INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 17 Control of Energy Conversion and Storage

CHAPTER 4

Torque-to-Pitch Handover and Gain Scheduling in a Wind Turbine Simulator

Pages 80–110

Edited by John D. Hedengren

How to cite this chapter. "Torque-to-Pitch Handover and Gain Scheduling in a Wind Turbine Simulator," in *Control of Energy Conversion and Storage*, John D. Hedengren, Ed., vol. 17 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 4, pp. 80–110.

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

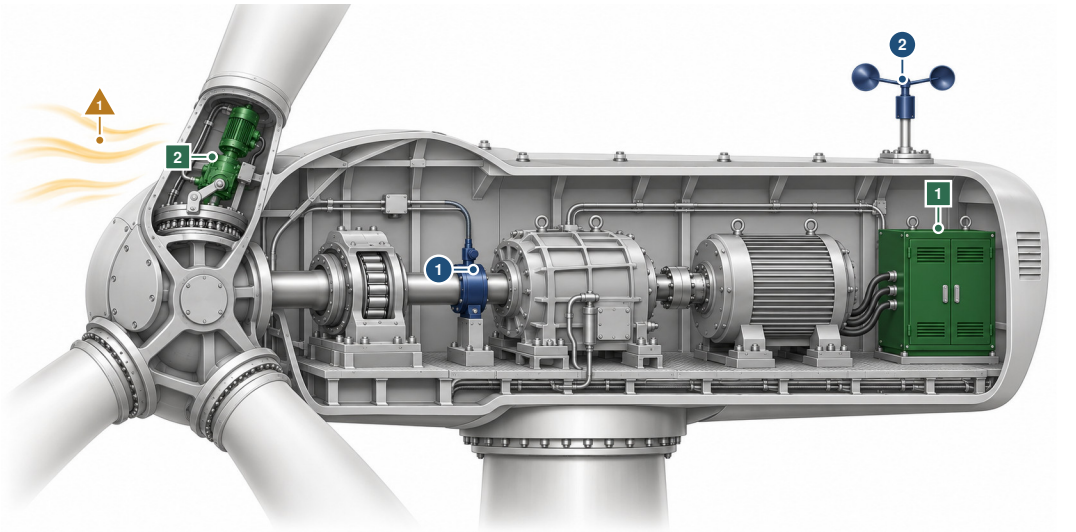
The app for this chapter
https://apmonitor.com/apps/wind-turbine/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v17-energy/#wind-turbine

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

CHAPTER 4

Torque-to-Pitch Handover and Gain Scheduling in a Wind Turbine Simulator



The controller moves 1 converter that sets the generator torque 2 drive that turns a blade, one of three

The controller reads 1 rotor speed 2 wind speed

What disturbs it 1 the wind, with its ramps, gusts and turbulence

Notice that the converter in the nacelle and the pitch drive in the hub act on the same rotor speed: torque holds it in light wind and pitch in strong wind. The chapter hands its controller the true wind at the rotor, where a real anemometer on the roof sits in the rotor's wake. Illustration, not to scale.

IN PLAIN TERMS **One Rotor Speed, Two Ways to Hold It**

The main idea. A large wind turbine must keep its blades turning at a safe speed while the wind keeps changing. In light wind the pull of the generator sets the speed. In strong wind the blades twist to spill the extra wind.

The challenge. The switch between the two must be smooth, or the gearbox takes a jolt each time the wind crosses over. The blades also grow about five times more sensitive as the wind rises, so a fixed setting moves them far more than needed in a gale.

Why it matters. Blade movement wears the bearings, and they must last twenty years.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	The first page, the fact box of the Introduction, the history box of Section 4.2, the two try-it boxes and the fact box of Section 4.6, and the key points
Study	65 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 4.6, Section 4.8, Section 4.9 and its open challenges

The control question A large wind turbine regulates one rotor speed with two actuators, generator torque below rated wind and collective pitch above it, and the pitch sensitivity of its rotor grows by a factor of 4.9 across the range above rated. How should the two actuators hand over, and what do a gain schedule and a predictive controller told the wind ahead buy and cost?

Learning objectives

- Derive the Region 2 speed setpoint and the torque-law constant from the power coefficient at its peak.
- Explain how a setpoint offset, the region index with hysteresis and dwell, and back-calculation hand the rotor from torque to pitch.
- Place the pitch loop by pole assignment and compute its loop gain with and without the gain schedule across Region 3.
- Judge a pitch tuning on blade travel and peak pitch rate as well as on speed error.
- Separate what a wind preview gives a predictive controller from what a tighter speed objective gives, and state its energy cost in turbulence.

Prerequisites Sections 1.6.2 (velocity form and back-calculation), 1.6.4 (scenarios, scorecard and data) and 1.4.2 (inertia as stored energy). Sections 1.6.3 and 1.4.6 are needed only for the predictive controller and its preview.

The app https://apmonitor.com/apps/wind-turbine/

See also Volume 7, Chapter 9 and Volume 7, Chapter 3 for loops scheduled on an operating variable. Volume 4, Chapter 10 for several actuators on one variable, Volume 4, Chapter 2 for a protective loop that rests on its clamp, and Volume 4, Chapter 3 for limiter selection around a scheduled governor.

4.1 Introduction

Variable-speed pitch-regulated wind turbines present a control problem whose difficulty is the objective rather than the plant. There is one measurement worth regulating, rotor speed, and there are two actuators that change it. Generator torque removes energy from the shaft through the converter. Collective blade pitch prevents the wind from putting energy into the shaft in the first place. Both slow the rotor, both are available at every instant, and driving them from the same target produces two integral controllers that settle into any combination satisfying their common setpoint. The supercritical carbon dioxide cycle of Volume 4, Chapter 10 meets the same arrangement with three actuators for one variable.

What separates them is the objective, and the objective changes with the wind. Below the rated wind speed there is less power in the flow than the machine is built to take, so nothing has to be limited and the goal is to capture as much as possible. Capture is maximized at one particular tip-speed ratio, so the rotor speed setpoint tracks the wind rather than staying constant, and generator torque is the actuator that makes the rotor follow it. Above rated wind there is more power available than the drivetrain and the generator can accept, and the goal reverses. The rotor speed setpoint is fixed at the speed rating, generator torque is held at the torque rating because that is what a rating means, and the only remaining way to keep the rotor inside its limits is to feather the blades.

The handover between those two arrangements carries three requirements. It has to be bumpless, because a step in generator torque is a fatigue load on the gearbox. It has to be free of chatter, because in turbulent wind the boundary is crossed many times a minute. The loop that takes over above rated also works against a plant whose gain is not constant. The derivative of aerodynamic torque with respect to pitch grows from -8.9 MN m/rad at fine pitch to -46.6 MN m/rad at 25 m/s in the model used here, so a pitch controller tuned where it first takes over runs at roughly five times its design loop gain at the top of the envelope. Gain scheduling exists to remove that variation, and what it costs when it is switched off is measured in Section 4.6.2.

The physics behind the model is a pair of conservation statements. Aerodynamic power is the kinetic energy flux through the swept area multiplied by a dimensionless power coefficient, and the rotor speed follows an angular momentum balance on a single lumped inertia, so the plant is a first-

order lag whose aerodynamic damping is itself a function of the operating point. The distinguishing dynamic property is a process gain that varies by almost a factor of five over one operating region while its sign never changes, which is the textbook case for gain scheduling [1, 2]. The experiments of this chapter cover split-range and sequenced actuators on one measurement, setpoint tracking against a computed setpoint, feedforward from a steady-state inverse, anti-windup on a loop that rests against a clamp for most of its life, actuator rate limits as a bandwidth ceiling no tuning can raise, hysteresis and dwell on discrete logic, and gain scheduling against a measured nonlinearity [3]. The behavior matters industrially because pitch travel is blade-bearing duty on a component that must last twenty years, and because a turbine that loses speed regulation above rated is a protection event rather than a production loss.

> **+ DID YOU KNOW? A Rotor Wider Than Three Airliners**
>
> By 2019 blades approached 80 m and tip heights often passed 200 m, so three of the largest passenger aircraft, each 80 m across the wings, could fit inside the swept area of one rotor. Wind then supplied about 5 % of the world's electricity [4]. The chapter's reference rotor has a radius of 63 m.

This chapter works with a wind turbine simulator whose model, region logic and measured numbers can be reproduced from the text alone. It separates three mechanisms that are usually taught together: the sequencing of two actuators onto one measurement, the region logic that selects torque operation and labels which actuator is in charge, and the gain schedule that keeps the second loop where it was designed. Their roles differ. A 0.3 % setpoint offset lets the continuously active pitch loop begin regulating before a region declaration, the region index selects rated torque once Region 3 is declared, and switching the gain schedule off buys speed regulation at more than twice the blade travel. Section 4.2 describes how turbine control arrived at this two-region form. Section 4.3 gives the five-state model, with a fitted power coefficient surface whose departures from the textbook correlation are stated and justified, two rate-limited actuators and the wind. Section 4.4 measures the gain nonlinearity of the pitch loop and the schedule that removes it. Section 4.5 describes the two loops, the region logic with its hysteresis and dwell timers, and two settings of a model predictive controller: the controller of Section 4.5.6, which separates what a wind preview contributes from what a tighter speed objective contributes and costs, and the scenario tree of Section 4.5.7, which plans on three winds inside the preview's declared error and measures what hedging against that error is worth. Section 4.6 reports the thirteen scenarios, Section 4.7 discusses the wind measurement the Region 2 setpoint relies on, Section 4.8 describes the sensors, actuators and controllers of real turbines, and Sections 4.9 and 4.10 state the limitations and the conclusions.

4.2 Background

The first generation of commercial machines avoided the problem described above with a fixed rotor speed and blades that stall, as the history box in this section tells, and paid for it with a tip-speed ratio that is right at one wind speed only and a loose power cap.

Variable speed and blade pitch were introduced to fix each of those in turn. Allowing the rotor to change speed lets the tip-speed ratio be held at its optimum through the whole of the low-wind range, which is worth several percent of annual energy on a typical site. Pitching the blades replaces stall with a commanded, repeatable limit on captured power. The two together are the configuration of essentially every megawatt-class machine built since the 1990s, and they are also what creates the two-actuator problem, because the converter and the pitch system are both authoritative over rotor speed. The objectives, the actuator assignments and the region boundaries were laid out early in a form that is still recognizable [5], and the pitch-regulated variable-speed configuration and its power limiting were described in detail shortly after [6].

> ⧖ HISTORY **From Stalling Blades to Pitch and Variable Speed (1970s to 2023)**
>
> The oil crisis of the 1970s rekindled interest in renewable energy and led to the commercial adoption of grid-connected wind turbines [4]. Most early turbines ran at constant speed. On a fixed-frequency grid an induction generator turns only a few percent above synchronous speed while the wind varies across a wide range [7]. In a fixed-pitch, stall-controlled machine the blades are shaped so that stall limits the power in high wind, no pitch mechanism is needed, and control means little more than starting and stopping. A fixed speed puts the power coefficient at its peak at a single wind speed only [8]. By 1999 more and more variable-speed turbines were going into wind farms [7]. By 2023 the industry had nearly 40 years of field experience, 20 of them with multi-megawatt turbines that had converged on three blades, pitch control and variable speed [9].

Two lines of work then developed in parallel. Below rated, the question was how to hold the optimum without measuring the wind. The law that became standard is a static one in which generator torque is made proportional to the square of the measured rotor speed, with the constant chosen so that the rotor is in equilibrium exactly at the optimal tip-speed ratio, and the practical questions became how much the constant is in error and how slowly the rotor converges under it [10]. Above rated, the question was how to place a pitch loop whose plant gain moves with operating point. The standard answer, from the design of commercial controllers, is a proportional-integral loop on rotor speed with both gains divided by a function of the pitch angle [11], later extended to use the same actuator for tower and drivetrain load reduction [12]. Surveys of the field [13] and the open reference designs published for research use [8, 14] both present the machine as a set of regions with different actuators, and the transition between them as a named piece of logic rather than an afterthought.

The wider setting is that wind now supplies a large enough share of generation that turbine controllers are grid assets as well as machine protection, and the open problems in the field are increasingly about loads, atmospheric inflow and system integration rather than about the loops themselves [4]. The loops themselves therefore remain the entry point to the subject, because a pitch controller that trades speed regulation against blade-bearing duty is the smallest complete example of the trade every later turbine question turns on.

4.3 Turbine Model

The application simulates a three-bladed upwind horizontal-axis turbine with the geometry and ratings of the NREL 5 MW offshore reference turbine [14], which is the standard open baseline for control studies. Figure 4.1 shows the physical arrangement and names every symbol that appears in the equations. Everything in the model lives on the low-speed shaft, so the 97:1 gearbox appears only in the conversion of the published rated generator torque of 43 093.55 N m at the high-speed shaft into the 4.180 074 MN m used here.

4.3.1 Variable Classification

The single controlled variable and the two manipulated variables are listed in Table 4.1, together with the disturbances, the five states and the principal derived outputs. Only one loop is scored, because there is only one process value being regulated.

4.3.2 Rotor Aerodynamics

Aerodynamic power is the product of the kinetic energy flux through the swept area and a dimensionless power coefficient that depends on the tip-speed ratio and the collective pitch angle, as shown in Equation (4.1).

$$P_{aero} = \tfrac{1}{2}\rho A V^3 C_P(\lambda, \beta), \qquad \lambda = \frac{\omega R}{\max(V, 0.3)} \tag{4.1}$$

where ρ is air density, $A = \pi R^2$ is the swept area, V is the wind speed reaching the rotor and λ is the ratio of blade tip speed to wind speed. Other chapters of this volume use λ for other quantities, as

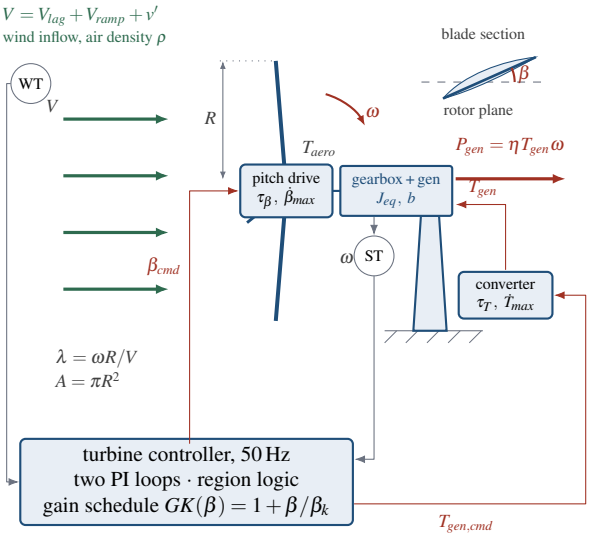


Figure 4.1 Rotor, drivetrain, actuators and instruments of the simulated machine. Rotor speed ω is measured exactly, and the same controller is given the true instantaneous wind speed V that reaches the rotor.

