INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 9 Model Predictive Control and Economic Optimization

CHAPTER 7

Reservoir Flood Control on an Ensemble Forecast With Scenario Trees

Pages 149–171

Edited by John D. Hedengren

How to cite this chapter. "Reservoir Flood Control on an Ensemble Forecast With Scenario Trees," in *Model Predictive Control and Economic Optimization*, John D. Hedengren, Ed., vol. 9 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 7, pp. 149–171.

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

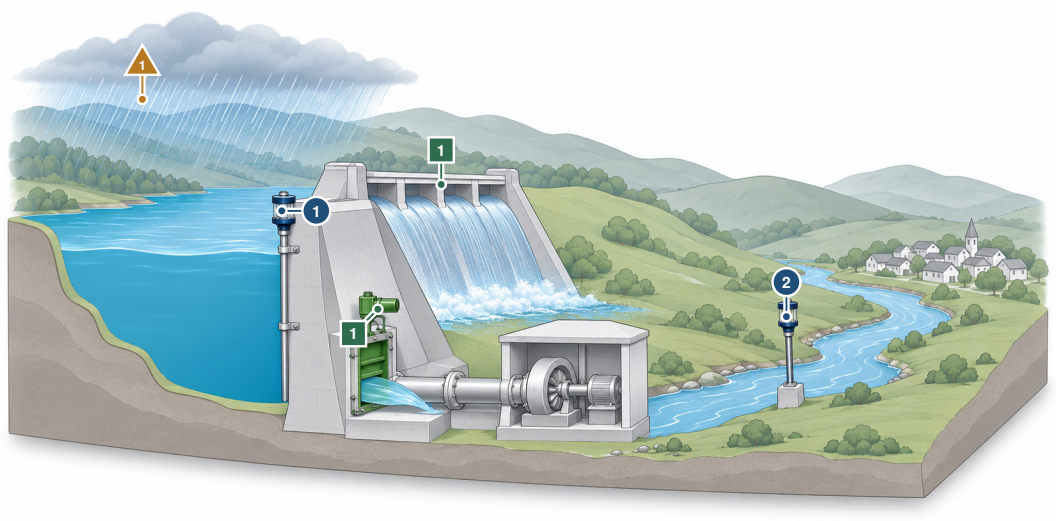
The app for this chapter
https://apmonitor.com/apps/reservoir-flood-control/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v09-mpc-optimization/#reservoir-flood-control

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

CHAPTER 7

Reservoir Flood Control on an Ensemble Forecast With Scenario Trees



The controller moves 1 the release, through the outlets and the spillway gates

The controller reads 1 lake level, read as storage 2 river gauge at the town

What disturbs it 1 storm rain on the hills above the lake

Water can leave the lake only as fast as the river below can carry it, so room for a storm must be made days before the rain. Illustration, not to scale.

IN PLAIN TERMS Making Room in a Lake Before the Rain

The main idea. A dam above a town keeps part of its lake empty in winter to catch a flood. Water can leave the lake only as fast as the river below can carry it. So the room must be made days before the storm.

The challenge. The operator decides on a forecast that gives several possible futures. Sometimes the promised storm never falls. Planning for the wettest future wastes water, and planning for a milder one can leave too little room.

Why it matters. Every dam that guards a town must choose how much flood risk to accept for the water and power it keeps.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the history in Section 7.2, the try-it segments in Sections 7.4 and 7.6, the two facts, and the key points
Study	50 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the results (Section 7.6), Technology in Practice (Section 7.7), Limitations (Section 7.8) and the open challenges

The control question A reservoir empties its flood pool only as fast as the channel below the dam carries water, so it must release days before a storm, on a forecast that sometimes promises a storm that never comes. What do a rule curve, a value-function policy and a scenario tree each cost and buy in flood safety, supply and energy?

Learning objectives

- Compute how fast a channel-limited release empties a flood pool, and the lead that rate demands.
- Explain why a tributary below the dam floods the town before the pool is full.
- Describe what a scenario tree commits today and how its chance constraint lets the wettest member go.
- Compare three deciders over a water year in days over the limit, supply and energy.
- Judge an energy result against a perfect-foresight bound at the run's own end storage.

Prerequisites Sections 1.4.6 (slack prices), 1.4.8 (terminal value), 1.4.9 (dynamic programming), 1.4.12 (scenario tree, chance constraint) and 1.6.4, and Chapter 5 for releasing ahead of a delay.

The app https://apmonitor.com/apps/reservoir-flood-control/

See also Volume 5, Chapter 4 (every useful move is a prediction), Volume 5, Chapter 5 (a quality fixed before it is measured, and a supervisor that says when no move can meet the limit) and Volume 5, Chapter 9 (a slow inventory refilled through a small, capped valve).

7.1 Introduction

On 13 February of the simulated water year a 440 mm atmospheric river begins to fall on the 2500 km^2 catchment above an 1100 hm^3 reservoir. Operated on its rule curve, the lake meets the storm at 798 hm^3, two cubic hectometers below its winter target, rises above its crest to 1137 hm^3, spills 111 hm^3 over it, and on 21 February the river at the town downstream breaks a levee declared to fail above 650 m^3/s with 666 m^3/s. On the same weather a chance-constrained scenario tree that reads the morning ensemble raises the release from 12 to 440 m^3/s between 8 and 12 February, meets the storm at 688 hm^3, fills to 1095 hm^3 without touching the crest, and never lets the town's river exceed 425 m^3/s against a channel limit of 450 m^3/s. What separates the two in February is what each does in the five days before the rain.

Reservoir operation has been posed as an optimization problem since the 1950s [1], and its methods have been reviewed repeatedly [2–4]. The control problem underneath it is general. An inventory receives an uncertain inflow, a downstream unit accepts only a limited rate, and the only defense against a large inflow is room made before it arrives. A stormwater basin, a surge tank ahead of a treatment plant, a buffer between a variable feed and a capacity-limited unit and a battery dispatched against a solar forecast share the shape. Volume 17, Chapter 10 dispatches such a battery against a forecast of the net load, and Volume 17, Chapter 11 plans the inventory of an air store with a water value. Chemical process control met the same structure in chance-constrained predictive control [5] and in multi-stage predictive control of a polymerization reactor on a tree of uncertain futures [6], both instances of multi-stage stochastic programming [7].

This chapter tests two beliefs. The first is that a reservoir floods its town because its flood pool is full. In case 2 the town's river first exceeds the channel limit on day 17.4 with the lake at 866 hm^3, 234 hm^3 below the crest, because the tributary between the dam and the town rises first and the dam's release reaches the town a day after it is made, and a flood pool enlarged from 300 to 400 hm^3 leaves the town's peak at the same 575 m^3/s. The second is that the cautious use of a forecast is to plan for its wettest member. In case 4 a storm that the forecasts promise from day 1 for day 10 never falls, and a rule that chases the wettest member sends 23 hm^3 through the gates for nothing, against 2 hm^3 for a rule that hedges on the median member.

These ideas usually meet a student in different courses. Stochastic dynamic programming is taught with water resources planning [4, 8], ensemble forecasts with hydrology [9], and scenario trees and chance constraints with optimization [7, 10]. This chapter runs a rule curve, a stochastic dynamic

program and a chance-constrained scenario tree on one declared reservoir and one ensemble forecast in a browser simulator, and measures what each one costs and buys in flood safety, supply and energy. The value function [1, 8], tree-based predictive control on an ensemble [11, 12] and the discarding of scenarios for a chance constraint [10] are established methods. The measurements are made on one declared plant and one weather year, and among them is the finding that the graded water year is passed by forecasting and not by emptying the lake. The simulator underneath has a 19-state model whose water balance closes to 10^{-10} hm^3, a value function on 73 000 states solved in under a second in a browser and checked against an independent Bellman backup, a tree linear program checked against brute force, and a measurement boundary that keeps every decider away from the true future. Ten cases carry the measurements, each with a completion rule in the app (Section 1.6.4), and the exercises at the end of the chapter are built on them.

Section 7.2 follows the domain from the operating diagram to ensemble forecasts. Section 7.3 declares the basin from the rain to the town, and Section 7.4 shows why the pool must be emptied before the storm. Section 7.5 sets out the deciders: the rule curve with its forecast pre-release, the stochastic dynamic program, the scenario tree with its chance constraint and fallbacks, and a perfect-foresight bound. Section 7.6 measures the ten cases, the last of them a graded water year, and verifies the code against its own equations and against published hydrology. Section 7.7 describes the gauges, gates and forecasts of real operation, Section 7.8 states what the model leaves out, and Section 7.9 draws the results together. The scenario tree, its shared first decisions and the chance constraint are introduced on the heated tank of Section 1.4.12, and the value function and its slope in Section 1.4.9.

7.2 Flood Pools, Value Functions and Ensemble Forecasts

Multipurpose dams have long been run from an operating diagram drawn when the dam was designed. A rule curve gives the target storage for each day of the year, the space between the target and the top of the dam is reserved for floods in the wet season and returned to supply and power in the dry season, and a flood rule releases as much as the channel downstream can carry whenever the lake rises into that reserve [4]. The diagram needs no forecast and is easy to audit.

> ⧖ HISTORY **1959: A Guide Curve Drawn Without Forecasts**
>
> At one flood-control reservoir in California, the water control manual was issued in 1959, with minor revisions in 1986, without the benefit of modern weather and streamflow forecasts. Its guide curve sets the storage during a flood and the safe releases soon after it that make room for the next one, and it rests on the historical pattern of wet winters and dry summers, not on the weather in sight. Decades later a group of agencies tested releases recommended from a 15-day ensemble streamflow forecast, as approved deviations from that manual. In the 2020 water year, the third driest in a 127-year record, the lake ended the winter with 19 percent more water stored [13]. This chapter's rule curve plays the part of such a diagram, and its trees play the forecast.

