INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 20 Motion Control, Robotics, and Autonomous Systems

CHAPTER 6

Choosing What Not to Reject in a Dynamic Positioning Simulator

Pages 130–157

Edited by John D. Hedengren

How to cite this chapter. "Choosing What Not to Reject in a Dynamic Positioning Simulator," in *Motion Control, Robotics, and Autonomous Systems*, John D. Hedengren, Ed., vol. 20 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 6, pp. 130–157.

The page numbers in this file are those of the printed volume.

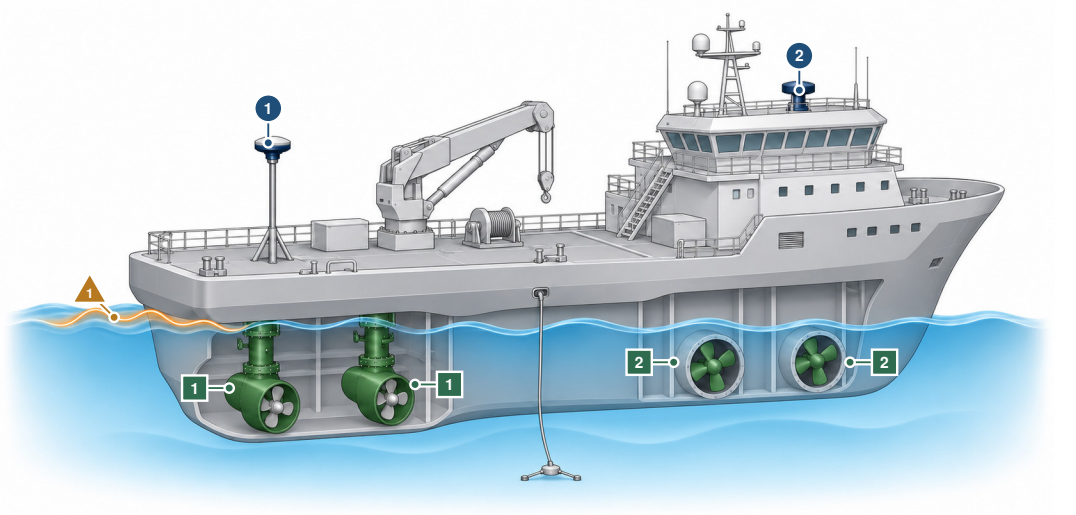
The app for this chapter
https://apmonitor.com/apps/dynamic-positioning/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v20-motion-robotics/#dynamic-positioning

Copyright © 2026 Advanced Process Solutions, LLC. All rights reserved. Produced by Advanced Process Solutions, LLC.

CHAPTER 6

Choosing What Not to Reject in a Dynamic Positioning Simulator



The controller moves 1 two azimuth thrusters aft, which turn to aim their thrust
2 two tunnel thrusters forward, which push sideways

The controller reads 1 position reference 2 heading sensor (gyro)

What disturbs it 1 waves, which rock the hull and push it slowly off station

Notice the line from the hull down to the wellhead the vessel must stay over, and the waves at the waterline: the thrusters should stop the slow drift and leave the rocking with each wave alone. Illustration, not to scale.

IN PLAIN TERMS **Holding a Ship Over One Spot**

The main idea. A work ship at sea often has no anchor down, so its propellers must keep pushing it back over one spot on the seabed. A computer reads where the ship is and tells the propellers how hard to push.

The challenge. The reading also shows the ship rocking with every wave, a motion no propeller can stop and none needs to. A computer that chases the rocking wears out the propellers and holds the ship no better.

Why it matters. Telling the slow drift from the rocking, and turning the bow into the weather, decide how well and how cheaply the ship stays on its spot.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the history of the first deep-water drill ship in the Introduction, the box on the first controllers (Section 6.2), the two facts and the try-it segments of scenarios 3 and 8 (Section 6.6), and the key points
Study	60 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the measured scenarios (Section 6.6), Technology in Practice (Section 6.7), Limitations (Section 6.8) and the open challenge

The control question The position reference of a vessel on thrusters adds the slow drift that the thrusters must stop to the wave motion that they cannot and need not stop. Which part of that measurement should the controllers be shown, and which decisions outside the three PID loops decide how well and how cheaply the vessel holds station?

Learning objectives

- Derive the inertias and the closed-loop frequency and damping of a position loop on a hull with no restoring force.
- Explain why an observer with bias and wave states, and not a low-pass filter, keeps wave motion out of the thruster commands.
- Compare raw, model-only and full-observer estimates on thruster travel, energy and position, and say which number carries the case.
- Compute the sway and heading capability of the thruster layout, and predict whether thrust, power or heading binds first.
- Evaluate a solved allocation, a heading supervisor and a redundant reference against a fixed split, a fixed heading and a single reference.

Prerequisites Sections 1.6.2 (velocity-form PID), 1.4.3 (position loops and integral action), 1.4.6 (bias states and wave filtering) and 1.4.7 (weighted allocation). Chapter 3 shows estimated-force feedforward on a simpler vehicle.

The app https://apmonitor.com/apps/dynamic-positioning/

See also Volume 5, Chapter 13 for a marine vehicle navigated by a Kalman filter on late position fixes, and Volume 5, Chapter 12 for a sensor whose reading does not say which way the error lies. Volume 4, Chapter 13 and Volume 4, Chapter 12 for allocation across actuators with limits and prices.

6.1 Introduction

A feedback loop is built on the premise that the difference between the measurement and the setpoint is an error, and on most plants the premise is safe. A dynamically positioned vessel is where it fails. The position reference on an offshore construction vessel reports two motions superimposed, the hull pushed off the wellhead by wind and current, and the same hull riding over individual waves. In the 4 m sea of this chapter's results the second reaches 2.39 m at an eleven second period while the first is not permitted to reach the 5 m watch circle at all, so the larger signal on the wire is the one that must be left alone. No thruster stops eight thousand tonnes moving with the sea, and a control system that tries destroys the machinery it depends on while holding station no better for the effort.

The plant underneath is as plain as the measurement is difficult. A vessel on thrusters has no anchor, no mooring and no hydrostatic restoring force in surge, sway or yaw, so all three controlled axes are double integrators carrying a little hydrodynamic damping, and the entire station-keeping capability of the vessel is the control system [1, 2]. The dynamics are trivial to write down and the measurement is the hard part, which is the reverse of the arrangement a first control course usually presents.

> ⧖ HISTORY **A Drill Ship Without Anchors (1961 to 1971)**
>
> Drilling into the deep ocean floor from a dynamically positioned drill ship was first done in March and April of 1961, with radar reflectors and acoustic transducers on taut line buoys as the reference. A single taut line, then in use on drill ships in 600 to 1000 feet of water, was tested and found to have too much error in deep water. The research drill ship accepted in 1968 listened instead to an acoustic beacon dropped on the sea floor, and a computer turned the reading into commands for four tunnel thrusters and two main propellers. It was the first use of tunnel thrusters as the restoring force for dynamic positioning, and the risk of losing the drill string was judged real enough to require redundancy wherever possible [3].

The misconception the simulator corrects is that wave filtering is a filter. A low-pass element on a signal on its way into a controller removes the wave-frequency motion and removes the phase margin with it, because a plant with no restoring force needs derivative action and derivative action needs phase. What industrial dynamic positioning does instead is carry a model of the motion it intends to keep: an observer runs a vessel model, a slowly varying bias state that absorbs the wind and the current, and a lightly damped oscillator tuned to the peak wave period, and divides one innovation among the three so that only the low-frequency estimate reaches the controllers [4, 5]. A block diagram cannot show why that matters, because the cost of getting it wrong is not visible in the controlled variable. It appears in the thruster commands, and this application measures it there. Section 1.4.6 names the idea for every machine in this volume, and Chapter 5 meets it again on a foiling boat that must choose which waves to ride over and which to follow.

Three general classes of problem sit behind the vessel. The first is measurement decomposition, in which part of a measured signal is real, periodic and outside the authority of the actuator, so that the correct response is a structured estimator rather than a filter [5, 6]. The second is disturbance estimation with feedforward, in which a large slow unmeasured load is identified inside the observer, where it can be found quickly without entering the loop, so the feedback can stay slow [7, 8]. The third is control allocation over an actuator set with more independent forces than controlled degrees of freedom, solved across aerospace, automotive and marine applications by the same weighted pseudo-inverse [9]. A reader who has met a level loop on a sloshing surface, or a vehicle with individual wheel actuators, has met one of the three already.

A dynamic positioning system is normally taught from its equations, and the equations are the part that transfers least well: the observer is written in three lines that look like a Kalman filter and hide the one decision the designer actually makes, which is how much of the measurement to believe. This chapter prices each structural decision separately, on an interactive vessel where the estimator, the allocation and the heading policy can each be switched to a naive alternative and the difference read off the same scorecard. Section 6.2 traces the problem from anchor spreads to the Kalman filter. Section 6.3 sets down the eighteen-state vessel, environment and thruster model, with the equations, the constants and the initial conditions in the detail another implementation needs to reproduce the trajectories, and Section 6.4 the observer, the three control loops, the allocation and the heading supervisor. Section 6.5 lists the twelve scenarios, which put the four structural decisions ahead of every gain. Section 6.6 puts a number against each choice on one plant. The sharpest is that the same sea and the same tuning cost 327.5 MN of thruster command travel over the twenty-five minutes from 300 to 1800 s with the raw measurement shown to the loops and 72.0 MN with the full observer, while the vessel holds station no better for the extra effort, 0.55 m off on average against 0.51 m. The most valuable of the four decisions turns out not to be in the control system at all. Section 6.7 names the position references, thrusters and computers of real vessels, and Sections 6.8 to 6.10 state what the model cannot represent and what the measurements say together. The vessel model, the dynamic positioning observer and the pseudo-inverse allocation are all standard results.

6.2 From Anchors to a Kalman Filter

Station keeping began as a mooring problem, solved by naval architects with an anchor spread rather than by control engineers, because the restoring force came from catenary chain and not from feedback. How deep-water drilling first held a ship on thrusters is told in the history box of the Introduction, and how its first controllers kept the sea out of the thrusters in the box below.

> ⧖ HISTORY **A Twenty-Second Lag in Front of the Loop (1967)**
>
> An equipment proposal of 1967 for that drill ship simulated the vessel in three degrees of freedom. The wave motion that leaked into the position data was represented by sine waves of plus or minus 70 feet at a 10 second period, roughly a sea state of 5 or greater. Each loop used rate, position and integral feedback, and the position inputs passed through simple 20 second lags to keep unnecessary demands off the thrusters. Its authors noted that these lags, added to the thruster lags, made a rather slow control loop. After three years at sea the operators asked for a program to which the current and the wind could be fed or sensed, so that the computer could anticipate losing station instead of reacting after moving off station [3].

The modern structure arrived between 1976 and 1983, as an estimation result rather than a control result. The Trondheim group posed dynamic positioning as a Kalman filtering and optimal control problem, with a vessel model, a wave-frequency model and a slowly varying bias estimated together and the controller closed on the low-frequency estimate alone [4, 8]. A British line of work reached the same structure from stochastic optimal control, treating wave-frequency motion as a colored disturbance the controller must not attempt to remove [10]. Both replaced the filter in front of the controller with a model inside an estimator, which made the loop fast enough to be useful and quiet enough to be affordable, and a survey three decades later still finds the same three blocks [2]. Later, a passive nonlinear observer removed the need to linearize about fixed headings [11], estimator tuning was scheduled against the sea state rather than fixed [6], and allocation grew from a fixed split written at commissioning into a solved optimization now standard in ships, aircraft and road vehicles alike [9].

The last piece of the history is not a control result. Dynamic positioning failures are dominated by loss of position reference, by thruster and power failures and by human factors rather than by loop tuning, and the industrial answer has been redundancy, independent reference systems and a formal vessel class rather than a better controller [12]. Scenario 10 is exactly that failure, included because no controller in the collection can solve it.

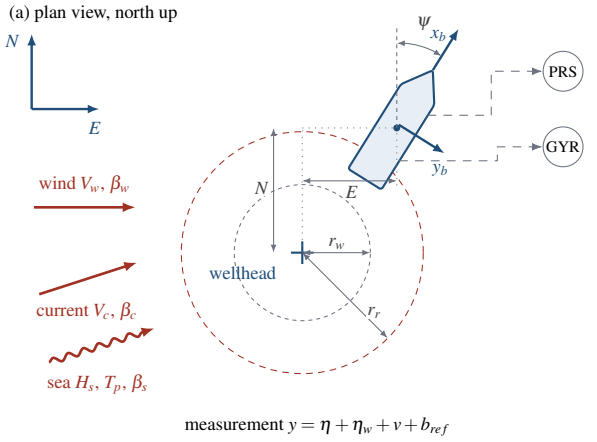
6.3 Vessel, Environment and Thruster Model

The vessel is a 90 m offshore construction ship of 8000 t displacement holding a wellhead. Two 220 kN azimuth thrusters sit 35 m aft of midships and two 120 kN tunnel thrusters sit forward, with 5200 kW of thruster power available between them. Work on deck stops when the vessel leaves a 5 m watch circle and the riser is disconnected at 9 m. Figure 6.1 defines the two reference frames, the thruster geometry and every symbol the equations use.

