INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 18 Control of Nuclear and Fusion Systems

CHAPTER 12

Tokamak Vertical Stabilization Simulator for Gain Limits and the Recoverable Region of an Unstable Plant

Pages 320–341

Edited by John D. Hedengren

How to cite this chapter. "Tokamak Vertical Stabilization Simulator for Gain Limits and the Recoverable Region of an Unstable Plant," in *Control of Nuclear and Fusion Systems*, John D. Hedengren, Ed., vol. 18 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 12, pp. 320–341.

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

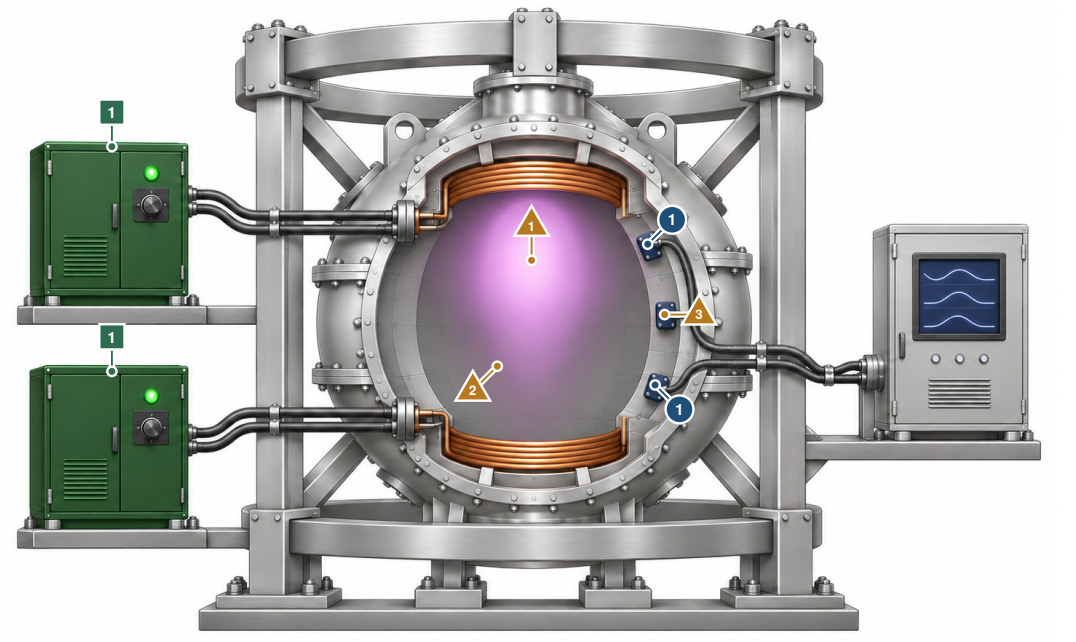
The app for this chapter
https://apmonitor.com/apps/tokamak-vertical/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v18-nuclear-fusion/#tokamak-vertical

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

CHAPTER 12

Tokamak Vertical Stabilization Simulator for Gain Limits and the Recoverable Region of an Unstable Plant



The controller moves 1 voltage of the upper and lower coil supplies

The controller reads 1 magnetic pickups on the vessel that sense the plasma's height and radius

What disturbs it 1 sudden kicks from edge bursts that jolt the plasma up or down

2 a drop in plasma pressure that shifts it inward 3 noise on the magnetic readings

The coils above and below the plasma are the controller's only push against a plasma that, once off center, falls faster and faster toward the wall. Illustration, not to scale.

IN PLAIN TERMS Holding Up a Plasma That Wants to Fall

The main idea. Inside a fusion machine, magnets stretch a ring of very hot gas into a tall shape that falls toward the wall once it moves off center. Coils above and below it push it back thousands of times each second, guided by magnetic sensors.

The challenge. Too gentle a push lets it drift away, and too hard a push, applied a little late, makes it shake. The coils also have a voltage limit, so a big enough jolt cannot be undone by any tuning.

Why it matters. If control is lost, the hot gas strikes the wall with its full electric current and puts large forces on the machine.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history boxes of Section 12.2, the fact boxes of Sections 12.3 and 12.5, the try-it boxes of Section 12.5, and the key points
Study	45 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 12.5, Section 12.6, Section 12.7 and the open challenges

The control question An elongated plasma falls toward the wall within milliseconds, and a pair of coils with a voltage limit and a latency must catch it. Which limits does the tuning set, which does the supply set, and when does a controller that knows the voltage limit recover less than one that does not?

Learning objectives

- Compute the wall stability margin and the growth rate from the force balance and the wall and coil circuits.
- Locate the gain floor and the gain ceiling of the position loop on runs, and explain how elongation, the wall and the latency move them.
- Compare a filtered differentiator with a model observer as the velocity estimate under magnetics noise.
- Predict the voltage each allocation delivers to the vertical loop when the radial loop saturates the shared supply.
- Explain why a supply limit inside a short predictive plan does not enlarge the recoverable region.

Prerequisites Sections 1.4.6 (floor, region and ceiling), 1.6.2 (velocity form) and 1.6.4 (running a case). Section 1.6.3 is needed only for the explicit MPC. Volume 7, Chapter 6 is a slower unstable plant to meet first.

The app https://apmonitor.com/apps/tokamak-vertical/

See also Volume 7, Chapter 6 for an unstable suspension bounded below by its gain inequalities and above by saturation and noise. Volume 7, Chapter 7 for the largest recoverable push on a finite rail. Volume 4, Chapter 13 for allocation against hard limits that switch. Volume 4, Chapter 12 for several actuators sharing one budget.

12.1 Introduction

Control courses usually treat an unstable plant as a problem of placing an unstable pole in the left half plane, after which the loop is considered solved. Real unstable plants are held against three limits that a pole placement calculation hides. There is a minimum gain, below which the actuator cannot push back as hard as the plant runs away. There is a maximum gain, set by the latency and the measurement noise. And because the actuator saturates, there is a region of states outside which the plant cannot be recovered whatever the gains, so that the size of the disturbance matters as much as the tuning [1, 2]. Section 1.4.6 derives the three limits for the reduced vertical mode of the example plasma P-1.

The vertical position of an elongated tokamak plasma is a clean industrial home for all three. The shaping field that elongates the plasma produces a destabilizing force proportional to the displacement, the conducting vessel wall resists any motion with induced eddy currents, and the instability grows at a rate set by the wall time constant and a wall stability margin [3, 4]. The control coils act through a supply limited in voltage and a digital loop with a finite cycle, and losing control produces a vertical displacement event that drives the plasma into the wall with its full current and large forces on the structure [5].

The simulator addresses the belief that an unstable plant is safe once it has a stabilizing gain. The class of problem is general: magnetic bearings and levitation (Volume 7, Chapter 6), balancing robots (Volume 7, Chapter 7), aircraft flown with relaxed static stability and rockets landing on thrust all have a gain floor, a gain ceiling and a recoverable region, and the design of each begins with the growth rate and the actuator authority rather than with a tuning rule [6].

This chapter works with a vertical stabilization simulator in which the elongation, the wall, the supply, the latency and the noise can be changed and their consequence read as a growth rate, a recoverable margin and a loss or survival. The rigid massless model, position and velocity control, priority allocation and explicit predictive control are all established methods. What the simulator adds for study is a two-circuit model whose open-loop growth rate is exact and whose unstable mode and recoverable region are computed live, with a model-based velocity observer beside the usual filtered differentiator, and an explicit model predictive controller whose voltage limit follows the

allocation, compared with the loop on identical runs. The cases find the floor, the ceiling and the region one at a time.

Section 12.2 traces vertical control from the first elongated machines to the controllable-displacement specification of ITER. Section 12.3 sets out the force balance, the wall and coil circuits, the growth rate and the recoverable region. Section 12.4 describes the position and velocity loop, the two velocity estimates, the radial loop with its two allocations and the explicit MPC. Section 12.5 reports the measured floor, ceiling and region, the offset and allocation cases, the graded challenge and the comparison of the MPC with the loop. Section 12.6 turns to the sensors, supplies and control systems of real machines, Section 12.7 states what the model leaves out, and Section 12.8 summarizes the results.

12.2 Vertical Control of Elongated Tokamaks

Early tokamaks had circular cross-sections and were vertically stable. As machines of the late 1970s and 1980s elongated their plasmas to raise the current and the pressure a given toroidal field could hold, vertical stability became a design constraint, and active control of the vertical position with coils and fast power supplies became standard. A review of that period established the elements still used: a passive conducting structure that slows the instability, an active loop that stabilizes it, and a stability margin that measures how much slowing the passive structure provides [3]. The relation between elongation, aspect ratio and the achievable vertical stability was mapped soon after [7].

> ⧖ HISTORY **Letting the Plasma Drift to Measure It (1982)**
>
> Reactor designs of the early 1980s chose elongated plasmas for a higher limit on plasma pressure and a larger volume, yet the forces that stretch the plasma also make it move unstably upward or downward. A 1982 study on a divertor tokamak measured how strong that instability was. Each plasma was formed in a stable shape and moved radially into the unstable range with the vertical loop still holding it. The vertical feedback was then removed, not by switching off the amplifier but by holding its input at zero, so the currents induced by the motion still flowed. The plasma drifted up or down at 0.2 to 1.0 m/s, followed over 30 to 350 ms, and the drift gave the growth rate. A feedback amplifier of only 100 kW, aided by passive structure, held plasmas of 170 to 310 kA [8].

Control-oriented plasma response models followed. Linearized models of the coupled plasma, vessel and coil circuits made it possible to design controllers for current, position and shape together [9], and practice on large machines grew into integrated magnetic control systems [10]. Definitions of the stability margin were examined carefully, because a margin computed with a rigid displacement model and one computed with a deformable plasma can differ [4].

> ⧖ HISTORY **A Recoverable Displacement Becomes a Requirement (2008 to 2010)**
>
> By 2008 the design of ITER asked how far its plasma could drift before full voltage could no longer bring it back. Experiments on several tokamaks disabled vertical control for varying intervals to measure this maximum controllable displacement, and the measured values came out close to the calculated ones or well below them. The machines' experience implied that ITER would need a recoverable displacement above about 5 % of its minor radius [11]. A pair of in-vessel coils for vertical stabilization was then designed from a simulation in which the feedback is turned off, the plasma drifts 16.5 cm upward, and the feedback is restored. The coils had to work behind the blanket in a radiation field of about 3000 MGy and to deliver 30 000 such pulses over the life of the machine [12, 13].

