INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 10 Control in Upstream and Midstream Oil and Gas

CHAPTER 2

A Managed Pressure Drilling Simulator to Teach Optimization on a Constraint Boundary

Pages 24–54

Edited by John D. Hedengren

How to cite this chapter. "A Managed Pressure Drilling Simulator to Teach Optimization on a Constraint Boundary," in *Control in Upstream and Midstream Oil and Gas*, John D. Hedengren, Ed., vol. 10 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 2, pp. 24–54.

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

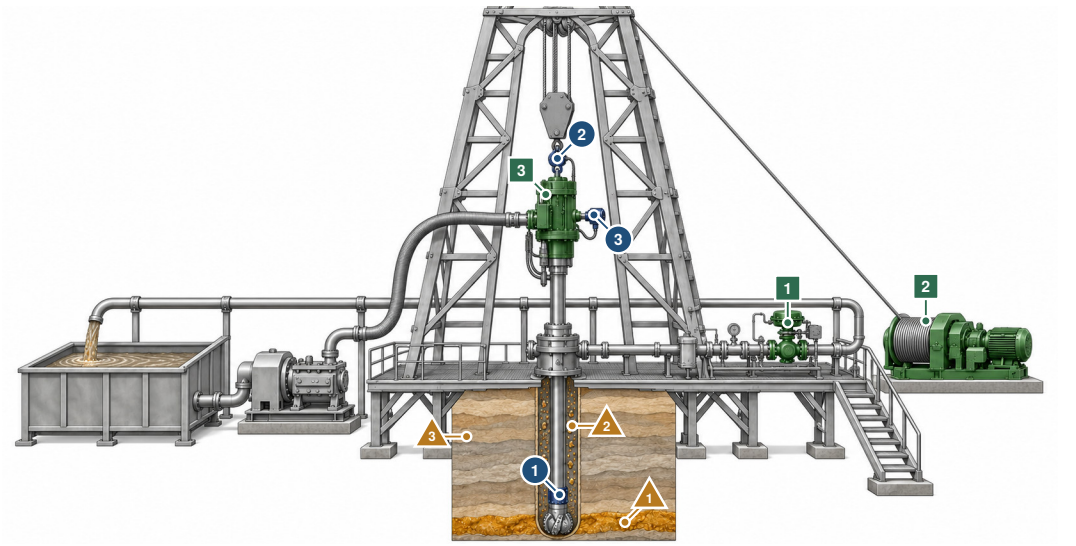
The app for this chapter
https://apmonitor.com/apps/drilling-mpd/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v10-oil-gas/#drilling-mpd

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

CHAPTER 2

A Managed Pressure Drilling Simulator to Teach Optimization on a Constraint Boundary



The controller moves 1 choke on the returning mud 2 drawworks that pays out the drilling line 3 top drive that turns the pipe

The controller reads 1 pressure at the bottom of the hole 2 hook load, which gives the weight on the bit 3 speed of the top drive

What disturbs it 1 a harder, more abrasive rock layer 2 rock cuttings carried up in the mud 3 pore and fracture pressure of the rock

The mud that lifts the cuttings up the annulus leaves the wellhead through the choke on the returns skid, so the cuttings and the choke both add to the pressure at the bit. Illustration, not to scale.

IN PLAIN TERMS **Drilling Fast Right at the Edge**

The main idea. A drill bit cuts rock faster when it is pushed harder, turned faster and pressed on less by the mud. A valve on the returning mud lets the crew lower that mud pressure.

The challenge. Every push toward speed brings the well closer to a limit: the pipe twists and jerks, rock chips pile up, or the mud pressure leaves the narrow range the rock allows. The best speed sits right on the edge of these limits.

Why it matters. In the chapter's simulator, a controller working along that edge drills 73 percent more hole than a cautious fixed setting.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the Did you know? box in the Introduction, the two History boxes in Section 2.2, the Try it boxes and the second Did you know? box in Section 2.6, and Key points in Section 2.9. No equations.
Study	70 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	25 min	This page, Section 2.6 with Figure 2.3, Section 2.7 and Section 2.8 with its Open challenge boxes.

The control question Weight on bit, rotary speed and a lower bottomhole pressure all raise the drilling rate, and each also pushes the well toward a vibration, cuttings or pressure limit. Where does the best feasible operating point sit, what holds it there, and what does operating away from it cost?

Learning objectives

- Explain why the best drilling rate lies on the boundary of the feasible region.
- Derive the cuttings ceiling and use it to pick the handle that raises the achievable rate.
- Compute the choke gain and authority and the ramp offset of the pressure loop.
- Locate the stick-slip corner and the bit-bounce band from G and the axial resonance.
- Separate a conservative limit from a detuned loop in the graded challenge.

Prerequisites Sections 1.4.2 (pressure column and choke), 1.4.8 (constraint control), 1.6.2 (velocity-form PID) and 1.6.4 (the apps). Section 1.6.3 is needed only for the choke MPC.

The app https://apmonitor.com/apps/drilling-mpd/

See also Volume 9, Chapter 8 and Volume 9, Chapter 10 for supervisory setpoints, Volume 9, Chapter 4 for an offset-free MPC, Volume 4, Chapter 7 for rate limits on a utility drill and Volume 4, Chapter 2 for a protective override.

2.1 Introduction

The rate of penetration in rotary drilling is limited by several mechanisms at once. Weight on bit raises the depth of cut until the cutters can no longer clear their own debris [1]. Rotary speed raises the rate sub-linearly, and weight and speed together determine whether the drillstring turns smoothly or enters a torsional limit cycle [2, 3]. The mud column that keeps the formation in the hole also presses each freshly sheared cutting back onto the rock, so overbalance reduces the drilling rate [4]. Every cubic meter of rock removed must be lifted up the annulus, and the mixture of rock and mud is heavier and rougher than mud alone, so drilling faster raises the pressure at the bottom of the hole and slows the drilling down again [5].

The objective and every constraint on it increase together with the two variables the operator manipulates, so the best operating point is not an interior maximum. It is a point on the boundary of a feasible region. At the default settings of the app of this chapter the objective peaks at 111.2 ft/hr with 35.0 klbf of weight and 250 RPM, and that peak is infeasible, because it loads the annulus to 2.29 % cuttings by volume against a limit of 2.2 %. The best point that satisfies every limit produces 106.7 ft/hr and the operating point supplied by the first scenario produces 67.6 ft/hr. The constraint costs 4 % of drilling rate and the conservative operating point costs 37 %, a factor of nine between the two losses.

The feasible region also changes during operation. Bit wear accumulates over hours and lowers the whole objective surface, so a limit that binds at the start of a bit run no longer binds at the end of it, and cuttings loading has a time constant near 340 s, two orders of magnitude slower than the torsional mode that governs stick-slip. An operating point selected from a static evaluation becomes suboptimal within the hour.

> ✦ DID YOU KNOW? **A Deep Well Takes Months**
>
> A gas exploration well finished in the United States in 1974 reached 9583 m. Drilling it took 504 days and 103 bits [6]. In 2016 a well of 6000 m took about three to six months to drill, and one of 8000 m more than a year [6].

The model rests on standard conservation statements, a momentum balance on two torsional inertias, a mass balance on the cuttings inventory in the annulus, and hydrostatic and frictional contributions to bottomhole pressure [7, 8]. The experiments built on the model show integrating dynamics, negative process gain, time-scale separation, anti-windup and supervisory constraint

control, the principles a process dynamics and control course develops with simpler plants [7, 9]. The behavior matters industrially because non-productive time is the largest addressable cost in a well and constraint-aware operation is the mechanism that reduces it [10].

This chapter works with a browser simulator of drilling and managed pressure, set out completely enough to be rebuilt from the text, and locates where its constrained optimum sits and what holds it there. The app gives the cost of a conservative limit as a measured number rather than an assertion. Section 2.2 reviews how the loops the app implements developed. Section 2.3 gives the thirteen-state model, whose steady-state evaluator and dynamic integration are required to agree. Section 2.4 maps the feasible region and derives a closed-form ceiling on the objective imposed by the cuttings constraint, $ROP \leq C_{max} Q \eta(\omega) / A_{bit}$. The ceiling depends on rotary speed alone and is reached at 106.7 ft/hr for every choke position above 52.1 % open, so two of the three manipulated variables contribute nothing beyond that point, while a 3.3 % increase in pump rate adds 4.2 % of feasible drilling rate. Section 2.5 describes the regulatory loops, the constraint-aware search transcribed from the executing code and the optional choke MPC. Section 2.6 reports the scenarios, the graded challenge and the comparison of pressure controllers. Section 2.7 sets the hardware of a real rig beside the model, Section 2.8 states the limitations, and Section 2.9 draws the results together.

2.2 Background

The first drilling automation was mechanical and it controlled exactly one variable. That device, the automatic driller, held weight on bit by paying out drilling line, and the History box below traces it to its first patents. Its structure survives unchanged in the app of this chapter. Weight is not something the rig can set directly, because the string is an elastic member and the bit is destroying the rock the weight rests on, so the manipulated variable is the rate at which line is paid out and the plant between it and the controlled variable is an integrator.

> ⧖ HISTORY **1927 and 1952: Machines That Hold the Weight on the Bit**
>
> A patent application of 1927 described a dummy driller, a device to keep a set force on a rotary drill hung in part from a cable [11]. A reversible motor paid the cable out or picked it up, and too much pull on the cable made the motor pay out until the strain fell. The hydraulic apparatus that already showed the strain on such cables was to work the automatic control as well [11]. A 1952 application, granted in 1956, let the tension in the cable regulate how fast the drawworks drum unspooled [12]. Its hydraulic pressure also read on a gauge calibrated as a weight indicator for the driller. A fall in that pressure meant too much weight on the bit, so the brake slowed the drum, the line paid out more slowly and the bit drilled off [12].

Everything above that loop arrived later as separate developments. Rate-of-penetration modeling became quantitative when the drilling variables were fitted as a multiple regression on offset-well data, which produced the multiplicative form still used today [4], and a mechanistic account for roller-cone bits, in which the rate is limited alternately by rock strength and by the removal of cuttings from beneath the bit, gave the founder behavior a physical basis [1]. Specific energy, the work spent per unit volume of rock destroyed, was proposed as the efficiency measure for rock drilling in 1965 [13] and became the standard real-time surveillance quantity on the rig floor forty years later, where it detects the moment at which extra weight stops making hole and starts grinding and decides when to trip a dull bit [14]. Volume 4, Chapter 7 uses specific energy in the same way, to read the soil a utility drill is cutting.

Drillstring vibration was studied on its own track. Torsional oscillation with the bit stopping completely and then releasing is a self-excited limit cycle driven by a bit torque that falls with speed, and active damping of that mode by feeding measured top-drive speed back through a compliant torque command was demonstrated on full-scale strings [2]. Later work reduced the mechanism to a bit-rock interaction law with a delay set by the depth of cut per revolution [3], and top-drive controllers that absorb torsional waves rather than reflecting them are now a standard option [15].

Hole cleaning followed a third track, with cuttings-transport models built from flow-loop experiments and organized around the annular velocity needed to keep a bed from forming [5].

Managed pressure drilling is the newest of the four developments and the one that made annular pressure a controlled variable. The second History box below follows it from mud density to an automatic choke. The vocabulary and the systems view that holds these pieces together came from a deliberate effort to import automation practice from aviation and process manufacturing into drilling [16], and the case for doing so rests on non-productive time, which is the fraction of rig days spent not making hole and is the largest addressable cost in a well [10]. The application implements the state that arrangement has reached, which is regulatory loops on weight, speed and pressure with a slow supervisory layer above them, rather than a single optimizer replacing them.

> ⧖ HISTORY **2007: The Choke Takes Over Well Pressure**
>
> Pressure at the bottom of a well was long set by the density of the mud, a method that is slow and inaccurate and cannot answer the brief swings of a pipe connection [17]. Managed pressure drilling grew out of underbalanced drilling. An industry committee defines it as an adaptive drilling process that precisely controls the annular pressure profile throughout the wellbore [18]. A rotating control device seals the top of the annulus and sends the returning mud through a choke, whose opening sets the back-pressure [17]. Narrow margins, especially in depleted reservoirs, call for accurate control, and in 2007 automatic control of downhole pressure by a topside choke was applied offshore in the North Sea [19]. Later work put the rate of penetration and the pressure into one controller that moves pump flow, choke, rotation and weight on bit together [20].

2.3 Process Model

The application simulates one vertical hole section drilled by one bit through one formation. Mud is pumped down the drill pipe, passes through the bit, and returns up the annulus between the pipe and the hole wall, through a choke that adds back-pressure and on to the pits. The layout, the measurement points and the symbol for every located quantity are given in Figure 2.1. Everything in the model is SI. Newtons, pascals, radians per second and meters per second are what the equations use, and the field units quoted in this chapter (klbf, RPM, psi, pounds per gallon, feet per hour) are display conversions that change no computed value.

2.3.1 Variable Classification

Three controlled variables, three manipulated variables, the disturbances, the thirteen states and the derived quantities the constraints act on are collected in Table 2.1. The rate of penetration is the objective and is deliberately not a controlled variable, because its achievable value is a property of the whole feasible region rather than of any one loop.

2.3.2 The Objective

The rate of penetration is a product of five factors, one for each mechanism that can slow the bit down, in the reduced form of the multiple-regression law of [4] with the foundering behavior of [1], as shown in Equation 2.1.

$$ROP = K\, f_W(W)\, f_N(\omega)\, f_h(h)\, f_{\Delta P}(P_b - P_{pore})\, f_C(C) \tag{2.1}$$

where K is the drillability coefficient of the rock, divided by the formation strength multiplier $1 + \Delta_{hard}/100$, and each factor is dimensionless and equals unity at its own best condition. The product form matters more than any individual factor. A gain in one factor multiplies the whole objective, which is why the back-pressure choke changes the drilling rate although no rock property changes.

The weight factor rises from a threshold and then founders, as shown in Equation 2.2.