> ✦ DID YOU KNOW? **Twenty Thousand Feet Down, Forty Feet Off**
>
> By 1971 the acoustic positioning system of the first deep-sea research drill ship had let it drill and core in water 20,000 feet deep, about six kilometers, and 3888 feet into the ocean floor, with an apparent excursion of only plus or minus 40 feet for four or five days at a time [3].

6.3.1 Classification of Variables

Four loops are declared. Three of them pair a position or heading with a generalized force demand, and the fourth reports the present heading against the cheapest one so that the supervisory decision is visible on the same time axis as the loops it displaces. Table 6.1 names every signal.



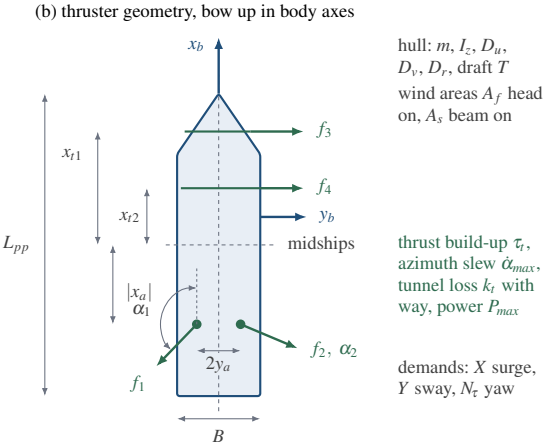


Figure 6.1 Dynamically positioned construction vessel and the geometry the parameters refer to. Panel (a) is the plan view in earth axes: the wellhead at the origin of N and E, the watch circle r_w inside which deck work continues, the disconnect radius r_r, the heading ψ from north, and the wind, current and sea each specified by the direction β they arrive from. The position reference and gyro report the low-frequency motion η, the wave-frequency motion η_w and noise v added together, plus any reference bias b_{ref}. Panel (b) is the same hull in body axes, with two azimuths of thrust f_1, f_2 at angles α_1, α_2, two tunnels of thrust f_3, f_4 on the longest levers on the vessel, and the three demands X, Y and N_τ the allocation resolves onto them.

Table 6.1 Classification of Every Signal. Entries Marked (adv.) Appear Only in the Advanced Control Set.

Symbol	Role	Default	Range
N	CV, north position	0 m	—
E	CV, east position	0 m	—
ψ	CV, heading	0°	—
r	output, radius off wellhead	0 m	—
N_{SP}	setpoint, north	0 m	±30 m
E_{SP}	setpoint, east	0 m	±30 m
ψ_{SP}	setpoint, heading	0°	0 to 359°
X	MV, surge demand	0 kN	±680 kN
Y	MV, sway demand	0 kN	±680 kN
N_τ	MV, yaw demand	0 MN m	±12 MN m
V_w, β_w	DV, wind speed, origin	8, 90°	0 to 30 m/s
V_c, β_c	DV, current speed, origin	0.3, 60°	0 to 2.5 m/s
H_s	DV, significant wave height	1.5 m	0 to 7 m
T_p	DV, peak wave period	9 s	4 to 18 s
β_s	DV, wave origin	95°	0 to 359°
b_{ref}	DV, reference bias (adv.)	0 m	±5 m
—	DV, thruster availability (adv.)	all four	0 or 1 each
N, E, ψ (LF)	states, low-frequency pose	0	—
u, v, r_ψ	states, body velocities	0	—
η_w (6)	states, wave motion and rates	0	—
$f_1..f_4$	states, thruster forces	balanced	±F per unit
α_1, α_2	states, azimuth angles	balanced	±180°
K_c^p	tuning, position gain	37	2 to 200 kN/m
τ_I^p	tuning, position reset	600	20 to 3000 s
τ_D^p	tuning, position rate	42	0 to 200 s
τ_f	tuning, derivative filter	9	0 to 60 s
K_c^h	tuning, heading gain	0.34	0.02 to 4 MN m/°
τ_I^h	tuning, heading reset	700	20 to 3000 s
τ_D^h	tuning, heading rate	34	0 to 200 s
K	tuning, observer gain	0.28	0.01 to 1.5/s
K_b	tuning, bias gain multiple	1.00	0 to 4
$\dot{\psi}_{SP}$	tuning, vane slew rate (adv.)	0.25	0.02 to 2 °/s
—	option, estimate	observer	raw, model, observer
—	option, allocation	minimum norm	minimum norm, fixed
—	option, heading	as set	as set, min. thrust
—	option, feedforward (adv.)	on	on, off

6.3.2 Rigid-Body Motion in Three Degrees of Freedom

The vessel is a rigid body in surge, sway and yaw. Body-axis momentum balances with linear damping and the rotating-frame terms give Equation (6.1).

$$\begin{aligned} m_u \frac{\mathrm{d}u}{\mathrm{d}t} &= X_{thr} + X_{wind} + X_{cur} + X_{drift} - D_u u + m_u v r_\psi \\ m_v \frac{\mathrm{d}v}{\mathrm{d}t} &= Y_{thr} + Y_{wind} + Y_{cur} + Y_{drift} - D_v v - m_v u r_\psi \\ I_z \frac{\mathrm{d}r_\psi}{\mathrm{d}t} &= N_{thr} + N_{wind} + N_{cur} + N_{drift} - D_r r_\psi \end{aligned} \tag{6.1}$$

where u and v are the body surge and sway velocities in m/s, r_ψ is the yaw rate in rad/s, D_u, D_v and D_r are the linear damping coefficients, and m_u, m_v and I_z are the inertias including added mass, given by Equation (6.2).

$$\begin{aligned} m_u &= m(1+a_u), \qquad m_v = m(1+a_v) \\ I_z &= m\,(k_g L_{pp})^2 (1+a_r) \end{aligned} \tag{6.2}$$

where m is the displacement, a_u, a_v and a_r are the added-mass fractions and k_g is the yaw radius of gyration as a fraction of L_{pp}. At the constants of Table 6.2, $m_u = 8.64 \times 10^6$ kg, $m_v = 1.48 \times 10^7$ kg and $I_z = 5.47 \times 10^9$ kg m^2, so the same force accelerates the hull 1.71 times faster along the bow than across it. Nothing on the right of (6.1) depends on position except through the environment, so there is no restoring term anywhere and each axis is an integrating plant driven through a second integration.

The pose in earth axes follows from the rotation through the heading, as shown in Equation (6.3).

$$\frac{\mathrm{d}N}{\mathrm{d}t} = u\cos\psi - v\sin\psi, \quad \frac{\mathrm{d}E}{\mathrm{d}t} = u\sin\psi + v\cos\psi, \quad \frac{\mathrm{d}\psi}{\mathrm{d}t} = r_\psi \tag{6.3}$$

The heading is wrapped to $(-\pi, \pi]$ after every step, as are the two azimuth angles.

6.3.3 Wind, Current and Wave Drift

The three environmental loads are algebraic in the disturbance settings and the heading. Wind is a dynamic pressure acting on two projected areas blended with the angle of attack, in the manner of a first-cut coefficient set [13], as shown in Equation (6.4).

$$\begin{aligned} q_a &= \tfrac{1}{2}\rho_a C_a V_w^2, \qquad \gamma_w = \beta_w - \psi \\ X_{wind} &= -q_a A_f \cos\gamma_w |\cos\gamma_w| \\ Y_{wind} &= -q_a A_s \sin\gamma_w |\sin\gamma_w| \\ N_{wind} &= 0.14\, L_{pp} \cos\gamma_w\, Y_{wind} \end{aligned} \tag{6.4}$$

where ρ_a is the air density, C_a the drag coefficient, A_f and A_s the head-on and beam-on projected areas and γ_w the wind angle relative to the bow, and the absolute values preserve the sign through the quadrant. The yaw moment acts through a lever that is a fixed fraction of the length and changes sign with the angle of attack, which is what makes the hull weathervane.

Current acts the same way on the underwater profile, driven by the water speed relative to the hull, as shown in Equation (6.5).

$$\begin{aligned} u_r &= -V_c\cos(\beta_c - \psi) - u, \;\; v_r = -V_c\sin(\beta_c - \psi) - v \\ X_{cur} &= \tfrac{1}{2}\rho_w C_w B T\, u_r|u_r| \\ Y_{cur} &= \tfrac{1}{2}\rho_w C_w L_{pp} T\, v_r|v_r| \\ N_{cur} &= 0.11\, L_{pp}\, Y_{cur} \end{aligned} \tag{6.5}$$

where ρ_w is the sea water density, C_w the hull drag coefficient, B the beam and T the draft. The relative velocity carries the vessel's own motion, so (6.5) supplies part of the damping as well as the load.

The sea contributes twice, and the two contributions must be kept apart. The steady part is the second-order drift force, which is proportional to the square of the significant wave height and is the part the thrusters genuinely have to carry [14], as shown in Equation (6.6).

$$\begin{aligned} X_{drift} &= -0.35\, k_d H_s^2 \cos(\beta_s - \psi) \\ Y_{drift} &= -k_d H_s^2 \sin(\beta_s - \psi) \\ N_{drift} &= 0.06\, L_{pp}\, Y_{drift} \end{aligned} \tag{6.6}$$

where k_d is the drift force per square meter of significant height in N/m^2. At $H_s = 4$ m this is 83.2 kN of beam force, slow enough that the bias estimator treats it as part of the steady environment.

6.3.4 Wave-Frequency Motion and the Measurement

The oscillatory part of the sea moves the hull without displacing it from station. It is represented by three lightly damped second-order oscillators at the peak wave frequency, driven by independent colored noise so that the seaway is irregular and the estimator cannot lock onto a sine, as shown in Equation (6.7).

$$\begin{aligned} \ddot{\eta}_{w,i} + 2\zeta\omega_0\dot{\eta}_{w,i} + \omega_0^2\eta_{w,i} &= 2\zeta\omega_0^2 g_i\, w_i(t) \\ g_N &= G_w H_s \cos(\beta_s - \psi) \\ g_E &= G_w H_s \sin(\beta_s - \psi) \\ g_\psi &= G_\psi H_s \sin(\beta_s - \psi) \end{aligned} \tag{6.7}$$

for $i \in \{N, E, \psi\}$, where $\omega_0 = 2\pi / \max(3\,\mathrm{s}, T_p)$ is the peak wave frequency, ζ is the motion damping ratio, G_w and G_ψ are the first-order motion gains in meters and degrees per meter of significant height, and w_N, w_E and w_ψ are unit-variance colored noise sequences of correlation time 2.2, 2.4 and 2.0 s. The factor 2ζ on the right cancels the resonator's own amplification at ω_0, so G_w means what its name says rather than that value times a quality factor of three. In the 4 m seaway used throughout the results the peak excursion over the observation window of Section 6.6 is 2.39 m (root mean square 0.76 m across the beam) against a watch circle of 5 m.

What the control system reads is the sum, as shown in Equation (6.8).

$$y = \eta + \eta_w + v + b_{ref} \tag{6.8}$$

where $\eta = (N, E, \psi)$ is the low-frequency pose, η_w is the wave-frequency motion of (6.7), v is instrument noise and b_{ref} is a stale differential correction applied to the north channel. Nothing in (6.8) says which term a given meter belongs to, and that is the whole of the control problem. The instrument noise v has reached the measurement since the repair recorded in Section 6.4.6, and every result in this chapter is obtained with it live at 0.35 m on position and 0.12° on the gyro, each a colored sequence with a 1.5 s correlation time. The position reference is sampled on the controller scan rather than the physics step, which is a zero-order hold of 0.5 s, and the controller acts on each sample at the instant it is taken, so no further dead time is applied.