The optimization of reservoir releases under uncertain inflow has early roots in hydroelectric storage. In 1955 Little treated the use of storage water in a hydroelectric system as a sequence of decisions under a random inflow [1]. Stochastic dynamic programs with inflow classes in a Markov chain were then set out as reservoir operation models [8], and reviews in 1985 and 2004 record how linear, dynamic and nonlinear programming and simulation were applied to single reservoirs and to systems of them [2, 3]. The control of a water reservoir under economic, social and environmental constraints was set out for the automatic control community in 2008 [14].

Forecasts changed what the optimizer could know. Ensemble streamflow prediction conditioned a sampling stochastic dynamic program on the forecast traces [15], and flood forecasting driven by ensembles from numerical weather prediction was reviewed in 2009 [9]. Short-term operation then moved to receding-horizon optimization over a tree of the ensemble members, with the first decisions shared by every member [11], and real-time flood control of a reservoir by tree-based predictive control has been reported with forecast uncertainty included [12]. The chance constraint

Table 7.1 Variables by Role, With Their Values on 17 February of the Challenge on the Suggested Settings

Variable	Role	Value	Range
Q_{town}	CV, true flow at the town	418 m^3/s	limit 450, levee 650
S	CV, storage, measured exactly	824 hm^3	125–1100
R	MV, release through the outlets	124 m^3/s	0–1780 m^3/s
Q_t, Q_g	split, turbines then gates	124, 0 m^3/s	$Q_t \leq 180$
Q_c	output, crest spill	0 m^3/s	above 1100 hm^3
P, P_L	DV, rain (teaching truth)	8.2 mm/d	storms declared
I	DV, inflow to the lake	1087 m^3/s	gauge ±2 %: 1064
L	DV, local tributary	348 m^3/s	teaching truth
$\hat{I}_{m,d}$, $\hat{L}_{m,d}$	measured, ensemble	$M = 7$, 10 days	issued daily
$T(d)$	SP, rule curve target	800 hm^3	800–1000
D, D_s	supply diverted, scheduled	6, 6 m^3/s	6–28 m^3/s
I_r	river below the diversion	118 m^3/s	$\geq Q_{env} = 6$
E	evaporation	0.03 hm^3/d	—
H, P_e	head, power	88 m, 94 MW	to 100 m, 156 MW
$n_{1..3}$, g	states, catchment stores	—	mm
W	state, water in the reach	—	$m^3/s\,d$

arrived from process control, where it states the probability with which an output limit must hold [5], and from optimization, where discarding a declared number of sampled scenarios gives a tractable approximation [10]. The hydrology underneath is older. Nash derived the unit hydrograph of a catchment by routing rain through a cascade of equal linear reservoirs [16], and the Muskingum storage law for a river reach was related to the physics of flood propagation by Cunge [17]. The simulator places the rule curve, the value function and the tree on one basin built from those two models.

> ✦ **DID YOU KNOW? Sixty-One Members on an Hourly Step**
>
> The streamflow ensembles prepared for one reservoir's forecast-based operation run 15 days ahead on an hourly step, and the hindcast used to test them carries 61 members [13]. This chapter's forecast issues $M = 7$ members for ten days, one value a day (Table 7.1).

7.3 A Declared Basin From Rain to the Town

The simulated basin is a declared teaching design with no real river, dam or operator behind it, as shown in Figure 7.1. Rain falls on a catchment of $A = 2500\,km^2$ above the dam and on a local catchment of $A_L = 600\,km^2$ between the dam and the town. The lake holds 1100 hm^3 at its uncontrolled crest and 120 hm^3 below its intake. The dam releases through a 180 m^3/s powerhouse and a gated spillway, a diversion below the powerhouse takes the supply schedule, an environmental flow stays in the river, and a one-day reach carries the rest to a town whose channel carries 450 m^3/s and whose levee fails above 650 m^3/s. The operator decides one release a day. The time base is the day, volumes are in cubic hectometers (1 hm^3 is 10^6 m^3), flows in m^3/s, rain in mm/d, power in MW, and 1 m^3/s held for a day is 0.0864 hm^3.

Each variable carries one role in Table 7.1. Nineteen states are integrated: the eight cascade and groundwater stores of the two catchments, the storage S, the water W in the reach, the energy generated, and eight running volumes that exist so the water balances can be checked exactly.

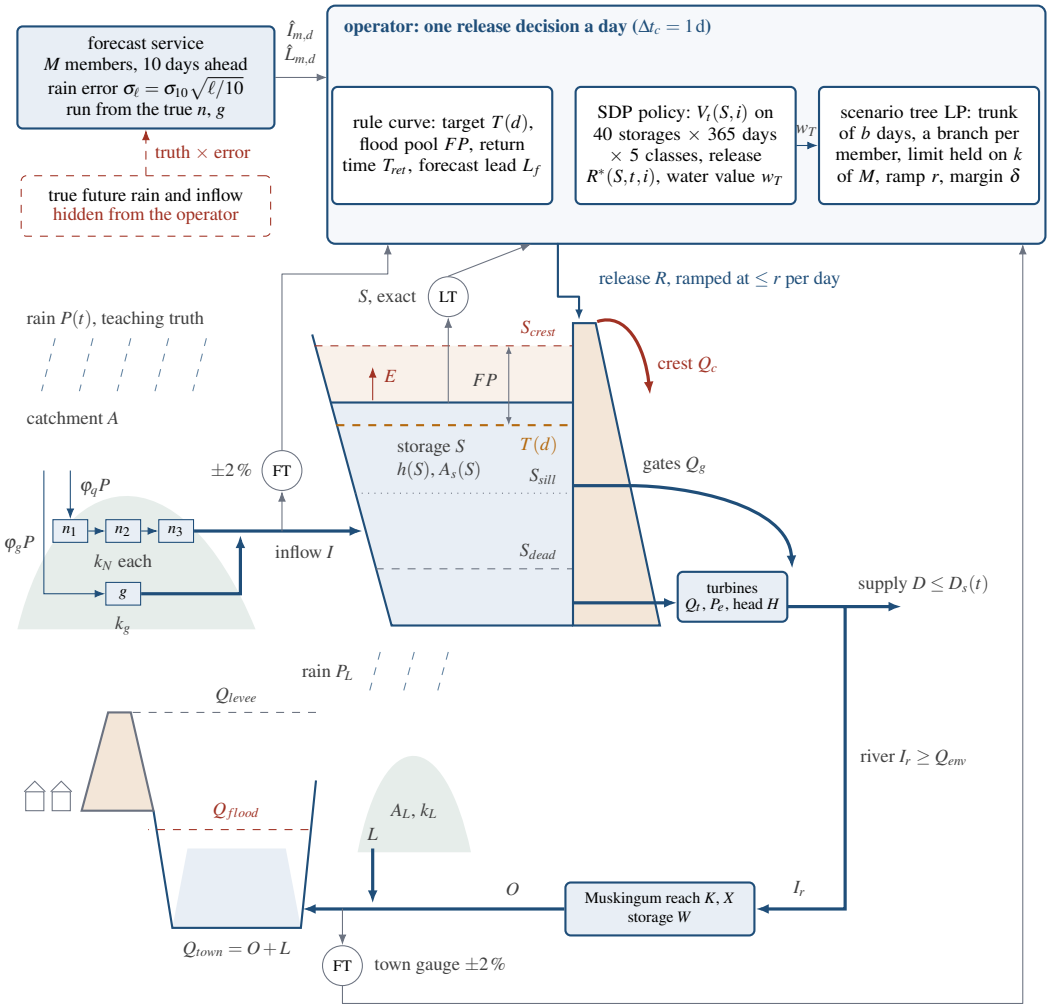


Figure 7.1 Reservoir, river and town as modeled, inside the daily decision loop, after the application's Basin view, with the river below the dam turned back so that the town sits at the lower left. Rain P (teaching truth) fills the cascade stores n_1 to n_3 (constant k_N) and the groundwater store g (k_g) of the catchment A, giving the inflow I. The lake holds storage S at depth $h(S)$ and area $A_s(S)$ between the dead pool S_{dead} and the crest S_{crest}, loses E to evaporation, and keeps the flood pool FP above the rule curve $T(d)$. The turbines pass Q_t at head H for power P_e, the gates Q_g above the sill S_{sill}, and the crest Q_c. The diversion takes D, the river keeps Q_{env}, the reach (K, X, storage W) gives O, and the tributary (A_L, k_L) adds L to make Q_{town} against Q_{flood} and Q_{levee}. The operator reads S exactly, the two gauges to ±2 % and the ensemble, never the true future. Not to scale.