Table 4.1 Variables, Their Roles and Their Ranges

Symbol	Role	Range	Unit
ω	CV, state	0–12.1 nominal	rad/s (rpm)
ω_{SP}	setpoint, computed	0.45–1.2633	rad/s
T_{gen}	MV 1, state	0–4.180074	MN·m
β	MV 2, state	0–90 (0–30 in AUTO)	deg
λ^*	operator setting	4–10	—
V_{lag}	state, wind lag	3–28	m/s
V_{ramp}	state, wind ramp integral	unbounded	m/s
$\bar{V}$	DV, mean wind setting	3–28	m/s
$\dot{V}$	DV, wind ramp rate	−0.4–+0.4	m/s per s
I_{turb}	DV, turbulence intensity	0–20	%
V	inflow seen by rotor	0.2–45	m/s
λ	output	$\omega R / \max(V, 0.3)$	—
C_P	output	0–0.482	—
T_{aero}	output	—	N·m
P_{gen}	output	0–5.0	MW
region	output, annunciation	2, 2.5, 3	—

Section 1.3 lists. The floor of 0.3 m/s on the denominator prevents a division by zero in a calm and has no effect anywhere the machine operates. The upper bound on C_P is the Betz limit of $16/27 \approx 0.593$, which follows from a momentum balance on the stream tube and holds for any rotor whatever [15]. The surface used here peaks at 0.482, which is 81 % of that ceiling.

Aerodynamic torque follows from power and speed, as shown in Equation (4.2).

$$T_{aero} = \frac{P_{aero}}{\max(\omega,\, \omega_\varepsilon)}, \qquad \omega_\varepsilon = 0.05\,\text{rad/s} \tag{4.2}$$

The speed floor ω_ε keeps the quotient finite at standstill. Dimensionally the expression is W divided by rad/s, which is N m as required.

The power coefficient is a fitted analytic surface of the family used for power system studies [16], given in Equation (4.3), with the pitch angle in degrees inside this expression only.

$$\lambda_t = s\lambda, \qquad \frac{1}{\lambda_i} = \frac{1}{\lambda_t} - c_8$$
$$C_P = k\left[c_1\left(\frac{c_2}{\lambda_i} - c_3(\beta + q\beta^2) - c_4\right)\mathrm{e}^{-c_5/\lambda_i} + c_6\lambda_t\right] \tag{4.3}$$

where $c_1 = 0.5176$, $c_2 = 116$, $c_3 = 0.40$, $c_4 = 6.5$, $c_5 = 12.25$, $c_6 = 0.0068$, $c_8 = 0.005$, $q = 0.08$, $s = 0.944705354917025$ and $k = 0.504333798658105$. Three departures from the textbook coefficient set are deliberate and each fixes a defect that would otherwise corrupt the lesson. The reciprocal $1/\lambda_i$ normally carries an additional $-0.035/(\beta^3+1)$ term and a $+0.08\beta$ term, which together make C_P non-monotonic in β at low tip-speed ratio, so that feathering a degree or two raises the modeled power and inverts the sign of the pitch loop. Both are dropped. The pitch loss is written $-c_3(\beta+q\beta^2)$ rather than $-c_3\beta$, because a linear term makes $|\partial C_P/\partial\beta|$ independent of pitch and would remove the nonlinearity the application exists to teach. The two scalings s and k place the optimum exactly at $\lambda^* = 7.55$ with $C_{P,max} = 0.482$, because the unscaled expression peaks at $\lambda = 7.13252543$ with a value of 0.95571624. The remaining coefficients were fitted so that the resulting steady pitch schedule reproduces the published fine-pitch schedule for this machine to 0.6° rms between 12 and 25 m/s and so that C_P falls to zero near $\lambda = 18.5$.

One term is added below the fitted surface. A stopped rotor in a gale is not torque-free, but every fit of this family goes to zero at low λ, which would leave a stopped rotor permanently stopped. A starting-torque floor $C_{P,start} = c_{q0}\lambda\,\mathrm{e}^{-(\lambda/\lambda_s)^2}$ with $c_{q0} = 0.015$ and $\lambda_s = 1.5$ is applied as a lower bound on C_P. At $\lambda = 3$ the floor is 0.0005 against a fitted 0.13, and at the optimum it is below 10^{-9}.

The resulting surface is shown in Figure 4.2. The two features that matter for control are both visible. There is a single maximum in λ at fixed β, which is what makes an optimal tip-speed ratio exist and what makes the Region 2 equilibrium stable. And the curves collapse toward zero as β grows, faster at each successive five degrees, which is the quadratic pitch term and the reason the pitch loop gain has to be scheduled.

4.3.3 Drivetrain

Blades, hub, shaft, gearbox and generator rotor are lumped into a single inertia referred to the low-speed shaft, giving the one rotor state shown in Equation (4.4).

$$J_{eq}\,\dot{\omega} = T_{aero}(\omega, V, \beta) - T_{gen} - b\,\omega \tag{4.4}$$

where $J_{eq} = 4.0 \times 10^7\,\mathrm{kg\,m^2}$ and $b = 1.0 \times 10^4\,\mathrm{N\,m\,s}$ is a viscous drag coefficient standing in for bearing and gear losses. Electrical power is the shaft power taken by the generator times a constant efficiency, $P_{gen} = \eta\,T_{gen}\,\omega$ with $\eta = 0.944$.

The open-loop rotor is stable, and its time constant is a strong function of where the machine is operating. Linearizing Equation (4.4) about a steady point gives $\tau_{OL} = J_{eq}/(|\partial T_{aero}/\partial\omega| + b)$, and $\partial T_{aero}/\partial\omega$ at the optimum is negative because past λ^* the power coefficient falls faster than the speed rises. Evaluated on the fitted surface, the aerodynamic damping is -1.54 MN m s at 6 m/s and -2.65 MN m s at rated speed and rated wind, so the open-loop time constant falls from 25.8 s to 15.0 s across Region 2. The plant is slow on the scale of a control loop, which is why the application runs with a 60 s default trend window while the magnetic suspension of Volume 7, Chapter 6 needs milliseconds. The energy this inertia stores at rated speed, in seconds of rated power, is the inertia constant of Section 1.4.2.

4.3.4 Actuators

Both actuators are first-order lags with a hard rate limit and end stops, as shown in Equations (4.5) and (4.6).

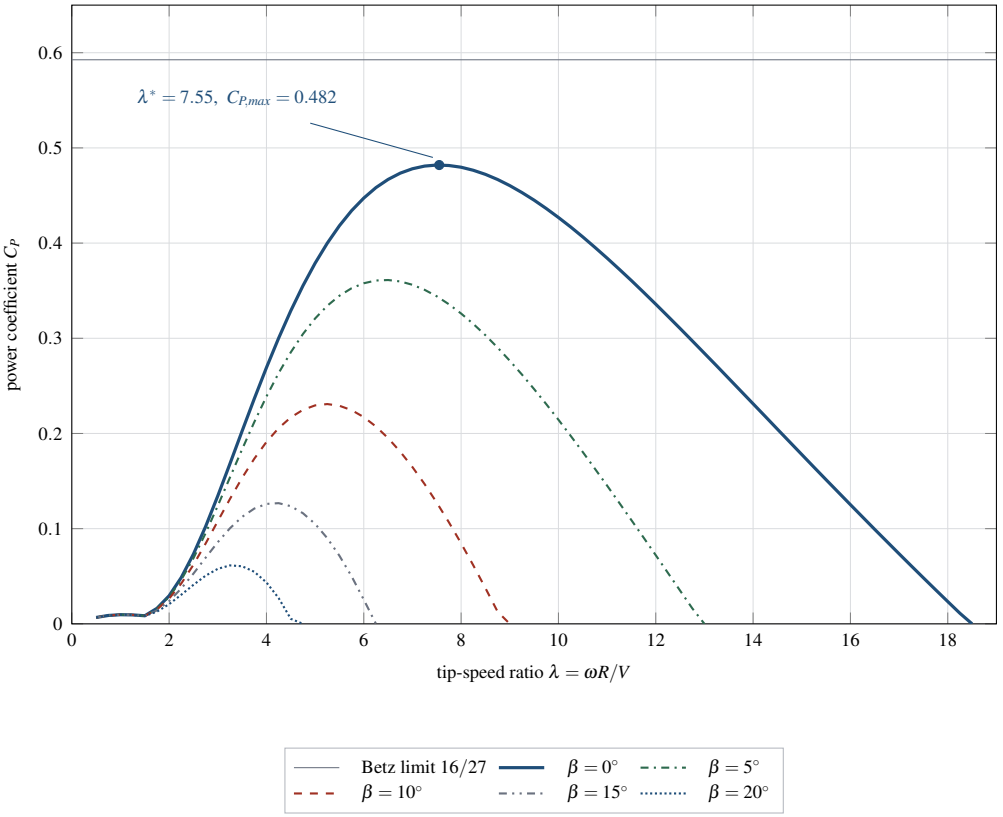


Figure 4.2 Power coefficient surface as implemented, at five collective pitch angles: 0° very thick solid, 5° dash-dotted, 10° dashed, 15° dash-dot-dotted and 20° densely dotted, under the thin horizontal Betz limit. The curves are evaluated from Equation (4.3) with the fitted coefficients, and the marker is the operating point the Region 2 controller holds.

$$\dot{\beta} = \mathrm{clamp}\left(\frac{\beta_{cmd} - \beta}{\tau_\beta},\ -\dot{\beta}_{max},\ +\dot{\beta}_{max}\right) \tag{4.5}$$

$$\dot{T}_{gen} = \mathrm{clamp}\left(\frac{T_{gen,cmd} - T_{gen}}{\tau_T},\ -\dot{T}_{max},\ +\dot{T}_{max}\right) \tag{4.6}$$

with $\tau_\beta = 0.25\,\mathrm{s}$, $\dot{\beta}_{max} = 8\,^\circ/\mathrm{s}$, $\tau_T = 0.05\,\mathrm{s}$ and $\dot{T}_{max} = 1.5\,\mathrm{MN\,m/s}$. The derivative is additionally set to zero when the state is against an end stop and the demand pushes further into it, so the right-hand side stays a pure function of the state and the fourth-order Runge-Kutta integrator sees no discontinuous jump inside a step. Pitch is bounded to $[0, 90^\circ]$ physically and generator torque to $[0, T_{rated}]$, and both bounds are re-imposed after each step.

The asymmetry between the two actuators is what makes several of the scenarios behave the way they do. Generator torque reaches a new value in about a tenth of a second, while pitch needs nearly three seconds to travel the 22° that 25 m/s requires, and pitch is the actuator that has to reject the fastest disturbances.

4.3.5 Wind

The wind reaching the rotor is a single scalar with three parts and two states, as shown in Equation (4.7).

Table 4.2 Machine Parameters and Where the Application Shows Them

Symbol	Meaning	Value	Unit
Rotor			
R	rotor radius	63	m
A	swept area, πR^2	12 469	m^2
ρ	air density	1.225	kg/m^3
J_{eq}	drivetrain inertia (LSS)	4.0×10^7	$kg{\cdot}m^2$
b	drivetrain damping[a]	1.0×10^4	N·m·s
λ^*	optimal tip-speed ratio[h]	7.55	—
$C_{P,max}$	peak power coefficient[h]	0.482	—
Ratings			
ω_{rated}	rated rotor speed	1.267109 (12.10)	rad/s (rpm)
ω_{min}	minimum rotor speed[a]	0.45 (4.30)	rad/s (rpm)
$0.997\,\omega_{rated}$	torque-loop setpoint cap	1.263308 (12.06)	rad/s (rpm)
T_{rated}	rated generator torque	4.180074	MN·m
η	generator efficiency[a]	0.944	—
V_{rated}	rated wind speed[h]	11.4	m/s
K	torque-law constant	2.1388×10^6	$N{\cdot}m{\cdot}s^2$
Actuators			
τ_β	pitch actuator lag	0.25	s
$\dot{\beta}_{max}$	pitch rate limit	8	deg/s
β_{max}	feather limit[a]	90	deg
β_{ctl}	pitch limit in AUTO[a]	30	deg
τ_T	torque actuator lag[a]	0.05	s
$\dot{T}_{max}$	torque rate limit[a]	1.5	MN·m/s
Limits and wind			
ω_{trip}	overspeed alarm[a]	1.20	× rated
τ_V	mean-wind setting lag[h]	2.0	s
ω_ε	torque speed floor	0.05	rad/s

[a] adjustable only in advanced mode. [h] not exposed for editing.

$$\dot{V}_{lag} = \frac{\bar{V} - V_{lag}}{\tau_V}, \quad \dot{V}_{ramp} = \dot{V}, \quad V = V_{lag} + V_{ramp} + v' \tag{4.7}$$

where $\bar{V}$ is the operator-set mean wind, $\tau_V = 2.0\,\text{s}$ so that a change in that setting is not a physical step in the wind, $\dot{V}$ is the operator-set ramp rate integrated independently, and v' is the turbulent component. The ramp integrator is frozen when the sum $V_{lag} + V_{ramp}$ leaves $[0.5, 40]$ m/s. Turbulence is a first-order filtered Gaussian process of time constant 6 s, scaled so that its stationary standard deviation equals the turbulence intensity times $\max(V_{lag} + V_{ramp}, 1)$, and it is generated once per controller step rather than inside the derivative function, because the derivative function is called four times per integration step and must be free of side effects. The random stream is a seeded mulberry32 generator, so a scenario with turbulence replays identically.

4.3.6 Parameters, Initialization and Numerics

Table 4.2 lists every parameter, its value and whether the application exposes it. Hidden parameters are used by the model but are not exposed for editing, which is what makes the low-wind scenarios a task for the reader rather than a demonstration, because a reader who does not already know λ^* has to find it from the reported C_P.

Every scenario begins on its own steady state. The initial speed and the initial generator torque are

$$\omega_0 = \text{clamp}\big(\lambda^*_{model} V_0 / R,\ \omega_{min},\ \omega_{rated}\big),$$
$$T_{gen,0} = \text{clamp}\big(T_{aero}(\omega_0, V_0, \beta_0) - b\,\omega_0,\ 0,\ T_{rated}\big),$$

and the initial pitch β_0 is the angle that makes the aerodynamic torque equal $T_{rated} + b\,\omega_0$, found by 60 bisection steps on the monotone C_P surface. The initialization uses the model's optimal ratio, not a subsequently edited controller setpoint. Before control commands are applied it gives $\dot{\omega} = 0$ at $t = 0$ in every region, and a verification test in the application confirms it to 10^{-12} rad/s^2 at 8 m/s.