6.3.5 Thrusters

A commanded force reaches each thruster through a first-order lag, and each azimuth swings toward its commanded angle no faster than its slew limit allows. Both appear in Equation (6.9).

$$\begin{aligned} \tau_t \frac{\mathrm{d}f_k}{\mathrm{d}t} &= f_k^{cmd} - f_k \\ \frac{\mathrm{d}\alpha_j}{\mathrm{d}t} &= \mathrm{clamp}\left(\frac{\alpha_j^{cmd} - \alpha_j}{0.5\,\mathrm{s}},\ \pm\dot{\alpha}_{max} \right) \end{aligned} \tag{6.9}$$

for $k = 1 \ldots 4$ and $j = 1, 2$, where the angular error is wrapped to $(-\pi, \pi]$ before the clamp. The forces are clipped to each machine's rating after every step. What the four thrusters produce in the body frame is Equation (6.10).

$$\begin{aligned} \ell &= \mathrm{clamp}(1 - k_t |u|,\ 0.25,\ 1) \\ X_{thr} &= f_1 \cos\alpha_1 + f_2 \cos\alpha_2 \\ Y_{thr} &= f_1 \sin\alpha_1 + f_2 \sin\alpha_2 + (f_3 + f_4)\,\ell \\ N_{thr} &= x_a\,(f_1 \sin\alpha_1 + f_2 \sin\alpha_2) \\ &\quad - y_a\,(f_1 \cos\alpha_1 - f_2 \cos\alpha_2) + (x_{t1} f_3 + x_{t2} f_4)\,\ell \end{aligned} \tag{6.10}$$

where ℓ is the tunnel efficiency, k_t is the tunnel loss per unit of forward speed, x_a is the signed azimuth position relative to midships and y_a is the azimuth offset from the centerline. At half a knot the tunnels retain 85.9 % of their rating and at two knots 43.4 %, so a tunnel thruster is worth appreciably less on a vessel that is moving than on one that is stopped. Shaft power follows the usual three-halves law of Equation (6.11).

$$P = k_P \sum_{k=1}^{4} |f_k|^{3/2} \tag{6.11}$$

where k_P is in kW per $\mathrm{N}^{1.5}$. All four machines at their ratings draw 5182 kW against 5200 kW installed, so the switchboard and the thrust run out at almost the same place. Equation (6.11) also prices the allocation directly, because 200 kN of sway spread over four thrusters costs 801 kW and the same 200 kN through two costs 1132 kW.

6.3.6 Parameters, Initialization and Numerics

Table 6.2 lists every constant. Four are hidden from the parameter panel (the two densities, the power coefficient and the allocation ridge) and the rest are adjustable, a deliberately thin concealment because what the student is asked to infer is structure and not a hidden number. The heading tolerance is read only by the heading badge above the animation, which turns amber while the heading is further than the tolerance from its setpoint.

Table 6.2 Model Parameters. Four Constants Are Hidden From the Parameter Panel and Used by the Model Only. Every Other Row Is Adjustable Over the Range Named in the Application.

Symbol	Meaning	Value	Unit	Visibility
Vessel				
L_{pp}	length between perpendiculars	90	m	adjustable
B	beam	20	m	adjustable
m	displacement	8.0×10^6	kg	adjustable
a_u	added mass, surge	0.08	—	adjustable
a_v	added mass, sway	0.85	—	adjustable
k_g	yaw radius of gyration	0.25	$\times L_{pp}$	adjustable
a_r	added inertia, yaw	0.35	—	adjustable
D_u	surge damping	9.0×10^4	N s/m	adjustable
D_v	sway damping	2.6×10^5	N s/m	adjustable
D_r	yaw damping	1.2×10^8	N m s	adjustable
Environment				
ρ_a	air density	1.225	kg/m^3	hidden
ρ_w	sea water density	1025	kg/m^3	hidden
A_f	wind area, head on	320	m^2	adjustable
A_s	wind area, beam on	900	m^2	adjustable
T	draft	6.0	m	adjustable
C_a	wind drag coefficient	0.85	—	adjustable
C_w	current drag coefficient	1.05	—	adjustable
k_d	wave drift per m^2 of H_s^2	5200	N/m^2	adjustable
G_w	wave motion gain	0.42	m/m	adjustable
G_ψ	wave yaw gain	0.22	°/m	adjustable
ζ	wave motion damping	0.16	—	adjustable
Thrusters				
F_a	azimuth rating, two aft	220	kN	adjustable
F_t	tunnel rating, two forward	120	kN	adjustable
x_a	azimuth position from midships	−35	m	adjustable
y_a	azimuth offset from centerline	6.5	m	adjustable
x_{t1}	forward tunnel position	38	m	adjustable
x_{t2}	second tunnel position	26	m	adjustable
τ_t	thrust build-up time constant	4.0	s	adjustable
$\dot{\alpha}_{max}$	azimuth turning rate	6.0	°/s	adjustable
k_P	power per unit thrust$^{3/2}$	1.79×10^{-5}	kW/N$^{3/2}$	hidden
P_{max}	thruster power available	5200	kW	adjustable
k_t	tunnel loss with way	0.55	s/m	adjustable
τ_α	azimuth angle filter	20	s	adjustable
ε	allocation ridge	0.025	—	hidden
Instrument and limits				
σ_p	position reference noise	0.35	m	adjustable
σ_g	gyro noise	0.12	°	adjustable
r_w	watch circle	5.0	m	adjustable
r_r	disconnect radius	9.0	m	adjustable
—	heading tolerance	3.0	°	adjustable, badge
Numerics				
Δt	integration step (RK4)	0.05	s	—
Δt_c	controller period	0.5	s	—
—	historian rate	2	Hz	—
—	random seed	20260916	—	—

Every scenario opens with the vessel on station and the thrusters carrying the environment: the pose starts at the setpoint, the velocities and wave states at zero, and the thruster forces and azimuth angles at the allocation solution for the steady demand at the initial heading. The control system starts from the same place. The bias estimate opens holding the environmental force at the initial heading in earth axes, and with the feedforward off the three integrals open holding the corresponding demand instead, so the first scan commands the thrust the vessel is already producing. In the first 200 s every automatic scenario except the capability limit and the challenge as loaded stays within 0.95 m of the

wellhead in its low-frequency position. An earlier version of the app started the other way, with the bias estimate and the integrals at zero, and a drive-off carried scenario 8 past the disconnect radius at 91 s. A fixed 0.05 s step of classical fourth-order Runge-Kutta carries the states forward, chosen against the wave period rather than the vessel, which puts two hundred steps in one wave at $T_p = 10\,\text{s}$. The control system executes on a 0.5 s period, the scan rate of an industrial dynamic positioning system, and the historian logs two samples per simulated second. Colored noise and random draws come from one seeded generator, so a run replays identically.

6.3.7 Assumptions

Detail is spent wherever the physics forces a decision on the controller and saved everywhere else. Each simplification below fails somewhere, and the place it fails is named with it.

Three degrees of freedom. Roll, pitch and heave are absent. On a construction vessel it is usually heave rather than horizontal position that stops a crane or a diving operation, and nothing in this application can represent that.

Blended area coefficients for wind and current. Equations (6.4) and (6.5) use two projected areas and a cosine blend with a fixed yaw lever. Measured coefficient sets come from a wind tunnel and are not smooth functions of heading [13], which matters precisely where this application makes its strongest claim, in choosing the cheapest heading.

A single-frequency sea. Equation (6.7) places all of the wave-frequency energy at the peak period, so a resonator tuned there leaves the tails of a real broad-banded spectrum in the position estimate, and the wave direction is applied in earth axes rather than resolved through the hull.

Drift proportional to H_s^2 with a fixed directional shape. The second-order force on a real hull depends on the period as well as the height and can change sign [14], so (6.6) is a magnitude estimate rather than a hull calculation.

Idealized machinery. Equations (6.9) and (6.10) carry a first-order build-up, an azimuth slew rate and the tunnel loss with forward speed. There are no forbidden azimuth zones, no thruster-hull or thruster-thruster interaction and no power management system shedding load before the switchboard trips. When the allocation asks for more than P_{max} the whole solution is scaled back, whereas a real power management system would decide which degree of freedom to surrender. The azimuth angles are filtered ahead of the solve rather than rate-constrained inside it, so the allocation is a two-stage approximation to the constrained problem rather than the constrained problem itself.

One position reference. A vessel of dynamic positioning class 2 or 3 carries three independent references and votes on them [12]. The application has one, which is what makes scenario 10 possible and its remedy impossible to demonstrate.

Repeatability that no control system has. Freezing the integration step and the scan period repeats one trajectory exactly wherever it is run, whereas a real dynamic positioning computer has jitter, dropped telegrams and a scan order that moves.

6.4 Estimator, Controllers and Allocation

The control system is an observer, three PID loops, an allocation and a supervisor above them, arranged as in Figure 6.2. All four execute on the same 0.5 s scan.

6.4.1 The Observer

The observer is a Kalman filter with the gains placed by hand, which is what most industrial dynamic positioning observers are [2]. It joins the bias state of Equation (1.10) and the wave filtering of Section 1.4.6 in one estimator, and Volume 5, Chapter 13 applies the same machinery to an underwater vehicle whose acoustic position fixes arrive late and occasionally wild. It carries a vessel model driven by the thrust actually produced, a bias state for the steady environment and a

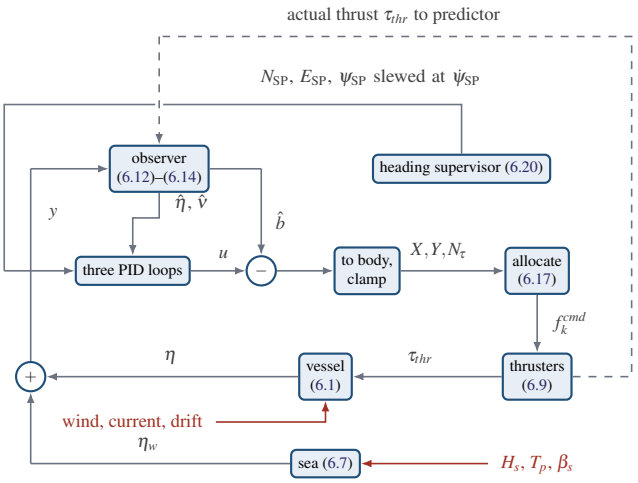


Figure 6.2 Control structure as implemented. One innovation from the position reference y is divided by the observer among the low-frequency pose $\hat{\eta}$, its velocity, the environmental bias $\hat{b}$ and the wave-frequency states, and only $\hat{\eta}$ reaches the three PID loops. The bias estimate is subtracted from the loop output as a feedforward (on every estimator setting, while its own switch is on), the result is rotated into body axes and clamped there, and the allocation resolves the three demands onto six thruster forces. The heading supervisor sits outside every loop.

wave-frequency resonator, is advanced by explicit Euler on the 0.5 s scan, and predicts by (6.12).

$$\begin{aligned} F_b &= (\hat{\tau}_N \cos\hat{\psi} + \hat{\tau}_E \sin\hat{\psi}) - D_u \hat{u}, \qquad \hat{a}_b = F_b/m_u \\ F_s &= (-\hat{\tau}_N \sin\hat{\psi} + \hat{\tau}_E \cos\hat{\psi}) - D_v \hat{w}, \qquad \hat{a}_s = F_s/m_v \\ \dot{\hat{V}}_N &= \hat{a}_b \cos\hat{\psi} - \hat{a}_s \sin\hat{\psi}, \qquad \dot{\hat{V}}_E = \hat{a}_b \sin\hat{\psi} + \hat{a}_s \cos\hat{\psi} \\ \dot{\hat{V}}_\psi &= \left(N_{thr} + \hat{b}_\psi - D_r \hat{V}_\psi\right)/I_z \end{aligned} \tag{6.12}$$

where $\hat{\tau}_N = X_{thr}\cos\hat{\psi} - Y_{thr}\sin\hat{\psi} + \hat{b}_N$ and $\hat{\tau}_E = X_{thr}\sin\hat{\psi} + Y_{thr}\cos\hat{\psi} + \hat{b}_E$ are the thrust and bias in earth axes, $\hat{u}$ and $\hat{w}$ are the estimated body surge and sway velocities obtained by rotating $\hat{V}_N$ and $\hat{V}_E$ into the body, and $\hat{V}$ denotes an estimated earth-frame velocity. The rotation in (6.12) is not cosmetic: the inertia is a body-frame matrix whose two entries differ by a factor of 1.71, so an earth-frame force divided directly by a body-frame mass is exact only with the bow due north. Hull damping belongs in the observer's model for the same reason, because without it every newton of drag has to be absorbed by the bias state, which then hunts.

The wave states are propagated as an undriven resonator at the same peak frequency the plant uses, as shown in Equation (6.13).

$$\ddot{\hat{\eta}}_w = -2\zeta\omega_0\dot{\hat{\eta}}_w - \omega_0^2\hat{\eta}_w \tag{6.13}$$

on each of the three axes. It has no input, because the estimator does not know the seaway, only the frequency at which part of the innovation belongs to it.

