INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 1 Process Dynamics, Modeling, and Identification

CHAPTER 8

Grid Frequency Simulator for Inertia, Droop Offset and Integral Restoration

Pages 181–203

Edited by John D. Hedengren

How to cite this chapter. "Grid Frequency Simulator for Inertia, Droop Offset and Integral Restoration," in *Process Dynamics, Modeling, and Identification*, John D. Hedengren, Ed., vol. 1 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 8, pp. 181–203.

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

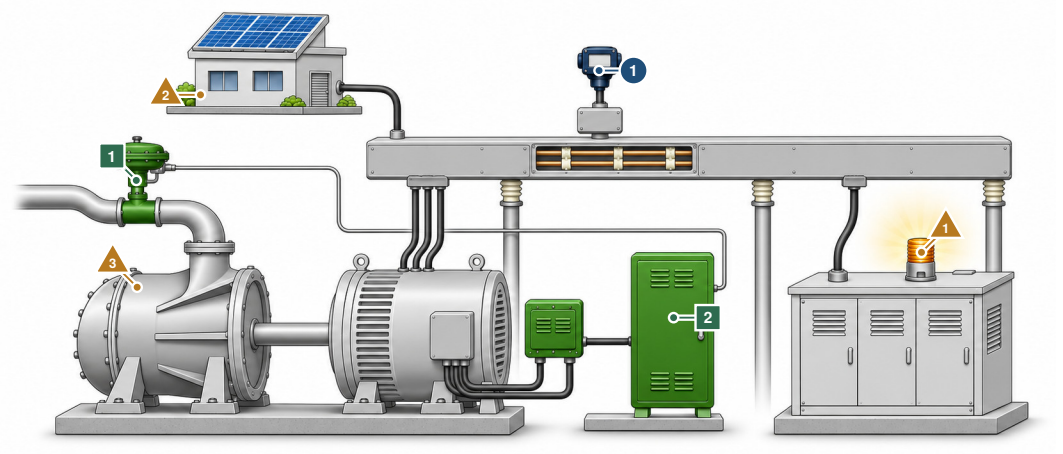
The app for this chapter
https://apmonitor.com/apps/grid-frequency/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v01-dynamics-modeling/#grid-frequency

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

CHAPTER 8

Grid Frequency Simulator for Inertia, Droop Offset and Integral Restoration



The controller moves 1 the steam valve, opened by the governor as the frequency falls
2 the generation command sent by the control center

The controller reads 1 the frequency of the bus

What disturbs it 1 a step in the load 2 a change in solar output 3 part of the generation lost when a unit trips

Notice that only the spinning turbine and generator store energy between the steam and the load, so every mismatch changes their speed and with it the frequency read on the bus. Illustration, not to scale.

IN PLAIN TERMS **Turning Together, Slowing Together**

The main idea. All the large power plants on a grid spin in step, and the speed they share is the grid's frequency. When more power is used than made, the machines give up some of their stored spin and slow down together.

The challenge. Each plant's governor, a controller on its turbine, opens the steam valve as the speed falls. This stops the fall but leaves the frequency a little low. A slower controller in a control center keeps adding power until the error is gone.

Why it matters. With fewer spinning machines the first fall comes faster, but the resting error and its cure stay the same.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact on lost generators (Section 8.1), the history of area control (Section 8.2), the try-it segment of scenario 7 (Section 8.4.1), the try-it of scenario 10 and the fact on a strong integral (Section 8.6.3) and the key points
Study	45 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the results (Section 8.6), Technology in Practice (Section 8.7), Limitations (Section 8.8) and the open challenge

The control question Grid frequency is the level of the kinetic energy stored in the spinning machines. After a load step, which parameters set how fast it falls, where droop arrests it and how integral action brings it back to 60 Hz?

Learning objectives

- Write the swing equation as an energy balance and explain why the frequency integrates a power imbalance.
- Compute the initial rate of change of frequency, the settled droop offset and the borrowed kinetic energy for a load step or a unit trip.
- Test the stability of the primary loop with the Routh condition and find the inertia below which it fails.
- Estimate the time constant of integral restoration and choose an integral gain that meets a restoration target inside a ramp limit.
- Explain why inertia, droop and integral gain each move one of the three numbers of a frequency event and leave the other two alone.

Prerequisites Sections 1.4.2 (balances), 1.4.3 (poles), 1.4.4 (integrating response) and 1.6.2 (velocity form), and Chapter 2 for the offset of a proportional loop.

The app https://apmonitor.com/apps/grid-frequency/

See also Volume 2, Chapter 5 for tuning a loop on an integrating inventory, Volume 2, Chapter 3 for a standing error from proportional control set beside one from an actuator bound. Volume 17, Chapter 7 for droop sharing, virtual inertia and a frequency-restoring integral loop on an islanded bus.

8.1 Introduction

An alternating-current power system stores almost no energy in electrical form. The energy that it does store is held in the rotors of the synchronized machines, and the size of that store is fixed by the speed at which the rotors turn. Generation and load are matched continuously by taking the difference out of those rotors or adding it to them, which slows the rotors or speeds them up, so the frequency of the interconnection is a direct measurement of an energy inventory. A tank whose outflow rises with head reaches a steady state without any control action, because the outflow is a function of the level. A rotating mass has no equivalent mechanism. A low rotor speed does not by itself raise generation or lower load beyond a small damping effect, so the inventory behaves as an integrator and the frequency ramps until a control action arrests it. It is the integrating tank of Section 1.4.4 with a small leak.

The first control action comes from the turbine governor, which is a proportional controller. Each governor measures shaft speed and opens its valve in proportion to the speed error, with a gain fixed by a setting called droop. The governor arrests a frequency excursion within a few seconds and leaves a steady-state offset, for the same reason that a proportional-only level controller leaves an offset, as in Chapter 2. A proportional controller requires a standing error to hold a standing output, and a load change that persists requires a valve position that also persists [1]. The offset is a structural property of the primary loop rather than a tuning defect, and no droop setting removes it.

The arrangement differs from most examples in a first control course in where the proportional controller is located. Droop is built into governor hardware and acts whether or not any computer is running, so the gain of the primary loop is a plant parameter rather than a value in a control system. The term that remains available to a control engineer is the integral term, supplied by automatic generation control (AGC) at a much slower cadence. The result is a two-tier structure in which the fast tier is not retuned by the operator and the slow tier has little effect during the fast event. Inertia sets the rate of change, the composite of droop and load damping sets the arrested offset, and integral action sets the final value, and no one of the three quantities depends on the parameters that set the other two.

The experiments of this chapter cover integrating process dynamics, proportional offset and its dependence on loop gain, the requirement for integral action to remove offset, disturbance rejection

without setpoint motion, actuator saturation against reserve and ramp limits, anti-windup, and the separation of a fast primary loop from a slow secondary loop in the manner of a cascade [1, 2]. The physical fundamentals are conservation of energy applied to the rotating store, in which the swing equation is a per-unit inventory balance, the frequency dependence of aggregate load as a driving force opposing the excursion, and a time-scale separation between a governor lag of 0.20 s and a secondary time constant of 3.5 s [3]. The behavior matters industrially because the rate of change of frequency and the frequency nadir set protection margins and reserve procurement on every interconnected system, and because the displacement of synchronous plant by converter-interfaced generation reduces the one parameter that sets both [4].

> ✦ DID YOU KNOW? **A Lost Generator Almost Every Day**
>
> A 2010 national laboratory report counts how often the sudden loss of a generator is recorded: almost daily in the very large Eastern Interconnection, about once a week in the smaller Texas Interconnection, and the very largest events rarely more than once or twice a year [5].

This chapter works with a single-area load frequency control simulator and states the model, the controller and the measured behavior in enough detail to reproduce them from the text. It separates the three time scales that a frequency event passes through: the rate of change of frequency, the arrested offset and the restored setpoint follow from three expressions with three independent parameter groups, and measured runs show that a change to one group moves one quantity. Section 8.2 describes how frequency regulation acquired its present two-tier form. Section 8.3 writes the three-state model of governor, turbine and rotating store, with its reserve and ramp limits and its protection flags. Section 8.4 derives the rate of change, the offset, the stability limit of the primary loop and the time constant of restoration. Section 8.5 states the secondary controller as the code executes it, and Section 8.6 reports the measured runs of the eleven scenarios. Section 8.7 sets the model beside the sensors, valves and control platforms of real grids. Sections 8.8 and 8.9 state the limits of the model and what the displacement of synchronous plant changes in it, and Section 8.10 sums up. Offset and integral action carry a physical meaning throughout, and a measured value in hertz.

8.2 Background

Early generating stations ran alone or in small groups, and a governor that returned speed exactly to its reference was both possible and desirable. The arrangement fails as soon as two machines are paralleled. Two isochronous governors on one bus each attempt to hold the same speed, the machine with the marginally higher reference takes as much load as its capacity allows, and the division of output between them is indeterminate. Droop was introduced to make that division determinate. A falling speed-output characteristic gives each machine exactly one output at any given frequency, so load divides between machines in inverse proportion to their droop settings, and the consequence is that the common frequency depends on total load.

Interconnection added a second requirement. Once control areas are tied together, a frequency error is shared across the whole interconnection while the schedule violation belongs to one area, and a regulator that responds to frequency alone leaves every area responding to every other area’s imbalance. Tie-line bias control resolves this by forming an area control error that combines the deviation of net tie-line flow from schedule with the frequency deviation weighted by a bias factor, so that an area responds fully to its own imbalance and, in the ideal case, not at all to a neighbor’s [6]. The resulting structure, proportional governing followed by integral restoration of an area control error, remains in use across successive generations of governor, telemetry and control-center hardware.

