligned}
$$

Hence, there are four input control channels, and the system can be discretized using the equations given above. The decomposition of a single cycle into four periods is arbitrary; the resolution could be easily set much higher, giving a better linearization of the trigonometric functions. It follows also that the case of azimuth hard limits can be considered. For instance, if the azimuth angle has to satisfy $|u_p| \leq \overline{u}_p$, it is reasonable to set $u_p = \overline{u}_p \cos(2\pi t/\Delta t)$, so that a smooth slewing through the range results. In complete analogy with the case of constant azimuth rate above, we see that the surge DOF obeys $m\ddot{X} = u_t \cos[\overline{u}_p \cos(2\pi t/\Delta t)]$, and so on. The resulting functions may benefit from unequal spacing in the decomposition. These options do not involve approximations on the basic scheme of sequential linearization.

Multi-input, multi-output control problems are suitable for the many optimal and robust control approaches available today. We study the discrete-time linear quadratic regulator here, which minimizes the cost function

$$J = \sum_{k=0}^{\infty} J_k = \sum_{k=0}^{\infty} \; x^T(t_k)Qx(t_k) + \qquad (7)$$
$$u^T(t_k)Ru(t_k) +$$
$$2x^T(t_k)Nu(t_k).$$

The solution of the discrete Riccati equation is readily computed, and provides the optimum gain such that $u(t_k) = -Kx(t_k)$.

As in the continuous time case, choices of Q and R affect the tradeoff between speed of response in regulation, and control action. One important point is that the LQR solution with diagonal Q and R allows all of the inputs to vary arbitrarily with respect to each other. This might not be practical or desired in the case of sequential linearization, because thrust reversals take power and are hard on equipment. Fortunately, some measure of fluctuation can be incorporated into the cost function J as follows: first, augment the discrete-time system with an additional state, this a delayed version of the last input, i.e., $x_{p+1}(t_{k+1}) = u_n(t_k)$. Here, p is the number of original states - not including the delay state, and n is the dimension of the control vector. Next, for the case of $n = 4$ and neglecting the indices for convenience, we write:

$$
\begin{aligned}
J \;=\;& [Q_{11}x_1^2 + \cdots + Q_{pp}x_p^2] + a[u_1^2 + \cdots + u_4^2] + \\
& b[(u_1 - u_2)^2 + (u_2 - u_3)^2 + (u_3 - u_4)^2 + \\
& \quad (x_{p+1} - u_1)^2], \text{ or,} \\
\;=\;& x^T Q x + \qquad\qquad\qquad\qquad\qquad (8) \\
& u^T \begin{bmatrix} a+2b & -b & 0 & 0 \\ -b & a+2b & -b & 0 \\ 0 & -b & a+2b & -b \\ 0 & 0 & -b & a+2b \end{bmatrix} u + \\
& 2x^T \begin{bmatrix} 0 & 0 & \cdots \\ \cdots & \cdots & \cdots \\ 0 & 0 & \cdots \\ -b & 0 & \cdots \end{bmatrix} u.
\end{aligned}
$$

It can be confirmed that the middle matrix R is positive definite for any positive value of b, so this arrangement allows one to systematically trade off performance for small differences between successive values of u. The additional delay state exists to complete the cycle; without it, there could be a large discontinuity between the n'th and the first input, at the end of the cycle.

Results are given for an example case of the Odyssey IV vehicle in Figure 4. The top three subplots show the open loop response to be lightly damped in heave and pitch, with a pure drift in the surge and heave directions. The discrete-time controller is designed with $\Delta t = 1.0$, $Q_{\dot{X}\dot{X}} = Q_{\dot{Y}\dot{Y}} = Q_{\dot{\phi}\dot{\phi}} = 10^{-4}$, $Q_{XX} = Q_{YY} = 10^4$, $Q_{\phi\phi} = 10^5$, $a = 10^{-4}$, and $b = 10^{-2}$. Response of the discrete-

time model to this control law, using sequential linearization, shows that although the damping has been increased, it still is lower than might be hoped. This level of performance is related to our selection of the parameter b, which scales the cost of changes in the thrust level. As simulated, the cost of these changes is larger than the cost of "uncorrelated,", i.e., diagonal R, control. The lower subplot confirms that the thrust level changes slowly. With $b = 0$, successive thrust commands are uncorrelated, and the damping performance can be much better.