The correction step takes one innovation per axis and divides it four ways, as shown in Equa-

tion (6.14).

$$
\begin{aligned}
&e = y - (\hat{\eta} + \hat{\eta}_w), \qquad L_2 = \frac{K^2}{3}, \qquad L_3 = \frac{K^3}{180} K_b \\
&\hat{\eta} \leftarrow \hat{\eta} + K\, e\, \Delta t_c, \qquad \hat{V} \leftarrow \hat{V} + L_2\, e\, \Delta t_c \\
&\dot{\hat{\eta}}_w \leftarrow \dot{\hat{\eta}}_w + 2.2\, K\, \omega_0\, e\, \Delta t_c \\
&\hat{b}_b \leftarrow \hat{b}_b + L_3\, m_u\, e_b\, \Delta t_c, \qquad \hat{b}_s \leftarrow \hat{b}_s + L_3\, m_v\, e_s\, \Delta t_c
\end{aligned}
\tag{6.14}
$$

where K is the single observer gain in 1/s, K_b is a multiplier on the nominal bias gain, and e_b and e_s are the north and east innovations rotated into body axes so that the two inertias scale the correct component. The two body-frame bias increments are rotated back into earth axes before they are added to $\hat{b}_N$ and $\hat{b}_E$, and the yaw bias is scaled by I_z in the same way. Three observer poles placed together at $\omega = K/3$ give $L_1 = K$, $L_2 = K^2/3$ and a third of $K^3/27$. The implementation uses $K^3/180$ instead, placing the bias pole 6.7 times slower, because a bias estimate that moves as fast as the position loop is a second controller working on the same error. The heading innovation is wrapped to $(-\pi, \pi]$ and the bias states are clamped at 1.2 MN and 60 MN m.

The estimator option selects what the controllers are given rather than whether the observer runs. The full setting passes $\hat{\eta}$ from (6.14). The model-only setting holds the wave states at zero, so the sea lands partly in the position estimate and partly in the bias. The raw setting overwrites $\hat{\eta}$ with the measurement y at the end of every scan. The observer runs underneath all three, so the bias estimate subtracted in (6.16) is available on every setting and the feedforward follows its own switch. Switching to the raw reference therefore changes what position the loops are shown and nothing else.

6.4.2 The Three Loops as Implemented

Each loop is the velocity-form PID of Section 1.6.2, written for this application as Equation (6.15). The output before limiting, which Chapter 1 writes v_k, is c_k here because v is the sway velocity, and the output bias u_b is zero on all three loops.

$$
\begin{aligned}
&e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c e_k \\
&\mathrm{PV}^f_k = \alpha \mathrm{PV}^f_{k-1} + (1-\alpha)\mathrm{PV}_k, \qquad \alpha = \mathrm{e}^{-\Delta t_c/\tau_f} \\
&D_k = -K_c \tau_D\, a\, \frac{\mathrm{PV}^f_k - \mathrm{PV}^f_{k-1}}{\Delta t_c} \\
&S_k = S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \\
&c_k = S_k + I_{k-1}, \qquad u_k = \mathrm{clamp}(c_k, u_{min}, u_{max}) \\
&I_k = I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - c_k}{T_t} \right] \Delta t_c
\end{aligned}
\tag{6.15}
$$

where $a = +1$ on all three loops, so the controller is reverse acting in the convention of Section 1.3 and a vessel north of its setpoint receives a southward demand. The derivative acts on the measurement after a first-order filter, so moving a setpoint does not kick the thrusters, and the filter uses the exact discrete update for a first-order lag, stable at any scan period. The integral is advanced after the clamp by back-calculation with the default tracking time $T_t = \tau_I$. There is no rate limit on the demand, because the rate limits in this plant are the thruster build-up τ_t and the azimuth slew $\dot{\alpha}_{max}$ and belong to the machinery. The back-calculation sees only the loop's own clamp, so while the body-frame clamp of (6.16), a thruster rating or the power scaling of Section 6.4.3 binds downstream, the integral keeps accumulating. In manual the heading loop is tracked on the same wrapped-error measurement it uses in automatic, which makes transfer to automatic bumpless on all three loops. When a run opens in automatic with the feedforward off, I starts at the steady demand of Section 6.3.6.

The north and east loops act on $\hat{N}$ and $\hat{E}$ from the observer with $u_{min}, u_{max} = \mp 680\,\mathrm{kN}$, which is the vessel's own sway capability rather than a round number, because a controller allowed to ask for

more than the machinery can produce is a controller whose anti-windup never engages. The heading loop acts on the wrapped heading error with $\mp 12\,\mathrm{MN\,m}$, and because the derivative in (6.15) is taken on the measurement, that wrapped error is supplied as PV_k against a setpoint of zero rather than the other way round, which is what makes τ_D^h effective at all. The three outputs are earth-frame forces, and the environment feedforward is applied there before the rotation, as shown in Equation (6.16).

$$\begin{aligned} X &= \mathrm{clamp}\big[(u_N - \hat{b}_N)\cos\psi + (u_E - \hat{b}_E)\sin\psi,\ \pm X_{lim}\big] \\ Y &= \mathrm{clamp}\big[-(u_N - \hat{b}_N)\sin\psi + (u_E - \hat{b}_E)\cos\psi,\ \pm X_{lim}\big] \\ N_\tau &= \mathrm{clamp}\big[u_\psi - \hat{b}_\psi,\ \pm N_{lim}\big] \end{aligned} \tag{6.16}$$

where $X_{lim} = 680\,\mathrm{kN}$ and $N_{lim} = 12\,\mathrm{MN\,m}$. The clamp is applied after the rotation and not before, because clamping the earth-frame pair separately would let 600 kN north and 600 kN east arrive as 849 kN of sway. Subtracting $\hat{b}$ in (6.16) is what allows τ_I^p to be as long as 600 s. The bias state finds the same force inside the estimator, where it is not in the control loop and can therefore be fast without destabilizing anything.

6.4.3 Thrust Allocation

Four machines produce six independent forces, an x and a y from each azimuth and a y from each tunnel, against three demands, so there is a three-dimensional family of solutions and the application offers two ways to pick one. The minimum-weighted-norm solution is the weighted least-norm allocation of Section 1.4.7, written for this vessel as (6.17) [9].

$$\mathbf{u} = W^{-1}B^T\left(BW^{-1}B^T\right)^{-1}\tau, \qquad \tau = [X,\ Y,\ N_\tau]^T \tag{6.17}$$

where $\mathbf{u} = [f_{x1}, f_{y1}, f_{x2}, f_{y2}, f_3, f_4]^T$, the configuration matrix is Equation (6.18),

$$B = \begin{bmatrix} 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & \ell & \ell \\ -y_a & x_a & y_a & x_a & x_{t1}\ell & x_{t2}\ell \end{bmatrix} \tag{6.18}$$

and $W = \mathrm{diag}(F_a^{-2}, F_a^{-2}, F_a^{-2}, F_a^{-2}, F_t^{-2}, F_t^{-2})$ divides each machine by the square of its own rating, which comes close to minimizing power, because the three-halves law of (6.11) makes 200 kN of sway spread over four machines cost 801 kW against 1132 kW through two. The 3×3 system is solved by Gaussian elimination with partial pivoting on every scan, and a failed thruster's entries of W are divided by an availability factor of 10^{-4} and its command is then set to zero, so the same arithmetic spreads its work across whatever remains with no rule written for the case. A constrained version of the same least-squares problem shares the push among the planted feet of the walking robot in Volume 4, Chapter 13, where a foot can push but never pull.

Three details of the azimuth solution are the difference between a working allocation and a limit cycle. An azimuth can produce any force by pointing at it or by pointing away and running astern, and always taking the first turns a demand that changes sign into a 180° slew during which the unit pushes everywhere except where it was asked. The implementation picks whichever representation is nearer the present angle, which is why f_1 and f_2 are signed. The commanded angle is then filtered and the forces re-solved with the angles held, for the reasons and by the arithmetic set out in Section 6.4.6. Finally, if (6.11) exceeds P_{max} the whole vector is scaled by $(P_{max}/P)^{2/3}$ rather than letting one machine absorb the shortfall.

The alternative is a fixed split of the kind written at commissioning. The azimuths take all of the surge and half of the sway, and the tunnels take the remaining sway and the whole yaw moment, as shown in Equation (6.19).

$$f_x = \frac{X}{2}, \quad f_y = \frac{Y}{4}, \quad f_3 = f_4 = \frac{1}{2\ell}\left(\frac{Y}{2} + \frac{N_\tau}{x_{t1}}\right) \tag{6.19}$$

with $f_1 = f_2 = |(f_x, f_y)|$ clipped at F_a and both azimuths pointed along $\operatorname{atan2}(f_y, f_x)$. Equation (6.19) is indistinguishable from (6.17) while all four machines are running. The share belonging to a machine that has stopped is simply not produced.

6.4.4 The Heading Supervisor

The heading is the one setpoint on the panel that the job does not fix, because a wellhead does not care which way the bow points, and both the wind area and the underwater area of (6.4) and (6.5) depend on it. That makes the environmental demand an objective rather than a load, as shown in Equation (6.20).

$$\begin{aligned} \boldsymbol{\tau}_{env}(\psi) &= -\left[\boldsymbol{\tau}_{wind}(\psi) + \boldsymbol{\tau}_{cur}(\psi) + \boldsymbol{\tau}_{drift}(\psi)\right] \\ \psi^{\star} &= \arg\min_{\psi} \| \boldsymbol{\tau}_{env}(\psi) \| \end{aligned} \tag{6.20}$$

where the norm is the magnitude of the surge and sway components. The minimum is scanned at 2° intervals every 30 s rather than solved, because the objective is not convex once the wind and the current disagree and the scan is cheap on a quantity that changes slowly. Figure 6.3 plots the objective for the conditions of scenario 8.

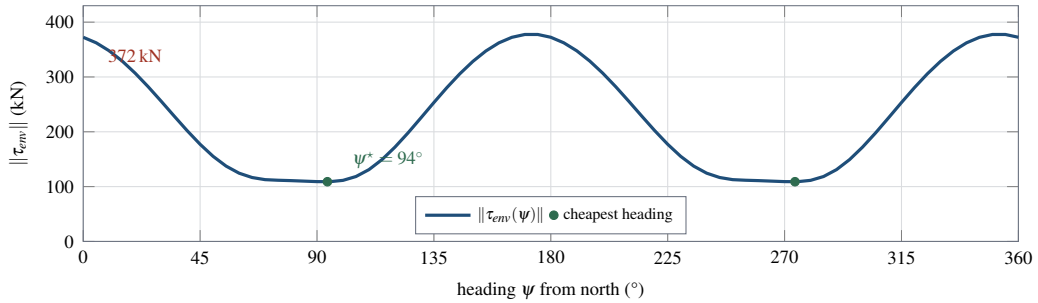


Figure 6.3 Environmental thrust demand against heading for the conditions of scenario 8, evaluated from Equation (6.20) at $V_w = 17\,\mathrm{m/s}$ from 90°, $V_c = 0.8\,\mathrm{m/s}$ from 75° and $H_s = 3.5\,\mathrm{m}$ from 95°. The demand falls from 372 kN beam on to 109 kN at $\psi^{\star} = 94°$, and the power of the exact minimum-norm solution of (6.17) from 2135 kW to 463 kW. The curve repeats at 180° because the projected-area model is symmetric fore and aft.

The supervisor does not command the heading setpoint to $\psi^{\star}$. It walks it there at $\dot{\psi}_{\mathrm{SP}} = 0.25\,^{\circ}/\mathrm{s}$, so the 94° turn of scenario 8 takes 375.5 s. Slewing a setpoint instead of stepping it is the general remedy whenever the transient of a supervisory move is itself the hazard, but on this vessel the slew is not what protects the position. At 2 °/s the setpoint arrives in 46.6 s and the worst excursion during the turn is 6.67 m against 6.46 m at the default rate, only 0.2 m more, because the yaw authority of the heading loop, not the setpoint, sets how fast the bow comes round.

6.4.5 Three Models Kept Apart

The simulation model is Equations (6.1) to (6.11) integrated over eighteen states. The controller's internal model is what appears in (6.12) to (6.14), which is the vessel inertias and damping, the peak wave frequency, and nothing else. It does not know the wind areas, the drag coefficients, the drift constant or the disturbance settings, and it never predicts forward. The heading supervisor is the one block that does hold the environmental model of (6.4), (6.5) and (6.6), which is realistic because a capability calculation on a real vessel is made from a measured coefficient set and a wind sensor, and it is worth noticing that the one place the application uses a plant model is outside every feedback loop.

