INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 18 Control of Nuclear and Fusion Systems

CHAPTER 13

Stellarator Heat Exhaust Simulator in Which One Impurity Both Protects and Poisons

Pages 342–364

Edited by John D. Hedengren

How to cite this chapter. "Stellarator Heat Exhaust Simulator in Which One Impurity Both Protects and Poisons," 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. 13, pp. 342–364.

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

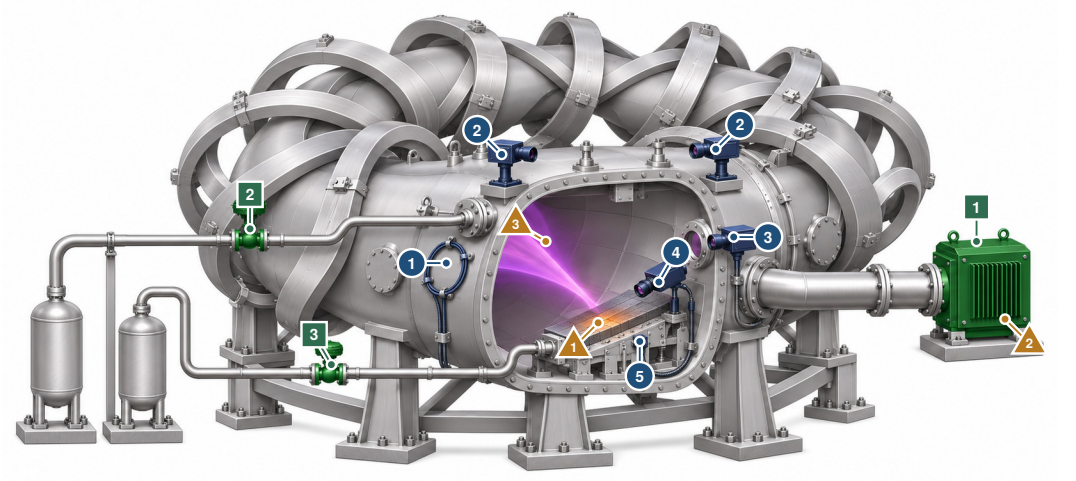
The app for this chapter
https://apmonitor.com/apps/stellarator-control/

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

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

CHAPTER 13

Stellarator Heat Exhaust Simulator in Which One Impurity Both Protects and Poisons



The controller moves 1 microwave heating 2 hydrogen fueling valve 3 neon seeding valve at the divertor

The controller reads 1 magnetic loop for the stored energy 2 interferometer for the density 3 bolometer for the radiated power 4 infrared camera on the target 5 thermocouple in the target

What disturbs it 1 loss of target cooling 2 heating sources trip 3 unannounced drop in confinement

Notice that the neon line ends beside the divertor target, so the gas that protects the warm strip enters at the edge, far from the core it can poison. Illustration, not to scale.

IN PLAIN TERMS One Gas That Protects and Poisons

The main idea. A stellarator is a fusion machine that holds very hot gas in a twisted magnetic cage for minutes. Heat leaking out lands on a few target plates, so a little neon is fed into the edge, where it gives most of that heat away as light.

The challenge. If too much neon reaches the center, the gas loses more heat than the heaters supply, and the run ends. Neon builds up slowly, so a controller that waits for the stored heat to drop acts too late.

Why it matters. The heat the machine can safely hold is set by what the plates and the neon can carry away.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history boxes of Section 13.2, the fact boxes of Sections 13.1 and 13.6, the try-it boxes of Sections 13.4 and 13.6, and the key points
Study	45 min	The whole chapter with its pause points and the twelve exercises
Practice	15 min	The first page, Section 13.6, Section 13.7, Section 13.8 and the open challenges

The control question A stellarator holds its plasma for minutes, and neon seeded into its edge protects the divertor targets from the exhaust but poisons the core if too much arrives there. How should three slow actuators share the limits of a plant that is stable but constrained?

Learning objectives

- Compute what more stored energy costs in heating under the ISS04 scaling.
- Compute the time constants of the neon inventories and the target element, and explain why the heat flux answers seeding slowly.
- Explain how a stored-energy loop hides an overseeding collapse, and how a core-radiation override on a low selector prevents it.
- Recognize an infeasible target after a source trip, and show how anti-windup on the available capacity removes the overshoot.
- Compare protected PI with an offset-free MPC on the heat flux and the neon each one uses.

Prerequisites Sections 1.4.6 (heat exhaust, detachment, radiative collapse and confinement scaling), 1.6.2 (velocity form) and 1.6.4 (running a case). Section 1.6.3 is needed only for the MPC. Chapter 12 is the fast, unstable counterpart.

The app https://apmonitor.com/apps/stellarator-control/

See also Volume 4, Chapter 3 for limiters that vote on one fuel valve with every integrator tracking the fuel commanded. Volume 4, Chapter 2 for a protective override that rests on its clamp in healthy operation. Volume 4, Chapter 6 for a vapor limit, not the product, that bounds the line speed. Volume 9, Chapter 4 for an offset-free MPC whose estimator carries the bias of a fuel no controller measures.



13.1 Introduction

A plant that is stable, slow and amply actuated looks easy to control until its limits are counted. In chemical plants the usual answer is a set of selectors, each giving an actuator to the constraint nearest its limit, and their systematic design is still an active topic [1]. A selector creates losing controllers, and a losing controller with integral action winds up unless it is told what the actuator is actually doing, which is why anti-windup and override design are treated together [2]. A jet engine's fuel valve, shared by a speed governor and its limiters, poses the same problem (Volume 4, Chapter 3). Model predictive control handles the same limits as constraints of one optimization, and an output-disturbance model makes it offset-free against a model that is wrong [3], as a boiler with an unmeasured fuel heating value shows (Volume 9, Chapter 4). Fired heaters, compressors (Volume 4, Chapter 2), reactors with a temperature limit and boilers with a drum-pressure limit all share this shape, as does the exhaust of a long-pulse fusion plasma.

A stellarator holds its plasma in a twisted magnetic field made entirely by external coils [4]. With no large plasma current there is nothing to drive and nothing to disrupt, and a stellarator can in principle run without stopping. Wendelstein 7-X (W7-X), a large optimized stellarator, has run discharges that injected 1.8 GJ under fully detached divertor conditions [5]. Over pulses that long, every watt of heating that the plasma does not radiate reaches the divertor targets, which in W7-X are built for 10 MW/m^2 in steady state [6]. Seeding an impurity such as neon into the edge radiates most of that power before it arrives, and the divertor plasma cools until it detaches [7]. The same neon that protects the targets dilutes and cools the core if it gets there, and a series of neon puffs in W7-X ended in exactly such a radiative collapse [7].

✦ DID YOU KNOW? Built for Half an Hour

Wendelstein 7-X is designed for discharges of up to 30 min and 18 GJ of heating energy, carried by superconducting coils and water-cooled plasma-facing components. Its operators raise the allowed heating energy per discharge in steps, from 1.3 GJ in one campaign to 2 GJ in the two that followed [5].

The simulator corrects the belief that a fusion plasma's performance is limited by the heating that can be bought. In this model, 20 % more stored energy costs about 60 % more heating, because the confinement time falls as the heating rises (Section 1.4.6), and every added megawatt leaves again as exhaust that the radiation loop and the targets must absorb. The binding limits are the target heat flux, the target surface temperature and the core radiation, and the controller's real task is to hold a narrow window between an attached divertor that overheats and a radiating layer that consumes the plasma. This makes the problem the opposite of the tokamak of Chapter 12, which stabilizes a fast unstable mode with one saturated actuator on a millisecond scale. Here three slow actuators (heating, hydrogen fueling and neon seeding) interact on a stable plant over minutes, and the lesson is coordination under constraints, protection logic and an infeasible target.

However, the teaching material on override control is drawn from process plants, and the fusion literature on radiated-power control is written for experts on a particular machine [8, 9]. This chapter works with a stellarator exhaust simulator, run in a web browser, in which the heating, the fueling and the seeding of a stellarator plasma are held by plain PI, by protected PI and by a constrained model predictive controller (MPC) on identical measurements, and it measures what each structure does against the exhaust limits. The two-zone balance, the ISS04 scaling [10], the two-point closure [11], the velocity-form PI law, the selectors, the bias-corrected MPC and the dual active-set quadratic program solver [12] are all established methods. What the simulator adds for study is an eight-state model whose energy, hydrogen and neon ledgers close to round-off and whose radiative collapse emerges from the balances without any scripted termination, a measurement packet that is the only input any controller reads, and every equation, constant and default that a reader needs to build the simulator again. Ten cases and a graded challenge each come with a rule that decides when they are done. The model is inspired by W7-X, but it is not a W7-X model, and no simulated result in this chapter is a W7-X result.

Section 13.2 traces the stellarator from its invention to the island divertor and radiated-power feedback. Section 13.3 sets out the two-zone power balance, the radiation of each zone, the radiating layer and detachment, the target element, the particle inventories and what the controllers measure. Section 13.4 describes the stable design point, the operating window and the unscripted collapse. Section 13.5 gives the PI loops, protected PI and the MPC. Section 13.6 measures the ten cases, the challenge and the paired comparison of the three structures, and Section 13.7 sets the loops beside the hardware of a real stellarator. Section 13.8 states what the model leaves out, and Section 13.9 summarizes the results.

13.2 Stellarators, Island Divertors and Radiated Power

The stellarator was conceived at Princeton's Project Matterhorn and described in full in 1958. A strong solenoidal field around a toroidal tube, given a rotational transform by a figure-eight tube or by additional helical windings, makes each field line trace out a closed toroidal surface [4]. Early stellarators were limited in plasma pressure and suffered high neoclassical transport and poor confinement of fast particles [13], and most fusion programs turned to the tokamak, whose transform comes from a current in the plasma itself. The current is also what a tokamak must drive, sustain and protect against disruption, which is why its vertical position needs the fast feedback of Chapter 12.

Optimization brought the stellarator back. Numerical design of the magnetic configuration reduced the neoclassical transport, raised the attainable pressure and improved the confinement of fast particles, and modular non-planar coils made such a configuration buildable [13]. W7-X, with its island divertor, is the result of that program, and the first history box below tells how it came into operation.

> ⧖ HISTORY **An Optimized Stellarator Comes On Line (2015 to 2018)**
>
> Wendelstein 7-X made its first plasma on 10 December 2015, after more than 15 years of construction. Its field comes from 50 non-planar and 20 planar superconducting coils, whose shape was found with coupled numerical codes for equilibrium, stability, transport, particle orbits and divertor loads. Its first plasmas, with five inboard limiters and no divertor, were reported in 2017 [14]. Its first operation with the island divertor followed, on ten identical uncooled graphite test divertor units, with the heating raised from six to ten gyrotrons [15]. The island divertor uses a chain of magnetic islands at the plasma boundary to guide the exhaust onto ten target units, and its stated purpose is steady-state exhaust for discharges of up to thirty minutes [6].

Radiative exhaust by impurity seeding was developed on tokamaks first. Nitrogen, neon and argon seeded into the divertor radiate power that would otherwise land on the strike lines, and the seeding rate is adjusted by feedback on a measure of the divertor state. Its extrapolation from present devices through ITER to a demonstration reactor was reviewed in 2013 [8]. Reduced models of the divertor, the two-point model above all, relate the target temperature and density to the power and density upstream and locate detachment below about 10 eV at the target [11]. The second history box below tells what the first seeding of W7-X's island divertor did.