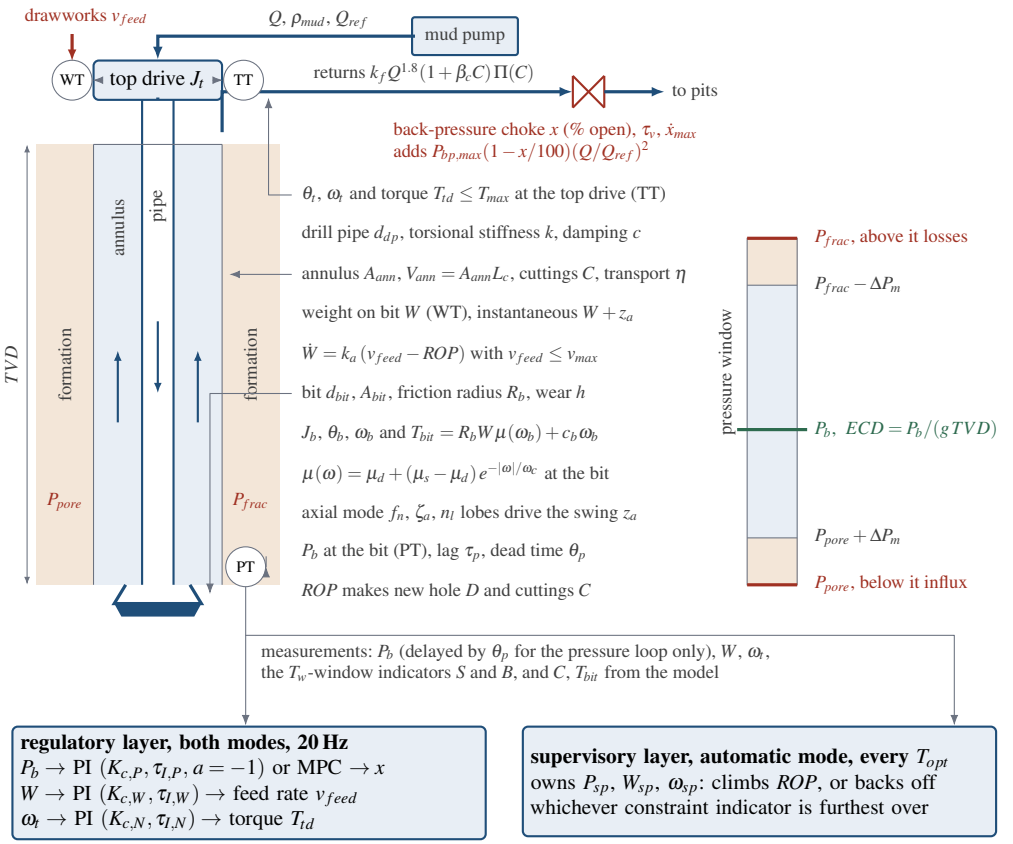


Figure 2.1 Well, surface circulating system and the two control layers. The weight and rotary loops run in both modes. In manual mode the operator sets the choke position and the three setpoints, and the pressure loop tracks the choke.

$$f_W = \chi^m \, \mathrm{e}^{m(1-\chi)}, \qquad \chi = \frac{W - W_t}{W_f - W_t} \tag{2.2}$$

where W_t is the weight below which the cutters do not engage and f_W is zero, W_f is the founder point and m sets how sharply the curve turns over. Differentiating gives $f'_W = m\chi^{m-1}\mathrm{e}^{m(1-\chi)}(1-\chi)$, which vanishes only at $\chi = 1$, so f_W peaks at exactly $W = W_f$ with the value 1 and the founder point is a declared parameter rather than an emergent property. The threshold also creates a region with no usable gradient, because f_W is identically zero below W_t and a gradient search that enters that region receives no direction, a consequence quantified in Section 2.8.

The remaining three factors are given by Equation 2.3.

$$f_N = (\omega/\omega_{ref})^{\gamma} \tag{2.3a}$$

$$f_h = \mathrm{e}^{-h/h_c} \tag{2.3b}$$

$$f_{\Delta P} = \exp\left[-\lambda \, \max(P_b - P_{pore}, 0)\right] \tag{2.3c}$$

$$f_C = \left[1 + (C/C_{50})^2\right]^{-1} \tag{2.3d}$$

Table 2.1 Every Signal in the Model, With Its Role and Range

Symbol	Role	Range (default)	Unit
P_b	CV, state	—	Pa
W	CV, state	≥ 0	N
ω_t	CV, state	—	rad/s
P_{sp}	setpoint	4200–6200 (5352)	psi
W_{sp}	setpoint	0–60 (25.0)	klbf
ω_{sp}	setpoint	0–250 (120)	RPM
x	MV, state	0–100 (65)	% open
v_{feed}	MV	0–0.030	m/s
T_{td}	MV	0–45	kN m
θ_t, θ_b	state, angles	—	rad
ω_b	state, bit speed	—	rad/s
$\bar{\omega}_b$	state, averaged bit speed	—	rad/s
z_a, $\dot{z}_a$	state, axial weight swing	—	N, N/s
C	state, cuttings fraction	0–0.30	—
h	state, bit wear	0–8	—
D	state, new hole	≥ 0	m
Δ_{hard}	DV, formation strength	-40 to $+80$ (0)	%
a_{abr}	DV, abrasiveness	0.4–2.5 (1)	$\times$
ΔP_{pore}	DV, pore shift	± 580 (0)	psi
ΔP_{frac}	DV, fracture shift	-870 to $+580$ (0)	psi
ROP	output, objective	—	m/s
S	output, stick-slip severity	—	—
B	output, bounce index	—	—
T_{bit}	output, bit torque	—	N m
ECD	output, circulating density	—	kg/m^3
MSE	output, specific energy	$\leq 5 \times 10^9$	Pa

where $\gamma < 1$ makes rotary speed sub-linear, h_c is the wear decay constant, λ is the chip hold-down coefficient and C_{50} is the cuttings concentration that halves the drilling rate. Equation 2.3c is the coupling that makes managed pressure drilling effective. Overbalance, the amount by which bottomhole pressure exceeds pore pressure, presses each cutting back onto the rock it was just sheared from, and the bit has to break it loose a second time. At $\lambda = 0.12/\text{MPa}$ the penalty is 7.9 % of drilling rate per 100 psi, which is the economic case for back-pressure control.

The rate of penetration is driven by an averaged bit speed rather than the instantaneous one, as shown in Equation 2.4.

$$\tau_R \frac{\mathrm{d}\bar{\omega}_b}{\mathrm{d}t} = \omega_b - \bar{\omega}_b \tag{2.4}$$

where $\tau_R = 2.0\,\text{s}$ is the depth-of-cut averaging time. Inside a stick-slip cycle the instantaneous bit speed runs from zero to several times surface speed within a second while rock removal integrates over many revolutions, so $\bar{\omega}_b$ is what appears in f_N, in the cuttings generation term and in the wear law, and it is what makes the dynamic model agree with the steady-state evaluator at an oscillating point.

2.3.3 Annular Hydraulics and the Choke

Bottomhole pressure is a hydrostatic head, a friction loss and a back-pressure, with a first-order lag between that sum and the pressure the bit actually sees, as shown in Equation 2.5.

$$\tau_p \frac{\mathrm{d}P_b}{\mathrm{d}t} = P_b^{ss}(C,x) - P_b \tag{2.5a}$$

$$\begin{aligned} P_b^{ss} = {} & \left[\rho_{mud} + C(\rho_c - \rho_{mud})\right] g\,TVD \\ & + k_f Q^{1.8}\,(1+\beta_c C)\,\Pi(C) \\ & + P_{bp,max}\left(1 - \frac{x}{100}\right)\left(\frac{Q}{Q_{ref}}\right)^2 \end{aligned} \tag{2.5b}$$

$$\Pi(C) = 1 + k_{pack}\,\max(C - C_{pack}, 0)/C_{pack} \tag{2.5c}$$

where ρ_c is the density of the cuttings, k_f is the annular friction coefficient, β_c is the extra friction each unit of cuttings loading contributes, Π is the pack-off multiplier and $P_{bp,max}$ is the back-pressure a fully shut choke produces at the sizing flow rate Q_{ref}. Three separate consequences follow from Equation 2.5b and each of them appears in a scenario. The choke position x is percent *open*, so closing it raises the pressure and the process gain of the pressure loop is $\partial P_b/\partial x = -40\,\mathrm{kPa}/\% = -5.80\,\mathrm{psi}/\%$ at the sizing flow. The back-pressure term carries $(Q/Q_{ref})^2$, so at the 174 gpm of the pack-off scenario the choke has only 78 psi of authority against 580 psi at full flow, and a reduced pump rate removes most of the choke authority. The cuttings appear twice, once as added weight and once as added friction, and above C_{pack} the multiplier Π climbs steeply, which is pack-off. Equivalent circulating density, $ECD = P_b/(g\,TVD)$, is reported because it compares directly against a pore or fracture gradient. Section 1.4.2 sets out the same column, friction and back-pressure for a producing well.

The choke itself is first order on its command with a stroking rate limit applied inside the controller, as shown in Equation 2.6.

$$\tau_v \frac{\mathrm{d}x}{\mathrm{d}t} = x_{cmd} - x, \qquad x \in [0, 100] \tag{2.6}$$

where τ_v is the stroke time constant and x is clamped to its travel after every integration step.

2.3.4 Cuttings in the Annulus

The annulus is treated as one well-mixed volume of length L_c carrying a single concentration, with generation set by the bit and removal by the flow, as shown in Equation 2.7.

$$V_{ann} \frac{\mathrm{d}C}{\mathrm{d}t} = ROP \cdot A_{bit} - C\,Q\,\eta(\bar{\omega}_b) \tag{2.7a}$$

$$\eta = \eta_{max} \frac{v_{ann}}{v_{ann} + v_0} \cdot \frac{1 + a_r\,\bar{\omega}_b/(\bar{\omega}_b + \omega_r)}{1 + a_r} \tag{2.7b}$$

where $A_{bit} = \pi d_{bit}^2/4$, $A_{ann} = \pi(d_{bit}^2 - d_{dp}^2)/4$, $V_{ann} = A_{ann}L_c$, $v_{ann} = Q/A_{ann}$ is the annular velocity, η is the transport efficiency, v_0 is the velocity at which half the ceiling efficiency is reached and a_r is the benefit a turning pipe gives by stirring the low side of the hole [5]. Both sides of Equation 2.7a are volumetric flow rates, so the equation is an inventory balance and $V_{ann}/(Q\eta)$ is its time constant, which is 340 s at the default flow and 120 RPM.

Equation 2.7a closes a loop back onto Equation 2.1 and that loop is what keeps the concentration finite. Drilling faster makes more cuttings, cuttings raise both the density and the friction in Equation 2.5b, higher bottomhole pressure means more overbalance, and chip hold-down in Equation 2.3c slows the drilling back down. In the pack-off scenario that negative feedback settles the concentration at 5.63 % rather than letting it run away, which is just above the pack-off threshold of 5.5 %.

2.3.5 Weight on Bit Dynamics

The drawworks pays line out at v_{feed} and the bit drills the hole away at ROP, and the string carries the difference as elastic strain, as shown in Equation 2.8.

$$\frac{\mathrm{d}W}{\mathrm{d}t} = k_a\,(v_{feed} - ROP), \qquad W \geq 0 \tag{2.8}$$

where k_a is the axial stiffness from the drawworks to the bit. Both terms are velocities and k_a is a force per unit length, so the right-hand side is a force per unit time. Equation 2.8 has no self-regulation and is the integrating process of the control literature [7]. A feed rate above the penetration rate raises the weight without bound and a feed rate below it drains the weight to zero, so the weight loop needs integral action to hold a setpoint and an autodriller manipulates a rate rather than a position.

2.3.6 Torsion and the Stick-Slip Indicator

The drillstring is two inertias joined by a twisting pipe, as shown in Equation 2.9.

$$J_t\,\dot{\omega}_t = T_{td} - k(\theta_t - \theta_b) - c(\omega_t - \omega_b) \tag{2.9a}$$

$$J_b\,\dot{\omega}_b = k(\theta_t - \theta_b) + c(\omega_t - \omega_b) - T_{bit} \tag{2.9b}$$

$$T_{bit} = R_b\,W\,\mu(\omega_b) + c_b\,\omega_b \tag{2.9c}$$

$$\mu(\omega) = \mu_d + (\mu_s - \mu_d)\,\mathrm{e}^{-|\omega|/\omega_c} \tag{2.9d}$$

with $\dot{\theta}_t = \omega_t$ and $\dot{\theta}_b = \omega_b$, where J_t and J_b are the top-drive and bit inertias, k and c are the torsional stiffness and damping of the pipe, R_b is the effective friction radius of the bit, c_b is viscous damping at the bit and ω_c is the velocity-weakening scale. Equation 2.9d contains the destabilizing mechanism. A bit that slows down meets a higher friction factor, which slows it further, so the bit torque has a negative slope with respect to bit speed wherever the exponential is still active. Setting that negative damping against the positive damping supplied by the pipe and the mud gives the dimensionless group of Equation 2.10.

$$G = \frac{R_b\,W\,(\mu_s - \mu_d)\,\mathrm{e}^{-\omega/\omega_c}/\omega_c}{c + c_b} \tag{2.10}$$

Both terms in Equation 2.10 are torques per unit of angular velocity, so G is dimensionless and $G = 1$ is the naive onset, the exact one if the top drive held its speed perfectly. The application uses $G_0 = 1.092$ as the boundary, the value at which the linearized string, two inertias, the twisting pipe and the speed loop of Section 2.5 at its default tuning, acquires a growing oscillatory mode. The onset is above unity because the speed loop lets the heavy top drive move and absorb part of the torsional energy, and it rises with the loop gain to 1.16 at four times the default. Near the onset the cycle grows slowly, so a run started at a steady state just above G_0 stays quiet for many minutes, and the instability is subcritical, so a cycle already running survives far below it, as Section 2.4 shows. The exponential is what gives the unstable region its shape. At 60 RPM it is 0.106 and at 250 RPM only 8.7×10^{-5}, so G grows with weight and collapses with rotary speed, and the unstable set is a corner at high weight and low rotary speed.

The severity itself is measured on the running simulation as the peak-to-peak bit speed over twice the mean surface speed within a moving window T_w. The map and the optimizer evaluate hundreds of candidate points per execution and cannot integrate the dynamic model at each one, so they use a correlation fitted to the same limit cycle, as shown in Equation 2.11.

$$\Lambda = \frac{R_b\,W\,(\mu_s - \mu_d)}{2\,\bar{\omega}_t\sqrt{kJ_b}} \tag{2.11a}$$

$$S = 1.343\,\Lambda^{0.712}\left[1 - \mathrm{e}^{-12(G - G_0)}\right] \quad \text{for } G > G_0 \tag{2.11b}$$

where Λ is the energy group obtained by winding the pipe up through the torque drop $R_b W(\mu_s - \mu_d)$ and releasing the stored energy into the bit inertia, and $S = 0$ below the onset. The map displays S as written, while feasibility is decided on the amplitude without the ramp factor, $1.343\,\Lambda^{0.712}$ for $G > G_0$. Equation 2.11 is a correlation and not a solution of Equations 2.9, and its accuracy is quantified in Section 2.4 rather than assumed.