The control-theoretic side matured in parallel. The conditions under which a massless plasma model is adequate for vertical stabilization were established [2], the multivariable structure of the vertical instability and its feedback stabilization were analyzed [6], and the effect of actuator saturation on the set of recoverable disturbances was treated directly, with a DIII-D application [1]. For ITER,

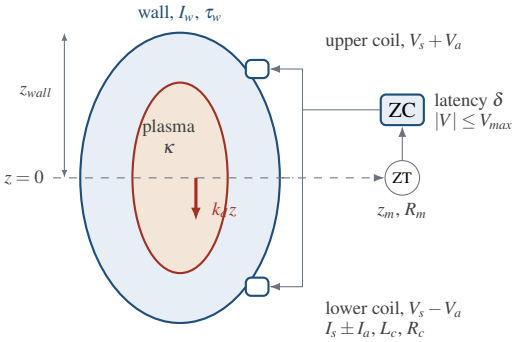


Figure 12.1 Vertical stabilization of an elongated plasma. The plasma of elongation κ, displaced by z from the midplane, feels the destabilizing force $k_d z$ and is held by the wall eddy current I_w with time constant τ_w and by the antisymmetric current I_a in the upper and lower in-vessel coils. The coils are driven with $V_s \pm V_a$, each within $\pm V_{max}$, the symmetric part V_s moving the plasma radially. The magnetics ZT measure z and R, the controller ZC acts after a latency δ, and the plasma is lost when it reaches the first wall at z_{wall}.

Table 12.1 Variables at the Operating Point

Symbol	Role	Nominal	Range
z	CV, vertical position	0 mm	SP $\pm$20
R	CV, radial position	0 mm	SP $\pm$30
V_a	MV, antisymmetric voltage	0 V	$\pm$10
V_s	MV, symmetric voltage	0 V	$\pm$10
I_w, I_a, I_s	states, currents	0	$\le$ 6 kA
R	state, radial position	0	—
κ	DV/parameter, elongation	1.80	1.50–2.02
n_z	DV, magnetics noise (1σ)	0.05 mm	0–1
—	DV, kicks (ELM, large, minor)	3, 20, 4 mm	events
—	DV, β drop	−15 mm	event

the maximum controllable displacement became a design requirement, as the second history box of this section tells. The forces on the vessel when control is lost are large enough that vertical displacement events are modeled in their own right [5].

12.3 Plasma, Wall and Coil Model

The model follows the rigid, massless displacement idea used in control analysis [2]. At elongation 1.8 the unstable growth rate with the coil voltage held is 356 s^{-1}, one e-fold every 2.8 ms, and the first wall is 150 mm from the plasma. The arrangement is shown in Figure 12.1.

The variables are classified in Table 12.1. Five states are carried: the wall current I_w, the antisymmetric and symmetric coil currents I_a and I_s, the radial position R, and a flag set at the loss of the plasma.

12.3.1 Force Balance and Circuits

The plasma is treated as a rigid current filament with negligible mass on these time scales, so the vertical forces on it balance at every instant, Equation (12.1).

$$k_d z + g_w I_w + g_c I_a = 0 \tag{12.1}$$

where k_d is the destabilizing stiffness, $g_w = 1.0\,\mathrm{N/A}$ and $g_c = 0.8\,\mathrm{N/A}$ the forces per ampere of wall and coil current. The position is therefore algebraic, $z = -(g_w I_w + g_c I_a)/k_d$. The wall and the coils are L-R circuits with mutual coupling, and a moving plasma induces a voltage in each, Equations (12.2) and (12.3).

$$L_w \dot{I}_w + M\dot{I}_a + g_w \dot{z} = -R_w I_w \tag{12.2}$$

$$M\dot{I}_w + L_c \dot{I}_a + g_c \dot{z} = V_a - R_c I_a \tag{12.3}$$

with $R_w = L_w/\tau_w$ and $M = k_M\sqrt{L_w L_c}$. Substituting $\dot{z}$ from the time derivative of Equation (12.1) gives a two-state linear system $\dot{x} = Fx + GV_a$ in $x = [I_w, I_a]^T$, solved at each step as a 2×2 linear system for the current rates.

12.3.2 Growth Rate and the Stability Margin

With the coils open-circuited the system has one unstable root, and its growth rate is exact, Equation (12.4).

$$\gamma\tau_w = \frac{1}{m}, \qquad m = \frac{g_w^2}{L_w k_d} - 1 \tag{12.4}$$

The wall stability margin m measures how much stiffer a perfectly conducting wall would be than the destabilizing push. The simulator ties it to elongation by $m(\kappa) = \max[0.02,\, 1.2(2.05 - \kappa)]$, so that $k_d = g_w^2/[L_w(1+m)]$ and a minor disruption scales m by $(1+\Delta m)$. At $\kappa = 1.8$, $m = 0.30$ and the wall alone would give $667\,\mathrm{s^{-1}}$. The coil circuit held at constant voltage also resists the motion, so the growth rate shown by the application, the unstable eigenvalue λ of F, is $356\,\mathrm{s^{-1}}$. It is $168\,\mathrm{s^{-1}}$ at $\kappa = 1.5$ and $1070\,\mathrm{s^{-1}}$ at 2.0.

12.3.3 The Recoverable Region

With w the left eigenvector of F for λ and $\xi = w^T x$ the unstable mode, the mode obeys Equation (12.5).

$$\dot{\xi} = \lambda\xi + bV_a, \qquad b = w^T G, \qquad |\xi| < \xi_{max} = \frac{|b|\,V_{max}}{\lambda} \tag{12.5}$$

Full voltage applied in the right direction can reverse the mode only while $|\xi| < \xi_{max}$, so the recoverable margin $1 - |\xi|/\xi_{max}$ is shown live. It grows smaller as λ rises and larger with the supply voltage, and it does not depend on the controller. It is Equation (1.10) written for the mode of the two-circuit model, with ξ in place of the position.

> **DID YOU KNOW? A Few Centimeters to Spare**
>
> On a research tokamak whose plasma grew vertically at $410\,\mathrm{s^{-1}}$, near the $356\,\mathrm{s^{-1}}$ of this chapter's plasma, the largest displacement that full control effort could still reverse was calculated at 0.8 cm, 3.7 % of the minor radius. That case was judged marginal, with a high likelihood of losing vertical control [11].

12.3.4 Radial Position, Limits and Disturbances

The symmetric coil current follows $L_c\dot{I}_s = V_s - R_c I_s$ and the radial position follows it through a first-order lag, $\tau_R\dot{R} = k_R I_s + R_{dist} - R$, with $k_R = 10\,\mathrm{mm/kA}$ and $\tau_R = 50\,\mathrm{ms}$. The coil currents $I_s \pm I_a$ are limited to 6 kA by scaling after each step, and the plasma is lost, with the simulation frozen, when $|z|$ reaches 150 mm. The magnetics add nearly white noise of 0.05 mm to z and half as much, correlated over 0.5 ms, to R. A kick of size Δz, which represents an edge localized mode (ELM) or a sawtooth crash, changes the wall current by $-k_d\Delta z/g_w$ at once, so the plasma is displaced instantly while the force balance still holds. A minor disruption cuts the margin by 40 % for 50 ms with a 4 mm kick, and a β drop moves the plasma and its radial equilibrium 15 mm inward.

The parameters are listed in Table 12.2.

Table 12.2 Plasma, Wall and Coil Parameters, With Hidden Ones Marked

Symbol	Meaning	Value	Class
τ_w	wall time constant	5.0 ms	visible
L_w	wall mode inductance	10 µH	hidden
g_w	wall force per ampere	1.0 N/A	hidden
g_c	coil force per ampere	0.8 N/A	visible
L_c, R_c	coil inductance, resistance	20 µH, 1.0 mΩ	visible
k_M	wall-coil coupling	0.30	visible
V_{max}	supply voltage limit	10 V	visible
—	coil current limit	6 kA	visible
δ	control and supply latency	100 µs	visible
k_R	radial shift per ampere	10 mm/kA	hidden
τ_R	radial response	50 ms	visible
z_{wall}	distance to first wall	150 mm	visible

12.3.5 Assumptions

The plasma is rigid and massless. Real plasmas deform as they move and the vertical mode couples to the shape, so a rigid-displacement margin can differ from a deformable one [4]. The massless idealization is adequate for stabilization analysis under conditions that have been stated [2].

One wall mode and one coil circuit represent the vessel and the coils. A real vessel has many eddy-current patterns, and the faster ones add delay to the coil action [9].

The margin is a fixed linear function of elongation. On a real machine it depends on the current profile, the pressure and the wall geometry, which is what a minor disruption changes.

The radial loop is a single lag coupled to the vertical loop only through the shared supply. The magnetics carry additive noise only, without pickup from the coils themselves or calibration drift.

12.4 Implemented Controller

Control executes each 100 µs and the equations are integrated with a 20 µs fourth-order Runge-Kutta step on the engine of Section 1.6.1. The voltages reach the coils after the latency δ through a delay line.

12.4.1 Position and Velocity Loop

The vertical loop uses the shared velocity-form PID of Section 1.6.2 on the measured position z_m, with the velocity term supplied as a live output bias recomputed at every scan, Equation (12.6).

$$\begin{aligned} e_k &= z_{SP} - z_{m,k}, \qquad s_k = s_{k-1} + K_p\,(e_k - e_{k-1}) \\ c_k &= s_k + I_{k-1} - K_d\,\hat{\dot{z}}_k \\ V_{a,k} &= \mathrm{clamp}\,(c_k, -V_{max}, V_{max}) \\ I_k &= I_{k-1} + \frac{K_p}{\tau_I}\,e_k\,\Delta t_c + \frac{\Delta t_c}{\tau_I}\,(V_{a,k} - c_k) \end{aligned} \tag{12.6}$$

where s_k accumulates the proportional increments without limit, c_k is the demanded voltage, I is the integral state, used at scan k as accumulated through scan $k-1$, and $\Delta t_c = 100\,\mu s$. In the symbols of Equations (1.13) to (1.16), s_k is S_k with no derivative increment, $-K_d\hat{\dot{z}}_k$ is the output bias u_b, c_k is v_k and $V_{a,k}$ is u_k. Anti-windup is the back-calculation term with a tracking time equal to τ_I (Section 1.6.2). Away from the voltage limit the recursion equals the positional law $V_a = K_p e - K_d\hat{\dot{z}} + (K_p/\tau_I)\int e\,dt$. The loop is reverse acting in the ISA sense, $a = +1$ in the convention of Section 1.3, because a positive antisymmetric voltage pushes the plasma up, so a plasma above the setpoint receives a negative voltage. The defaults are $K_p = 2.0\,\mathrm{V/mm}$ and $K_d = 1.0\,\mathrm{V}$ per m/s, and the integral is off, which the

interface represents by an integral time of 200 ms. The integral update and its back-calculation are then skipped. Transfer between manual and automatic is bumpless as Section 1.6.2 describes.

12.4.2 Two Velocity Estimates

The differentiator estimate is the first difference of z_m passed through a first-order filter of $\tau_f = 0.3\,\text{ms}$. The model observer avoids differentiating noise. It reads the wall current from the force balance, $\hat{I}_w = -(k_d z_m + g_c I_a)/g_w$, and computes the drift velocity from the circuit equations with the applied voltage set to zero, Equation (12.7).

$$\hat{\dot{z}} = -\frac{g_w \dot{\hat{I}}_w + g_c \dot{I}_a}{k_d}, \qquad [\dot{\hat{I}}_w, \dot{I}_a] = F\,[\hat{I}_w, I_a]^T \tag{12.7}$$

The constants in the force balance and in F are the declared values at the selected elongation, so a plant changed in the parameters, a faster wall or a minor disruption is not followed by the observer. Leaving out the instantaneous response to the coil's own voltage keeps the velocity term from chasing its own output through the latency. A counterfactual observer that includes it, run on the same engine with the noise at zero, sustains an oscillation after an ELM for velocity gains above about 1.3 V per m/s.