> ⧖ HISTORY **Four Neon Puffs and a Radiative Collapse (2019 to 2020)**
>
> Impurity seeding in an island divertor was first shown on Wendelstein 7-X. Neon and nitrogen raised the radiated fraction to 80 % and more and cut the divertor heat loads by two thirds, at a price in stored energy. Neon went on radiating for tens of seconds after each puff ended, because it recycles from the carbon wall. In one discharge neon was puffed for 200 ms at 5, 7, 9 and 11 s. Each puff moved the radiating zone further inward, and after the fourth, about 8×10^{19} atoms in all, the radiated power approached the heating and a radiative collapse ended the plasma [7]. Detachment by density alone was later held stable with efficient particle exhaust: the peak divertor heat flux fell from above 5 to below 0.5 MW/m^2 as the radiated fraction rose from 0.35 to above 0.8 [6].

Feedback control of the radiated power is the current answer. A controller designed from identified frequency responses of the bolometer signal to nitrogen and neon injection now holds chosen radiated-power levels in W7-X, attached or detached, and kept the divertor detached through discharges longer than 300 s [9]. The energy confinement against which all of this is measured follows the ISS04 scaling, fitted across the international stellarator database with a renormalization factor for each configuration [10]. The density at which a stellarator plasma collapses radiatively scales with the square root of the heating power [16]. The coronal cooling rates that set how much an impurity radiates come from atomic data, and the public Open-ADAS effective coefficients rest on a generalized collisional–radiative model [17].

13.3 The Plasma and Exhaust Model

The simulated device has the plasma volume and field of W7-X, 30 m^3 and 2.5 T [15], a major radius of 5.5 m and a minor radius of 0.53 m, with five field periods and ten island divertor units, as shown in Figure 13.1. Ten microwave sources of 0.8 MW heat a hydrogen plasma at the second harmonic. A gas valve fuels it, and a neon inlet beside one divertor unit seeds the edge. The magnetic configuration is a preset chosen before the run (Standard, High iota or High mirror), and the coils are not actuators. Inside the boundary are the confined plasma as two zones, a core inside $r/a = 0.8$ and an edge shell that includes the island scrape-off layer (SOL), the hydrogen and neon inventories, the radiating layer and divertor plasma of the SOL, the bulk temperature of the hottest target element, the actuators and the diagnostics. Plasma current, magnetohydrodynamics, neutral beams, pellets, profile shapes, turbulence and fusion reactions are outside it.

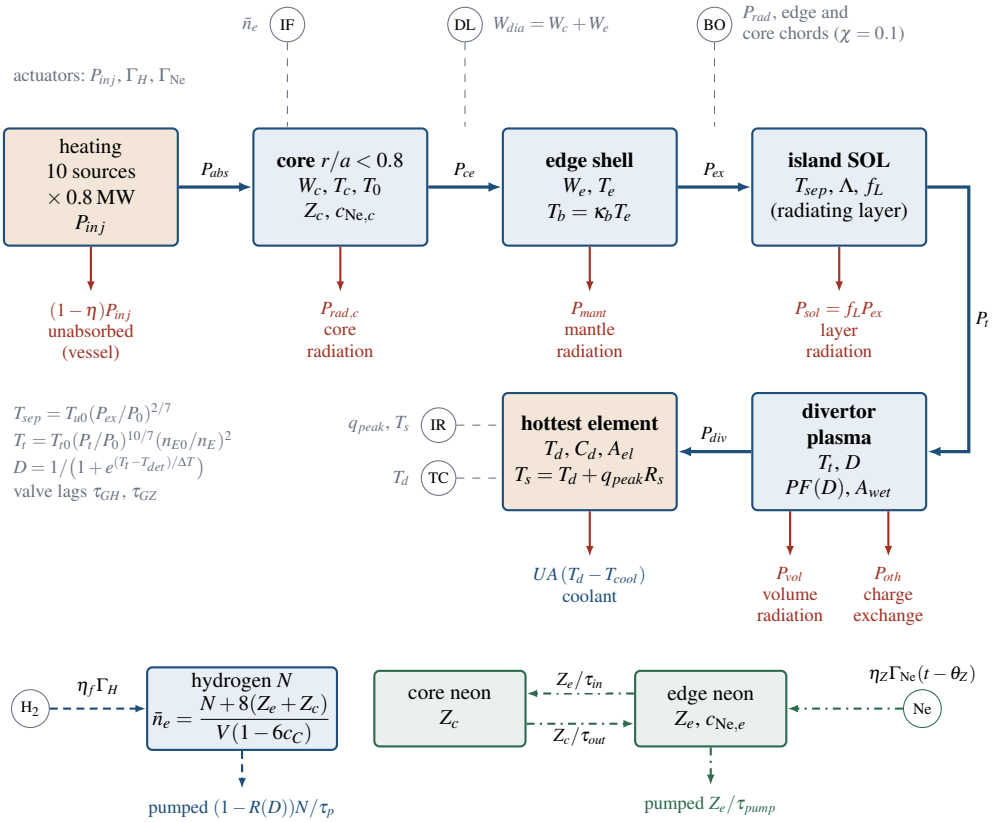


Figure 13.1 Power ledger, neon path and diagnostics of the simulated stellarator. Injected heating P_{inj} is absorbed as P_{abs} in the core, conducted as P_{ce} to the edge shell, leaves the shell as P_{ex} into the island SOL, crosses the radiating layer as P_t to the divertor plasma, and reaches the targets as P_{div} with peak heat flux q_{peak} on the hottest element. The short solid arrows below the zones (red) are the ledger's destinations, each nonnegative. Neon (dash-dotted arrows) enters near one divertor unit after the inlet delay θ_Z, exchanges between the edge and the core inventories and is pumped from the edge, and hydrogen (densely dashed arrows) is fueled and pumped. The diagnostics, joined by thin dashed taps, are the diamagnetic loop (DL), the interferometer (IF), the bolometers (BO), the infrared camera (IR) and a thermocouple (TC). After the app's Exhaust view, not to scale.

Table 13.1 classifies the variables. Eight states are integrated: six physical (the hydrogen ion inventory N, the core and shell thermal energies W_c and W_e, the edge and core neon inventories Z_e and Z_c, and the hottest element's bulk temperature T_d) and two actuator lags (the fueling and seeding flows that reach the plasma, G_H and G_Z). The heating slew, the seeding inlet delay and the instrument filters are discrete states of the actuator and sensor models. The model works in SI units with plasma temperatures in eV.

13.3.1 Two Zones and the Power Balance

The two zone energies follow the power balance of Equations (13.1) and (13.2), in which the core absorbs the heating, radiates, and conducts the rest to the shell, while the shell radiates its own share

Table 13.1 Classification of the Variables and Their Case 1 Design Values

Variable	Role	Design	Range
W	CV, stored energy	0.600 MJ	SP 0.2–0.9 MJ
$\bar{n}_e$	CV, line-averaged density	6.00×10^{19} m^{-3}	SP 2–12 $\times 10^{19}$
P_{rad} or f_{rad}	CV, radiation	2.41 MW, 48 %	SP 0–7.5 MW
q_{peak}	CV (cascade), constraint	4.50 MW/m^2	limit 10
T_s	constraint, target surface	480 °C	limit 1100
P_{inj}	MV, heating	5.03 MW	0–8, 4 MW/s
Γ_H	MV, hydrogen fueling	2.17×10^{21} s^{-1}	0–5 $\times 10^{21}$
Γ_{Ne}	MV, neon seeding	0.034×10^{18} s^{-1}	0–2 $\times 10^{19}$
f_{ren}	DV (hidden), confinement	0.65	$\times$0.5–1.2
UA	DV, target cooling	1000 W/K	20–100 %
n_{src}	DV, sources available	10	trip of 2
W_c, W_e	states, zone energies	0.529, 0.071 MJ	—
N	state, hydrogen ions	1.63×10^{21}	—
Z_e, Z_c	states, neon	8.6×10^{17}, 6.4×10^{17}	≥ 0
T_d	state, element bulk	255 °C	—
T_t, D	algebraic, divertor	21.4 eV, 0.000	$D \in [0,1]$

and passes the remainder into the island SOL,

$$\frac{\mathrm{d}W_c}{\mathrm{d}t} = P_{abs} - P_{ce} - P_{rad,c}, \tag{13.1}$$

$$\frac{\mathrm{d}W_e}{\mathrm{d}t} = P_{ce} - P_{mant} - P_{ex}, \tag{13.2}$$

where $P_{abs} = \eta(\bar{n}_e)P_{inj}$ is the absorbed heating, P_{ce} the power from core to shell, $P_{rad,c}$ the core radiation, P_{mant} the radiation of the shell (the mantle) and P_{ex} the exhaust into the SOL. The unabsorbed power $(1-\eta)P_{inj}$ goes to the vessel and is booked as its own destination. Second-harmonic X-mode heating at 140 GHz is absorbed only below its cutoff density $n_{cut} = 1.2 \times 10^{20}\,\mathrm{m}^{-3}$, and the absorption is the smooth step of Equation (13.3),

$$\eta(\bar{n}_e) = 0.10 + \frac{0.85}{1 + \exp\left[(\bar{n}_e/n_{cut} - 0.93)/0.025\right]}, \tag{13.3}$$

which gives 0.950 at the design density and 0.149 at the cutoff.

The temperatures follow from the energies, $T_c = W_c/(3\bar{n}_e V_c e)$ and $T_e = W_e/(3n_E V_e e)$, with the core volume $V_c = 0.64V$, the shell volume $V_e = 0.36V$, the edge density $n_E = 0.7\bar{n}_e$ and the elementary charge e. The temperature falls linearly across each zone, in the core from $T_0 = \max(2T_c - T_b, T_c)$ on the axis to $T_b = \kappa_b T_e$ at $r/a = 0.8$, and in the shell from T_b to the separatrix temperature T_{sep}, with the profile factor $\kappa_b = 1.6$. The core rests on its boundary. It loses $P_{ce} = (W_c - W_b)/\tau_c$ with $W_b = 3e\bar{n}_e V_c T_b$, so a shell that cools pulls the core down with it. The shell passes $P_{ex} = W_e/\tau_e$ into the SOL.

Confinement follows ISS04 times the configuration factor f_{ren}, as shown in Equation (13.4),

$$\tau_E = f_{ren}\, 0.134\, a^{2.28} R^{0.64} P^{-0.61} \bar{n}_{19}^{0.54} B^{0.84} \iota_{2/3}^{0.41}, \tag{13.4}$$

with P in MW, $\bar{n}_{19}$ in $10^{19}\,\mathrm{m}^{-3}$, a and R in m, B in T and the rotational transform $\iota_{2/3} = 0.9$ [10]. The power degradation is read from the stored energy, not from the heating. Inverting $W = \tau_E(P)P = C(\bar{n}_e)P^{0.39}$, the scaling of Equation (1.12), gives the loss power $P_W = (W/C)^{1/0.39}$ that the present energy implies, and $\tau_E = W/P_W$. A heating step therefore changes the losses only as the energy changes. The split is $\tau_c = c_1\tau_E$ and $\tau_e = c_2\tau_E$ with $c_2 = w_E = 0.12$ and $c_1 = 1 - w_E(1+k)$, $k = \kappa_b V_c/(0.7V_e) = 4.06$, so $c_1 = 0.392$. With no shell radiation, this split makes the steady stored energy exactly $\tau_E(P_{abs} - P_{rad,c})$. The factor f_{ren} is 0.65 for Standard, 0.75 for High iota and 0.58 for High mirror, teaching values, the Standard one chosen so that 5 MW holds 0.60 MJ at $6 \times 10^{19}\,\mathrm{m}^{-3}$. The

presets also change the wetted area and the peaking factors of Section 13.3.3, to 1.2 m^2 and 3.5→1.7 for High iota and 1.9 m^2 and 2.6→1.4 for High mirror.