Integration uses fourth-order Runge-Kutta at a fixed step of $dt = 0.01$ s. The controller executes every 0.02 s, which is the 50 Hz cadence of a real turbine controller, and the historian logs at 10 Hz. The three clocks, the order of work within a step and the independence of the results from the playback speed are those of Section 1.6.1. Rotor speed is clipped to non-negative values after each step and to eight times rated as a numerical guard, and the overspeed annunciator latches at $1.20\,\omega_{rated}$ and clears below $1.15\,\omega_{rated}$. The alarm is annunciation only, because the manual scenario exists to show where an unprotected rotor comes to rest.

4.3.7 Assumptions

The model is reduced deliberately, and each reduction removes something that matters somewhere.

One power coefficient surface stands for the entire rotor. There is no blade element momentum theory, no dynamic inflow, no stall hysteresis and no unsteady aerodynamics, so anything that depends on how fast the flow field itself rearranges is outside what the model can say. The first second after a large pitch step and deep stall at very low λ are both in that category.

The surface is fitted, not derived. It reproduces the reference machine over λ from 3 to 12 and β from 0 to about 25°, which covers Regions 2 and 3, and is an extrapolation outside that box. At rated wind and rated speed with $\beta = 0$ the surface gives 5.410 MW of aerodynamic power against 5.297 MW of rated mechanical power, so the modeled machine reaches its rating 2.1 % early. That residual is the price of one smooth analytic expression fitted to a real rotor.

One rotor-speed state. The drivetrain torsional mode near 1.7 Hz on the real machine is absent, which matters because it is exactly the mode an over-bandwidth pitch loop excites in practice [17].

No tower. Fast pitch action drives the tower fore-aft mode on a real turbine, and that coupling is the main reason commercial pitch loops are scheduled conservatively [11]. It is not represented, so the cost of an unscheduled loop is measured here as blade travel. A margin against tower motion cannot be quantified without the omitted dynamics.

The wind is one scalar at the hub. There is no shear, no tower shadow, no rotational sampling and no yaw error, and the turbulence is a band-limited first-order process rather than a Kaimal spectrum.

The electrical side is a constant. Generator power is shaft power times 0.944, so losses do not vary with load and there are no converter or grid dynamics of any kind.

Air density and inertia are constant, and the pitch limit of 30° in AUTO is a controller setting rather than a physical stop. It exists to keep the integrator out of the region where the fitted C_P has already reached zero and the pitch loop would have no gain at all. In MANUAL the blades still reach 90°.

4.4 Pitch Sensitivity and the Gain Schedule

The reason this machine needs a gain schedule is visible in one derivative. Table 4.3 lists the steady pitch angle that holds rated speed at rated torque, and the pitch sensitivity of aerodynamic torque and of aerodynamic power at that point, computed from Equations (4.1) to (4.3) at $\omega = \omega_{rated}$.

The sensitivity grows by a factor of 4.9 between the bottom and the top of Region 3. A proportional-integral loop whose gain $K_{c,\beta}$ is fixed therefore has a loop gain product $K_{c,\beta}\,|\partial T_{aero}/\partial\beta|$ that grows with it, from 33.6 where the tuning was placed, at the fine-pitch sensitivity of 8.92 MN m/rad, to 175.6 at 25 m/s with the default tuning, a factor of 5.2. The standard correction [11] divides both the proportional and the integral gain by a linear function of the pitch angle itself, as shown in Equation (4.8).

$$GK(\beta) = 1 + \frac{\beta}{\beta_k}, \qquad K_{c,\beta}^{\mathrm{eff}} = \frac{K_{c,\beta}}{GK(\beta)}, \qquad \tau_{I,\beta}^{\mathrm{eff}} = \tau_{I,\beta} \tag{4.8}$$

Table 4.3 Steady Pitch and Pitch Sensitivity Across Region 3

V (m/s)	β (deg)	$\partial T_{aero}/\partial\beta$ (MN·m/rad)	$\partial P_{aero}/\partial\beta$ (MW/rad)
11.4	0.48	−9.59	−12.16
15.0	10.95	−31.13	−39.45
18.0	15.27	−41.36	−52.41
21.0	18.51	−46.42	−58.81
25.0	21.97	−46.61	−59.05

where $\beta_k = 4.98°$ was fitted to this model's own sensitivity curve, close to the 6.30° published for the real 5 MW machine. Dividing K_c and K_c/τ_I by the same factor leaves the effective integral time untouched, so the schedule moves the bandwidth of the loop and not the shape of its response. Both effective values are logged as separate columns, and a 120 s run of scenario 7 ends with $GK = 5.266$, $K_{c,\beta}^{\text{eff}} = 0.7155\,\text{s}$ and $\tau_{I,\beta}^{\text{eff}} = 2.3333\,\text{s}$ against a set value of 2.3333 s, which confirms the claim from the data rather than from the code. The scheduled loop gain product at 25 m/s is 32.4 against 33.6 at fine pitch, so the schedule holds the loop within 4 % of where it was placed.

Scheduling a gain against a measured operating point is a well-studied device with well-known failure modes, chiefly that the schedule variable must move slowly compared with the loop it is scheduling [18]. Pitch angle satisfies that here because the blades move at most 8 °/s against a loop placed at $\omega_n = 0.6\,\text{rad/s}$. The riderless bicycle of Volume 7, Chapter 9 schedules its loop on forward speed and the turbofan governor of Volume 4, Chapter 3 on corrected core speed, two other schedule variables chosen because they move more slowly than the loop.

4.5 Controller Design

4.5.1 Velocity Form PID Algorithm

Both loops use the shared velocity-form proportional-integral controller of Section 1.6.2 on an exact 0.02 s cycle [19]. For the scheduled pitch loop the gain moves during a run, so the increment must difference the fully formed proportional terms, as in Equation (4.9).

$$\begin{aligned} e_k &= a(\text{SP}_k - \text{PV}_k), \qquad a = -1 \\ P_k &= K_{c,k} e_k, \qquad \Delta P_k = P_k - P_{k-1} \\ \Delta I_k &= (K_{c,k}/\tau_I) e_k \Delta t \end{aligned} \tag{4.9}$$

For constant gain this reduces to $K_c(e_k - e_{k-1})$. With scheduling, the term $K_{c,k}e_k - K_{c,k-1}e_{k-1}$ also carries the gain change. The implementation applies the command before advancing its integral state, as in Equation (4.10).

$$\begin{aligned} s_k &= s_{k-1} + \Delta P_k \\ c_k &= u_{bias,k} + s_k + I_{k-1} \\ u_k &= \text{clamp}(c_k, u_{min}, u_{max}) \\ I_k &= I_{k-1} + \Delta I_k + (u_k - c_k)\Delta t / T_t \end{aligned} \tag{4.10}$$

These are Equations (1.11) to (1.13) of Chapter 1 with $s_k = S_k$ and $c_k = v_k$, and $T_t = \tau_I$ by default. The running sum is never clipped, and back-calculation acts on the integral alone, as Section 1.6.2 explains. The pitch and torque tracking times are 2.3333 and 12 s, respectively, at default tuning. Actuator rate limits act in the plant dynamics, separately from the controller output clamp. Both derivative gains are zero. In manual, and for torque while Region 3 fixes its command, the tracking path sets $I = u_{applied} - u_{bias} - s$ using the current error. This aligns the controller output with the applied command for transfer back to automatic regulation, and later error and integral evolution may move it.

The direction flag is the one place a sign error would propagate everywhere. More generator torque slows the rotor and more pitch slows the rotor, so the process gain of both loops is negative. Both loops are therefore direct acting, their output rising as the measured speed rises, and both carry $a = -1$ in the convention of Section 1.3. Setting that flag once on the shared controller is the entire sign handling in the application.

4.5.2 Region 2 Torque Control

Below rated the rotor speed setpoint is the speed that puts the machine at the peak of the C_P curve, capped just below the speed rating, as shown in Equation (4.11).

$$\omega_{SP} = \text{clamp}\left(\frac{\lambda^* V}{R},\ \omega_{min},\ 0.997\,\omega_{rated}\right) \tag{4.11}$$

Generator torque is then a feedforward plus proportional-integral trim, as shown in Equation (4.12).

$$T_{gen,cmd} = K\omega_{SP}^2 + K_{c,T}(\omega - \omega_{SP}) + \frac{K_{c,T}}{\tau_{I,T}}\int(\omega - \omega_{SP})\,dt \tag{4.12}$$

where the bias term is the exact steady-state aerodynamic torque at the optimum. Substituting $\lambda = \lambda^*$ into Equations (4.1) and (4.2) gives $T_{aero} = K\omega^2$ with $K = \rho\pi R^5 C_{P,max}/(2\lambda^{*3}) = 2.1388 \times 10^6\,\text{N m s}^2$, and a verification test in the application confirms the identity to a relative 10^{-9} at four rotor speeds.

The same identity is the basis of the alternative Region 2 law offered as an advanced option, $T_{gen} = K\omega^2$ applied directly [10]. That law carries no speed setpoint and no wind measurement, which is the reason it is the baseline in the reference designs, and its two known weaknesses are that convergence to the optimum is limited by the rotor time constant and that the result is only as good as K. Both laws take K, and Equation (4.11) takes R, from the declared values of Table 4.2 once when the application loads, so an air density or rotor radius edited in the application detunes them as it would a commissioned controller, and with a 70 m rotor at 6 m/s the $K\omega^2$ law settles at $\lambda = 8.87$ and the proportional-integral law at $\lambda = 8.39$. The law does use the measured speed ω. What it does without is the wind and a speed setpoint.

🕮 **Pause.** Try Exercise 4.3 before reading on.

4.5.3 Region 3 and the Handover

Above rated, generator torque is held at T_{rated} and the pitch loop regulates speed against the rating, as shown in Equation (4.13).

$$\begin{aligned} \beta_{cmd,k} &= \text{clamp}\left(K_{c,\beta}^{\text{eff}} e_{\omega,k} + I_{k-1},\ 0,\ \beta_{ctl,max}\right) \\ e_{\omega,k} &= \omega_k - \omega_{rated} \end{aligned} \tag{4.13}$$

Pitch control remains active in every region with a setpoint at ω_{rated}, while the torque setpoint is capped at $0.997\omega_{rated}$. This offset prevents both loops from seeking the same zero-error condition. Below rated, pitch rests at its lower limit and its integral follows the tracking law of Equation (4.10), as the protective override of Volume 4, Chapter 2 rests on its lower clamp through healthy operation. The tracking state need not lie exactly at the boundary on every scan. As speed rises, pitch can begin feathering before Region 3 is declared. Once that declaration occurs, an explicit selector commands rated generator torque and tracks the torque controller to that command. The onset of pitch regulation and the torque selector are distinct parts of the handover.

🕮 **Pause.** Try Exercise 4.1 before reading on.

4.5.4 Region Index, Hysteresis and Dwell

The region index selects rated torque in Region 3. Outside Region 3 the selected PI or $K\omega^2$ law controls torque. The index also drives annunciation. Pitch remains active throughout, and its gain schedule depends on blade angle and the schedule setting, rather than on the index. Region logic carries hysteresis on both edges and dwell timers on all but one transition. Region 3 is entered when $\beta_{cmd} \geq 0.5^\circ$, or when $T_{gen,cmd} \geq 0.999\,T_{rated}$ and $\omega \geq 1.002\,\omega_{rated}$ together. It is left only when $\beta_{cmd} \leq 0.1^\circ$ and $\omega \leq 0.990\,\omega_{rated}$ together. Those two conditions cannot both hold, so the boundary is a dead band and not a threshold. The 2 against 2.5 label is decided by rotor speed rather than by the setpoint, with 2.5 declared above $0.96\,\omega_{rated}$ and 2 below $0.90\,\omega_{rated}$, because the setpoint follows the wind and in 14 % turbulence it crosses the cap several times a second while the rotor speed is filtered by 4×10^7 kg m^2 of inertia.

Dwell timers sit on top of the hysteresis. Leaving Region 3 requires 2.0 s since the last change, and any transition between 2 and 2.5 requires 6.0 s. Entry into Region 3 is never delayed, because that transition selects rated torque while pitch is already regulating speed. A verification test runs the wind ramp from 7 to 22 m/s and requires between one and four region changes over 130 s, and the measured value is two.

4.5.5 Tuning and Defaults

The pitch loop is placed by pole assignment on the integrating rotor. Treating $J_{eq}\dot{\omega} = (\partial T_{aero}/\partial\beta)\beta$ as the plant and requiring a closed-loop damping ratio $\zeta = 0.7$ at $\omega_n = 0.6$ rad/s gives $K_{c,\beta} = 2\zeta\omega_n J_{eq}/|\partial T_{aero}/\partial\beta| = 3.775$ s and $\tau_{I,\beta} = 2\zeta/\omega_n = 2.333$ s at the fine-pitch sensitivity of 8.9 MN m/rad. The shipped defaults are $K_{c,\beta} = 3.7673$ s, displayed as 22.6 °/rpm, and $\tau_{I,\beta} = 2.3333$ s. The torque loop is tuned for a closed-loop rotor time constant near 8 s, with $K_{c,T} = 3.0 \times 10^6$ N m s/rad displayed as 314 kNm/rpm and $\tau_{I,T} = 12$ s. The gain schedule is on by default and the Region 2 law defaults to the proportional-integral form.

4.5.6 Model Predictive Controller

A second controller structure can take the place of both loops and of the torque selector of the region logic. The model predictive controller (MPC) acts only in AUTO and only when selected, and the scheduled loops remain the default of the application and of every scenario except scenarios 11 and 12. Every $T_s = 0.1$ s, or five controller cycles, it solves one quadratic program for the pitch and torque commands over $N_p = 50$ steps, or 5 s, in 12 move blocks that grow from four single steps to one of 15 steps, with the last block taking the remaining 5. The first move is applied and held until the next plan, as Section 1.6.3 describes [20].