12.4.3 Radial Loop and Allocation

The radial loop is a PI on R_m with $K_c = 1.0\,\text{V/mm}$ and $\tau_I = 60\,\text{ms}$ setting V_s. The upper and lower coils receive $V_s \pm V_a$, and two allocations are offered, Equation (12.8).

$$\begin{aligned} \text{clip:} \quad & V_u = \text{clamp}(V_s + V_a),\; V_l = \text{clamp}(V_s - V_a) \\ \text{priority:} \quad & V_a' = \text{clamp}(V_a) \\ & V_s' = \text{clamp}\big(V_s, \pm(V_{max} - |V_a'|)\big) \end{aligned} \tag{12.8}$$

Clipping each coil on its own lets a saturated radial demand pin both supplies at the same limit, which reduces the antisymmetric voltage, half the difference of the two, to a fraction of what the vertical loop asks for. Priority allocation serves the vertical demand first and gives the radial loop what remains, on the reasoning that a lost vertical position is the costlier failure [5]. Sharing a request among actuators that each have a hard limit is the subject of Volume 4, Chapter 13. The back-calculation of each loop sees only its own $\pm V_{max}$ clamp, not the allocation or the latency that follow it, so an active integral, such as the radial one while priority allocation cuts V_s back, can keep winding while the coils receive less than its loop demanded.

12.4.4 Explicit Model Predictive Control

A second vertical controller can replace the position and velocity loop in automatic mode, while the radial PI and the allocation stay as they are. It is a model predictive controller (MPC) on the vertical channel whose quadratic program has been solved in advance, unlike the online MPC option of Section 1.6.3. Because the program's linear term and constraint limits move affinely with the measured state and the box, its minimizer is an affine function of them on each of a finite number of polyhedral regions, one for each optimal combination of active limits, and a decision reduces to finding the region and evaluating its law [14]. An explicit MPC of this kind has been proposed for the vertical stabilization of ITER [15], and Volume 7, Chapter 6 uses one on a suspended ball. The loop remains the default, and only case 11 selects the MPC.

The controller's model is Equations (12.1) to (12.3) with the declared constants of Table 12.2 and the stiffness k_d of the elongation the operator sets, without the margin drop of a minor disruption, and it is rebuilt whenever the elongation changes. Eliminating I_w with the force balance writes it in the measured position and coil current, $x = [z, I_a]^T$, as $\dot{x} = F_z x + G_u(V_a + d_c) + G_w d_w$, where d_w and d_c are constant voltage errors on the wall and coil circuits. The declared latency equals one control period,

so the voltage applied over the next 100 µs is the command u_p decided at the previous execution, and u_p joins the state. An exact zero-order hold over $T_s = \Delta t_c = 100\,\mu\text{s}$ gives Equation (12.9).

$$\xi_{k+1} = A\,\xi_k + B\,u_k, \qquad A = \begin{bmatrix} \Phi & \Gamma_u \\ 0 & 0 \end{bmatrix}, \qquad B = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \tag{12.9}$$

where $\xi = [z, I_a, u_p]^T$, $\Phi = e^{F_z T_s}$ and Γ_u is the held voltage's column. At $\kappa = 1.8$ the channel's roots are $+356$ and $-77\ \text{s}^{-1}$, the first being the growth rate of the application. The model is written from the chapter's own linear equations rather than identified, and an independent derivation from the simulator's circuit function reproduces A to 10^{-12} and the controller gain below to a relative 3×10^{-10} at $\kappa = 1.8$.

Each execution plans three single-period voltage moves, a 0.3 ms horizon, and applies the first, as shown in Equation (12.10).

$$\begin{aligned} &\min_{u_0,u_1,u_2} \sum_{k=1}^{2} Q_z\,\tilde{z}_k^2 + \tilde{\xi}_3^T P\,\tilde{\xi}_3 + R\sum_{k=0}^{2}\tilde{u}_k^2 \\ &\text{subject to } \tilde{\xi}_{k+1} = A\tilde{\xi}_k + B\tilde{u}_k, \quad \tilde{\xi}_0 = [z_m, I_a, u_p]^T - \xi_t \\ &\qquad -b \le u_t + \tilde{u}_k \le b, \quad k = 0,1,2 \end{aligned} \tag{12.10}$$

where the tilde marks a deviation from the target $\xi_t = [z_{SP}, I_{a,t}, u_t]^T$, b is the half-width of the voltage box and P is the solution of the discrete algebraic Riccati equation for A, B, $\text{diag}(Q_z, 0, 0)$ and R. Charging the infinite-horizon cost after the last move, without a terminal set, is the usual stabilizing choice for a receding horizon [16]. The weights follow Bryson's rule on the 0.5 mm band of the vertical loop and the declared 10 V supply, $Q_z = 4 \times 10^6\,\text{m}^{-2}$ and $R = 0.01\,\text{V}^{-2}$, and were fixed before any comparison was run. With no limit active the first move is the LQR's, $\tilde{u}_0 = -K\tilde{\xi}_0$, with $K = (11.5\,\text{V/mm}, 0.0104\,\text{V/A}, 0.785)$ at $\kappa = 1.8$ and $(9.90\,\text{V/mm}, 0.0110\,\text{V/A}, 0.891)$ at 1.9, so its position gain is five to six times the loop's default K_p and the third entry feeds back the voltage still on its way. The closed loop of $A - BK$ has a spectral radius of 0.958.

The only limits are the box on every move, all hard, and because only inputs are constrained the program is feasible whenever $b \ge 0$. The box is what the allocation leaves the vertical channel after the radial command V_s of the same scan. Under priority allocation $b = V_{max}$. Under clipping it is the largest antisymmetric voltage the coil-by-coil clip can deliver with a command inside $\pm V_{max}$, $b = V_{max} - |V_s|/2$, and a planned voltage u is commanded as u while $|u| \le V_{max} - |V_s|$ and as $\text{sign}(u)(2|u| - V_{max} + |V_s|)$ beyond, which puts one coil at its limit and delivers exactly u. On 20 000 random supplies and radial commands the clip's reach equals $V_{max} - |V_s|/2$ to 7×10^{-15} V, and in closed loop the antisymmetric voltage reaching the coils one period later equals the plan's first move to 7×10^{-11} V. A first version that gave the MPC $V_{max} - |V_s|$ under clipping was lost in four clipped cases that both forms of the loop hold, because it claimed less than the clip itself delivers. The clip, the priority rule, the coil current limit and the first wall act on the MPC's command exactly as on the loop's.

The parameter of the program is $\theta = (z_m - z_{SP},\, I_a - I_{a,t},\, u_p - u_t,\, -b - u_t,\, b - u_t)$, so the box enters only the right-hand sides and one table serves every supply limit, radial share and allocation rule. Three moves, each free or at one of two edges, give $3^3 = 27$ combinations, and all 27 are stored at every elongation from 1.50 to 2.02. The table is complete, since every θ with a nonempty box lies in one of the regions, so a state leaves it only through a number that is not finite. Restricting the pending voltage to one the supply could have delivered, 21 regions have an interior on the default plant at every elongation (21 to 23 over every parameter setting), and 13 were met in 112 321 closed-loop decisions. Each lookup checks the stored regions in sequence, starting with the unconstrained law, against the signs of its multipliers and the six box rows. The row test is relative, 10^{-12} times the largest of 1, $|b|$

and the unconstrained moves, because far from the midplane the unconstrained move reaches 10^3 V, and a fixed 10^{-12} V test misses 2 of those 112 321 decisions. On 199 321 states, 87 000 with a limit active and 15 000 at region boundaries, the lookup agrees with the apps' shared dual active-set solver, run online on the identical program, to 1.2×10^{-11} V in the first move. Rebuilding the design and the table for a new elongation takes 0.38 ms at the median in Node.js, inside the execution that sees the change.

The MPC is offset-free by a target shift [17]. At each execution the measured z_m and I_a are compared with their one-step prediction, the two errors are mapped onto d_w and d_c through the disturbance columns of the discretized model, and each estimate is moved through a first-order filter of 50 ms, ten declared wall times, and clamped to ± 4 V and ± 5 V. The target is the declared equilibrium that holds the setpoint against them, Equation (12.11).

$$I_{w,t} = \frac{d_w}{R_w}, \quad I_{a,t} = -\frac{k_d z_{SP} + g_w I_{w,t}}{g_c}, \quad u_t = R_c I_{a,t} - d_c \tag{12.11}$$

With the coil force per ampere at 0.7 on case 9's 10 mm step and no noise, the MPC ends 0.000 mm from the setpoint and the loop with the differentiator 0.58 mm away.

A region not found, or an online solve that is not optimal, rejects the decision. The loop, which tracks the applied command at every execution, then takes the coils, its offset cleared and its voltage ramped from the last command to its own law over 20 decisions (2 ms), the record is marked, and after 20 decisions found in a row the MPC returns over another 2 ms. A switch to the loop during a run is handed over the same way. Misses forced with the noise off give steps of 0.27 V in the recovery from an ELM and 0.49 V at the reversal after a large kick, and no measured run missed. The plant integrates the run's parameters and the disruption's margin drop, while the MPC and the observer hold the declared model, so any parameter a learner edits becomes model mismatch for both. The MPC's estimator runs under every controller and writes nothing the loop reads, so adding the MPC changed no earlier run under the loop.

12.5 Results

Table 12.3 lists the twelve cases in the order the application presents them, with the point each teaches and what completes it. A case is credited when the learner has performed the listed actions and the simulation has reached the listed state, and the credits accumulate toward certificates of progress and of completion. The challenge is judged on the case as given, so a change to a model parameter or disturbance setting, or firing an extra disturbance, voids the attempt.

The values in Table 12.4 were read from runs at the default seed of the published app (Section 1.6.4), each case selected, changed and then run from $t = 0$. An rms position is the application's running value for the true z, sampled at every control scan from the start of a run of the stated length, except in case 5, where it covers the last 0.3 s of 0.6 s runs. Peaks and settling times are taken at the integration step. Rerunning each closed-loop configuration with its settings carried through a reset gives the same run. With other noise seeds the loss times of cases 3, 4 and 8 move by up to 47 ms. The suggested challenge settings pass at all six seeds tested, while priority allocation alone passes at three and fails at three.

12.5.1 Growth Rate and the Gain Floor

Held at the midplane with 0.05 mm of magnetics noise, the true position wanders 0.023 mm rms over 0.2 s, and nothing in the trends suggests an instability growing at 356 s^{-1}. With the coil voltage held at zero, a 0.1 mm displacement reaches the wall in 20.0 ms at elongation 1.8, 42.3 ms at 1.5 and 6.7 ms at 2.0, about seven e-folding times in each case. Choosing the elongation before Run or carrying it through Reset initializes the same 0.1 mm displacement.