13.3.2 Radiation in Each Zone

Each zone radiates from intrinsic carbon at a fixed fraction $c_C = 0.015$ of the electron density and from seeded neon, with the cooling rate averaged over the zone's linear temperature range, as shown in Equations (13.5) and (13.6),

$$P_{rad,c} = \bar{n}_e^2 V_c \Big[c_C \langle L_C \rangle + c_{\mathrm{Ne},c} \langle L_{\mathrm{Ne}} \rangle \Big]_{T_b}^{T_0} + P_{br}, \tag{13.5}$$

$$P_{mant} = n_E^2 V_e \Big[c_C \langle L_C \rangle + c_{\mathrm{Ne},e} \langle L_{\mathrm{Ne}} \rangle \Big]_{T_{sep}}^{T_b}, \tag{13.6}$$

where $\langle L \rangle_a^b = (b-a)^{-1} \int_a^b L dT$, $c_{\mathrm{Ne},c} = Z_c/(\bar{n}_e V_c)$ and $c_{\mathrm{Ne},e} = Z_e/(n_E V_e)$ are the neon fractions, and the hydrogen bremsstrahlung is $P_{br} = 5.35 \times 10^{-37}\, \bar{n}_e n_H \sqrt{T_c/1000}\, V_c$ in W with $n_H = \bar{n}_e(1 - 6c_C - 8c_{\mathrm{Ne},c})$. The cooling rates $L_{\mathrm{Ne}}(T)$ and $L_C(T)$, in W m^3, are two-bump analytic forms fitted to coronal rates computed from the Open-ADAS adf11 files at 10^{20} m^{-3} (https://open.adas.ac.uk/) [17]. The fits stay within a factor of 1.36 of the table for neon between 5 eV and 3 keV and within 1.61 for carbon between 2 eV and 3 keV. Neon radiates most near 30 eV (5.8×10^{-32} W m^3) and about fifty times less between 200 eV and 1 keV, which is why it radiates usefully in the cool edge and harmfully only once the radiating range reaches into the confined plasma.

13.3.3 Radiating Layer, Two-Point Closure and Detachment

The SOL receives P_{ex} at the separatrix temperature $T_{sep} = T_{u0}(P_{ex}/P_0)^{2/7}$ of the two-point model, with $T_{u0} = 110$ eV and $P_0 = 3$ MW [11]. A conduction-limited radiating layer takes the fraction f_L of it, as shown in Equation (13.7),

$$\Lambda = 10^{31} K_L \frac{n_{E,19}^2 T_{sep}^2 \big[c_C I_C + c_{\mathrm{Ne},e} I_{\mathrm{Ne}} \big]}{P_{ex,MW}^2},$$

$$f_L = f_{max}\big(1 - e^{-\Lambda/2 f_{max}}\big), \tag{13.7}$$

where $I(T_{sep}) = \int_0^{T_{sep}} \sqrt{T'} L dT'$ in W m^3 $\mathrm{eV}^{3/2}$, $n_{E,19}$ is the edge density in 10^{19} m^{-3}, $P_{ex,MW}$ the exhaust in MW, $K_L = 3 \times 10^{-4}$ and $f_{max} = 0.85$. For small Λ the layer radiates $\Lambda/2$ of the power that enters it, the slope of the conduction-limited layer, and it saturates smoothly instead of reaching 100 % with infinite slope. The layer radiates $P_{sol} = f_L P_{ex}$ and passes $P_t = P_{ex} - P_{sol}$ to the divertor plasma, whose temperature and detachment degree follow Equation (13.8),

$$T_t = T_{t0}\left(\frac{P_t}{P_0}\right)^{10/7} \left(\frac{n_{E0}}{n_E}\right)^2, \qquad D = \frac{1}{1 + e^{(T_t - T_{det})/\Delta T}}, \tag{13.8}$$

with $T_{t0} = 30$ eV, $n_{E0} = 4.2 \times 10^{19}$ m^{-3}, $T_{det} = 5$ eV and $\Delta T = 1.5$ eV. Detachment acts three ways. The divertor plasma radiates $P_{vol} = D f_{vol} P_t$ with $f_{vol} = 0.5$ and loses $P_{oth} = (0.05 + 0.15D)P_t$ to charge exchange and the main-chamber wall, so the targets receive $P_{div} = P_t - P_{vol} - P_{oth}$. The peaking factor falls from 3.0 to 1.5, $PF = 3.0 - 1.5D$, so the heat spreads. The recycling falls from 0.95 to 0.92, so the pumping improves. The peak heat flux on the wetted area $A_{wet} = 1.5\,\mathrm{m}^2$ is $q_{peak} = PF\, P_{div}/A_{wet}$. Every destination is nonnegative because $f_L < 1$ and $f_{vol} + 0.05 + 0.15 < 1$.

13.3.4 Target Element, Particles and Terminations

The hottest target element of area $A_{el} = 0.05\,\mathrm{m}^2$ receives $q_{peak} A_{el}$, which is part of P_{div}, and its bulk temperature follows Equation (13.9),

$$C_d \frac{\mathrm{d}T_d}{\mathrm{d}t} = q_{peak} A_{el} - UA\,(T_d - T_{cool}), \quad T_s = T_d + q_{peak} R_s, \tag{13.9}$$

with $C_d = 10^4\,\mathrm{J/K}$, $UA = 1000\,\mathrm{W/K}$ (a 10 s time constant), $T_{cool} = 30\,^\circ\mathrm{C}$ and the surface resistance R_s of 50 K per MW/m^2. The infrared camera sees the surface T_s, the thermocouple the bulk.

The hydrogen and neon inventories follow Equations (13.10) to (13.12),

$$\frac{\mathrm{d}N}{\mathrm{d}t} = \eta_f G_H - \big(1 - R(D)\big)\frac{N}{\tau_p}, \tag{13.10}$$

$$\frac{\mathrm{d}Z_e}{\mathrm{d}t} = \eta_Z G_Z - \frac{Z_e}{\tau_{pump}} - \frac{Z_e}{\tau_{in}} + \frac{Z_c}{\tau_{out}}, \tag{13.11}$$

$$\frac{\mathrm{d}Z_c}{\mathrm{d}t} = \frac{Z_e}{\tau_{in}} - \frac{Z_c}{\tau_{out}}, \tag{13.12}$$

with the valve lags $\tau_{GH}\dot{G}_H = \Gamma_H - G_H$ and $\tau_{GZ}\dot{G}_Z = \Gamma_{\mathrm{Ne}}(t - \theta_Z) - G_Z$, the recycling $R = 0.95 - 0.03D$, $\tau_p = 0.15\,\mathrm{s}$ (an effective particle time of 3 s), $\eta_f = 0.25$, $\eta_Z = 1$, $\tau_{pump} = 25\,\mathrm{s}$, $\tau_{in} = 8\,\mathrm{s}$, $\tau_{out} = 6\,\mathrm{s}$, $\theta_Z = 0.3\,\mathrm{s}$ and both lags 0.2 s. Quasi-neutrality with a stated average neon charge of 8 and fully stripped carbon gives the density $\bar{n}_e = (N + 8(Z_e + Z_c))/(V(1 - 6c_C))$ and $Z_{eff} = 1 + 30c_C + 56c_{\mathrm{Ne},c}$. Inventories never go negative, and a clamp, if one acts, is booked in the ledger.

Four terminations latch once, each with a historian mark and a named cause. A radiative collapse is declared when the radiation inside the separatrix, $P_{rad,c} + P_{mant}$, has exceeded P_{abs} for 0.3 s and W has fallen below half its largest value of the last ten seconds, or when T_c falls below 60 eV. The same test with $\eta < 0.5$ is a heating cutoff. A divertor interlock trips every source when T_s stays above 1200 °C for 1 s, and the end of the programmed pulse is a success. These tests only detect a collapse that the balances have already produced. Nothing in the derivative scripts it.

13.3.5 Parameters, Initialization and Numerics

Table 13.2 lists the parameters and their provenance. Every one is hidden from the learner except through the disturbance sliders, which scale f_{ren} (*teaching truth*, never told to a controller) and UA, add a delay to the diamagnetic signal or the bolometers, and scale the noise. A run starts lined out, by damped Newton iteration on the steady balances, at the operating point the case's controls define: each loop's setpoint when it is in Automatic and its manual station when it is in Manual, or the heat-flux setpoint when the cascade is on. The plant steps every 5 ms with fourth-order Runge–Kutta, and every ledger flow is accumulated with the same four stage weights, so the energy, hydrogen and neon books close to round-off. Two guards keep the zone energies nonnegative through a collapse. Every outflow of a zone (its radiation, and P_{ce} or P_{ex} leaving it) is multiplied by the gate $g(W) = x(2 - x)$ with $x = W/10\,\mathrm{kJ}$, which equals 1 at and above 10 kJ and falls to 0 at zero energy, so an emptied zone stops radiating. The least shell energy of any case that does not collapse is 0.0495 MJ (case 9 under the MPC), almost five times the gate. A step that would carry W_c or W_e below zero, or move either by more than half of $\max(W, 10\,\mathrm{kJ})$, is halved recursively, at most 14 times, each half adding to the ledger, so every held case integrates exactly as plain Runge–Kutta does.

13.3.6 What the Controllers Measure

Every controller receives the same packet every 20 ms: the diamagnetic stored energy $W_{dia} = W_c + W_e$ (2 ms lag), the interferometer density (1 ms lag), the bolometer total and an edge and a core chord (20 ms filter and a 20 ms base delay), the infrared peak heat flux and surface temperature (30 ms delay, saturating at 20 MW/m^2 and 1500 °C), the thermocouple (1 s lag), and the readbacks of the injected heating, the available source count and the two valve flows. The edge chord reads $P_{rad,e} + 0.1P_{rad,c}$ and the core chord $P_{rad,c} + 0.1P_{rad,e}$, a stated crosstalk $\chi = 0.1$. Each channel carries seeded Gaussian noise of standard deviation 4 kJ, $0.04 \times 10^{19}\,\mathrm{m^{-3}}$, 0.03 MW, 0.08 MW/m^2, 6 K and 1 K at the default level. The radiated fraction is measured radiation over measured heating, and the inferred average temperature $W_{dia}/(3\bar{n}_e V e)$ is labeled as an inference. The zone energies apart, the neon inventories, the absorbed fraction, f_{ren}, UA, D and the case's event schedule never enter the packet.