Analytical treatment of the resulting multi-area problem followed the availability of state-space methods, with optimal control formulations of the megawatt-frequency problem appearing around 1970 and establishing the low-order linear model that is still the standard vehicle for load frequency control studies [7]. Later work condensed the multi-machine response into a single equivalent inertia

and a single reheat time constant to predict frequency nadir directly [8], and successive surveys have catalogued the controller structures applied to the same model [9, 10]. The model has also been carried outward from the transmission system, because droop is now the standard primary control law for parallel inverters in islanded microgrids [11], the setting of Volume 17, Chapter 7, where a droop-governed genset and a battery inverter share an islanded load and an integral loop restores frequency.

Current work in the field concerns the inertia coefficient itself. Converter-interfaced wind and solar generation displaces synchronous plant without contributing rotational energy to the system, so the inertia constant falls while the disturbances it must absorb do not [4, 12]. System operators in small or weakly interconnected systems reached the operational limits of this first and now constrain dispatch to keep inertia above a floor [13], and the stability of a grid dominated by converters is an active research program [14]. That literature concerns the coefficient $2H$ in the model of this chapter, and the analysis of Section 8.4 states which measured quantities that coefficient does and does not set.

> ⧖ HISTORY **Judging the Second Tier, 1962 to 1973**
>
> A committee of North American operators formed in 1962 to recommend an informal operating structure for joining the Eastern and Western Interconnections, which were joined in 1967 and separated a short time later over transmission flow problems. In the 1960s most energy management systems were analog, and a dispatcher relied on the telephone and on strip charts of frequency, tie-line flows and area control error. After the 1965 blackout the industry began to apply digital computers. The first control performance criteria were tested in 1971 and put in place in 1973, and their 10-minute interval was chosen because it was the division most often marked on the strip charts. The attention given to this secondary tier cost the industry much of its memory of why governor response matters [15].

8.3 Process Model

The simulated plant is a single control area reduced to one equivalent machine, one equivalent governor and one equivalent turbine, supplying one aggregate load and receiving one converter-interfaced renewable infeed. That arrangement is drawn in Figure 8.1, which carries every symbol the equations below use. The reduction to a single equivalent machine follows the low-order system frequency response model that is the standard vehicle for load frequency control studies [8, 16].

8.3.1 Per-Unit Conventions

Every power in the model is expressed per unit of a system base S_B equal to the aggregate rating of the synchronous fleet, and every power symbol carrying a Δ is a deviation from the scheduled dispatch. Frequency is normalized differently. The state that the model integrates is the per-unit frequency deviation

$$\Delta f = \frac{f - f_0}{f_0}, \tag{8.1}$$

where f is the system frequency in Hz and $f_0 = 60\,\text{Hz}$ is the nominal value, so Δf is dimensionless and a full hertz of error is 1.667×10^{-2} pu. The inertia constant H is the stored kinetic energy of all synchronized machines at rated speed divided by S_B, which gives it units of seconds and the interpretation of how long the fleet could supply rated output from its rotors alone. The load damping constant D carries units of per-unit power for each per unit of frequency, so $D = 1$ means a 1 % fall in frequency reduces load by 1 % of base. The droop R carries the reciprocal units, per-unit frequency for each per unit of power, so $R = 0.05$ means a 5 % frequency change drives a governor from no load to full load and the primary gain is $1/R = 20$ pu/pu. Rates of change of frequency are reported in Hz/s rather than in per unit, because relay settings are written in those units.

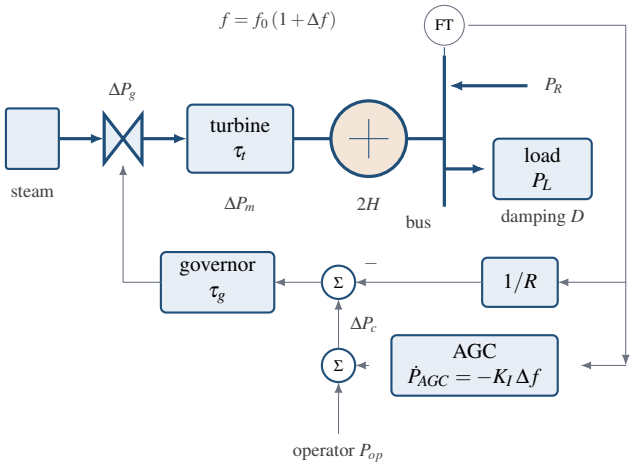


Figure 8.1 Single-area load frequency control arrangement modeled by the application. The governor valve sets ΔP_g, the turbine converts it to mechanical power ΔP_m through the lag τ_t, and any mismatch with the bus load P_L less the renewable infeed P_R is taken from the rotating mass characterized by $2H$ and D. Frequency measured at the bus returns through the droop gain $1/R$ as primary control and through the integrator K_I as secondary control.

8.3.2 Governing Equations

The energy inventory is written as the swing equation with a damping term, as shown in (8.2), the balance of Equation (1.1) applied to the rotating store.

$$2H \frac{\mathrm{d}\Delta f}{\mathrm{d}t} = (\Delta P_m - P_{loss}) - \Delta P_L - D\Delta f \tag{8.2}$$

where ΔP_m is the deviation of turbine mechanical power in pu, P_{loss} is mechanical power removed by a tripped unit in pu, ΔP_L is the deviation of net load in pu and $D\Delta f$ is the self-regulation of the load. Both sides are in per unit of power because $2H$ carries units of seconds and the derivative carries units of reciprocal seconds. Equation (8.2) is conservation of energy applied to the rotating store, with the stored kinetic energy linearized about rated speed, so it reads as an inventory balance in which the level is Δf itself [8]. The only self-correcting term is $D\Delta f$, and it is a property of the connected load rather than of any controller.

The steam path cannot fill instantly, so mechanical power follows valve position through a first-order lag, as shown in (8.3).

$$\tau_t \frac{\mathrm{d}\Delta P_m}{\mathrm{d}t} = \Delta P_g - \Delta P_m \tag{8.3}$$

where ΔP_g is the governor valve output in pu and τ_t is the turbine time constant in seconds.

The governor is the primary control loop and it is written as a lag driven by the difference between the commanded output and the droop feedback, as shown in (8.4).

$$\tau_g \frac{\mathrm{d}\Delta P_g}{\mathrm{d}t} = \Delta P_c - \Delta P_g - \frac{1}{R}\Delta f \tag{8.4}$$

where ΔP_c is the aggregate generation command in pu, τ_g is the governor time constant in seconds and $1/R$ is the primary gain in pu/pu. The term $-(1/R)\Delta f$ is the whole of primary control. It is a pure proportional action on the frequency error, it belongs to the plant model rather than to the controller, and it operates identically whether or not the secondary loop is in service.

Table 8.1 Variable Classification

Symbol	Role	Unit	Range
f	CV, measured	Hz	59.0–61.0 typical
f_{SP}	setpoint	Hz	59.8–60.2
Δf	state	pu	$\pm 1.7 \times 10^{-2}$
ΔP_m	state	pu	± 0.65 reachable
ΔP_g	state	pu	± 0.65 reachable
ΔP_c	MV	pu	-0.30–0.30
P_{op}	MV, manual	% of base	-30–30
P_{AGC}	MV, secondary	pu	signed, unclamped
d_L	DV	%	-10–15
d_R	DV	%	-15–15
d_G	DV	%	0–50
σ_R	DV	%	0–5
RoCoF	derived	Hz/s	± 2 typical
E_{kin}	derived	pu s	signed

Frequency and its rate of change are reported from the state through (8.5).

$$f = f_0\,(1 + \Delta f)\,, \qquad \mathrm{RoCoF} = f_0\,\frac{\mathrm{d}\Delta f}{\mathrm{d}t} \tag{8.5}$$

Net load collects the three disturbance channels and the renewable variability into one signal, as shown in (8.6).

$$\Delta P_L = \frac{d_L}{100} - \frac{d_R}{100} - n(t), \qquad P_{loss} = \frac{d_G}{100} \tag{8.6}$$

where d_L is the commanded load change in percent of base, d_R is the commanded renewable output change in percent of base, d_G is the tripped generation in percent of base and $n(t)$ is band-limited renewable variability in pu. A renewable drop enters net load with a positive sign, because net load is load less infeed, so the swing equation does not distinguish a loss of renewable output from an equal rise in load. A unit trip enters separately as P_{loss} so that mechanical power and renewable infeed remain separately reported quantities.

Renewable variability is an Ornstein-Uhlenbeck sequence generated by the framework's colored-noise primitive, updated once per controller sample with the recursion in (8.7).

$$n_k = a\,n_{k-1} + \sqrt{1 - a^2}\;\sigma_R\,\xi_k, \qquad a = \mathrm{e}^{-\Delta t_c/\tau_R} \tag{8.7}$$

where ξ_k is a unit normal deviate from a seeded mulberry32 generator, σ_R is the requested standard deviation in pu, entered as a percentage of base, and τ_R is the correlation time. The innovation is scaled so that the stationary standard deviation equals σ_R exactly. The value is held constant across the integration step, because the derivative function is evaluated four times per step by the integrator and must be a pure function of its arguments.

8.3.3 Variables and Parameters

Table 8.1 sorts the exposed signals into controlled, manipulated, disturbance, state and derived quantities, and Table 8.2 lists the parameters with their defaults, their adjustable ranges and, where a parameter is fixed by construction, that status.

The eleven parameters above the rule are adjustable while the simulation runs. The six below the rule are fixed in the application. Changing the first five would put the reported axes and the declared actuator limits out of step with the model, and T_t belongs to the secondary controller of Section 8.5 rather than to the plant. The defaults are the customary low-order load frequency control study set [8], and the composite grid stiffness they produce is $1/R + D = 21\,\mathrm{pu/pu}$.