The model is written rather than identified. It is Equation (4.14), the application's own equations evaluated with the declared values of Table 4.2, which are copied when the application loads, so a parameter that a reader edits is model mismatch.

$$\begin{aligned} J_{eq}\dot{\omega} &= T_{aero}(\omega, V, \beta) - T_{gen} - b\,\omega + d \\ \tau_\beta\dot{\beta} &= \beta_{cmd} - \beta, \qquad \tau_T\dot{T}_{gen} = T_{gen,cmd} - T_{gen} \end{aligned} \tag{4.14}$$

Here d is an unmeasured torque disturbance, and the rate limits and end stops of Equations (4.5) and (4.6) become constraints. At every decision the model is linearized at the measured speed ω_0 and pitch β_0 and the first previewed wind V_0, with $a_\omega = \partial T_{aero}/\partial\omega$ and $a_\beta = \partial T_{aero}/\partial\beta$ from finite differences, as Equation (4.15) states.

$$\begin{aligned} \dot{x} &= \begin{bmatrix} \frac{a_\omega - b}{J_{eq}} & \frac{a_\beta}{J_{eq}} & -\frac{1}{J_{eq}} \\ 0 & -\frac{1}{\tau_\beta} & 0 \\ 0 & 0 & -\frac{1}{\tau_T} \end{bmatrix} x + \begin{bmatrix} 0 & 0 \\ \frac{1}{\tau_\beta} & 0 \\ 0 & \frac{1}{\tau_T} \end{bmatrix} u + \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} f_k \\ f_k &= \frac{T_{aero}(\omega_0, V_k, \beta_0) - \hat{T}_{gen} - b\,\omega_0 + \hat{d}}{J_{eq}} \end{aligned} \tag{4.15}$$

The states $x = (\delta\omega,\ \delta\beta,\ \delta T_{gen})$ are deviations from the measured speed and pitch and the estimated torque $\hat{T}_{gen}$, the inputs $u = (\beta_{cmd} - \beta_0,\ T_{gen,cmd} - \hat{T}_{gen})$ are measured from where the actuators are, and the system is sampled at T_s with an exact zero-order hold. The previewed wind V_k of period k enters through the forcing f_k, evaluated before the solve, so the program stays linear. Electrical power is linearized as $\hat{P} = \eta(\hat{T}_{gen}\omega_0 + \hat{T}_{gen}\,\delta\omega + \omega_0\,\delta T_{gen})$. At rated speed the linearization gives $a_\beta = -9.60$, -31.14 and -46.61 MN m/rad at 11.4, 15 and 25 m/s, as in Table 4.3. Open loop from 8 to 22 m/s, its 5 s speed response to pitch and torque steps is within 6.89 % of the plant's, but it predicts a 14 to 20 m/s wind step 12.5 % low, 3.52 against 4.03 rpm, because it keeps the sensitivities of the moment.

The plan tracks the operating point that each previewed wind calls for and has no regions. The speed target is $\omega_{t,k} = \text{clamp}(\lambda^* V_k/R,\ \omega_{min},\ \omega_{rated})$, the formula of Equation (4.11) capped at the rating itself, with fine pitch and the balancing torque as the input targets $\beta_{t,k}$ and $T_{t,k}$. Where that torque would exceed T_{rated}, the speed target rises toward rated speed at fine pitch, and at rated speed the pitch target is the angle that sheds the excess, found by bisection on the declared C_P surface with $\hat{d}$ in the balance. Each plan minimizes Equation (4.16), in which β_c and T_c are the pitch and torque commands and j counts the move blocks.

$$
\begin{aligned}
\mathscr{J} = &\sum_{k=1}^{N_p} w_\omega(\hat{\omega}_k - \omega_{t,k})^2 + \sum_{o=1}^{2}\left(\kappa_o\varepsilon_o^2 + \nu_o\varepsilon_o\right) \\
&+ \sum_{k=0}^{N_p-1}\left[w_\beta(\beta_{c,k} - \beta_{t,k})^2 + w_T(T_{c,k} - T_{t,k})^2\right] \\
&+ \mu\sum_{j=1}^{12}\left[\left(\frac{\Delta\beta_{c,j}}{m_\beta}\right)^2 + \left(\frac{\Delta T_{c,j}}{m_T}\right)^2\right] \\
\text{s.t. } &\hat{\omega}_k \le \omega_{max} + \varepsilon_1, \quad \hat{P}_k \le P_{max} + \varepsilon_2, \quad \varepsilon_o \ge 0 \\
&0 \le \beta_{c,j} \le \min(\beta_{ctl}, \beta_{max}), \quad 0 \le T_{c,j} \le T_{rated} \\
&|\Delta\beta_{c,j}| \le \Delta\beta_{lim}, \quad |\Delta T_{c,j}| \le \dot{T}_{max}T_s
\end{aligned}
\tag{4.16}
$$

Speed enters in rpm with $w_\omega = 4$ per rpm^2, the inverse square of the torque loop's 0.5 rpm band, and $w_\beta = 0.04$ per deg^2, $w_T = 0.33$ per (MN·m)^2, $m_\beta = 1°$, $m_T = 0.35$ MN m and $\mu = 3$. The soft ceilings are $\omega_{max} = 1.15\,\omega_{rated}$, the overspeed alarm's reset level read from its live setting, and $P_{max} = \eta T_{rated}\omega_{rated} = 5.0$ MW, with slack prices $\kappa_1 = 10^3$ per rpm^2, $\nu_1 = 10^2$ per rpm, $\kappa_2 = 10^2$ per MW^2 and $\nu_2 = 10$ per MW, and a ceiling already crossed is re-entered along a 1 s funnel. Neither decides a result, because scenario 11 peaks at 100.35 % of rated speed without the power ceiling and at 100.29 % without the speed ceiling or with both. The command boxes are hard and read from the live settings, as the loops' clamps read them. The torque bound is hard rather than soft because the rating is the loops' clamp and the plant's end stop. The cap $\Delta\beta_{lim} = 0.95\,\dot{\beta}_{max}\tau_\beta(1 - \mathrm{e}^{-T_s/\tau_\beta}) = 0.626°$ keeps a held step through the lag below 7.6 °/s, 95 % of the limit, and a long move made of such steps averages 6.26 °/s. The torque cap is the declared rate, and there is no terminal weight. A dual active-set method [21] solves each program, and any plan without a reported optimum, with a value that is not finite or with a command outside its box is discarded.

A steady-state Kalman filter on the speed and a random-walk disturbance d makes the plan offset-free [22], the arrangement of Section 1.6.3. With $G = (60/2\pi)\,10^6\,\Delta t/J_{eq} = 4.77 \times 10^{-3}$ rpm per MN·m per 0.02 s step, the transition matrix $\left[\begin{smallmatrix}1 & G\\ 0 & 1\end{smallmatrix}\right]$, a measurement variance $r = (0.01\,\text{rpm})^2$ and a disturbance intensity $\gamma r/G^2$ with $\gamma = 1$, its gain is $L = (0.769,\ 100.6)$. The prediction is a Heun step of the declared rotor equation driven by the measured pitch, the hub wind and the torque that the commands produce through the declared lag, because the loops read no torque measurement. It runs under both structures, so a switch starts from a current estimate.

The wind over the horizon comes from an instrument with three settings. With the preview off the plan holds the hub wind measured now, the number the Region 2 setpoint already reads. The perfect

preview, a teaching bound, is the wind that will arrive, projected from the wind states, the settings and any scripted change within 5 s through the declared 2 s lag, with the turbulence drawn ahead from the plant's own generator and queued, so the plant's wind is bit-identical under both structures. The LIDAR setting, the default, adds a declared error $e_k = (\ell_k/5\,\mathrm{s})\,\sigma_L\,n$, with $\ell_k = \min\big((k+\frac{1}{2})T_s,\ 5\,\mathrm{s}\big)$, $\sigma_L = \sqrt{\sigma_u^2 + (0.01\,V)^2}$, σ_u the turbulence standard deviation of Section 4.3, and n one draw per plan of a unit-variance noise with a 2 s correlation time from the instrument's own seeded stream. The error is zero at the rotor, and at 5 s it combines the evolution of the air on its way in with a line-of-sight measurement that is not the rotor-effective wind, as in LIDAR-assisted predictive control [23]. The wave buoy of Chapter 5 plans on a forecast of its wave force in the same way, and Section 1.4.6 discusses what such a forecast is worth.

The move price and γ were tuned once, on scenario 11 with the LIDAR, over nine values of μ and five of γ, scoring the IAE relative to the loops'. No setting, and no preview, met the bound of no more travel than the loops in move-scale units (least 22.1 against 17.5), because the loops hold torque still above rated, so the bound is reported rather than imposed. The largest μ within 10 % of the best score, 0.4286 at $\mu = 1$, is $\mu = 3$ at 0.4314. Two earlier definitions of the LIDAR error, a fixed 1.0 m/s and 5 % of the wind, gave scenario 1 106.8 and 33.3 MN m of torque travel in 180 s, against 6.65 with the shipped error. Under those errors scenario 11 peaks at 100.67 and 100.51 %, so the shipped definition removes laminar jitter and does not buy the headline result.

In MANUAL the MPC plans nothing, and the manual overspeed case is identical under both structures. A rejected plan hands control to the loops, which track the applied commands every cycle, marks the historian, and returns control with the next valid plan, and the first plan is solved at $t = 0$. A switch to the MPC moves the commands by at most one capped step, and a switch back, or a fallback, moves the pitch command by at most 0.09°, but above rated the loops' constant-torque law commands the rating at once, a step of up to 0.667 MN m at 120 s of scenario 4 against 0.062 MN m for the MPC's own largest step in the 5 s before, which the actuator follows at its rate. The application ships no trip and no feather action, and the overspeed alarm, the actuator limits, the AUTO pitch limit and the torque rating act under the MPC exactly as under the loops. With the hydraulics slowed to 0.5 °/s in scenario 11 the alarm latches under both, at 33.98 s under the loops and 35.18 s under the MPC.

The plant integrates the model with the run's parameters, the instruments are the exact speed, pitch and hub wind and the preview, and the controller reads only what the loops read, the declared table, its own state and the preview. A test that gives a second plant eleven other parameters and overwrites its states with the first plant's every step finds the commands bit-identical over 6000 steps, and 24 of 24 mutated copies that read a true state, a run parameter or extra future wind fail the tests. A plan takes about 0.2 ms in Chrome, 172 to 210 µs, and the whole controller about 3 % of a core at ten times real time, at 59.8 to 60.0 frames per second.

4.5.7 Scenario Tree on the Declared Preview Error

The MPC takes the previewed wind as exact, although the LIDAR declares an error with it. A secondary setting of the MPC, its wind scenarios, keeps that plan as the default (One) and adds a scenario tree (Tree of 3), the multi-stage form of MPC [24, 25], that plans on the forecast and on the forecast plus and minus its declared error. Each wind keeps a plan of its own after the first move, because the next decision will see which wind arrived, and only the move applied before that decision is shared, the condition known as non-anticipativity. The objective is the expected cost over equally likely winds [26], not the worst case of min-max feedback MPC. Every scenario except scenario 12 ships the setting at One.

The winds are $V_k^{(0)} = V_k$ and $V_k^{(1,2)} = V_k \pm \bar{e}_k$, where $\bar{e}_k = (\ell_k/5\,\mathrm{s})\,\sigma_L$ is the one-standard-deviation envelope of the LIDAR error $e_k = \bar{e}_k n$, computed from the channels the preview reads and never from the error it draws. Five seconds ahead $\bar{e}$ is 1.95 m/s at 14 m/s in 14 % turbulence, and over ten seeds of scenario 12 the error the LIDAR makes five seconds ahead, divided by $\bar{e}$, has a standard deviation

of 1.01. With the preview off the envelope surrounds the held hub wind, and with the perfect preview it is zero. Three equally weighted winds one standard deviation apart carry two thirds of the declared variance, a spread fixed before any run.

Each wind has its own copy of Equation (4.15) at the same linearization point, with its own forcing and targets, which gives nine states and six inputs over the same 50 steps and 12 move blocks, and the plan solves Equation (4.17).

$$\min \frac{1}{3}\sum_{i=0}^{2} \mathscr{J}^{(i)} \quad \text{s.t.} \quad \beta_{c,1}^{(i)} = \beta_{c,1}^{(0)}, \;\; T_{c,1}^{(i)} = T_{c,1}^{(0)}, \;\; i = 1,2 \tag{4.17}$$

Here $\mathscr{J}^{(i)}$ is Equation (4.16) on the wind $V^{(i)}$, with slacks of its own and every weight, price, cap and box unchanged, and each equality enters the solver as two inequality rows. The program has 78 variables and 602 rows against the one-wind plan's 26 and 198. With no declared error the branches coincide, and the move equals the one-wind plan's within 3.1×10^{-9} degrees over 7395 decisions in four scenarios. A plan is also rejected when a command of any branch leaves its box or the three first moves differ by more than 10^{-6}.

The hedge is small by construction. When every branch carries the same model and the same weights and no constraint is active, minimizing over each branch's later moves leaves one quadratic in the shared move whose linear term is the mean of the branches' terms, so the shared move is the mean of the first moves of three one-wind plans made on $V^{(0)}$, $V^{(1)}$ and $V^{(2)}$. The plans on the higher and the lower wind nearly cancel, and the remainder comes from the curvature of aerodynamic torque in the wind and from the fine-pitch bound. Although a soft ceiling or a box is active at every decision of scenario 12, the tree's first pitch move lies within 0.01° of that mean at 1136 of the 1199 decisions, and it exceeds the one-wind plan's by 0.027° on average, where a one-wind plan told $V + \bar{e}$ moves 0.144° more.

A tree plan takes 2.9 ms in Chrome against 0.21 ms for one wind, about fourteen times as long, and the controller uses 3.2 % of a core at 1× and 32 % at the 10× top speed, with 60 frames per second at both on a host carrying other work. Because the plan exceeds the 2 ms threshold the apps set for a plan, a frame guard stops stepping a frame after 25 ms of wall time while the tree plans, so a slower computer runs below the requested speed with the same steps and the same logged rows. A switch between one wind and the tree in the middle of a gust, with the turbulence and the preview off, moves the pitch command by exactly the 0.061° step that the plan takes without the switch.

4.6 Results

Table 4.4 lists the thirteen set-ups. Each pins only the quantities it is about and leaves everything else at the defaults of Table 4.2, so a scenario differs from the base machine in exactly the settings its row names.

Table 4.5 reports the scorecard for the declared horizons. Runs use the shipped engine with its 0.01 s physics step, exact control timing and the default seed, which is the app's 20260824 unless the scenario declares its own. Scenario selection restores its settings, and comparisons change only the stated setting before the first step. IAE integrates absolute rotor-speed error against the logged torque-loop setpoint, in radians, and travel is the sampled total variation of actual generator torque at 10 Hz. Saturation uses the logged controller flag, matching the torque trend shading. Deliberately holding rated torque in Region 3 is excluded because pitch regulates speed there. The rating remains visible on the torque strip. Physical torque-bound occupancy can be computed separately from the CSV.

Rows 11 and 12 run under the MPC of Section 4.5.6 with the LIDAR preview, as the scenarios load (row 12 with the scenario tree of Section 4.5.7), and every other row runs under the scheduled loops or, in scenario 10, in MANUAL.