Table 12.3 Scenarios and Completion Requirements

#	Scenario and kind	Teaching point	Required to complete
1	Held at the midplane *Experiment*	a quiet unstable loop	After the first 0.1 s, fire the 3 mm ELM kick, and the loop holds the plasma within 0.5 mm of the setpoint for 20 ms.
2	Let go (manual) *Demonstration*	the growth rate	With the coil voltage held at zero, let the plasma reach the first wall at $\kappa = 1.8$, then reset and repeat with κ set to 1.5 and to 2.0 before each run.
3	Too little gain *Demonstration*	the gain floor	Lose the plasma at $K_p = 0.3\,\mathrm{V/mm}$, then reset and run 0.2 s with the plasma held at 0.5 V/mm.
4	Stretch it *Demonstration*	the floor moves with κ	Run $K_p = 0.5\,\mathrm{V/mm}$ through the change from $\kappa = 1.8$ to 2.0 at 0.1 s until the plasma is lost or $t = 0.2$ s, then reset and run 0.2 s at 1.0 V/mm.
5	Too much gain *Experiment*	the ceiling from latency	With 300 µs of latency, run 0.1 s at $K_p = 10\,\mathrm{V/mm}$, then 50 ms at 5 and 50 ms at 2 V/mm, then 0.1 s at 10 V/mm with the latency set back to 100 µs.
6	Velocity from a noisy position *Experiment*	differentiator against observer	With 0.3 mm of noise and $K_d = 2$ V per m/s, run 0.1 s each on the differentiator, on the model observer, and with $K_d = 0$.
7	The recoverable region *Experiment*	the size of the region	Run past the 20 mm kick at 50 ms (to loss or 0.15 s) at $\kappa = 1.9$ as given, with $K_p = 10\,\mathrm{V/mm}$, at $\kappa = 1.8$, and at $\kappa = 1.9$ with $V_{max} = 15$ V, each change made before its run.
8	A faster wall *Demonstration*	the passive structure	Run $K_p = 0.6\,\mathrm{V/mm}$ through the halving of τ_w at 0.1 s until the plasma is lost or $t = 0.2$ s, then reset and run 0.2 s at 2 V/mm.
9	Move it up *Experiment*	offset and integral	Run to 0.15 s through the setpoint step to 10 mm at 50 ms with no integral action, then reset and run 0.15 s with $\tau_I = 20$ ms.
10	Two jobs, one supply *Experiment*	allocation	Run with clipped allocation past the β drop and kick until the plasma is lost or $t = 0.15$ s, then reset and run 0.15 s with vertical priority.
11	Explicit MPC against PD *Experiment*	a limit in the plan	Run under the explicit MPC through the ELM at 50 ms and the 19.7 mm kick at 0.15 s, then reset, select the loop before 50 ms and run the same case to 0.25 s.
12	Challenge at $\kappa = 1.95$ *Challenge*	graded run	Keep the plasma for 1.2 s, with $\lvert z - z_{SP}\rvert \le 4$ mm and $\lvert R - R_{SP}\rvert \le 1$ mm over the last 0.4 s, on the case as given.[a]

[a] As given, the plasma is lost at 661.0 ms. Priority allocation alone passes at the default seed, with 3.94 mm vertically and 0.80 mm radially over the last 0.4 s. The suggested settings ($K_p = 2\,\mathrm{V/mm}$, observer at $K_d = 1.5$ V per m/s, priority) pass with 3.23 mm and 0.18 mm. Before-Run and Reset settings give the same run. The verdict checks each 20 µs physics step and each kick instant.

> **TRY IT IN THE APP** **2 · Let Go (MANUAL)**
>
> Load scenario 2 and press Run. The loop is in manual with the coil voltage held at zero, and the plasma starts 0.1 mm above the midplane. Watch the vertical position grow faster and faster until the plasma reaches the wall at 20.0 ms. Press Reset, set the elongation to 1.5 before Run, and the same fall takes 42.3 ms.

The loop has a floor. At a position gain of 0.3 V/mm with no velocity term the plasma drifts away and is lost at 166.6 ms, a slow failure that looks like drift rather than instability. The floor at $\kappa = 1.8$, the gain at which a noise-free kick neither grows nor decays with the 100 µs latency, is about 0.45 V/mm. At 0.4 the plasma survives the first 0.2 s within 0.5 mm but is lost at 482 ms, and at 0.5 it holds to 0.023 mm rms over 0.2 s. The floor moves with the plasma. When the elongation is raised from 1.8 to 2.0 during a run at 0.5 V/mm, the growth rate triples to $1070\,\mathrm{s}^{-1}$, the floor rises to about 0.70 V/mm, and the plasma is lost 56 ms after the change, while 1.0 V/mm holds it to 0.023 mm rms. Halving the wall time raises the growth rate from 356 to $680\,\mathrm{s}^{-1}$ and the floor from about 0.45 to about 1.03 V/mm, so 0.6 V/mm loses the plasma at 158 ms and 2 V/mm holds it to 0.021 mm rms.

Pause. Try Exercise 12.1 before reading on.

Table 12.4 Measured Results

Case and condition	Result
1, noise 0.05 mm, 0.2 s	true z 0.023 mm rms
2, from 0.1 mm, κ 1.5 / 1.8 / 2.0	wall at 42.3 / 20.0 / 6.7 ms
3, $K_p = 0.3$	lost at 166.6 ms; floor $\approx$ 0.45 V/mm
3, $K_p = 0.4$	lost at 482 ms
3, $K_p = 0.5$, 0.2 s	0.023 mm rms
4, κ 1.8 to 2.0, $K_p = 0.5$	lost 56 ms later; floor $\approx$ 0.70 V/mm
4, $K_p = 1.0$, 0.2 s	0.023 mm rms
5, $\delta = 300\,\mu$s, K_p 10 / 5 / 2	1.45 / 0.063 / 0.022 mm rms
5, $K_p = 10$, saturation	64 % of the time
5, $\delta = 100\,\mu$s, $K_p = 10$	0.059 mm rms
6, noise 0.3 mm, 0.35 s, differentiator	0.216 mm, 33.9 J, 31,000 V/s
6, observer	0.116 mm, 9.1 J, 9,800 V/s
6, no velocity term	0.104 mm, 6.1 J
1, 3 mm ELM, K_d 0 / observer 1	overshoot 0.52 / 0.10 mm
7, $\kappa = 1.9$, 20 mm kick	margin 7 %, default gains lost at 67.9 ms
7, $K_p = 10$	lost at 58.9 ms
7, $\kappa = 1.8$	margin 28 %, held, peak 23.1 mm
7, $V_{max} = 15$ V	margin 38 %, held, peak 21.4 mm
8, τ_w halved, $K_p = 0.6$	γ 356 to 680 s^{-1}, lost at 158 ms
8, floor at $\tau_w = 2.5$ ms	$\approx$ 1.03 V/mm
8, $K_p = 2$, 0.2 s	0.021 mm rms
9, SP 10 mm, no integral	settles 0.51 mm high, peak 12.3 mm
9, $\tau_I = 20$ ms	within 0.1 mm in 41 ms (34 ms logged)
10, clip	lost at 72.7 ms
10, priority	held, peak 10.7 mm

12.5.2 The Gain Ceiling and the Velocity Estimate

With 300 µs of latency, a position gain of 10 V/mm settles into a limit cycle of 1.45 mm rms with the supply saturated 64 % of the time. The plasma is still held, with no reserve. At 5 V/mm it holds to 0.063 mm and at 2 V/mm to 0.022 mm, and cutting the latency to 100 µs lets 10 V/mm hold to 0.059 mm. The window between the floor and the ceiling, 0.45 V/mm to somewhere between 5 and 10 V/mm at this latency, is what the loop is designed within. For the reduced mode, Equation (1.11) shows the same fall of the ceiling as the latency grows.

🕮 **Pause.** Try Exercise 12.2 before reading on.

With the noise raised to 0.3 mm, a velocity gain of 2 V per m/s on the filtered differentiator leaves the position at 0.216 mm rms over a 0.35 s run, moves the voltage about 31 000 V/s summed over the control scans and delivers 33.9 J to the coils. The model observer gives 0.116 mm, about 9800 V/s and 9.1 J. With no velocity term at all the position is 0.104 mm rms for 6.1 J, so on noise alone the velocity term does not pay. It pays on kicks. With the noise at zero, a 3 mm ELM overshoots by 0.52 mm and settles within 0.15 mm in 5.3 ms without a velocity term, and overshoots by 0.10 mm and settles in 1.5 ms with the observer at a velocity gain of 1.

12.5.3 The Recoverable Region

At $\kappa = 1.9$ a 20 mm kick leaves a recoverable margin of 7 %. The voltage saturates and the plasma is lost at 67.9 ms, and raising the gain to 10 V/mm only loses it sooner, at 58.9 ms. Lowering the elongation to 1.8 cuts the growth rate from 542 to 356 s^{-1}, the margin after the same kick is 28 %, and the plasma is held with a peak of 23.1 mm. Raising the supply to 15 V at $\kappa = 1.9$ enlarges the region instead, to a margin of 38 %, and also holds it, with a peak of 21.4 mm. The kick at $\kappa = 1.9$ lands inside the region, so the tuning still decides whether the plasma is kept: a stronger velocity term, which reverses the voltage earlier, also holds it on the 10 V supply, and so does a weaker position gain. Gain decides how the loop behaves inside the region, while the growth rate and the supply voltage decide how large the region is, which is the reasoning behind specifying a maximum controllable displacement for a machine [1, 11].