2.3.7 The Axial Mode and the Bounce Index

The bottomhole surface develops a pattern of n_l lobes, so a bit turning at ω_b is excited at $n_l \omega_b$. The dynamic part of the weight on bit is the response of a damped second-order axial mode to that excitation, written in newtons so that the equation is the unit-mass normal form, as shown in Equation 2.12.

$$\ddot{z}_a = \omega_n^2 \left[F_0 \sin(n_l \theta_b) - z_a \right] - 2\zeta_a \omega_n \dot{z}_a \tag{2.12a}$$

$$F_0 = \varepsilon_0 \sqrt{W / W_{ref}}\, W \tag{2.12b}$$

where $\omega_n = 2\pi f_n$, ζ_a is the axial damping ratio, ε_0 is the lobe excitation fraction at the reference weight and W_{ref} is that reference. The instantaneous weight on bit is $W + z_a$. The steady amplitude of Equation 2.12 is $F_0 M(r)$ with the textbook magnification factor, which gives the closed-form bounce index of Equation 2.13.

$$B = \varepsilon_0 \sqrt{W / W_{ref}}\, M(r), \quad M(r) = \left[(1 - r^2)^2 + (2\zeta_a r)^2 \right]^{-1/2} \tag{2.13}$$

where $r = n_l \omega / \omega_n$ is the excitation frequency ratio. At $B = 1$ the instantaneous weight reaches zero and the bit leaves bottom. The resonance sits at $\omega = \omega_n / n_l$, which with $f_n = 5.00\,\text{Hz}$ and $n_l = 3$ is exactly 100 RPM, and the index there is 0.833 at 25 klbf against 0.216 at 120 RPM, 0.079 at 150 RPM and 0.019 at 250 RPM. The two vibration constraints differ in shape. Stick-slip occupies a corner because Equation 2.10 couples weight and speed through one exponential, while bit bounce occupies a horizontal band because a resonance depends on frequency and only weakly on weight. The response to stick-slip is to raise rotary speed, and the response to bounce is to move the speed out of the band in whichever direction is nearer.

Pause. Try Exercise 2.3 before reading on.

2.3.8 Bit Wear and Specific Energy

Wear accumulates as a power law in weight and speed, as shown in Equation 2.14.

$$\frac{\mathrm{d}h}{\mathrm{d}t} = \dot{h}_0\, a_{abr} \left(\frac{W}{W_{ref}} \right)^{1.2} \left(\frac{\bar{\omega}_b}{\omega_{ref}} \right)^{0.8}, \qquad h \in [0, 8] \tag{2.14}$$

where $\dot{h}_0$ is the wear rate at the reference weight and speed and a_{abr} is the abrasiveness of the formation. The derivative is forced to zero once h reaches 8. Mechanical specific energy, the quantity a driller uses to decide whether the bit is cutting or rubbing [13], is the axial work plus the rotary work per unit volume removed, as shown in Equation 2.15.

$$MSE = \frac{W}{A_{bit}} + \frac{T_{bit}\, \bar{\omega}_b}{A_{bit}\, ROP} \tag{2.15}$$

The second term is a power divided by a volumetric removal rate and both terms are pressures. It is capped at 5×10^9 Pa so that a bit turning and making no hole logs a large finite number rather than an infinity. MSE does not track the drilling rate, because at steady state stepping the weight from 25 klbf to the founder point at 35 klbf raises the rate of penetration from 67.6 ft/hr to 79.9 ft/hr and the specific energy from 28.5 kpsi to 32.8 kpsi at the same time.

Table 2.2 Model Constants, Their Defaults and Their Adjustable Ranges

Symbol	Meaning	Value	Range	Unit	Symbol	Meaning	Value	Range	Unit
Well and geometry					*Cuttings transport*				
TVD	true vertical depth	3000	1500–5000	m	η_{max}	transport efficiency ceiling	0.90	0.30–1.00	—
d_{bit}	bit diameter	0.2159	0.152–0.311	m	v_0	half-transport velocity	0.80	0.10–2.50	m/s
$d_{dp}^{\dagger}$	pipe outside diameter	0.1270	0.089–0.168	m	a_r	rotation cleaning benefit	0.60	0–2.0	—
L_c	cuttings transport length	200	60–900	m	ω_r	half-benefit rotary speed	8.0	1–25	rad/s
$g^{\dagger}$	gravity	9.81	9.7–9.9	m/s^2	C_{pack}	pack-off concentration	0.055	0.02–0.15	—
Mud and hydraulics					k_{pack}	pack-off friction gain	4.0	0–15	—
ρ_{mud}	mud density	1140	950–1800	kg/m^3	*Drillstring*				
$\rho_c^{\dagger}$	cuttings density	2500	2100–2900	kg/m^3	J_t	top drive inertia	2000	300–6000	kg m^2
Q	pump flow rate	0.030	0.008–0.055	m^3/s	J_b	bit and BHA inertia	120	30–500	kg m^2
$Q_{ref}^{\dagger}$	choke sizing flow	0.030	0.010–0.055	m^3/s	k	torsional stiffness	500	100–2500	N m/rad
k_f	annular friction coefficient	6.5	2–14	10^8 SI	c	string torsional damping	30	0–200	N m s/rad
β_c	cuttings friction loading	12	0–40	—	c_b	bit viscous damping	30	0–200	N m s/rad
$P_{bp,max}$	back-pressure, choke shut	4.0	1–9	MPa	R_b	bit friction radius	0.10	0.03–0.20	m
τ_p	bottomhole pressure lag	1.5	0.2–8.0	s	μ_s	static friction factor	0.80	0.30–1.40	—
θ_p	pressure dead time	0.0	0.0–4.0	s	μ_d	dynamic friction factor	0.25	0.05–0.70	—
τ_v	choke stroke constant	2.0	0.3–8.0	s	ω_c	velocity-weakening scale	2.80	0.50–10.0	rad/s
$\dot{x}_{max}$	choke rate limit	20	2–60	%/s	k_a	string axial stiffness	227	50–700	kN/m
Formation					$\tau_R^{\dagger}$	depth-of-cut averaging	2.0	0.2–10	s
P_{pore}	pore pressure	34.0	24–44	MPa	*Axial mode*				
P_{frac}	fracture pressure	41.5	28–55	MPa	f_n	axial natural frequency	5.00	2.0–9.0	Hz
Rate of penetration					ζ_a	axial damping ratio	0.060	0.015–0.30	—
K	drillability coefficient	0.011	0.002–0.030	m/s	ε_0	lobe excitation fraction	0.100	0.02–0.30	—
W_t	cutting threshold weight	22.2	0–60	kN	$W_{ref}^{\dagger}$	reference weight	111.2	40–220	kN
W_f	founder point	155.7	60–250	kN	n_l	bottomhole pattern lobes	3	2–5	—
m	founder sharpness	3	1.5–6	—	*Rig, detector and optimizer*				
$\omega_{ref}^{\dagger}$	reference rotary speed	12.566	5–25	rad/s	v_{max}	maximum feed rate	0.030	0.005–0.060	m/s
γ	rotary speed exponent	0.60	0.30–1.00	—	T_{max}	top drive torque limit	45	15–80	kN m
h_c	wear decay constant	6.0	2.0–20.0	—	$T_w^{\dagger}$	vibration detector window	20.0	4–60	s
λ	chip hold-down	0.12	0–0.40	1/MPa	T_{opt}	optimizer interval	2.0	0.5–10	s
C_{50}	half-rate cuttings loading	0.050	0.010–0.150	—	$\dot{W}_{sp}$	weight setpoint rate limit	1.5	0.1–8.0	kN/s
$\dot{h}_0$	wear rate at reference	0.50	0–3.0	10^{-3}/s	$\dot{\omega}_{sp}$	rotary setpoint rate limit	0.35	0.02–2.0	rad/s^2
					$\dot{P}_{sp}$	pressure setpoint rate limit	10.0	0.5–40	kPa/s

† not operator adjustable. k_f is in Pa·(s/m^3)$^{1.8}$.

2.3.9 Parameters, Initialization and Numerics

Table 2.2 collects every constant with its default and its adjustable range. Eight are not operator adjustable, because they are universal constants, internal reference values that only non-dimensionalize a ratio, or instrument configuration rather than plant properties.

Five further quantities are exposed as operator-adjustable limits rather than as parameters, because they are engineering choices and not properties of the well. They are the stick-slip limit $S_{max} = 0.40$ on 0.20 to 3.00, the bounce limit $B_{max} = 0.50$ on 0.10 to 1.50, the cuttings limit $C_{max} = 2.2\,\%$ by volume on 0.5 to 8, the bit torque limit $T_{lim} = 20\,\mathrm{kN\,m}$ on 4 to 55, and the pressure window margin $\Delta P_m = 0.7\,\mathrm{MPa}$ on 0.2 to 2.4 MPa, applied at both ends of the window. Only pack-off at C_{pack}, losses above P_{frac}, influx below P_{pore} and the rig limits v_{max} and T_{max} are physics. The distinction is the point of the graded challenge of Section 2.6.7.

Every run begins at a genuine steady state. The initializer takes the weight and rotary setpoints, takes the choke from the operator-set position in manual mode or from the position that puts P_b on setpoint in automatic mode, and then solves the implicit cuttings balance. Because ROP is strictly decreasing in C, both directly through Equation 2.3d and again through the head and friction the cuttings add to Equation 2.5b, the residual $ROP(C)A_{bit} - CQ\eta$ is strictly decreasing and has exactly one root, found by twenty bisections on $[0, 0.30]$. The pipe is initialized already wound up at $\theta_t = T_{bit}/k$, and the axial mode is placed on its steady orbit with $z_a = -F_0 M \sin\phi$ and $\dot{z}_a = F_0 M n_l \omega_b \cos\phi$ where $\phi = \mathrm{atan2}(2\zeta_a r,\ 1 - r^2)$, because at the default 120 RPM a state starting

at rest would ring at 1.5 times the eventual amplitude and the bounce indicator would report it.

The engine and its three clocks are those of Section 1.6.1. Integration is fixed-step fourth-order Runge-Kutta at $\Delta t = 0.005$ s, sixteen steps per cycle of the fastest thing in the model, which is the axial excitation at 12.5 Hz when the rotary speed is at the top of its range. Regulatory execution comes every 0.05 s, exactly ten physics steps, and supervisory execution every $T_{opt} = 2.0$ s, exactly forty regulatory executions. Logging runs at 20 samples per second, and the seed is 20260824, although the model carries no noise at all, so every run repeats exactly.

2.3.10 Assumptions

Each simplification below is stated with the place it fails.

One vertical section. There is no directional work, no sliding, no torque and drag along a curved wellbore and no drag from pipe lying on the low side of the hole. Those terms dominate the torsional problem in a deviated well, and a stick-slip severity computed here would not transfer to one.

Five factors in the drilling rate. A field rate-of-penetration law is regressed per field from offset wells and carries eight or more factors, including hydraulic horsepower at the bit and compaction with depth [4]. The values here are plausible for a soft-to-medium 8.5 in hole section and are not portable.

Two torsional degrees of freedom and one axial mode. A real bottomhole assembly has dozens of coupled modes. Lateral vibration, forward whirl and backward whirl are absent entirely, as is the axial-torsional coupling that makes real bit bounce and real stick-slip feed each other.

Bit bounce is detected and never fed back. The instantaneous weight fluctuation is computed, flagged and used as a constraint, but the bit is never taken off bottom. A real bouncing bit loses contact, so its average depth of cut and its own torque both fall, and it destroys cutters while doing so. Everything the application reports about bit bounce is therefore an indicator and not a consequence.

One well-mixed cuttings concentration. Real hole cleaning is a bed that forms on the low side and whose behavior depends on inclination more than on anything modeled here [5]. The lumped volume uses the residence time of the lower 200 m of annulus but applies the concentration to the whole column when computing *ECD*, which compresses a real bottoms-up time of half an hour into a transient short enough to observe in one session.

A linear choke. Back-pressure goes as $(1 - x/100)$ rather than as an orifice law $\Delta P \propto Q^2/x^2$, so the loop has one process gain to teach rather than three, and the very high gain a real choke shows near the shut end is absent.

Quasi-steady hydraulics. There is no surge and swab, no gas migration, no compressibility wave along the annulus, no cuttings slip velocity and no temperature effect on mud density. One first-order lag τ_p is the only dynamics between the choke position and the bottomhole pressure, and the optional dead time θ_p delays only the measurement the pressure loop reads.

Blind wear. Wear follows Equation 2.14 and nothing else. It does not accelerate with vibration, the bit never fails suddenly, balls up or loses a cone, and tripping for a new bit is free here while in a well it costs hours.

2.4 Feasible Region Analysis

Everything about a candidate operating point is available without integrating anything. Given a weight, a rotary speed and a choke position, the bisection described above returns the steady cuttings concentration, and from it the bottomhole pressure, the drilling rate, the bit torque, the two vibration indicators and the specific energy. That evaluator paints the map, the supervisory optimizer climbs it, and the verification suite compares the integrated model against it. Marching the dynamic model for 1300 s from a deliberately wrong start, no cuttings in the annulus and 60 % of the target weight, and then averaging over a further 60 s, reproduces the evaluator's drilling rate, pressure and concentration at four operating points to within 2.4 %.

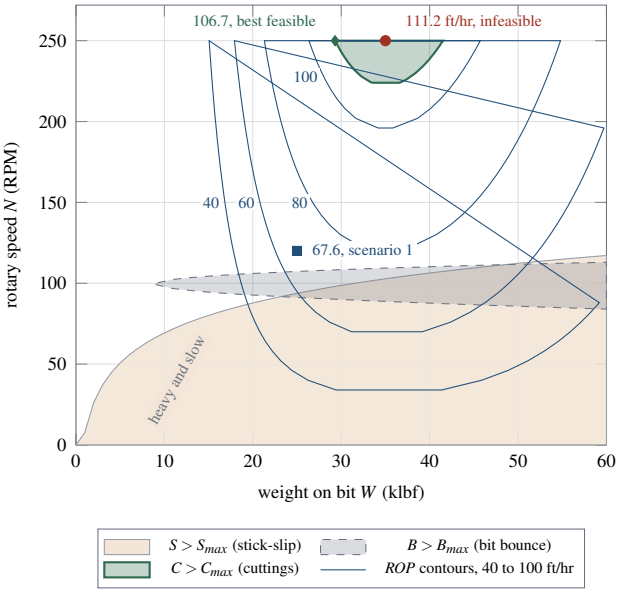


Figure 2.2 Objective contours and the three binding constraints in the weight and rotary speed plane, computed from Equations 2.1 to 2.13 at the scenario-1 choke position of 65 % open with a new bit. Stick-slip forbids the shaded corner with a thin solid outline, bit bounce the band with a dashed outline and the cuttings limit the lens with a thick outline. The circle marks the unconstrained peak, the diamond the best feasible point and the square the scenario-1 point. The unconstrained peak sits inside the cuttings lens, so the constrained optimum lies on the lower edge of the lens.

2.4.1 Shape of the Feasible Region

Figure 2.2 draws the objective and the three constraints that carve into it at the choke position of the first scenario. The objective contours are closed loops, because f_W rises and then founders while f_N increases monotonically, so every level set has a light branch and a heavy branch. The three forbidden sets have three different shapes. Stick-slip occupies the region below a boundary running from 92 RPM at 23 klbf to 117 RPM at 60 klbf. Bit bounce occupies a horizontal stripe from 84 RPM to 113 RPM at the heaviest weight, narrowing to nothing below 9 klbf where even the full resonant magnification cannot swing the weight by half of its mean. The cuttings limit occupies a lens in the top right that contains the unconstrained peak, and that lens is the one that binds.