Table 8.2 Model Parameters, Defaults and Adjustable Ranges

Symbol	Meaning	Default	Unit	Range
H	inertia constant	5.0	s	1.0–10.0
D	load damping	1.0	pu/pu	0–3
R	governor droop	0.05	–	0.02–0.20
τ_g	governor time constant	0.20	s	0.05–1.00
τ_t	turbine time constant	0.50	s	0.10–2.00
$\dot{P}_{max}$	generation ramp limit	0.10	pu/s	0.02–0.50
P_{L0}	scheduled load	1.00	pu	0.70–1.20
P_{R0}	renewable infeed	0.20	pu	0–0.60
Δf_n	normal band ($\pm$)	0.10	Hz	0.02–0.30
f_{UFLS}	load-shed relay	59.5	Hz	58.0–59.8
f_{OF}	over-frequency trip	60.5	Hz	60.2–62.0
f_0	nominal frequency	60.0	Hz	fixed
$\Delta P_{c,max}$	reserve up	0.30	pu	fixed
$\Delta P_{c,min}$	reserve down	−0.30	pu	fixed
P_{rated}	turbine rating	1.30	pu	fixed
τ_R	renewable noise correlation	2.0	s	fixed
T_t	anti-windup tracking time	1.0	s	fixed

8.3.4 Initialization and Numerics

Every scenario begins from an exact equilibrium rather than from zero, by solving (8.2) to (8.4) with all derivatives set to zero for whatever command and disturbance the scenario has pinned, as shown in (8.8).

$$\Delta f_0 = \frac{\Delta P_c - \Delta P_{L,\mathrm{tot}}}{1/R + D}, \qquad \Delta P_{m,0} = \Delta P_{g,0} = \Delta P_c - \frac{\Delta f_0}{R} \tag{8.8}$$

where $\Delta P_{L,\mathrm{tot}} = \Delta P_L + P_{loss}$ evaluated with the noise term set to zero. With the default scenario settings this gives $\Delta f_0 = 0$ and a reported frequency of exactly 60.000 Hz, so an undisturbed baseline precedes the event at $t = 2$ s.

The engine of Section 1.6.1 advances the three states with the fixed-step Runge-Kutta method of fourth order and a step of $\Delta t = 0.02$ s, which is a factor of ten below the fastest mode in the model at $\tau_g = 0.20$ s. The secondary controller executes on its own clock at $\Delta t_c = 0.10$ s, first at $t = 0$ and then at every multiple of Δt_c. Data are logged every third integration step, at intervals of 0.06 s. The protection logic of the next paragraph runs after the integration and before the outputs, within the step order of Figure 1.10. The random sequence is seeded at 20260824, so a scenario with variability enabled replays identically.

Two conditions are latched after each step and annunciated. The frequency falling below f_{UFLS} raises an under-frequency load-shedding flag and marks the record, and rising above f_{OF} raises an over-frequency trip flag, each clearing with 0.1 Hz of hysteresis. Neither actuates. A real under-frequency scheme would disconnect a block of feeders and change the model discontinuously, which would replace the transient under study with a step, so the relay is deliberately a flag rather than an event.

8.3.5 Assumptions

The model is a single-area, single-machine equivalent, so inter-machine swings, inter-area oscillations and any spatial variation of frequency are outside it. Real events show frequency differing by tenths of a hertz across an interconnection for the first several seconds, and the model does not reproduce that.

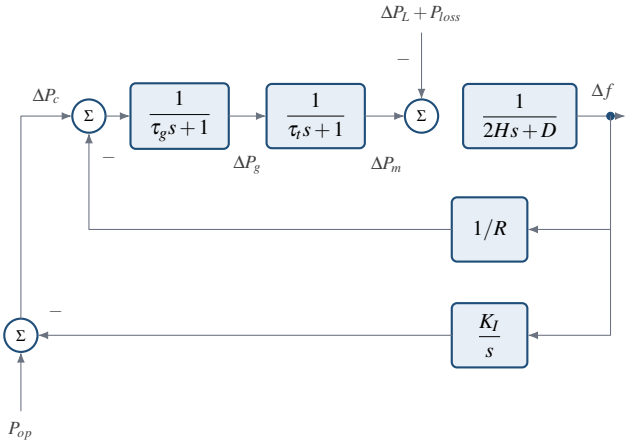


Figure 8.2 Primary and secondary loops around the frequency inventory. The governor droop path $1/R$ closes a proportional loop through two lags inside the plant, while the secondary path K_I/s adds the only integral action in the system and enters at the same summing point as the operator's command P_{op}.

There is no network. Voltages, angles, line impedances and losses are absent, so power injected at the bus arrives at the load unchanged and the model represents neither transmission limits nor voltage stability.

Everything is linearized about the scheduled dispatch. The swing equation uses $2H$ rather than the exact $2H(1+\Delta f)$, and the damping is a single constant. The linearization is accurate for a few percent of imbalance and progressively inaccurate beyond about 10 %, although the disturbance range extends to a 50 % generation loss.

A tripped unit removes mechanical power but not inertia. On a real system the surviving fleet loses that machine's stored energy as well, so a real trip produces a worse rate of change of frequency and a deeper nadir than this model predicts, and the discrepancy grows on a system already short of inertia.

The turbine is a single 0.5 s lag. A reheat steam unit adds a second stage of several seconds, a hydro unit has transient droop and an inverse response from water inertia, treated in Chapter 7, and a gas turbine behaves differently again. Any of these would slow primary response substantially, which is why real primary reserve is specified over tens of seconds rather than the five seconds of this model.

Load damping is one constant. Real frequency-dependent load is a mixture of induction motors, power electronics and thermostatically controlled devices whose aggregate behavior varies with time of day and season.

The generation command is a single aggregate quantity with one reserve limit and one ramp rate. Real reserve is a stack of individual units with individual ramp rates, of which only part is under automatic control at any moment.

Frequency measurement is perfect. There is no transducer noise, no filter lag, no dead band and no telemetry delay, although all four are present in real AGC and are the reason real secondary control runs on a multi-second cycle. The only stochastic element in the application is the renewable disturbance itself.

8.4 Dynamic Analysis

The three quantities that characterize a frequency event, the initial rate of change, the arrested offset and the final value, are set by three disjoint groups of parameters. The block diagram that produces them is shown in Figure 8.2.

8.4.1 Rate of Change of Frequency and Inertia

At the instant a disturbance is applied, ΔP_m has not moved because it is the output of two lags, and Δf is still zero because it is a state. Substituting both into (8.2) and (8.5) gives (8.9).

$$\mathrm{RoCoF}|_{t=0^+} = -\frac{f_0\,(\Delta P_L + P_{loss})}{2H} \tag{8.9}$$

Neither R nor D nor K_I appears. A 5 % load step on the default system therefore produces exactly $-60 \times 0.05/10 = -0.300\,\mathrm{Hz/s}$, and the application measures $-0.299\,\mathrm{Hz/s}$. The residual is the first 0.02 s integration step, because the peak is evaluated after every step and the first evaluation follows the step that applies the load, by which time the depressed frequency has already shed a little load through D. Cutting the inertia to $H = 2\,\mathrm{s}$ multiplies the measured peak by 2.49 to $-0.746\,\mathrm{Hz/s}$ against a predicted factor of 2.5, and doubling it to $H = 10\,\mathrm{s}$ gives $-0.150\,\mathrm{Hz/s}$ exactly.

> **TRY IT IN THE APP** **7 · Low Inertia ($H = 2\,\mathrm{s}$)**
>
> Load the case and press Run. The 5 % load step lands at 2 s. Watch the RoCoF strip. Its peak reaches −0.746 Hz/s, against −0.299 Hz/s in scenario 2, and the nadir falls to 59.658 Hz. Then watch the dial: it still settles at 59.857 Hz, the value of scenario 2.

8.4.2 Steady-State Offset and Composite Stiffness

Setting all three derivatives to zero collapses the model to a single algebraic relation, as shown in (8.10).

$$\Delta f_{ss} = \frac{\Delta P_c - \Delta P_L - P_{loss}}{1/R + D} \xrightarrow{\Delta P_c = 0} -\frac{\Delta P_L + P_{loss}}{1/R + D} \tag{8.10}$$

where the denominator $1/R + D$ is the composite stiffness of the area in pu of power per pu of frequency. Two contributions make up that stiffness and only one of them is a control action. The governors supply $1/R = 20$ and the load supplies $D = 1$ by drawing less power at a depressed frequency. The inertia constant H has dropped out, because a store that is neither filling nor emptying does not enter a steady-state balance whatever its size.

The stiffness is finite, so a finite imbalance forces a finite error. A 5 % load step gives $-0.05/21 = -2.381 \times 10^{-3}$ pu, which is $-0.143\,\mathrm{Hz}$ and a reported frequency of 59.857 Hz. The application settles at 59.857 Hz and its own verification test asserts agreement with (8.10) to 10^{-6} pu. Raising droop from 5 % to 12 % lowers the primary gain to 8.33 and more than doubles the measured offset to 0.321 Hz, while lowering droop to 2 % raises the gain to 50 and reduces the offset to 0.059 Hz. Raising the gain reduces the offset and does not remove it, which is a measured demonstration of the requirement for integral action [1].

Pause. Try Exercise 8.2 before reading on.

8.4.3 Stability of the Primary Loop

Eliminating ΔP_m and ΔP_g from (8.2) to (8.4) with $\Delta P_c = 0$ gives the third-order characteristic polynomial in (8.11), whose roots are the poles of the primary loop in the sense of Section 1.4.3.

$$2H\tau_g\tau_t\,s^3 + [2H(\tau_g + \tau_t) + D\tau_g\tau_t]\,s^2 + [2H + D(\tau_g + \tau_t)]\,s + \left(D + \tfrac{1}{R}\right) = 0 \tag{8.11}$$

The Routh criterion on a cubic requires the product of the two middle coefficients to exceed the product of the outer two, which for the default τ_g and τ_t reduces to (8.12).

$$(1.4H + 0.1D)\,(2H + 0.7D) > 0.2H\left(D + \tfrac{1}{R}\right) \tag{8.12}$$

At the defaults the left side is 75.97 against 21 on the right, which is a wide margin. Holding $D = 1$ and $1/R = 20$ and reducing inertia, the inequality fails below $H = 1.05\,\text{s}$, and the primary loop is then unstable rather than merely oscillatory. The mechanism is the same as high proportional gain on any plant with two lags in series [3], and it is the reason the strong-droop scenario at $R = 2\,\%$ is oscillatory while settling at the smallest offset of the three droop cases.