The upper figures also show the result of a continuous-time simulation with the discrete-time control law implemented with a zero-order-hold. There are slight discrepancies between the continuous and discrete model responses because the continuous model fully accounts for the continuous change in pitch angle through $[0, 2\pi]$ over Δt.

An important implementation point is that thruster jet dynamics will almost certainly play a role in the system performance. The effect is neglected in this work, and should be considered more carefully, although published data from thrusters of a size similar to that we will put on the new Sea Grant vehicle confirm that rise times of one- to two-hundred milliseconds are common (21,22). Surprisingly little attention has been given to speed of response in azimuthing propellers, but Stettler shows that extremely small time scales are suitable, incurring virtually no delay in the force response (23). Where possible, these effects should be modelled.

4. CONCLUSION

By making the assumption of a constant azimuth rate, or of a given cyclic azimuth trajectory, the fundamental nonlinearity in the basic vehicle control problem is eliminated, and the plant obtains a time-varying gain matrix. This admits a clean discretization, recovering normal controllability of the plant, and hence allowing any one of a number of optimal control design techniques. The advantages of the approach include explicit bounds for robustness and performance, as with other linear systems, and the extension to multiple actuators.

REFERENCES

1. SORENSEN, A.J., and STRAND, J.P., 'Positioning of small-waterplane-area marine constructions with roll and pitch damping', Control Eng. Practice, 8(2), 205, 2000.

2. SORENSEN, A.J., SAGATUN, S.I., and FOSSEN, T.I., 'Design of a dynamic positioning system using model-based control', Control Eng. Practice, 4(3), 359, 1996.

3. LIANG, C.C., and CHENG, W.H., 'The optimum control of thruster system for dynamically positioned vessels', Ocean Eng., 31, 97, 2004.

4. JOHANSEN, T.A., FOSSEN, T.I., and BERGE, S.P., 'Constrained nonlinear control allocation with singularity avoidance using sequential quadratic programming', IEEE Trans. Control Systems Technology, 12(1), 211, 2004.

5. WEBSTER, W.C., and SOUSA, J., 'Optimum allocation for multiple thrusters', Proc. Int. Offshore and Polar Eng. Conf., volume 1, 83, 1999.

6. SINDING, P., and ANDERSEN, S.V., 'Force allocation strategy for dynamic positioning', Proc. Int. Offshore and Polar Eng. Conf., volume 1, 346, 1998.

7. SORDALEN, O.J., 'Optimal thrust allocation for marine vessels', Control Eng. Practice, 5(9), 1223, 1997.

8. DESSET, S., CHRYSSOSTOMIDIS, C., DAMUS, R., HOVER, F.S., MORASH, J., and POLIDORO, V., 'Closer to Deep Underwater Science with ODYSSEY IV Class Hovering Autonomous Underwater Vehicle', Proc. Oceans 2005, to appear.

9. RAJKUMAR, V., and MOHLER, R.R., 'Nonlinear control methods for power systems - A comparison', IEEE Trans. on Control Systems Technology, 3(2), 231, 1995.

10. SLOTINE, J.-J.E., and LI, W., Applied nonlinear control, Englewood Cliffs, New Jersey: Prentice-Hall, 1991.

11. MURRAY, R.M., and SASTRY, S.S., 'Nonholonomic motion planning - Steering using sinusoids', IEEE Trans. Automatic Control, 38(5), 700, 1993.

