INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 18 Control of Nuclear and Fusion Systems

CHAPTER 14

Summary and Outlook

Pages 365–378

Edited by John D. Hedengren

How to cite this chapter. "Summary and Outlook," 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. 14, pp. 365–378.

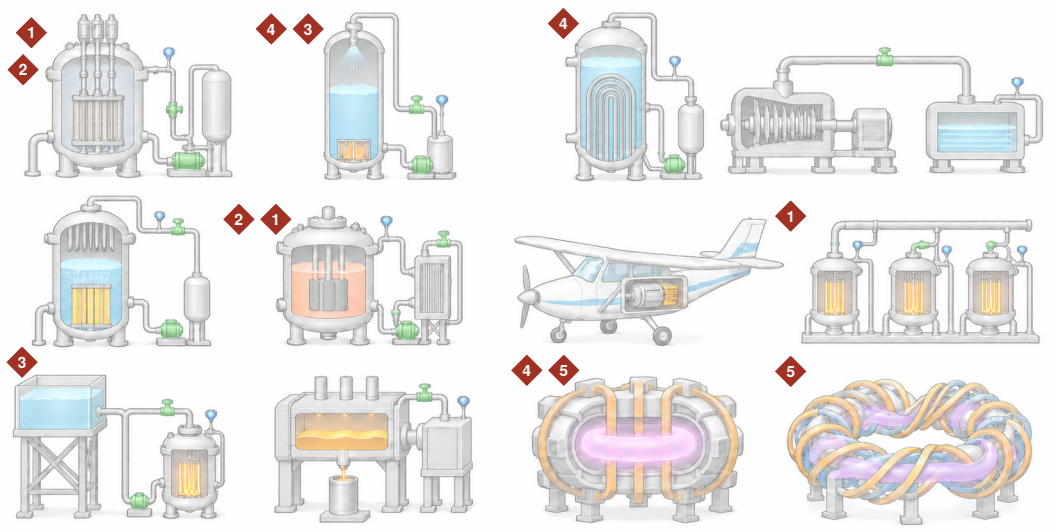
The page numbers in this file are those of the printed volume.

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

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

CHAPTER 14

Summary and Outlook



Open gaps 1 moving coefficients 2 unseen margins 3 hidden failures 4 layers that act 5 model error

Notice that the reactor with its rods, the salt reactor, the pressurizer and the tokamak each carry two gaps, and that no gap marks the steam generator, the boiling water reactor, the aircraft or the melter. Illustration, not to scale.

IN PLAIN TERMS **What Twelve Nuclear and Fusion Plants Taught**

The main idea. Each chapter of this book ran one nuclear or fusion plant in a simulator, from the rods of a reactor to the hot edge of a fusion machine. In most of them the physics already steadied the plant, and the controller removed what was left.

The challenge. Some limits arrive faster than any reading can warn, and some margins never show on a gauge. A loop can hold its reading steady while the plant quietly spends its reserve.

Why it matters. The best plants measured the cause of a disturbance, let the protective rules act first, and tracked the margin left.

Twelve chapters controlled twelve plants, from a pressurized water reactor's control rods to the heat exhaust of a stellarator. Each plant was read in three parts in Section 1.4.7: physics that helps or hurts, a limit set by protection or equipment, and the control layers between them. Table 1.1 listed the three parts for every chapter. This chapter adds what each arrangement achieved and where it failed, draws a short guide for choosing a method, looks at where practice is heading, reports how the volume itself was made with AI agents and what they got wrong, and lists the problems that remain open.

14.1 What the Chapters Showed

Tables 14.1 and 14.2 compare the chapters. The second column condenses Table 1.1. The other columns name a structure that worked, one that failed, and the measured result with the table or section of the chapter that reports it.

The columns make five points, and Figure 14.1 draws one pair of results from each row.

Table 14.1 Physics, Limits and Layers, Chapters 2 to 7

Chapter	Physics and limit	What worked	What failed	Measured result
2, PWR rods	Feedback follows the load and leaves the offset to the rods. A thermowell, a deadband, an insertion limit.	Mismatch feedforward on a turbine ramp. A flux cascade on a dropped rod. A 0.8 °C deadband against noise.	$K_c = 8$, past the thermowell's margin. A 0.2 °C band with a minimum rod speed. Any tuning near the insertion limit.	Ramp IAE 41.1 against 10.1 °C min, rod travel 123 against 30.7 %. Dropped rod peak 5.93, 3.58 and 2.23 °C (Table 2.8). Bank at 26.8 % against a 30 % limit (Section 2.5.5).
3, pressurizer	An insurge compresses the bubble faster than full spray condenses. Relief valve and level limits.	Overrides first, then the split-range settings. The coolant's own compliance.	Tuning alone in the challenge. A high gain with no deadband, which cuts the peak at 8.9 times the travel.	Peak within 20 kPa for three architectures, time off level 1014 against 40 s (Table 3.4). Challenge worst error 1203 shipped against 27 kPa suggested (Table 3.5).
4, steam generator	Shrink and swell. The zero's crossing moves from 38 to 65 s as load falls. Two level trips.	Three elements with a fast inner flow loop. MPC forecasting the voids from the steam flow.	One element at a higher gain, which cycles at $K_c = 3.0$. The challenge as shipped, which trips high-high.	IAE 11486, 9094 and 6521 mm s for one, two and three elements (Table 4.6). Pickup peak 167.6 against 122.1 mm under MPC (Table 4.7).
5, steam dump	The shut-in secondary reaches its safety valve in 14.80 s against a 4 s RTD.	Turbine impulse pressure as feedforward. Two low-selected overrides and a reactor runback.	Feedback alone, at seven times the gain or with an instantaneous RTD.	Peak 83.37 feedback, 79.74 with feedforward, 76.39 bar coordinated (Table 5.4). Gain alone 81.73, instant RTD 82.87 bar (Table 5.5).
6, boiling water reactor	Voids are moderator, feedback and a level error at once. Level trips and a power-flow map.	Three-element feed. The six tunings and the feedforward gain restored together.	One element at the same gains. Gain scheduling switched on over the original gains.	Ramp worst level +287.4 against −63.0 mm (Table 6.6). Challenge worst level 410 as given, 47 corrected, 3097 mm scheduled, a scram (Table 6.8).
7, molten salt reactor	The pump sets the delayed neutron fraction, 419 pcm of 650 at design flow. A transport delay.	The fuel pump on outlet temperature. A lagged load feedforward. MPC with a margin limit.	The secondary flow pairing, with 16 s of dead time. A static feedforward, which raised the peak.	Peak 7.05 feedback, 7.42 static and 2.64 °C lagged (Table 7.7). Margin case 476.76 against 0 s below 415 pcm, IAE 299.24 against 2324.30 °C s (Table 7.10).