7.3.1 Rain and Runoff

Each catchment is a cascade of three equal linear reservoirs of quick runoff beside a groundwater store that feeds the base flow, as shown in Equation (7.1).

$$\begin{aligned} \frac{dn_1}{dt} &= \varphi_q P - \frac{n_1}{k_N}, & \frac{dn_i}{dt} &= \frac{n_{i-1} - n_i}{k_N} \quad (i = 2,3) \\ \frac{dg}{dt} &= \varphi_g P - \frac{g}{k_g}, & I &= \frac{A}{86.4}\left(\frac{n_3}{k_N} + \frac{g}{k_g}\right) \end{aligned} \tag{7.1}$$

where the stores are in mm, P is in mm/d, $A/86.4$ converts mm/d over A in km^2 to m^3/s, and the local catchment has the same form with P_L, A_L, $k_L = 1.0\,\mathrm{d}$ and outflow L. Of the rain, $\varphi_q = 0.40$ runs off quickly and $\varphi_g = 0.25$ recharges the groundwater, which drains with $k_g = 40\,\mathrm{d}$. The response of the cascade to a unit impulse of quick runoff is Nash's gamma-shaped unit hydrograph, Equation (7.2) [16].

$$u(t) = \frac{1}{k_N \Gamma(n)} \left(\frac{t}{k_N}\right)^{n-1} \mathrm{e}^{-t/k_N} \tag{7.2}$$

With $n = 3$ and $k_N = 1.5\,\mathrm{d}$ it peaks at $(n-1)k_N = 3\,\mathrm{d}$ with a mean lag of $nk_N = 4.5\,\mathrm{d}$, while the smaller local catchment peaks at 2 d. The tributary below the dam therefore rises a day before the reservoir's own inflow on the same storm.

7.3.2 The Lake, Its Outlets and Its Crest

The storage integrates the inflow less the three outflows and the lake's evaporation, the depth at the dam follows a valley's storage curve, and the turbines produce power on the head above a fixed tailwater, as shown in Equation (7.3).

$$\begin{aligned} \frac{\mathrm{d}S}{\mathrm{d}t} &= 0.0864\,(I - Q_t - Q_g - Q_c) - E(S,t) \\ h &= 6.69\,S^{0.4}, \qquad H = h - 10, \qquad A_s = \frac{S}{0.4\,h} \\ E &= 10^{-3} e_{m(t)} A_s, \qquad P_e = 9.81 \times 10^{-3}\,\eta\, Q_t H \end{aligned} \tag{7.3}$$

where h and H are in m, A_s in km^2, $e_{m(t)}$ is the month's evaporation of 0.8 to 7.0 mm/d, $\eta = 0.88$, the factor 9.81×10^{-3} is the weight of water in MN per m^3, and P_e is in MW. At the crest the head is 100.1 m and the full powerhouse makes 155.6 MW. The release R follows the operator's command as a straight line in time, from the release at the last decision toward the new target at the declared rate, and is capped by what the outlets pass at the present storage, as shown in Equation (7.4).

$$\begin{aligned} R &= \min\left(R_{cmd}(t),\ R_{cap}(S)\right) \\ Q_t &= \min(R, Q_{t,max}), \qquad Q_g = R - Q_t \\ R_{cap} &= c(S)\left[Q_{t,max} + \gamma \min\left(Q_{s,full},\ a\,(S - S_{sill})^{+}\right)\right] \\ Q_c &= C_w \left(h(S) - h(S_{crest})\right)^{3/2} \quad (S > S_{crest}) \end{aligned} \tag{7.4}$$

where $a = Q_{s,full}/(S_{crest} - S_{sill}) = 2.29\,\mathrm{m^3/s}$ per hm^3 makes the gated rating linear from the sill at 400 hm^3 to 1600 m^3/s at the crest, γ is 1 or 0.5 after a gate fault, $c(S)$ falls from 1 to 0 over the 5 hm^3 above the dead pool as the intake uncovers, and $C_w = 255\,\mathrm{m^{1.5}/s}$ is a 150 m weir. The turbines take the release first and the gates the rest, and the crest spills on its own.

7.3.3 Supply, the Reach and the Town

The diversion takes the month's scheduled supply D_s (6 to 28 m^3/s, 462 hm^3 a year) from the dam's outflow above the environmental flow, and the river below it is routed by Muskingum's storage law to the town, as shown in Equation (7.5).

$$\begin{aligned} D &= \min\left(D_s,\ \max(0,\ R + Q_c - Q_{env})\right) \\ I_r &= R + Q_c - D \\ W &= K\,[X I_r + (1 - X)\,O], \qquad \frac{\mathrm{d}W}{\mathrm{d}t} = I_r - O \\ Q_{town} &= O + L \end{aligned} \tag{7.5}$$

where $K = 1\,\mathrm{d}$ and $X = 0.2$. Solved for O, the law gives $\mathrm{d}W/\mathrm{d}t = (I_r - W/K)/(1-X)$ and $O = (W/K - X I_r)/(1-X)$, integrated at the physics step. A sudden rise in I_r makes O dip by $X/(1-X) =$

0.25 of the step before it rises, a property of the law itself, which the ramped release keeps small. The daily form of the same law has coefficients 0.231, 0.538 and 0.231, which are all positive because $2KX = 0.4 \le \Delta t = 1 \le 2K(1-X) = 1.6$, the range of the step discussed for the method by Cunge [17].

7.3.4 Weather and the Forecast Service

A case's rain is drawn when it loads. A colored-noise index x with a two-day correlation becomes a day's rain on the season's mean, and the local catchment uses 0.9 of the same index plus its own noise, as shown in Equation (7.6).

$$P(d) = \bar{P}(d)\, w\, \frac{\max(0,\, x_d - \theta)^2}{\kappa(\theta)} \tag{7.6}$$

where $\bar{P}$ is interpolated from monthly means of 70, 130, 170, 180, 150, 120, 70, 35, 12, 4, 4 and 20 mm (October first, 965 mm a year), w is a case's wetness, $\theta = 0.3$, and $\kappa(\theta) = 0.30206$ is the mean of $\max(0, x-\theta)^2$ for a unit normal, so the long-run mean is $\bar{P}w$. Storms are declared by each case as three days carrying 30, 45 and 25 % of their total, on the local catchment at 0.9 of that unless the case says otherwise. Over 2000 generated years the mean is 966 mm, and seed 365, which every case uses, gives 919 mm with no natural inflow above 249 m^3/s.

Each morning the forecast service issues M members for the next ten days. Member m scales the true rain of lead day ℓ by a mean-one lognormal factor whose spread grows with the lead, and runs the declared runoff model from the catchments' true state that morning, as shown in Equation (7.7).

$$\begin{aligned} \hat{P}_{m,\ell} &= P(d+\ell-1)\, \exp\!\left(\sigma_\ell z_{m,\ell} - \sigma_\ell^2/2\right) \\ \sigma_\ell &= \sigma_{10}\sqrt{\ell/10} \\ z_{m,\ell} &= 0.7\, z_{m,\ell-1} + \sqrt{0.51}\, \varepsilon_{m,\ell} \end{aligned} \tag{7.7}$$

where ε is a unit normal deviate and $\sigma_{10} = 0.6$. The members' daily inflows $\hat{I}_{m,d}$ and local flows $\hat{L}_{m,d}$ are the ensemble. A forecast bust is a declared storm added to every member, scaled by 0.15, 0.4, 0.7, 0.9, 1.1, 1.35 and 1.7 for the seven members, on every morning before its reveal day, and absent from the truth. Over 355 mornings of seed 365 the relative mean error of the ensemble mean is 0.0000, 0.0037 and 0.0002 at leads of 1, 5 and 10 days, and the members' relative spread grows from 0.002 to 0.060 and 0.151.

7.3.5 What the Operator Measures

The deciders read one packet each morning: the storage exactly, the inflow and town gauges with a colored error of standard deviation 2 % and correlation time 0.5 d, the dam's own exact records of its release, outflow, diversion and river flow, the gates' status, the public supply schedule and the ensemble. The rule estimates the tributary as the town gauge less the output of its own Muskingum model run on its own record of I_r. The rain, the true inflow ahead and the tributary are never in the packet, and a check in the application's own test suite rejects any version of the controller, the dynamic program or the tree that reads anything else. The exported file writes each gauge beside the true flow it reads, under a `truth_` prefix that keeps a replay from using the true values.

7.3.6 Constants, Starting State and Integration

Table 7.2 lists the constants. Every value is a declared teaching value chosen so that the basin is consistent, not one measured on a real river. The channel limit stands for a reach's rated capacity from a hydraulic study, the levee's failure flow for a levee's design flow, the environmental flow for the instream requirement of a dam's license or permit, the dead pool for an intake's invert, and the ramp limit for a permit's ramping rate, and the application's model dialog says so. Each case starts at its declared day of the water year, storage and release, with the catchment stores run through a

Table 7.2 Constants. B and A Mark Controls in the Beginner and Advanced Panels, F Marks Fixed Plant Constants

Symbol	Meaning	Value	Unit	
A, A_L	catchments, above the dam and local	2500, 600	km^2	F
φ_q, φ_g	quick and groundwater shares	0.40, 0.25	—	F
k_N, k_L	cascade constants, above and local	1.5, 1.0	d	A
k_g	groundwater recession	40	d	F
S_{dead}, S_{sill}	intake, gate sill	120, 400	hm^3	F
S_{crest}	uncontrolled crest	1100	hm^3	F
$Q_{t,max}, \eta$	powerhouse, efficiency	180, 0.88	m^3/s, —	F
$Q_{s,full}$	gated spillway at the crest	1600	m^3/s	F
C_w	crest weir	255	$m^{1.5}/s$	F
Q_{flood}	channel limit at the town	450	m^3/s	B
Q_{levee}	levee failure flow	650	m^3/s	A
Q_{env}	environmental flow	6	m^3/s	B
K, X	Muskingum reach	1.0, 0.2	d, —	A
FP	winter flood pool of the rule	300	hm^3	B
T_{ret}	return time of the rule	10	d	B
r	ramp limit on the release	150	m^3/s per d	B
δ	planning margin below the limit	7	%	A
L_f	pre-release lead of the rule	0	d	B
M, σ_{10}	members, rain error at 10 days	7, 0.6	—	A
b	tree trunk	2	d	A
k	members held	6	—	B
c_F	flood damage over the limit	5×10^4	$/($m^3$/s d)	A
c_S, c_E	supply, environmental shortfall	0.10, 0.20	$/$m^3$	A
p_t	energy value by month	34–56	$/MWh	A
—	SDP grid	40×365×5	—	A
—	gauge error, correlation time	2, 0.5	%, d	F

year of the previous year's weather and the reach in steady flow at the starting release. The engine of Section 1.6.1 integrates by fourth-order Runge-Kutta on a step of 0.1 d, the decision is made every 1 d, and a row is logged every 0.2 d. Rain is constant within a day, so the derivatives read the day from the start of the step and stay pure.