12. POMET, J.B., 'Explicit design of time-varying stabilizing control laws for a class of controllable systems without drift', Systems & Control Letters, 18(2), 147, 1992.

13. PETTERSEN, K.Y., and FOSSEN, T.I., 'Underactuated dynamic positioning of a ship - Experimental results', IEEE Trans. Control System Technology, 8(5), 856, 2000.

14. JUANG, Y.T., KUO, T.S., and HSU, C.F., 'Stability robustness analysis of digital control systems in state-space models', Int. J. Control, 46(5), 1547, 1987.

15. MARTIN, J.M., and HEWER, G.A., 'Smallest destabilizing perturbations for linear systems', Int. J. Control, 45(5), 1495, 1987.

16. SHAKED, U., 'Guaranteed stability margins for the discrete-time linear quadratic optimal regulator', IEEE Trans. Automatic Control, 31(2), 162, 1986.

17. KOLLA, S.R., and FARISON, J.B., 'Reduced-order dynamic compensator design for stability robustness of linear discrete-time systems', IEEE Trans. Automatic Control, 36(9), 1077, 1991.

18. KOLLA, R., and FARISON, J.B., 'Improved robust stability bounds for discrete-time linear regulators with computational delays', Automatica, 26(3), 619, 1990.

19. ISHIHARA, T., 'Robust stability bounds for a class of discrete-time regulators with computation delays', Automatica, 24(5), 697, 1988.

20. ASTROM, K.J., and WITTENMARK, B., Computer controlled systems, Englewood Cliffs, New Jersey: Prentice-Hall, 1984.

21. BACHMAYER, R., WHITCOMB, L.L., and GROSENBAUGH, M.A., 'An accurate four-quadrant nonlinear dynamical model for marine thrusters: theory and experimental validation', IEEE J. Oceanic Eng., 25(1), 146, 2000.

22. HEALEY, A.J., ROCK, S.M., CODY, S., MILES, D., and BROWN, J.P., 'Toward an improved understanding of thruster dynamics for underwater vehicles', IEEE J. Oceanic Eng., 20(4), 354, 1995.