8.4.4 Restoration by Integral Action

Secondary control acts far more slowly than the governors, so the primary loop can be treated as quasi-static and (8.10) used as its gain. Combining that with the integral law $\dot{P}_{AGC} = -K_I \Delta f$ gives the first-order closed-loop response in (8.13), whose time constant is (8.14).

$$\frac{\mathrm{d}\Delta P_c}{\mathrm{d}t} = -\frac{K_I}{1/R + D}\left(\Delta P_c - \Delta P_L - P_{loss}\right) \tag{8.13}$$

$$\tau_{sec} \approx \frac{1/R + D}{K_I} \tag{8.14}$$

At the default $K_I = 6$ pu/s per pu the secondary time constant is $21/6 = 3.5\,\text{s}$, which is seven times the turbine lag of 0.5 s and 17.5 times the governor lag of 0.20 s. The separation is deliberate and is the same time-scale argument that governs the choice of a cascade inner loop [1], as in the level-to-flow cascade of Volume 3, Chapter 2. Raising K_I above approximately 25 places the secondary loop on the time scale of the lags below it and the two loops then interact. Integral gain barely moves the nadir, because whatever command the integrator has accumulated by the arrest still has to pass through the governor and turbine lags. On the challenge case the nadir is 59.781 Hz at $K_I = 1$ and 59.797 Hz at $K_I = 20$, while at $K_I = 25$ the second swing falls to 59.655 Hz, below the first, so a large enough K_I makes the nadir worse.

The steady state of the complete system follows directly and is what the manual scenario is constructed to demonstrate. Setting $\dot{P}_{AGC} = 0$ requires $\Delta f = 0$, and (8.10) with $\Delta f_{ss} = 0$ then requires $\Delta P_c = \Delta P_L + P_{loss}$. The generation command must equal the imbalance exactly, because at 60.000 Hz the droop term contributes nothing.

🕮 **Pause.** Try Exercise 8.3 before reading on.

8.5 Controller Design

The primary controller is (8.4), and it is a plant parameter rather than an adjustable controller setting. The implemented code is the secondary layer alone. It carries integral action with an optional proportional term and no derivative term, so it is coded separately from the shared PID block of Section 1.6.2, which parameterizes integral action as K_c/τ_I and cannot express a pure integral law at $K_c = 0$. The code adds a secondary contribution to the operator's own setting, which the secondary loop never writes. It otherwise follows the shared block, with the limits on the applied command and anti-windup on the integral term alone. Away from the limits the law it executes is the incremental, or velocity, form of the PID law [17], in which each sample adds a change to the previous command.

🕮 **Pause.** Try Exercise 8.1 before reading on.

8.5.1 Discrete Algorithm

At each controller sample of period $\Delta t_c = 0.10\,\text{s}$ the following is executed in order, with k indexing the sample. The error is formed in per-unit frequency, as shown in (8.15).

$$e_k = \frac{f_{SP} - f_0}{f_0} - \Delta f_k \tag{8.15}$$

A positive e_k means the frequency is below setpoint and generation must rise, so the command falls as the measured frequency rises and the loop is reverse acting in the ISA sense. The two per-sample

increments are given by (8.16).

$$\begin{aligned} \Delta u_k^P &= K_p\,(e_k - e_{k-1}) \\ \Delta u_k^I &= K_I\,e_k\,\Delta t_c \end{aligned} \tag{8.16}$$

with K_p in pu/pu and K_I in pu per second per pu. There is no derivative increment, and none is appropriate, because the measurement enters an inventory whose own dynamics already differentiate the imbalance.

The operator contribution is clipped to the reserve limits, the secondary contribution is added, and the demanded command passes through the reserve clamp and then the ramp limit, as shown in (8.17).

$$\begin{aligned} P_{op,k} &= \mathrm{clamp}\big(\Delta P_{c,k}^{man},\, \Delta P_{c,min},\, \Delta P_{c,max}\big) \\ S_k &= K_p\,e_k + I_{k-1} \\ c_k &= P_{op,k} + S_k \\ v_k &= \mathrm{clamp}(c_k,\, \Delta P_{c,min},\, \Delta P_{c,max}) \\ \delta_k &= \mathrm{clamp}\big(v_k - \Delta P_{c,k-1},\, -\dot{P}_{max}\Delta t_c,\, \dot{P}_{max}\Delta t_c\big) \\ \Delta P_{c,k} &= \Delta P_{c,k-1} + \delta_k \\ I_k &= I_{k-1} + \Delta u_k^I + \frac{\Delta t_c}{T_t}\left(\Delta P_{c,k} - c_k\right) \end{aligned} \tag{8.17}$$

where $\Delta P_{c,k}^{man}$ is the operator input, S_k is the secondary contribution, c_k is the command the algorithm demands, δ_k is the change in command that the ramp rate permits within one controller sample and $T_t = 1.0\,\mathrm{s}$ is the tracking time of the back-calculation. In the symbols of Equations (1.24) to (1.27), $P_{op,k}$ takes the place of the bias u_b, c_k is v_k, $\Delta P_{c,k}$ is u_k, K_I replaces K_c/τ_I, and the reserve clamp followed by the ramp limit is the saturation. The S_k of this chapter is the whole secondary contribution, the running sum of Chapter 1 plus I_{k-1}. The integral used at sample k was accumulated through sample $k-1$. For an unchanged operator input, subtracting the demand at sample $k-1$ from the demand at sample k gives the recursion actually executed, (8.18).

$$c_k = c_{k-1} + \Delta u_k^P + \Delta u_{k-1}^I + \frac{\Delta t_c}{T_t}\left(\Delta P_{c,k-1} - c_{k-1}\right) \tag{8.18}$$

Each demand is the previous demand plus one proportional increment, one integral increment computed from the previous error and a back-calculation correction. Away from the limits $\Delta P_{c,k} = c_k$, the last term vanishes, and the recursion is algebraically identical to the positional law $\Delta P_c = P_{op} + K_p e + K_I \sum e\,\Delta t_c$. At a limit the two differ, as described below.

When the secondary loop is out of service, whether because the application is in manual or because the secondary mode is droop only, the secondary contribution is zero and, as in the manual tracking of Section 1.6.2, the integral state is solved for the value that, with the proportional term of the next sample, reproduces the applied command, as shown in (8.19).

$$S_k = 0, \qquad I_k = \Delta P_{c,k} - P_{op,k} - K_p e_k \tag{8.19}$$

At the first sample after the secondary loop is switched on, the third line of (8.17) then gives $c_{k+1} = \Delta P_{c,k} + K_p(e_{k+1} - e_k)$ for an unchanged operator input, which is the previously applied command plus one proportional increment, and exactly that command at the default $K_p = 0$. The first integral increment reaches the command one sample later. Switching the secondary loop off is bumpless as well. At the first sample without it, the operator input is set once to the command last applied, so the operator takes over the whole command, as units keep their last basepoints when automatic generation control is suspended.

8.5.2 Saturation and Windup

Anti-windup is the back-calculation term alone. While the reserve limit or the ramp limit refuses part of a demand, each sample removes $\Delta t_c / T_t = 0.1$ of the refused part from the integral, so the demand is pulled back toward the command actually applied [17]. With $K_I = 10$ on the unit trip of scenario 3 the command rides the ramp limit for 1.6 s, from 2.72 to 4.32 s, and the frequency overshoots to 60.173 Hz. A counterfactual run with T_t made effectively infinite rides the limit for 1.8 s and overshoots to 60.207 Hz, and at $K_I = 15$ the two peaks are 60.245 and 60.317 Hz. The operator component stays at zero throughout, and the AGC component carries the whole 0.200 pu at the end of the run. A second unit trip on top of the scripted one makes a 40 % loss, which pins the command on its 0.30 pu reserve limit with the frequency held at 59.714 Hz. By $t = 20$ s the integral has settled at 0.329 pu, where the tracking term balances $K_I e$, against 0.729 pu and still rising with T_t effectively infinite. Removing the loss at $t = 20$ s drives the frequency to 61.467 Hz while the command ramps back to zero, against 61.479 Hz without the tracking term, so that peak is set by the ramp limit. The windup shows in the recovery instead: the frequency is back within 0.05 Hz at $t = 29.5$ s, against 33.7 s without the tracking term.

8.5.3 Tuning and Defaults

The integral gain K_I is adjustable over 0 to 40 pu/s per pu in steps of 0.5 with a default of 6. The optional proportional gain K_p is adjustable over 0 to 20 pu/pu with a default of 0. A nonzero K_p acts as additional droop routed through the command path, so it stiffens the arrested frequency in the same way a smaller R would and it still leaves an offset. The default secondary mode is droop only, so the offset is observed before any mechanism that removes it is introduced.

The operator command spans ± 30 % of base in steps of 0.1 %, and the frequency setpoint is adjustable over 59.8 to 60.2 Hz in steps of 0.005 Hz, which corresponds to the time-error correction bias that real interconnections apply [6].

8.5.4 Implementation

The app has the controls and the data export of Section 1.6.4. Manual and automatic modes are provided, with the aggregate generation command as the manipulated variable in manual and the integral law of (8.17) in automatic. The logged record covers the three states, both components of the command, the demanded command, the saturation and rate-limit flags, the achieved command rate, the borrowed kinetic energy, the running nadir and peak rate of change, both protection flags, the plant, dispatch and limit parameters and every disturbance channel, at 0.06 s intervals, which is what a reproduction of any table in this chapter requires.

8.6 Results

Eleven scenarios are supplied with the application and are listed in Table 8.3. Every one starts from an exact equilibrium at 60.000 Hz and applies its disturbance as a timed event at $t = 2$ s, apart from the variability case whose persistent component lands at $t = 10$ s and the handover case which switches secondary control on at $t = 10$ s.