6.4.6 What the Repaired Measurement Path Changed

Four implementation defects in an earlier version of the application, all repaired in the version this chapter describes, show what each part of the measurement path and the allocation is for, and each was large enough in effect to state.

The position reference noise σ_p and the gyro noise σ_g of Table 6.2 were declared as model parameters and read from the disturbance set, so both resolved to an undefined quantity and no noise reached (6.8) at any setting. Reading them from the parameter set makes both live, and two arguments move with them: the case for the derivative filter τ_f no longer rests on the wave-frequency content of the measurement alone, because instrument noise through $K_c^p \tau_D^p$ by itself drives the surge demand to its clamp in flat calm, and the case for the wave filter is stronger rather than weaker, because the same estimator rejects both motions and the travel ratio rises from 3.3 on a quiet reference to 4.5 on the realistic one. Second, the heading loop was handed the wrapped heading error as its setpoint and a constant zero as its measurement, so the derivative in (6.15) was taken on a signal that never moved and τ_D^h had no effect at any value. Passing the error as the measurement instead makes the loop genuinely three-term, and a 20° heading step in the calm water of scenario 2, made at 300 s, now peaks at 24.1° against 35.7° with the term removed. The manual tracking call kept the old arrangement for a time, which kicked the yaw demand by 1.25 MN m per degree of heading error on transfer to automatic, and now uses the same one. Third, a declared and unread reference lag was deleted, leaving the 0.5 s zero-order hold of the controller scan as the only reference delay the model claims.

The fourth defect was in the allocation and was hidden by the first. Equation (6.17) was solved afresh every scan with the azimuth angles read straight off atan2 of the result, and the direction of a small vector is badly conditioned. The behavior is reproduced here by the current model with the angle filter and the allocation ridge described below both set to zero, a counterfactual run because the ridge is hidden. In the calm water of scenario 2 the port azimuth demand then has a median magnitude of 36 kN, much of it noise through the derivative term, so the demanded direction moves a median of 3.2° per scan and 51.5 % of scans ask for more than the 3° a 6 °/s unit can deliver in half a second. The unit never arrives: over the observation window the port azimuth turns through 7115°, nearly twenty full revolutions, with the rate limiter saturated 66 % of the time, and the heading loop inherits a limit cycle from it, so that a 5 m north step made at 300 s peaks anywhere between 5.95 m and 7.77 m across five seeds, with the heading swinging as far as 10.9°. Setting only the adjustable angle filter to zero in the application gives the same picture, 7048° with the limiter saturated 65 % of the time. A magnitude guard cannot address it, because the magnitude clears any threshold long before the direction stops being noise.

The repair is the one an industrial allocation makes. The commanded azimuth angle is passed through a first-order lag whose time constant is scheduled inversely with the demand magnitude against a reference of 11 kN, so a direction that is mostly noise is filtered hard and one carrying real load is barely filtered at all. The schedule matters: a fixed time constant also delays the redistribution after a thruster trip and costs scenario 9 its point. The forces are then re-solved from the same weighted pseudo-inverse with both angles held, four unknowns against the same three demands, so the tunnels and the azimuth magnitudes make up what the lagging angles cannot. That solve carries a ridge ε of 2.5 % of each axis's free authority, because beam on both azimuths sit near 90°, the surge row collapses, and an exact solve would make a few kilonewtons of surge out of two enormous opposing forces. With the filter at zero the angles are those of the free solution and the forces fall short of it only by the ridge, which beam on in the conditions of scenario 8 delivers 358 kN of the 372 kN of sway asked and leaves the rest to the loops. With the repair as shipped, azimuth travel in calm water falls to 2636° and saturation to 3.2 %, the 5 m step peaks between 5.89 m and 6.70 m across the same five seeds with the heading inside 1.7°, and station keeping improves everywhere. The thruster command travel reported below rises, from 42.8 MN to 72.0 MN in the seaway of scenario 4 over the observation window while the mean offset falls from 0.64 to 0.51 m, because churn once absorbed by

Table 6.3 The Twelve Scenarios and What Each One Isolates

#	Scenario	What it isolates	Kind
1	The environment by hand	a balance is not control	Demonstration
2	Calm water, three loops	derivative action on an undamped plant	Experiment
3	The sea arrives, no filter	chasing a motion that is not an error	Demonstration
4	The same sea, filtered	what the wave model is worth	Experiment
5	Where the waves come back	the observer gain as the estimation trade	Experiment
6	The bias estimate against the integral	disturbance feedforward	Experiment
7	A squall	an unmeasured load found by being pushed	Experiment
8	The heading is free	supervisory choice of a free variable	Experiment
9	A thruster trips	solved allocation against a written one	Experiment
10	The position reference lies	a fault no controller can answer	Demonstration
11	The capability limit	heading capability, not thrust, binds	Demonstration
12	Challenge: a working shift	all four structural decisions at once	Challenge

unphysical spinning now appears where the metric can see it. Three further behaviors of the earlier version have since been repaired in the application and are described in this chapter as they now act: the start-up of Section 6.3.6, the feedforward the raw setting used to withhold (Section 6.4.1) and the heading-loop transfer kick.

6.5 Scenarios

Twelve set-ups are arranged so that the plant is met before the control system and each structural decision is broken alone. Table 6.3 lists them with the mechanism each one isolates. Each carries a briefing and a completion rule in the manner of Section 1.6.4, the challenge is judged on the case as given, so a changed model parameter or disturbance voids the attempt, and the Exercises at the end of the chapter use the scenarios by number.

The reference tuning is $K_c^p = 37\,\mathrm{kN/m}$, $\tau_I^p = 600\,\mathrm{s}$, $\tau_D^p = 42\,\mathrm{s}$, $\tau_f = 9\,\mathrm{s}$, $K_c^h = 0.34\,\mathrm{MN\,m/^\circ}$, $\tau_I^h = 700\,\mathrm{s}$, $\tau_D^h = 34\,\mathrm{s}$, $K = 0.28/\mathrm{s}$ and $K_b = 1.00$. The weak set that opens the manual and challenge scenarios is roughly a quarter of the position gain, a fifth of the derivative time, no derivative filter and an observer gain of 0.9/s.

Scenario 1 is worth running once. The manual helper computes the steady environmental demand of (6.20) at the present heading, which is 164 kN of sway at 24 % of what the thrusters can produce, and sets the three manual demands to it, rounded to 1 kN and 0.1 MN m. Nothing measures the position, so nothing corrects the residual. The heading wanders 9.2° in the first five minutes, the balance computed for the original heading no longer matches the load, and the vessel leaves the watch circle at 267 s, before the weather has changed at all. The wind then freshens at 300 s and veers 30° at 700 s, the vessel passes the disconnect radius at 343 s, and it is 124 m away at twenty minutes. Eight thousand tonnes with no restoring force is a pure double integrator, and every scenario after the first is about what feedback costs rather than what it is for.

6.6 Results

The values that follow come from the application's own scorecard and its exported samples, at the constants of Table 6.2, with each setting made before the run and the shipped seed. Accumulating measures such as thruster command travel and energy are quoted with the horizon over which they accumulated, and averages are quoted with their window, because both continue to move after a transient ends. Peaks and crossing times are taken at the 0.05 s physics step.

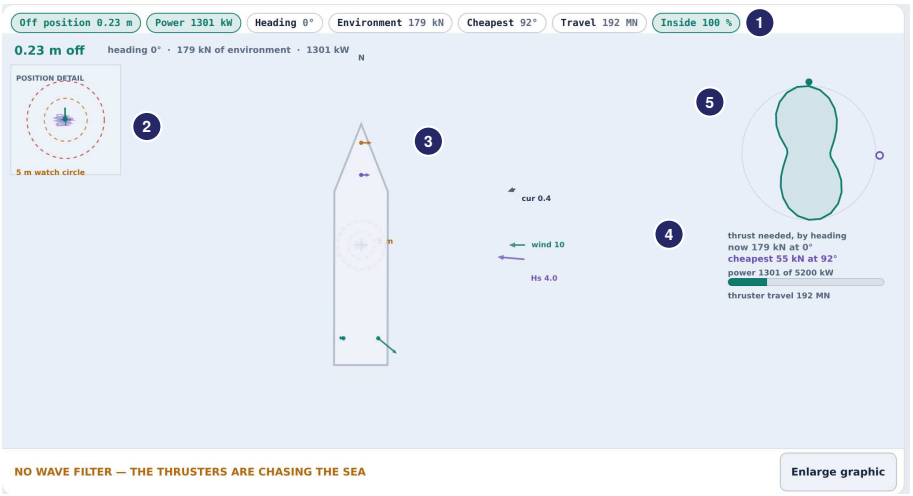


Figure 6.4 Scenario 3 in progress, with the controllers looking straight at the position reference in a 4 m sea at $T_p = 11$ s, cropped to the plant view. (1) The badges give the frame at $t = 900$ s, with the vessel 0.23 m off station on 1301 kW and the command travel at 192 MN, on the way to 327.5 MN over the 300 to 1800 s window used throughout Section 6.6, against 72.0 MN in the identical sea with the wave filter in. (2) The position detail draws the watch circle. (3) The hull carries its thruster arrows, and (4) the current, the wind and the sea arrive from the side. (5) The thrust needed by heading is the objective of the heading supervisor. The banner at the foot names the cause. The trend strips, in which the demands reverse at the wave period, and the scorecard are cut.

6.6.1 The Cost of Chasing the Sea

Scenarios 3 and 4 are the same seaway with the same tuning and the same environment feedforward, differing only in what the controllers are shown. The environment is $H_s = 4$ m at $T_p = 11$ s from 95°, with 10 m/s of wind and 0.4 m/s of current, which is a large sea and a modest steady load on purpose, because a vessel that cannot hold the environment anyway would teach the wrong lesson twice. Table 6.4 reports the three estimator settings, the model-only column being scenario 4 with that setting chosen before the run. Figure 6.5 draws four of its rows as bars, and Figure 6.4 shows scenario 3 in progress.

> **TRY IT IN THE APP** **3 · The Sea Arrives, With No Filter**
>
> Load the case and press Run. Watch the surge and sway demand strips rather than the position: they reverse at the wave period for as long as the sea lasts. Then load scenario 4 and compare the Thruster travel tiles at the same time. Over the chapter's window the travel is 327.5 MN against 72.0 MN, while the vessel sits about half a meter off station in both.

The travel figure is the one to put in a business case, because it is the only one that moves by a factor. Over the observation window the thruster commands travel 327.5 MN with the raw reference and 72.0 MN with the wave model, almost all of the difference in reversals at the wave period on machinery designed for hours of steady running. Energy is not the argument at 9 %, and a reader shown only the fuel bill would conclude the wave filter was optional. The vessel is not better positioned for any of it either. It sits 0.55 m off on average against 0.51 m, because a force that reverses every five seconds does nothing at all to eight thousand tonnes. An earlier version of the application gave 1.34 m against 0.63 m, a gap that came mostly from the feedforward its raw setting withheld and from the start-up of Section 6.3.6, both since repaired.

Table 6.4 The Same 4 m Seaway and the Same Tuning at Three Estimator Settings. Every Figure Is Taken Over the Same Observation Window, the 1500 s From $t = 300$ s to $t = 1800$ s, With Travel and Energy Accumulated Across It, Mean Power as Energy Over the Window, and the Rest Averaged or Maximized Across It. The Environment Feedforward Is On in All Three Columns.

Quantity	Raw	Model only	Observer
Surge demand, standard deviation (kN)	96	72	65
Thruster command travel (MN)	327.5	91.1	72.0
Thruster energy (kW h)	452	461	412
Mean thruster power (kW)	1085	1107	989
Mean off-position, low frequency (m)	0.55	0.51	0.51
Peak off-position (m)	3.02	3.19	3.29
Mean estimate error (m)	0.77	0.55	0.41
Inside the watch circle (%)	100.0	100.0	100.0

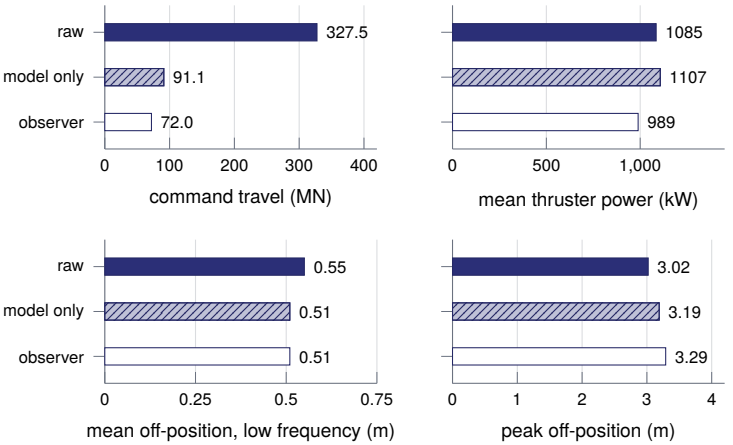


Figure 6.5 Four rows of Table 6.4 drawn as bars for the three estimator settings in the same 4 m sea over the window from 300 to 1800 s: solid for the raw reference, hatched for the model only and white for the full observer. Notice that only the command travel moves by a factor, while the mean power and the mean off-position hardly change and the peak off-position is a little larger with the observer.