7.3.7 Assumptions

The runoff model is linear, so a fixed share of rain runs off whatever the soil's wetness, and the zero mean error of the forecast relies on that linearity. Rain is uniform over each catchment and constant within a day. The lake is one well-mixed storage with no wind setup, and the reach is one Muskingum element with no backwater or floodplain storage. The forecast service knows the catchments' true state and errs only in the rain, while real ensembles also err in timing and in the initial state. The levee's failure latches a stop, and the breach is not routed. The water stage drawn at the town is a schematic stage, linear in the true flow between the channel's bank-full line and the levee's crest, not a rating curve. The powerhouse is a static map from flow and head to power: the water column, the turbine and its governor are left out, and Volume 1, Chapter 7 treats them.

7.4 Why the Pool Must Be Emptied Before the Storm

Three facts make the problem one of timing. The first is a rate. With the planning margin the dam may send $(1-\delta)Q_{flood} - L + D_s$ toward the town, and at about 400 m^3/s the channel takes 34.6 hm^3 a day, so emptying a 300 hm^3 flood pool with no inflow at all takes 8.7 days. The second is an order of arrival. The local catchment's unit hydrograph peaks a day before the reservoir's inflow, and a release reaches the town a day after it is made because the reach's centroid lag equals K. By the time the tributary is high the channel has no room left for the dam. The third is the forecast. A storm seen ten days ahead is seen through a rain error of $\sigma_{10} = 0.6$ on its logarithm, which in case 7's setting of $\sigma_{10} = 0.9$ puts the wettest member's ten-day rain at 2.01 times the median member's.

The cases make each fact visible. On run of river in case 1 the lake rises only 31 hm^3 through a 300 mm storm and the town receives the dam's outflow plus the tributary, so the levee fails on day 16.8 at 660 m^3/s, while the rule curve holds the same storm with 137 hm^3 of its pool and a town peak of 444 m^3/s. In case 3 the rule is told the forecast, counting the storm's volume above the channel as if it were already in the lake. With two days of lead the town peaks at 552 m^3/s, with three at 525, with four at 489, and with five at 425 m^3/s, under the limit, so the second storm of a wet December needs five days of warning. A dam can be operated well only by acting on water that has not yet arrived. Chapter 5 meets the same need on a canal, where each gate releases a published order ahead of the demand by the travel time of the water wave.

> **TRY IT IN THE APP** **1 · Run of River**
>
> Load the scenario and press Run. Watch the town's river against the channel's limit on the top trend: on run of river the levee fails on day 16.8 at 660 m^3/s, and the run pauses there. Then choose Rule curve as the Strategy (Auto), press Reset and run again. The town now peaks at 444 m^3/s (Table 7.4).

7.5 Deciding the Morning's Release

The decision is made once a day, at $t = 0, 1, 2, \ldots$ days, at the start of the physics step, and handed to the plant as a ramp. The rules and the dynamic program move the release toward their target at the full ramp rate r, and the tree's plan is a straight line to its first boundary value, which its own ramp rows keep within r. In Manual the release station sets the target.

7.5.1 The Rule Curve

The rule curve's target $T(d)$ is 1000 hm^3 from May to September, falls linearly to $S_{crest} - FP$ by 15 November, holds there to the end of February and refills through March and April. The rule returns the lake to the curve over T_{ret} days and switches to the channel's limit whenever the lake is more than 5 hm^3 into its flood pool, as shown in Equation (7.8).

$$\begin{aligned} R &= R_{ch} \quad \text{if } S + V_f - T(d) > 5\,\text{hm}^3 \\ R &= \min\left(I_m + \frac{S - T(d)}{0.0864\,T_{ret}},\ R_{ch}\right) \quad \text{otherwise} \\ R_{ch} &= (1-\delta)\,Q_{flood} - \hat{L} + D_s \end{aligned} \tag{7.8}$$

where I_m is the inflow gauge and $\hat{L}$ the tributary estimate. While the lake is more than 30 hm^3 above the dead pool, R is held at least at $D_s + Q_{env}$ unless R_{ch} is lower. Without a forecast $V_f = 0$. With a pre-release lead L_f, $V_f = \sum_{\ell=1}^{L_f} 0.0864 \max_m \left(\hat{I}_{m,\ell} - R_{ch,m,\ell}\right)^+$ counts, day by day, the largest excess among the members of the forecast inflow over what the channel lets the dam release with that member's tributary (or the median excess, to hedge), and the rule's tributary is then the largest (or median) forecast local flow over the next three days. A margin of 5 % let the rule touch 453 m^3/s while emptying case 1's pool, which is why the margin is 7 %.

7.5.2 Stochastic Dynamic Programming

The dynamic program turns the year into a value function on the storage, the day and the inflow class. Each day the inflow falls in one of five classes of its half-month, split at the 35th, 65th, 85th and 96th percentiles so that the top class holds the storm days, and the next day's class follows a counted transition matrix $P_{s(t)}(i,j)$. The recursion is the stochastic form of Equation (1.11), run backward over 365 daily stages, as shown in Equation (7.9).

$$\begin{aligned} V_t(S,i) &= \max_{R \le R_{cap}(S)} \left[r_t + \textstyle\sum_j P_{s(t)}(i,j)\, V_{t+1}(S',j) \right] \\ S' &= S + 0.0864\,(\bar{q}_{s,i} - R) - E(S,t) \end{aligned} \tag{7.9}$$

where $r_t = r_t(S,i,R)$ is the day's reward, $\bar{q}_{s,i}$ the class's mean inflow, a release the storage cannot supply is cut at the dead pool, and storage above the crest leaves over it within the day. The reward is the day's energy at the month's price less the costs of a shortfall, as shown in Equation (7.10).

$$\begin{aligned} r_t = {} & 24\,p_t P_e\big(\min(R, Q_{t,max}),\ \tfrac{S+S'}{2}\big) \\ & - 0.0864\left[c_S(D_s - D) + c_E(Q_{env} - I_r)^+\right] \\ & - c_F(Q_{town} - Q_{flood})^+ - 40\,c_F(Q_{town} - Q_{levee})^+ \end{aligned} \tag{7.10}$$

where $Q_{town} = I_r + \bar{l}_{s,i}$ uses the class's mean tributary flow without routing, and reaching the dead pool costs \$2 M a day. Storage is discretized to 40 levels from 120 to 1100 hm^3, interpolated linearly, and the release to a grid of 2 to 50 m^3/s steps cut at the outlets' capacity. The statistics are counted from a declared record of 100 synthetic years of the same weather generator. Its storm days are one-day bursts drawn at 1/120 a day from November to April, each 30 % of a lognormal 230 mm total, so the record holds 1.66 bursts a year of 71 mm on average and a mean inflow of 59.3 m^3/s. The year is repeated, V_0 of one pass becoming V_{365} of the next, until the yearly increment is the same at every state to within \$1. It converges in four years, with spans of 5.3×10^7, 2.2×10^3, 3.2 and 4.9×10^{-3} dollars and a gain of \$15.73 M a year. The policy's release $R^*(S,t,i)$ is read each morning at the storage, the day and the class of the ensemble mean for that day. The water value $w(t,S,i) = \sum_j P(i,j)\,\partial V_t(S,j)/\partial S$ is the price the tree puts on stored water, the slope of the value function that Section 1.4.9 reads as a costate.