Table 13.2 Parameters, Standard Configuration

Symbol	Meaning	Value	Source
R, a, V, B	geometry and field	5.5, 0.53, 30, 2.5	W7-X [15]
ISS04	form and exponents	Eq. (13.4)	fitted [10]
f_{ren}	configuration factor	0.65 (0.75, 0.58)	teaching, calibrated
c_C	carbon fraction	0.015	teaching ($Z_{eff} \approx 1.5$ [6])
L_{Ne}, L_C	cooling rates	fits	Open-ADAS, $\times 1.36$, $\times 1.61$
κ_b, w_E	profile, shell share	1.6, 0.12	teaching
K_L, f_{max}	layer scale, saturation	3×10^{-4}, 0.85	teaching, calibrated
T_{u0}, T_{t0}	two-point prefactors	110, 30 eV	teaching
$T_{det}, \Delta T$	detachment	5, 1.5 eV	teaching
A_{wet}, PF	wetted area, peaking	1.5 m^2, 3.0→1.5	teaching
q_{lim}	heat-flux limit	10 $\mathrm{MW/m}^2$	published [6]
$T_{s,lim}$, trip	surface limit, interlock	1100, 1200 °C	teaching
A_{el}, C_d, UA	hottest element	0.05, 10^4, 1000	teaching
τ_p, R, η_f	hydrogen	0.15 s, 0.95→0.92, 0.25	teaching
$\tau_{in}, \tau_{out}, \tau_{pump}$	neon transport	8, 6, 25 s	teaching
θ_Z, lags	seeding inlet	0.3, 0.2 s	teaching
n_{cut}	X2 cutoff	1.2×10^{20} m^{-3}	physical
sources	heating	10×0.8 MW, 4 MW/s	teaching
valves	full scale	5×10^{21}, 2×10^{19} s^{-1}	teaching

13.3.7 Assumptions

The two zones keep the power balance, the loss of stored energy when the edge cools and the inward-moving radiating layer, and they omit transport equations, profile shapes and transport barriers. ISS04 is a cross-machine fit with a scatter of tens of percent, and the presets' factors are teaching values. Coronal cooling rates understate radiation at the low edge temperatures where the plasma is far from coronal equilibrium, so K_L is calibrated rather than derived: it makes the unseeded design radiate 28.4 % of its heating with carbon alone, against 0.35 measured attached [6]. The detachment closure is smooth and single-branched by construction, while real detachment can be abrupt, hysteretic and three-dimensional. One target element stands for ten units, so strike-line motion and asymmetries are absent. Neon is one species with two transport times and no neoclassical accumulation, and the absorption depends on density only, while real second-harmonic absorption also falls at low electron temperature.

13.4 Dynamic Character

13.4.1 A Stable Design Point With Slow Neon

At the design point of case 1, 5.03 MW is injected at 6.00×10^{19} m^{-3}, and the plasma stores 0.600 MJ (0.529 MJ in the core at an average of 955 eV and 0.071 MJ in the shell at 327 eV). Of the heating, 2.41 MW (48 %) is radiated, the cascade holds q_{peak} at 4.50 $\mathrm{MW/m}^2$, and the divertor plasma sits at 21.4 eV, attached. The ISS04 estimate typed independently from Equation (13.4) gives $\tau_E = 130.3$ ms and $\tau_E P_L = 0.6032$ MJ with $P_L = P_{abs} - P_{rad,c} = 4.629$ MW, 0.53 % above the model's 0.600 MJ, the difference being the shell's own radiation. The linearized eight-state plant, with the three actuators held, has the eigenvalues −0.0215, −0.100, −0.310, −0.334, −5.0, −5.0, −20.5 and −152.3 s^{-1}, all real and negative. The slowest, 46.5 s, belongs to the neon inventories, and the fastest, 6.6 ms, to the shell. The plant is stable, and its difficulty lies elsewhere. The neon that protects the target answers on a time scale about 350 times the energy confinement time.

With seeding off, the stored energy follows ISS04 times f_{ren} within 5.4 % over 2–8 MW and 3–10×10^{19} m^{-3}, the largest departure at 2 MW and 10×10^{19} m^{-3}, where carbon radiation in the shell is largest relative to the heating. The plasma holds up to 11.2–11.7 $\times 10^{19}$ m^{-3} at 3–8 MW before the absorption cutoff ends it, 1.3 to 2.0 times the Sudo density limit [16], which is an order-of-magnitude check and not a fit.

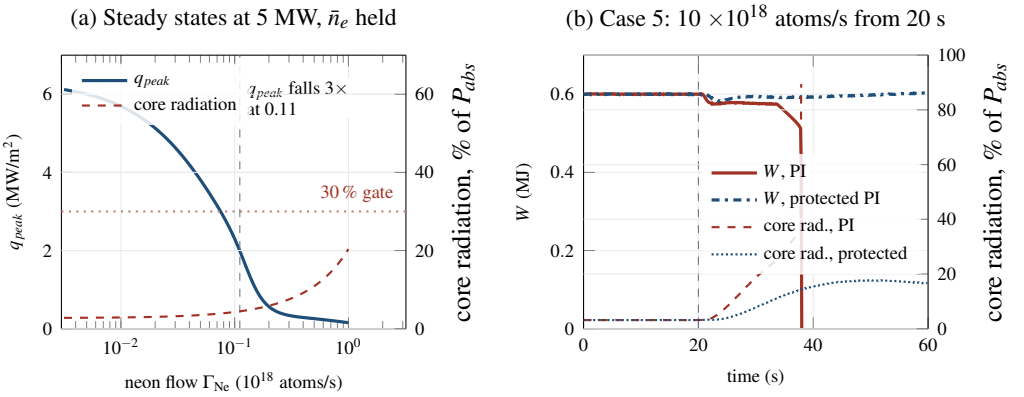


Figure 13.2 The operating window between an attached and a collapsing plasma. (a) Steady states of the model at 5 MW and 6×10^{19} m^{-3} against the neon flow: the peak heat flux (solid) falls smoothly and the core radiation (dashed) rises slowly. (b) Case 5 with the valve held at 10×10^{18} atoms/s from 20 s: under PI (stored energy solid, core radiation dashed) the heating rises to its ceiling and the plasma collapses 18.0 s after the valve opened, while protected PI's core-radiation override (stored energy dash-dotted, core radiation dotted) keeps the core under its limit and completes the pulse.

13.4.2 Detachment and the Operating Window

At fixed 5 MW and a held density, more neon lowers the peak heat flux smoothly and monotonically, as shown in Figure 13.2(a). Over 161 seeding rates from 0 to 10^{18} s^{-1}, q_{peak} falls from 6.33 MW/m^2 (unseeded, 28.4 % radiated, $T_t = 34.8$ eV) to 0.153 MW/m^2 (90.8 % radiated), and no step between neighboring rates moves it by more than 2.2 % of its unseeded value. It falls by three at 1.11×10^{17} s^{-1}, where the core radiates 4.4 % of the absorbed heating, well inside the app's design gate of 30 %. Integrated for 600 s from an attached and from a deeply detached start at 0.5, 1.2 and 3×10^{17} s^{-1}, the plant settles to the same q_{peak} within 9×10^{-5} MW/m^2, so the closure has one branch. The price is stored energy, 0.513 MJ at 10^{18} s^{-1} against 0.601 MJ unseeded, and core radiation, which reaches 20.4 % of P_{abs} at that flow. In the first seeding experiments on W7-X the stored energy fell by about 13 % at maximum radiation [7].

13.4.3 A Collapse Nobody Scripted

Holding the neon inlet at 10^{19} atoms/s from the design point, with the heating and fueling of case 1 held fixed, collapses the plasma without any termination code. The stored energy falls below 90 % after 3.1 s, and at 9.74 s the radiation inside the separatrix overtakes the absorbed heating and the shell's energy drops through the 10 kJ gate. Half the stored energy is gone at 9.80 s. By then the inlet has sent 9.8×10^{19} atoms, the core holds 1.99 % neon and Z_{eff} is 2.56. W7-X collapsed after about 8×10^{19} atoms with a core neon fraction of about 1.3 % inferred from a Z_{eff} of 2.2 [7]. Changing τ_{in} from 4 to 16 s moves the shell's emptying only from 9.1 to 10.2 s (and the loss of half the stored energy from 9.2 to 10.8 s), so the delay is of the order of τ_{in} but is set by the inventory the inlet must deliver, not by the transport time. This matters for control. A loop that waits for the stored energy to fall acts after the inventory is already in the plasma.

After that instant the model is a reduced account of a plasma that is ending. In case 5 under PI the gate holds both zone energies above 0.05 kJ, the energy left at the latch is 0.0024 MJ, and the collapse is declared at 37.985 s with 9.73 MW radiated inside the separatrix against 7.60 MW absorbed. Halving the step to 2.5 ms gives the same latch time.

TRY IT IN THE APP 5 · Seed Too Much

Load scenario 5 and press Run. Plain PI holds the stored energy, and at 20 s the neon valve opens to 10×10^{18} atoms/s and stays there. Watch the heating climb to its 8 MW ceiling while the stored energy barely moves. Once no heating is left, the plasma collapses, 18.0 s after the valve opened. Switch Controller to Protected PI, press Reset, and the same pulse completes.

13.5 The Control System as Implemented

Three loops run every 20 ms on the packet of Section 13.3: the stored energy (WIC) on the heating, the density (NIC) on the fueling, and the radiation (RIC), either the total radiated power or the measured fraction, on the seeding. Each has its own Manual/Automatic switch. Every PI is the velocity-form PI of Section 1.6.2, with the tracking time equal to the integral time. The shipped tunings are $K_{c,W} = 5\,\mathrm{MW/MJ}$ with $\tau_{I,W} = 0.5\,\mathrm{s}$, $K_{c,n} = 0.5 \times 10^{21}\ \mathrm{s^{-1}}$ per $10^{19}\ \mathrm{m^{-3}}$ with $\tau_{I,n} = 3\,\mathrm{s}$, and $K_{c,R} = 0.3 \times 10^{18}\ \mathrm{s^{-1}}$ per MW with $\tau_{I,R} = 10\,\mathrm{s}$, all reverse acting ($a = +1$ in the convention of Section 1.3, since each measurement rises with its actuator) and without derivative action. Setpoint ramps of 0.05 MJ/s and $0.5 \times 10^{19}\ \mathrm{m^{-3}/s}$ act on the stored-energy and density references in every structure, and the panel shows them in every structure. The actuators then apply their own limits: the heating clamps to the available capacity and slews at 4 MW/s, the valves clamp to their full scale, and the seeding command reaches the plasma 0.3 s later through a 0.2 s lag.

13.5.1 Protected PI

Protected PI adds four elements to the same loops. A heat-flux cascade reads the infrared peak and sets the radiation setpoint, with a direct-acting PI ($a = -1$, since more radiation lowers the heat flux) of $K_{c,q} = 0.2\,\mathrm{MW}$ per $\mathrm{MW/m^2}$ and $\tau_{I,q} = 5\,\mathrm{s}$ whose output is bounded between 30 and 90 % of the injected heating. Two protectors and the WIC request the heating, and a low selector takes the smallest, as shown in Equation (13.13),

$$P_{inj} = \min(u_W, u_q, u_{T_s}), \qquad \Gamma_{\mathrm{Ne}} = \min(u_R, u_{core}), \tag{13.13}$$

where u_q holds the infrared peak at $8\,\mathrm{MW/m^2}$ ($K_c = 0.5\,\mathrm{MW}$ per $\mathrm{MW/m^2}$, $\tau_I = 1\,\mathrm{s}$), u_{T_s} holds the surface at 1050 °C ($K_c = 0.005\,\mathrm{MW/K}$, $\tau_I = 4\,\mathrm{s}$) and u_{core} holds the core chord at 20 % of the injected heating ($K_c = 20 \times 10^{18}\ \mathrm{s^{-1}}$ per unit fraction, $\tau_I = 5\,\mathrm{s}$). The density setpoint is capped at $10 \times 10^{19}\ \mathrm{m^{-3}}$, 0.83 of the cutoff. The three protectors are reverse acting like the loops they override. With anti-windup on, every losing selector input carries external reset, its integral set so that its output equals the selected value plus its own proportional part, so a protector far from its limit stays well above the selected value and noise alone cannot hand it the actuator. The heating PI's ceiling is then the available capacity from the source readback. With anti-windup off, the losers run free and the heating PI believes in the installed 8 MW. Every PI starts by tracking its own actuator at the lined-out point, so the first update is bumpless.