2.4.2 A Closed-Form Ceiling on the Objective

The cuttings constraint is not an ordinary boundary in the plane. Setting $\frac{\mathrm{d}C}{\mathrm{d}t} = 0$ in Equation 2.7a gives $C = ROP\,A_{bit}/(Q\eta)$, so requiring $C \le C_{max}$ is exactly a bound on the objective itself, as shown in Equation 2.16.

$$ROP \le \frac{C_{max}\,Q\,\eta(\omega)}{A_{bit}} \tag{2.16}$$

Weight on bit does not appear in Equation 2.16 and neither does the choke. The ceiling depends on rotary speed alone, through the transport efficiency, and η saturates, so the ceiling rises from 92.4 ft/hr at 60 RPM to 99.9 ft/hr at 120 RPM and only 106.7 ft/hr at 250 RPM. Two consequences follow and both are visible on the application.

🕮 **Pause.** Try Exercise 2.2 before reading on.

The first is that the choke stops contributing beyond a crossover position. The unconstrained peak grows steadily as the choke opens, from 89.4 ft/hr with the choke shut to 111.2 ft/hr at 65 % open and

124.1 ft/hr wide open, because Equation 2.3c multiplies the whole surface. The best feasible point follows it only until the peak crosses into the lens at $x = 52.1\,\%$ open, and past that it is pinned at 106.7 ft/hr. Managed pressure drilling adds 19 % of drilling rate between a shut choke and 52 % open and nothing after that, while the operating point that delivers the ceiling slides along the boundary from 35.0 klbf to 24.7 klbf as the choke opens fully.

The second is that the binding constraint identifies its own remedy, and the remedy is not one of the three manipulated variables. Raising the pump rate from 476 gpm to 491 gpm, an increase of 3.3 %, lifts the ceiling to 111.7 ft/hr, above the unconstrained peak, so the whole plane becomes feasible and the achievable rate rises to 111.2 ft/hr. A further increase in pump rate reduces the drilling rate, because annular friction in Equation 2.5b raises bottomhole pressure and chip hold-down removes the gain, and at 602 gpm the achievable rate has fallen to 105.9 ft/hr. Flow rate is a fourth degree of freedom with an interior optimum of its own, and it is held fixed so that the three degrees of freedom without one can be studied.

2.4.3 Dependence of the Binding Constraint on the Formation

Scaling the whole objective by a constant cannot move an unconstrained optimum, and it does move a constrained one. Firing the hard-stringer event multiplies K by $1/1.45$, and the achievable rate at the best point falls from 106.7 ft/hr to 85.5 ft/hr while the concentration there falls to 1.76 %, under the limit, so the cuttings constraint releases and the answer moves back to the founder point at 35.0 klbf and 250 RPM. Wear does the same thing continuously, and on a bit worn to $h = 6$ of 8 the best point makes 51.4 ft/hr at the founder point with 1.06 % cuttings and is unconstrained. Nothing about the plant changed in either case.

2.4.4 Accuracy of the Vibration Indicators

The bounce index of Equation 2.13 is the exact steady amplitude of Equation 2.12 and is treated as such by the verification suite, which checks the dynamic model against it wherever the bit speed is steady enough for the excitation frequency to mean anything. The stick-slip severity of Equation 2.11 is a correlation, and the difference matters enough to quantify.

The onset G_0 comes from the linearization rather than from a fit. The two-inertia string of Equation 2.9 with the proportional-integral speed loop at its default tuning has a pair of eigenvalues that crosses into the right half-plane at $G = 1.092$, with a period of 2.99 s, at every rotary speed. The integrated model agrees: a settled run kicked by 10^{-6} rad/s of bit speed decays at $G = 1.08$ and grows at $G = 1.10$ at 60 RPM, 80 RPM and 100 RPM. The growth is slow near the onset, with an e-folding time of about 100 s at $G = 1.13$ and 33 s at $G = 1.22$, so a run started at a steady state can take a long time to show the cycle. With the weight and speed chosen and the run started, cells at 70 RPM to 85 RPM with G from 1.18 to 1.23 pass a severity of 0.40 after 487 s to 713 s, cells with G from 1.13 to 1.17 after 863 s to 1659 s, and 15.0 klbf at 80 RPM ($G = 1.096$) not within 2400 s. Every one of these cells lies above G_0 and the map marks it.

Driving the dynamic model at 96 points spanning 9 klbf to 58.5 klbf and 40 RPM to 110 RPM, each started on weight with an empty annulus and frozen wear and run for 600 s, produces a limit cycle at 74 of them. The windowed severity of a developed cycle wanders by tens of percent from one window to the next, so it is averaged over the last 300 s of each run. Comparing that average with the displayed value of Equation 2.11 gives a median relative error of 9.2 %, a ninetieth percentile of 17.6 % and a worst case of 21.1 % over the 72 points at which G exceeds the onset by more than 15 %. That is the accuracy a published field severity chart has. The application's own boundary test, twelve runs of 160 s at points at least 7 % from G_0 across the whole adjustable range, puts every simulated point on the predicted side of G_0.

Inside the 15 % band above the onset the ramp factor in Equation 2.11 lowers the displayed value while the dynamic model, once its cycle has grown, runs at full amplitude. At 36 klbf and 100 RPM, where $G = 1.245$, the simulation settles at a severity of 1.81 against a displayed 1.66, and at 49.5 klbf and 110 RPM ($G = 1.178$) at 1.98 against 1.49. Feasibility is therefore decided on the unramped

amplitude, 1.97 and 2.31 at those points. With that rule the map and the simulation agree about the 0.40 limit at 95 of the 96 points. The exception is 22.5 klbf at 90 RPM, where the map marks stick-slip and the 600 s run has not yet grown its cycle. A run started at the steady state there develops it at 1629 s, so the map is conservative for short runs near the onset and right for long ones.

A cycle once started also survives far below the onset. Moving from the scenario-1 point to 9 klbf and 70 RPM ($G = 0.96$) ten seconds into the run starts a cycle within 6 s that is still at a severity of 1.18 after a further 900 s, while the same point chosen before the run starts never leaves steady rotation. The boundary on the map therefore marks where steady rotation loses stability and not where a running cycle stops.

2.5 Controller Design

2.5.1 The Regulatory Layer

The default three proportional-integral loops run at 20 Hz on the shared velocity-form algorithm of Section 1.6.2, with derivative action set to zero throughout. The weight and rotary loops run in both modes, and the pressure loop runs in automatic mode and tracks the operator's choke in manual mode. Written in the loops' own symbols, with executions indexed by k and separated by $\Delta t = 0.05$ s, the algorithm is Equation 2.17.

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k) \tag{2.17a}$$

$$\Delta u_k^P = K_c\,e_k - K_c\,e_{k-1}$$

$$\Delta u_k^I = \frac{K_c}{\tau_I}\,e_k\,\Delta t$$

$$q_k = q_{k-1} + \Delta u_k^P$$

$$u_k^* = q_k + I_{k-1} \tag{2.17b}$$

$$u_k = \mathrm{slew}\big(\mathrm{sat}_{[u_{min},u_{max}]}(u_k^*),\ u_{k-1},\ \dot{u}_{max},\ \Delta t\big) \tag{2.17c}$$

$$I_k = I_{k-1} + \Delta u_k^I + \frac{\Delta t}{T_t}\,(u_k - u_k^*) \tag{2.17d}$$

where $a = \pm 1$ is the loop action, q_k is the accumulator that carries the proportional increments without limit (the running sum S_k of Equation (1.17)), u^* is the unclamped output the algorithm demands (the v_k of Equation (1.18)) and $T_t = \tau_I$ is the tracking time of the back-calculation. The output bias is zero and the initial actuator value is carried in I. Anti-windup is the back-calculation term of (2.17d), which corrects the integral alone and never truncates q_k, as Section 1.6.2 explains [21]. At the default $\tau_{I,P} = 4$ s each execution removes $1/80$ of the excess. The tracking term acts in the reported runs mainly in the challenge as delivered, where the drawworks feed sits at zero, its lower clamp, for 68 % of the run. In manual mode the pressure loop sets its integral to $I_k = u_{man} - q_k$, and in automatic mode the operator's choke value follows the controller output after every execution, so transfer is bumpless in both directions with no dedicated transfer logic.

The pressure loop is the loop with negative process gain. Its manipulated variable is the choke position in percent open and opening the choke lowers the pressure, so the loop is declared with $a = -1$. That single declaration is the only place a minus sign appears, and the loop is otherwise identical to the other two. Its output is slew limited at $\dot{x}_{max} = 20\,\%/\mathrm{s}$ before it reaches the stroke lag of Equation 2.6.

Table 2.3 lists the six tuning constants, and three derived quantities follow from them. The pressure loop has $K_c|K_p| = 2.0$ with $\tau_I = 4$ s, giving a velocity constant $K_v = K_c|K_p|/\tau_I = 0.5/\mathrm{s}$ against a plant lag of $\tau_p + \tau_v = 3.5$ s. The weight loop closes on an integrating plant, so its proportional gain sets a first-order closed-loop time constant of $1/(k_a K_{c,W}) = 2.5$ s with $\tau_I = 15$ s well outside it. The rotary loop sees the top-drive inertia through $J_t/K_{c,N} = 4.0$ s, slower than the 3.1 s natural period $2\pi\sqrt{J_b/k}$

Table 2.3 Controller Settings and Their Adjustable Ranges

Symbol	Meaning	Default	Range
$K_{c,P}$	choke per pressure error	5.0×10^{-5} %/Pa	0–2 %/psi
$\tau_{I,P}$	pressure integral time	4.0 s	0.5–60
$K_{c,W}$	feed rate per weight error	1.762×10^{-6}	0–400 (ft/hr)/klbf
$\tau_{I,W}$	weight integral time	15 s	2–300
$K_{c,N}$	torque per speed error	500 N m s/rad	0–3000
$\tau_{I,N}$	rotary integral time	3.0 s	0.2–30

of the torsional mode, so it regulates the mean surface speed and leaves the torsional mode to the string. Stiffening it does not cure stick-slip. At 40 klbf and 60 RPM, $K_{c,N} = 1000\,\mathrm{N\,m\,s/rad}$ lowers the peak severity from 4.80 to 2.90 and starts the cycle sooner (flagged at 14.3 s against 20.2 s), 2000 gives 2.94 and 250 raises it to 6.47, and the cycle persists at all four gains. Field soft-torque systems tune the top drive to absorb the torsional wave, which a plain PI speed loop does not do [2, 15].

The instrument chain is short. Bottomhole pressure passes through a delay line of adjustable dead time θ_p, which is the only transport delay in the model, defaults to zero and is read only by the pressure loop, and weight and surface speed are read exactly with no noise and no bias anywhere. The two vibration indicators come from windowed peak-to-peak detectors of length $T_w = 20\,\mathrm{s}$, which collect the extremes of their signal at every physics step and write the envelope into a ring buffer every 0.05 s, because storing the instantaneous value would alias badly against an axial excitation that reaches 12.5 Hz. Both buffers are pre-filled with the starting value so that a run beginning at a settled point reads zero swing rather than climbing out of a fictitious transient. Those two indicators are slow by construction, so an override on either acts on information up to 20 s old.

2.5.2 The Supervisory Layer

Above the three loops, and only in automatic mode, is a constraint-aware hill climber that owns the three setpoints. It is the third shape of constraint control in Section 1.4.8, a supervisor that moves setpoints toward the active limit and backs away from it. It executes once every T_{opt} seconds of simulated time and takes exactly one of three actions per execution. Every execution begins by measuring how far past its own limit each of the six constraints currently is, as shown in Equation 2.18.

$$[v_1, v_2, v_3, v_4] = \left[\frac{S}{S_{max}}, \frac{B}{B_{max}}, \frac{C}{C_{max}}, \frac{T_{bit}(W, \bar{\omega}_t)}{T_{lim}}\right] - 1 \tag{2.18a}$$

$$v_5 = (P_b - P_{high})/s, \qquad v_6 = (P_{low} - P_b)/s \tag{2.18b}$$

where P_b is the undelayed bottomhole pressure, $P_{low} = P_{pore} + \Delta P_m$, $P_{high} = P_{frac} - \Delta P_m$, $s = \max(\Delta P_m, 0.1\,\mathrm{MPa})$ normalizes the two pressure entries to the same scale as the four ratios, and $\bar{\omega}_t$ is the window mean of surface speed rather than the instantaneous value, because inside a stick-slip cycle the torque swings from nothing to the breakaway value several times a second. Every entry of v is negative when its constraint has room and positive when it does not, and all six are comparable. The most-violated constraint is $j = \arg\max_i v_i$.

Three states follow from v_j, as shown in Equation 2.19.

$$\begin{cases} \text{back off} & v_j > 0 \\ \text{hold} & -0.05 < v_j \le 0 \\ \text{climb} & v_j \le -0.05 \end{cases} \tag{2.19}$$

The hold band in the middle is what makes an override usable on a slow constraint. The cuttings concentration has a time constant of 340 s, so a back-off that ran until the indicator read clear

would keep moving for minutes after the move that fixed it, which is integral windup against a lagging measurement. Backing off is also rate-scaled by how far over the limit the indicator is, through an urgency factor $\kappa = \mathrm{clamp}(v_j/0.15,\ 0.15,\ 1)$, so the move sizes are $\Delta W = \dot{W}_{sp} T_{opt}\,\kappa$ and $\Delta\omega = \dot{\omega}_{sp} T_{opt}\,\kappa$. The setpoint that moves depends on the selected constraint. Stick-slip, cuttings and torque all reduce weight, with the weight setpoint floored at the cutting threshold W_t, below which f_W is identically zero and the climb would find no gradient. Bit bounce moves the rotary speed, and the direction is decided by evaluating the closed form 0.4 rad/s either side and taking the lower, because below the resonance the escape is to slow down and above it to speed up. The two pressure entries move the pressure setpoint by $3\dot{P}_{sp} T_{opt}$.

When nothing is close, the optimizer climbs. It estimates the gradient of the objective by central differences on the steady-state evaluator at the live choke position, with perturbations of 2 % of each adjustable range, as shown in Equation 2.20.

$$g_W = ROP(W_{sp}{+}h_W,\ \omega_{sp}) - ROP(W_{sp}{-}h_W,\ \omega_{sp}) \tag{2.20a}$$

$$g_\omega = ROP(W_{sp},\ \omega_{sp}{+}h_\omega) - ROP(W_{sp},\ \omega_{sp}{-}h_\omega) \tag{2.20b}$$

$$W_{sp} \leftarrow W_{sp} + \dot{W}_{sp} T_{opt}\, g_W / \sqrt{g_W^2 + g_\omega^2} \tag{2.20c}$$

$$\omega_{sp} \leftarrow \omega_{sp} + \dot{\omega}_{sp} T_{opt}\, g_\omega / \sqrt{g_W^2 + g_\omega^2} \tag{2.20d}$$

where $h_\omega = 0.524\,\mathrm{rad/s}$ and h_W starts at 5.34 kN (1.2 klbf) and doubles until g_W is nonzero, so that a probe straddling the dead zone below W_t widens until it reaches weights that cut. Normalizing by the gradient magnitude means the move is always at the rate limit and the direction alone carries the information, so the setpoint trends move at the declared rate limits and the limits can be read from the recorded data.