> ✦ DID YOU KNOW? **A Quieter Loop With a Worse Worst Moment**
>
> From the simulator: the full observer does not buy a smaller worst moment. Over the window its peak off-position is 3.29 m, against 3.02 m on the raw reference that costs four and a half times the command travel. The observer saves the machinery, not the worst moment (Table 6.4).

The measurement carries two motions the thrusters cannot answer rather than one, and the wave filter rejects both. With σ_p and σ_g set to zero before the run the same comparison gives 82.8 MN against 25.4 MN, a factor of 3.3 against the 4.5 measured with the instruments as specified, so wave-frequency motion alone accounts for the smaller part of the difference. That is a stronger claim than a noise-free reference can support, and a more honest one, because an instrument resolving position to a third of a meter is the realistic case. The silenced generators draw no random numbers, so the sea in those runs is a different realization of the same seaway, with a 2.21 m peak.

The mechanism is the derivative term. With the raw reference the position handed to (6.15) carries the full wave motion, 2.39 m at its peak at $\omega_0 = 0.571\,\mathrm{rad/s}$, a rate amplitude of 1.37 m/s. The derivative filter at $\tau_f = 9$ s attenuates that by $1/\sqrt{1+(\omega_0\tau_f)^2} = 0.19$, and $K_c^p \tau_D^p = 1.55\,\mathrm{MN\,s/m}$ turns the remainder into about 400 kN of commanded amplitude at the peak. The same product acting on the instrument noise is what makes τ_f itself compulsory rather than merely useful: in the calm

water of scenario 2 with the raw reference and $\tau_f = 0$ the surge demand runs from one 680 kN clamp to the other, at a clamp 20 % of the time, with a standard deviation of 504 kN over 300 to 1800 s, against 91 kN at $\tau_f = 9$ s and 2 kN with the noise removed as well.

Removing the derivative term altogether removes the loop with it, which is what scenario 2 exists to show. In its calm water a 5 m north step made at 300 s is 63 % complete at 27.6 s and 95 % at 79.2 s, and the low-frequency position peaks at 6.29 m over the following ten minutes against 5.65 m with the instruments silenced. Stepped at other instants between 200 and 900 s the peak ranges from 6.10 to 6.39 m, so the noise costs between eight and fifteen points of overshoot and little else. At $\tau_D^p = 0$, set before the run, the same step first overshoots to 10.7 m at 53 s, and the swing then grows into a sustained oscillation of about 15 m either side of the setpoint with a period near 102 s, still running an hour later. The derivative term is compulsory on this plant and it is also the term that converts both wave motion and instrument noise into thruster commands, which is precisely why the answer has to be an estimator and not a filter.

🕮 **Pause.** Try Exercise 6.1 before reading on.

The intermediate setting is the useful control. A vessel model and a bias estimate with no wave model give 91.1 MN of travel, capturing 92 % of the reduction from 327.5 MN to 72.0 MN and holding station as well as the full observer to within a centimeter. Almost all of the benefit is in having any model at all, and the remaining 19.1 MN is what the wave states are worth. That is a weaker case for the wave states than the same comparison on a noise-free reference, where the model-only setting captures only 60 %, and it is the honest one.

6.6.2 The Observer Gain

Scenario 5 walks the observer gain across its range in the same heavy sea, raising it to 1.2/s at 300 s, dropping it to 0.05/s at 900 s and restoring the tuned 0.28/s at 1500 s. Held at the tuned gain over 300 to 1800 s (scenario 4), the estimate sits 0.41 m from the true low-frequency position on average, the vessel holds 0.81 m from the wellhead and the commands travel 2.88 MN per minute. At $K = 1.2/\text{s}$ the travel is 11.93 MN per minute in scenario 5 over 300 to 900 s and 12.18 MN held over the whole window, nearly all the way back to the 13.10 MN of no filter at all, because an estimator that believes the measurement completely is not an estimator and the measurement contains the sea and the instrument noise together. In the ten minutes at $K = 0.05/\text{s}$ the estimate is smooth and 1.78 m from the truth on average, 2.30 m at worst, and the vessel wanders to 4.14 m while the trace on the screen reports that everything is fine. Held at 0.05 from the start of scenario 4, where the estimator opens on the true load and nothing changes it, the estimate averages only 0.21 m from the truth over 300 to 1800 s. The drift in scenario 5 is the low gain failing to correct the velocity error the high-gain spell leaves behind, 0.19 m/s in the last ten seconds before the switch against 0.03 m/s over the first 300 s at the tuned gain.

The price of filtering is stated cleanly by scenario 10, which steps the position reference 3 m at $t = 400$ s. Measured as the estimate's offset from the true position, the estimate covers 63 % of that step in 5.05 s at the tuned gain and in 17.35 s with the gain set to 0.05/s before the run. Anything the estimator puts into the wave states and the vessel model it does not put into the position, so an estimator tuned for a heavy sea is slower to notice a real excursion. That is the trade in every filtering problem, and it is why industrial observers schedule the gain against the sea state rather than fixing it [6].

6.6.3 Bias Feedforward Against Integral Action

Scenario 6 removes the feedforward at 200 s and restores it at 900 s. Until then the bias estimate holds the environmental force it opened with, and the controller subtracts it before the loops see any error at all. When the load changes the estimate has to find the change, and after the squall of scenario 7 it covers 63 % of it in 237 s. With the feedforward off the vessel is pushed 4.38 m downwind by 387 s and is still 1.78 m out when the feedforward returns at 900 s, because the only mechanism left that

can find a steady force is an integral with a ten-minute time constant, long deliberately, because a fast integral on a plant with this little damping rings. The return is not clean. The integral still holds most of the force it found, so the force is counted twice: the vessel crosses inside half a meter at 919 s, carries on to 3.28 m off the other way at 1073 s, and is still 1.41 m out at 1500 s while the same slow integral unwinds.

Scenario 7 measures what the arrangement is worth when something happens. A squall takes the wind from 12 to 24 m/s and the sea to 4 m, which raises the environmental demand at zero heading from 164 to 417 kN and the mean thruster power from 15 % of installed over 100 to 300 s to 51 % over 500 to 1500 s. The vessel is pushed 6.7 m off station before the estimator works out what has happened, and is back inside a meter 6.9 min after the squall arrives. With the feedforward off before the run the integrals carry the original load and the same squall takes the vessel 7.7 m off, enough with the wave motion on top to pass the 9 m disconnect radius at 435 s, and it is still 3.4 m off at 900 s and not back within a meter until 1195 s, a quarter of an hour after the squall. The bias converges to within 2 % of the true environment about ten and a half minutes after the squall, which places the feedforward exactly where it belongs. It is fast enough to matter and not fast enough to be the answer on its own.

6.6.4 The Heading Is Worth More Than the Gains

Scenario 8 holds the vessel beam on to 17 m/s of wind, 0.8 m/s of current and 3.5 m of sea, an environment of 372 kN that the exact minimum-norm solution carries on 2135 kW. The scripted selection of the minimum-thrust heading at 400 s slews the setpoint through 94° in 375.5 s at 0.25 °/s, and the heading has held within 5° of the cheapest one for two minutes by 981 s. Afterward the environment asks for 109 kN and the mean thruster power over 1400 to 2400 s is 738 kW against 2297 kW over 100 to 400 s before the turn, a saving of 68 %. Nothing about the control system changed. The vessel presents about a third of its wind area and a fifth of its underwater area to the weather, which is the whole of the mechanism and is visible in Figure 6.3.

> **TRY IT IN THE APP** **8 · The Heading Is Free**
>
> Load the case and press Run. At 400 s the minimum-thrust heading is selected: watch the strip Heading against the cheapest one as the setpoint walks through 94° in 375.5 s. Compare the Power tile before the turn and after 1400 s. The chapter's means are 2297 kW and 738 kW, with the control system unchanged.

This is the result the application is organized around. A single supervisory choice over a variable the job leaves free is worth a factor of three in thruster power, and no position loop can make it, because it has no way to know that the setpoint it is not being given is free. Weather optimal positioning of this kind has been designed as a feedback law and demonstrated experimentally [15], and the version here is the simpler supervisory form.

6.6.5 A Solved Allocation Against a Written One

Scenario 9 trips the port azimuth at 400 s in a brisk sea, and Table 6.5 compares the two allocations, the fixed split chosen before the run, over the ten minutes from 300 to 900 s that bracket the trip.

Table 6.5 Port Azimuth Failure at $t = 400$ s. Both Columns Use the Same Tuning and the Same Environment, and Every Statistic Is Taken Over the 600 s From $t = 300$ s to $t = 900$ s. The Median Starboard Azimuth Thrust Over 200 to 400 s, Before the Trip, Is 65 kN.

Quantity	Minimum norm	Fixed split
Peak off-position (m)	1.87	5.40
Mean off-position (m)	0.70	2.86
Median starboard azimuth after the trip (kN)	128	105

The pseudo-inverse solves the same problem with one fewer column and hands the work to whatever is left, and nobody wrote a rule for it, because a failed machine is a weight of zero in (6.17) and the arithmetic redistributes on its own. The fixed split of (6.19) does have a rule, the rule is still followed with one azimuth gone, and the share belonging to the missing machine is simply not produced, so the remaining azimuth never rises past a median of 105 kN against the 128 kN the pseudo-inverse asks of it. The vessel is four times as far off, sustained, in an environment that happens to be almost all sway. Placing the weather on the bow makes the same failure considerably more expensive, because the surge axis has only the two azimuths.

6.6.6 A Fault No Controller Can Answer

Scenario 10 applies 3 m of reference bias to the north channel at 400 s and removes it at 1200 s. The vessel drives 3 m the other way, exactly the bias (3.18 m south on average over 1000 to 1200 s), after a 5.9 m overshoot at 429 s on the way, to correct an error that does not exist, and every trace available on the bridge agrees that it is on position. Filtering does not help: a steady offset is precisely what a bias estimator and an integral are built to find, and they find it perfectly. The vessel is more than a meter off the wellhead from 410 s to 1340 s, 15.5 min spent off station for a reason nobody aboard could have seen. The remedy is not in the control system. It is three independent position references and a voting scheme that rejects the one that disagrees, which is why a vessel of dynamic positioning class 2 carries them, and loss of position reference is among the dominant contributors to real incidents [12].

🕮 **Pause.** Try Exercise 6.2 before reading on.

6.6.7 What Actually Runs Out

Scenario 11 holds the vessel beam on while the weather builds through 25 m/s at 300 s, 29 m/s at 700 s and a current of 1.6 m/s at 1100 s. Three things could bind, the sway thrust, the thruster power, or the heading, and the answer is settled before the first scheduled event. The heading leaves a 3° band at 17 s, the vessel leaves the watch circle at 50 s and passes the disconnect radius at 75 s, while the power sits between 54 and 68 % throughout and never approaches its limit.

The arithmetic follows from (6.18) alone and is the result an operator can carry off this application and use. With both tunnels at their 120 kN rating and the azimuths free, the vessel can make $2F_a + 2F_t = 680\,\text{kN}$ of sway if it does not care where the bow points. Requiring zero net yaw moment couples the two, because $x_a \Sigma f_{y,azi} + x_{t1} f_3 + x_{t2} f_4 = 0$ forces the azimuths to produce $(38 f_3 + 26 f_4)/35 = 219\,\text{kN}$ of sway purely to cancel the tunnels, leaving 459 kN of useful beam thrust. Requiring the vessel also to cancel the 3.1 MN m of yaw moment the environment itself applies takes the azimuth share down to 130 kN and the total to 370 kN, against a beam demand of 541 kN. The vessel therefore cannot hold sway and heading at once, the heading loses first, and once the heading has gone the projected areas in (6.4) and (6.5) change and the rest follows, at 21 m off station at two and a half minutes, 320 m at eighteen and 904 m at thirty.

🕮 **Pause.** Try Exercise 6.3 before reading on.