The physics does most of the work, and the controller removes what is left. With every rod parked, the reactor of Chapter 2 took its power from 100 to 89.9 % within eight minutes of a load cut and left the coolant 13.9 °C above program (Section 2.5.1). The aircraft reactor followed a takeoff with nobody driving, lagging by 90.9 points of rated power and then overshooting (Table 8.4). Natural circulation removed the decay heat of Chapter 10 with no pump. In each case the controller's task was the residue: an offset, an overshoot, an inventory. The same physics can turn. At a moderator coefficient of $+0.5\,\text{pcm}/^\circ\text{C}$ the rod loop's process gain rose by a factor of 2.8, and the shipped tuning swung 12.6 °C peak to peak (Section 2.5.6). A graphite coefficient of $+6\,\text{pcm}/^\circ\text{C}$ made the salt reactor unstable (Table 7.9). A design that rests on a physical property must keep measuring it.

The measurement or the transport, not the core, sets the gain. The rod loop's ultimate gain was 9.01 with a three-second thermowell and 37.3 without it (Table 2.3). The salt reactor's gain limit was set by its transport delay, and its secondary flow pairing waited 16 s before the outlet moved

Table 14.2 Physics, Limits and Layers, Chapters 8 to 13

Chapter	Physics and limit	What worked	What failed	Measured result
8, aircraft reactor	The core follows a takeoff unaided, then overshoots. Relative gain 0.86 at takeoff, 3.58 at idle.	Preview of the published mission. A coordinated predictive controller on the go-around.	The same controller with the preview off. Two independent loops on the go-around.	Takeoff IAE 6358 without preview, 729 % s with it (Table 8.4). Go-around peak outlet 859 two loops, 815 °C coordinated (Table 8.5).
9, three SMRs	The header holds 7.55 s of steam, the modules answer in 52.5 s. A coolant ceiling that moves.	Coordinated allocation with a 2 s bypass. A margin-proportional split. A turbine runback.	Master–slave on a large ramp. An equal megawatt split. A scram with no runback.	Ramp: master–slave 2.647 to 3.874 MPa and the customer lost, coordinated 3.166 to 3.449 MPa (Table 9.6). Margins 8.5, 4.9, 0.0 against 5.41, 5.66, 4.59 °C (Section 9.6).
10, decay heat	Natural circulation removes a power-law heat. The pool pays 165.2 kg per kelvin of cooldown avoided.	A margin override on a 2.0 MW auxiliary train with a setpoint pair.	A latched-open condenser. Throttling alone, with no auxiliary train.	Pool at 8 h 2.67, 7.61 and 31.76 t, tubes uncovered at 3.93, 6.26 h and never (Table 10.5). Challenge 69.9 and 120.8 minutes over the limit, overrides none (Table 10.7).
11, waste melter	A cold cap insulates the pool. Joule heating unstable at a fixed tap. Coverage and off-gas limits.	Power regulation. Feed under coverage control with overrides on a low selector. Feedforward to the exhauster.	A fixed voltage tap, which walks the pool. Independent loops, which overrun the coverage ceiling.	Glass 1968 against 2444 kg, 188.3 against 0 min over the ceiling, worst pressure 90.0 against 54.2 Pa with feedforward (Table 11.5).
12, tokamak plasma	An unstable mode at 356 s^{-1}. A gain floor, a latency ceiling and a recoverable region.	A gain between floor and ceiling. A model observer for velocity. Priority for the vertical loop.	$K_p = 0.3$, below the floor. 10 V/mm at 300 µs latency. Clipped allocation.	$K_p = 0.3$ lost at 166.6 ms. Limit cycle of 1.45 mm rms. Coil energy 33.9 against 9.1 J. Clipping lost at 72.7 ms, priority held (Table 12.4).
13, stellarator exhaust	Neon protects the targets and poisons the core. Stored energy costs heating.	A core-radiation override. MPC raising seeding with heating. A lower heat-flux setpoint.	Plain PI, which hides a collapse for 18.0 s. Protected PI at its shipped setpoint in the challenge.	Challenge in the envelope 86.7 PI with an interlock, 66.4 protected and 100 % MPC. Peak heat flux 4.53 to 4.66 against 6.43 to 7.53 MW/m^2 (Table 13.4).

at all (Table 7.4). A single-element drum loop stopped improving at $K_c = 2.0$ and cycled at 3.0, on a disturbance that needed a five-point valve move (Section 4.6). The tokamak's position loop settled into a 1.45 mm limit cycle at 10 V/mm with 300 µs of latency and held to 0.059 mm at 100 µs (Table 12.4). In every one of these plants the first number to find is the lag between the plant and the controller.

A disturbance faster than its measurement needs a measurement of the disturbance. The steam dump of Chapter 5 has fifteen seconds. Seven times the gain reached 81.73 bar and an instantaneous RTD 82.87 bar, while the turbine impulse pressure wired to the same loop reached 79.74 bar and lifted no safety valve (Tables 5.4 and 5.5). Steam flow in a three-element loop, the power mismatch on a turbine ramp, a lagged load signal in the salt reactor and the published mission of the aircraft did the same job. The point of entry decided the structure. On the turbine ramp the mismatch feedforward cut the rod travel from 123 to 30.7 %, and on the dropped rod the flux cascade gave the smallest peak, 2.23 against 3.58 °C, because the excore detector saw the rod first (Table 2.8).