Table 4.4 Scenarios and What Each One Changes

#	Name	Setting
1	Steady 6 m/s	Region 2 at the C_P peak, no turbulence
2	Wind ramp through rated	7 to 22 m/s at 0.15 m/s per s, $t = 10$ to 110 s
3	Step gust	9 m/s, +6 m/s at $t = 25$ s, back at 70 s
4	Turbulent wind	11 m/s mean, 14% intensity, seeded
5	Pitch rate limit	$\dot{\beta}_{max} = 2$ deg/s, 13 to 19 m/s step
6	Generator torque limit	$T_{rated} = 2.8$ MN·m at 10.5 m/s
7	Gain schedule off	21 m/s, 14% turbulence, +4 m/s at $t = 30$ s
8	Doubled rotor inertia	$J_{eq} = 8.0 \times 10^7$, 8 to 14 m/s step
9	Region-transition hunting	10.9 m/s mean, 9% turbulence, three mean shifts
10	Manual overspeed	MANUAL, 16 m/s, $T_{gen} = 0.4$ MN·m, $\beta = 0$
11	MPC or PI loops in a gust	MPC with LIDAR preview, 14 m/s, +6 m/s at $t = 30$ s, back at 45 s
12	Plan on one wind or three	MPC tree with LIDAR, the gust of 11 in 14% turbulence, seeded
13	Challenge	weak gains, schedule off, ramp to 18 m/s

Table 4.5 Measured Scorecard Values at the End of Each Run

#	t (s)	IAE (rad)	MV travel (MN·m)	Sat. (%)	Peak ω (% rated)	Peak $\dot{\beta}$ (deg/s)	P (MW)
1	60	0.000	0.00	0.0	57	0.0	0.75
2	130	2.911	2.68	0.0	101	3.0	5.00
3	120	4.473	3.44	0.0	106	8.0	2.42
4	120	4.819	14.72	2.2	105	7.1	4.90
5	90	1.292	0.00	0.0	111	2.0	5.00
6	90	0.789	0.00	0.0	100	4.6	3.35
7	120	0.865	0.00	0.0	102	8.0	5.04
8	120	1.838	2.22	0.0	104	6.0	5.00
9	180	5.652	11.58	0.3	104	7.2	4.97
10	200	566.646	3.63	0.0	345	0.0	1.65
11	120	0.447	1.07	0.0	100	2.5	4.97
12	120	0.913	6.21	0.0	102	6.1	4.53
13	200	25.981	9.77	0.0	126	1.8	5.13

Above the torque setpoint cap, a rotor held at rating still differs from the logged target by 0.0363 rpm. That offset contributes to IAE, but it cannot in general be subtracted as a common floor because the metric takes absolute error. Pitch regulation should also be assessed against rated rotor speed.

4.6.1 Handover Through Rated Wind

Scenario 2 ramps wind from 7 to 22 m/s at 0.15 m/s per second. The target reaches its cap at 10.54 m/s. Region 2.5 is declared at 34.76 s, with wind 10.714 m/s, and Region 3 at 37.06 s, with wind 11.059 m/s. At the latter change actual generator torque is about 2.96 MN·m. The selector now

commands 4.18 MN·m and the rate-limited actuator catches up. Thus the region change can occur before actual torque reaches its rating. Pitch was already active and supplies the entry condition. At 130 s the blades have feathered to about 19.4 degrees, power is 5.00 MW, and speed is within 0.001 rpm of rating. The logged 0.0363 rpm residual is the setpoint offset.

> **TRY IT IN THE APP** **2 · Wind Ramp Through Rated**
>
> Load scenario 2 and press Run in Auto, at 5 times speed if you like. Watch the region badge. It turns to 2.5 at 34.76 s and to 3 at 37.06 s, while the generator torque is still about 2.96 MN m. By 130 s the blades have feathered to about 19.4° and the power holds 5.00 MW.

At steady 6 m/s, scenario 1 holds 6.87 rpm and $C_P = 0.482$, producing 0.746 MW electrical from about 1.65 MW available wind power. Changing the controller target ratio to 5.5 and running 240 s gives $\lambda = 5.500$, $C_P = 0.4183$ and 0.649 MW, a 13.0% loss for the same wind and machine.

4.6.2 Cost of Disabling the Gain Schedule

Scenario 7 starts at 21 m/s with 14% turbulence and adds 4 m/s at 30 s. The same seeded 120 s input is used with scheduling off and on. Without the schedule, IAE is 0.865 rad and peak speed 102% of rating. With it, IAE is 3.058 rad and peak speed 107%. Terminal logged errors are 0.14 and 0.26 rpm, respectively. This endpoint is one turbulence realization, not a steady error. For the ideal integrating plant, increasing both PI gains at fixed integral time raises both natural frequency and damping ratio with the square root of gain. A faster trace therefore need not ring in this reduced model.

Actual pitch travel is 255.2 degrees without scheduling versus 108.4 degrees with it, a factor of 2.35. Command travel is 350.9 versus 118.4 degrees. Peak pitch rate reaches the 8 deg/s constraint without scheduling, but is 5.2 deg/s with it. This measures actuator duty, not a quantified fatigue life or a stability margin against the omitted tower and drivetrain modes.

> **TRY IT IN THE APP** **7 · Gain Schedule Off, High Wind**
>
> Load scenario 7 and run it for 120 s as it loads, with the Pitch gain schedule off, and note that the Peak pitch rate reaches the 8 °/s limit. Reset, switch Pitch gain schedule to On, change nothing else and run again. The speed IAE rises from 0.865 to 3.058 rad, but the blade travel falls from 255.2 to 108.4°.

Pause. Try Exercise 4.2 before reading on.

4.6.3 Actuator and Rating Limits

Scenario 5 reduces the pitch rate limit from 8 to 2 deg/s, one quarter of the default, and steps wind from 13 to 19 m/s. Peak speed is 111% of rating. Raising pitch gain from 3.7673 to 6 with the same 2 deg/s actuator changes the response: IAE rises from 1.292 to 7.452 rad and peak speed from 111 to 116%. It cannot increase actuator speed. Restoring 8 deg/s at default tuning reduces IAE to 0.749 rad and peak speed to 106%, with peak actual pitch rate 3.6 deg/s. Gain therefore still affects command history under a rate constraint, and "more gain changes nothing" would be an incorrect description.

Scenario 6 derates torque to 2.8 MN·m at 10.5 m/s. It starts in Region 3, settles at 3.13 degrees of pitch and 12.10 rpm, and produces 3.35 MW. Actual torque is at its bound throughout, but both the scorecard saturation value and its controller flag remain zero, because pitch has control authority.

Scenario 8 doubles inertia to 8×10^7 kg m^2. It peaks at 104% rated speed with a 6.0 deg/s peak pitch rate, and the scorecard settling value is 10.7 s. Doubling inertia halves acceleration for a fixed imbalance. At fixed PI gains, the ideal integrating-plant natural frequency scales as $J^{-1/2}$, so the poles do not simply halve their frequency. Restoring a chosen pole placement requires gains proportional to inertia.

Table 4.6 Scenario 11 Over 0 to 120 s, With the Ten-Seed Peak Speed and Energy of Scenarios 4 and 9 Over 0 to 180 s

	PI loops	PI, sched. off	MPC, preview off	MPC, LIDAR	MPC, perfect
Peak ω (% rated)	105.59	101.56	102.39	100.29	100.34
Speed IAE (rad)	1.036	0.558	0.595	0.447	0.468
Pitch travel (deg)	17.5	17.0	16.9	20.5	17.4
Torque travel (MN·m)	0.00	0.00	0.72	1.07	0.73
Peak P (MW)	5.28	5.08	5.01	5.00	5.00
Energy (kWh)	166.8	166.7	166.1	165.8	166.1
Mean solve, headless (µs)	—	—	161	179	158
Ten turbulence seeds, 0 to 180 s					
Sc. 4 peak ω (%)	104.72	—	102.42	101.32	100.97
Sc. 4 energy (kWh)	216.50	—	209.44	205.26	209.28
Sc. 9 peak ω (%)	103.65	—	101.74	101.36	100.83
Sc. 9 energy (kWh)	225.83	—	221.20	219.93	221.98

4.6.4 Region Transitions in Turbulence

Scenario 4 produces twelve region changes in 120 s, and scenario 9 produces fourteen in 180 s and IAE 5.652 rad. The pitch loop stays active, but crossing Region 3 changes the torque selector and its tracking mode. Actual torque responds through its lag and rate constraint, so a region change should not be described as a purely cosmetic label. The hysteresis and dwell limit repeated selection near the boundary. These counts alone do not quantify fatigue or predict operation without that logic.

4.6.5 Predictive Control Against the Scheduled Loops

Scenario 11 holds the machine at rated in a laminar 14 m/s, raises the mean wind to 20 m/s at 30 s and returns it at 45 s, and it loads the MPC with the LIDAR preview. Table 4.6 compares it over 120 s with the shipped loops, whose tuning is the challenge's suggested tuning, the loops with the schedule off and the MPC under the other two previews, in the same wind, with no case retuned and 1200 plans per run, none rejected. The MPC with the LIDAR is told the wind ahead and the loops are not, so the like-for-like comparison is the preview-off column, given only the hub wind that the loops also read.

Under the loops the wind at the rotor starts to rise at 30.03 s, and the speed peaks at 105.59 % of rating at 33.12 s while the power reaches 5.28 MW, as shown in Figure 4.3. Told of the gust, the MPC raises its pitch command 0.5° above its line-out at 28.0 s, holds 11.4° at 30 s against the loops' 9.0°, and keeps the speed at or below 100.29 %, a value reached at 45.65 s as the gust passes, with the power at its 5.00 MW rating. With the preview off it peaks at 102.39 % at 31.93 s. The preview is therefore worth 2.1 of the 5.3 points, 40 % of the gain over the loops, and its share stays between 29 and 41 % for gusts of 3, 6 and 9 m/s, a gust from 12 m/s, a gust lasting 3 s, and 5 or 9 % turbulence over ten seeds, where the LIDAR peaks below the preview off at every seed.

The remaining 3.2 points come from how tightly the MPC's weights hold the speed, not from its model reading the hub wind. Given no wind reading at all, with a 30 s average of the hub wind in its model and filter, the MPC with the preview off peaks at 102.50 %, and with its torque held at the rating at 102.21 %, while its move price moves the peak to 101.26 % at $\mu = 0.3$ and 104.16 % at $\mu = 30$. A shipped option of the loops that is as tight does as well on this gust. With the schedule off the loops peak at 101.56 % with an IAE of 0.558 rad and 17.0° of pitch travel. The tight loops pay elsewhere, because on the gust from below rated of scenario 3 and in the turbulence of scenarios 4 and 9 the schedule off and the doubled pitch gains ride the 8 °/s limit for 0.90 to 5.56 s and still reach 102.59 to 104.74 %, where the MPC with the preview off reaches 100.61, 103.02 and 101.41 % without touching the limit. In turbulence the hub-wind reading does matter, and on scenario 9 the MPC without it peaks at 102.48 % with an IAE of 2.60 rad, against 101.41 % and 2.03 rad.

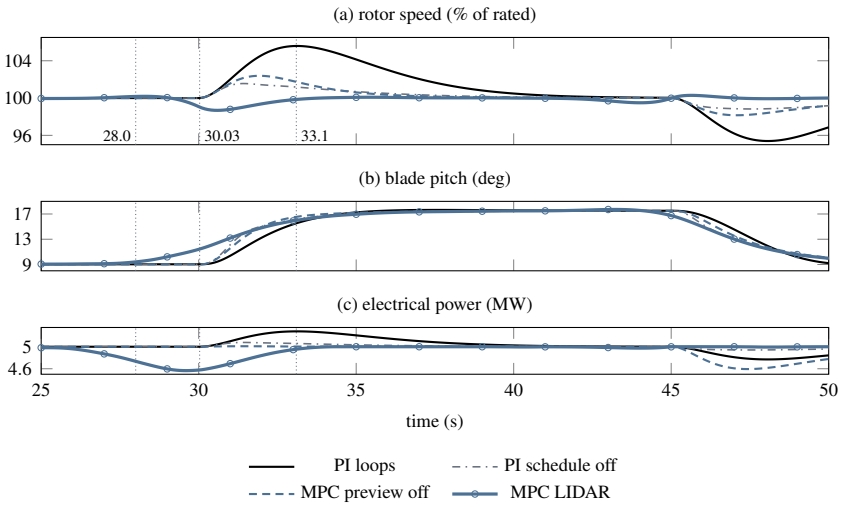


Figure 4.3 Scenario 11 from 25 to 50 s under the scheduled loops (black, solid), the loops with the schedule off (gray, dash-dotted), the MPC with the preview off (blue, densely dashed) and the MPC with the LIDAR preview (blue, very thick, with circles). The thin dotted vertical marks are the LIDAR plan's first pitch move of 0.5° at 28.0 s, the wind reaching the rotor at 30.03 s and the loops' peak at 33.1 s. (a) Rotor speed. (b) Blade pitch. (c) Electrical power.

The preview has a cost before the gust arrives. The plan feathers early and lowers its torque command to 3.83 MN m to keep the speed on target, so the power falls to 4.57 MW at 29.6 s while the loops hold 5.00 MW, as Figure 4.3c shows. The LIDAR's small error keeps the plan moving, so the MPC travels 20.5° of pitch against 17.5 and 1.07 MN m of torque where the loops hold rated torque still, and it makes 0.6 % less energy, 165.8 against 166.8 kWh, of which 0.41 kWh is the loops' power above the rating. A gust a second early or two seconds late still peaks at 100.19 to 100.49 %, and one announced that never comes costs 9.3° of needless pitch travel.

The energy losses in turbulence are larger, as the lower block of Table 4.6 shows. Over ten seeds the MPC with the LIDAR makes 5.19 % less energy than the loops in scenario 4, 3.26 % less with the preview off and 3.34 % with the perfect preview, and 2.61 % less in scenario 9. At the scenario's own seed, as it loads, the loss in scenario 4 is 5.8 %. In scenario 4, 10.07 of the 11.24 kWh are lost while the hub wind is below rated, where the MPC holds its speed target and keeps the blades pitched beyond 0.5° for 59.75 s of each run against the loops' 8.72 s, capturing 217.0 against 228.2 kWh from the wind, while the loops let the rotor speed up in the gusts and give that energy back at fine pitch and rated torque. The loops' power above the rating is only 0.89 kWh of the difference, and $\mu = 30$ does not recover it. The MPC was tuned on the laminar scenario 11 with energy outside the score, and this loss appeared afterward. The MPC therefore trades energy for speed regulation and is not the better controller overall. It also moves the torque 2.4 and 1.8 times as far in scenarios 3 and 9, raises scenario 9's power standard deviation from 0.42 to 0.52 MW, travels 6.65 MN m of torque in the steady wind of scenario 1, and has an IAE of 1.624 against 1.272 rad in scenario 7 at a third of the unscheduled loops' pitch travel. Figure 4.4 draws the peak speeds and the energy of Table 4.6 as bars.