13.5.2 Model Predictive Control

The MPC moves the three actuators together every 0.2 s, a three-input case of the formulation in Section 1.6.3. Its model is the declared nominal plant, the equations of Section 13.3 with the preset's published values frozen at import, plus a 0.1 s heating lag. It never reads a running case's parameters. Each decision linearizes the model about the current estimate by central differences, discretized with an exact zero-order hold in scaled states (the raw states span 21 orders of magnitude) with its affine term, and the seeding delay is carried as two delay states (0.4 s against the 0.3 s inlet). The plan

covers 150 steps (30 s) in 12 graduated move blocks and minimizes

$$J = \sum_{k=1}^{N} \left[Q_W (W_k - W_{SP})^2 + Q_n (\bar{n}_k - \bar{n}_{SP})^2 + \ell \Gamma_{\mathrm{Ne},k} \right] + \sum_j \Delta u_j^{\mathsf{T}} S \Delta u_j + \sum_i (\rho_i s_i^2 + \sigma_i s_i), \quad (13.14)$$

with $Q_W = 300$ per MJ^2, $Q_n = 3$, $S = \mathrm{diag}(0.05, 0.5, 0.02)$ and the neon price $\ell = 0.02$ in the scaled units of MW, 10^{21} and 10^{18} s^{-1}. Four soft bands, each with one slack s_i judged every fifth step, keep $q_{peak} \le q_{SP}$, $T_s \le 1050\,^{\circ}$C, the core chord below 20 % of the heating and $\bar{n}_e \le 10 \times 10^{19}$ m^{-3}, priced (ρ, σ) of $(10^3, 30)$, $(10^2, 10)$, $(10^3, 30)$ and $(10^3, 10)$. The actuators keep hard boxes, with the heating ceiling from the source readback, and per-decision move limits of 0.8 MW, 10^{21} and 2×10^{18} s^{-1}. The linear neon price pulls the seeding down until the heat-flux band holds it up, so the MPC seeds what the heat flux needs. The program is a strictly convex quadratic program solved by the dual active-set solver that the apps share [12].

The estimate is the nominal model run forward every scan from the applied heating, the fueling command and the seeding the inlet is delivering, plus five output biases on W_{dia}, $\bar{n}_e$, q_{peak}, T_s and the core chord, each updated as $b \leftarrow b + \lambda (y_{meas} - y_{model} - b)$ with $\lambda = 1 - e^{-0.02/1}$, a one-second time constant. This output-disturbance form makes the plan offset-free [3], but it corrects only levels, not gains. A plan reaches the actuators only when the solver calls it optimal and all of its values are finite and within their boxes. Otherwise protected PI takes over, labeled as a fallback, and the PI integrals, which track the applied commands while the MPC acts, make the hand-over bumpless. The protectors of Equation (13.13) remain in force as a safety layer over the MPC's commands. With a loop in Manual the MPC stands by, which is not a fallback.

Three models are kept apart in this design. The simulation model is Section 13.3 with the case's true f_{ren}, UA and disturbances. The control model is the same equations with the nominal preset values, which a hidden confinement drop makes wrong. The PI loops use no model, and their tunings were set by measurement on the plant.

13.6 Ten Cases and a Challenge, Measured

The cases appear in Table 13.3. Every number in this section was measured without a display on the app's engine (Section 1.6.1), at the default seed, stepping exactly as the page does, on the shipped case unless the sentence names a change. The scorecard judges the true plant at every 5 ms step after the first 5 s. The operating envelope requires W within 5 % of its setpoint, $q_{peak} \le 10\,\mathrm{MW/m^2}$, $T_s \le 1100\,^{\circ}$C, the core radiation at most 25 % of P_{abs} and no termination. The stored-energy error is $\int |W - W_{SP}|/W_{SP}\, dt$ over the whole run, in seconds.

13.6.1 The Ledger and the Price of Stored Energy

Over the 300 s of case 1, 1510.8 MJ is injected. Of it, 675.0 MJ (44.7 %) lands on the targets and 647.8 MJ is radiated in the SOL layer, while the core radiates 46.2 MJ, the shell 30.8 MJ, other edge losses take 35.5 MJ and 75.5 MJ of microwaves is never absorbed. The energy, hydrogen and neon residuals are 7.5×10^{-15}, 8.9×10^{-15} and 2.7×10^{-15} of their throughputs, and the books close to 2.4×10^{-14} over the challenge and to 6.6×10^{-15} through the collapse of case 5, where the plasma's remaining energy is booked as lost at termination.

Raising W_{SP} to 0.70 MJ under plain PI with the radiated fraction held (case 2) takes the heating from 5.04 to 7.40 MW, and q_{peak} climbs from 4.50 to 7.01 MW/m^2 within 6.8 s of the step and stays at 6.64 MW/m^2 at 120 s. Holding a fraction lets the target take its share of every added megawatt. The same step in case 1, with the cascade in charge, peaks at 7.11 MW/m^2 5.8 s after the step and returns to 4.54 MW/m^2 at 120 s as the seeding rises and the heating settles at 7.51 MW. With the seeding loop in Manual the extra exhaust has nowhere to go, the heat-flux protector takes the heating 4.6 s

Table 13.3 The Ten Cases and the Challenge

#	Structure and events	Teaching point
1	Protected PI, 300 s, no event	every megawatt has a destination
2	PI, W_{SP} 0.60→0.70 MJ at 20 s	stored energy is paid in exhaust
3	PI, $\bar{n}_{e,SP}$ 6→8 at 20 s	density moves confinement, radiation and detachment
4	PI, seeding by hand from 0.03	detachment is slow
5	PI, seeding held at 10 from 20 s	delayed collapse, the core override
6	PI, anti-windup off, W_{SP} 0.68, two sources out 30–90 s	infeasible target, windup
7	Protected PI, $q_{SP} = 7$, cooling halved at 30 s	thermal inertia, surface override
8	PI, $K_{c,W} = 30$, W_{dia} 250 ms late at 20 s	delay under a high gain
9	MPC, W_{SP} 0.50, hidden $f_{ren} \times 0.75$ at 30 s	mismatch, offset-free estimate
10	Protected PI, W_{SP} ramp to 0.85 at 2 kJ/s	which limit binds first
•	Challenge, 360 s: W program, trip at 150, cooling 40 % at 200, sources back at 300 s	margin for the bad day

Units: W in MJ, $\bar{n}_e$ in 10^{19} m^{-3}, seeding in 10^{18} atoms/s, q_{SP} in MW/m^2.

after the step, and W settles at 0.675 MJ, short of its target. Performance was bought with exhaust. Raising the density to 8×10^{19} m^{-3} instead (case 3) lowers the heating from 5.03 to 3.58 MW through the ISS04 density dependence, cools the divertor plasma from 21.7 to 4.3 eV ($D = 0.61$) so the seeding loop closes its valve, cuts q_{peak} from 4.54 to 0.89 MW/m^2 and lifts the core radiation from 3.2 to 7.9 % of P_{abs}.

🕮 **Pause.** Try Exercise 13.2 before reading on.

13.6.2 Detachment by Hand and the Delayed Collapse

Raising the seeding by hand from 0.03 to 0.2×10^{18} atoms/s (case 4) cuts q_{peak} from 4.68 to a third in 52.8 s and to 0.85 MW/m^2 by 180 s, with core radiation at 5.3 %. At 0.4×10^{18} atoms/s the third arrives after 17.1 s. The heat flux follows the neon inventory, which builds over the 25 s edge removal time, not the valve.

With the valve held at 10×10^{18} atoms/s from 20 s (case 5), plain PI hides the problem. The WIC answers the falling stored energy by raising the heating to its 8 MW ceiling, the core neon reaches 3.92 %, and the plasma collapses at 37.985 s, 18.0 s after the valve opened, as shown in Figure 13.2(b). Protected PI's core-radiation override takes the seeding at 20.01 s, holds the core neon at or below 2.34 % and completes the 120 s pulse with the heating never above 7.19 MW. The override acts on a measured symptom, the core chord, before the stored energy falls. The screenshot of Figure 13.3 is case 5 under PI four seconds before the collapse.

🕮 **Pause.** Try Exercise 13.1 before reading on.

13.6.3 Windup, Cooling Loss and Delay

Two tripped sources in case 6 leave 6.40 MW, and no heating level lines the plant out at the 0.68 MJ target, so the target is infeasible and W sits at 0.662 MJ. With anti-windup off the heating integral winds to 8.00 MW, and when the sources return at 90 s the stored energy overshoots to 0.717 MJ (5.4 %) and q_{peak} to 9.53 MW/m^2. With anti-windup on the integral stops at the readback's 6.40 MW, the overshoot is 0.2 % and q_{peak} peaks at 7.85 MW/m^2.

> **TRY IT IN THE APP** **6 · Lose Two Heating Sources (PI)**
>
> Load scenario 6 and press Run. Two of the ten sources trip at 30 s and return at 90 s, with anti-windup off. Watch the stored energy overshoot to 0.717 MJ when the sources come back. Set Anti-windup to On, press Reset and run again: the overshoot falls to 0.2 %.

Halving the target cooling in case 7 doubles C_d/UA from 10 to 20 s. With the cascade at 7 MW/m^2 the bulk climbs for tens of seconds, the surface protector first takes the heating at 69.1 s, the surface

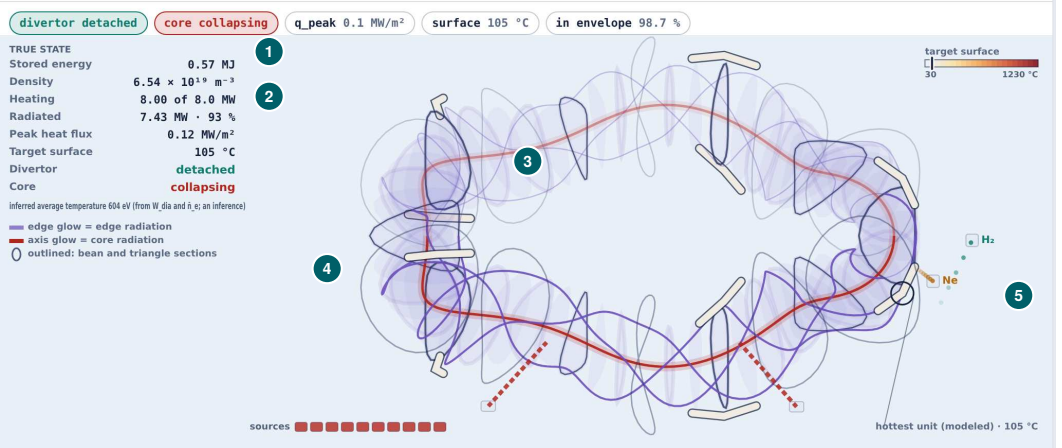


Figure 13.3 Case 5 under plain PI at $t = 34$ s, 14 s after the neon valve opened to 10×10^{18} atoms/s, on the Machine view. (1) The badges read the divertor detached and the core collapsing. (2) The true-state readout: 7.43 MW is radiated (93 %), $q_{peak} = 0.12\,\mathrm{MW/m^2}$ and the surface 105 °C. (3) The axis glow is red, core radiation. (4) The edge glows at full strength, edge radiation. (5) The neon inlet beside one divertor unit. The core radiates 26 % of the absorbed heating with 2.98 % neon, and the heating has climbed to its 8.00 MW ceiling while the stored energy has fallen to 0.571 MJ. The plasma collapses four seconds later.

peaks at 1061 °C, and W is 0.631 MJ at 240 s. Lowering q_{SP} to 5 MW/m^2 at 60 s keeps the protector out, the surface peaking at 1004 °C with W at 0.680 MJ. Radiation, not heating, takes the load.