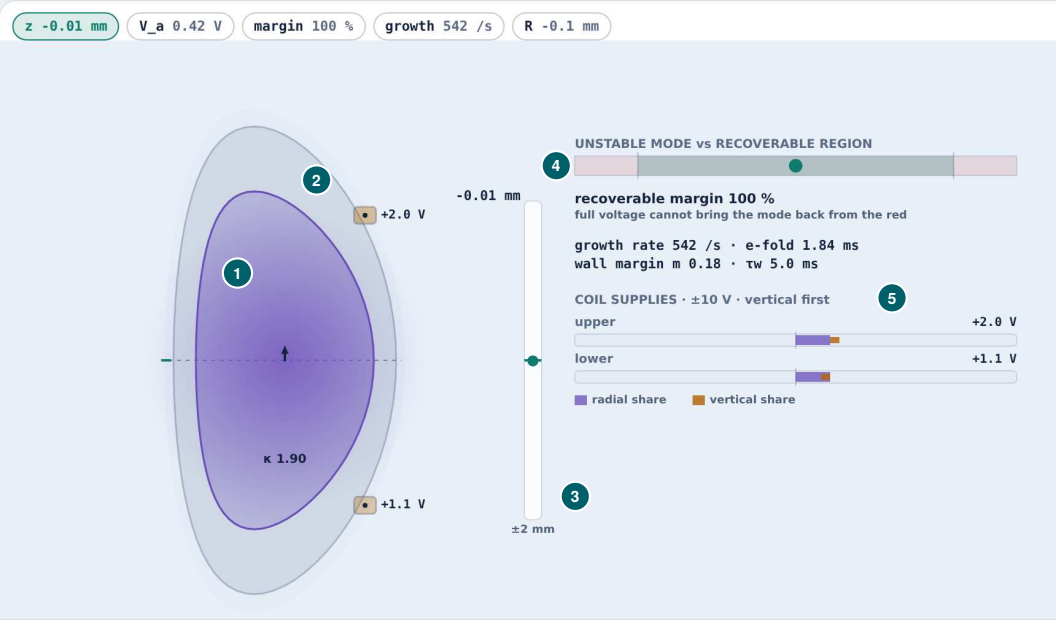


Figure 12.2 The vertical stabilization application running scenario 10 with priority allocation set before the run, at $t = 0.25$ s, 200 ms after a β drop and 195 ms after a 10 mm vertical kick at elongation 1.9, cropped to the plant view. The cross-section shows the vessel and the plasma at its elongation (1), the two in-vessel coils with their voltages (2) and a zoomed ±2 mm position gauge (3). Beside it the unstable mode is drawn against the recoverable region (4), above the growth rate of 542 s^{-1} and wall margin of 0.18, and the two coil supplies are split into radial and vertical shares (5). The plasma is held, back at the midplane.

> ✦ DID YOU KNOW? **More Gain, Lost Sooner**
>
> From the simulator: at elongation 1.9 a 20 mm kick leaves a recoverable margin of 7 %. The default gains lose the plasma at 67.9 ms, and a position gain of 10 V/mm, five times the default, loses it sooner, at 58.9 ms (Table 12.4).

12.5.4 Offset and Allocation

Moving the setpoint 10 mm up shows why vertical controllers carry a slow integral or a separate position loop. Holding the plasma off the midplane takes about −0.96 kA in the coils and −1.0 V across their resistance, so the proportional loop settles 0.51 mm above the setpoint after a peak of 12.3 mm. A 20 ms integral time brings the true position within 0.1 mm of the setpoint for good 41 ms after the step, while the trend logged at 2 kHz is inside that band from 34 ms, the noise carrying the plasma to the edge of the band in between.

When a β drop throws the plasma 15 mm inward and a 10 mm vertical kick follows 5 ms later at $\kappa = 1.9$, the radial loop asks for more than the supply has. With each coil clipped on its own, the plasma is lost at 72.7 ms, by which time the radial position has recovered only from 15 to 4.9 mm inward. With the vertical demand served first it is held with a peak of 10.7 mm (Table 12.4; Figure 12.2 shows the plasma held), and the radial position is back within 1 mm for good about 107 ms after the drop. The recoverable margin reads 53 % at the kick in both cases, because it assumes the vertical loop has every volt.

> **TRY IT IN THE APP** **10 · Two Jobs, One Supply**
>
> Load scenario 10 and press Run. At 50 ms a β drop throws the plasma inward and a 10 mm vertical kick follows. With each coil clipped on its own the plasma is lost at 72.7 ms. Press Reset, set the coil voltage allocation to Vertical first before Run, and the plasma is held with a peak of 10.7 mm.

12.5.5 The Graded Challenge

The challenge runs 1.2 s at $\kappa = 1.95$ with 0.2 mm of noise, two ELMs, a minor disruption, a β drop with a 10 mm kick and two more ELMs, starting from $K_p = 0.8\,\mathrm{V/mm}$, a differentiator velocity gain of 3 and clipped allocation. It is met when the plasma is kept and, over the last 0.4 s, the vertical position stays within 4 mm and the radial position within 1 mm. The verdict checks every 20 µs physics step and each instantaneous kick, so brief ELM excursions count even when they fall between the 2 kHz log samples. As loaded the plasma is lost at 661.0 ms, 11 ms after the β drop. Priority allocation alone keeps the plasma and meets the target at the default seed, with 3.94 mm vertically and 0.80 mm radially, although the differentiated noise saturates the vertical supply 22 % of the logged time. It passes at two of five additional seeds and fails at the other three, including a 5.02 mm vertical excursion at seed 3. Every tested combination of the observer, a velocity gain of 1.5 and $K_p = 2\,\mathrm{V/mm}$ that keeps clipped allocation loses the plasma within 0.3 ms of the original time. With $K_p = 2\,\mathrm{V/mm}$, the observer at a velocity gain of 1.5 and priority allocation, the last 0.4 s stay within 3.23 mm and 0.18 mm, including the ELM kicks. The supply saturates 0.1 % of the logged time and the position is 0.34 mm rms over the whole run. These suggested settings pass at all six seeds tested, with worst vertical errors of 3.17 mm to 3.33 mm.

12.5.6 Explicit MPC Against the Position and Velocity Loop

Case 11 holds the plasma of case 7 at $\kappa = 1.9$ under the explicit MPC, fires a 3 mm ELM at 50 ms and a 19.7 mm kick at 150 ms, a size chosen between the largest kicks the MPC and the loop recover. It is credited when the case has been run under the MPC and then, after Reset, under the loop selected before the ELM. After the ELM the MPC is back within 0.5 mm of the setpoint in 0.42 ms, against 1.02 ms for the loop with the differentiator and 0.62 ms with the observer, read at the control scans. The large kick leaves a recoverable margin of 7 % under each controller, and each holds the full −10 V for 7.5 to 7.8 ms. The MPC's command turns positive 7.82 ms after the kick with the plasma 3.8 mm from the midplane, the loop's at 7.72 ms with the plasma at 7.7 mm. Within the next 4 ms the MPC's unstable mode swings to 1.14 times ξ_{max} on the other side, past the edge of the recoverable region, and the plasma is lost at 168.0 ms, while the loop's mode reaches 0.93 times ξ_{max} and the plasma is held with a largest excursion of 34.2 mm, as shown in Figure 12.3.

Every row of Tables 12.5 to 12.7 uses one engine, noise seed, case script and duration, the loop at each case's own gains with each velocity estimate and both allocations, and the MPC at case 11's settings with nothing retuned per case. Two further rows isolate the MPC's ingredients. The loop retuned once by the MPC's own rule takes the LQR's gains on z and on the declared drift velocity at the elongation in force ($K_p = 8.55\,\mathrm{V/mm}$ and $K_d = 2.02\,\mathrm{V}$ per m/s at $\kappa = 1.9$), without the latency term, inside the MPC's box. The MPC with no box rows is the same law, estimator and target with the box removed from the program and the first move clipped afterwards, the strongest classical form of the MPC's own gain.

The MPC loses the headline comparison, the largest recoverable kick. At $\kappa = 1.9$ it recovers kicks up to 19.63 mm with clipping against 19.91 mm and 19.86 mm for the loop with each estimate, and at 1.8 up to 25.15 mm against 25.65 and 25.50 mm, as listed in Table 12.5 and drawn in Figure 12.4. The loop retuned to the LQR's gains recovers 19.63 mm and 24.81 mm with clipping (bisected to 0.05 mm). The plan sees 0.3 ms, and its terminal cost prices the motion after that as if any voltage were available, so the MPC holds −10 V until the plasma is nearer the midplane and the latency carries the swing past the region's far edge. Its knowledge of the box is worth 0.05 to 0.13 mm of recoverable kick over the same law without the rows. Held over a 2.1 ms horizon the same 27-region

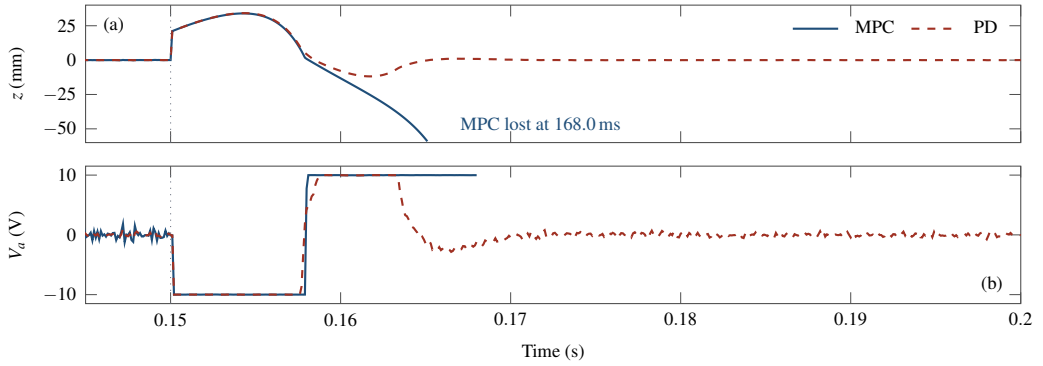


Figure 12.3 Case 11 to $t = 0.200$ s, run first with the explicit MPC and then, after Reset, with the loop and its differentiator. (a) Vertical position. (b) Antisymmetric voltage at the coils. The MPC is the solid line and the loop the dashed line. After the 19.7 mm kick at the dotted line both hold −10 V. The MPC reverses later and closer to the midplane, and its plasma crosses to the wall on the other side, while the loop's settles.

Table 12.5 Largest Recoverable Kick (mm) at 50 ms, Default Noise

κ, allocation	PD, diff.	PD, obs.	MPC	No rows	Held[a]
1.9, clip	19.91	19.86	19.63	19.58	20.19
1.9, priority	19.92	19.87	19.68	19.60	20.19
1.8, clip	25.65	25.50	25.15	25.03	26.27
1.8, priority	25.67	25.52	25.21	25.08	26.31

Bisection to 0.01 mm on case 7 without its scripted kick, 0.25 s runs. "No rows" is the MPC's law without its box rows, the first move clipped.
[a]Counterfactual: the same three moves held over 0.1, 0.4 and 1.6 ms, a 2.1 ms horizon, not offered in the application.

structure sees the braking, recovers 20.19 mm at $\kappa = 1.9$ and holds case 11, and over 5.1 ms it reaches 20.35 mm. No controller tested recovers more than about 20.4 mm at $\kappa = 1.9$ on a 10 V supply, so the supply's voltage limit decides how large a displacement can be recovered under every controller, and the look-ahead decides only how much of that region is used. Halving the integration step changes none of the rankings (19.91, 19.86 and 19.63 mm).

🕮 **Pause.** Try Exercise 12.3 before reading on.