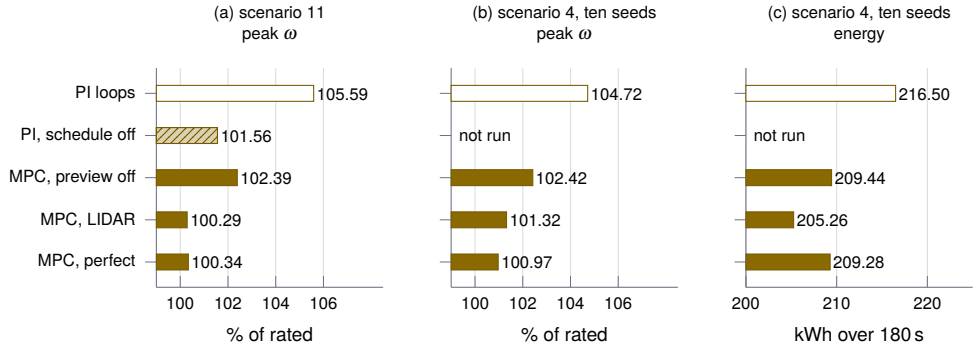


Figure 4.4 Table 4.6 as a chart: (a) the peak rotor speed in the scripted gust of scenario 11, (b) the peak speed and (c) the energy over ten turbulence seeds of scenario 4. The white bar is the shipped loops, the hatched bar the loops with the schedule off, solid bars the predictive controller under its three previews. The axes of (a) and (b) start at 99 % of rated and that of (c) at 200 kWh, so bar lengths show differences, not ratios. Notice that every predictive bar is shorter than the loops' in (b) and in (c): the controller that holds the speed tighter in turbulence also makes less energy.

> ✦ **DID YOU KNOW?** **Tighter Speed, Less Energy**
>
> From the simulator: over ten turbulence seeds of scenario 4, the predictive controller with the LIDAR preview held the peak rotor speed to 101.32 % of rated against the loops' 104.72 %, yet made 205.26 kWh against 216.50 kWh (Table 4.6). Better speed regulation did not bring more energy.

The same settings help where the loops lack authority. Over 180 s the MPC peaks at 100.32 % against 105.92 % on the gust of scenario 3, at 100.11 % against 111.27 % with the 2 °/s actuator of scenario 5, which its model still takes to be 8 °/s, and at 100.72 % against 104.17 % with the doubled inertia of scenario 8. After a Reset into AUTO from scenario 10, with the blades flat at 16 m/s, it overshoots to 112.68 % against 109.27 %, and the pitch move cap accounts for all of it, because it holds the command to 6.26 °/s while the loops run on the limit. With the cap at the full 8 °/s average the MPC peaks at 110.60 %, and with no cap at 108.75 %. Across the eleven scenarios that ran in AUTO before scenario 12 was added and the three previews, 66 000 plans were solved in runs of 200 s, none was rejected, and no command left its live bounds.

4.6.6 Planning on Three Winds

Scenario 12 repeats the gust of scenario 11 in 14 % turbulence and loads the MPC with the LIDAR and the tree. Table 4.7 compares seven structures over 0 to 120 s on the same ten turbulence seeds, with the MPC rows at the scenario's settings and nothing retuned. The hub wind hashes identical under the loops and the three LIDAR plans, and so do those plans' previews. Two rows are rivals on the same preview, the one-wind plan given the forecast plus its declared error, which is the cheapest use of the tree's information and is not offered in the application, and the loops with the LIDAR as a collective pitch feedforward 0.3 s ahead, a lead chosen on this scenario. The application has no tower and shows no load, so the thrust column is an estimate made outside it by ideal momentum theory, with the run's $C_P = 4a(1-a)^2$ solved for the axial induction $a \le 1/3$ and $F = \frac{1}{2}\rho\pi R^2 V^2\, 4a(1-a)$, a stand-in for tower load and nothing more.

Paired by seed, the tree peaks 0.10 points below the one-wind plan (by 0.02 to 0.47) with 2.7 % less estimated thrust, 1.0 % less pitch travel and 0.6 % less energy, each at all ten seeds. At the scenario's own seed the difference is 0.05 points, 101.99 against 102.04 % at 40.24 s, which the scene cannot show and the trends can, as shown in Figure 4.5. The one-wind plan on $V + \bar{e}$ peaks 0.52 points below the tree at all ten seeds, with 7.1 % less pitch travel and 1.4 % less estimated thrust (at

Table 4.7 Scenario 12 Over 0 to 120 s, Mean of Ten Turbulence Seeds

	Peak ω (% rated)	Thrust est. (MN)	Pitch (deg)	Torque (MN·m)	Energy (kWh)
PI loops	107.85	0.755	169.5	1.50	166.6
PI, LIDAR feedforward	101.76	0.671	270.1	1.37	166.3
MPC, one wind	101.76	0.589	124.0	9.50	160.8
MPC, tree	101.66	0.573	122.7	10.08	159.9
MPC, one wind on $V + \bar{e}$	101.14	0.565	113.9	11.33	154.8
MPC, one wind, preview off	104.00	0.674	211.0	16.02	161.9
MPC, tree, preview off	103.51	0.654	205.2	14.60	161.7

Every MPC row solves 1200 plans a run and rejects none. The thrust is an estimate by momentum theory that the application does not compute.

six seeds), but it makes 3.2 % less energy than the tree, a stronger hedge rather than a better one. At equal energy neither rule dominates, because told a quarter of the error the one-wind plan makes 159.7 kWh, peaks at 101.63 % and travels 121.0° against the tree's 159.9 kWh, 101.66 % and 122.7°, but carries 0.583 against 0.573 MN of estimated thrust. The feedforward loops match the one-wind plan's 101.76 % with 2.2 times its pitch travel and 3.39 s a run on the 8 °/s rate limit, which no MPC row touches. With the preview off the tree peaks 0.49 points below the one-wind plan at all ten seeds. Halving the integration step to 0.005 s leaves every paired verdict unchanged.

The declared error describes turbulence and not the timing or the size of a gust. With the plant's gust a second early, one or two seconds late, 2 m/s stronger or absent while the LIDAR reports the scripted one, the tree stays 0.08 to 0.15 points below the one-wind plan, as under the announced gust, and in the laminar wind of scenario 11 it equals the one-wind plan to three decimals under all six forecasts. In twelve worst cases, among them extreme wind, slowed hydraulics, a derated generator, a lowered alarm and model mismatch, 27 600 tree plans were solved with none rejected and no command outside its box.

4.6.7 Manual Operation and the Graded Challenge

Scenario 10 commands 0.4 MN·m and zero pitch at 16 m/s. From its balanced initial plant state, actuator motion creates a torque imbalance and the rotor accelerates to 41.77 rpm, or 345% rated, at $\lambda = 17.22$. The alarm latches at 1.60 s and the simulation continues. It is an annunciator, not a protective shutdown model. A real protection system is outside the modeled control layer.

The challenge checks the logged speed error from 132 s onward, Region 3 at the final sample, peak speed below 120% of rating, and peak pitch rate strictly below the configured actuator limit. The challenge inherits weak torque and pitch gains with scheduling off and a ramp from 9 to about 18 m/s under 8% turbulence. At 200 s this tuning has IAE 25.981 rad, 9.77 MN·m torque travel and peak speed 126% rated. Its worst logged speed error from 132 s onward is 1.086 rpm, so it fails both speed and overspeed requirements. Suggested tuning, applied before running, gives IAE 4.621 rad, 5.26 MN·m travel, peak speed 104%, peak pitch rate 6.4 deg/s and worst error 0.336 rpm in that window. It passes the challenge predicate, with 82.2% less IAE and 46.2% less torque travel.

With only the controller switched to the MPC the challenge also passes, with a worst error of 0.22 rpm from 132 s, a 102 % peak, 2.7 °/s of peak pitch rate and an IAE of 2.288 rad. The predicate scores no energy, of which the MPC makes 1.5 % less than the suggested loops, 0.2576 against 0.2614 MWh, with 19.42 against 5.26 MN m of torque travel. The challenge and its rules are unchanged.

4.7 Wind Measurement in the Region 2 Setpoint

Equation (4.11) needs the wind speed, and the application supplies the true instantaneous value including the turbulent component, sampled at the same 50 Hz the controller runs at. That is a perfect rotor-effective anemometer, and it is precisely the measurement a real turbine cannot make. A nacelle

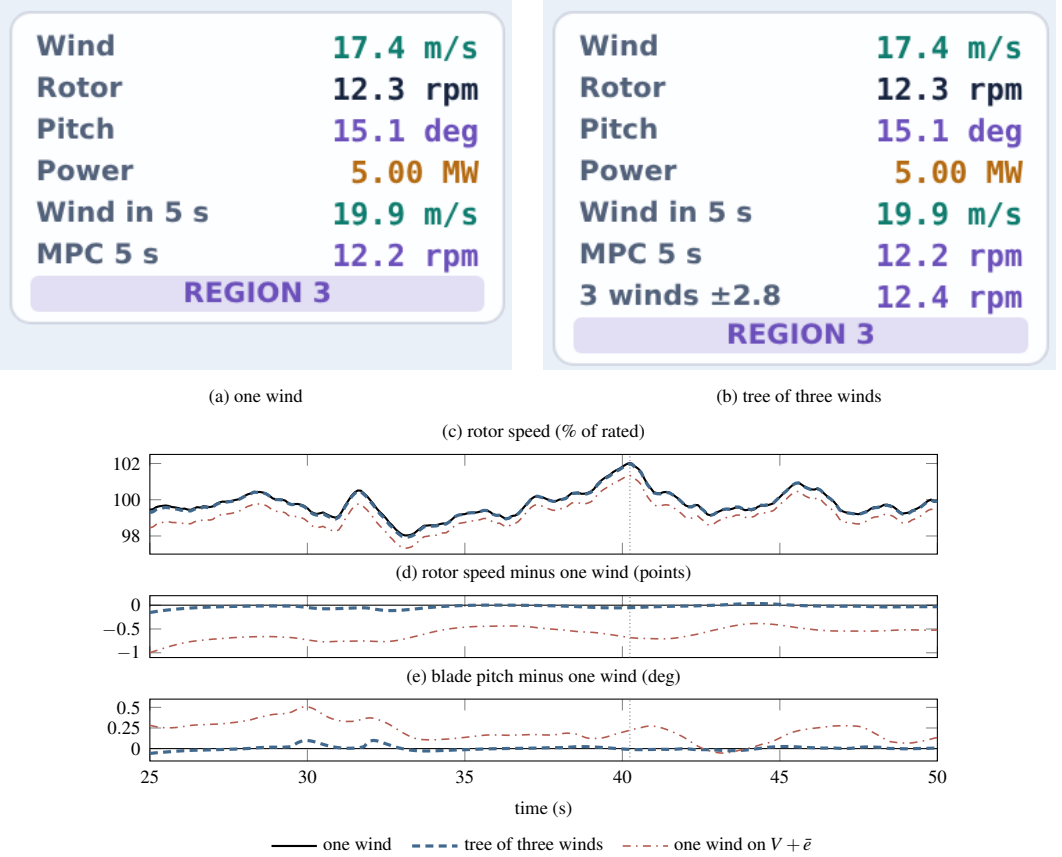


Figure 4.5 Scenario 12 at its own seed under the MPC with the LIDAR. (a), (b) The operating box at 40.3 s, just after both plans peak, where the tree adds the branches' spread 5 s ahead and the highest speed any branch predicts. (c) to (e) Rotor speed and the speed and pitch differences from the one-wind plan (black, solid) for the tree (blue, very thick, densely dashed) and the one-wind plan on $V+\bar{e}$ (red, dash-dotted), with the thin dotted line at the 40.24 s peak.

anemometer sits in the wake of the rotor it is supposed to be measuring, it reports one point rather than an average over a 12 469 m^2 disc, and the quantity the setpoint requires is the rotor-effective wind speed, which is a spatial average weighted by the aerodynamic loading. Lidar instruments measure the inflow ahead of the rotor and are an active research topic for exactly this reason [13].

The consequence has to be stated with the results. The Region 2 loop as written is a setpoint-tracking loop given information no machine has, and it converges faster than any real implementation would. The advanced option that replaces it, $T_{gen} = K\omega^2$, is the wind-free alternative and is what the reference designs use below rated [10, 14]. Running the two against each other in the same wind is the cleanest way in the application to show what a measurement is worth, and it is also the cleanest way to show what a feedforward term costs when its constant is wrong, because K depends on $C_{P,max}$ and λ^{*3} and both are properties of a rotor that erodes, ices and accumulates dirt over its life.

The MPC of Section 4.5.6 makes the information question explicit, because with the LIDAR preview it is told the next five seconds of wind and the loops are not, so every comparison in Section 4.6.5 carries the preview-off column, given the same hub wind that Equation (4.11) reads. In turbulence the preview's advantage rests on the instrument model, which hands the LIDAR the turbulence that will arrive with an error that is zero at the rotor. Told only the scripted mean wind,

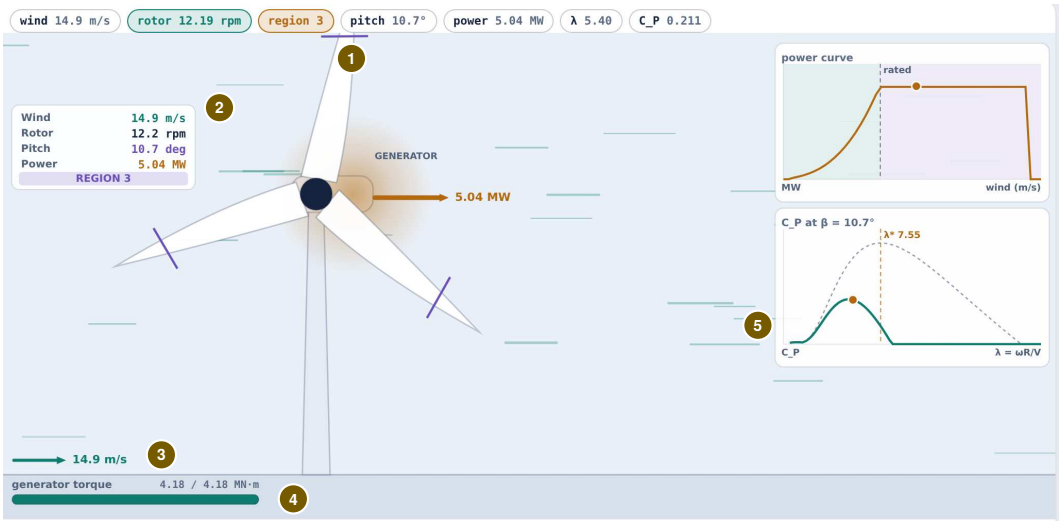


Figure 4.6 Wind ramp through rated, paused at $t = 62.7$ s, cropped to the plant view. The region changes came at 34.8 and 37.1 s, after which generator torque holds 4.18 MN m and the blades feather to 10.7° while rotor speed holds 12.19 rpm against its 12.10 rpm rating and power holds 5.04 MW. Callouts: (1) the badge row with the rotor speed and the region, (2) the readout box of wind, rotor speed, pitch and power, (3) the wind arrow, (4) the generator torque bar at its rating, and (5) the power curve and the power coefficient curve at the present pitch, with λ^* marked.

with the turbulence ahead decaying from its present value and a 1 % error, the MPC does worse than with the preview off, peaking over ten seeds at 102.15 against 101.74 % on scenario 9 and 103.21 against 102.42 % on scenario 4, and an instrument that sees the coming turbulence only through a 6 s moving average also loses, at 102.70 and 103.84 %. At three times the declared turbulence error plus 5 % of the wind, scenario 9's speed IAE is 5.686 rad against the loops' 5.652, with 2.6 times their pitch travel. In laminar wind only the 1 % part acts, which is why the scripted gust of scenario 11 is where the preview's contribution is robust.