7.5.3 The Scenario Tree

Each morning the tree solves one linear program over $N = 10$ days and the M members. Its decisions are the release at each day boundary, $B_{m,j}$, with the release ramping linearly between boundaries so that the day's mean release is $(B_{m,d} + B_{m,d+1})/2$, and the turbine flows $Q_{t,m,d}$. For the first $b = 2$ days the decisions are one shared variable for every member, because they are taken before the members can be told apart (the non-anticipativity constraint of Equation (1.13)), and from day 3 each member has its own branch, as shown in Equation (7.11).

$$\begin{aligned} \max\ & \textstyle\sum_{d<b} p_d k_H Q_{t,d} + \frac{1}{M}\sum_m J_m \\ J_m = {} & \textstyle\sum_{d\ge b} p_d k_H Q_{t,m,d} + w_T S_{m,N} \\ & - c_m \textstyle\sum_j x_{m,j} - c_\sigma \sum_d \sigma_{m,d} \\ \text{s.t. } & B_{m,j} = B_j\ (j \le b), \qquad |B_{m,j} - B_{m,j-1}| \le r \\ S_{m,j} = {} & S_0 + \textstyle\sum_{d<j}\left[0.0864\,(\hat{I}_{m,d} - \bar{B}_{m,d}) - E_d - \sigma_{m,d}\right] \\ & S_{dead} + 5 \le S_{m,j} \le S_{crest} \\ & D_s + Q_{env} \le B_{m,j} \le Q_{t,max} + a\,(S_{m,j} - S_{sill}) \\ & 0 \le Q_{t,m,d} \le \min\left(Q_{t,max},\ \bar{B}_{m,d}\right) \\ & O_{m,j} + \max(\hat{L}_{m,j-1}, \hat{L}_{m,j}) \le (1-\delta)\,Q_{flood} + x_{m,j} \end{aligned} \tag{7.11}$$

where $\bar{B}_{m,d} = (B_{m,d} + B_{m,d+1})/2$ is the day's mean release, k_H the energy of 1 m^3/s for a day at today's head, p_d the price, and $O_{m,j}$ the reach outflow at boundary j, an affine function of the boundary releases because the Muskingum law of Equation (7.5) is integrated exactly for a release that ramps linearly through each day. The rating row applies where the member's storage cannot fall below the sill, and $B_{m,j} \le Q_{t,max}$ applies elsewhere. The crest variable $\sigma_{m,d}$ exists only on the days the ramp lets that member's lake reach the crest. It costs $c_\sigma = 2c_F/0.0864 = \$1.16 \times 10^6$ per hm^3, twice the damage the same water would do over the town's limit, and it is never credited with a negative water value, so the plan spills only where the gates at their rating and ramp cannot hold the lake. The terminal price w_T, the terminal value of Section 1.4.8, is the dynamic program's water

value ten days ahead at today's storage and class, one linear price for every member's end storage. The tree is solved by the bounded-variable simplex method of the apps' shared linear programming module, after a presolve that drops rows the ramp makes unreachable. On the 365 mornings of the challenge with the suggested settings the program has 193 to 291 rows and 126 to 224 columns.

🕮 **Pause.** Try Exercise 7.2 before reading on.

7.5.4 The Chance Constraint and the Fallbacks

The chance-constrained tree holds the channel on k of the M members, the chance constraint of Section 1.4.12. Before the solve, the $M-k$ members with the largest forecast volume reaching the town are let go, and their excess $x_{m,j}$, a slack priced linearly as in Section 1.4.6, carries a token cost $c_m = 0.02\,c_F$ so that their branches stay sensible, while every other member has no $x_{m,j}$ at all. This is the scenario-discarding approach of Campi and Garatti [10] applied to seven members, which is far too few for its probabilistic guarantee, so k/M is a declared fraction of the ensemble and not a measured probability. With $k = M$ the program is the scenario tree in which every member is held. When the members to be held cannot all be held, the program is solved again with every member soft at the full damage cost $c_m = c_F$, which minimizes the expected damage while the energy and a positive water value still count, and a negative water value is taken as zero for any soft member so that it cannot buy flooding. If that fails the storage bounds and the supply floor are relaxed, and if no solve succeeds the rule curve decides that day. The Forecast view labels every morning on which no plan held the required members.

🕮 **Pause.** Try Exercise 7.3 before reading on.

7.5.5 A Perfect-Foresight Bound and Three Models

The challenge measures energy against the most any operator could make from the year's realized inflow, as shown in Equation (7.12).

$$\begin{aligned} E_{PF}(S_f) &= \max_{Q_t} \sum_t 24\,P_e\big(Q_{t,t},\ \tfrac{S_t+S_{t+1}}{2}\big) \\ \text{s.t. } S_{t+1} &= \min\big(S_{crest},\ S_t + V_{in,t} - E_t - 0.0864\,Q_{t,t}\big) \\ 0 &\le Q_{t,t} \le Q_{t,max}, \qquad S_N \ge S_f \end{aligned} \tag{7.12}$$

where $V_{in,t}$ is the day's recorded inflow volume and S_f is the run's own end storage. It relaxes the channel, the ramp and the supply, so it bounds every policy that ends at or above S_f. A forward dynamic program solves it on 1 and 0.5 hm^3 storage grids, and the two answers are extrapolated to zero spacing. It is a teaching bound, never a controller. Three models are kept apart. The plant is everything in Section 7.3. The dynamic program's model is a daily Markov chain of class means with no routing and no ramp. The tree's model is the plant's own storage, rating, ramp and Muskingum laws driven by the forecast members, with energy valued at today's head.

7.6 Measurements on Ten Cases and a Water Year

The cases appear in Table 7.3, and the measured comparisons in Table 7.4. All values below come from headless runs of the application's own engine on each case's declared weather (seed 365), advanced through the same step sequence the browser uses, and a setting that a row changes is applied before Run, which primes the run exactly as Reset would (Section 1.6.4). A day over the limit is a day on which the true town flow exceeded 450 m^3/s at any physics step, the pool used is the largest storage above the rule curve, and gate volumes are integrated over time.

7.6.1 Routing, Lead and a Forecast That Busts

Case 2's two storms show the order of arrival at work. The first storm leaves 92 hm^3 in the pool when the second begins on day 16, the town exceeds the limit from day 17.4 and peaks at 575 m^3/s on day

Table 7.3 The Ten Cases

#	Strategy, start and weather	Teaching point
1	Run of river, 1 Dec, 300 mm on day 12	storage is flood control
2	Rule, wet December, 260 and 360 mm ten days apart	the tributary wins
3	Rule + forecast, lead 2 d, case 2's weather	the lead a channel needs
4	Rule + forecast, lead 5 d, 450 mm bust for day 10	what chasing costs
5	SDP policy, 1 Nov, 380 mm on day 40	a season's table
6	Scenario tree, 420 mm on day 12	trunk and branches
7	Scenario tree, mid March, $\sigma_{10} = 0.9$	letting a member go
8	Rule, ramp 40 m^3/s a day, cloudburst on the tributary	permit against flood rule
9	Rule, $FP = 100\,hm^3$, 360 mm on day 1	an infeasible flood
10	Challenge: rule, 1 Oct, storms 20 Dec and 13 Feb, bust for 19 Jan, dry spring	the whole year

Table 7.4 Measured Comparisons on the Cases' Own Windows

Case	Run	Town (m^3/s)	Days over	Pool (hm^3)	Gates (hm^3)	Energy (GWh)
1	run of river[a]	732	7	31	209	57.2
	rule curve	444	0	137	196	63.9
2	rule, FP 300	575	4	304	465	130.2
	rule, FP 400	575	4	273	515	131.7
3	lead 2 d	552	4	322	430	130.0
	lead 5 d	425	0	307	444	133.5
4	wettest member	314	0	2	23	23.4
	median member	185	0	2	2	24.8
	tree, all held	429	0	0	66	46.0
	tree, $k = 6$	279	0	0	13	55.1
5	SDP policy	530	4	38	171	224.4
	rule curve	452	2	310	368	164.9
6	scenario tree	434	0	115	139	123.7
	rule curve	442	0	285	301	93.6
7	tree, all held	412	0	173	64	67.6
	tree, $k = 6$	340	0	215	24	78.0
8	rule, ramp 40	642	5	157	220	64.5
	rule, ramp 150	565	3	157	219	60.2
	tree, ramp 40	429	0	249	105	56.7
	tree, Q_{env} 60	459	1	242	87	58.0
9	rule curve[a]	752	7	139	156	65.9
	manual 300 m^3/s	560	7	120	258	89.9
	tree, perfect forecast[a]	651	8	105	269	85.9

[a]The levee failed. The browser pauses there, and the measurement runs on to the end of the window.

19.3, and only then does the lake fill, reaching 1104 hm^3 and spilling over the crest for 1.7 days from day 24.3. With a 400 hm^3 pool the lake stays at 973 hm^3 or below and the town's peak is unchanged, as Table 7.4 shows. In case 3 the lake's highest level falls only from 1122 to 1107 hm^3 between a two-day and a five-day lead, so what the lead buys is not volume but timing, room made while the channel is still empty. In case 4 a 450 mm storm promised for day 10 vanishes from the forecasts on day 9. Chasing the wettest member, which put it at 765 mm, the rule sends 23 hm^3 through the gates and the lake stands at 755 hm^3 on day 10, against 786 hm^3 when hedging on the median member and 801 hm^3 with no forecast. At the case's head 1 hm^3 through the turbines is worth 0.21 GWh, so the chase costs water and not safety.

🕮 **Pause.** Try Exercise 7.1 before reading on.

7.6.2 A Season's Table and a Storm's Plan

The dynamic program draws the lake down early. In case 5 it holds the turbines at their full 180 m^3/s for 55 % of the 40 days before the storm, opens no gate, and stands at 559 hm^3 on day 40, when the storm falls, where the rule holds 794 hm^3. It makes 224.4 GWh over 90 days against 164.9 GWh,

uses 38 hm^3 of its pool against 310 hm^3, and ends at 636 hm^3 against 797 hm^3, so part of its energy is water the rule kept. Its table, however, neither routes the river nor sees the ramp, and its storm release puts the town at 530 m^3/s on day 44.6, over the limit on 4 days, while the rule touches 452 m^3/s. Across the flood season the policy releases early at and below the rule's target, a mean of 225.8 m^3/s against the rule's 128.1 m^3/s from 100 hm^3 below the target up to it, and less than the rule above the target, 276.7 m^3/s against 380.0 m^3/s, where the rule releases at the channel's limit. The water value of Figure 7.2(b) shows why. It is positive all year at low storage and rises to 87 k$ per hm^3 at 305 hm^3 in June, when summer supply and energy depend on the water, and it turns negative in the winter near the crest, reaching −72 k$ per hm^3 at 1000 hm^3 on 1 February in the middle class. Over the 4800 grid points of half-month, storage and class it is negative at 448, never below 472 hm^3.