Separately, and on every climbing execution, the pressure setpoint walks down toward a floor, as shown in Equation 2.21.

$$P_1 = \min\left(P_{low} + 0.15\,\mathrm{MPa},\ P_{high}\right) \tag{2.21a}$$

$$P_2 = P_b\big|_{x=100} + 0.10\,\mathrm{MPa} \tag{2.21b}$$

$$P_{floor} = \max(P_1,\ P_2) \tag{2.21c}$$

Equation 2.21b is the term that binds in practice. A wide-open choke adds no back-pressure at all, so the lowest bottomhole pressure the well can reach is whatever the mud column and the annular friction produce on their own, and that minimum rises as the annulus loads up. Enforcing the floor on every execution rather than only while climbing keeps the pressure loop off its limit when cuttings push the achievable minimum upward underneath it. On a well with a light mud column the choke travel, and not the formation, limits what managed pressure drilling can add.

🕮 **Pause.** Try Exercise 2.1 before reading on.

Two further pieces complete the layer. On every execution, in both modes, a 26×22 sweep of the plane at the live choke position finds the highest feasible drilling rate, and a second pass over the points within 2 % of it takes the one with the lowest specific energy. The second pass exists because Equation 2.16 makes the cuttings boundary a contour of equally fast points with a light branch and a heavy branch, and without a tie-break the selected point alternates between them as the grid shifts. On every manual to automatic transition the three supervisory targets are re-seeded from the operator's setpoints, so the supervisory layer is bumpless in the same way the regulatory loops are, and a test that drives the transition measures a setpoint change of zero against a tolerance of 1 %.

2.5.3 Choke Model Predictive Control

Model predictive control (MPC) can replace the pressure PI loop as an option, while the weight and rotary loops, the supervisory layer and its protections stay as they are. The arrangement is the one of Section 1.6.3. The controller plans with a frozen nominal model in the pressure p and a baseline b, both in MPa, an estimated opening v and a command u, both in percentage points, and the current pump flow Q in m^3/s, which it takes as known exactly. With $a = 4(Q/0.030)^2$ MPa, the model is Equation 2.22.

$$\dot{v} = (u - v)/2, \quad \dot{p} = [b + a(1 - v/100) - p]/1.5, \quad \dot{b} = 0. \tag{2.22}$$

Its nominal gain is -40000 Pa per percentage point, and an independent step propagation agrees with it within 1.57×10^{-7} Pa. Every $h = 0.05$ s an exact discretization predicts (p^-, v^-, b^-) from the emitted command. The measured pressure p_m then sets $p = p_m$, $v = v^-$ and $b = b^- + g(p_m - p^-)$, with $g = (1 - e^{-h/5})/(1 - e^{-h/1.5})$. At initialization v is the known command and $b = p_m - a(1 - v/100)$. A change of pump flow adds $650(Q_{new}^{1.8} - Q_{old}^{1.8})$ MPa to b. This is a disturbance observer [22] with a filtered bias in place of a Kalman gain, and it reads no cuttings loading, wear, actual valve state or live model coefficient. The offset-free controller of Volume 9, Chapter 4 carries its disturbance in a Kalman estimator instead.

A decision is made every 1 s. The prediction holds each command over its block and spans 60 s in blocks of $(1, 1, 2, 3, 4, 6, 8, 10, 12, 13)$ seconds. With the currently issued target r in MPa, the block changes Δu_j and one global slack $\varepsilon \geq 0$ in MPa, the shared receding-horizon program [23] minimizes Equation 2.23, a case of Equation (1.21).

$$\sum_{k=1}^{60} (p_k - r)^2 + 0.001 \sum_{j=1}^{10} (\Delta u_j)^2 + 1000\varepsilon^2 + 10\varepsilon. \tag{2.23}$$

The move weight is in MPa^2 per squared percentage point, and the linear slack price is 10 MPa. The hard bounds are $0 \leq u \leq 100$ and $|\Delta u_j| \leq \dot{x}_{max}(1\,\text{s})$. The current limits, taken as known exactly, impose $P_{low}/10^6 - \varepsilon \leq p_k \leq P_{high}/10^6 + \varepsilon$ across the horizon. The emitted commands are also slewed every 0.05 s at 20 percentage points/s by default, and the predictions leave this extra ramp out. No future target or event enters them. A plan is accepted only if it is optimal, of the right dimensions, finite and within every planned bound. On a rejection the PI loop, which has tracked the applied command throughout, resumes until the solver recovers. In manual mode the controller follows the requested command. A forced failure of 12 s produces no jump at the first handover.

2.6 Results

The app follows the pattern of Section 1.6.4. Manual and automatic modes are provided, the drilling rate and each constraint indicator are displayed over the weight and speed plane with the feasible region marked, and the measured and best feasible drilling rates are displayed together. Both SI and field units are supported, and the exported file holds the 20 Hz samples, including the constraint indicators, the supervisory setpoints and the binding constraint at every sample.

The eleven original scenarios appear in Table 2.4 together with the measured performance of each, and Figure 2.3 sets the hole drilled and the rate against the best feasible rate side by side. A scenario fixes only the quantities its own lesson requires, so the change it makes is visible against the baseline.

Performance values were recorded after the elapsed time shown and appear in the right-hand half of the same table. For these original cases, the three standard loop metrics are omitted and quoted below only where they carry information. The integral of absolute error on the pressure loop accumulates in pascal-seconds against a setpoint the supervisory layer is deliberately moving, so it is quoted below as a mean absolute error in psi with its run length attached, and overshoot and settling time are produced by step-response calculations with no step to work with in most of the runs. Every run starts from a steady state at the scenario's own settings, so the first seconds carry no start-up transient.

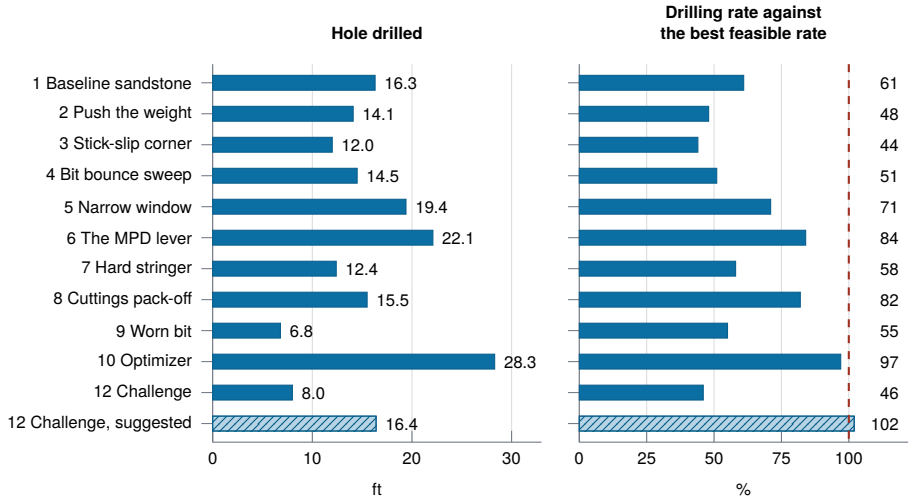


Figure 2.3 The scenarios of Table 2.4 side by side. Solid bars are runs of 900 s. The hatched bar is the suggested challenge, which stops itself at 473 s once 5.0 m is drilled. The dashed line marks the best feasible rate. Notice that only the optimizer and the suggested challenge come near that line, while the delivered challenge, automatic but held back by its conservative limits, stays far below it, and that the narrow window and the MPD lever drill more than the baseline because they carry less overbalance.

Table 2.4 The Original Scenarios, the Setting Each One Changes, and the Measured Performance of Each Run

#	Name	Setting	t (s)	Hole (ft)	ROP vs best (%)	MSE (kpsi)	Min margin (psi)	Peak S
1	Baseline sandstone	manual, 25 klbf, 120 RPM, choke 65 %, nothing scheduled	900	16.3	61	30.4	421	0.00
2	Push the weight	W_{sp} stepped 15, 25, 35, 45, 55 klbf every 90 s	900	14.1	48	82.5	354	0.08
3	Stick-slip corner	40 klbf and 60 RPM, heavy and slow	900	12.0	44	34.1	402	4.80
4	Bit bounce sweep	ω_{sp} stepped 160, 130, 110, 100, 90, 85 RPM	900	14.5	51	26.2	409	1.70
5	Narrow window	$P_{pore} = 31.0$ MPa, $P_{frac} = 33.2$ MPa, mud 1000 kg/m^3	900	19.4	71	25.6	139	0.00
6	The MPD lever	choke stepped 10, 35, 60, 85, 100 % at fixed W and ω	900	22.1	84	32.1	241	0.00
7	Hard stringer	$\Delta_{hard} \to 45$ %, $a_{abr} \to 1.6$ at $t = 150$ s	900	12.4	58	42.9	392	0.00
8	Cuttings pack-off	$Q = 0.011$ m^3/s, restored at $t = 420$ s	900	15.5	82	54.9	332	0.00
9	Worn bit	identical to 1 with $h = 6$ of 8	900	6.8	55	72.7	355	0.00
10	Optimizer	automatic from the settings of scenario 1	900	28.3	97	47.4	263	0.16
12	Challenge	automatic with every limit set conservatively and the pressure loop detuned	900	8.0	46	102.8	309	0.07
12*	Challenge, suggested	$C_{max} = 2.4$ %, $B_{max} = 0.50$, $S_{max} = 0.40$, $\Delta P_m = 0.7$ MPa, $K_{c,P} = 5 \times 10^{-5}$, $\tau_{I,P} = 4$ s	473	16.4	102	49.7	263	0.16

* the run stops itself the moment the 5.0 m interval is drilled, which is why its elapsed time is the score

2.6.1 Weight Steps and the Founder Point

The second scenario steps the weight through five values 90 s apart. At steady state the evaluator gives 26.2 ft/hr, 67.6 ft/hr, 79.9 ft/hr, 71.9 ft/hr and 55.1 ft/hr at 15 klbf, 25 klbf, 35 klbf, 45 klbf and 55 klbf, and the run reads 26.1 ft/hr, 74.0 ft/hr, 87.6 ft/hr, 73.8 ft/hr and 52.1 ft/hr at the end of each step, higher at the middle weights because 90 s is a quarter of the cuttings time constant and the annulus has not yet loaded. The peak is at the founder point in both, where $\chi = 1$ in Equation 2.2, and

past it the extra weight costs drilling rate outright. Two other quantities move during the same sequence and neither moves with the objective. Bit torque climbs monotonically from 2.09 kN m to 6.64 kN m, because Equation 2.9c is proportional to weight and carries no founder term. Specific energy goes 47.1 kpsi, 28.5 kpsi, 32.8 kpsi, 45.8 kpsi and 71.9 kpsi at steady state (47.3 kpsi, 26.1 kpsi, 29.9 kpsi, 44.7 kpsi and 76.0 kpsi in the run), cheapest at 25 klbf rather than at the founder point, so between 25 klbf and 35 klbf the drilling rate rises by 18 % and the specific energy rises by 15 % with it, in both. Whether that is a good trade depends on whether rig time or bit life is the binding cost, and specific energy is the number that distinguishes them [14].

2.6.2 Stick-Slip and Bit Bounce

The third scenario sits at 40 klbf and 60 RPM, and the bit stops and then breaks free at up to 418 RPM, seven times the speed the top drive is holding. The cycle is irregular, with a median period of 5.1 s and a tenth-to-ninetieth percentile range of 3.9 s to 13.5 s, longer than the 3.1 s natural period of the torsional mode. Severe stick-slip is flagged at 20.2 s, and the severity averages 3.39 over the last 600 s and peaks at 4.80 against a limit of 0.40. The cost shows up in two places at once. The run makes 12.0 ft of hole in 900 s against 16.3 ft for the baseline, and the top drive spends 1.06×10^6 N m of total variation on its torque command against 0.12 N m in the baseline, a factor of 8.9×10^6. A torsional limit cycle is a fatigue problem for the string and a duty problem for the drive, and the drilling rate is not the largest cost. Speed cures it reliably only before it starts. Raised a minute into the run, to 150 RPM the severity settles near 1.5, at 200 RPM near 1.3, and only 250 RPM quenches it, although 200 RPM applied at the first flag, 20 s in, still does. A run started at 150 RPM never stick-slips.

> **TRY IT IN THE APP** **3 · Stick-slip corner (heavy and slow)**
>
> Load the case and press Run in Manual. Watch the Rotary speed strip, where the surface trace and the bit trace share one plot. The top drive holds its 60 RPM while the bit stops and then breaks free at up to 418 RPM, over and over at irregular intervals. Severe stick-slip is flagged at 20.2 s.

The fourth scenario sweeps the rotary speed down through the axial resonance at 160 RPM, 130 RPM, 110 RPM, 100 RPM, 90 RPM and 85 RPM. The bounce index peaks at 100 RPM exactly, where Equation 2.13 gives 0.833 against 0.216 at 120 RPM, and the application reports 0.83 there. The sweep also shows that the two vibration limits are neighbors, because at 90 RPM and 25 klbf the run leaves the bounce stripe and enters the stick-slip corner. Severe stick-slip is flagged at 398.9 s, the severity reads 1.69 at the end of the run with a peak of 1.70, and the reported constraint is stick-slip rather than bounce. Escaping a resonance by slowing down works only when the stick-slip region is elsewhere.

2.6.3 Managed Pressure Scenarios

The sixth scenario holds the weight at 30 klbf and the speed at 150 RPM and moves nothing but the choke. Opening it from 10 % to 100 % takes bottomhole pressure from 5678 psi at 90 s to 5177 psi at 430 s, equivalent circulating density from 11.10 ppg to 10.12 ppg, and the drilling rate from 69.3 ft/hr to 97.3 ft/hr. The choke alone adds 40 % of drilling rate with no other change, and the scenario separates a manipulated variable from a disturbance, because the same overbalance would conventionally have been inherited from a mud density chosen days earlier.

The fifth scenario is the case that motivated the commercial adoption of managed pressure drilling. Pore pressure in the depleted sand is 4496 psi and the rock fractures at 4815 psi, a window 319 psi wide of which the two margins take 203 psi. Cutting the mud to 8.35 ppg means the column alone reaches only 4268 psi of hydrostatic head, so the well is held inside its window by the choke rather than by the mud. The feasible band is choke positions of 77.8 % to 99.9 % open, about a fifth of the travel, with the minimum achievable pressure at a wide-open choke sitting 101 psi above pore pressure. The run holds 139 psi of minimum window margin over 900 s and makes 19.4 ft, more than

the baseline, because the reduced overbalance of the light mud adds more drilling rate through chip hold-down than the narrow window removes.