Heading capability rather than installed thrust or power is therefore what binds, which is why a capability plot is drawn per heading. Once the tunnels are pinned there is no recovery: with the minimum-thrust heading chosen before the run, the setpoint walks toward the weather but the bow stalls near 28°, the disconnect radius is reached at 68 s against 75 s, and the vessel is 816 m off at thirty minutes. The operational answer is the turn of scenario 8, made before the weather arrives, while the vessel still has yaw authority to turn with.

6.6.8 The Graded Challenge

The challenge is a forty-minute shift opening with the weak tuning and four things set the wrong way: a model-only estimate with no wave states, a fixed allocation, no environment feedforward and a heading nobody chose. The sea builds at 300 s, a squall arrives at 700 s, the port azimuth trips

at 1400 s and the squall passes at 1900 s. The target, judged at 2400 s, is to hold within 3 m of the wellhead throughout the last third, to stay inside the watch circle for at least 97 % of the shift, never to run over the power limit, and to finish under 130 MN of thruster travel. As loaded the shift fails badly, 23.0 m off at worst in the last third and 20.4 % inside the circle, while meeting the travel and power clauses on 73 MN, which is why travel alone is not the target. No one of the four structural changes passes on its own, and the four together on the weak gains still fail at 10.2 m and 46 % inside.

With all four decisions right the shift finishes on 102.5 MN of travel, 100.0 % inside the watch circle, a peak thruster power of 60 % of installed and 2.64 m off station at worst over the last third. Against the target these are margins of 0.36 m, 3.0 points and 28 MN, and the application gives the same run whether the suggested settings are revealed before the run or carried through Reset. Putting each one back the wrong way, after revealing the suggested settings and before the run, separates their contributions. Leaving the heading fixed leaves the vessel 9.9 m off on average over the last third and outside the watch circle for 40 % of the shift. Leaving the wave model out gives 140.7 MN of travel against a target of 130 MN, and 3.22 m at worst, and the raw reference gives 459.2 MN. The minimum-norm allocation is what makes the trip survivable, because the fixed split finishes 53.7 % inside the circle and 4.71 m off at worst, with the power peaking at 99.7 % of installed. The power clause is a guard rather than a test, because the allocation scales the thrust back before the switchboard is overloaded, so no line-up measured here spends any time over it. The feedforward is the one of the four that could be omitted, still a pass at 100.0 % inside but only just, 2.98 m at worst against the 3 m limit on the shipped seed. None of the four is a gain, and the tuning that is left is nearly what it was.

6.7 Technology in Practice

The vessel of Section 6.3 is read by one position reference and a gyro, and moved by two azimuths and two tunnel thrusters. Figure 6.6 draws its loops with real devices in their place.

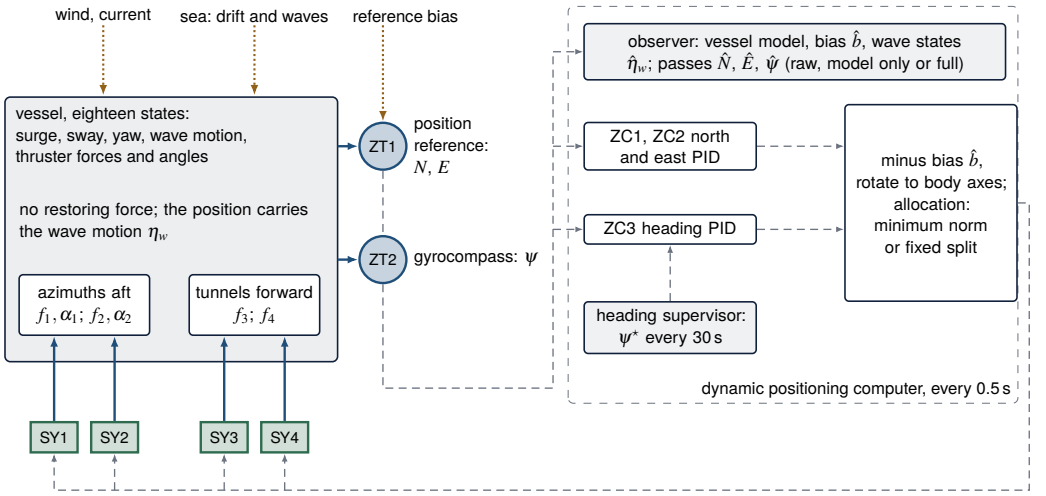
What Measures It

The reference of Equation (6.8), with its 0.35 m of colored noise, its 0.5 s hold and its optional bias, stands for one real position reference. The first deep-water drill ship carried two acoustic systems in parallel, and the redundancy was used whenever one of them was being repaired [3]. A vessel of class 2 or 3 now carries at least three references built on different principles, and three gyrocompasses where its control system depends on the heading, beside sensors for its motions and the wind [16], where the model has one reference and one gyro.

> **TECHNOLOGY** **Sensor: Hydroacoustic Position Reference**
>
> A battery beacon on the sea floor sends short sound pulses, and hydrophones at the four corners of a square in the hull hear each pulse at slightly different times. The differences are changed to north and east offsets from the beacon, and two small gyros correct them for roll and pitch. On the first deep-water drill ship, thruster cavitation drowned the pulses at high speed, beacons faded after 40 to 60 hours when their floats took on water, and one wiring fault left the computer holding station on the last acceptable signal it had received [3]. On a trend such a fault is a reading that freezes or jumps while the vessel moves, the reference bias of scenario 10 in another form.

What Moves It

The model drives each thruster through a 4 s build-up, turns the azimuths at 6 °/s and takes thrust from the tunnels as the vessel gathers way, so a lost or weakened unit is a weight in the allocation of Section 6.4.3. The first deep-water drill ship chose fixed propellers in tunnels for reliability and for protection from the drilling work and from storms. Fixed thrusters need more connected power than steerable ones, since they cannot point their thrust exactly where it is wanted, but its designers found the power actually used not appreciably different [3].



SY1, SY2: azimuth thruster drives, 220 kN each, thrust lag 4 s, turning at 6 °/s. SY3, SY4: tunnel thruster drives, 120 kN each, thrust lag 4 s, less thrust with way on. All four share 5200 kW.

Figure 6.6 The loops of the chapter with real devices in place of the schematic's symbols. ZT1 is the position reference (on a real vessel several, built on different principles, in the model one with 0.35 m of noise) and ZT2 the gyrocompass. One dynamic positioning computer (dashed frame) runs the observer, the north, east and heading PID loops (ZC1 to ZC3), the heading supervisor and the allocation every 0.5 s. SY1 and SY2 drive the two azimuths aft, SY3 and SY4 the two tunnels forward. Notice that the loops see only the observer's estimate, never the reference itself unless the raw setting is chosen, and that the reference bias enters at the instrument, where no loop can tell it from a real move.

TECHNOLOGY Final Element: Tunnel Thruster

A tunnel thruster is a fixed propeller in a tube across the hull that pushes the bow or the stern sideways. On the first deep-water drill ship each was a four-bladed propeller of 59 inches turned through a right-angle gear by an electric motor, for about 18,000 pounds of thrust. The long bow tunnels cut that thrust until the blades were reshaped, cavitation set in at full speed, and the bow units, which ran longer and harder than the stern ones, failed several times in service [3]. Thrust is lost as well to interference between thrusters and the hull [16]. On a trend a lost thruster shows as a step in the others' commands, as the starboard azimuth of scenario 9 rises from a median of 65 kN to 128 kN (Table 6.5).

Where the Controller Runs

Supervisory and lower-level control loops run with cycle times from milliseconds to minutes [17]. The chapter's observer, three PID loops and allocation share one 0.5 s scan, and the heading supervisor rescans every 30 s on top of them. In 1968 that work was done by a computer fed with digital angles from the acoustic receivers [3]. In class 2 it is done by at least two independent computers that transfer control automatically after a detected failure, each on its own uninterruptible supply, and their software also runs the consequence analysis that checks the thrust left after the worst failure [16].

What Is Modernizing It

The guideline of 1994 was replaced for vessels built from June 2017, and one flag administration recommends the new operational requirements for existing vessels as well [18]. The computers of a redundant system now talk over networks, and a regulator has warned that a dual bus network can be silenced by one faulty module where a dual star network need not share a medium at all [19]. Test platforms join a six-degree-of-freedom simulation, with the frequency-dependent hydrodynamics the

chapter's three-degree model leaves out, to a scale model tested at the same Froude number, and use the pair to find a vessel's operational time window [20].

> **IN PRACTICE** **Equipment Class and the Worst Failure**
>
> An international guideline of 1994 sorts dynamic positioning systems into three equipment classes by their worst failure. In class 1 a single fault may lose position. In class 2 no single fault in an active component, such as a generator, a thruster or a switchboard, may lose it, and class 3 adds the loss of everything in one compartment to fire or flooding. Classes 2 and 3 carry at least three position references, not all of one type, and three gyrocompasses. A consequence analysis warns when the thrusters left after the worst failure could no longer hold the vessel, the work stops before the weather reaches that point, and position keeping after single failures is tested every year [16].

6.8 Limitations

The application teaches an estimator structure and does not size a vessel. Blended area coefficients, a single-frequency sea and no thruster interaction put a real hull's capability footprint, its consequence analysis and its dynamic positioning class out of reach, and the capability arithmetic of Section 6.6 is exact for (6.18) and approximate for any real vessel. Three modeling choices in Section 6.3.7 cost more than the rest: the absent heave, which is what most often stops offshore construction work, the smooth coefficient sets, and the single position reference. Three implementation choices cost almost as much. The observer is hand tuned with one gain and a fixed pole pattern rather than carrying a covariance, so the adaptation industrial systems perform against the sea state [6] is left to a slider. Neither the passive nonlinear observer that removes the linearization about a fixed heading [11] nor the constrained allocation [9] is implemented, so the power limit scales the whole solution back rather than choosing a degree of freedom to surrender. The azimuth angles are settled by a filter outside the solve rather than a rate constraint inside it, and every automatic run opens with the estimator already converged on the true load (Section 6.3.6), a state a commissioned system reaches only by letting its bias estimate converge before it takes control.

> **OPEN CHALLENGE** **One Fault in a Redundant Network**
>
> Redundancy protects a vessel only if no single fault can reach both halves. A safety alert of 2013 reports a class 2 diving support vessel whose positioning system failed while divers were working below, so that the vessel drifted off position. The probable cause was a single fault that blocked the system's internal data communications: one module with access to both buses of a dual bus network can block both. The regulator asked owners and suppliers who claim class 2 or better to check their communication architecture [19]. Progress would be designs and tests that show, before a vessel works, that no single fault can silence every path between the references, the computers and the thrusters.

6.9 Discussion

Dynamic positioning is what makes deep-water work possible at all, and the control system on a mid-size construction vessel is the arrangement this chapter describes, a hand-tuned Kalman filter, three PID loops and a pseudo-inverse allocation, because that is what most of them are [2]. The measurements reported above say where the engineering value in that arrangement sits, and it is not where a tuning exercise would put it. Three PID loops account for none of the factor of 4.5 in thruster travel, none of the 68 % saving in thruster power and none of the difference between surviving a thruster failure and drifting off station. The estimator, the allocation and one supervisory decision account for all of it.

Of the three general questions named in the introduction, the third is the rarest and the most valuable, and it is the one a control course is least likely to raise. A variable the process leaves free is an optimization variable, and the gap between its present value and its best one is frequently larger than anything feedback can recover.

The model of Section 6.3 also supports methods beyond the loops implemented in it. Equations (6.1) to (6.10) are already the prediction model a nonlinear model predictive controller would carry, and the constraints a PID loop can meet only by clipping are what a constrained formulation represents naturally, including in the allocation, where a constrained solve handles the saturation, the azimuth rate and the forbidden zones that the pseudo-inverse and the filter of Section 6.4.6 only approximate [9].

6.10 Conclusion

An offshore vessel holding station on thrusters serves in this chapter as an interactive laboratory for the problem of deciding what part of a measurement is an error. The measured comparisons are collected in Section 6.6 and are not repeated here. What they say together is that the three PID loops in the middle of a dynamic positioning system account for almost none of the difference between a good one and a poor one. The estimator that decides what the measurement means, the allocation that turns three demands into six machines, and the one supervisory choice over a variable the job leaves free account for all of it, and the largest of the three is the one no loop can make. Three defects in the measurement path and one in the allocation were found and repaired along the way, three more in the controller since, and the wave-filter argument is stronger on the repaired engine than it was on the quiet one. Further work would add heave, a broad-banded sea, a voting scheme over three position references, and a power management system that surrenders a degree of freedom rather than scaling the whole allocation. The vessel is an accessible browser platform for a lesson that arrives before any tuning does, which is that part of a measurement can be real and still not be an error.