The MPC is the fastest controller on every small kick, as listed in Table 12.6. After the ELM it is back within 0.5 mm in 0.52 ms against 1.32 and 1.02 ms, at 1.4 J of coil energy against 0.9 and 1.2 J, after the minor disruption in 0.52 ms against 1.52 and 0.82 ms, and at $\kappa = 1.95$ in 0.32 ms against 1.02 and 0.42 ms. The no-rows law gives the same numbers on all of these cases and on the noise and the challenge, so the difference from the loop is the LQR gain with the latency in its model, not the knowledge of the limit. The loop given the LQR's gains without that latency term limit-cycles at $\kappa = 1.95$, saturated for 76.2 ms of the 0.3 s run. The MPC loses the noise case. Under 0.3 mm of magnetics noise its position is 0.241 mm rms against 0.119 mm for the loop with the observer, for 87.9 J against 16.0 J and 14.6 ms saturated against none, because its position gain of 11.5 V/mm passes the noise to the coils. Taken at every integration step instead of the scans, the same runs give 0.212 against 0.117 mm and 81.5 against 16.0 J. In case 10 with clipping every controller is lost, the MPC 0.3 ms sooner, and with priority every controller holds a peak near 10.7 mm. In the challenge under priority allocation the MPC meets the target with 3.17 mm vertically against 3.23 mm for the suggested loop, at 27 % more coil energy (5082 J against 3991 J).

The worst-case table runs 53 cases under each controller and both allocations, as summarized in Table 12.7. The MPC loses 18 of 106 runs against 24 and 26 for the loop. It holds 8 runs a form of the loop loses, cases 3, 4 and 8 under both allocations, where the shipped low gains are the lesson

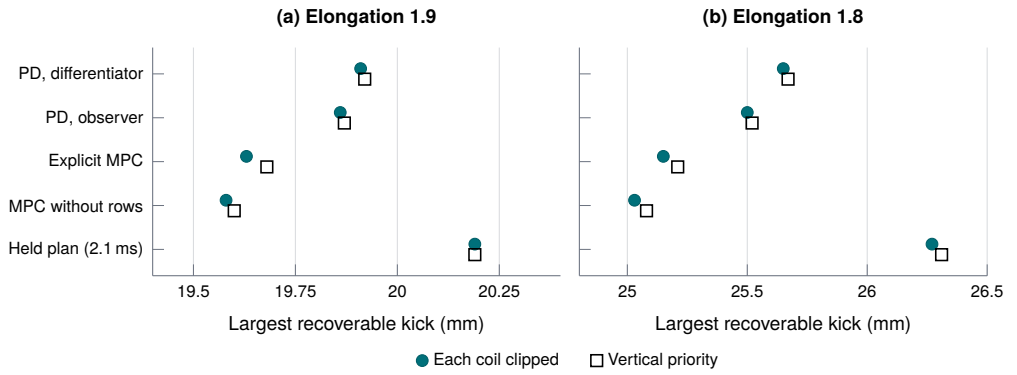


Figure 12.4 The largest kick at 50 ms that each controller recovers, from Table 12.5, at elongation 1.9 (a) and 1.8 (b). Filled circles: each coil clipped. Open squares: vertical priority. The last row is the counterfactual plan held over 2.1 ms, which the application does not offer. Notice that the explicit MPC recovers less than either form of the loop, that removing its box rows costs only a little more, and that the allocation barely matters here: the look-ahead, not the knowledge of the limit, decides how much of the region is used.

Table 12.6 Explicit MPC Against the Loop on the Shipped Cases

Case	Measure	PD, diff.	PD, obs.	MPC	PD, LQR	No rows
1, 3 mm ELM, no noise	back within 0.5 mm (ms)	1.32	1.02	0.52	1.02	0.52
	peak $\lvert V_a \rvert$ (V) / coil energy (J)	10.00 / 0.9	7.69 / 1.2	10.00 / 1.4	10.00 / 3.2	10.00 / 1.4
1, minor disruption, no noise	back (ms) / saturated (ms)	1.52 / 0.3	0.82 / 0.1	0.52 / 0.5	50.22 / 20.6	0.52 / 0.5
1 at $\kappa = 1.95$, ELM, no noise	back (ms) / saturated (ms)	1.02 / 0.2	0.42 / 0.0	0.32 / 0.3	199.92 / 76.2	0.32 / 0.3
6, noise 0.3 mm, 0.6 s, clip	rms (mm) / energy (J)	0.218 / 58.7	0.119 / 16.0	0.241 / 87.9	0.428 / 152.5	0.241 / 87.9
	saturated (ms)	0.5	0.0	14.6	68.4	14.6
10, clip	lost at (ms)	72.7	72.7	72.4	74.7	72.1
10, priority	back (ms) / peak (mm)	61.7 / 10.70	56.9 / 10.67	56.7 / 10.71	280.0 / 10.77	56.7 / 10.71
Challenge, priority[a]	last 0.4 s, z / R (mm)	3.94 / 0.80	3.23 / 0.18	3.17 / 0.20	4.39 / 0.62[b]	3.17 / 0.20
	coil energy (J) / saturated (ms)	11,415 / 270.9	3,991 / 2.2	5,082 / 2.8	11,380 / 288.3	5,082 / 2.8
11	ELM back (ms) / 19.7 mm kick	1.02 / held	0.62 / held	0.42 / lost 168.0	48.72 / lost 165.8	0.42 / lost 166.2

"Back" runs from the event until $\lvert z - z_{SP} \rvert$ stays within 0.5 mm, read at the 10 kHz control scans, "saturated" counts control periods with the vertical voltage at its limit or cut by the allocation (under the MPC also a plan on its box), and the rms is the application's running value over the run. Every MPC run decided every 100 µs with no rejected decision, and the online solve reproduces every MPC entry to the printed digit. "PD, LQR" is the loop with the LQR's gains and "No rows" the MPC's law without its box rows. [a]The differentiator column is the shipped tuning ($K_p = 0.8\,\mathrm{V/mm}$, $K_d = 3$) and the observer column the suggested one ($K_p = 2$, $K_d = 1.5$). With clipping every controller loses the plasma at 660.5 to 661.2 ms. [b]Fails the 1 mm radial target.

and the MPC's gain is fixed by its weights, and case 1 with a 1 ms wall, which the observer built on the declared 5 ms wall loses at 33.3 ms. In 12 runs every controller is lost and the MPC 0.2 to 2.4 ms sooner (case 7, case 10 and the challenge with clipping, and a learner's $g_w = 3$, $g_c = 0.1\,\mathrm{N/A}$, $L_c = 100\,\mu\mathrm{H}$ or $k_M = 0.8$). At the ends of the parameter ranges it is never lost where the loop holds. Between the ends, a sweep of 74 mismatched runs finds three where it is: case 9's 10 mm step with $g_c = 0.6\,\mathrm{N/A}$, $k_M = 0.5$ or $\tau_w = 2\,\mathrm{ms}$ is lost by the MPC at 66.1, 62.9 and 57.7 ms and by the loop with the observer, which holds the same declared model, while the loop with the differentiator holds it within 0.69 mm. The application states both the noise loss and these mismatch losses beside the controller.

In Node.js a lookup takes 0.08 µs at the median closed-loop decision and the online solve 7.2 µs.

Table 12.7 Worst Case Over 53 Cases and Both Allocations (106 Runs Each)

Output	PD, diff.	PD, obs.	MPC
runs lost	24	26	18
largest $\lvert z\rvert$ in a run held (mm)	24.66	23.07	24.28
largest coil current (kA)	6.00	6.00	6.00
longest at the coil current limit (ms)	22.4	22.4	20.7
longest at $\pm V_{max}$, run held (ms)	380.4	380.4	365.8
rejected MPC decisions	—	—	0

The ten scenarios and the challenge as shipped, the challenge with its suggested gains, each disturbance and a combined sequence, all 13 parameters at both ends of their ranges, elongations 1.5 and 2.02, 1 mm of noise, setpoint steps of 20 mm vertically and 30 mm radially, case 10 at $\kappa = 2.02$ and case 7 at 15 V.

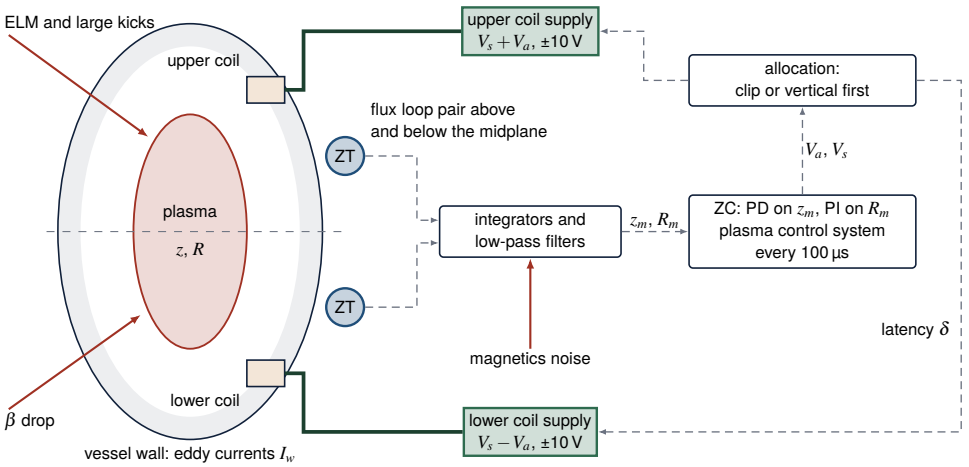


Figure 12.5 The loop of Figure 12.1 with the hardware of a tokamak. A pair of flux loops above and below the midplane (ZT) feeds integrators and filters, and the plasma control system ZC computes the position and velocity law and the radial PI every 100 µs. The allocation splits the two demands between the upper and lower coil supplies, fast converters limited to ±10 V, after the latency δ. Thick lines carry coil power, thin dashed arrows signals and solid arrows disturbances.

In the Chrome browser, on the application's own frame loop at 1×, the top speed of 10 000 decisions a second, the whole controller execution takes 1.1 to 1.3 % of one core under the lookup, 8.1 to 8.2 % under the online solve and 0.6 to 0.7 % under the loop, whose execution also runs the MPC's estimator. The host carried other computations while these times were measured, and the online figure lies within a factor of two of a tenth of a core.

12.6 Technology in Practice

The magnetics ZT, the controller ZC and the coil voltages of Figure 12.1 stand, on a tokamak, for magnetic sensors on the vessel, a real-time plasma control computer and a pair of coil power supplies. Figure 12.5 draws the loop with that hardware in place of the schematic's symbols.

What Measures It

No instrument touches the plasma. Its position is inferred from magnetic sensors and from the currents in the active coils. A learned controller on one research tokamak read 34 flux loops, 38 magnetic probes and 19 coil currents at 10 kHz, through a filtering stage modeled as a time delay with Gaussian noise [18]. On another machine the position from flux loop pairs fed the proportional term of a PID controller and their voltages the derivative term, and the integral term was never used

in its first operation [19], as in the chapter's default. The chapter's ZT is the true position plus white noise, and its latency δ stands for every delay from plasma to coil.

> **TECHNOLOGY Sensor: Flux Loops and Magnetic Probes**
>
> A tokamak never sees its plasma directly. Wire loops and small coils on the vessel pick up the changing magnetic field as an induced voltage. Pairs of flux loops above and below the midplane give the position: their flux difference is zero for a centered plasma and changes sign with the direction of a shift. The flux is the time integral of the measured voltage, and the voltage itself gives the velocity [19]. Integration brings drift. A repeatable linear drift is subtracted as a known slope and offset, while drift that does not repeat looks like slow plasma motion [20]. The coils' own fields must be corrected for, and one estimator run without those corrections erred by 1 cm or more in 38 % of its samples [19]. The chapter's ZT adds white noise only.