A 250 ms delay on the diamagnetic signal (case 8) under the aggressive $K_{c,W} = 30\,\mathrm{MW/MJ}$ swings W between 0.562 and 0.638 MJ and the heating between 4.18 and 5.97 MW from 50 to 90 s. At the shipped 5 MW/MJ the same delay leaves W within 0.599–0.602 MJ. The same delay on the bolometers moves the radiation loop by less than 0.03 MW at any tested gain, because the 0.5 s seeding path and the slow neon inventory dominate, so case 8 delays the stored-energy signal and the bolometer delay is left as a disturbance slider.

13.6.4 Model Mismatch, Maximum Performance and the Challenge

In case 9 the confinement falls by a quarter at 30 s, relaxing over 0.5 s, and no controller is told. Every structure removes the offset: W returns within 0.01 MJ of 0.50 MJ and stays for 10 s from 36.4 s under the MPC, 38.6 s under protected PI and 39.8 s under PI, the MPC's stored-energy bias settling at −0.168 MJ. The difference is the exhaust. The MPC raises seeding with heating and holds q_{peak} at or below 4.53 MW/m^2 after the change, against 7.53 under protected PI and 7.32 MW/m^2 under PI.

🕮 **Pause.** Try Exercise 13.3 before reading on.

In case 10 the stored-energy setpoint ramps toward 0.85 MJ, more than the machine holds. Under protected PI the heating saturates at 82.4 s and W peaks at 0.720 MJ, with q_{peak} held at 4.50 MW/m^2 and 5.23 MW radiated, so actuator authority binds first. Lowering q_{SP} to 1 MW/m^2 costs 0.011 MJ (0.709 MJ). Plain PI reaches 0.725 MJ only by letting the target reach 9.87 MW/m^2 and 1018 °C. The MPC reaches 0.753 MJ, but by running the density at 6.62×10^{19} m^{-3}, 10 % over its setpoint, a trade its weights allow and the scorecard does not credit.

The challenge runs 360 s: the program raises W from 0.60 to 0.66 MJ from 30 s, two sources trip at 150 s, the cooling falls to 40 % at 200 s, and the sources return at 300 s. A pass needs 95 % of the judged time in the envelope, judged against the program, and a completed pulse. As shipped, protected PI with the cascade at 7 MW/m^2 spends 66.4 % in the envelope, losing 119 s to the stored

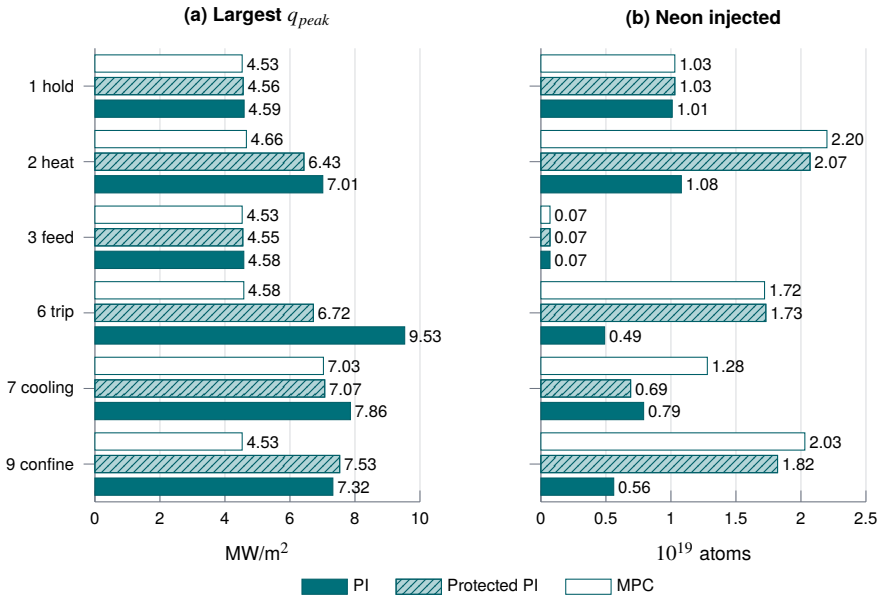


Figure 13.4 The paired comparison of Table 13.4 for the six cases whose rows report the heat flux and the neon: (a) the largest peak heat flux, against a limit of 10 MW/m^2 at the right edge of the axis, and (b) the neon injected. Solid: PI. Hatched: protected PI. White: MPC. Notice that the MPC keeps the largest heat flux near 4.5 MW/m^2 in cases 2, 6 and 9, where protected PI lets it reach 6.43 to 7.53 MW/m^2, and that it does so with more neon than protected PI in cases 7 and 9.

energy after the surface protector takes the heating at 221.2 s, with the surface peaking at 1074 °C. With the suggested q_{SP} of 4.5 MW/m^2 it passes at 100 %, the surface peaking at 822 °C. Plain PI trips the divertor interlock at 267.7 s, after 86.7 % of the time judged up to then, and the MPC passes at 100 % with the shipped settings while its soft surface band keeps the surface at or below 1060 °C.

13.6.5 The Paired Comparison

Table 13.4 lists the three structures on identical seeds, diagnostics, limits and event schedules, and Figure 13.4 draws its heat flux and neon columns. Protected PI uses less neon than the MPC in cases 7 and 9, and plain PI uses the least neon wherever it does not overheat the target. Cases 1 and 3 are ties in the envelope. In cases 4 and 5 a loop is in Manual, the MPC stands by and the runs are identical to protected PI. Case 8 ships a deliberately aggressive PI gain and is not a fair comparison. The MPC's advantage comes from coordination, raising seeding together with heating, which holds q_{peak} at 4.53–4.66 MW/m^2 in cases 2, 6 and 9 against 6.43–7.53 MW/m^2 for protected PI, and from its surface band in case 7 and the challenge. Its quadratic program solves in a median of 1.71 ms in Node.js (95th percentile 1.85 ms, largest 2.6 ms) over the 900 solves of case 9, and in a median of 1.7 ms in Chrome, where the 360 s challenge needs 18.0 s at the fastest playback, twenty times real time, under either structure. At that speed the five-minute case 1 takes 15.0 s of wall time.

+ DID YOU KNOW? Losing Cooling Costs the Most Neon

From the simulator: when the target cooling is halved in case 7, plain PI keeps the plant in the operating envelope only 22.3 % of the time and protected PI 77.2 %. The MPC stays in it throughout, at 100 %, and injects the most neon, 1.28×10^{19} atoms against 0.69 under protected PI and 0.79 under PI (Table 13.4).

Table 13.4 Paired Comparison: Time in the Envelope (%), Largest q_{peak} (MW/m^2) and Neon Injected (10^{19} atoms)

Case	PI	Protected PI	MPC
1 hold	100 / 4.59 / 1.01	100 / 4.56 / 1.03	100 / 4.53 / 1.03
2 heat	97.0 / 7.01 / 1.08	97.0 / 6.43 / 2.07	97.9 / 4.66 / 2.20
3 feed	100 / 4.58 / 0.07	100 / 4.55 / 0.07	100 / 4.53 / 0.07
6 trip[a]	99.8 / 9.53 / 0.49	100 / 6.72 / 1.73	100 / 4.58 / 1.72
7 cooling	22.3 / 7.86 / 0.79	77.2 / 7.07 / 0.69	100 / 7.03 / 1.28
9 confine	97.5 / 7.32 / 0.56	97.2 / 7.53 / 1.82	97.5 / 4.53 / 2.03
10 max W (MJ)	0.725	0.720	0.753[b]
Challenge, shipped	86.7, interlock	66.4, fail	100, pass
Challenge, $q_{SP} = 4.5$	—	100, pass	100, pass

[a] Case 6 ships with anti-windup off, which applies to the PI column. [b] With the density 10 % over its setpoint. Challenge figures are judged against the program.

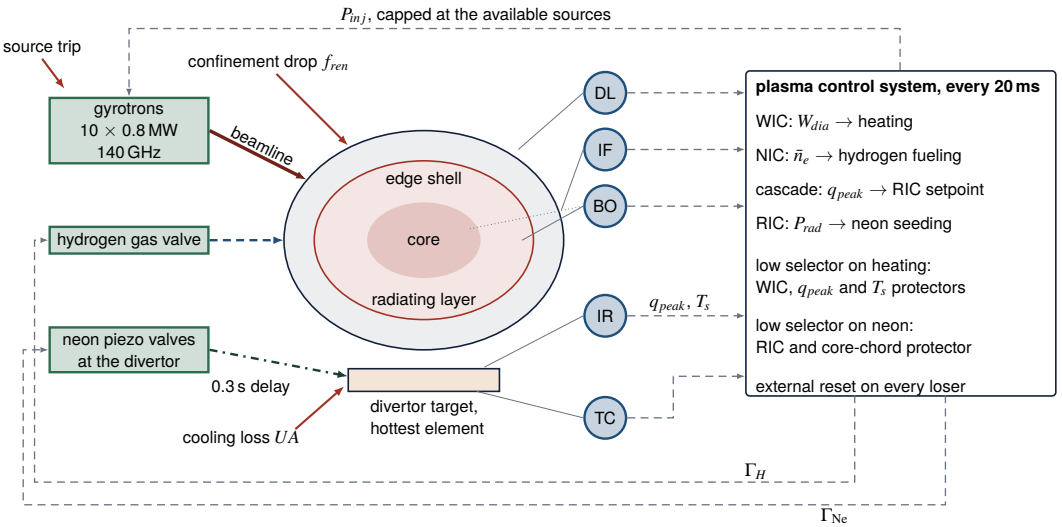


Figure 13.5 The loops of Figure 13.1 with the hardware of a stellarator. Diamagnetic loops (DL), an interferometer (IF), bolometer cameras with an edge and a core chord (BO), infrared cameras on the targets (IR) and a thermocouple (TC) feed the plasma control system, which runs the three PI loops, the heat-flux cascade and the two low selectors every 20 ms. Its commands go to the gyrotrons, the hydrogen valve and the neon valves at the divertor. Thick solid line: microwave power. Dash-dotted: neon. Densely dashed: hydrogen. Thin dashed arrows: signals. Solid arrows: disturbances.

13.7 Technology in Practice

The diagnostics DL, IF, BO, IR and TC and the three actuators of Figure 13.1 stand, on a real stellarator, for magnetic loops, an interferometer, bolometer cameras, infrared cameras on the targets, gyrotrons and gas valves. Figure 13.5 draws the chapter's loops with that hardware in place of the schematic's symbols.

What Measures It

The stored energy comes from diamagnetic loops, and the density from an interferometer as a line integral over a sight line [5, 15]. The target heat flux is not measured directly. Infrared cameras in the 8 to 10 µm or 3 to 5 µm band record the surface temperature, and the heat flux follows from the heat diffusion equation solved for the bulk of the tile. Ten divertor modules need ten camera systems,

and a detached target can fall below the system's detection limit of 0.4 MW/m^2 [15]. The chapter's infrared channel is a 30 ms delay that saturates at 20 MW/m^2, on one element that stands for all ten.