23. STETTLER, J.W., HOVER, F.S., and TRIANTAFYLLOU, M.S., 'Investigating the steady and unsteady maneuvering dynamics of an azimuthing podded propulsor', SNAME Maritime Technology Conference and Transactions, to appear, 2005.
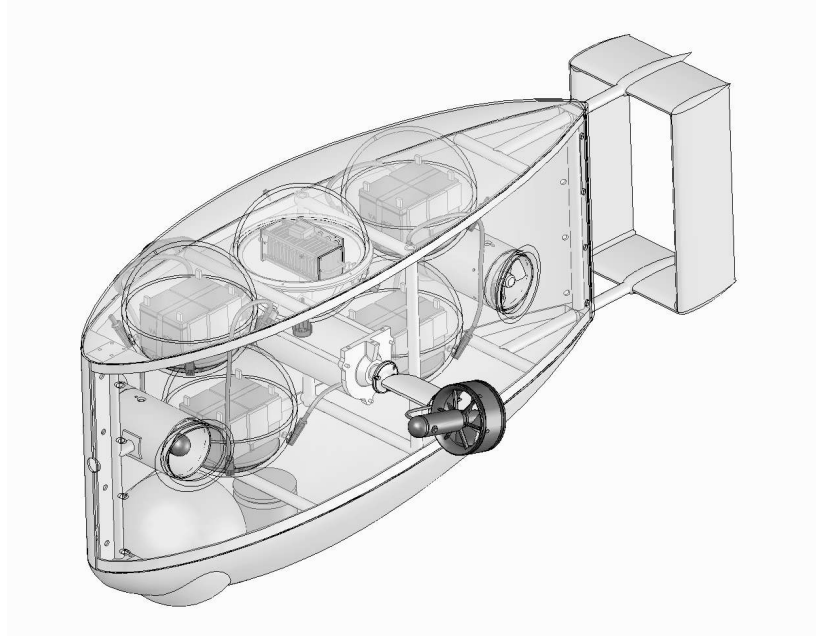
Figure 1: The MIT Sea Grant Odyssey IV vehicle. The vehicle comprises a streamlined main body with aft stabilizing fins. The two surge thrusters are actively controlled in pitch through a drive assembly near the center of the body; hence, the thrusters actuate vehicle surge, heave, and pitch.
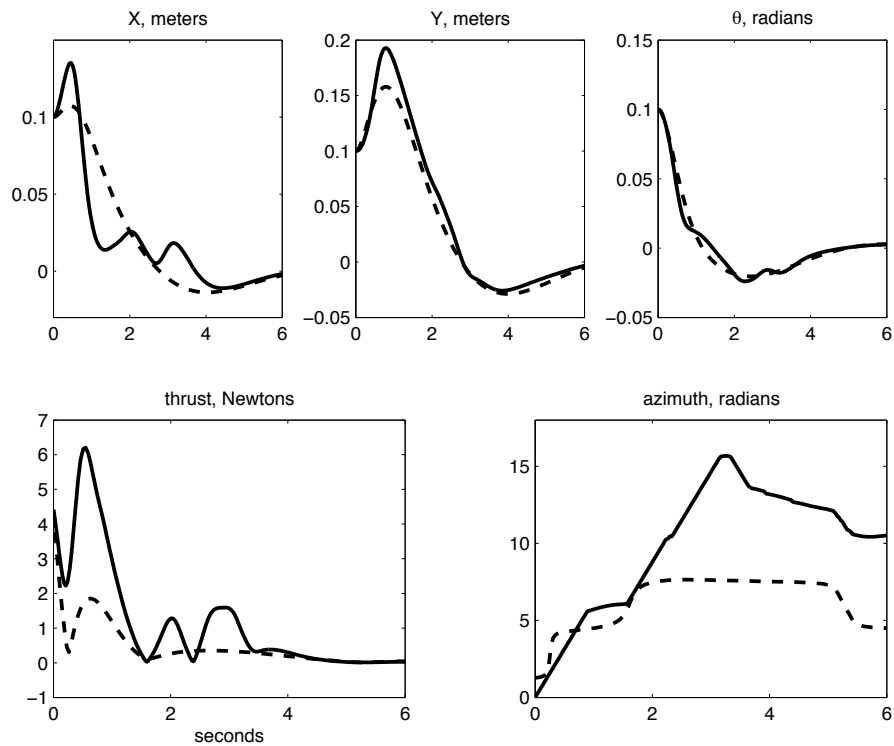