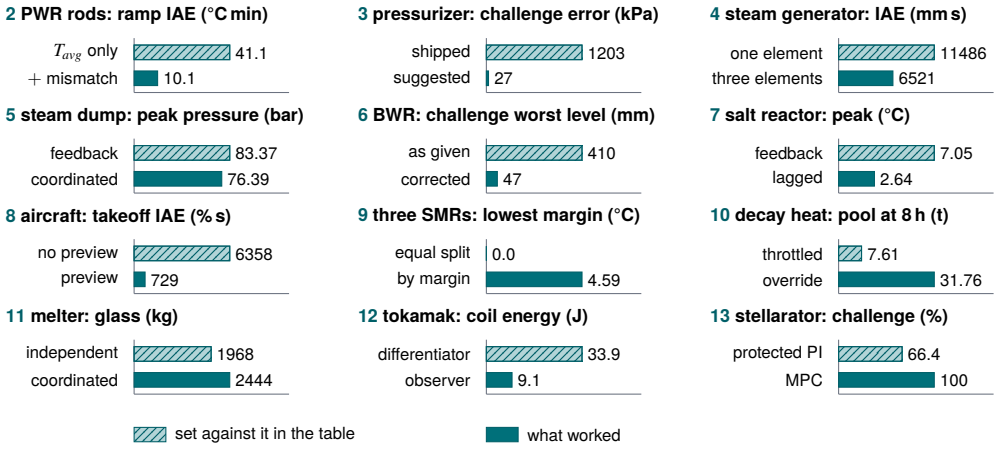


Figure 14.1 One pair of results from each chapter, with the numbers as the "Measured result" column of Tables 14.1 and 14.2 prints them. Each panel has its own scale from zero. A larger value is better in the panels of Chapters 9, 10, 11 and 13, and a smaller one in the others. The three-module panel sets the lowest of the three coolant margins, and the equal split's 0.0 °C draws no bar. Notice that in most rows the better result came from a change of structure or measurement, not of tuning alone.

The limiting layer is structure, and tuning cannot replace it. Across the pressurizer's three architectures the insurge peak moved by 20 kPa. In the challenge the suggested split-range settings on the two-loop structure still ended 320 kPa low, and the same settings with the overrides passed (Section 3.7). The coordinated steam dump peaked at 76.39 bar, and its reactor runback was worth 2.71 bar of its 3.35 bar advantage (Section 5.6). On the steam header the turbine runback was the one change both necessary and sufficient in the challenge (Section 9.6). The tokamak kept a plasma under priority allocation that clipping lost (Table 12.4), and the stellarator's core-radiation override took the seeding at 20.01 s and completed a pulse that plain PI lost at 37.985 s (Section 13.6). A layer that switched on without its tuning made things worse: gain scheduling over the boiling water reactor's original gains scrammed the plant (Table 6.8).

The binding limit is often a margin or a resource that no setpoint shows. The rod bank of Chapter 2 passed through its insertion limit at 26.8 % while the temperature was on program, because nothing in the objective mentioned shutdown margin. The salt reactor's predictive controller kept the delayed neutron fraction above 415 pcm for the whole margin case at 7.8 times the temperature error (Table 7.10). The condenser pool of Chapter 10 is spent and does not come back, and two passing configurations one setpoint apart finished 2.97 t of water apart (Table 10.7). An equal megawatt split on the steam header left one module at 0.0 °C of margin with the header pressure back in its band (Section 9.6). A pressurizer with a spray valve stuck open held its pressure within 4 kPa by running every heater bank, and only two scorecard numbers, the fighting time and the heater energy, showed that the loop had spent its upward authority (Section 3.7). A controlled variable on its setpoint said little about the margin left.

14.2 Choosing a Method

The chapters suggest an order of questions for a nuclear or fusion plant that is new to the reader. Each names the chapter where its answer decided the result, and Figure 14.2 draws them as a chart.

1. *What does the plant do with no controller?* Run it with the actuators parked and size the residue, as the rod and aircraft chapters do. The residue is the controller's job.

2. *Does the helpful physics keep its sign over the fuel cycle, the load range and the fouling?* If not, schedule, retune or design across the range, and measure the range first (Chapter 2, Chapter 7).
3. *What lag sets the gain ceiling?* A thermowell, a transport delay, a moving zero or a latency each bounds the gain before the core does (Chapter 4, Chapter 12).
4. *How many seconds to the nearest protection limit, and which signals act inside them?* Count them with Equation (1.8). If no measurement of the effect is fast enough, measure the cause (Chapter 5).
5. *Where does each disturbance enter?* Feedforward from a disturbance that arrives through the heat path, a cascade on one that arrives as reactivity (Chapter 2).
6. *Which limit must win, and what enforces it?* Build the override, runback or priority allocation before tuning the loop beneath it (Chapter 3, Chapter 9, Chapter 12).
7. *Which margin or resource does no setpoint show?* Put bank margin, pool inventory, delayed neutron margin or actuator authority on the scorecard, and decide whether it becomes a constraint (Chapter 10, Chapter 7).

A predictive controller paid where it carried a forecast or coordinated actuators that the loops ran apart. Forecasting the voids from the measured steam flow, it held a pickup to 122.1 mm where three-element control reached 167.6 mm (Table 4.7). It cut the boiling water reactor's level dip on a pressure step from 114.9 to 17.7 mm, at 2.5 times the integrated pressure error (Table 6.10). Raising seeding with heating, it held the stellarator's peak heat flux near 4.5 MW/m^2 and passed the challenge as shipped (Table 13.4). It lost or tied elsewhere. On the pressurizer's insurge it used less spray and heat than the overrides for a 14.7 % higher IAE (Section 3.7.5). On the steam header it tied coordinated allocation, 71.61 against 72.56 MPa s (Table 9.7). In the melter it used less energy and made less glass on the loaded scrubber, and no controller won every constraint (Table 11.6). The tokamak's explicit MPC settled an ELM in 0.52 ms against 1.32 ms but recovered kicks only up to 19.63 mm against 19.91 mm, because it reversed the voltage late (Table 12.5). Volume 9 treats that choice in full.

14.3 Where Practice Is Heading

Nuclear plants that follow the grid. Nuclear plants are commonly run at full output whenever online, although they can ramp, follow load and provide reserves. A unit commitment and dispatch model that represents xenon transients and the change of core reactivity over the fuel cycle found, for a system with large shares of solar and wind power, that flexible nuclear operation lowered operating costs, raised the reactor owner's revenue and reduced the curtailment of renewables [1]. The rod chapter shows what such a schedule must respect. Its sixteen-hour maneuver drove the bank fully out for thirteen minutes on the return to full power, and a cut to 40 % crossed the insertion limit (Section 2.5.5). In the editor's view the xenon, boron and insertion limits will move from the rod controller into the dispatch that asks for the maneuver.