What Moves It

Coils whose radial field pushes the plasma up or down act together with eddy currents in the passive structure, and a 2010 design for ITER adds a pair of coils inside the vessel [12]. On one machine, adding two inboard coils to an outboard set roughly doubled the recoverable displacement. That displacement grows in proportion to the saturation voltage and falls roughly as the inverse of the growth rate, and a supply fast against the growth time makes its own dynamics unimportant [11]. Equation (12.5) carries the same two dependences for the chapter's two coils.

> **TECHNOLOGY Final Element: The Coil Power Supply**
>
> Fast power converters drive the coils. The thyristor supplies of one research tokamak were modeled for control design as a fixed bias and a fixed time delay identified from data. When a coil current changes polarity under too low a voltage request, such a supply can stick at zero output current for an extended time [18]. The supply planned for ITER's in-vessel vertical coils is a high-power chopper with energy storage and a 1 ms response. Two units are interleaved in series with the upper and lower coils, which are connected in anti-series, and every coil turn is brought out of the vessel so that a faulted turn can be left out of the circuit [12]. The chapter's supply is a ±10 V clamp behind a 100 µs delay line.

Where the Controller Runs

A plasma control system runs the loop on real-time computers. On the research tokamak above a compiled control policy ran at 10 kHz, the 100 µs period of the chapter's controller, after a conventional controller had formed the plasma [18]. ITER's plasma control system follows a pulse schedule, a supervisory layer changes the control in real time on plasma or plant events, and a central interlock system backs it up with simpler, robust algorithms. Every algorithm is first tested on a simulation platform [21]. The chapter's loop is one PID block in such a system, with the allocation rule between it and the supplies.

> **IN PRACTICE Size the Region Against the Noise**
>
> Vertical control is sized against noise. On operating tokamaks the standard deviation of the measured vertical position typically runs at 0.5 to 1 % of the minor radius, from power supplies, instruments, aliasing, crosstalk and plasma activity. Control turned marginal when the maximum controllable displacement was only 5 to 8 times that noise, and was lost near twice it [11]. From such statistics ITER took the empirical rule that the controllable displacement exceed a tenth of the minor radius, and operating machines keep it two to three times the largest expected perturbation [22].

What Is Modernizing It

Reinforcement learning has produced one controller that commands all of a tokamak's control coils. Trained on a simulator and run without further tuning, it held an elongation of 1.9 with a vertical growth rate of 1.4 kHz [18]. Data-driven estimators fit many signals at once: a linear model of 140

real-time magnetics signals reached a coefficient of determination of 0.906 on equilibria from 727 discharges [19]. And control systems are being designed to forecast approaching stability limits with physics-based and machine-learning predictors, so that a loss of vertical control can be mitigated [21]. The chapter's model observer is the simplest model-based estimate of this kind.

12.7 Limitations

The rigid massless plasma omits deformation, which matters at high elongation and for the definition of the margin [4].

One wall mode stands in for the vessel. A multi-mode vessel adds delay between the coil and the plasma and lowers the achievable gain ceiling beyond what latency alone produces [9].

> **OPEN CHALLENGE** **Calculated Margins That Miss the Measurement**
>
> Models of the vessel and coils still predict the recoverable displacement poorly. On one tokamak the measured value was 0.15 to 0.24 m against about 0.40 m calculated, a gap far larger than the noise, attributed to complex non-axisymmetric passive structures and to the plasma striking the first wall. On two others the measurements agreed reasonably with the calculations but with significant variability, which reinforced the need for margin in ITER's design. Progress would be an understanding of non-axisymmetric structures and wall contact good enough to predict the measured value [11].

The margin law and the radial response are simple fits rather than equilibrium calculations, so absolute numbers carry less meaning than trends.

The magnetics noise is additive. On real machines the measurement also picks up the control coils, which couples the actuator into the sensor.

The disruption at the wall ends the simulation. Halo currents and the forces of the vertical displacement event are outside the model [5].

> **OPEN CHALLENGE** **Seeing a Loss of Control Coming**
>
> ITER's design allows only 1 to 10 unmitigated vertical displacement events at the highest current and stored energy, a much stricter limit than the one on disruptions [22]. Its control system is to forecast approaching stability limits from many measurements and real-time calculations, and to call for impurity injection when it predicts a loss of vertical control [21]. The hard part is knowing how close the plasma is to the controllability boundary while it changes. Progress would be reliable, fast tools that assess that proximity in real time from reconstructed equilibria and other measurements, with warning long enough to reduce the elongation before the injection [22].

The explicit MPC's model holds the declared constants, so a learner's parameter change, the margin drop of a minor disruption, a faster wall, a latency other than 100 µs and the radial dynamics, which enter only through the box, are absent from its predictions, and the coil current limit stays a plant protection outside the program. Its plan covers 0.3 ms, the three single-period moves chosen before any comparison was run, and that horizon is why it recovers smaller large kicks than the loop. Its position gain makes it noisier than the loop with the observer.

12.8 Conclusion

An elongated tokamak plasma is an unstable plant held against a floor, a ceiling and a region, and the simulator of this chapter measures all three. Released at elongation 1.8 a 0.1 mm displacement reaches the wall in 20.0 ms, and at 2.0 in 6.7 ms. The position gain that holds the plasma must exceed about 0.45 V/mm at 1.8, about 0.70 V/mm at 2.0 and about 1.03 V/mm with the wall time halved, while with 300 µs of latency 10 V/mm settles into a 1.45 mm limit cycle. A model observer that reads the wall current from the force balance cuts the coil energy under 0.3 mm of noise from 33.9 J to 9.1 J against a filtered differentiator. A 20 mm kick at elongation 1.9 is lost at the default gains and

sooner at five times the position gain, and held with a 15 V supply or at elongation 1.8, and when the radial loop saturates the supply, priority allocation keeps a plasma that clipping loses 23 ms after the β drop. In the challenge that allocation alone meets the target at the default seed, while the suggested observer settings also pass at all five additional seeds tested. An explicit MPC with the voltage box in a 0.3 ms plan settles a 3 mm ELM in 0.52 ms against 1.32 ms but recovers kicks only up to 19.63 mm at elongation 1.9 against 19.91 mm for the loop, because it reverses the voltage late.

The model is rigid and massless with one wall mode and additive noise. A fuller model would add deformable plasma response, several wall modes, coil pickup in the magnetics and halo currents after loss. Within those limits the app lets a reader hold an unstable plant against hard limits and see which of them a tuning can move.

KEY POINTS

- An elongated plasma falls into the wall within milliseconds unless the coils push it back.
- The position gain must exceed a floor, about 0.45 V/mm at elongation 1.8, and the latency sets a ceiling above it.
- Gain decides how the loop behaves inside the recoverable region, while the growth rate and the supply voltage decide how large the region is.
- When the radial loop saturates the shared supply, priority allocation keeps a plasma that clipping loses.
- An explicit MPC with the voltage box in a 0.3 ms plan settles small kicks fastest but recovers smaller large kicks than the loop.

Where else this applies. An unstable plant held between a gain floor, a gain ceiling and a recoverable region is also the suspended ball of Volume 7, Chapter 6 and the balancing pole of Volume 7, Chapter 7.

References

[1] E. Schuster, M. Walker, D. Humphreys, and M. Krstić, “Plasma vertical stabilization with actuation constraints in the DIII-D tokamak,” *Automatica*, vol. 41, no. 7, pp. 1173–1179, 2005.

[2] J. W. Helton, K. J. McGown, and M. Walker, “Conditions for stabilization of the tokamak plasma vertical instability using only a massless plasma analysis,” *Automatica*, vol. 46, no. 11, pp. 1762–1772, 2010.

[3] E. Lazarus, J. Lister, and G. Neilson, “Control of the vertical instability in tokamaks,” *Nuclear Fusion*, vol. 30, no. 1, pp. 111–141, 1990.

[4] A. Portone, “The stability margin of elongated plasmas,” *Nuclear Fusion*, vol. 45, no. 8, pp. 926–932, 2005.

[5] C. Clauser, S. Jardin, and N. Ferraro, “Vertical forces during vertical displacement events in an ITER plasma and the role of halo currents,” *Nuclear Fusion*, vol. 59, no. 12, p. 126037, 2019.

[6] M. Walker and D. Humphreys, “On feedback stabilization of the tokamak plasma vertical instability,” *Automatica*, vol. 45, no. 3, pp. 665–674, 2009.

[7] R. Stambaugh, L. Lao, and E. Lazarus, “Relation of vertical stability and aspect ratio in tokamaks,” *Nuclear Fusion*, vol. 32, no. 9, pp. 1642–1646, 1992.

[8] H. Takahashi, K. Bol, H. Maeda, M. Okabayashi, and M. Reusch, “Experimental determination of the vertical instability strength in PDX tokamak,” Princeton Plasma Physics Laboratory, Princeton, NJ, Tech. Rep. PPPL-1911, Jul. 1982.

[9] R. Albanese and F. Villone, “The linearized CREATE-L plasma response model for the control of current, position and shape in tokamaks,” *Nuclear Fusion*, vol. 38, no. 5, pp. 723–738, 1998.

[10] G. De Tommasi, R. Albanese, G. Ambrosino, M. Ariola, P. J. Lomas, A. Pironti, F. Sartori, and J.-E. Contributors, “Current, Position, and Shape Control in Tokamaks,” *Fusion Science and Technology*, vol. 59, no. 3, pp. 486–498, 2011.

[11] D. Humphreys, T. Casper, N. Eidietis, M. Ferrara, D. Gates, I. Hutchinson, G. Jackson, E. Kolemen, J. Leuer, J. Lister, L. LoDestro, W. Meyer, L. Pearlstein, A. Portone, F. Sartori, M. Walker, A. Welander, and S. Wolfe, “Experimental vertical stability studies for ITER performance and design guidance,” *Nuclear Fusion*, vol. 49, no. 11, p. 115003, 2009.

[12] C. Neumeyer *et al.*, “Design of the ITER in-vessel coils,” Princeton Plasma Physics Laboratory, Princeton, NJ, Tech. Rep. PPPL-4576, Nov. 2010. [Online]. Available: https://www.osti.gov/biblio/1001665

[13] C. Neumeyer, “Design of the ITER In-Vessel Coils,” *Fusion Science and Technology*, 2011.

[14] A. Bemporad, M. Morari, V. Dua, and E. N. Pistikopoulos, "The explicit linear quadratic regulator for constrained systems," *Automatica*, vol. 38, no. 1, pp. 3–20, 2002.

[15] S. Gerkšič and G. de Tommasi, "Vertical stabilization of ITER plasma using explicit model predictive control," *Fusion Engineering and Design*, vol. 88, no. 6–8, pp. 1082–1086, 2013.

[16] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.