The scenario tree is the storm's plan. In case 6 it lowers the lake from 800 to 664 hm^3 by day 12, mostly through the turbines, uses 115 hm^3 of its pool, holds the town at 434 m^3/s and makes 123.7 GWh in 40 days. The rule waits for the storm, uses 285 hm^3 of the pool, peaks the town at 442 m^3/s, makes 93.6 GWh and sends 301 hm^3 through the gates to the tree's 139 hm^3, and the two end within 10 hm^3 of each other. A morning of the challenge, 9 February, is drawn in Figure 7.2(c) and (d). Its seven members put the inflow peak between 536 and 1523 m^3/s, and the inflow that came peaked at 1121 m^3/s. The trunk commits 198 m^3/s and then 162 m^3/s, identical in every branch. Over days 3 to 6 the held members' releases fall to the 12 m^3/s floor of supply and environmental flow as their tributaries fill the channel, and five of the seven plans, four held members and the let-go one, end the horizon with the lake at the crest. The let-go wettest member releases up to 355 m^3/s and plans the town at 642 m^3/s. The crest bound is active on those members at day 10, so the trunk's release is what keeps them under it.

7.6.3 The Chance Constraint and the Permit

Holding every member means planning for the wettest. In case 7, a spring storm with $\sigma_{10} = 0.9$, the tree that holds all seven sends 64 hm^3 through the gates and makes 67.6 GWh. Holding six, it sends 24 hm^3 and makes 78.0 GWh, and the town peaks at 340 m^3/s against 412 m^3/s. Both end at 1082 hm^3. The all-member tree finds no plan that holds the wettest member on 6 of its 30 mornings and weighs that member's excess at the full damage cost, while the chance-constrained tree weighs it at 2 % of it, which is what frees the water. On the bust of case 4 the same contrast is 13 against 66 hm^3 through the gates and 55.1 against 46.0 GWh. The storm that came was no wetter than the six held members, and the measurement says nothing about the year in which it is the seventh.

> **TRY IT IN THE APP 7 · The Chance Constraint**
>
> Load the scenario and run it on Scenario MPC, which holds every member. Then choose Chance-constrained MPC with Members held (k) at 6, press Reset and run again. Over the case's window the gates pass 24 hm^3 instead of 64 hm^3, the energy rises from 67.6 to 78.0 GWh, and the town peaks lower (Table 7.4).

Case 8 puts a permit against the flood rule. The lake starts 160 hm^3 into its pool with the release near the channel's limit, and a cloudburst falls mostly on the tributary, which peaks at 394 m^3/s. The rule at a ramp of 40 m^3/s a day lets the town reach 642 m^3/s, 5 days over the limit, and at 150 m^3/s a day still 565 m^3/s, because it starts the cut only when the gauge rises. The tree at 40 m^3/s a day starts the cut before the rain, holds 249 hm^3 in the pool and keeps the town at 429 m^3/s, although on 9 of its 20 mornings no plan held every member. Raising the environmental flow to 60 m^3/s puts a floor under that cut, and the town reaches 459 m^3/s on 1 day.

7.6.4 A Flood No Release Holds

Case 9 is the infeasible variant. The flood pool is at its smallest, 100 hm^3, the lake at the top of the rule curve, and a 360 mm atmospheric river begins on day 1. The rule fills the lake to 1139 hm^3, spills over the crest, and the levee fails on day 7.1 with the town peaking at 752 m^3/s. Constant releases

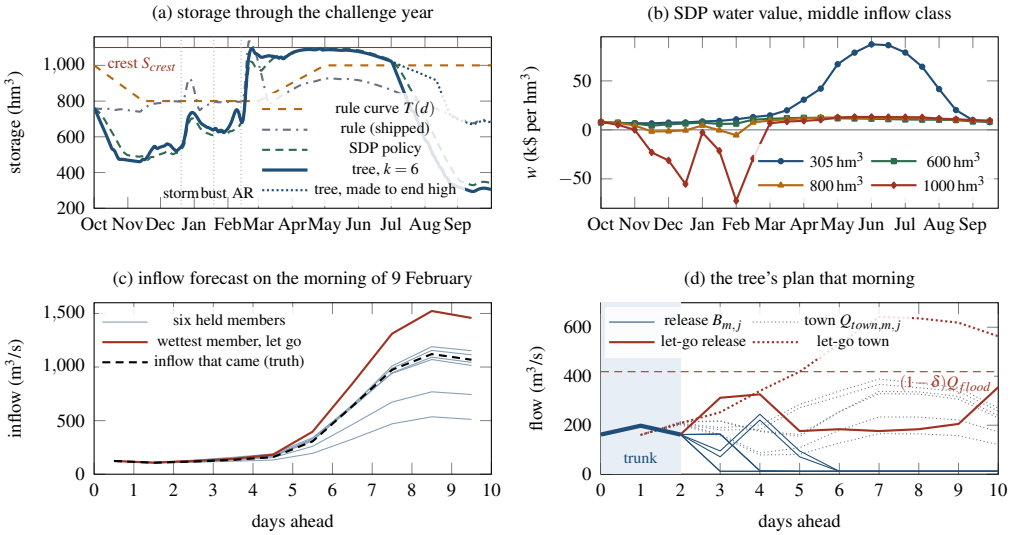


Figure 7.2 Three deciders measured on the challenge. (a) Storage through the water year under the shipped rule curve (dash-dotted), the dynamic program's policy (densely dashed), the suggested chance-constrained tree ($k = 6$, heavy solid) and the same tree made to end the year high by raising its terminal water value by 2000 \$ per hm^3 from 1 July (dotted), with the rule curve's target dashed. Vertical dotted lines mark the 20 December storm, the bust promised for 19 January and the 13 February atmospheric river (AR). (b) The dynamic program's water value w at four storages on the first day of each half-month, in the middle inflow class, one marker shape per storage. (c) The ensemble issued on the morning of 9 February (six held members thin, the wettest heavy) and the daily inflow that then came (dashed, teaching truth, never read by the tree). (d) The tree's plan that morning: boundary releases $B_{m,j}$ (solid), with the two-day trunk shared (heaviest line), and each member's town flow (dotted) against the planning limit of 418.5 m^3/s. Six members are held, and the wettest is let go (heavy lines).

held from the storm's first day give peaks of 744, 643, 560, 660 and 759 m^3/s for 100, 200, 300, 400 and 500 m^3/s. The scenario tree peaks at 663 m^3/s, and on a perfect forecast, a teaching bound, at 651 m^3/s. A wider search through the plant does no better: 245 two-step schedules reach 522 m^3/s at best and a coordinate descent on the first 12 daily releases 503 m^3/s. This is a search, not a proof of infeasibility. Linear programs with perfect foresight that check the town only at day ends or let the crest overflow instantly do reach 450 m^3/s, but their schedules peak at 554 to 671 m^3/s when run through the plant.

7.6.5 The Water Year

The challenge runs from 1 October with a 280 mm storm on 20 December, a 300 mm storm forecast for 19 January that never falls, a 440 mm atmospheric river on 13 February and rain at 35 % of its mean for 92 days from 1 March. Its target is no day over the 450 m^3/s limit, supply of at least 95 % of the schedule, and energy of at least 80 % of the perfect-foresight energy for the same year ending at the run's own storage. The verdict holds only for the case's own weather and water: moving the rain slider, pressing an event button or using the helper that sets the lake on the rule curve voids it, the last because it moves water outside the balance. Table 7.5 lists eight configurations, and the storage paths of four appear in Figure 7.2(a). The shipped rule curve fails in February, as the opening paragraph describes, with 400.4 GWh, 75.6 % of its bound. The dynamic program makes 89.5 % but puts the town over the limit on 9 days. Both trees pass, and the suggested chance-constrained tree, holding six members, makes 528.9 GWh, 89.5 % of 590.8 GWh, with a town peak of 425 m^3/s and

Table 7.5 The Challenge Water Year at Seed 365. Share of the Perfect-Foresight Energy at the Run's Own End Storage

Configuration	Energy (GWh)	Share (%)	End (hm^3)	Days over	Town (m^3/s)	Seeds passed
rule curve (shipped)[a]	400.4	75.6	687	8	710	0/7
rule, lead 5 d, wettest[a]	399.1	75.4	687	5	651	0/7
SDP policy	524.6	89.5	341	9	600	1/7
tree, all seven held	528.8	89.5	305	0	428	7/7
tree, $k = 6$ (suggested)	528.9	89.5	305	0	425	7/7
tree, $k = 5$	532.8	90.4	318	4	525	6/7
tree, $k = 6$, made to end high	469.7	88.5	681	0	425	—
run of river[a,b]	377.4	76.3	861	16	962	0/7

[a]The levee failed. [b]Supply 45.0 %, the others 100.0 %. All start at 760 hm^3.