> ⧖ HISTORY **1979 to 2011: Instruments for the Worst Day**
>
> The second history box of Chapter 1 ended in 1979, with a commission that found instruments not designed to follow an accident. The accident at Fukushima Daiichi began on 11 March 2011. There the loss of instrumentation and control left the operators with little indication of actual plant conditions, and the accident guidelines did not cover the loss of the instruments that display the key parameters. Not knowing the state of the spent fuel pools led to an effort to add water to them. The report of the Director General of the International Atomic Energy Agency concluded that the instrumentation and control needed in accidents beyond the design basis must remain operable, and that strategies for manual control of vital equipment need to improve [2]. In the editor's reading the lesson of 1979 was learned again, for a plant without power.

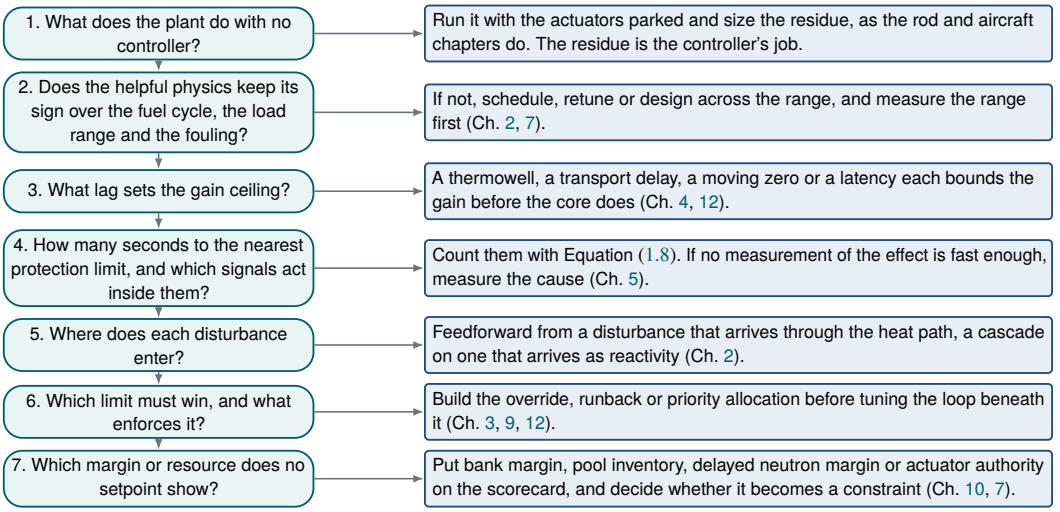


Figure 14.2 The seven questions of this section in order, with the list's answer and the chapters where it decided the result. Notice that the plant with no controller and the lag come before any question about structure, and that the last question is about what no setpoint shows.

Rod control that adapts to the core. From 2008 to 2015 an adaptive reactor control system was installed in several pressurized water reactors in Germany. It receives a new set of reactivity coefficients for every core reload, determined for the fuel cycle as a function of the boron concentration, and balances reactivity, xenon included, to keep a margin to the bank limits that assure shutdown reactivity and the insertion limits. Its more precise control allows reduced dead bands near full power [3]. The same guide notes that unplanned maneuvering can move the power and burnup profile of a core away from its predictions, so the core physics calculations may need to be checked again after each maneuver. The rod loop of Chapter 2 runs on one fixed deadband, and nothing in its objective mentions shutdown margin. In the editor's view the next step is coefficients estimated from plant data between reloads, with their uncertainty stated.

Autonomy for multimodule plants. A framework for the autonomous control of advanced reactors starts from multiple units on one site and from remote small modular reactors, whose economics depend on lower operating and maintenance costs. It calls for diagnosis, planning, reconfiguration and self-validation inside the control system, and it describes an autonomous control system for nuclear power as a goal not yet realized [4]. The steam header of Chapter 9 is a small instance of that problem. Its allocator splits load against margins it estimates from measurements, and its runback sheds turbine load after a scram so that the industrial customer keeps its steam. In the editor's view an allocator of this kind will have to report the uncertainty of each module's estimated limit before an autonomous system can be trusted to act on it.

Instruments watched while they run. Online condition monitoring detects and diagnoses abnormalities from the process signals while the plant operates. Its intent is to mine the recorded data for changes in the relationships between measured parameters, which the indicators of a typical control system do not include. Shown that a sensor has not drifted since its last calibration, a plant can extend the calibration period [5]. The scorecards of this volume did the same on a small scale. The fighting time and the heater energy of Chapter 3 showed a stuck spray valve that the pressure trend hid, and the pool inventory of Chapter 10 separated two configurations that the temperature scorecard could

not. In the editor's view such measures of spent authority belong beside the drift checks of an online monitoring system.

Plasma control that must work on the first discharge. A review of plasma control for ITER finds its design driven by the nuclear mission, by aggressive commissioning constraints and by a limited number of discharges, with little time to tune algorithms on the machine. It names actuator sharing and fault handling among the novel problems [6]. Priority allocation in Chapter 12 is a two-loop case of actuator sharing, and the stellarator's low selectors hand the heating to a protector and the seeding to a core override as limits approach. In the editor's view fusion controllers will be designed and tuned almost entirely in simulation, which makes the model's stated recoverable region the quantity to verify.

Learned controllers and predictors. A controller trained by deep reinforcement learning has commanded the full set of control coils of a tokamak, producing and holding a diverse set of plasma shapes while meeting physical and operational constraints, with a notable reduction in design effort [7]. A deep learning method has forecast disruptions and made reliable predictions for machines other than the one it was trained on, which a large reactor needs because it cannot afford training disruptions [8]. The vertical loop of Chapter 12 gives the yardstick. Any controller, learned or written, is bounded by the same gain floor, latency ceiling and recoverable region. In the editor's view a learned controller will be accepted on the kicks and latencies it was not trained on.