2.6.4 Cuttings Pack-Off

The eighth scenario drills at 35 klbf and 200 RPM with the pumps at 174 gpm, an annular velocity of 90 ft/min. Cuttings are made faster than they are carried away, and the case opens at its steady loading of 5.63 %, past the pack-off threshold of 5.5 %, where the friction multiplier of Equation 2.5c has begun to climb steeply. The concentration is finite rather than running away, and the mechanism that holds it is the chip hold-down coupling of Equation 2.3c, because higher pressure reduces the rock removal rate. Restoring the pumps to 476 gpm at $t = 420$ s first raises bottomhole pressure by 378 psi on the added friction of the loaded annulus and takes the concentration below the pack-off threshold within 6 s. The annulus then cleans out slowly, to 2.43 % at 900 s with the drilling rate at 81.8 ft/hr, and the concentration falls below the 2.2 % limit only at 1030 s. The reduced flow also removes choke authority, 78 psi against 580 psi at full flow, so for the first seven minutes of the run the choke can do little about the pressure.

> **TRY IT IN THE APP** **8 · Cuttings pack-off (fast hole, slow pumps)**
>
> Load the case and press Run in Manual. Read the cuttings trace on the Constraint indicators strip: the annulus opens at 5.63 % by volume, past the pack-off threshold of 5.5 %. At 420 s the pumps come back to 476 gpm. Watch the Bottomhole pressure strip rise by 378 psi at once, then the annulus clean out slowly, to 2.43 % at 900 s.

2.6.5 Bit Wear Effects

The ninth scenario is scenario 1 with $h = 6$ of 8. It makes 6.8 ft in 900 s against 16.3 ft for a sharp bit, at a specific energy of 72.7 kpsi against 30.4 kpsi. The bit spends more than twice as much energy per unit volume of rock to make hole at 42 % of the rate, which is the signature that separates a dull bit from hard rock and the criterion that triggers a trip for a new bit [14].

2.6.6 Supervisory Optimization

The tenth scenario takes the same well and the same opening settings as the first and switches the supervisory layer on, and Figure 2.4 is that run frozen near its end. Over 900 s it makes 28.3 ft of hole against 16.3 ft, an improvement of 73 %, and finishes at 97 % of the best feasible rate against 61 %. What it does to get there is visible in three separate movements. The rotary setpoint runs up to the top of its range at 250 RPM by 52 s, the weight setpoint reaches the founder point at 35.0 klbf by 60 s and is later pulled back by the cuttings override, ending at 28.3 klbf, and the pressure setpoint walks down from 5352 psi to 5193 psi by 110 s, where the choke reaches 97.5 % open and Equation 2.21 stops it. The floor then rises with the cuttings loading and carries the setpoint back to 5225 psi. The binding constraint at the end is cuttings and the concentration reads 2.19 % against a limit of 2.2 %. The scenario demonstrates that a constrained optimum is a point on a boundary and that a supervisory controller stopping short of the boundary gives up drilling rate. The supervisory layer of Volume 9, Chapter 8 meets the same structure, with a lowest-cost point on a constraint that moves with the weather.

Two features of the trends require interpretation. From 30 s to 427 s the measured drilling rate sits above the best-feasible rate, peaking at 140 ft/hr, which is not an inconsistency, because the best-feasible rate is a steady-state quantity and the annulus takes about five minutes to load up (the concentration first passes its limit at 297 s), so the well genuinely drills faster than any point that will be feasible once it has. The second feature is hunting along the boundary. The binding indicator is the cuttings concentration, which answers a weight move only through the cuttings inventory, so the override acts on a measurement that lags its own action by minutes and the result is a slow cycle along the boundary rather than a direct arrival at it. After 300 s the weight setpoint moves between

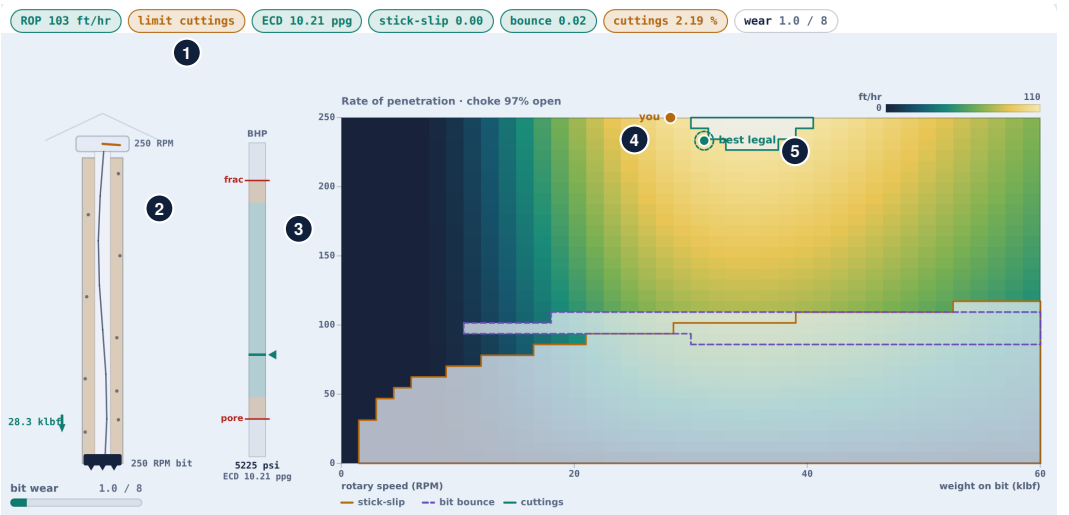


Figure 2.4 The supervisory optimization run of scenario 10 at $t = 891$ s, cropped to the plant view. (1) The chip that names the binding constraint, cuttings. (2) The drill string, with the weight and the surface and bit speeds. (3) The bottomhole pressure gauge between its fracture and pore marks. (4) The operating point on the drilling-rate map, which is evaluated at the choke position the optimizer has reached, 97 % open. (5) The best feasible point on the cuttings boundary. The operating point and the best feasible point have almost converged at 102.7 ft/hr against 106.2 ft/hr with cuttings as the binding constraint.

22.6 klbf and 34.8 klbf and the concentration between 2.08 % and 2.27 %, above its limit 42 % of the time. Every constraint controller acting on a lagging measurement behaves this way, and the alternative, backing off far enough that the lag never matters, is the conservatism the run quantifies.

2.6.7 The Graded Challenge

The challenge asks for 5.0 m of new hole within 900 s while the pressure loop holds bottomhole pressure within 17 psi of a setpoint the optimizer is moving and no constraint flag ever sets. Both conditions are checked at every physics step until the run stops itself, so a single excursion fails the attempt even after it clears. The challenge is delivered with two independent problems. Every constraint limit has been set well inside where it needs to be, with $C_{max} = 0.8\,\%$, $B_{max} = 0.20$, $S_{max} = 0.30$ and $\Delta P_m = 2.0$ MPa on a window only 1088 psi wide, and the pressure loop has been detuned to $K_{c,P} = 8 \times 10^{-6}\,\%/\text{Pa}$ and $\tau_{I,P} = 30$ s.

Only one of those is a control problem. The limits are an engineering choice, and the cuttings limit in particular has been set below the 1.49 % the default operating point already carries, so the cuttings flag is set from the first step and the override takes the weight setpoint down to its floor at $W_t = 5.0$ klbf within 60 s. The annulus then drains, the climb resumes near 300 s, and the weight setpoint cycles between 5.0 klbf and 26.9 klbf for the rest of the run. As delivered the run makes 8.0 ft (2.44 m) in 900 s with the cuttings and bounce flags set, ends at a specific energy of 102.8 kpsi and reaches a worst pressure error of 49.1 psi, so it fails. The detuned loop is the tuning problem, and its cost appears against a moving setpoint. A loop tracking a ramp settles with a standing offset of $\dot{P}_{sp}/K_v$, and the delivered tuning gives $K_v = 0.0107/\text{s}$ against 0.5/s for the default, so the predicted offset falls from 136 psi to 2.9 psi when the gain is multiplied by 6.25 and the integral time divided by 7.5. The two problems cost different things. Loosening only the limits drills the interval in 484 s with a worst pressure error of 94.4 psi, and retuning only the loop leaves 2.44 m at 900 s. Applying the suggested settings, which set the cuttings limit to 2.4 % because the concentration reaches 2.32 %

Table 2.5 Full 900 s Choke Comparison. Entries Within Each Cell Are Shipped PI / MPC / Fixed SIMC PI

Case	Pressure IAE (MPa s)	Outside known band (s)	Command travel (percentage points)
Pore rise	9.036 / 13.947 / 8.899	4.825 / 6.280 / 4.840	184.876 / 189.716 / 159.353
Fracture fall	5.762 / 32.289 / 5.091	44.080 / 42.945 / 43.830	159.903 / 193.732 / 122.204
Rate push	3.021 / 7.374 / 2.679	0.000 / 0.000 / 0.000	118.966 / 42.331 / 88.477
Challenge, shipped	85.065 / 6.959 / 2.883	0.000 / 0.000 / 0.000	26.084 / 42.416 / 63.738
Challenge, suggested	2.663 / 7.163 / 2.353	0.000 / 0.000 / 0.000	105.777 / 36.999 / 78.082
Shifting window	6.984 / 24.721 / 6.754	25.600 / 63.135 / 25.345	230.928 / 224.508 / 181.592
Gain/lag mismatch	13.153 / 27.834 / 17.801	34.255 / 46.500 / 36.990	355.727 / 219.467 / 378.930
Delay mismatch	83.200 / 33.084 / 82.733	78.895 / 62.530 / 75.830	2286.034 / 275.660 / 1935.749

before the interval is drilled, completes the interval in 473 s with no flag set, a worst pressure error of 4.9 psi, 12.1 psi inside the band, and a mean absolute error of 0.80 psi, against 49.1 psi and 13.7 psi for the delivered tuning, a peak stick-slip of 0.16 and 263 psi of minimum window margin. With the default 2.2 % cuttings limit the same settings set the cuttings flag before the interval is complete. The target requires operating at a limit rather than away from it, which is reasonable for a pressure limit measured continuously with a lag of a few seconds and not for a stick-slip limit measured over a 20 s window.

2.6.8 Pressure-Controller Comparison

Table 2.5 compares three pressure controllers on eight cases. Every run lasts 900 s at seed 20260824, with the same instruments and settings for the three controllers in each case and no MPC tuning per case. The third controller is a fixed PI loop derived from a model by the SIMC half rule [24]: $\tau = 2.75$ s, $\theta = 0.75$ s and a chosen $\lambda = 1$ s give $K_c = \tau/[40000(\lambda + \theta)] = 3.92857 \times 10^{-5}$ %/Pa and $\tau_I = 2.75$ s. No tuning search was used.

Scenario 11, the shifting window, raises the pore pressure by 0.9 MPa from 60 to 120 s and lowers the fracture pressure by 1 MPa from 180 to 240 s, with no notice to either controller, and the narrowest band is 0.4 MPa wide. The MPC loses on IAE and on time outside the band (Table 2.5). The supervisor feeds back on the pressure each controller achieves, so the two runs receive different targets (Figure 2.5). At 235 s the PI pressure and target are 32.072 and 32.070 MPa, against 32.110 and 32.310 MPa under the MPC. The upper margin there is 32.100 MPa, which is not the actual fracture pressure of 32.8 MPa, and neither run crosses an actual formation edge.

On the rate push and the suggested challenge the MPC trades a higher IAE for less choke travel. It loses under gain and lag mismatch, with an authority of 7 MPa, a pressure lag of 4 s and a valve lag of 5 s. With a transmitter delay of 2 s it wins on IAE, 33.084 against 83.200 MPa s. The fixed SIMC loop has a lower IAE than the MPC in all six nominal cases. The penalties, the slew limit and the supervisor feedback interact, so these data do not identify their separate effects.

> ✦ **DID YOU KNOW? A Two-Second Delay Keeps the Choke Moving**
>
> From the simulator (Table 2.5). With a transmitter delay of 2 s, the shipped PI loop moves the choke command through 2286.034 percentage points in 900 s and the fixed SIMC loop through 1935.749, while the MPC moves it through 275.660. Without the delay, on the rate push, the same PI loop travels 118.966.

The pore-rise case sets the fracture pressure to 34.5 MPa and raises the pore pressure by 1.6 MPa at 60 s. The pressure then stays below the actual pore pressure for 3.280 s under PI and 3.255 s under MPC, while 0.04922 and 0.04885 m of hole are drilled. The fracture-fall case uses a mud of 950 kg/m^3 and a pore pressure of 29.8 MPa and lowers the fracture pressure by 2 MPa at 15 s. The MPC shortens the time above the actual fracture pressure from 20.700 s to 3.040 s and the hole drilled meanwhile from 0.16336 m to 0.02318 m, while its IAE is worse and the peak excess rises from

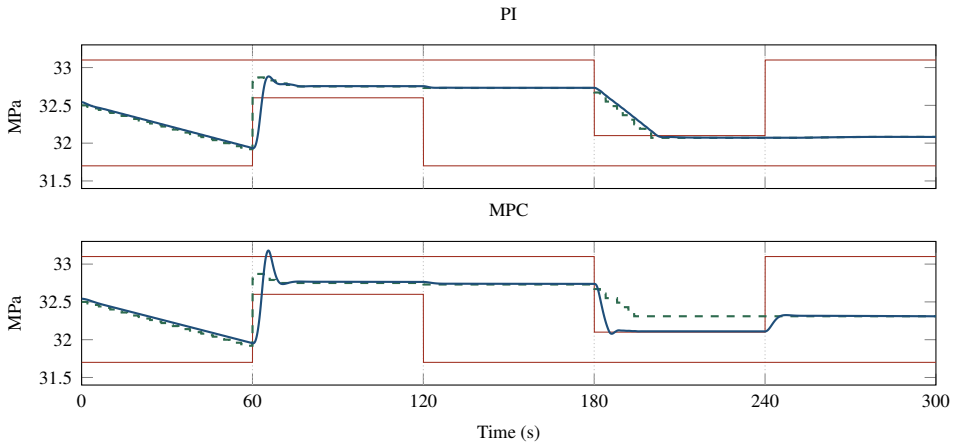


Figure 2.5 Bottomhole pressure in the shifting-window scenario over 300 s, under PI (top) and under MPC (bottom). The thick solid line is the pressure, the thick dashed line the current issued target and the thin solid lines the known band. Dotted vertical lines mark the four window events, and both panels share their axes. Under MPC the tracking error and the band violation remain.