8.6.1 Primary Response

Measured performance for the seven scenarios that test primary control alone is given in Table 8.4, each value recorded at the end of a 30 s run. Overshoot and settling time are omitted because the frequency setpoint never steps in these scenarios, which leaves both quantities undefined, and the application leaves both blank.

Three comparisons in that table establish the separation derived in Section 8.4. Scenarios 2, 6 and 7 differ only in inertia and give peak rates of change of -0.299, -0.150 and -0.746 Hz/s with nadirs of 59.779, 59.833 and 59.658 Hz, and an identical settled frequency of 59.857 Hz in all three. Scenarios 2, 8 and 9 differ only in droop and give settled frequencies of 59.857, 59.679 and

Table 8.3 Scenario Definitions

#	Scenario	Mode / secondary	Change from default
1	`load-step-2`	manual	$d_L = +2\%$
2	`load-step-5`	auto / droop	$d_L = +5\%$
3	`gen-trip`	auto / AGC	$d_G = 20\%$
4	`renewable-drop`	auto / droop	$d_R = -8\%$
5	`renewable-noise`	auto / AGC	$\sigma_R = 1.5\%$, $d_R = -6\%$
6	`high-inertia`	auto / droop	$H = 10\,\text{s}$, $d_L = +5\%$
7	`low-inertia`	auto / droop	$H = 2\,\text{s}$, $d_L = +5\%$
8	`weak-droop`	auto / droop	$R = 12\%$, $d_L = +5\%$
9	`strong-droop`	auto / droop	$R = 2\%$, $d_L = +5\%$
10	`agc-handover`	auto / both	$d_L = +5\%$, AGC at 10 s
11	`challenge`	auto / AGC	$K_I = 1$, $d_L = +5\%$

Table 8.4 Measured Primary Response at $t = 30\,\text{s}$. Settled Frequency Is the Measured Value Compared With (8.10)

#	IAE Hz s	$\lvert e\rvert_{max}$ Hz	nadir Hz	peak RoCoF Hz/s	settled f Hz
1	1.613	0.088	59.912	−0.120	59.943
2	4.032	0.221	59.779	−0.299	59.857
4	6.450	0.353	59.647	−0.479	59.771
6	3.963	0.167	59.833	−0.150	59.857
7	4.072	0.342	59.658	−0.746	59.857
8	8.866	0.356	59.644	−0.299	59.679
9	1.681	0.144	59.856	−0.299	59.941

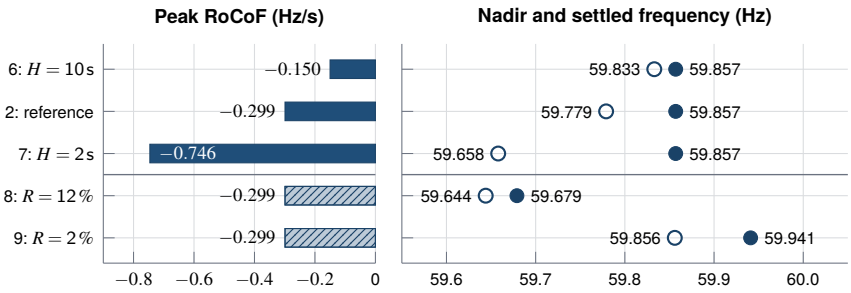


Figure 8.3 The droop-only load steps of Table 8.4: scenarios 6, 2 and 7 change the inertia (solid bars), scenarios 8 and 9 the droop (hatched bars), and each value is printed beside its mark. On the right the open circles are the nadirs and the filled circles the settled frequencies. Notice that inertia moves the peak rate of change and leaves the settled frequency at 59.857, while droop moves the settled frequency and leaves the peak rate at −0.299.

59.941 Hz against the predictions $-0.05/21$, $-0.05/9.33$ and $-0.05/51$ pu, with an identical peak rate of change of $-0.299\,\text{Hz/s}$ in all three. Inertia moves the transient and leaves the steady state unchanged, droop moves the steady state and leaves the transient unchanged, and the manipulated variable travel is 0 pu in every one of these runs because no command is issued. Scenario 4 confirms the disturbance equivalence, because an 8 % renewable loss settles at 59.771 Hz, which is $-0.08/21$ pu below nominal and identical to the result of an 8 % load rise. A load step of the same size added in the same run once the frequency has settled starts at the same −0.479 Hz/s and falls the same 0.353 Hz, from 59.771 to 59.418 Hz. Figure 8.3 sets the five droop-only load steps side by side.

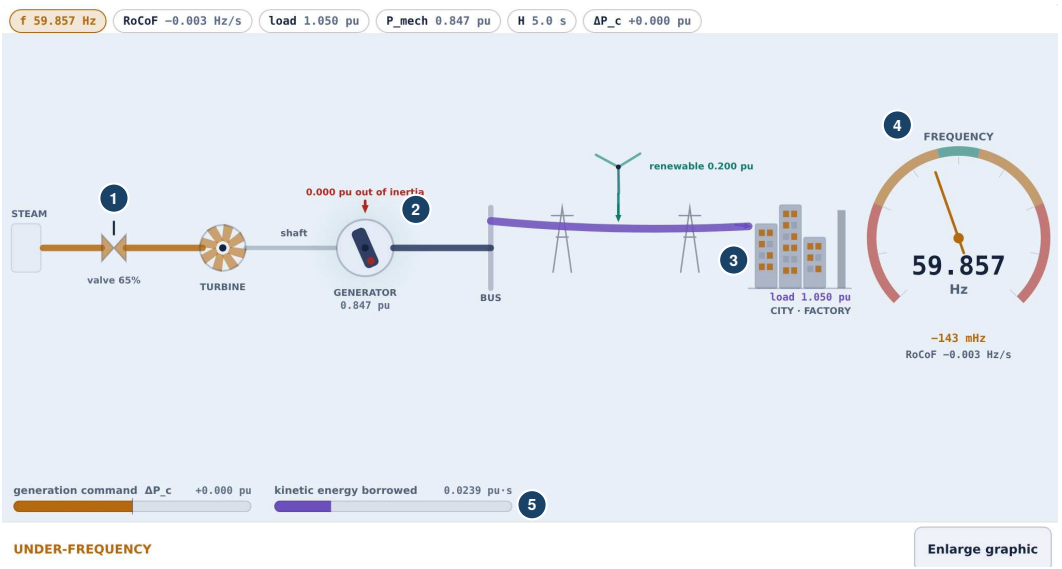


Figure 8.4 Reference scenario of a 5 % load step arrested by governor droop alone, on the Plant view cropped from the application window at $t = 9.7$ s (the clock is in the cut header). The step lands at $t = 2$ s, and the frequency falls at –0.299 Hz/s to a nadir of 59.779 Hz, values that this view does not print (Table 8.4, row 2). The governor valve (1) admits steam to the turbine, the generator (2) gives 0.847 pu and no longer draws on inertia, and the load (3) is the stepped one. The dial (4) reads 59.857 Hz, the 0.143 Hz offset below nominal that (8.10) predicts for an imbalance of 0.05 pu against a stiffness of 21 pu/pu. The borrowed kinetic energy (5) reads 0.0239 pu s. The generation command bar reads zero, inside the ±0.30 pu reserve limits of Table 8.2, and the under-frequency warning at the foot is shown because the frequency is outside the 0.10 Hz normal band.

The manual scenario states the requirement on secondary control as a measured quantity. A 2 % load step arrests at 59.943 Hz, and returning the frequency to 60.000 Hz requires a generation command of 0.020 pu, which is the imbalance and nothing else. A smaller command leaves part of the balance to the governors and leaves a smaller offset, and a larger command overshoots, because at nominal frequency the droop term contributes nothing and the balance closes on the command alone. Commands of 0.015 and 0.025 pu settle at 59.986 and 60.014 Hz.

8.6.2 The Reference Case

The reference scenario is a 5 % load step with droop only, shown in Figure 8.4, captured at $t = 9.7$ s. The frequency response passes through three phases in sequence. From $t = 2$ s the fall is almost straight at the inertial rate, the governors arrest it at $t = 3.22$ s at a nadir of 59.779 Hz, and the turbine lag carries the recovery past the eventual steady state to 59.884 Hz at $t = 4.98$ s, in a decaying oscillation that stays within 5 mHz of 59.857 Hz from $t = 7.3$ s. Mechanical power rises from 0.800 to 0.848 pu against a load that stepped from 1.000 to 1.050 pu, and the remaining 0.0024 pu of the imbalance is closed by the load itself, which draws that much less power at the depressed frequency. The stored energy borrowed from the rotors settles at 0.0238 pu s, which is $2H\,\Delta f$ at the settled offset, and it is held for as long as the offset lasts. The generation command stays at 0.000 pu throughout and the manipulated variable travel is 0 pu, because the secondary loop is out of service in this configuration.

Table 8.5 Measured Secondary Control. Values Are Recorded at the End of Each Run, at 30 s for Scenarios 3 and 10, 60 s for Scenario 5 and 16 s for Scenario 11

#	Secondary	IAE Hz s	MV travel pu	ramp lim. %	final f Hz
3	AGC, $K_I = 6$	2.001	0.200	0.0	60.000
3	AGC, $K_I = 10$	1.855	0.290	5.3	60.000
3	droop only	16.126	0.000	0.0	59.429
5	AGC, $K_I = 6$	3.264	0.326	0.0	60.038
5	AGC, $K_I = 2$	3.422	0.114	0.0	60.046
5	AGC, $K_I = 25$	13.005	4.729	59.8	59.900
5	droop only	9.316	0.000	0.0	59.862
10	AGC from 10 s	1.662	0.050	0.0	60.000
11	AGC, $K_I = 1$	1.505	0.025	0.0	59.927
11	AGC, $K_I = 8$	0.443	0.059	0.0	60.000

8.6.3 Secondary Response

Four scenarios engage automatic generation control, and the measured runs are compared in Table 8.5. Each row uses the scenario as loaded, with any change of mode or gain made before the run starts.