Exhaust under feedback. Detachment is required to keep heat and particle fluxes on the divertor tolerable, and feedback control of the location of an impurity emission front has been demonstrated with real-time diagnostics, a dynamic characterization of the divertor plasma and an offline controller design [9]. The stellarator of Chapter 13 adds the core to that problem, since the same neon that detaches the divertor collapsed the plasma 18.0 s after a valve was held open. In the editor's view exhaust control will be designed as a constrained multivariable problem with core radiation as a hard limit.

Long pulses with the exhaust under feedback. Wendelstein 7-X is designed for discharges of up to 30 min and 18 GJ of heating energy. A new feedback loop holds its total radiation at a target by changing the voltage of the seeding valve from bolometer measurements, and a second raises the remaining microwave sources in real time when one of them trips. With both, a discharge ran 363 s with nearly constant heating and radiation and 1.8 GJ of heating energy, until arcing caused by cooling problems in a heating beamline ended it [10]. The same report notes that feedback is not fundamentally required for plasma control in a stellarator, whose magnetic configuration is steady by design. Chapter 13 works with the same two actuators, the seeding valve and the heating, and adds the core limit that a long pulse must respect. In the editor's view the pulses of a reactor-scale stellarator will depend on how the exhaust loops share those actuators.

Figure 14.3 sets these developments in three threads.

14.4 Engineering With AI Agents

An AI agent is a language model that works toward a goal in a loop. It plans, calls tools such as a compiler, a simulator or a browser, reads the results and continues [11]. In the editor's view, agents can now carry out much of the work that used to fill the hours of a control project. They derive a model from balances or fit one to plant data. They write the simulator and its tests, code the controller, tune it against stated criteria, script commissioning checks, write the documentation and watch running loops for degradation. The apps and chapters of this series were produced that way under the editor's direction, and the account that follows this section reports what the agents got wrong in this volume and how it was caught.

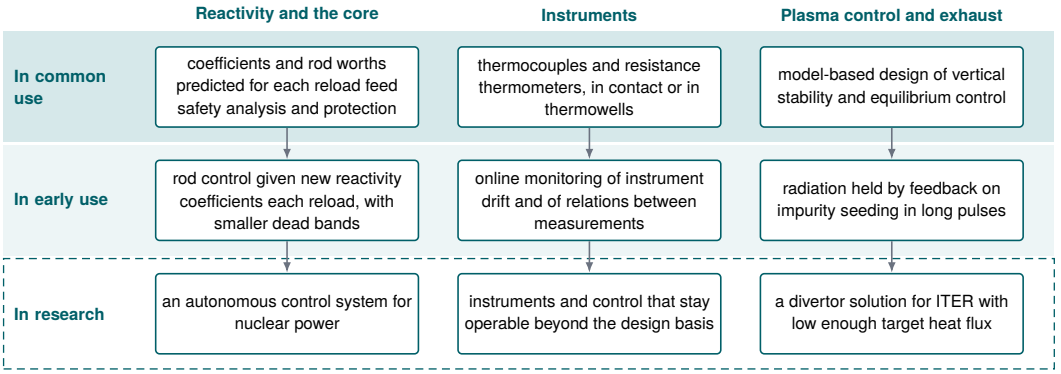


Figure 14.3 Developments in three threads. The top lane holds what the cited source describes as usual or routine practice, or as the most common, the middle lane what it reports in service at some plants or used operationally on one machine, and the bottom lane what it reports as a goal or a need not yet met, in the editor's reading. Sources by thread: reactivity and the core [3, 4], instruments [2, 5] and plasma control and exhaust [6, 10]. Notice that each thread moves from what is computed or measured before operation toward what must be known while the plant runs.

In the editor's view, this changes where an engineer's value lies. When a model, a controller and a report can each be produced in minutes, producing them is no longer the scarce skill. The work moves to five things.

Quality control. An agent's output arrives as a claim: this model fits, this tuning settles in four minutes, this code implements the equation. The engineer checks claims against the plant or the simulator. A reported number is measured again on a fresh run at the stated seed. The code is read for what it clamps, what it initializes and what it reads, because a display that shows the controller's model in place of the plant looks correct until the two differ. A tuning is tried on a disturbance the agent did not test. The most useful checks are independent of the author, made by a second agent or a person who has not seen the first answer, and they are checks that can fail.

Supervision. The engineer sets the scope of each task and decides what an agent may touch. Reading a historian and proposing a tuning is one level of access. Writing a setpoint, downloading a configuration or changing an interlock is another, and it passes through the plant's management of change with a named approver. Supervision also fixes the order of work: a baseline before a change, one change at a time and a comparison that shows what moved.

Creativity. An agent answers the question it is given. Choosing the question remains the engineer's task: which loop is worth improving, whether a disturbance calls for feedforward, a cascade or a change to the process, which step test the operation can tolerate, and which comparison would show that a predictive controller earns its place over a retuned PID loop. The experiment that separates two explanations is designed by someone who can imagine both.

Process insight. A model contains what its author put in. The equipment does more. A valve sticks, the dead time of an analyzer grows when its sample flow drops, a thermowell lags, a heat exchanger fouls and an operator intervenes at the same hour each night. The engineer who knows the unit recognizes a result that is too good, asks what the model left out and knows which measurement to distrust first.

Accountability. Responsibility does not transfer to a tool. Functional safety standards assign the activities of the safety lifecycle to persons and organizations of stated competence, and they require an assessment that includes someone independent of the design team [12]. An agent may draft the hazard analysis or the test procedure. A person owns the decision to accept it, to start up and to keep running.

These five are the professional form of the habit introduced in Section 1.7: check before trusting, and know enough to check.

14.4.1 How This Volume Was Made

Each of the twelve apps was built and tested by agents. An agent then wrote the paper on it, keeping verification notes that trace each number to a run. Eight of the twelve referee reports call themselves self-reviews, and the other four say they were written against the draft and applied or describe a reviewer's stance. None names a reviewer independent of the writing. Seven carry a later pass on predictive-control material, two of them called independent. Eleven papers were measured again in a consistency round and an app-fix round, and the twelfth was written after both. The defect logs hold 125 entries with a severity, in eleven of the twelve logs, since the tokamak's log numbers its findings without one. Of these, 30 are wrong output or behavior, 71 documentation or content, 15 cosmetic or display and 9 something else, such as a completion rule or a design question. Some round notes list fixes without a severity, so the count is a floor. The editor set the scope of each round and what each agent could change. Four cases show what went wrong and what caught it. In each, a layer between the physics and the limit was not doing what its text said.