579.111 to 588.859 kPa. In scenario 11 the cuttings limit, which the comparison does not score, is broken for 410.17 s under PI and 410.59 s under MPC, over 3.84795 m in both. In the shipped challenge the bounce limit is broken for 33.30 and 33.15 s, over 0.14840 and 0.14814 m. With the suggested settings all three controllers pass the challenge, reaching 5 m at 473.25, 473.60 and 473.22 s, and with the shipped limits all three fail. The grade concerns the run up to 5 m, while the metrics of the table continue to 900 s.

IAE and travel are computed from the 20 Hz samples, and the limit times and depths on the 5 ms physics grid. Halving that step for all 24 full runs leaves the rankings and grades unchanged. Every MPC run makes 900 solves with no rejection. In a browser the mean times of the solver, of a complete decision and of a complete control scan are 227, 373 and 782 μs at speed 1 and 107, 184 and 570 μs at speed 20. All control work, including the loops kept and the ROP search, takes 0.121 and 1.251% of a core at those two speeds.

2.7 Technology in Practice

Figure 2.6 draws the three regulatory loops and the supervisory layer of Section 2.5 with the hardware a rig installs.

What Measures It

A rotary rig hoists the drill string, turns it, presses the bit down on the rock and circulates mud down the inside of the string, through the bit and up the annulus with the cuttings [6]. The weight on bit is inferred from the pull on the drilling line. A hydraulic load cell turns a force into a liquid pressure that an ordinary pressure transmitter reports as weight [25], and the 1952 feed control read the line tension the same way, on a gauge calibrated as a weight indicator [12]. The return flow is usually metered with a Coriolis mass flowmeter, which needs a continuous flow and a closed system that can take its pressure drop of 3 to 5 psi [18]. Bottomhole pressure comes from a sensor near the bit. The chapter reads weight and surface speed exactly, with no noise and no bias.

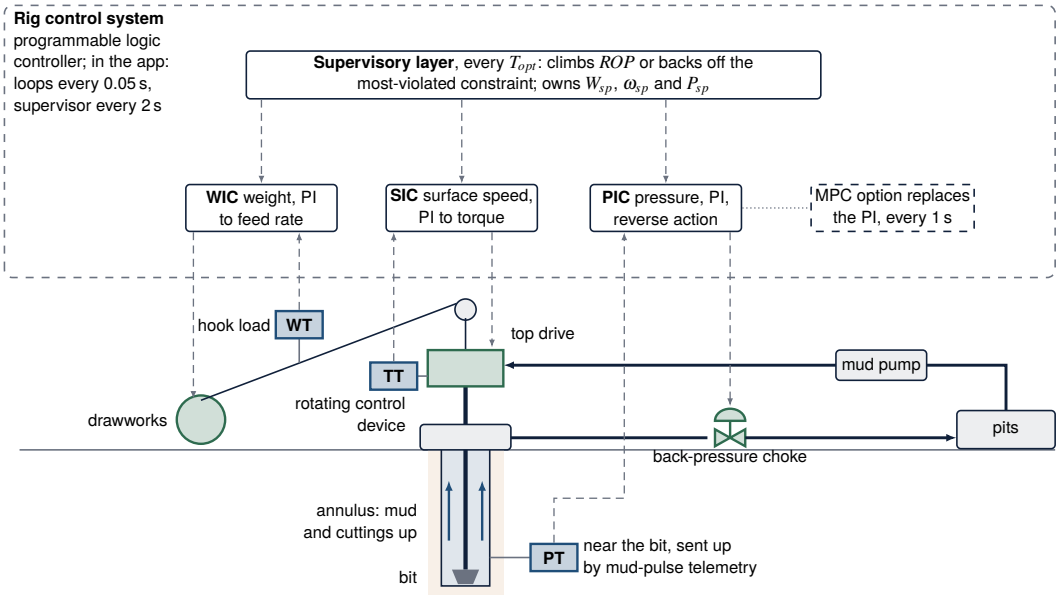


Figure 2.6 The loops of the drilling simulator with their hardware. Lettered boxes are sensing devices: the downhole pressure sensor (PT), the hook load cell on the drilling line (WT) and the top-drive speed sensor (TT). The drawworks drum, the top drive and the back-pressure choke are the final elements, and dashed arrows are signals. Notice that the pressure reading travels up through the mud before any controller sees it, and that the supervisory layer moves only setpoints. The vibration and cuttings indicators the supervisor reads and the setpoint rate limits are left out. Not to scale, and not a piping diagram.

> **TECHNOLOGY** **Sensor: Bottomhole Pressure Sent Up the Mud**
>
> A pressure sensor near the bit sends its reading to the rig through the mud itself. A valve downhole briefly restricts the flow, and the pressure pulses travel up the drill string to the standpipe pressure sensors [6]. Positive- and negative-pulse tools send 0.5 to 3.0 bits per second and continuous-wave tools up to 18, and the rate falls with depth to 1.5 to 3.0 bits per second at 10 000 m. Aerated mud is more compressible and carries the pulses worse. Most such tools work below 150 °C, and hotter holes make them unstable [6]. The chapter reads bottomhole pressure exactly, through a dead time θ_p that is zero by default.

What Moves It

The drawworks pays out drilling line and so sets the weight, the top drive turns the string, and the choke on the returns sets the back-pressure. In the chapter the feed is clamped at $v_{max} = 0.030\,\mathrm{m/s}$ and the top-drive torque at $T_{max} = 45\,\mathrm{kN\,m}$, and the choke is the only final element with a lag of its own in the model.

> **TECHNOLOGY** **Final Element: The Choke Manifold and Its Seal**
>
> A rotating control device seals the annulus around the turning pipe and sends the returns to the managed pressure choke manifold, or to the main flow line for conventional drilling [18]. The manifold controls the well pressure, not the flow rate, and when an influx is detected the choke is moved automatically to raise the surface back-pressure [18]. A back-pressure pump often accompanies the choke. A test model has to reproduce choke plugging and choke swapping [17], and a control system has to handle a blocked choke [19]. The chapter's choke has a 2 s stroke lag, a slew limit of 20 %/s and a linear characteristic.

Where the Controller Runs

Programmable logic controllers serve as local controllers within a supervisory scheme, with programmed functions that include logic, timing and three-mode PID control [26]. Supervisory and lower-level loops of such systems run with cycle times from milliseconds to minutes [26]. The app runs its three PI loops every 0.05 s, the choke MPC every 1 s and the supervisory layer every $T_{opt} = 2.0\,\text{s}$, all inside that range. One field stick-slip prevention system is a PI-type speed controller tuned to damp the torsional oscillation, with all of its software in a standard programmable logic controller [15]. The chapter's structure maps onto three PID blocks, one with reverse action for the choke, under a supervisory sequence that owns their setpoints.

> **IN PRACTICE** **The Safe Drilling Margin Is a Rule**
>
> In United States federal offshore waters, 30 CFR 250.414 asks each drilling prognosis for a planned safe drilling margin between the estimated pore pressure and the lesser of the estimated fracture gradient and the casing shoe pressure integrity test. The equivalent downhole mud weight must exceed the pore pressure and, unless a justified alternative is approved, sit at least 0.5 pound per gallon below the lower of those two limits [27]. Under 30 CFR 250.427 the margin must be kept while drilling. When it cannot be, the operator suspends drilling and submits remedial actions, or notifies the regulator and acts under API Bulletin 92L [28]. The chapter's window margin ΔP_m, 0.7 MPa at each end by default, plays this part.

What Is Modernizing It

The first development is high-speed telemetry. Wired drill pipe carries data at 57,000 bits per second and can send commands down to the tools for closed-loop control, although it is still expensive and a single failed connection stops the link [6]. With it, rate of penetration and bottomhole pressure have been combined in one controller for a managed pressure well [20]. The second is open standards. A 2013 review expected open standards that enable automation to reshape the industry's business model, as they did for industrial automation in the 1990s, and found collaboration on an open environment for drilling data [10]. The third is testing in simulation. A control system for automated managed pressure drilling should pass representative test scenarios in simulation before it is used in the field [17], which is the role the app of this chapter plays on a smaller scale.

2.8 Limitations

The predictor of the choke MPC omits the evolution of the cuttings loading, the wear and future supervisory actions. Unknown formation limits and mismatch in the parameters or in the sensor delay remain limitations. A priced operating band cannot guarantee well control, and the plant models no influx, gas migration or fracture flow.

> **OPEN CHALLENGE** **Seeing a Kick in Time**
>
> A kick is formation fluid that enters the well when the bottomhole pressure falls below the formation pressure, and a late detection can lead to an uncontrolled blowout [18]. Managed pressure drilling usually watches the return flow with a Coriolis mass flowmeter, but the meter reads poorly until its line is full or when the returning mud carries gas, which a 2020 review names as a need for further research. The same review finds no kick circulation method that suits every drilling operation and no single rule for the response [18]. The chapter's plant has no influx at all. Progress would be a detector and a response that hold across muds, depths and rigs.

Section 2.3 lists the modeling assumptions and they are not restated here. Four further shortcomings are properties of the implementation.

The stick-slip boundary on the map is the linear stability limit G_0 of the model at the default speed-loop tuning, and it describes steady rotation only. Just above it a steady start can stay quiet for half an hour or more before the cycle appears, so the map is conservative for short runs there, and a

cycle already running survives far below it, as quantified in Section 2.4. The map keeps G_0 when the speed-loop gain is changed, although the onset moves with it. Inside the band above the onset the displayed severity understates the developed cycle, and feasibility there is decided on the unramped amplitude. Marking the region where a running cycle persists would show the hysteresis on the map as well.

The hill climber of Equation 2.20 has no memory of why it backed off. The back-off is floored at the cutting threshold $W_t = 5.0\,\mathrm{klbf}$ and the weight probe widens until it finds a gradient, so the climber cannot stall in the region where f_W is identically zero. A limit set below the loading of the operating point it starts from still produces a cycle rather than a steady compromise. In the challenge as delivered the weight setpoint is at its floor within 60 s, the climb resumes near 300 s, and from then on the setpoint swings between 5.0 klbf and 26.9 klbf while the cuttings loading swings between 0.70 % and 0.92 % around its 0.8 % limit. The cycle comes from climbing at a fixed rate against a measurement that lags by minutes, the mechanism behind the hunting in the optimizer scenario.

Bit bounce is detected and does not act. The application does not show what a bouncing bit does to a bottomhole assembly, only that the operating point is one at which bounce would happen, and the constraint is therefore an assertion rather than a demonstration.

Flow rate is a parameter rather than a manipulated variable. Section 2.4 shows that it has an interior optimum near 491 gpm and that moving it 3.3 % does more for the achievable drilling rate than the entire upper half of the choke travel, so the most valuable extension is to promote it to a fourth degree of freedom and let the optimizer find that optimum too.

The documented model also supports control methods beyond the implemented PI loops, choke MPC and gradient search. Model predictive control over Equations 2.1 to 2.14 would coordinate weight, rotary speed and choke position in one constrained optimization, treat the pressure window and the cuttings ceiling as explicit constraints rather than override triggers, and act ahead of the 340 s cuttings lag instead of waiting inside the hold band of Equation 2.19. Moving horizon estimation or Kalman filtering would reconstruct the cuttings concentration and the downhole pressure from surface measurements, the soft-sensing role that simplified hydraulics models already fill on managed pressure wells [29] and that the downhole card fills for the pump of Chapter 3. Combined rate-of-penetration and pressure control of this form has been demonstrated on a test well with high-speed telemetry [20]. The APMonitor framework implements the nonlinear estimation and predictive control formulations these extensions require [30], and the Gekko Optimization Suite provides the same capabilities in Python [31].

> **OPEN CHALLENGE** **Pressure Waves the Controller Does Not See**
>
> Pressure controllers for managed pressure drilling usually rest on simplified models that ignore pressure waves in part or in full [17]. Those waves act over tens of seconds, or minutes when gas has entered the well, and if the design leaves them out they can degrade the control and even destabilize it. In field data a pump pressure transient reached the choke about 4 s later, at the speed of sound in the mud [17]. The chapter's hydraulics are quasi-steady, with one lag. Progress would be controllers designed on simple models and proven on models that keep the waves before they reach a rig.

2.9 Conclusion

This chapter set out a thirteen-state drilling and managed pressure simulator, with its two control layers and its verification, closely enough for the whole of it to be built again from the text. The rate of penetration is a five-factor product whose weight term founders at a declared point, and every constraint on it increases with the same two factors that increase the objective, so the optimum is a point on a boundary. At the default settings that peak is 111.2 ft/hr and infeasible, the best feasible point makes 106.7 ft/hr, and the operating point the first scenario supplies makes 67.6 ft/hr. The cuttings constraint turns out to be an explicit ceiling on the objective, $C_{max} Q \eta(\omega)/A_{bit}$, which is set by rotary speed and pump rate and carries no dependence at all on weight on bit, the variable adjusted

first in practice. It pins the achievable rate at 106.7 ft/hr for every choke position above 52.1 % open, while a 3.3 % increase in pump rate lifts it above the unconstrained peak. The implemented supervisory layer climbs the objective gradient above three proportional-integral loops and hands the offending setpoint to the most-violated constraint, and it drills 28.3 ft in 900 s against 16.3 ft for the fixed manual point, finishing at 97 % of the feasible ceiling with cuttings as the binding constraint. The stick-slip severity on the map is a correlation with a median error of 9.2 % away from the onset, feasibility is decided on its unramped amplitude, and the onset itself is the linear stability limit of the string at $G = 1.092$, where a disturbance grows slowly enough that a steady start can stay quiet for many minutes. The optional choke MPC reduces the choke travel on the rate push from 118.966 to 42.331 percentage points while raising the IAE from 3.021 to 7.374 MPa s.

A fuller app would feed bit bounce back into the depth of cut, promote flow rate to a manipulated variable, add lateral vibration modes, and replace the hill climber with a model predictive formulation for comparison. As it stands, the app locates an economic optimum on a constraint boundary and measures the cost of operating away from one.

KEY POINTS

- Pushing harder and turning faster speed up the drilling, but every limit grows with them, so the best point sits on the edge of the safe region.
- At the default settings the unconstrained peak of 111.2 ft/hr is infeasible and the best feasible point makes 106.7 ft/hr.
- The cuttings constraint is a ceiling on the drilling rate set by rotary speed and pump rate, not by weight on bit.
- The supervisory layer drills 28.3 ft in 900 s against 16.3 ft for the fixed manual point.

Where else this applies. Volume 9, Chapter 8 moves a supervisory setpoint along a constraint that shifts with the weather.

References

[1] T. M. Warren, “Penetration-rate performance of roller-cone bits,” *SPE Drilling Engineering*, vol. 2, no. 1, pp. 9–18, 1987.

[2] J. D. Jansen and L. van den Steen, “Active damping of self-excited torsional vibrations in oil well drillstrings,” *Journal of Sound and Vibration*, vol. 179, no. 4, pp. 647–668, 1995.

[3] T. Richard, C. Germay, and E. Detournay, “A simplified model to explore the root cause of stick-slip vibrations in drilling systems with drag bits,” *Journal of Sound and Vibration*, vol. 305, no. 3, pp. 432–456, 2007.