4.8 Technology in Practice

The loops of Section 4.5 read one speed and one wind and move two actuators, the converter that sets the generator torque and the pitch system that turns the blades. Figure 4.7 draws them with that hardware in place of the schematic's symbols.

What Measures It

Rotor speed is the one controlled variable. In the reference design it is measured on the generator, after the gearbox, and filtered before either controller sees it [14]. A classic speed sensor is the tachogenerator, a small direct-current generator whose voltage is proportional to its shaft speed [27]. The wind is harder to read. A nacelle anemometer sits in the wake of its own rotor, as Section 4.7 explains, which is one reason the reference designs use the wind-free torque law below rated. The chapter hands its controller the exact rotor speed and the true rotor wind, so a filter's lag and an anemometer's error are both absent from its results.

turbine controller: SIC-101, SIC-102 and SY-103 every 0.02 s (50 Hz). The converter follows its command through a 0.05 s lag and 1.5 MN m/s; the pitch drives through 0.25 s and 8 °/s.

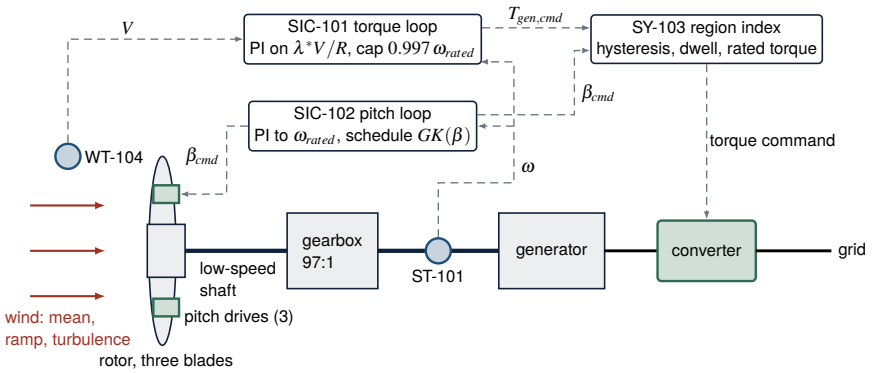


Figure 4.7 The torque and pitch loops drawn with their hardware. The speed sensor ST-101 on the generator shaft feeds both loops. The torque loop SIC-101 tracks the speed that the rotor wind WT-104 calls for, capped just below rated, and the pitch loop SIC-102 regulates to rated speed through the three pitch drives, with its gain divided by the schedule $GK(\beta)$. The region index SY-103 reads the pitch command and passes the torque loop's command below Region 3 and rated torque in it. The predictive controller, which can replace both loops and the selector, is left out. Notice that both loops act on the same speed and that only the selector decides which one sets the torque.

> **TECHNOLOGY** **Sensor: Generator Speed for Both Loops**
>
> In a utility-scale turbine both the torque and the pitch controller use the generator speed as their sole feedback input. The reference design filters it with a single-pole low-pass filter whose corner, 0.25 Hz, is about a quarter of the blade's first edgewise frequency, so that blade vibration does not excite the loops [14]. A research turbine reads its shaft and pitch positions with absolute position encoders sampled at 100 Hz [8]. Losing the speed signal is the dangerous failure. A drive's speed loop then reads a stopped shaft and runs the motor away to full speed, and a large propulsion drive carries two tachogenerators for redundancy [27]. On a trend the reading falls to zero while the true speed rises.

What Moves It

Below rated the converter moves the rotor. It sets the generator torque electronically, and on a research turbine the torque loop has a rated bandwidth of 500 rad/s [8], so torque is the fast actuator. The chapter's converter follows its command through a 0.05 s lag and a 1.5 MN m/s rate limit. Above rated the pitch system takes over, and its rate limit, not its gain, sets how fast it can meet a gust, as scenario 5 of Section 4.6 shows.

> **TECHNOLOGY** **Final Element: The Blade Pitch System**
>
> Each blade turns on a large grease-lubricated rolling bearing, most often a four-point contact ball bearing, driven through spur gears by geared motors or by hydraulic cylinders [28]. On a research turbine each blade has its own electromechanical servo that reaches 18 °/s, and the servo follows a pitch rate command with a fitted time constant of 1/60 of a second [8]. The reference design limits pitch to 8 °/s between 0 and 90° [14], the limits the chapter adopts with a 0.25 s lag. A slow pitch system shows on a trend as a pitch rate pinned at its limit and a rotor speed that overshoots, as in scenario 5.

Where the Controller Runs

A research turbine's control computer collects its measurements and closes its loops at a cycle rate of 100 Hz [8]. In the reference design the generator torque is a tabulated function of the filtered

generator speed across five regions, and the pitch loop is a gain-scheduled proportional-integral controller whose integral is saturated at the pitch limits for a fast transition between Regions 2 and 3 [14]. In process plants, programmable logic controllers (PLCs) close local feedback loops under a supervisory system, at cycle times from milliseconds to minutes [29]. The chapter's controller runs every 0.02 s, and it holds the handover in three places: the setpoint offset, the region index with its hysteresis and dwell, and back-calculation on the clamped loop.

> **IN PRACTICE** **Pitch Bearings Pay for Pitch Travel**
>
> Pitch bearing loads come from aeroelastic simulation of the design load cases of IEC 61400-1. The damage with the most critical consequence is a cracked bearing ring, found at bolt holes or stiffness changes, since it risks the blade detaching, and replacing a pitch bearing is costly because the blade must come off the hub [28]. Small oscillations wear the raceway, as standstill marks or false brinelling. Protection runs of a few degrees, and moves of 30° or more once a day, spread the grease again [28]. The pitch travel the chapter reports is the duty these bearings carry.

What Is Modernizing It

Three developments touch the loops of this chapter. Individual pitch control turns each blade on its own rather than all three together, to reduce fatigue loads, at the possible cost of slightly less power and more cyclic loading of the pitch bearing raceways [28]. Turbine controllers are moving toward feedforward that watches the incoming flow and flies the rotor in response [9], the role the LIDAR preview plays in Section 4.5.6. And because every wind turbine is digitally controlled, its operation has to be secure against cyber-physical attack [9], as Ethernet, Internet Protocol and wireless devices raise the likelihood of incidents across control systems [29].

4.9 Limitations

The model omits the two dynamics that set the real bandwidth limit on a Region 3 pitch loop. There is no drivetrain torsion and no tower motion, so the argument that an unscheduled loop is dangerous rather than merely expensive is made in prose here and cannot be demonstrated in the application. What the application can demonstrate is the actuator duty and the time spent against the rate limit. The omitted modes are close enough to the loop to matter. The first tower fore-aft mode of the reference machine is at 0.32 Hz [14] and the first drivetrain torsional mode is near 1.7 Hz, against a pitch loop placed at $\omega_n = 0.6\,\mathrm{rad/s}$, or 0.095 Hz. The scheduled tuning therefore sits a factor of 3.4 below the tower mode, which is the separation a commercial design aims for. On an integrating plant under proportional-integral control ω_n scales as the square root of the gain, so the unscheduled loop at five times the gain moves to 0.21 Hz and the separation falls to 1.5. That is where a real machine would begin to excite the tower, and the application cannot show it.

> **OPEN CHALLENGE** **Larger, Softer Rotors**
>
> Wind turbines are the largest flexible rotating machines in the world, and they must run for 20 years or more under constantly changing loads. As blades grow they flex into and out of the wind, and the rotor meets its own shed vorticity, which calls design assumptions into question. Aeroelastic stability has rarely driven blade design so far, but that may change for future large, flexible rotors [4]. The chapter's rigid rotor, without tower or drivetrain modes, cannot show this. Progress would be pitch controllers designed and verified with the flexible modes in the loop.

The power coefficient surface is an approximation with three stated departures from the textbook coefficient set, and it should not be presented to readers as the standard correlation. Its 2.1 % error at the rated point is small enough that the rating lands where a reader expects and large enough to be worth stating.

The app's Learn texts state the same 0.1-degree exit threshold, 2 s Region 3 exit dwell and 6 s Region 2/2.5 dwell as the controller, and the derated-torque scenario distinguishes operation at the rated bound from the controller saturation flag.

The MPC's model keeps the sensitivities of the moment over its horizon, although a gust from 9 to 15 m/s more than triples $|\partial T_{aero}/\partial\beta|$, omits the rate limits within a step, and inherits every omission of Section 4.3, so its tighter speed regulation is not shown to be acceptable to the tower or the drivetrain. Its LIDAR is a declared instrument, handed the coming turbulence with an error that vanishes at the rotor, which is kinder than a real nacelle LIDAR, and its turbulent results depend on that, as Section 4.7 shows. Its pitch move cap is conservative and is the whole cause of its larger overshoot after a Reset into AUTO in scenario 10, a fallback above rated in turbulence steps the torque command by up to 0.667 MN m, and no preview setting met the travel bound of the tuning pass. Most of all, it captures 2.6 to 5.2 % less energy than the loops over ten seeds of turbulence near rated, so it is a speed regulator and not an energy optimizer, and an economic objective that prices energy against speed error and actuator duty is the extension this comparison points to.

> **OPEN CHALLENGE** **Predictive Control on a Real Turbine**
>
> A model predictive controller can respect pitch limits, structural load limits and generator overspeed limits at once, and a lidar in the hub or on the nacelle can preview the wind and further reduce structural loads. Such controllers are computationally demanding, however, and by 2023 they had been demonstrated on real turbines only in very limited experiments [9]. The chapter's predictive controller plans on a declared model with a declared lidar. Progress would be field tests that measure its energy, pitch travel and loads against the scheduled loops on one machine.

The scenario tree plans on three winds one declared standard deviation apart with equal weights, which carry two thirds of the declared variance and nothing about a gust's timing or size, and it prices its soft limits by their expected value rather than holding them on every branch. Its thrust advantage is measured on a momentum-theory estimate of a machine without a tower, its cheapest rival is a stronger hedge that gives up a further 3.2 % of the energy, and each of its plans costs about fourteen times the one-wind plan's solve.

Finally, the loop is given a measurement no turbine has, as Section 4.7 describes, and rotor speed is read with no noise, no lag and no quantization. Wind is the only source of randomness in the plant, and the LIDAR preview of Section 4.5.6 draws its error from a separate seeded stream.

4.10 Conclusion

This chapter studied a wind turbine control simulator built around the handover between two actuators that reach the same measurement. The model is five states, a rotor speed on a lumped 4×10^7 kg m^2 inertia, two rate-limited actuator states and two wind states, driven by a fitted power coefficient surface whose peak is placed at $\lambda^* = 7.55$ with $C_{P,max} = 0.482$ and whose three departures from the textbook coefficient set are stated. Below rated a proportional-integral loop on generator torque holds the rotor at $\lambda^* V/R$ and captures 0.482 of the available power, and detuning the target tip-speed ratio to 5.5 costs 13.0 % of the electrical output at 6 m/s. Above rated the torque is frozen at 4.18 MN m and collective pitch regulates speed, with the pitch onset assisted by a 0.3 % setpoint offset and rated torque selected explicitly by the region index. The wind ramp crosses both boundaries with a terminal speed error of 0.04 rpm against the logged setpoint, and actual actuator motion remains continuous through its physical dynamics. The pitch sensitivity of aerodynamic torque grows by a factor of 4.9 across Region 3, and switching the schedule off improves IAE from 3.058 to 0.865 while raising blade travel from 108.4° to 255.2° and putting the actuator on its rate limit, so the fixed high gain buys speed regulation at the price of pitch travel. Margin against the omitted structural modes is an engineering concern beyond these simulations.

The model holds air density and inertia constant, and it contains no drivetrain torsion, no tower motion and no blade flexibility, so the loads that set the real bandwidth limit on a pitch loop are outside its reach. Adding a two-mass drivetrain, a fore-aft tower mode, individual pitch and a rotor-effective wind estimator would each bring the simulator closer to a commercial controller. Those additions would change what the application can demonstrate, not what it teaches. The predictive controller added as an option shows that a five-second preview supplies about two fifths of its gust advantage, that the rest is a tighter speed objective, which the loops match on that gust with the schedule off but not in turbulence, where the same objective costs 2.6 to 5.2 % of the energy over ten seeds. Its scenario-tree setting lowers a turbulent gust's peak by 0.10 points and the estimated thrust by 2.7 % for 0.6 % of the energy, because with one model in every branch its shared move stays close to the mean of three one-wind moves, whose outer branches nearly cancel. The wind turbine simulator puts operating regions, sequenced actuators and gain scheduling on a machine whose nonlinearity can be measured.

KEY POINTS

- A large wind turbine holds its rotor speed with two controls, the generator's torque in light wind and the angle of the blades in strong wind.
- A 0.3 % setpoint offset lets the pitch loop begin regulating before Region 3 is declared, and the region index then selects rated torque.
- The pitch sensitivity grows by a factor of 4.9 across Region 3, and switching the schedule off buys speed regulation at more than twice the blade travel.
- A five-second wind preview supplies about two fifths of the predictive controller's gust advantage, and in turbulence the controller captures 2.6 to 5.2 % less energy than the loops.

Where else this applies. A gain scheduled on a slowly moving variable also keeps the loops of Volume 7, Chapter 9 and Volume 4, Chapter 3 at their design point, and Volume 4, Chapter 10 shares one variable among three actuators.