Overrides that never acted. The pressurizer's texts said the high-pressure override held the charging valve shut through the peak, and counted it among three things that win the challenge. Measured, the selector was active for 0.0 s in the challenge, and an override setpoint raised out of reach reproduced the run to the last digit. The insurge raises the level, so the level loop has already shut the charging valve before the pressure reaches the override. The steam dump's level override never selected in twelve scenarios, because on a load rejection the pressure only rises and the level cannot swell. Agents measuring the apps' claims found both. The first was confirmed by raising the override out of reach and seeing nothing move, the second by lowering its setpoint until it acted. The texts now say when each override acts.

A pressure loop that sprayed while the pressure was low. After an outsurge 1 MPa below setpoint, the pressurizer's back-calculation with a 30 s tracking time wound the integral to $+594\,\%$ against a proportional and derivative sum of $-825\,\%$. As the error shrank the output crossed the deadband, and the spray was commanded for 589 s with the pressure more than 70 kPa low. A consistency round that measured every scenario again found it. Tracking at the integral time cut the figure to 12.2 s, and the same fix was then needed on the override, which had taken the charging valve with the pressure 836 kPa below its own setpoint.

Controllers that knew the plant. Six entries in four logs record a controller or an operator aid that read the run's hidden parameters or the plant's true state. The aircraft reactor's dynamic matrix controller and preview feedforward solved their models on the run's parameters and scaled the prompt jump by the true neutron power. With the instruments untouched, 16 of 44 parameters scaled by 1.25 moved its gains, and the fouling scheduled in the challenge reached the feedforward as a better model. Fixing it raised the takeoff power error from 634 to 729 % s, so the predictive controller now trails the two loops there. The steam header's allocator read each module's hidden thermal ceiling, and a hidden change moved its split by up to 57.310 MW. The tokamak's observer read the margin drop of a minor disruption that no instrument measures. Each entry records the same probe: move a hidden value with the measurements held and ask whether the controller moves.

Causes the runs did not support. The book round tested each chapter's causal sentences by run and corrected six in four chapters without adding a number. The boiling water reactor's runback dip was said to come from fewer voids plus a feedwater cut. At the dip the core void had risen from 40.3 to 43.2 % and the vessel had gained 1.5 t, so the dip is the steam held below the surface falling with steam production. The steam dump's long time off program was said to be overcooling, yet every one of its 114.8 seconds had the coolant above its reference. The tokamak's 20 mm kick was said to be lost at any gain, and a velocity gain of 2 V per m/s held it on the same supply. That last claim had reached the app's briefing and Chapter 1.

The book round. The twelve chapter records list 140 rows of numeric verification, each a stated value or group of values set beside its reproduction. Eight report no drift, and four name a printed value that a run does not reproduce exactly. On the steam header the margins of the margin-proportional split had moved after a fix to the allocator, and a later repair round measured them again. The writer of Chapter 1 read the apps' action flags and found five places in three chapters that named a loop's action against the convention of Section 1.3, and its own reactor example first settled 17 % low until it was checked against its closed form. Three statements of Chapter 1 were corrected from the chapter reports. Every chapter pass logged at least one defect in its app. The melter's pass found that after an exhauster trip the low selector and the pressure feedforward chattered, moving the feed command between about 630 and 1800 kg/h from one scan to the next. The editor approved a repair of the model and its results were regenerated. Eleven chapters of this volume were then synchronized with the repaired apps and papers, and the melter's exercises were run again. With the repair the loaded-scrubber ranking reversed, and coordinated control now makes more glass than the predictive controller.

Every error above was caught by a check that could fail: a layer switched off to see what moved, a scenario measured again, a hidden value moved while the measurements were held, a causal sentence tested against the variable it names. None was caught by rereading the text that contained it. Four working rules follow. Disable each protective layer once and confirm that it acts on the event it guards. Check what an integral does while its output is clamped. Give a controller only what the plant measures and the design declares. Test every causal sentence with the variable it names.

14.5 Gaps in the Control of Nuclear and Fusion Systems

Five gaps belong to this volume's topic, and the gaps common to agent-era practice follow them in Section 14.5.1. In the editor's view none is closed. Each item gives the chapters' evidence, then the state of the literature with its sources. Why each gap is hard and what progress would look like are the editor's view. Figure 14.4 places the gaps by the chapters that met them and by the maturity of the method available, and the same numbers mark the plants in the picture that opens this chapter.

1. *Coefficients that move.* The rod loop's process gain rose 2.8 times when the moderator coefficient went slightly positive, and the salt reactor became unstable at a graphite coefficient of $+6\,\mathrm{pcm}/^{\circ}\mathrm{C}$. The steam header now estimates each module's ceiling from measurements because the fouling moves it. A guide to flexible operation describes rod control that receives a new set of reactivity coefficients for each core reload, and notes that unplanned maneuvering can move the core away from the predictions that its physics calculations rest on [3]. The gap is hard because a coefficient shows only through the loop it changes, and a licensed plant cannot be stepped at will to measure it. Online estimates of reactivity coefficients and thermal limits that state their own uncertainty, and a licensed way to schedule a loop on them, are not routine.

2. *Margins that the controller cannot see.* The rod bank crossed its insertion limit with the temperature on program, and the salt reactor held its delayed neutron margin only by accepting 7.8 times the temperature error. The adaptive rod control of Section 14.3 balances reactivity, xenon included, to keep a margin to the bank limits that assure shutdown reactivity and the insertion

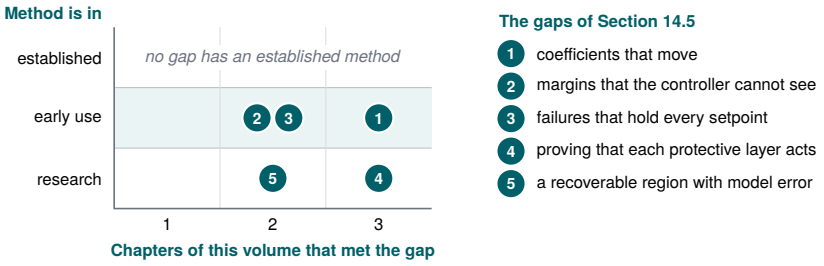


Figure 14.4 The five gaps of this section. Across is the number of chapters that each item names as evidence: three for gap 1 (Chapters 2, 7 and 9), two for gap 2 (2 and 7), two for gap 3 (3 and 10), three for gap 4 (3, 5 and 12) and two for gap 5 (12 and 13). Up is the maturity of the method available, the editor's reading of the sources cited in each item and in the Open challenge box. Notice that the two gaps about proof, of a protective layer and of a recoverable region, sit lowest.

limits [3]. The gap is hard because the margins of a safety analysis are computed in advance for assumed conditions, while a controller acts on the present ones. Methods that put the margins of a safety analysis into a control layer's objective, while keeping that layer outside the safety system, are not settled.