The unit trip of scenario 3 passes through all three tiers in one run. The frequency falls at −1.198 Hz/s to a nadir of 59.144 Hz at $t = 3.14$ s and spends 1.5 s below the 59.5 Hz load-shedding setting, so the flag picks up with secondary control in service. Secondary control then restores 60.000 Hz with the command at 0.200 pu, which is the mechanical power that was lost. The frequency is back within 0.05 Hz 6.4 s after the trip, and the command rate peaks at 0.0855 pu/s, 85.5 % of the ramp limit, without reaching it. With droop only the frequency settles at 59.429 Hz, the $-0.20/21$ pu offset, after a nadir of 59.118 Hz. Raising K_I to 10 puts the command on its ramp limit for 1.6 s, overshoots the frequency to 60.173 Hz and delays the return within 0.05 Hz to 7.5 s after the trip, so a faster integrator buys nothing once the ramp limit binds.

The handover of scenario 10 transfers with no step in the command. The command is still exactly zero after the first execution with secondary control in service at $t = 10.0$ s, and it takes its first increment of 0.001 438 pu at the execution at $t = 10.1$ s, which is one integral increment $K_I e \Delta t_c = 6 \times 2.397 \times 10^{-3} \times 0.1$ and confirms the holding rule of (8.19). The frequency is back within 0.05 Hz 3.1 s after the handover and within 0.02 Hz after 6.1 s, against a secondary time constant of 3.5 s. Halving K_I to 3 lengthens the time to 0.02 Hz to 13.1 s. Doubling it to 12 first reaches 0.02 Hz after 2.2 s, then rings to 60.016 and 59.979 Hz and stays inside only from 5.1 s.

> **TRY IT IN THE APP** **10 · Secondary Control Off → On**
>
> Load the case and press Run. The load step at 2 s settles with its droop offset, and at 10 s the scenario switches secondary control on. Watch the generation command trend at that instant: it does not jump, and it takes its first step of 0.001 438 pu at 10.1 s. The frequency is back within 0.05 Hz 3.1 s after the handover.

The variability case of scenario 5 separates the two parts of a renewable disturbance. Over one minute on the shipped seed, droop alone leaves the persistent 6 % loss as a 0.171 Hz offset and gives nearly three times the IAE of integral action without any command travel. Integral action removes that offset and can only add travel against the zero-mean part, so tripling K_I from 2 to 6 nearly triples the travel and lowers IAE by 5 %, while $K_I = 25$ spends 59.8 % of the run on the ramp limit, swings to 60.47 Hz and ends worse than droop alone. The ranking is not a property of the seed. Over five other seeds the IAE is 8.50 to 9.46 Hz s for droop alone, 2.99 to 3.70 Hz s at $K_I = 2$, 2.44 to 3.40 Hz s at $K_I = 6$ and 10.52 to 12.07 Hz s at $K_I = 25$.

DID YOU KNOW? A Stronger Integral Can Do Worse Than None

From the simulator: under renewable variability (scenario 5) an integral gain of 25 spends 59.8 % of the minute on the ramp limit and ends with an IAE of 13.005 Hz s, worse than the 9.316 Hz s of droop alone, while the default gain of 6 gives 3.264 Hz s (Table 8.5).

Three results for the load steps follow directly from the analysis of Section 8.4. The manipulated variable travel for the unit trip is 0.200 pu and for the handover is 0.050 pu, which are the imbalances themselves, so the integrator settles on the required correction and no travel is wasted. The challenge at $K_I = 1$ gives a secondary time constant of $21/1 = 21$ s and is still 0.073 Hz away at $t = 16$ s, against the $0.143 \exp(-14/21) = 0.073$ Hz that (8.14) predicts, while at the suggested $K_I = 8$ the time constant is 2.6 s and the residual is below the reported resolution. The nadir is 59.781 and 59.788 Hz in those two runs. At $K_I = 8$ the integrator has already commanded 0.018 pu by the nadir at $t = 3.1$ s, but that command reaches mechanical power through the two lags and raises the nadir by only 0.007 Hz. Integral gain sets the settling time and has almost no effect before the arrest.

The challenge asks for the frequency within ± 0.02 Hz of 60 Hz from $t = 10$ s to the end of its 16 s run, a nadir above 59.75 Hz and no time on the 0.10 pu/s ramp limit, with the model parameters and disturbances as loaded, so changing either voids the attempt. With the shipped $K_I = 1$ the worst error after $t = 10$ s is 0.099 Hz, five times the band, although the nadir and the ramp clause pass. With the suggested $K_I = 8$ the worst error is 0.008 Hz, a margin of 0.012 Hz, and the command never reaches the ramp limit. The gain applied before Run and the same gain carried through Reset give one run.

The scorecard leaves overshoot and settling time blank, because the frequency setpoint never steps, and it reports the time on the ramp limit and the time on the reserve limit as separate tiles. No run in Table 8.5 reaches the reserve limit.

8.7 Technology in Practice

The model of Section 8.3 reads one frequency and moves one aggregate generation command through two tiers of control. Figure 8.5 draws both tiers with plant and control-center devices in place of the schematic's symbols.

What Measures It

A real governor reads the speed of its own shaft, while the control center reads the frequency of the bus, and the two readings serve different time scales. The chapter's measurement is perfect, with no filter lag, dead band or telemetry delay, so its peak rate of change, read after the first 0.02 s step, is close to the theoretical value. A relay that averages over several hundred milliseconds may not be triggered by that initial rate if the rate is lower soon after [18]. Until recently the frequency nadir could not be measured reliably in the field and was studied only in simulation [5].

TECHNOLOGY Sensor: Frequency for the Governor and the Control Center

The flyball governors of early plants have largely been replaced by electronic sensors that measure frequency quickly and signal the unit to change its output [19]. Frequency is computed from the voltage over a measuring window. A European guideline asks for a window of 100 to 200 ms and an accuracy of 10 mHz for control at a generating unit, and for 500 ms and 1 mHz for central generation control [20]. Because of that filtering a relay will most likely not see the full initial rate of change, and in the first few cycles after an event different bus-bars read different rates before they converge [18]. Near noisy industrial load the computed rate jumps from sample to sample until it is averaged [20].

What Moves It

The governor of Equation (8.4) and the turbine of Equation (8.3) stand for the valves and steam paths of a whole fleet, with lags of 0.20 and 0.50 s and a ramp limit of 0.10 pu/s on the command, faster than most real units. Typical inertia constants run from 2 to 7 s, lowest in hydro plants and highest

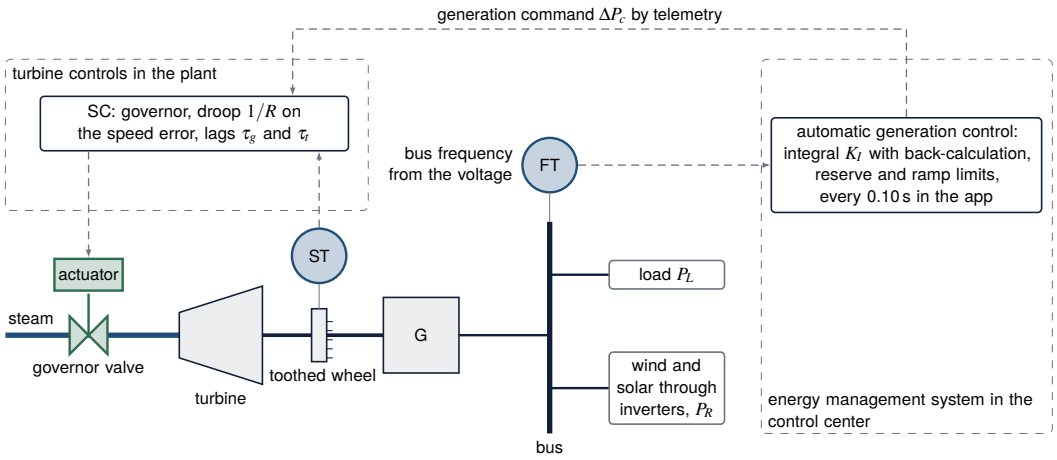


Figure 8.5 The two tiers of the chapter with plant and control-center devices in place of the schematic's symbols. In each plant the governor (SC) reads the shaft speed from a toothed wheel (ST) and moves the governor valve, while the control center reads the bus frequency (FT) and sends the generation command that shifts every governor's setting. The model replaces ST by the bus frequency and the many units by one equivalent machine. Notice that the fast proportional tier never leaves the plant, while the integral tier closes through the frequency of the whole grid.

per unit of capacity in gas plants, and one estimate for the Western Interconnection is about 3.9 s, so the default H of 5 s sits inside that range. A synchronized generator gives the same inertia at any output, but it can give primary response only from head room [19]. In the model the reserve limit of Table 8.2 bounds the generation command, the head room that secondary control can call on.

> **TECHNOLOGY** **Final Element: Turbine Valves That Need Headroom**
>
> When the frequency falls, a governor opens the valves that admit steam or water to its turbine, but the valves and the other mechanical parts take time to move. A typical primary response rises at about 0.3 % of rating per second in slower units and up to 2 % per second in fast units, including certain gas turbines, and in a worked example the output starts to climb a little over half a second after the loss [19]. With a 5 % governor setting a fall of 0.3 Hz raises output by 10 %, but only from a unit running at 90 % of its maximum or less, because a unit at full output has no head room [5].

Where the Controller Runs

Droop sits in each plant's turbine controls. A generating facility is run by a distributed control system that talks to the grid's supervisory system, and its loops run with cycle times from milliseconds to minutes [21]. Automatic generation control runs in the energy management system of the system operations center [15]. Its performance is judged on area control error averaged over one minute for the first control performance standard and over ten minutes for the second [5]. The app runs the integral law with back-calculation of Equation (8.17) every 0.10 s, faster than the multi-second cycle of a real secondary loop.