KEY POINTS

- A vessel on thrusters has no anchor and no restoring force, so its position reference shows both the slow drift the thrusters must stop and the wave motion they cannot and need not stop.
- Shown the raw reference, the thruster commands travel 327.5 MN over the window against 72.0 MN with the full observer, and the vessel holds station no better.
- Choosing the heading with the least environmental load saves 68 % of the thruster power in scenario 8, a choice no position loop can make.
- After a thruster trip the solved allocation keeps the vessel 0.70 m off on average against 2.86 m for a fixed split.

Where else this applies. Chapter 5 chooses which waves a foiling boat rides over and which it follows, and Volume 5, Chapter 13 filters the late position fixes of an underwater vehicle.

References

[1] T. I. Fossen, *Handbook of Marine Craft Hydrodynamics and Motion Control*. Chichester, UK: John Wiley & Sons, 2011.

[2] A. J. Sørensen, "A survey of dynamic positioning control systems," *Annual Reviews in Control*, vol. 35, no. 1, pp. 123–136, 2011.

[3] Deep Sea Drilling Project, "Dynamic positioning," Deep Sea Drilling Project, Tech. Rep., 1971.

[4] J. G. Balchen, N. A. Jenssen, E. Mathisen, and S. Sælid, "A dynamic positioning system based on Kalman filtering and optimal control," *Modeling, Identification and Control*, vol. 1, no. 3, pp. 135–163, 1980.

[5] T. I. Fossen and T. Pérez, "Kalman filtering for positioning and heading control of ships and offshore rigs," *IEEE Control Systems*, vol. 29, no. 6, pp. 32–46, 2009.

[6] T. D. Nguyen, A. J. Sørensen, and S. T. Quek, "Design of hybrid controller for dynamic positioning from calm to extreme sea conditions," *Automatica*, vol. 43, no. 5, pp. 768–785, 2007.

[7] K. J. Åström and R. M. Murray, *Feedback Systems: An Introduction for Scientists and Engineers*, 2nd ed. Princeton, NJ: Princeton University Press, 2021.

[8] S. Saelid, N. Jenssen, and J. Balchen, "Design and analysis of a dynamic positioning system based on Kalman filtering and optimal control," *IEEE Transactions on Automatic Control*, vol. 28, no. 3, pp. 331–339, 1983.

[9] T. A. Johansen and T. I. Fossen, "Control allocation—a survey," *Automatica*, vol. 49, no. 5, pp. 1087–1103, 2013.

[10] M. J. Grimble, R. J. Patton, and D. A. Wise, "The design of dynamic ship positioning control systems using stochastic optimal control theory," *Optimal Control Applications and Methods*, vol. 1, no. 2, pp. 167–202, 1980.

[11] T. I. Fossen and J. P. Strand, "Passive nonlinear observer design for ships using Lyapunov methods: full-scale experiments with a supply vessel," *Automatica*, vol. 35, no. 1, pp. 3–16, 1999.

[12] H. Chen, T. Moan, and H. Verhoeven, "Safety of dynamic positioning operations on mobile offshore drilling units," *Reliability Engineering & System Safety*, vol. 93, no. 7, pp. 1072–1090, 2008.

[13] W. Blendermann, "Parameter identification of wind loads on ships," *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 51, no. 3, pp. 339–351, 1994.

[14] J. A. Pinkster, "Mean and low frequency wave drifting forces on floating structures," *Ocean Engineering*, vol. 6, no. 6, pp. 593–615, 1979.

[15] T. I. Fossen and J. P. Strand, "Nonlinear passive weather optimal positioning control (WOPC) system for ships and rigs: experimental results," *Automatica*, vol. 37, no. 5, pp. 701–715, 2001.

[16] International Maritime Organization, "Guidelines for vessels with dynamic positioning systems," MSC/Circ.645, Maritime Safety Committee, London, 1994, adopted 6 June 1994.

[17] K. Stouffer, M. Pease, C. Tang, T. Zimmerman, V. Pillitteri, S. Lightman, A. Hahn, S. Saravia, A. Sherule, and M. Thompson, "Guide to operational technology (OT) security," National Institute of Standards and Technology, Gaithersburg, MD, Special Publication 800-82r3, 2023.

[18] Commissariat aux affaires maritimes, "Ships fitted with dynamic positioning systems," Circular CAM 03/2020, Luxembourg, 2020. [Online]. Available: https://maritime.public.lu/dam-assets/circulaires/003-Dynamic-Positioning-Systems.pdf

[19] Health and Safety Executive, "Warning to offshore industry on blocking of data communications in dynamic positioning systems," Safety alert OSD 1-2013, Offshore Safety Division, 2013. [Online]. Available: https://www.hse.gov.uk/safetybulletins/dynamic-positioning-systems.htm

[20] C. Hu, Q. Shi, X. Li, and X. Guo, "Development of a hybrid simulation and experiment test platform for dynamic positioning vessels," *Ships and Offshore Structures*, vol. 19, no. 10, pp. 1548–1562, 2024.

6.11 Exercises

6.1 ○ Concept check · 2 min

A designer proposes to keep the sea out of the thruster commands by placing a first-order low-pass filter, with a time constant of about one wave period, on the position reference ahead of the three PID loops. Why does an industrial dynamic positioning system use an observer with a wave model instead?

(a) A low-pass filter cannot attenuate a motion as slow as an eleven-second wave
(b) The filter's phase lag removes the damping that the derivative term must supply
(c) A low-pass filter also removes the steady environmental load the loops must carry
(d) The observer executes on a faster clock than any filter placed in the loop

6.2 ○ Concept check · 1 min

In scenario 10 the differential correction goes stale at 400 s and the north reading moves 3 m away from the truth. Which change would have kept the vessel on the wellhead?

(a) A higher observer gain, so the estimate follows the reference sooner
(b) A lower observer gain, so the estimate ignores the sudden jump
(c) Independent position references with voting that rejects the one that disagrees
(d) A shorter position integral time, so the offset is removed sooner

6.3 ○ Concept check · 2 min

In scenario 11 the weather builds on the beam with the heading as set. The beam demand is 541 kN and the four thrusters can make 680 kN of sway between them. Which limit binds first?

(a) The thruster power, because the wind load grows with the square of the wind speed
(b) The sway thrust, because the beam demand exceeds the installed sway at the start
(c) The position integral, which winds up against the demand clamp of the north loop
(d) The heading, because the pinned tunnels cannot hold the sway and the yaw at once

6.4 Theory and calculation · 8 min

The vessel of Table 6.2 has $m = 8.0 \times 10^6$ kg, $a_u = 0.08$, $a_v = 0.85$, $k_g = 0.25$, $L_{pp} = 90$ m, $a_r = 0.35$ and $D_u = 9.0 \times 10^4$ N s/m. At $\psi = 0$ the north loop acts in surge with $K_c^p = 37$ kN/m and $\tau_D^p = 42$ s.

(a) Find m_u, m_v and I_z from Equation (6.2), and the ratio m_v/m_u. *Checkpoint:* 8.64×10^6 kg, 1.48×10^7 kg, 5.47×10^9 kg m^2 and 1.71.

(b) Treat the north loop as a proportional-derivative law on the true position acting on m_u with damping D_u, and ignore the derivative filter, the integral, the observer and the thruster lag. Find the natural frequency, its period and the damping ratio. *Checkpoint:* 0.0654 rad/s, 96 s and 1.45.

(c) Find the damping ratio with $\tau_D^p = 0$, and compare the period with the sustained oscillation of about 102 s that scenario 2 shows at $\tau_D^p = 0$. *Checkpoint:* 0.080. The oscillation is a little slower than 96 s.

6.5 Theory and calculation · 6 min

In scenario 3 the raw position reference carries wave motion of 2.39 m peak at $T_p = 11$ s. The position loop has $K_c^p = 37$ kN/m, $\tau_D^p = 42$ s and a derivative filter $\tau_f = 9$ s.

(a) Find the peak wave frequency ω_0 and the peak rate of the wave motion. *Checkpoint:* 0.571 rad/s and 1.37 m/s.

(b) Find the gain of the derivative filter at ω_0 and the amplitude of the commanded force that the derivative term makes from the wave motion. *Checkpoint:* 0.19 and about 405 kN.

(c) Repeat (b) with $\tau_f = 0$ and compare the result with the 680 kN demand clamp. *Checkpoint:* About 2.12 MN, about three times the clamp.

6.6 Theory and calculation · 7 min

The two azimuths are rated 220 kN each at $x_a = -35$ m and the two tunnels 120 kN each at $x_{t1} = 38$ m and $x_{t2} = 26$ m. The vessel is stopped, so $\ell = 1$. The azimuths can point their thrust anywhere.

(a) Find the most sway the four machines can make if the yaw moment does not matter, and the most they can make with zero net yaw moment while both tunnels push at their rating. *Checkpoint:* 680 kN, and 459 kN with the azimuths making 219 kN.

(b) The environment also applies a yaw moment of 3.1 MN m in the same sense as the azimuths' moment, so the thrusters together must make 3.1 MN m in the sense of the tunnels' moment to cancel it. Find the sway left, and compare it with the beam demand of 541 kN in scenario 11. *Checkpoint:* About 371 kN, short of the demand by about 170 kN.

(c) With $k_P = 1.79 \times 10^{-5}$ kW per N$^{1.5}$, find the power of all four machines at their ratings and say whether power can be the binding limit. *Checkpoint:* 5182 kW against 5200 kW installed.

6.7 App challenge · 8 min

Load scenario 3, *The sea arrives, with no filter*, as shipped (Position estimate on Raw), default seed. Run to 1800 s at speed 20 and export the CSV. Read `metric_travel_kN` and `metric_energy_kWh` in the rows where `Time (min)` is 5 and 30. Then load scenario 4, *The same sea, filtered*, as shipped (Wave filter) and repeat.

(a) Find the thruster command travel over 300 to 1800 s for each run, in MN, within 1 MN. *Checkpoint:* 327.5 MN on Raw and 72.0 MN with the wave filter.

(b) Find the energy over the same window for each run, within 2 kW h, and the ratio of Raw to Wave filter for travel and for energy. *Checkpoint:* 452 kW h and 412 kW h; ratios 4.5 and 1.10.

6.8 App challenge · 8 min

Load scenario 8, *The heading is free*, as shipped, default seed. Run to 2400 s at speed 20 and read the Cheapest heading tile. Export the CSV.

(a) Average `power_kW` over the rows with `Time (min)` from 1.66 to 6.67 (100 to 400 s, before the turn), within 10 kW. *Checkpoint:* About 2296 kW.

(b) Average `power_kW` over the rows from 23.33 to 40 (1400 to 2400 s), within 10 kW, and find the saving as a percentage of the power before the turn. *Checkpoint:* About 738 kW, a saving of 68 %.

(c) Explain the saving from Figure 6.3 and say which setting of the control loops changed during the run. *Checkpoint:* Cheapest heading 94°. No loop setting changed.

6.9 App challenge · 8 min

Load scenario 9, *A thruster trips*, as shipped (Thrust allocation on Minimum norm), default seed. Run to 900 s at speed 20 and export the CSV. Press Reset, set Thrust allocation to Fixed split before Run, run to 900 s again and export a second file. For each file use the rows with `Time (min)` from 5 to 15 (300 to 900 s).

(a) Find the mean and the largest value of `radius_m`, within 0.05 m. *Checkpoint:* Minimum norm 0.70 m and 1.87 m; fixed split 2.86 m and 5.40 m.

(b) Find the median of `f2_kN`, the starboard azimuth, over the rows after the trip (`Time (min)` above 6.67), within 2 kN, and explain the difference between the two runs. *Checkpoint:* 128 kN against 105 kN.

6.10 Application area · 3 min

A vessel manager reviews the comparison of scenarios 3 and 4 and notes that the wave filter saves only about 9 % of the thruster energy. The manager asks whether the wave filter is worth keeping. Which number makes the case, and what does it stand for on the vessel?

6.11 Application area · 4 min

Every automatic run in the application starts with the bias estimate already holding the true environmental load. On a real vessel the DP system is started with the vessel already on station and then switched to automatic. What should the operator wait for before relying on automatic control, and what does this chapter's estimator say about how long that takes?

6.12 Review the agent · 5 min

An AI assistant, shown scenario 4 at the tuned observer gain, writes: "The thrusters are still busy in this sea. Set the Observer gain to 0.05 and leave it there in every weather. The estimate on the North strip becomes almost flat, which shows the vessel is now holding station better, and the thruster travel falls further." Find the flaw in this advice and say how you would check it in the application.