3. *Failures that hold every setpoint.* Two lost heater banks changed nothing on the pressurizer's trends until an outsurge asked for heat. A quarter of the decay-heat pool was lost with no temperature showing it. Online monitoring mines recorded data for changes in the relationships between measured parameters, relationships that the indicators of a typical control system do not include [5]. The gap is hard because a loop that holds its setpoint hides the effort it spends, and that effort sets off no alarm. A standard set of authority measures, such as fighting time, actuator reserve and inventory used, computed beside the error, is not yet common practice.

4. *Proving that each protective layer acts.* The level cap on the steam dump never selects in any of its chapter's twelve scenarios, because a load rejection there can only shrink the level (Section 5.6). The pressurizer's high-pressure override takes the charging valve in an early proportional hand-over, before the pressure reaches its setpoint (Section 3.5.2). The tokamak's verdict checks every physics step, because brief excursions fall between the 2 kHz log samples. Test generation that forces every override, runback and interlock to act, and checks the record at the rate of the event, is not routine. The Open challenge box after this list takes the gap further.

5. *A recoverable region with model error.* The vertical loop's region was computed for a rigid, massless plasma with one wall mode, and the stellarator's model is calibrated, not validated. On one tokamak the largest controllable displacement, measured by switching vertical control off for varying intervals, was 0.15 to 0.24 m against about 0.40 m calculated, a difference far beyond the noise and more likely due to errors in modeling the non-axisymmetric structures and to the plasma striking the wall [13]. ITER permits only 1 to 10 unmitigated vertical displacement events at its highest current and stored energy [6]. The gap is hard because the measurement that settles the region is the event the design must avoid. A recoverable region stated with the model's uncertainty, which a controller designed in simulation could be held to on the first discharge, is missing.

⚠ **OPEN CHALLENGE** **Showing That a Protective Layer Acts**

Three chapters met gap 4, as many as any gap. A guide to nuclear instrumentation and control notes that digital systems may be too complex to test in all aspects, and that a common cause failure that leaves the outputs as they are is not revealed until a safety action is demanded [5]. A review of plasma control for ITER requires every control algorithm to be verified in simulation under the expected noise, disturbances and key off-normal events before use [6]. In the editor's view the gap is hard because a layer written for a rare event is exercised only by that event. Progress would be a test that shows each layer acting on its own event, run again after every change.

14.5.1 Gaps Common to Agent-Era Practice

The final chapter of each volume lists gaps specific to its topic. Seven more are common to every topic in the series. In the editor's view none of them is closed.

1. *Verifying agent-written control code and models before they reach a plant.* Tests show that a controller behaves on the cases tested. For controllers learned from data, guarantees of stability and constraint satisfaction rest on assumptions about the model and its uncertainty, and open problems remain [14]. In the editor's view, code written by an agent has no comparable body of method yet. Practice needs a routine path from an agent's draft to code a plant can trust: independent derivation, tests of what the code reads and clamps, and trials with hardware in the loop.

2. *Trustworthy models from sparse or poor plant data.* Plant records are short, taken in closed loop, noisy and marked by unmeasured disturbances. A fitting routine returns a model from any record. What is missing is a dependable statement, delivered with the model, of where it can be trusted, and a method that says so when the data cannot identify a parameter.

3. *Safety cases and standards for agents in commissioning and operation.* Functional safety standards define a lifecycle of specified requirements, verification and assigned human roles [12]. In the editor's view, how an agent that writes configuration, runs commissioning checks or advises an operator is qualified, limited and audited under them is not yet settled.

4. *Security of agent access to control systems.* An agent with credentials to a control system is a new path into it. Industrial security standards state the requirements a control system must meet at each security level [15]. An agent that reads documents, messages or web pages can also be steered by text placed there. In the editor's view, least privilege, read-only access by default and human approval of every write are a starting point and not a solution.

5. *Benchmarks and reproducibility.* An agent asked twice may deliver two different controllers, and a result reported without its plant, seed and scoring cannot be checked. Control needs shared benchmark plants with fixed disturbances and published measures, on which agent-built and conventional designs are compared like for like. The apps of this series, with a default seed and an exported record of every run, are a start.

6. *Transfer from simulation to hardware.* A controller developed on a simulator inherits what the simulator omits. The robotics literature documents the gap between simulated and real performance and the methods used to narrow it, such as randomizing the simulated plant and adapting on the real one [16]. Process plants add slow dynamics, rare upsets and little freedom to experiment, so each transfer still needs a staged plan from the simulator to a bench rig to the unit.

7. *Keeping engineers' judgment sharp when agents do the routine work.* Automation removes the routine work through which skill is built and leaves the person the hard cases, an irony described for process operators in 1983 [17]. The same now applies to engineers. An engineer who has never tuned a loop by hand or closed a balance on paper is poorly placed to review an agent that

has. Deliberate practice, of the kind the exercises in this book provide, has to be designed into education and into work.

KEY POINTS

- The physics does most of the work, and the controller removes what is left.
- In every one of these plants the first number to find is the lag between the plant and the controller.
- A disturbance faster than its measurement needs a measurement of the disturbance.
- The limiting layer is structure, and tuning cannot replace it.
- A controlled variable on its setpoint said little about the margin left.