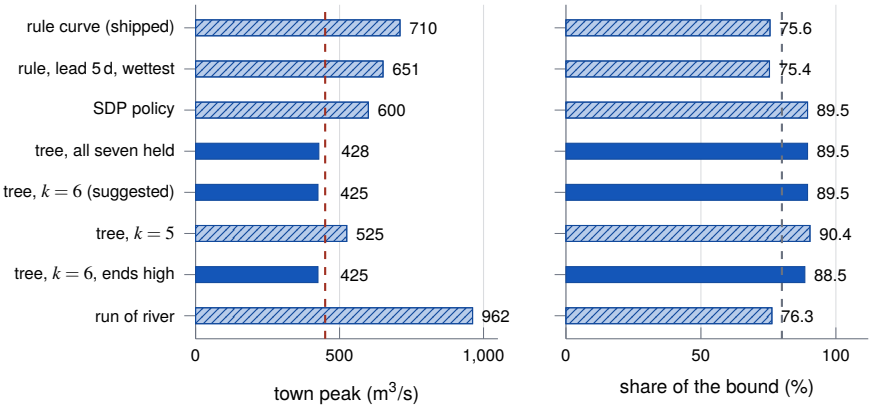


Figure 7.3 The challenge water year redrawn from Table 7.5: the town's peak flow against the 450 m^3/s channel limit (red dashed line, left) and the share of the perfect-foresight energy against the 80 % target (gray dashed line, right). Dark bars pass the challenge, hatched bars fail it. Notice that the shares of the dynamic program, the trees and the tree that holds five members lie within two points of each other, while only the trees that hold six or seven members keep the town under its limit.

the supply in full. Holding five members puts the town over the limit on 4 days. Over seven weather seeds the suggested configuration passes 7 times with 84.7 to 91.8 %, the rule passes none, and the dynamic program passes once. On 9 of the 365 mornings, between 6 and 15 February, no plan held six members and the tree minimized the expected damage instead, and on every such morning of the shipped cases the plan's expected excess is within 2 % of the least any plan reaches.

Both trees end the year at 305 hm^3 from a start of 760 hm^3, while the rule ends at 687 hm^3, which raises the question of whether the pass was bought by emptying the lake. It was not. Made to end the year at 681 hm^3, by raising its terminal water value by 2000 $ per hm^3 from 1 July, the same tree makes 469.7 GWh, 88.5 % of the 530.8 GWh its own end storage allows, and still passes with the town at 425 m^3/s. The drawdown is worth one point of the share. The comparison at the run's own end storage was kept because a common floor is not a bound. Perfect foresight from 760 hm^3 gives 515.5 GWh, which the suggested run exceeds, at 102.6 %. The trees spend the lake in the late summer, when energy is dearest, because the water value says the winter refills it, and run on into a second year they refill it to 1096 hm^3 with the supply in full and make 398.1 GWh against the rule's 352.5 GWh. Figure 7.3 draws the town peaks and energy shares of Table 7.5 side by side.

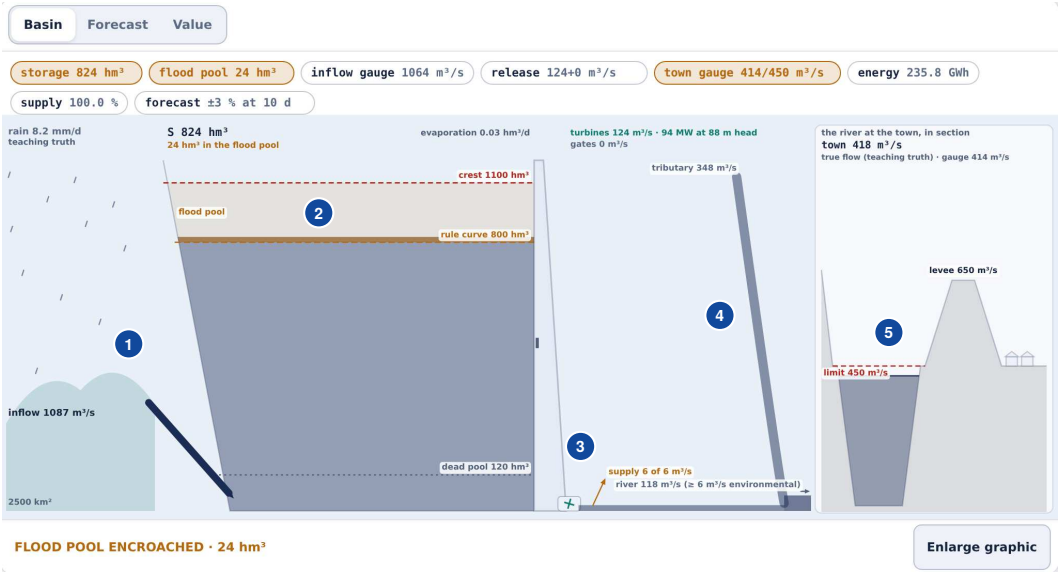


Figure 7.4 The challenge on the suggested settings (chance-constrained tree, $k = 6$) at day 139.0, 17 February, the state of Table 7.1, cropped to the Basin view. (1) The atmospheric river brings 1087 m^3/s into the lake (the gauge reads 1064). (2) The lake stands at 824 hm^3, 24 hm^3 into its flood pool above the 800 hm^3 rule curve, so the status strip reads FLOOD POOL ENCROACHED. (3) The dam releases 124 m^3/s through the turbines (94 MW at an 88 m head) with the gates shut, and 118 m^3/s stays in the river below the 6 m^3/s diversion. (4) The tributary brings 348 m^3/s. (5) The town's true flow is 418 m^3/s (gauge 414) against the 450 m^3/s limit and the 650 m^3/s levee. The chips read 235.8 GWh and the supply in full, and the year ends with no day over the limit and a town peak of 425 m^3/s (Table 7.5).

+ DID YOU KNOW? One Member More Let Go Buys Little

From the simulator: over the challenge year the tree that holds five of the seven members makes 532.8 GWh against 528.9 GWh for six, but puts the town over the limit on 4 days with a peak of 525 m^3/s and passes 6 of 7 seeds, against 0 days, 425 m^3/s and 7 of 7 (Table 7.5).

The state of the challenge on 17 February, four days after the atmospheric river began, is shown in Figure 7.4.

7.6.6 Code Verification and Agreement

Code verification asks whether the program solves its own equations. At every 0.1 d step of the challenge year the storage equals its start plus inflow less outflow and evaporation to 1.3×10^{-10}, 3.4×10^{-11} and 6.3×10^{-11} hm^3 under the rule, the dynamic program and the suggested tree, and the outflow equals the diversion plus the river and the reach's inflow less outflow its storage change to 2.6×10^{-10} hm^3. Halving the step changes the challenge year's energy by -1.3×10^{-5} under the rule and 2.8×10^{-4} under the suggested tree, and leaves every verdict and the energy ranking of the seven strategies unchanged. The perfect-foresight extrapolation from 1 and 0.5 hm^3 agrees with the one from 0.5 and 0.25 hm^3 to 2.9×10^{-4}. A Bellman backup recomputed with independent reward and transition code at all 73 000 states equals the stored value function to the dollar with the same decision at every state, and one more year of recursion changes no decision. Simulated for 2900 years on its own Markov model, the policy averages \$15.97 M a year against the reported gain

of \$15.73 M, a difference of 1.5 % from the storage interpolation. An independent brute force on a three-day, three-member tree with a binding rating and a let-go member equals the linear program to 0.000 on a 0.05 m^3/s grid, and in closed loop every plan stays inside the ramp to 8×10^{-12} m^3/s. With k of 5 and 4 the plan exceeds the channel on exactly $M - k$ members. Runs logged at 1 and 30 days a second are identical.

Agreement with the published hydrology is claimed separately, and more narrowly. The plant's cascade, driven by a unit impulse, matches Nash's Equation (7.2) at 1.5, 3 and 6 d to 2.0×10^{-6}, 3.8×10^{-7} and -1.7×10^{-7} relative, with its peak at 3.00 d, unit volume, a first moment of 4.5000 d and a second of 27.000 d^2, the moments of nk_N and $n(n+1)k_N^2$ that Nash gives. The reach's response to a unit step from rest matches its analytic solution, $O(t) = 1 - e^{-t/K(1-X)}/(1-X)$, including the initial dip of -0.250, at 0.1, 0.5 and 2.0 d to 10^{-6}, and on a 450 m^3/s hydrograph the volume closes to 6.4×10^{-14} m^3/s d with a centroid lag of 1.000 00 d, equal to K as Muskingum's law requires. Because no river or dam data entered the model, these checks validate it against no real basin. The value function takes 735 to 746 ms to solve on first load in a desktop browser, a tree solve averages 4.8 to 4.9 ms, and the water year plays in 12.5 to 12.6 s at 30 days a second.

7.7 Technology in Practice

What Measures It

A flood-control decision reads the lake, the rivers and the forecast. Lake stage is the index from which reservoir contents are computed, and a lake gauge is held to the same accuracy as a river gauge. Satellite telemetry is now the standard link from a gauging station to the office: it delivers the stage in near real time and alerts the hydrographer to extreme or abnormal events and to instrument malfunctions [18]. The inflow gauge and the town gauge of Figure 7.1 stand for two such stations, and the exactly measured storage S for a lake gauge read through the storage curve.

> **TECHNOLOGY** **Sensor: The River Gauge**
>
> A river gauge measures stage, the height of the water surface, with a float in a stilling well, a gas-purge bubbler read by a pressure transducer, a submerged pressure transducer, or a radar that times a pulse reflected from the water. The flow follows from the stage through the site's stage-discharge relation. A whole station must be accurate to 0.01 foot or 0.2 percent of the effective stage, whichever is greater, and its recorder stores a value every 5 or 15 minutes or every hour. A clogged intake makes the well lag the river, high velocity beside a pier can draw the reading down by 0.5 foot or more, and radar is upset by waves and floating debris [18]. The chapter's two gauges carry a colored 2 % error and none of these faults.