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

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

[4] P. Veers, K. Dykes, E. Lantz, S. Barth, C. L. Bottasso, O. Carlson, A. Clifton, J. Green, P. Green, H. Holttinen, D. Laird, V. Lehtomäki, J. K. Lundquist, J. Manwell, M. Marquis, C. Meneveau, P. Moriarty, X. Munduate, M. Muskulus, J. Naughton, L. Pao, J. Paquette, J. Peinke, A. Robertson, J. Sanz Rodrigo, A. M. Sempreviva, J. C. Smith, A. Tuohy, and R. Wiser, "Grand challenges in the science of wind energy," *Science*, vol. 366, no. 6464, p. eaau2027, 2019.

[5] W. E. Leithead, S. de la Salle, and D. Reardon, "Role and objectives of control for wind turbines," *IEE Proceedings C Generation, Transmission and Distribution*, vol. 138, no. 2, pp. 135–148, 1991.

[6] E. Muljadi and C. P. Butterfield, "Pitch-controlled variable-speed wind turbine generation," *IEEE Transactions on Industry Applications*, vol. 37, no. 1, pp. 240–246, 2001.

[7] ——, "Pitch-controlled variable-speed wind turbine generation," National Renewable Energy Laboratory, Golden, CO, Conference Paper NREL/CP-500-27143, 2000, presented at the 1999 IEEE Industry Applications Society Annual Meeting. OSTI 753780.

[8] A. D. Wright and L. J. Fingersh, "Advanced control design for wind turbines; part I: Control design, implementation, and initial tests," National Renewable Energy Laboratory, Tech. Rep. NREL/TP-500-42437, 2008.

[9] P. Veers, C. L. Bottasso, L. Manuel, J. Naughton, L. Pao, J. Paquette, A. Robertson, M. Robinson, S. Ananthan, T. Barlas *et al.*, "Grand challenges in the design, manufacture, and operation of future wind turbine systems," *Wind Energy Science*, vol. 8, no. 7, pp. 1071–1131, 2023.

[10] K. E. Johnson, L. Y. Pao, M. J. Balas, and L. J. Fingersh, "Control of variable-speed wind turbines: standard and adaptive techniques for maximizing energy capture," *IEEE Control Systems*, vol. 26, no. 3, pp. 70–81, 2006.

[11] E. A. Bossanyi, “The design of closed loop controllers for wind turbines,” *Wind Energy*, vol. 3, no. 3, pp. 149–163, 2000.

[12] F. D. Bianchi, R. J. Mantz, and H. De Battista, *Wind Turbine Control Systems*, ser. Advances in Industrial Control. Springer London, 2007.

[13] L. Y. Pao and K. E. Johnson, “Control of wind turbines,” *IEEE Control Systems*, vol. 31, no. 2, pp. 44–62, 2011.

[14] J. Jonkman, S. Butterfield, W. Musial, and G. Scott, “Definition of a 5-MW reference wind turbine for offshore system development,” National Renewable Energy Laboratory, Tech. Rep. NREL/TP-500-38060, 2009.

[15] A. Betz, “The maximum of the theoretically possible exploitation of wind by means of a wind motor,” *Wind Engineering*, vol. 37, no. 4, pp. 441–446, 2013.

[16] J. G. Slootweg, S. W. H. de Haan, H. Polinder, and W. L. Kling, “General model for representing variable speed wind turbines in power system dynamics simulations,” *IEEE Transactions on Power Systems*, vol. 18, no. 1, pp. 144–151, 2003.

[17] B. Boukhezzar and H. Siguerdidjane, “Nonlinear control of a variable-speed wind turbine using a two-mass model,” *IEEE Transactions on Energy Conversion*, vol. 26, no. 1, pp. 149–162, 2011.

[18] W. J. Rugh and J. S. Shamma, “Research on gain scheduling,” *Automatica*, vol. 36, no. 10, pp. 1401–1425, 2000.

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

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

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

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

[23] D. Schlipf, D. J. Schlipf, and M. Kühn, “Nonlinear model predictive control of wind turbines using LIDAR,” *Wind Energy*, vol. 16, no. 7, pp. 1107–1129, 2013.

[24] P. O. M. Scokaert and D. Q. Mayne, “Min-max feedback model predictive control for constrained linear systems,” *IEEE Transactions on Automatic Control*, vol. 43, no. 8, pp. 1136–1142, 1998.

[25] S. Lucia, T. Finkler, and S. Engell, “Multi-stage nonlinear model predictive control applied to a semi-batch polymerization reactor under uncertainty,” *Journal of Process Control*, vol. 23, no. 9, pp. 1306–1319, 2013.

[26] D. Bernardini and A. Bemporad, “Stabilizing model predictive control of stochastic constrained linear systems,” *IEEE Transactions on Automatic Control*, vol. 57, no. 6, pp. 1468–1480, 2012.

[27] T. R. Kuphaldt, *Lessons in Industrial Instrumentation*, version 2.33 ed. Open textbook, Creative Commons Attribution 4.0 license, ibiblio.org/kuphaldt/socratic/sinst, 2024.

[28] M. Stammler, O. Menck, Y. Guo, and J. Keller, “Wind turbine design guideline DG03: Yaw and pitch bearings,” National Renewable Energy Laboratory, Golden, CO, Tech. Rep. NREL/TP-5000-89161, 2024.

[29] 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.

4.11 Exercises

4.1 ○ Concept check · 2 min

In the shipped loops the torque loop’s speed setpoint is capped at $0.997\,\omega_{rated}$, while the pitch loop regulates to ω_{rated} itself. What does this 0.3 % offset do?

(a) It leaves the overspeed alarm a margin, so the alarm does not latch during the handover

(b) It makes up for the generator efficiency, so the electrical power settles at exactly its rating

(c) It keeps the logged speed error inside the scorecard band while the wind sits near rated

(d) It keeps the two integrators off one zero-error point, so pitch can start before Region 3

4.2 Concept check · 2 min

Scenario 7 is run for 120 s with the pitch gain schedule off and then, after Reset, with it on, on the same seeded wind. Off: speed IAE 0.865 rad, blade travel 255.2°, peak pitch rate 8.0 °/s. On: 3.058 rad, 108.4° and 5.2 °/s. Which conclusion does the chapter support?

(a) The fixed gain buys tighter speed with more than twice the blade travel and a pitch rate on its limit

(b) The fixed gain is the better tuning above rated, because a lower speed IAE means less fatigue for the rotor

(c) The scheduled loop is close to instability at high wind, which is why its speed IAE is three times higher

(d) The two runs cannot be compared, because the turbulence differs between the first run and the Reset

4.3 Concept check · 1 min

The advanced Region 2 option replaces the proportional-integral law of Equation (4.12) with $T_{gen} = K\omega^2$. What does that law do without?

(a) A rotor speed measurement, because the constant K already holds the optimal ratio

(b) A wind measurement and a speed setpoint, because the aerodynamics settle the rotor at λ^*

(c) Slow convergence to the optimum, because it skips the integral action of the proportional-integral law

(d) A correct constant, because the rotor finds its optimum by itself as the blades erode over time

4.4 Theory and calculation · 6 min

The default machine has $\rho = 1.225\,\mathrm{kg/m^3}$, $R = 63\,\mathrm{m}$, $\lambda^* = 7.55$, $C_{P,max} = 0.482$, $\eta = 0.944$, $b = 1.0 \times 10^4\,\mathrm{N\,m\,s}$ and $\omega_{rated} = 1.267\,109\,\mathrm{rad/s}$. The wind is a steady 6 m/s.

(a) Find the Region 2 speed setpoint of Equation (4.11) in rad/s and in rpm. *Checkpoint:* 0.719 rad/s, 6.87 rpm.

(b) Find the torque-law constant K, the aerodynamic torque $K\omega^2$ at the setpoint, and the electrical power $\eta\,(K\omega^2 - b\,\omega)\,\omega$. *Checkpoint:* $2.139 \times 10^6\,\mathrm{N\,m\,s^2}$, 1.106 MN m, 0.746 MW.

(c) Find the wind speed at which the setpoint reaches its cap $0.997\,\omega_{rated}$. *Checkpoint:* 10.54 m/s.

4.5 Theory and calculation · 8 min

The pitch loop sees the integrating plant $J_{eq}\dot{\omega} = (\partial T_{aero}/\partial\beta)\,\beta$ with $J_{eq} = 4.0 \times 10^7\,\mathrm{kg\,m^2}$. The sensitivity $|\partial T_{aero}/\partial\beta|$ is 8.915 MN m/rad at fine pitch and 46.61 MN m/rad at 25 m/s, where the steady pitch is 21.97°. The schedule knee is $\beta_k = 4.98°$ and the shipped tuning is $K_{c,\beta} = 3.7673\,\mathrm{s}$, $\tau_{I,\beta} = 2.3333\,\mathrm{s}$.

(a) Place the loop at $\zeta = 0.7$ and $\omega_n = 0.6\,\mathrm{rad/s}$ at fine pitch: find $K_{c,\beta}$ and $\tau_{I,\beta}$. *Checkpoint:* 3.769 s and 2.333 s.

(b) With the shipped gain, find the loop gain product $K_{c,\beta}|\partial T_{aero}/\partial\beta|$, in MN m s/rad, at fine pitch and at 25 m/s, then the divisor GK of Equation (4.8) at 21.97° and the scheduled product there. *Checkpoint:* 33.6 and 175.6, $GK = 5.41$, 32.4.

(c) For the unscheduled loop at 25 m/s, find the natural frequency, in rad/s and Hz, and the damping ratio. *Checkpoint:* 1.37 rad/s or 0.218 Hz, $\zeta = 1.60$.

4.6 Theory and calculation · 6 min

The rotor of the default machine has $J_{eq} = 4.0 \times 10^7\,\mathrm{kg\,m^2}$, $\omega_{rated} = 1.267\,109\,\mathrm{rad/s}$, a rating of 5.0 MW and $b = 1.0 \times 10^4\,\mathrm{N\,m\,s}$.

(a) Find the kinetic energy at rated speed and express it in seconds of rated power. *Checkpoint:* 32.1 MJ, 6.42 s.

(b) The aerodynamic damping at rated speed and rated wind is −2.65 MN m s. Find the open-loop time constant. *Checkpoint:* 15.0 s.

(c) Scenario 8 doubles the inertia. With the shipped pitch tuning ($K_{c,\beta} = 3.7673$ s, $\tau_{I,\beta} = 2.3333$ s, $k = 8.915$ MN m/rad), find the natural frequency and damping ratio of the pitch loop, and the gain that restores the design, in seconds and as the app displays it in °/rpm, which is six times the value in seconds. *Checkpoint:* 0.424 rad/s, $\zeta = 0.495$; 7.535 s or 45.2 °/rpm.

4.7 App challenge · 6 min

Load scenario 1, Steady 6 m/s (Region 2), with the default seed. Run 60 s at 5× and export the CSV file. Then press the reload button, set Target tip-speed ratio to 5.50 before Run, run 240 s at 5× and export again. Read the last row of each file.

(a) Read `generator_power_W` at 60 s of the first run. *Checkpoint:* 0.746 MW, within 0.002 MW.

(b) Read `tip_speed_ratio`, `Cp` and `generator_power_W` at 240 s of the second run. *Checkpoint:* 5.50, 0.418 and 0.649 MW, within 0.002 MW.

(c) Express the power lost by the detuned ratio as a percentage of the power at $\lambda^* = 7.55$. *Checkpoint:* 13.0 %, within 0.3 points.

4.8 App challenge · 8 min

Load scenario 7, Gain schedule off, high wind, with the default seed. Run 120 s at 10× with Pitch gain schedule Off as loaded. Read the speed IAE on the scorecard and the Peak pitch rate tile, and export the CSV file. Then press Reset, switch Pitch gain schedule to On before Run, change nothing else, run 120 s again and repeat the readings.

(a) Give the speed IAE and the peak pitch rate of each run. *Checkpoint:* Off: 0.865 rad and 8.0 °/s. On: 3.058 rad and 5.2 °/s. IAE within 0.005 rad.

(b) From each file, sum the absolute changes of `pitch_deg` from row to row to get the blade travel, and find the ratio of the two. *Checkpoint:* 255.2° and 108.4°, within 0.5°, a ratio of 2.35.

(c) Read `ctrl_gain_divisor` and `ctrl_tauI_pitch_eff_s` in the last row of the second file. *Checkpoint:* $GK = 5.27$, within 0.01, and 2.3333 s.

4.9 App challenge · 8 min

Load scenario 11, MPC or PI loops in a gust, with the default seed. The mean wind steps from 14 to 20 m/s at 30 s and back at 45 s. Run each case to 60 s at 2×, export the CSV file, skip its comment lines that start with #, and find the largest value of `overspeed_pct`. The Peak speed tile rounds to a whole percent.

(a) Run the scenario as loaded, under the MPC with the LIDAR preview. *Checkpoint:* 100.29 %, within 0.02 points.

(b) Press Reset, switch Controller to PI loops before Run, and run again. *Checkpoint:* 105.59 %, within 0.02 points.

(c) Press the reload button, set MPC wind preview to Off before Run, and run again. Then find the share of the gain of (a) over (b) that the preview supplies. *Checkpoint:* 102.39 %, within 0.02 points, and a share of about 40 %.

4.10 Application area · 3 min

The first tower fore-aft mode of the reference machine is at 0.32 Hz and the first drivetrain torsional mode near 1.7 Hz. The scheduled pitch loop is placed at $\omega_n = 0.6\,\mathrm{rad/s}$, and the unscheduled loop at the top of Region 3 moves to about 0.21 Hz.

(a) Find the separation between the tower mode and the loop, as a ratio of frequencies, for the scheduled and for the unscheduled loop. *Checkpoint:* 3.4 and 1.5.

(b) Say why the app cannot show what the smaller separation costs, and what it measures instead. *Checkpoint:* No tower or drivetrain states. It measures blade travel and time on the rate limit.

4.11 Application area · 4 min

The Region 2 law of Equation (4.11) computes its setpoint from the wind speed, and the app hands it the true wind at the rotor. The reference designs use $T_{gen} = K\omega^2$ with $K = \rho\pi R^5 C_{P,max}/(2\lambda^{*3})$ instead.

(a) Give two reasons a nacelle anemometer cannot supply the wind that the setpoint needs. *Checkpoint:* It sits in the rotor's wake, and it reads one point, not the average over the disc.

(b) Over several years the blades erode, $C_{P,max}$ falls by 4 % and λ^* rises from 7.55 to 7.8. By what factor does the correct K change, and does a $K\omega^2$ law that keeps the commissioned K run the rotor above or below its new optimal ratio? *Checkpoint:* A factor of 0.871, about 13 % lower. Below.

4.12 Review the agent · 5 min

An AI assistant reviews the chapter's scenario 11 and writes: "The MPC holds the gust peak to 100.29 % of rated speed against the PI loops' 105.59 %, so the MPC is the better controller for this turbine and should replace the loops." Find the flaw in the reasoning and say how you would check the claim in the app.