References

[1] J. D. Jenkins, Z. Zhou, R. Ponciroli, R. B. Vilim, F. Ganda, F. de Sisternes, and A. Botterud, "The benefits of nuclear flexibility in power system operations with renewable energy," *Applied Energy*, vol. 222, pp. 872–884, 2018.

[2] International Atomic Energy Agency, *The Fukushima Daiichi Accident*. Vienna, Austria: International Atomic Energy Agency, 2015, sTI/PUB/1710, report by the Director General with technical volumes 1–5.

[3] ——, *Non-baseload Operation in Nuclear Power Plants: Load Following and Frequency Control Modes of Flexible Operation*, ser. IAEA Nuclear Energy Series. Vienna: International Atomic Energy Agency, 2018, no. NP-T-3.23.

[4] R. T. Wood, B. R. Upadhyaya, and D. C. Floyd, "An autonomous control framework for advanced reactors," *Nuclear Engineering and Technology*, vol. 49, no. 5, pp. 896–904, 2017.

[5] International Atomic Energy Agency, *Core Knowledge on Instrumentation and Control Systems in Nuclear Power Plants*, ser. IAEA Nuclear Energy Series. Vienna: International Atomic Energy Agency, 2011, no. NP-T-3.12.

[6] D. Humphreys, G. Ambrosino, P. de Vries, F. Felici, S. H. Kim, G. Jackson, A. Kallenbach, E. Kolemen, J. Lister, D. Moreau, A. Pironti, G. Raupp, O. Sauter, E. Schuster, J. Snipes, W. Treutterer, M. Walker, A. Welander, A. Winter, and L. Zabeo, "Novel aspects of plasma control in ITER," *Physics of Plasmas*, vol. 22, no. 2, p. 021806, 2015.

[7] J. Degrave, F. Felici, J. Buchli, M. Neunert, B. Tracey, F. Carpanese, T. Ewalds, R. Hafner, A. Abdolmaleki, D. de las Casas, C. Donner, L. Fritz, C. Galperti, A. Huber, J. Keeling, M. Tsimpoukelli, J. Kay, A. Merle, J.-M. Moret, S. Noury, F. Pesamosca, D. Pfau, O. Sauter, C. Sommariva, S. Coda, B. Duval, A. Fasoli, P. Kohli, K. Kavukcuoglu, D. Hassabis, and M. Riedmiller, "Magnetic control of tokamak plasmas through deep reinforcement learning," *Nature*, vol. 602, no. 7897, pp. 414–419, 2022.

[8] J. Kates-Harbeck, A. Svyatkovskiy, and W. Tang, "Predicting disruptive instabilities in controlled fusion plasmas through deep learning," *Nature*, vol. 568, no. 7753, pp. 526–531, 2019.

[9] T. Ravensbergen, M. van Berkel, A. Perek, C. Galperti, B. P. Duval, O. Février, R. J. R. van Kampen, F. Felici, J. T. Lammers, C. Theiler, J. Schoukens, B. Linehan, M. Komm, S. Henderson, D. Brida, and M. R. de Baar, "Real-time feedback control of the impurity emission front in tokamak divertor plasmas," *Nature Communications*, vol. 12, no. 1, p. 1105, 2021.

[10] 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.

[11] L. Wang, C. Ma, X. Feng, Z. Zhang, H. Yang, J. Zhang, Z. Chen, J. Tang, X. Chen, Y. Lin, W. X. Zhao, Z. Wei, and J. Wen, "A survey on large language model based autonomous agents," *Frontiers of Computer Science*, vol. 18, no. 6, p. 186345, 2024.

[12] International Electrotechnical Commission, "Functional safety – Safety instrumented systems for the process industry sector – Part 1: Framework, definitions, system, hardware and application programming requirements," IEC 61511-1:2016, Edition 2.0, 2016.

[13] 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.

[14] L. Brunke, M. Greeff, A. W. Hall, Z. Yuan, S. Zhou, J. Panerati, and A. P. Schoellig, "Safe learning in robotics: From learning-based control to safe reinforcement learning," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, pp. 411–444, 2022.

[15] International Electrotechnical Commission, "Industrial communication networks – Network and system security – Part 3-3: System security requirements and security levels," IEC 62443-3-3:2013, Edition 1.0, 2013.

[16] W. Zhao, J. P. Queralta, and T. Westerlund, "Sim-to-real transfer in deep reinforcement learning for robotics: A survey," in *Proc. 2020 IEEE Symposium Series on Computational Intelligence (SSCI)*, 2020, pp. 737–744.

[17] L. Bainbridge, "Ironies of automation," *Automatica*, vol. 19, no. 6, pp. 775–779, 1983.

14.6 Exercises

14.1 Discussion · 8 min

An agent reviewing Chapter 5 proposes to replace the coordinated structure with the feedback-only dump loop at a seven times higher gain and a faster temperature sensor, "because it has fewer parts to maintain". Using Tables 5.4 and 5.5, say what the proposal misses and what you would keep.

14.2 Discussion · 8 min

The rod loop of Chapter 2 and the pump loop of Chapter 7 both rest on reactivity feedback that helps. Using the moderator coefficient case of Chapter 2 and the graphite coefficient sweep of Table 7.9, discuss how you would keep these loops tuned over a fuel cycle, and what you would measure.

14.3 Discussion · 8 min

Section 14.4.1 reports two overrides in this volume that never acted on the event they were written for. Write a test plan, in five to seven steps, that an agent could run on any app or plant model to show that each override, runback and interlock acts when it should and stays out when it should.

14.4 Discussion · 8 min

In Chapter 10, Chapter 7 and Chapter 9 the binding limit is a margin or a resource rather than a setpoint. For each, say what quantity belongs on the scorecard, and argue whether it should become a constraint in a predictive controller or stay a number for the operator.

14.5 Discussion · 8 min

A vendor offers a learned controller for the vertical position loop of Chapter 12, trained in simulation. Using the floor, the ceiling and the recoverable region of that chapter and of Section 1.4.6, list what you would ask the vendor to show before the controller is tried on a plasma.

14.6 Discussion · 10 min

Section 14.4.1 reports six entries in which a controller or operator aid read the run's hidden parameters or the plant's true state. An agent is asked to commission the module allocator of Chapter 9 on a new plant. Write a procedure, in six to eight steps, that keeps the allocator honest, states what the agent may read and write, and says what a person must approve.