[4] A. T. Bourgoyne, Jr. and F. S. Young, Jr., “A multiple regression approach to optimal drilling and abnormal pressure detection,” *Society of Petroleum Engineers Journal*, vol. 14, no. 4, pp. 371–384, 1974.

[5] T. I. Larsen, A. A. Pilehvari, and J. J. Azar, “Development of a new cuttings-transport model for high-angle wellbores including horizontal wells,” *SPE Drilling & Completion*, vol. 12, no. 2, pp. 129–135, 1997.

[6] T. Ma, P. Chen, and J. Zhao, “Overview on vertical and directional drilling technologies for the exploration and exploitation of deep petroleum resources,” *Geomechanics and Geophysics for Geo-Energy and Geo-Resources*, vol. 2, no. 4, pp. 365–395, 2016.

[7] 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.

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

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

[10] J. D. Macpherson, J. P. de Wardt, F. Florence, C. D. Chapman, M. Zamora, M. L. Laing, and F. P. Iversen, “Drilling-systems automation: Current state, initiatives, and potential impact,” *SPE Drilling & Completion*, vol. 28, no. 4, pp. 296–308, 2013.

[11] W. A. Loomis, “Dummy driller,” U.S. Patent 1 779 480, Oct. 28, 1930.

[12] J. G. Abraham, “Drilling control,” U.S. Patent 2 759 702, Aug. 21, 1956.

[13] R. Teale, “The concept of specific energy in rock drilling,” *International Journal of Rock Mechanics and Mining Sciences & Geomechanics Abstracts*, vol. 2, no. 1, pp. 57–73, 1965.

[14] F. E. Dupriest and W. L. Koederitz, “Maximizing drill rates with real-time surveillance of mechanical specific energy,” in *SPE/IADC Drilling Conference, Amsterdam, The Netherlands*. SPE, 2005.

[15] Å. Kyllingstad and P. J. Nessjøen, “A new stick-slip prevention system,” in *SPE/IADC Drilling Conference and Exhibition, Amsterdam, The Netherlands*. SPE, 2009.

[16] J. Thorogood, W. Aldred, F. Florence, and F. Iversen, “Drilling automation: Technologies, terminology, and parallels with other industries,” *SPE Drilling & Completion*, vol. 25, no. 4, pp. 419–425, 2010.

[17] S. Naderi Lordejani, M. H. Abbasi, N. Velmurugan, C. Berg, J. Å. Stakvik, B. Besselink, L. Iapichino, F. Di Meglio, W. H. A. Schilders, and N. van de Wouw, “Modeling and numerical implementation of managed-pressure-drilling systems for the assessment of pressure-control systems,” *SPE Drilling & Completion*, vol. 35, no. 4, pp. 598–619, 2020.

[18] M. M. Huque, S. Imtiaz, A. Rahman, and M. Hossain, “Kick detection and remedial action in managed pressure drilling: a review,” *SN Applied Sciences*, vol. 2, no. 7, p. 1178, 2020.

[19] J.-M. Godhavn, “Control requirements for automatic managed pressure drilling system,” *SPE Drilling & Completion*, vol. 25, no. 3, pp. 336–345, 2010.

[20] R. Asgharzadeh Shishavan, C. Hubbell, H. Perez, J. Hedengren, and D. Pixton, “Combined rate of penetration and pressure regulation for drilling optimization by use of high-speed telemetry,” *SPE Drilling & Completion*, vol. 30, no. 1, pp. 17–26, 2015.

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

[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. 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.

[24] S. Skogestad, “Simple analytic rules for model reduction and PID controller tuning,” *Journal of Process Control*, vol. 13, no. 4, pp. 291–309, 2003.

[25] 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.

[26] 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.

[27] U.S. Bureau of Safety and Environmental Enforcement, “30 CFR 250.414, what must my drilling prognosis include?” Code of Federal Regulations, Title 30, Part 250, Subpart D, edition of July 1, 2025, govinfo.gov, 2025.

[28] ——, “30 CFR 250.427, what are the requirements for pressure integrity tests?” Code of Federal Regulations, Title 30, Part 250, Subpart D, edition of July 1, 2025, govinfo.gov, 2025.

[29] G.-O. Kaasa, Ø. N. Stamnes, O. M. Aamo, and L. S. Imsland, “Simplified hydraulics model used for intelligent estimation of downhole pressure for a managed-pressure-drilling control system,” *SPE Drilling & Completion*, vol. 27, no. 1, pp. 127–138, 2012.

[30] J. D. Hedengren, R. Asgharzadeh Shishavan, K. M. Powell, and T. F. Edgar, “Nonlinear modeling, estimation and predictive control in APMonitor,” *Computers & Chemical Engineering*, vol. 70, pp. 133–148, 2014.

[31] L. D. R. Beal, D. C. Hill, R. A. Martin, and J. D. Hedengren, “GEKKO optimization suite,” *Processes*, vol. 6, no. 8, p. 106, 2018.

2.10 Exercises

2.1 Concept check · 2 min

On a climbing execution in automatic mode, the supervisory layer also moves the bottomhole pressure setpoint. Which way does it move it, and what stops it?

(a) Up toward the fracture margin, until the high-pressure entry of v becomes the largest
(b) Nowhere, because the pressure setpoint belongs to the pressure loop and not to the supervisor
(c) Down toward the pore pressure itself, until an influx is flagged at the bit
(d) Down toward a floor set by the pore margin or by the wide-open choke, whichever is higher

2.2 Concept check · 2 min

Equation (2.7a) is a balance on the cuttings carried in the annulus. At steady state, why can no setting of the weight on bit raise the drilling rate above the cuttings ceiling at a given rotary speed and pump rate?

(a) The steady balance gives $C = ROP A_{bit}/(Q\eta)$, and η holds the speed and the flow but no weight
(b) The weight factor founders at W_f, so any weight above the founder point lowers the rate
(c) Added weight raises bottomhole pressure, and chip hold-down removes the gain it makes
(d) The weight loop saturates at its maximum feed rate before the ceiling can be reached

2.3 Concept check · 1 min

Which description of the two vibration constraints on the weight and rotary speed map is right?

(a) Both forbid a corner at high weight and low speed, so a higher rotary speed escapes both
(b) Stick-slip forbids a corner at high weight and low speed, bit bounce a band of speeds around a resonance
(c) Bit bounce forbids a corner at high weight, and stick-slip a band of speeds around 100 RPM
(d) Both forbid bands of rotary speed, whose width grows in proportion to the weight

2.4 Theory and calculation · 7 min

The default well has a 0.2159 m bit, 0.1270 m pipe, a cuttings transport length $L_c = 200\,\mathrm{m}$ and a pump rate $Q = 0.030\,\mathrm{m^3/s}$. Transport uses $\eta_{max} = 0.90$, $v_0 = 0.80\,\mathrm{m/s}$, $a_r = 0.60$ and $\omega_r = 8.0\,\mathrm{rad/s}$, and the cuttings limit is $C_{max} = 2.2\,\%$. Take $\bar{\omega}_b$ equal to the surface speed.

(a) Find the annular area and the annular velocity. *Checkpoint:* $0.023\,94\,\mathrm{m^2}$ and 1.253 m/s.

(b) Find the transport efficiency and the cuttings ceiling on the drilling rate at 120 RPM, in ft/hr. *Checkpoint:* 0.469 and 99.9 ft/hr.

(c) Find the time constant of the cuttings inventory at 120 RPM, and the ceiling at 250 RPM. *Checkpoint:* 340 s and 106.7 ft/hr.

2.5 Theory and calculation · 6 min

The choke adds $P_{bp,max}(1 - x/100)(Q/Q_{ref})^2$ with $P_{bp,max} = 4.0\,\mathrm{MPa}$ and $Q_{ref} = 0.030\,\mathrm{m^3/s}$. The supervisor walks the pressure setpoint at $\dot{P}_{sp} = 10\,\mathrm{kPa/s}$.

(a) Find the process gain $\partial P_b/\partial x$ in psi per percent at the sizing flow, and the whole authority of the choke at the $0.011\,\mathrm{m^3/s}$ of the pack-off scenario. *Checkpoint:* −5.80 psi per percent; 78 psi.

(b) Find the velocity constant $K_v = K_c|K_p|/\tau_I$ and the standing offset $\dot{P}_{sp}/K_v$ behind the ramp for the default tuning ($K_c = 5 \times 10^{-5}\,\%/\mathrm{Pa}$, $\tau_I = 4\,\mathrm{s}$) and for the challenge as delivered ($K_c = 8 \times 10^{-6}\,\%/\mathrm{Pa}$, $\tau_I = 30\,\mathrm{s}$). Compare both with the 17 psi band of the challenge. *Checkpoint:* 0.5/s and 2.9 psi; 0.0107/s and 136 psi.

2.6 Theory and calculation · 7 min

The string uses $R_b = 0.10\,\mathrm{m}$, $\mu_s = 0.80$, $\mu_d = 0.25$, $\omega_c = 2.80\,\mathrm{rad/s}$ and $c = c_b = 30\,\mathrm{N\,m\,s/rad}$, with the onset $G_0 = 1.092$. The axial mode has $f_n = 5.00\,\mathrm{Hz}$, $\zeta_a = 0.060$, $n_l = 3$, $\varepsilon_0 = 0.100$ and $W_{ref} = 111.2\,\mathrm{kN}$. One klbf is 4.448 kN.

(a) Evaluate G of Equation (2.10) at the scenario-3 point, 40 klbf and 60 RPM. *Checkpoint:* $G = 6.18$.

(b) Find the rotary speed above which steady rotation is stable at 40 klbf. *Checkpoint:* 106.3 RPM.

(c) Find the rotary speed of the axial resonance and the bounce index there at 25 klbf, and compare it with the limit $B_{max} = 0.50$. *Checkpoint:* 100 RPM and 0.833.

2.7 App challenge · 8 min

Load scenario 6, The MPD lever. Weight and rotary speed are held at 30 klbf and 150 RPM while the choke opens in steps from 10 % to wide open. Change no setting, use the default seed, run for at least 430 s at speed 20 and export the CSV file.

(a) Read `bhp_psi` and `rop_ftphr` in the rows at 90 and 430 s, and find the ratio of the two drilling rates. *Checkpoint:* 5678 psi and 69.3 ft/hr; 5177 psi and 97.3 ft/hr; ratio 1.404.

(b) Predict the ratio from chip hold-down alone, with $\lambda = 0.12/\text{MPa}$. *Checkpoint:* 1.513.

(c) Read `cuttings_frac` and `bit_wear` in the same rows and account for the difference with Equations (2.3b) and (2.3d) ($h_c = 6$, $C_{50} = 0.050$). *Checkpoint:* Factors 0.968 and 0.959; product 1.404.

Accept the pressures within 1 psi and the rates within 0.1 ft/hr.

2.8 App challenge · 8 min

Load scenario 10, Optimizer against a fixed manual point, which starts in Auto from the settings of scenario 1. Change no setting, use the default seed, run 900 s at speed 20 and export the CSV file. Then load scenario 1, Baseline sandstone, and run it 900 s unchanged.

(a) Read the Hole drilled and ROP vs best legal tiles of both runs, and the binding constraint of scenario 10 at the end (the limit badge, or the column `opt_binding`). *Checkpoint:* 28.3 ft and 97 % against 16.3 ft and 61 %; cuttings.

(b) In the file of scenario 10, find when the rotary speed setpoint (`opt_omega_sp_radps`) reaches 250 RPM, when the weight setpoint (`opt_wob_sp_n`) first reads 35.0 klbf to one decimal, and the lowest pressure setpoint (`opt_bhp_sp_pa`) with its time. *Checkpoint:* 52 s; about 58 s; 5193 psi at 110 s.

(c) Find the interval in which `rop_mps` first exceeds `rop_best_mps`, and explain why that is not an error. *Checkpoint:* From 30 to 427 s.

2.9 App challenge · 9 min

Load the challenge, Drill 5 m in under 15 minutes. Use the default seed and speed 20, and export the CSV file of each run. The worst pressure error is the largest $|\texttt{bhp_pa} - \texttt{bhp_sp_pa}|$ in the file.

(a) Run the case as delivered until it stops itself. Read the Hole drilled tile, the worst pressure error and the flag columns that read 1 at the start. *Checkpoint:* 8.0 ft at 900 s; 49.1 psi; `flag_packoff` from 0 s and `flag_bounce` from 0.15 s.

(b) Press the reload button, press Reveal suggested tuning before Run, and run until the case stops itself. Read the elapsed time, the worst pressure error, the Peak stick-slip and Min window margin tiles, and whether any flag column reads 1. *Checkpoint:* 473 s; 4.9 psi; 0.16 and 263 psi; no flag.

(c) Press the reload button, press Reveal suggested tuning, set Cuttings loading limit back to 2.2 %, and run. Find when the first flag sets, and say whether the attempt passes. *Checkpoint:* `flag_packoff` at 297 s; it fails.

Accept the pressure errors within 0.1 psi.

2.10 Application area · 4 min

Scenario 5, Narrow pressure window, drills a depleted sand at a true vertical depth of 3000 m with a pore pressure of 31.0 MPa, a fracture pressure of 33.2 MPa and a mud of 1000 kg/m^3. The margin is 0.7 MPa at each end of the window.

(a) Find the hydrostatic head of the mud column, in psi, and compare it with the pore pressure. *Checkpoint:* 4268 psi against 4496 psi, 228 psi short.

(b) Find the width of the window and the part left after the two margins. *Checkpoint:* 319 psi, of which 116 psi is left.

(c) The pumps stop for a pipe connection. Say what happens to the bottomhole pressure, what a rig does about it, and whether the app can show the consequence.

2.11 Application area · 4 min

Scenario 9, Worn bit, is scenario 1 with the bit worn to $h = 6$ of 8. In 900 s it makes 6.8 ft against 16.3 ft for a sharp bit, and the MSE tile ends at 72.7 kpsi against 30.4 kpsi.

(a) Compare the ratio of the holes with the wear factor $f_h = \mathrm{e}^{-h/h_c}$, $h_c = 6$, and explain the difference. *Checkpoint:* 0.417 against 0.368.

(b) Compare the ratio of the specific energies with the inverse ratio of the holes, and explain the match from Equation (2.15). *Checkpoint:* 2.39 against 2.40.

(c) A driller sees the specific energy more than double at constant weight and speed. Say what else separates a dull bit from a hard stringer.

2.12 Review the agent · 5 min

An AI assistant is asked how to drill faster in scenario 1 and replies as follows. “Chip hold-down multiplies the whole drilling-rate surface. Opening the choke from 65 % to wide open lifts the unconstrained peak from 111.2 to 124.1 ft/hr, a gain of 11.6 %. The best feasible point, 106.7 ft/hr, sits on the same surface, so it rises by the same share to about 119 ft/hr. Open the choke fully and let the optimizer climb to it.” Find the flaw and say how you would check it.