> TECHNOLOGY **Sensor: Resistive Bolometers**
>
> A bolometer measures radiation as heat. In the stellarator's system each detector is a blackened gold absorber 5 µm thick, with a platinum sensor and reference resistor behind it in a Wheatstone bridge. The bridge voltage and its rate of change give the absorbed power through a cooling time of about 0.1 s, calibrated before each discharge, sampled every 0.8 or 1.6 ms. Two pinhole cameras with 32 and 24 channels view one cross section. The foil also absorbs stray microwaves, which had to be suppressed about 300-fold, and thermal offset drift matters most in long pulses [18]. A plane far from the radiating zone reads low: radiation concentrated in the divertor may be substantially underestimated [7]. The chapter's BO is a 20 ms filter and delay with two chords that see each other's radiation.

What Moves It

The neon reaches the divertor through gas inlet boxes of five piezoelectric valves mounted behind a horizontal target plate, whose capillary nozzles end at the plate's surface and feed the gas straight into the divertor plasma. The valves open for as little as 2 ms, from reservoirs at 2 mbar to 60 bar [15]. In the chapter the seeding command reaches the plasma 0.3 s later through a 0.2 s valve lag.

> TECHNOLOGY **Final Element: Gyrotrons**
>
> The heating is microwaves at 140 GHz from gyrotrons, tubes in which an electron beam gives its energy to a wave in a resonator. The original tubes of the stellarator are rated at 1 MW. A 1.5 MW-class prototype was held to 1.3 MW by parasitic oscillations and run routinely at 1.1 MW. Single tubes fail without warning in long pulses, so a power stabilization system now raises the output of the others in real time and holds the total. In a 363 s discharge, cooling problems in the beamline caused growing power fluctuations from 270 s and arcing that ended the discharge [5]. The chapter's ten 0.8 MW sources slew at 4 MW/s, and its trip takes two of them out at once.

Where the Controller Runs

Feedback is not fundamentally required to control a stellarator plasma, yet the long pulses of Wendelstein 7-X now run on three real-time loops: radiated power from the bolometers on the seeding valve's voltage, a stabilization of the total heating power, and density control [5]. In the control system designed for ITER a pulse schedule drives nominal operation, a supervisory layer changes the control in real time on plasma or plant events, and multiple actuators call for shared actuator management [19]. The sources read give no cycle time for these loops. The chapter's PI loops run every 20 ms and its MPC every 0.2 s, and its two low selectors with external reset are that shared management in its simplest form.

> IN PRACTICE **Keep the Fueling Loop Out of the Seeding Test**
>
> Density feedback and impurity seeding disturb each other. In the first neon seeding of an island divertor, density changes caused by the seeding had always made the fueling feedback answer with unwanted hydrogen puffs that covered the effect of the neon, so its density target was lowered before the seeding began. Density lost to seeding can instead be made up by feedback-controlled fueling through the divertor [7]. A review of ITER's control makes the general point: core density and core radiation are tightly coupled goals, and where such goals are not joined in one algorithm, managing the actuators they share becomes an important challenge [20].

What Is Modernizing It

Radiated-power feedback has turned seeding from a programmed puff into a regulated variable. Neon led the radiated power through steps from 3.5 to 6.2 MW close to its target, and the same loop kept

a 363 s discharge detached [5]. A tokamak has gone a step further and given each limit its own impurity: argon at the outer midplane regulated the core radiation while nitrogen in the divertor held the peak target heat flux below 5 MW/m^2 with 23 MW of input power [20], the split this chapter's core-radiation override makes with one gas. Fueling is changing too. A screw-type pellet injector built for discharges of up to 30 min delivers pellets of 2 to 3 mm at 400 to 600 m/s and up to 5 Hz, and is to be joined to the real-time density control [5].

13.8 Limitations

The model is a teaching model assembled from public literature and calibrated to a few published measurements, not a W7-X model and not plant data, and its cases are not W7-X results. Every limit, setpoint and tuning is a teaching value. The 10 MW/m^2 heat-flux limit is the published rating of plasma-facing components [6], while the 1100 °C surface limit and the 1200 °C interlock stand in for values that come from the target material and a machine's protection system. Passing a case in simulation demonstrates no safety property. The two-zone profile, the coronal fits with a calibrated layer scale, the smooth single-branched detachment, one target element, a neon species with two transport times and a density-only absorption are the reductions of Section 13.3, and each one removes a behavior that real devices show.

> ⚠ **OPEN CHALLENGE** **A Divertor Regime Not Yet Settled**
>
> For ITER, a fully consistent way to run the divertor with a low enough target heat flux had not been found when its control was reviewed in 2015. The answer is expected to combine enough divertor radiation with partial detachment, but even the controlled variable is open: the local radiated power in core and divertor, the measured or inferred target heat flux, a local temperature or density in the divertor, or a line ratio that stands for the degree of detachment. Fueling and impurity gas or pellets are the actuators. Experiments on present machines have begun to test such schemes, and progress would be one that carries over to ITER [20].

The collapse is resolved only until the shell empties. The zone gate of Section 13.3 is a numerical guard rather than physics, and what follows the gate in a real device (a radiating front reaching the axis, a heating that is no longer absorbed at low temperature) is outside the model. The setpoint ramps act in every structure, so the steps of cases 2 and 3 reach the loops as ramps of 2 and 4 s.

The MPC's model is the nominal plant linearized every 0.2 s, and its biases correct levels but not gains, so a large hidden change is met by feedback rather than by prediction. Its horizon of 30 s is shorter than the 46.5 s neon mode, and the plan sees the far tail of an inventory change only through its last block. Nothing is guaranteed beyond the horizon, and the fallback is protected PI rather than a certified controller. The PI tunings were set by measurement on this plant and are not claimed optimal.

> ⚠ **OPEN CHALLENGE** **Measuring Radiation in a Twisted Plasma**
>
> In a stellarator the three-dimensional field makes an accurate total radiated power hard to derive from one cross section [15]. Tomography in one plane found an up-down asymmetry that three-dimensional edge modeling did not capture [18]. Imaging bolometers are one way forward. One installed in 2021 views the divertor region over 40 by 68 degrees with more than 500 channels, and its upgrade has more than 1400. Progress would be the radiated power measured in three dimensions, which four imaging bolometers combined have achieved on another helical device [21].

13.9 Conclusion

A long-pulse stellarator turns heat exhaust into a constrained control problem, and this simulator presents it with ledgers that close to 10^{-14} of the throughput. Its stored energy costs heating at the ISS04 rate, its divertor detaches smoothly as neon is added, q_{peak} falling by three with the core radiating 4.4 % of the absorbed heating, and a held neon flow ends the pulse in a collapse that no code

scripts and whose delay is set by the delivered inventory. Plain PI hides that collapse behind rising heating for 18.0 s, protected PI's core override prevents it, and anti-windup cuts the overshoot after a source trip from 5.4 to 0.2 %. The MPC holds the heat flux near 4.5 MW/m^2 where protected PI lets it reach 6.4–7.5 MW/m^2, and protected PI uses less neon in two of the cases. The model is reduced and calibrated rather than validated, and its collapse is followed only until the edge empties. A fuller model would add a radial transport equation, non-coronal edge radiation, a second impurity and a temperature-dependent absorption, and an identified rather than declared MPC model in the manner of current radiated-power control. The simulator shows that the highest sustainable performance of a long pulse is set by its exhaust and its actuators, not by the heating it can buy.

KEY POINTS

- Neon at the edge protects the divertor targets, but too much of it in the core ends the pulse in a radiative collapse.
- In this model 20 % more stored energy costs about 60 % more heating, and every added megawatt leaves again as exhaust.
- Plain PI hides an overseeding collapse behind rising heating for 18.0 s, and protected PI's core-radiation override prevents it.
- Anti-windup on the available capacity cuts the overshoot after a source trip from 5.4 to 0.2 %.
- The MPC holds the heat flux near 4.5 MW/m^2 by raising seeding with heating, where protected PI lets it reach 6.4–7.5 MW/m^2.

Where else this applies. A protective override that takes a shared actuator through a low selector while the losers track the winner also guards the fuel valve of Volume 4, Chapter 3 and the compressor of Volume 4, Chapter 2.

References

[1] D. Krishnamoorthy and S. Skogestad, "Systematic design of active constraint switching using selectors," *Computers & Chemical Engineering*, vol. 143, p. 107106, 2020.

[2] M. V. Kothare, P. J. Campo, M. Morari, and C. N. Nett, "A unified framework for the study of anti-windup designs," *Automatica*, vol. 30, no. 12, pp. 1869–1883, 1994.

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

[4] L. Spitzer, "The stellarator concept," *The Physics of Fluids*, vol. 1, no. 4, pp. 253–264, 1958.

[5] O. Grulke, G. Acton, J. Adamek, D. Aggelis, R.-M. Alamo-Calderon, C. Albert *et al.*, "Overview of Wendelstein 7-X high-performance operation," *Nuclear Fusion*, vol. 66, no. 11, p. 116003, 2026.

[6] O. Schmitz, Y. Feng, M. Jakubowski, R. König, M. Krychowiak, M. Otte *et al.*, "Stable heat and particle flux detachment with efficient particle exhaust in the island divertor of Wendelstein 7-X," *Nuclear Fusion*, 2020, published online 3 September 2020.

[7] F. Effenberg, S. Brezinsek, Y. Feng, R. König, M. Krychowiak, M. Jakubowski *et al.*, "First demonstration of radiative power exhaust with impurity seeding in the island divertor at Wendelstein 7-X," *Nuclear Fusion*, vol. 59, no. 10, p. 106020, 2019.

[8] A. Kallenbach, M. Bernert, R. Dux, L. Casali, T. Eich, L. Giannone *et al.*, "Impurity seeding for tokamak power exhaust: from present devices via ITER to DEMO," *Plasma Physics and Controlled Fusion*, vol. 55, no. 12, p. 124041, 2013.

[9] A. Tsikouras, F. Reimold, G. Partesotti, M. Krychowiak, T. Schröder, M. van Berkel, and the W7-X Team, "Improved feedback control of the radiated power in Wendelstein 7-X using system identification," *Plasma Physics and Controlled Fusion*, vol. 68, no. 8, p. 085004, 2026.

[10] H. Yamada, J. H. Harris, A. Dinklage, E. Ascasibar, F. Sano, S. Okamura *et al.*, "Characterization of energy confinement in net-current free plasmas using the extended International Stellarator Database," *Nuclear Fusion*, vol. 45, no. 12, pp. 1684–1693, 2005.

[11] P. C. Stangeby, "Basic physical processes and reduced models for plasma detachment," *Plasma Physics and Controlled Fusion*, vol. 60, no. 4, p. 044022, 2018.

[12] D. Goldfarb and A. Idnani, "A numerically stable dual method for solving strictly convex quadratic programs," *Mathematical Programming*, vol. 27, no. 1, pp. 1–33, 1983.

[13] G. Grieger, W. Lotz, P. Merkel, J. Nührenberg, J. Sapper, E. Strumberger, H. Wobig *et al.*, "Physics optimization of stellarators," *Physics of Fluids B: Plasma Physics*, vol. 4, no. 7, pp. 2081–2091, 1992.