What Is Modernizing It

Fast frequency response replaces part of the mechanical response with electronics. In Texas, relays on volunteered industrial loads drop them in less than half a second, against the less than 2 % per second of most primary response, and since 2018 new utility-scale wind and solar plants in the United States must be able to respond to frequency. Grid-forming inverters, already used in zero-inertia microgrids typically smaller than 10 MW, can form a grid's voltage and frequency without

synchronous machines [19]. Measurement is changing as well: monitoring that now records the nadir in the field lets planners judge frequency response by the nadir rather than by the settled value [5].

> **IN PRACTICE** **A Governor Free to Respond**
>
> A North American operating policy stated that every generating unit should have an operating governor free to respond to frequency [15]. In service some owners disable their governors or run at maximum output, and others let plant controls restore the scheduled output within seconds, which overrides the governor. On an interconnection's frequency trace this shows as a slope that keeps falling, more slowly, after the frequency first seems arrested. A dead band keeps governors out of normal regulation, so that primary reserve is held for emergencies. The median frequency response of the Eastern and Western Interconnections fell by about a quarter from 2002 to 2008 [5].

8.8 Limitations

The model carries no tie lines, so it does not represent area control error, frequency bias or the inadvertent interchange accounting that occupies most of the practical literature on automatic generation control [6, 9]. The model is the single-area reduction, and the extension to tie-line bias control is not implemented.

A generator trip removes mechanical power but not inertia, which understates the severity of the contingency it is meant to represent. The under-frequency relay annunciates and does not shed, so a nadir below 59.5 Hz is a flag rather than an event and the model never shows the discontinuity that would follow in practice. Frequency measurement is perfect, which removes the filtering, dead band and telemetry delay that force real secondary control onto a multi-second cycle and that dominate its tuning.

The disturbance range extends to a 50 % generation loss, which is well outside the range in which the linear swing equation is accurate. The model integrates that case and reports values, and those values are not predictions of the response of a real system.

> **OPEN CHALLENGE** **How Little Inertia Is Enough**
>
> Inverter-based wind, solar and storage lower a grid's inertia, yet they answer a frequency drop far faster than a turbine valve, which lowers the inertia needed. How much wind and solar this arrangement can carry has not been examined in detail, and the cost, or the need, of a grid that runs with almost no inertia has not been analyzed. How to shape fast response from wind and solar, which need not copy a generator, also needs research [19]. The link between the share of inverters and the rate of change of frequency is not simple, so withstand limits must allow for a plant's whole life [18]. Progress would be studies that find the share at which new approaches are needed [19].

8.9 Discussion

Converter-interfaced generation changes exactly one term in (8.2). Wind and solar plant connected through power electronics supply energy without coupling any rotating mass to system frequency, so displacing synchronous plant reduces H while leaving ΔP_L and P_{loss} untouched [4, 12]. Equation (8.9) then says the rate of change of frequency after a given contingency rises in inverse proportion, and (8.10) says the settled frequency does not move at all. Scenarios 6 and 7 make that pair of statements in one experiment, with the peak rate rising by a factor of 4.98 from -0.150 to $-0.746\,\mathrm{Hz/s}$, against the factor of 5 that (8.9) predicts, and the settled frequency staying at 59.857 Hz to three decimal places.

The consequences are operational rather than theoretical. Relays that disconnect distributed generation on rate of change of frequency are commonly set between 0.5 and 1.0 Hz/s, and a contingency within those settings on a system with high inertia can exceed them on the same system with low inertia, which converts a contained event into a cascading one. Small and weakly interconnected systems reached that limit first, and inertia is now published as an operational quantity

and defended by a dispatch constraint [4, 13]. The stability analysis of systems in which converters supply most of the power is being reformulated, because the assumption under revision is that frequency is a slow, system-wide variable set by rotating mass [14, 16]. Equation (8.12) illustrates the same point on the model of this chapter, because with $1/R = 20$ and the default lags the primary loop itself loses stability below $H = 1.05\,\mathrm{s}$, which is a factor of five below the default value.

8.10 Conclusion

This chapter studied a single-area load frequency control simulator with three states, a governor whose droop feedback is a plant parameter rather than a controller setting, a first-order turbine, reserve and ramp limits on the aggregate generation command, and eleven scenarios. It separated three quantities that a frequency event produces and showed that each belongs to a different group of parameters, with the initial rate of change set by inertia alone at $-f_0(\Delta P_L + P_{loss})/2H$, the arrested offset set by the composite stiffness $1/R + D$ alone at 0.143 Hz for a 5 % step, and restoration to nominal available only from integral action.

Measured runs confirm the separation, with peak rates of change of -0.150, -0.299 and $-0.746\,\mathrm{Hz/s}$ at H of 10, 5 and 2 s against an identical settled frequency of 59.857 Hz, and settled frequencies of 59.941, 59.857 and 59.679 Hz at droop of 2, 5 and 12 % against an identical peak rate of $-0.299\,\mathrm{Hz/s}$. Integral action then restores the unit trip to 60.000 Hz with a command of exactly the 0.200 pu that was lost and without reaching the ramp limit, and the graded challenge is met at $K_I = 8$ with a worst error of 0.008 Hz after $t = 10\,\mathrm{s}$, while integral gains from 1 to 20 move the nadir by less than 0.02 Hz. The app gives a setting for studying proportional offset, integral restoration and two-time-scale control on an integrating process.

KEY POINTS

- Grid frequency is the level of the kinetic energy stored in the spinning machines, so a power imbalance makes it ramp until a control action arrests it.
- Inertia alone sets the initial rate of change, $-f_0(\Delta P_L + P_{loss})/2H$, and the composite stiffness $1/R + D$ alone sets the arrested offset, 0.143 Hz for a 5 % step.
- Droop is proportional action built into governor hardware, so no droop setting removes the offset. Only the integral action of automatic generation control restores 60 Hz.
- Integral gain sets the settling time and has almost no effect before the arrest: gains from 1 to 20 move the nadir by less than 0.02 Hz.

Where else this applies. Chapter 2 leaves the same proportional offset on a tank, Volume 3, Chapter 2 separates a fast inner loop from a slow outer one in the same way, and Volume 17, Chapter 7 restores an island's frequency with droop and an integral loop.

References

[1] D. E. Seborg, T. F. Edgar, D. A. Mellichamp, and F. J. Doyle, III, *Process Dynamics and Control*, 4th ed. Hoboken, NJ: John Wiley & Sons, 2016.

[2] B. W. Bequette, *Process Control: Modeling, Design, and Simulation*. Upper Saddle River, NJ: Prentice Hall, 2003.

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

[4] P. Tielens and D. Van Hertem, "The relevance of inertia in power systems," *Renewable and Sustainable Energy Reviews*, vol. 55, pp. 999–1009, 2016.

[5] J. H. Eto, J. Undrill, P. Mackin, R. Daschmans, B. Williams, B. Haney, R. Hunt, J. Ellis, H. Illian, C. Martinez, M. O'Malley, K. Coughlin, and K. H. LaCommare, "Use of frequency response metrics to assess the planning and operating requirements for reliable integration of variable renewable generation," Lawrence Berkeley National Laboratory, Berkeley, CA, Report LBNL-4142E, Dec. 2010.

[6] N. Cohn, "Some aspects of tie-line bias control on interconnected power systems," *Transactions of the American Institute of Electrical Engineers. Part III: Power Apparatus and Systems*, vol. 75, no. 3, 1956.

[7] C. Fosha and O. Elgerd, "The megawatt-frequency control problem: A new approach via optimal control theory," *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-89, no. 4, pp. 563–577, 1970.

[8] P. M. Anderson and M. Mirheydar, "A low-order system frequency response model," *IEEE Transactions on Power Systems*, vol. 5, no. 3, pp. 720–729, 1990.

[9] Ibraheem, P. Kumar, and D. P. Kothari, "Recent philosophies of automatic generation control strategies in power systems," *IEEE Transactions on Power Systems*, vol. 20, no. 1, pp. 346–357, 2005.

[10] S. K. Pandey, S. R. Mohanty, and N. Kishor, "A literature survey on load–frequency control for conventional and distribution generation power systems," *Renewable and Sustainable Energy Reviews*, vol. 25, pp. 318–334, 2013.

[11] J. W. Simpson-Porco, F. Dörfler, and F. Bullo, "Synchronization and power sharing for droop-controlled inverters in islanded microgrids," *Automatica*, vol. 49, no. 9, pp. 2603–2611, 2013.

[12] A. Ulbig, T. S. Borsche, and G. Andersson, "Impact of low rotational inertia on power system stability and operation," *IFAC Proceedings Volumes*, vol. 47, no. 3, pp. 7290–7297, 2014.

[13] J. O'Sullivan, A. Rogers, D. Flynn, P. Smith, A. Mullane, and M. O'Malley, "Studying the maximum instantaneous non-synchronous generation in an island system—frequency stability challenges in ireland," *IEEE Transactions on Power Systems*, vol. 29, no. 6, pp. 2943–2951, 2014.

[14] F. Milano, F. Dörfler, G. Hug, D. J. Hill, and G. Verbič, "Foundations and challenges of low-inertia systems (invited paper)," in *2018 Power Systems Computation Conference (PSCC)*. IEEE, 2018, pp. 1–25.

[15] H. F. Illian, "Frequency control performance measurement and requirements," Lawrence Berkeley National Laboratory, Berkeley, CA, Report LBNL-4145E, Dec. 2010.

[16] H. Bevrani, *Robust Power System Frequency Control*, ser. Power Electronics and Power Systems. Cham, Switzerland: Springer International Publishing, 2014.

[17] K. J. Åström and T. Hägglund, *Advanced PID Control*. Research Triangle Park, NC: ISA – The Instrumentation, Systems, and Automation Society, 2006.