Figure 2: Comparison of the system closed-loop response with a direct LQR controller implemented in continuous time (dashed lines), and with the simplest nonlinear thruster controller (solid lines).
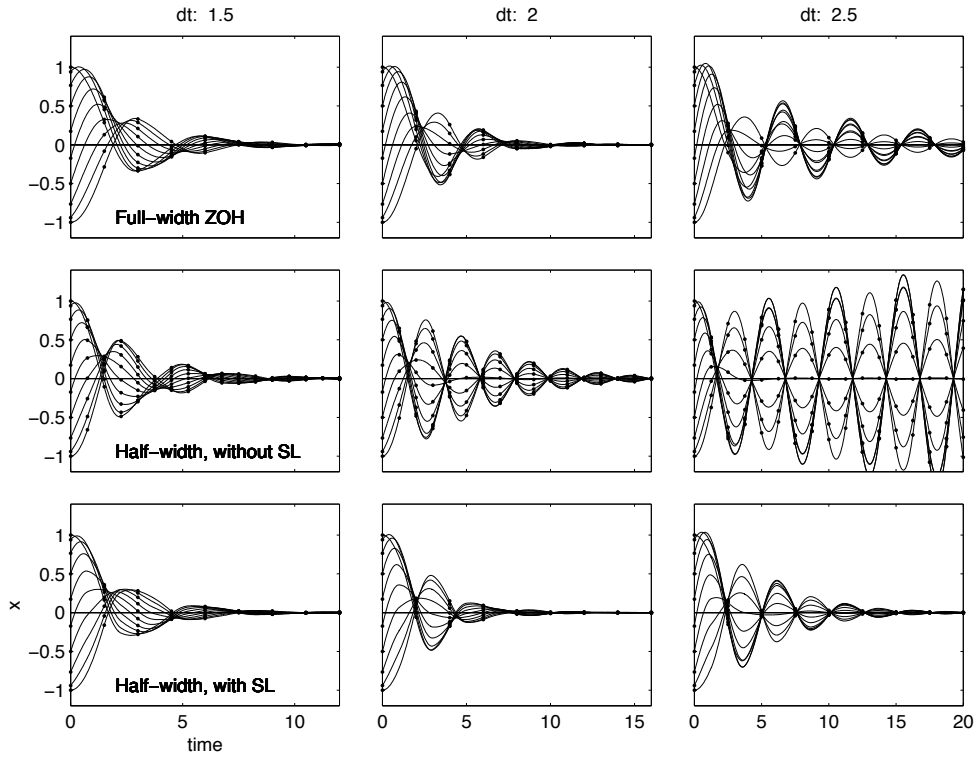
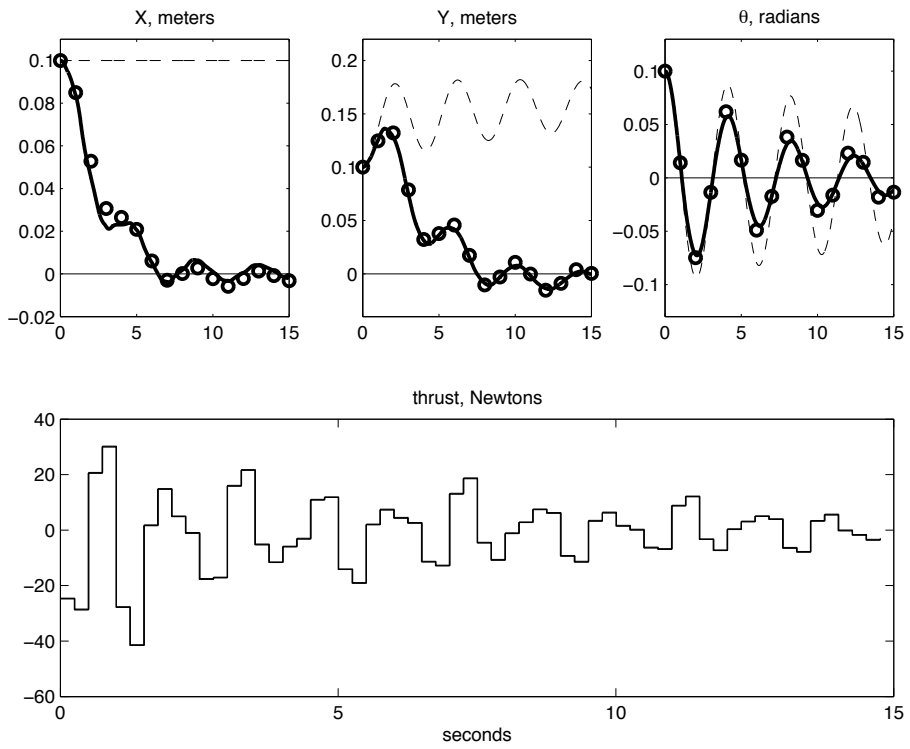Figure 3: Examples of control of the undamped oscillator (natural frequency $1 rad/s$).

Figure 4: Odyssey IV system response with the split-ZOH discrete-time LQR control. Circles show the discrete-time plant, solid line shows the effect of the split-ZOH control action on a continuous model, and the dashed-line is the continuous-time open-loop response.