[17] G. Pannocchia and J. B. Rawlings, "Disturbance models for offset-free model-predictive control," *AIChE Journal*, vol. 49, no. 2, pp. 426–437, 2003.

[18] J. Degrave, F. Felici, J. Buchli, M. Neunert, B. Tracey, F. Carpanese *et al.*, "Magnetic control of tokamak plasmas through deep reinforcement learning," *Nature*, vol. 602, no. 7897, pp. 414–419, 2022.

[19] M. S. Parsons, A. H. Rasmussen, I. G. Stewart, C. Hansen, S. A. Sabbagh, S. Munaretto, J. W. Berkery, and S. P. Gerhardt, "A data-driven approach to real-time vertical position estimation for NSTX-U vertical stability control," *Plasma Physics and Controlled Fusion*, vol. 68, no. 2, p. 025008, 2026.

[20] J. D. King, E. J. Strait, R. L. Boivin, D. Taussig, M. G. Watkins, J. M. Hanson *et al.*, "An upgrade of the magnetic diagnostic system of the DIII-D tokamak for non-axisymmetric measurements," *Review of Scientific Instruments*, vol. 85, no. 8, p. 083503, 2014.

[21] J. A. Snipes, R. Albanese, G. Ambrosino, R. Ambrosino *et al.*, "Overview of the preliminary design of the ITER plasma control system," *Nuclear Fusion*, vol. 57, no. 12, p. 125001, 2017.

[22] D. Humphreys, G. Ambrosino, P. de Vries, F. Felici, S. H. Kim, G. Jackson *et al.*, "Novel aspects of plasma control in ITER," *Physics of Plasmas*, vol. 22, no. 2, p. 021806, 2015.

12.9 Exercises

12.1 Concept check · 1 min

In scenario 3 the position loop runs at $K_p = 0.3\,\mathrm{V/mm}$ with no velocity term at elongation 1.8. The plasma wanders for about a sixth of a second and then runs into the wall at 166.6 ms. Which explanation is right?

(a) The gain is below the floor, so the coil pushes back more weakly than the plasma runs away

(b) The gain is above the ceiling, so the latency turns the loop into a slowly growing oscillation

(c) The magnetics noise winds up the integral until the supply saturates and the loop lets go

(d) The supply is too small, so the first burst of noise carries the plasma out of the region

12.2 Concept check · 2 min

In scenario 5 the latency is 300 µs and the loop runs at $K_p = 10\,\mathrm{V/mm}$ with no velocity term. The position settles into a limit cycle of 1.45 mm rms with the supply saturated 64 % of the time. With the latency set back to 100 µs the same gain holds the plasma to 0.059 mm. What sets this behavior?

(a) The filtered differentiator turns the magnetics noise into voltage that saturates the supply

(b) The latency removes phase near the crossover that the high gain reaches, so the loop is above its ceiling

(c) The longer latency lifts the gain floor to near 10 V/mm, so the plasma is only barely held

(d) The longer latency shrinks the recoverable region, so each noise sample leaves it for a moment

12.3 Concept check · 2 min

In scenario 11 the explicit MPC brings a 3 mm ELM back within 0.5 mm in 0.42 ms, against 1.02 ms for the loop with the differentiator. Yet it loses the 19.7 mm kick at 168.0 ms, while the loop holds it. Which explanation does the chapter support?

(a) The table has no region for so large a state, so the loop takes over too late to save it

(b) Its position gain of 11.5 V/mm passes the noise to the coils and wastes the supply

(c) Its 0.3 ms plan prices the later motion as if any voltage were available, so it reverses late

(d) Under clipping its box is $V_{max} - |V_s|/2$, so it has less voltage than the loop after the kick

12.4 Theory and calculation · 6 min

The wall time constant is $\tau_w = 5.0\,\text{ms}$, and the margin follows $m(\kappa) = \max[0.02,\, 1.2(2.05 - \kappa)]$, as given in Section 12.3.

(a) Find the margin and the wall-only growth rate of Equation (12.4) at $\kappa = 1.8$ and at $\kappa = 2.0$. *Checkpoint:* $m = 0.30$ and $667\,\text{s}^{-1}$. $m = 0.06$ and $3333\,\text{s}^{-1}$.

(b) The app shows 356 and $1070\,\text{s}^{-1}$ at these elongations with the coil voltage held. Find the ratio of the rate at 2.0 to the rate at 1.8 for the wall alone and for the app. *Checkpoint:* 5.0 for the wall alone, 3.0 for the app.

(c) Treat the growth as a single exponential at the app's rate. Estimate the time for a 0.1 mm displacement to reach the wall 150 mm away at each elongation, and compare with case 2 in Table 12.4. *Checkpoint:* 20.5 ms and 6.83 ms, against 20.0 ms and 6.7 ms measured.

12.5 Theory and calculation · 5 min

Scenario 9 moves the vertical setpoint to 10 mm at $\kappa = 1.8$, where $m = 0.30$. Use $g_w = 1.0\,\text{N/A}$, $g_c = 0.8\,\text{N/A}$, $L_w = 10\,\mu\text{H}$ and $R_c = 1.0\,\text{m}\Omega$ from Table 12.2, and take the plasma at rest.

(a) Find k_d, and the steady coil current and voltage that hold the plasma at $z = 10\,\text{mm}$. *Checkpoint:* 7.69×10^4 N/m, –0.96 kA and –0.96 V.

(b) With the integral off and $K_p = 2.0\,\text{V/mm}$, find where the plasma settles. *Checkpoint:* 10.51 mm, 0.51 mm above the setpoint.

(c) Find the position gain that would bring this offset down to 0.1 mm with no integral action. *Checkpoint:* 9.71 V/mm.

12.6 Theory and calculation · 6 min

Each coil supply is limited to ±10 V, and the coils receive $V_u = V_s + V_a$ and $V_l = V_s - V_a$, so that $V_a = (V_u - V_l)/2$. After a β drop the radial loop demands $V_s = +10\,\text{V}$, its own limit, and the vertical loop demands $V_a = -6\,\text{V}$.

(a) Find the two coil voltages and the delivered V_a and V_s when each coil is clipped on its own (Equation (12.8)). *Checkpoint:* $V_u = 4\,\text{V}$, $V_l = 10\,\text{V}$, $V_a = -3\,\text{V}$, $V_s = 7\,\text{V}$.

(b) Repeat with vertical priority. *Checkpoint:* $V_a = -6\,\text{V}$, $V_s = 4\,\text{V}$, $V_u = -2\,\text{V}$, $V_l = 10\,\text{V}$.

(c) With $V_s = +10\,\text{V}$ and every command inside ±10 V, find the largest $|V_a|$ that clipping can deliver and the command that delivers it. Compare with the box the explicit MPC is given under clipping. *Checkpoint:* 5 V, from a command of –10 V. The box $V_{max} - |V_s|/2$ is also 5 V.

12.7 App challenge · 8 min

Load scenario 3, Too little gain, at the default seed. It runs at $K_p = 0.3\,\text{V/mm}$ with no velocity term at elongation 1.8. Read the Plasma tile after each run. Where a value at 0.2 s is asked for, export the CSV file and read the column `metric_max_abs_z_mm` or `metric_rms_z_mm` in the row at `time_s` 0.2.

(a) Run the case as loaded until the plasma is lost. *Checkpoint:* Lost at 166.6 ms, within 0.5 ms.

(b) Press the reload button, set Position gain to 0.40 before Run, and run to 0.6 s. Read the largest excursion at 0.2 s and the Plasma tile at the end. *Checkpoint:* 0.48 mm at 0.2 s, within 0.02 mm. Lost at 482.4 ms, within 1 ms.

(c) Reload, set Position gain to 0.50 before Run, and run to 0.2 s. Read the rms position at 0.2 s. *Checkpoint:* 0.023 mm, within 0.002 mm, held.

Explain why the run at 0.40 looks healthy at 0.2 s, using the floor of about 0.45 V/mm.

12.8 App challenge · 8 min

Load scenario 6, Velocity from a noisy position, at the default seed. The noise is 0.3 mm and the velocity gain 2 V per m/s on the filtered differentiator. Make three separate runs of 0.35 s, pressing the reload button before each, and export the CSV file of each. Read the columns `metric_rms_z_mm` and `metric_coil_energy_J` in the row at `time_s` 0.35.

(a) Run the case as loaded. *Checkpoint:* 0.216 mm and 33.9 J.

(b) Set Velocity estimate to Model observer before Run. *Checkpoint:* 0.116 mm and 9.1 J.

(c) Set Velocity gain to 0 before Run. *Checkpoint:* 0.104 mm and 6.1 J.

Give each value within 0.003 mm and 0.3 J. Which velocity estimate would you keep on this machine, and why not drop the velocity term altogether?

12.9 App challenge · 9 min

Load the graded challenge, listed last as Challenge, a bad afternoon at κ 1.95, at the default seed. It starts at $K_p = 0.8\,\mathrm{V/mm}$ with a differentiator velocity gain of 3 and clipped allocation. For each run read the Plasma tile and the challenge verdict, and press the reload button before the next run.

(a) Run the case as loaded for 1.2 s. *Checkpoint:* Lost at 661.0 ms, within 0.5 ms.

(b) Before Run set Position gain to 2, Velocity estimate to Model observer and Velocity gain to 1.5, keep Clip each coil, and run again. *Checkpoint:* Lost at 661.2 ms, within 0.5 ms.

(c) Before Run change only Coil voltage allocation to Vertical first, and run the full 1.2 s. *Checkpoint:* Held. Vertical 3.94 mm and radial 0.80 mm over the last 0.4 s, within 0.02 mm: a pass.

Say which of the three threats of the challenge decides it.

12.10 Application area · 4 min

In scenario 7 a 20 mm kick at $\kappa = 1.9$ leaves a recoverable margin of 7 % on the 10 V supply. The margin is $1 - |\xi|/\xi_{max}$ with ξ_{max} from Equation (12.5).

(a) Predict the margin after the same kick with a 15 V supply. *Checkpoint:* 38 %.

(b) Find the supply voltage that would leave 50 % margin after this kick, and name one other rating that a larger supply does not provide by itself. *Checkpoint:* 18.6 V. The coil current limit, among others.

12.11 Application area · 3 min

In scenario 10 a β drop saturates the radial loop and a 10 mm vertical kick follows 5 ms later at $\kappa = 1.9$. With each coil clipped on its own the plasma is lost at 72.7 ms, with the radial position recovered only to 4.9 mm inward. With vertical priority it is held, with a peak of 10.7 mm. Why do real machines serve the vertical loop first, and what does the recoverable margin, which reads 53 % at the kick in both runs, fail to tell the operator?

12.12 Review the agent · 5 min

An AI assistant reviews scenario 7 and writes: “At $\kappa = 1.9$ the 20 mm kick lands outside the recoverable region. The plasma is lost at 67.9 ms, and five times the position gain only loses it sooner, at 58.9 ms. No tuning can save it. The only fixes are a 15 V supply or a lower elongation, so the larger supply should be ordered.” Find the flaw and say how you would check it.