[18] ENTSO-E, "Rate of change of frequency (RoCoF) withstand capability: ENTSO-E guidance document for national implementation for network codes on grid connection," European Network of Transmission System Operators for Electricity, Brussels, Guidance document, Jan. 2018. [Online]. Available: https://eepublicdownloads.entsoe.eu/clean-documents/Network%20codes%20documents/NC%20RfG/IGD_RoCoF_withstand_capability_final.pdf

[19] P. Denholm, T. Mai, R. W. Kenyon, B. Kroposki, and M. O'Malley, "Inertia and the power grid: A guide without the spin," National Renewable Energy Laboratory, Golden, CO, Technical Report NREL/TP-6A20-73856, May 2020.

[20] ENTSO-E RG-CE System Protection & Dynamics Sub Group, "Frequency measurement requirements and usage," European Network of Transmission System Operators for Electricity, Brussels, Report Final Version 7, 2018. [Online]. Available: https://eepublicdownloads.entsoe.eu/clean-documents/SOC%20documents/Regional_Groups_Continental_Europe/2018/TF_Freq_Meas_v7.pdf

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

8.11 Exercises

8.1 ⌕ Concept check · 2 min

The app holds the grid at 60 Hz with two control actions on the same frequency error, the governor droop and automatic generation control (AGC). Taken together, what kind of controller do they form, and where does each term act?

(a) A PID controller in the control center, in which the governor droop supplies the derivative term

(b) A split-range controller that hands the load from one group of turbines to the other as the error grows

(c) A Smith predictor in the AGC computer that compensates the delay of the telemetry from the governors

(d) A proportional term built into the governor hardware plus an integral term computed by AGC

8.2 Concept check · 1 min

Scenario 6 repeats the 5 % load step of scenario 2 with the inertia constant doubled from 5 to 10 s. Which of the three numbers of the frequency event stays exactly where it was?

(a) The settled frequency, because H drops out of the steady-state balance
(b) The initial rate of change, because the governors act before inertia matters
(c) The nadir, because the turbine lag alone sets how far the frequency falls
(d) None of them, because every number of the event scales with the stored energy

8.3 Concept check · 2 min

In scenario 1 the app is in Manual and a 2 % load step arrests at 59.943 Hz under droop alone. Which generation command returns the frequency to 60.000 Hz and holds it there?

(a) About 0.019 pu, the share the governors carry at the offset, since load damping supplies the rest
(b) Exactly 0.020 pu, because at 60.000 Hz neither droop nor load damping contributes
(c) About 0.021 pu, the imbalance plus the load that damping shed at 59.943 Hz
(d) About 0.040 pu, twice the step, to overcome the governors pulling back as the frequency rises

8.4 Theory and calculation · 6 min

Scenario 4 removes 8 % of base from the renewable infeed at $t = 2\,\mathrm{s}$ with droop only and the default parameters ($H = 5\,\mathrm{s}$, $D = 1$, $R = 0.05$, $f_0 = 60\,\mathrm{Hz}$).

(a) Find the initial rate of change of frequency and the settled frequency. *Checkpoint:* –0.480 Hz/s and 59.771 Hz.

(b) At the settled offset, how much of the 0.08 pu imbalance does the extra mechanical power carry, and how much does the load itself shed? *Checkpoint:* 0.0762 pu and 0.0038 pu.

(c) Find the kinetic energy borrowed from the rotors at the settled offset, taken as $2H\,|\Delta f|$, in pu s. *Checkpoint:* 0.0381 pu s.

8.5 Theory and calculation · 8 min

Use the characteristic polynomial of Equation (8.11) with $\tau_g = 0.20\,\mathrm{s}$, $\tau_t = 0.50\,\mathrm{s}$ and $D = 1$. A cubic $a_3 s^3 + a_2 s^2 + a_1 s + a_0$ with positive coefficients is stable when $a_2 a_1 > a_3 a_0$.

(a) Write the four coefficients in H, D and R for these lags, and evaluate both sides of the condition at $H = 5\,\mathrm{s}$, $R = 0.05$. *Checkpoint:* $a_2 a_1 = 75.97$ against $a_3 a_0 = 21$.

(b) Find the inertia constant below which the primary loop is unstable at $R = 2\,\%$. *Checkpoint:* $H = 3.21\,\mathrm{s}$.

(c) The completion rules of scenarios 7 and 9 ask the reader to lower H to at most 1.04 and 3.15 s. Say why these two values. *Checkpoint:* Each lies just below the threshold of its droop, 1.05 and 3.21 s.

8.6 Theory and calculation · 6 min

The graded challenge applies the 5 % load step at $t = 2\,\mathrm{s}$ with AGC in service from the start, and asks for the frequency within $\pm 0.02\,\mathrm{Hz}$ of 60 Hz from $t = 10\,\mathrm{s}$ on. Treat the restoration as the first-order response of Equation (8.14), decaying from the 0.143 Hz droop offset from $t = 2\,\mathrm{s}$.

(a) Find the secondary time constant at the shipped $K_I = 1$ and at the suggested $K_I = 8$. *Checkpoint:* 21 s and 2.6 s.

(b) Estimate the offset left at $t = 16\,\mathrm{s}$ at $K_I = 1$. *Checkpoint:* 0.073 Hz.

(c) Find the smallest K_I for which the estimate is inside 0.02 Hz at $t = 10\,\mathrm{s}$. *Checkpoint:* About 5.2.

8.7 App challenge · 7 min

Load scenario 2, +5% load step, droop only. The step lands at $t = 2\,\text{s}$ with Secondary control at Droop only. Use the default seed and run each case for 30 s.

(a) Run the scenario as shipped. Read the Peak RoCoF, Nadir and Offset now tiles at $t = 30\,\text{s}$, and compare the first and last with Equations (8.9) and (8.10). *Checkpoint:* −0.299 Hz/s, 59.779 Hz and 0.143 Hz, each within 0.002, against -0.300 and 0.143.

(b) In Show Model, Parameters, set Governor droop to 12 %, press Reset and run again. Read the same three tiles and the IAE tile. *Checkpoint:* −0.299 Hz/s, 59.644 Hz, 0.321 Hz and 8.87 Hz s, within 0.01 on the IAE.

(c) Which of the three numbers did the droop change, and by what factor did the offset grow? Compare that factor with the ratio of the two stiffnesses. *Checkpoint:* Nadir and offset moved, the rate did not. Offset 2.25 times, and $21/9.33 = 2.25$.

8.8 App challenge · 7 min

Load scenario 10, Secondary control off → on. A 5 % load step lands at $t = 2\,\text{s}$ with droop only, and a scripted event switches Secondary control to Droop + AGC at $t = 10\,\text{s}$. Use the default seed, run each case for 30 s and export the CSV file.

(a) Run the scenario as shipped. In column `total_command_pu`, read the command in the rows at 10.08 and 10.14 s, and find the first logged time after which the frequency stays within 0.02 Hz of 60 Hz. *Checkpoint:* 0 pu and 0.001 438 pu, and 16.1 s, within 0.1 s.

(b) Set AGC integral gain to 3, press Reset and run again. Find the same time. *Checkpoint:* 23.1 s, within 0.1 s.

(c) Express both times from the handover, and compare them with the time that Equation (8.14) predicts for the 0.143 Hz offset to fall to 0.02 Hz. *Checkpoint:* 6.1 and 13.1 s against 6.9 and 13.8 s.

8.9 App challenge · 8 min

Load scenario 3, Generator trip (−20 % of output). A unit carrying 0.20 pu trips at $t = 2\,\text{s}$ with Droop + AGC at $K_I = 6$. Use the default seed and run each case for 30 s.

(a) Run the scenario as shipped. Read the Nadir, Ramp limited and Offset now tiles and the MV travel on the scorecard. *Checkpoint:* 59.144 Hz, 0.0 %, 0.000 Hz and 0.200 pu.

(b) Press the reload button, set AGC integral gain to 10 and run again. Read the Ramp limited tile, the MV travel and the last value of column `peak_frequency_Hz` in the exported CSV file. *Checkpoint:* 5.3 %, 0.290 pu and 60.173 Hz, each within one in the last digit.

(c) The frequency is back within 0.05 Hz of 60 Hz 6.4 s after the trip at $K_I = 6$ and 7.5 s after it at $K_I = 10$. Explain why the faster integrator is slower here. *Checkpoint:* The command rides the 0.10 pu/s ramp limit, and the excess demand overshoots.

8.10 Application area · 4 min

Relays that disconnect distributed generation on the rate of change of frequency are set between 0.5 and 1.0 Hz/s. Use Equation (8.9).

(a) What smallest inertia constant keeps the initial rate of change of a 5 % load step within 0.5 Hz/s? What would the 20 % trip of scenario 3 need? *Checkpoint:* 3.0 s, and 12 s.

(b) The model keeps H when a unit trips (Section 8.8). Suppose the tripped unit also held 20 % of the stored energy, so H falls from 5 to 4 s on the same base. Find the initial rate of change, and compare it with the model's. *Checkpoint:* −1.50 Hz/s against −1.20 Hz/s.

8.11 Application area · 3 min

Two generating units of equal rating run in parallel on one bus, one with 4 % droop and one with 5 %.

(a) What share of a load increase does each unit take once the frequency has settled? *Checkpoint:* 55.6 % and 44.4 %.

(b) What would happen if both governors were isochronous instead? *Checkpoint:* The division of load between them is indeterminate.

8.12 Review the agent · 5 min

A planner asks an AI assistant for an integral gain that passes the graded challenge with margin. It replies: "Use $K_I = 30$. The secondary time constant is $(1/R + D)/K_I = 21/30 = 0.7\,\mathrm{s}$, so the 0.143 Hz offset is gone in about 1.4 s, long before $t = 10\,\mathrm{s}$. Integral action cannot act before the arrest, so the nadir stays at about 59.78 Hz, and the gain only shortens restoration." Find the flaw and say how to check it.