[14] T. Klinger, A. Alonso, S. Bozhenkov, R. Burhenn, A. Dinklage, G. Fuchert *et al.*, "Performance and properties of the first plasmas of Wendelstein 7-X," *Plasma Physics and Controlled Fusion*, vol. 59, no. 1, p. 014018, 2017.

[15] T. S. Pedersen, R. König, M. Krychowiak, M. Jakubowski, J. Baldzuhn, S. Bozhenkov *et al.*, "First results from divertor operation in Wendelstein 7-X," *Plasma Physics and Controlled Fusion*, vol. 61, no. 1, p. 014035, 2019.

[16] S. Sudo, Y. Takeiri, H. Zushi, F. Sano, K. Itoh, K. Kondo, and A. Iiyoshi, "Scalings of energy confinement and density limit in stellarator/heliotron devices," *Nuclear Fusion*, vol. 30, no. 1, pp. 11–21, 1990.

[17] H. P. Summers, W. J. Dickson, M. G. O'Mullane, N. R. Badnell, A. D. Whiteford, D. H. Brooks *et al.*, "Ionization state, excited populations and emission of impurities in dynamic finite density plasmas: I. The generalized collisional–radiative model for light elements," *Plasma Physics and Controlled Fusion*, vol. 48, no. 2, pp. 263–293, 2006.

[18] D. Zhang, R. Burhenn, C. D. Beidler, Y. Feng, H. Thomsen, C. Brandt *et al.*, "Bolometer tomography on Wendelstein 7-X for study of radiation asymmetry," *Nuclear Fusion*, vol. 61, no. 11, p. 116043, 2021.

[19] 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.

[20] 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.

[21] G. Partesotti, G. A. Wurden, F. Reimold, B. J. Peterson, F. Federici, K. Mukai *et al.*, "The infrared imaging video bolometer at Wendelstein 7-X," *Review of Scientific Instruments*, vol. 96, no. 6, p. 063503, 2025.

13.10 Exercises

13.1 ○ Concept check · 1 min

In case 5 under plain PI the neon valve is held at 10×10^{18} atoms/s from 20 s while the seeding loop is in Manual. The stored energy stays near 0.57 MJ for more than ten seconds and then collapses. What holds it there for so long?

(a) The neon needs the inlet delay and several transport times to reach the core, so the core radiation does not change until then

(b) The radiation loop closes the neon valve until the core chord reads twenty percent of the heating

(c) The fueling loop lowers the density, so each neon atom in the core radiates less than at the design point

(d) The stored-energy loop raises the heating toward its ceiling, which offsets the growing radiation until no heating is left

13.2 ○ Concept check · 2 min

In case 2 three plain PI loops hold the stored energy, the density and the measured radiated fraction at 48 %. The stored-energy setpoint steps from 0.60 to 0.70 MJ, and the peak heat flux settles near 6.6 MW/m^2 instead of the 4.5 MW/m^2 it had before. Why?

(a) The heating rises by nearly half, and the unradiated share of every added megawatt reaches the divertor

(b) The seeding valve saturates, so the radiated fraction falls below its setpoint and more power reaches the target

(c) The larger stored energy raises the peaking factor of the target toward its attached value of 3.0

(d) The density rises with the heating, and the denser edge heats the divertor plasma through the two-point closure

13.3 Concept check · 2 min

In case 9 the confinement falls by a quarter at 30 s and the MPC is not told. Its model keeps the nominal confinement factor. Which statement describes what happens?

(a) The MPC leaves a steady stored-energy offset, because the gain of its model is wrong after the change

(b) The output bias on the stored energy absorbs the steady error, so the offset goes, but a transient remains because a bias corrects levels and not gains

(c) The MPC estimates the new confinement factor from the diamagnetic signal, so its model is correct again within a second

(d) The MPC rejects its plans after the change and hands the heating to protected PI, whose integral removes the offset

13.4 Theory and calculation · 6 min

The neon inventories follow Equations (13.11) and (13.12) with $\eta_Z = 1$, $\tau_{pump} = 25\,\text{s}$, $\tau_{in} = 8\,\text{s}$ and $\tau_{out} = 6\,\text{s}$.

(a) Write the two equations as $\dot{z} = Az + bG_Z$ with $z = (Z_e, Z_c)$, and find the two eigenvalues of A and their time constants. *Checkpoint:* -0.0215 and $-0.310\,\text{s}^{-1}$, 46.5 s and 3.22 s.

(b) Find the steady inventories for a seeding flow of 0.034×10^{18} atoms/s, and compare them with Table 13.1. *Checkpoint:* $Z_e = 8.5 \times 10^{17}$ and $Z_c = 6.4 \times 10^{17}$, against 8.6×10^{17} and 6.4×10^{17} in the table.

(c) In case 4 the heat flux takes 52.8 s to fall to a third after the valve is raised. Which part of the seeding path sets that time: the 0.3 s inlet delay, the 0.2 s valve lag, or one of the two modes of part (a)? *Checkpoint:* The slow mode, 46.5 s.

13.5 Theory and calculation · 6 min

The hottest target element follows Equation (13.9) with $C_d = 10^4\,\text{J/K}$, $A_{el} = 0.05\,\text{m}^2$, $R_s = 50\,\text{K}$ per MW/m^2, $T_{cool} = 30\,^\circ\text{C}$ and a nominal $UA = 1000\,\text{W/K}$. The surface protector holds T_s at 1050 °C.

(a) Compute the steady bulk and surface temperatures at $q_{peak} = 4.5\,\text{MW/m}^2$ with nominal cooling, and compare them with Table 13.1. *Checkpoint:* 255 °C and 480 °C, as in the table.

(b) Show that the steady surface temperature is $T_{cool} + q_{peak}(A_{el}/UA + R_s)$. With the cooling at 40 %, as in the challenge after 200 s, find the bulk time constant and the largest steady q_{peak} that keeps the surface at or below the protector's setpoint. *Checkpoint:* 25 s, $5.83\,\text{MW/m}^2$.

(c) With the cooling at 40 %, compute the steady surface temperature at $7\,\text{MW/m}^2$ and at $4.5\,\text{MW/m}^2$. *Checkpoint:* 1255 °C and 817.5 °C.

13.6 Theory and calculation · 6 min

Under the ISS04 scaling with $\tau_E \propto P^{-0.61}\bar{n}^{0.54}$, the stored energy at a fixed density grows as $P^{0.39}$, where P is the loss power $P_{abs} - P_{rad,c}$.

(a) Find the heating factor that 20 % more stored energy needs. Then predict the heating after the case 2 step from 0.60 to 0.70 MJ, starting from 5.04 MW, if the absorbed and radiated shares stay fixed. *Checkpoint:* 1.60, 7.48 MW, against 7.40 MW measured.

(b) In case 3 the density steps from 6 to $8 \times 10^{19}\,\text{m}^{-3}$ at a held stored energy. Predict the loss power after the step from the design value of 4.629 MW, and then the injected heating, with $\eta = 0.950$ and the core radiating 7.9 % of the absorbed heating after the step. *Checkpoint:* 3.11 MW, 3.55 MW, against 3.58 MW measured.

(c) After the case 3 step the divertor plasma is at 4.3 eV. Compute the detachment degree from Equation (13.8), and compare it with the design point at 21.4 eV. *Checkpoint:* 0.61 against 0.000.

13.7 App challenge · 8 min

Load scenario 5, Seed too much, and change no setting. Plain PI runs with the seeding loop in Manual, and the neon valve is held at 10×10^{18} atoms/s from 20 s. Use the default seed and the 20× playback speed.

(a) Run until the pulse ends. Read the Discharge tile of the scorecard, and find in the exported CSV file the first row in which `P_inj_MW` reaches its ceiling. *Checkpoint:* Radiative collapse at 38.0 s. The heating reaches 8.00 MW at 33.7 s, within 0.2 s.

(b) Switch Controller to Protected PI, press Reset and run the full 120 s. Read the Discharge tile and the line "won" on the seeding card, and find in the exported file the largest `P_inj_MW` and the largest `cNe_core_pct`. *Checkpoint:* Pulse completed, seeding won by the core-radiation protector, 7.18 MW and 2.34 %, each within 0.02.

Explain from the two runs why the override acts on a measured symptom rather than on the stored energy.

13.8 App challenge · 8 min

Load scenario 6, Lose two heating sources (PI). Two of the ten sources trip at 30 s and return at 90 s, and Anti-windup ships Off. Use the default seed, the 20× playback speed and runs of 150 s.

(a) Run as shipped. In the exported CSV file read `P_req_MW` and `P_inj_MW` at 60 s, and find the largest `W_MJ` after 90 s. Read the Peak heat flux (max) tile. *Checkpoint:* 8.00 MW asked against 6.40 MW applied, 0.716 MJ, within 0.002 MJ, 9.53 MW/m^2 at 90 s.

(b) Set Anti-windup to On, press Reset and run again. Read the same three values. *Checkpoint:* 6.40 MW asked and applied, 0.681 MJ, within 0.002 MJ, 7.85 MW/m^2 at 4 s.

Express each stored-energy peak as an overshoot of the 0.68 MJ setpoint, in percent of the setpoint, and explain why the heat-flux tile of run (b) names a time before the trip.

13.9 App challenge · 9 min

Load the challenge, The long pulse. It runs 360 s: the stored-energy program rises to 0.66 MJ from 30 s, two sources trip at 150 s, the cooling falls to 40 % at 200 s and the sources return at 300 s. Protected PI ships with the Peak heat-flux setpoint at 7 MW/m^2. Use the default seed and the 20× playback speed, and read the Challenge target and Target surface (max) tiles at 360 s after each run.

(a) Run as shipped. *Checkpoint:* Not met, 66.4 % in the envelope, most lost to the stored energy (119 s), 1074 °C.

(b) Set Peak heat-flux setpoint to 4.5 MW/m^2, press Reset and run. Read the Neon injected tile as well, for both runs. *Checkpoint:* Met, 100.0 %, 822 °C, 3.37×10^{19} atoms against 0.65×10^{19}.

(c) Set the Peak heat-flux setpoint to the bound of Exercise 13.5(b), rounded down to the slider's step of 0.1 MW/m^2, press Reset and run. *Checkpoint:* 5.8 MW/m^2: met, 100.0 %, 1047 °C.

13.10 Application area · 3 min

The model uses three temperature levels on the target surface: the protector holds 1050 °C, the operating envelope ends at 1100 °C, and the interlock trips every source after 1 s above 1200 °C. Explain what each level is for, which of them a control engineer may tune, and two features of a real divertor target that the single-element model leaves out and that would move these numbers.

13.11 Application area · 4 min

In case 10 the stored-energy setpoint ramps toward 0.85 MJ. Under protected PI the heating saturates and the stored energy peaks at 0.720 MJ with the heat flux held at 4.5 MW/m^2. A planning engineer must choose one upgrade to raise the sustainable stored energy: more microwave sources, targets rated above 10 MW/m^2, or a larger allowance for core radiation. Which does the case support, and what will the upgrade cost in exhaust?

13.12 Review the agent · 5 min

A learner asks an AI assistant how to pass the challenge without spending more neon. The assistant replies: “Keep the Peak heat-flux setpoint at 7 MW/m^2 so the cascade seeds little neon. The run fails because the surface protector is too slow when the cooling falls, so raise the Heat-flux protector gain from 0.5 to 2.0. A faster protector will hold the surface and the stored energy together.” Find the flaw, and say how you would check it with Equation (13.9) and one run.