What Moves It

A dam releases through its powerhouse, its outlet works and its spillway. A spillway needs gates or another movable crest when an uncontrolled crest cannot pass the required flood, and also when it must release storage below the normal water surface. Radial gates and wheel-mounted slide gates are the most common on large spillways [19]. In the chapter's plant the turbines take the release first and the gated spillway the rest, with a rating that is linear from the sill to the crest (Equation (7.4)), and the crest spills on its own.

> **TECHNOLOGY** **Final Element: The Radial Spillway Gate**
>
> A radial gate is a steel cylindrical segment carried on radial arms that turn on trunnion pins. Its face is centered on the pins, so the whole water load passes through them and the hoist lifts only part of the gate's weight. The small effort suits simple automatic control, and the gates of one spillway can open step by step as the lake rises. The choice of gate weighs ice and debris, whether an operator attends in a flood, and whether electric power is available. A human or mechanical failure to open gates raises the lake above its expected level [19]. The chapter's gate fault halves the gated capacity, $\gamma = 0.5$.

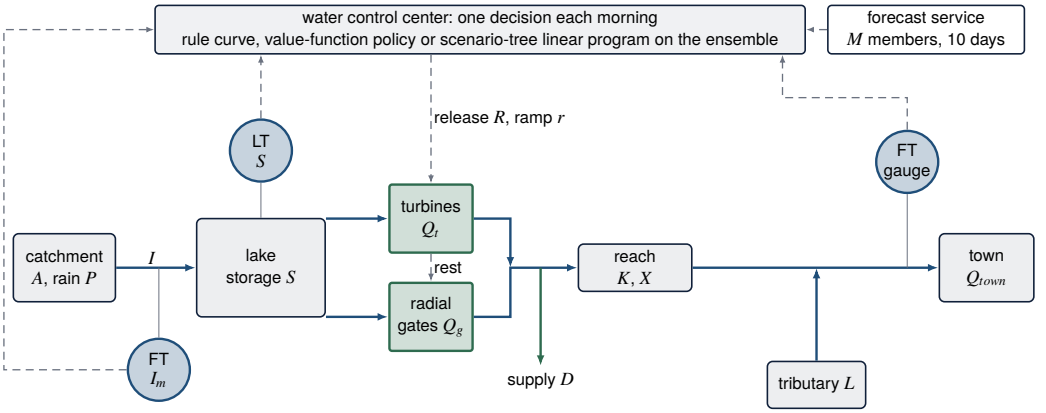


Figure 7.5 The basin's decision loop on field hardware, drawn from Sections 7.3 and 7.5. Green boxes are the turbines and the radial spillway gates, blue circles the inflow gauge, the lake gauge and the town gauge, dashed arrows signals. Notice that the town gauge sits below the tributary, a day's travel from the dam, so only the forecast lets the morning's release act before the town's river rises.

Where the Controller Runs

A water control plan is a well-defined set of rules and tools that guides daily decisions to release or retain water in a reservoir's flood pool. In one forecast-based study, the streamflow forecast was issued at least daily, a decision-support model recommended the flood releases, and that model's computational step was later moved from daily to hourly [13]. River gauges record every 5 or 15 minutes or every hour [18]. The chapter's deciders act once a day on the morning's ensemble, while the plant is integrated every 0.1 d, and the scenario tree is one linear program solved each morning, as Figure 7.5 draws it.

> **IN PRACTICE** **Hard Constraints Every Plan Must Meet**
>
> When one reservoir's water control plan was compared with four forecast-based alternatives, every plan had to satisfy the same hard, inviolable constraints: the limits on the release's rate of change, the maximum release schedule, the instream minimum flow, and no call for forecast updates more often than they are issued. Each plan should also avoid spillway flow as far as possible. Under the existing plan no forecast is used, and storage above the guide curve is always evacuated as quickly as feasible [13]. The chapter's ramp limit r, environmental flow Q_{env} and one ensemble each morning stand for these rows.

What Is Modernizing It

Three developments are changing reservoir operation. Ensemble streamflow forecasts have been operational at one river forecast center since 2012, driven by three days of forecaster precipitation followed by twelve days of a global weather ensemble. In two test years an operation added a variable buffer pool to the guide curve and let a 15-day ensemble forecast recommend the flood releases [13]. Observation is improving as well: reconnaissance fills major gaps in observations over the ocean inside atmospheric rivers, storms that can be forecast with a long lead [13], and radar stage sensors, which need no stilling well, are now in common use [18]. The chapter's chance-constrained tree belongs to the same family of forecast-based rules, on a declared basin and with seven members instead of an operational ensemble.

7.8 Limitations

The basin is declared, and every limit is a teaching value. The runoff is linear, rain is uniform over a catchment, the forecast errs only in the rain, and the cases run on one weather seed. At other seeds

three answers change: at seed 2 the rule in case 1 touches 469 m^3/s for a day and a five-day lead in case 3 peaks at 485 m^3/s, 3 days over the limit, and at seed 1 the wettest and median runs of case 4 coincide.

> **OPEN CHALLENGE** **Ensembles That Run Low in a Flood**
>
> Tested on years of winter hindcasts, one reservoir's ensemble put the observed three-day inflow in its upper tail at leads of one to six days, a bias toward under-forecasting that shrinks with the lead, and its wettest forecasts at one to three days looked too narrow. The study asked for methods that produce ensembles capturing more of the meteorological and hydrological uncertainty [13]. This chapter's members err only in the rain, with no mean error. Progress would be ensembles whose spread matches their error in the floods that drive release decisions.

The dynamic program has three known mismatches with the plant. Its 100-year record holds one-day bursts of 71 mm, 1.66 a year, not the three-day atmospheric rivers of 260 to 440 mm that the cases declare, so it learns a milder winter than the plant has, and the application's model dialog and instructions say so. Trained on full three-day storms instead, with a mean record inflow of 73.4 m^3/s rather than 59.3 m^3/s, its gain falls from \$15.73 M to −\$4.20 M a year, and the suggested tree that prices water with it draws the challenge's lake to 171 hm^3 by 1 November and makes 66.3 % of its bound, so that record was not adopted. A reader who compares case 5's storm with the record's bursts sees the mismatch directly. The table is daily and ignores the reach and the ramp, which is why its storm release in case 5 overshoots the channel, and it releases early at and below the rule's target, not above it.

The tree prices every member's end storage at one linear water value taken at today's storage, so on 9 February it values five members' lakes at the crest at the 10 128 \$ per hm^3 of a lake at 752 hm^3. Its let-go members carry a token 2 % of the damage cost, its fallback on a truly infeasible morning minimizes expected damage and is labeled as such, and seven members are far too few for the guarantee of the scenario-discarding method. Case 9's infeasibility rests on a search, not a proof. The band of the forecast issued five days earlier is remembered by the scene for display and never read by a decider, and a capture made without a display shows none. A simulated pass is not a flood-safety assessment, an operating permit or a certification of any decision rule.

> **OPEN CHALLENGE** **How Much Risk a Release Should Accept**
>
> One forecast-based release rule runs every ensemble member through the reservoir and limits the fraction of members whose storage exceeds a threshold, by a risk curve over the 15-day forecast. The study recommended optimizing those risk-tolerance curves by more sophisticated methods, exploring multi-objective approaches, and generating synthetic ensembles to test the rule over a wider range of hydrology [13]. This chapter's k of M plays the same role, a declared fraction rather than a measured probability. Progress would be a risk level chosen by a stated method and checked against long records.

7.9 Conclusion

A flood-control reservoir is a plant whose safety is decided before the storm, and this chapter measured that on a declared basin with Nash cascades, a Muskingum reach, a tributary below the dam and a seven-member ensemble forecast. A wet December floods the town with the lake 234 hm^3 below its crest, and a rule told the forecast needs five days of lead to hold the second storm. The dynamic program draws the lake to 559 hm^3 before a storm where the rule holds 794 hm^3, but its season's table puts the town at 530 m^3/s. The scenario tree uses 115 hm^3 of pool where the rule uses 285 hm^3 and makes 123.7 GWh against 93.6 GWh, and letting the wettest member go buys 10.4 GWh in a spring storm. Over the water year the suggested tree keeps the town at 425 m^3/s and makes 89.5 % of the perfect-foresight energy, and 88.5 % when made to end the year high. The dynamic program's record is milder than the cases' storms, the tree prices stored water at one linear value, and the infeasible case

rests on a search. A fuller model would add soil moisture to the runoff, timing errors to the forecast, a ramp state to the dynamic program and a concave terminal value to the tree. The measurements show that the room a dam makes for a flood is made from a forecast, days before the flood arrives.

KEY POINTS

- A dam makes room for a flood only as fast as the channel below carries water, so it must act on water that has not yet arrived.
- A wet December floods the town with the lake 234 hm^3 below its crest, because the tributary below the dam rises first.
- Letting the wettest member go buys 10.4 GWh in a spring storm.
- Over the water year the suggested tree keeps the town at 425 m^3/s and makes 89.5 % of the perfect-foresight energy.

Where else this applies. An irrigation canal (Chapter 5) also releases ahead of a delay: an order released ahead moves its levels about 2 cm, against up to 16 cm at its start (Table 5.6).

References

[1] J. D. C. Little, "The use of storage water in a hydroelectric system," *Journal of the Operations Research Society of America*, vol. 3, no. 2, pp. 187–197, 1955.

[2] W. W.-G. Yeh, "Reservoir management and operations models: A state-of-the-art review," *Water Resources Research*, vol. 21, no. 12, pp. 1797–1818, 1985.

[3] J. W. Labadie, "Optimal operation of multireservoir systems: State-of-the-art review," *Journal of Water Resources Planning and Management*, vol. 130, no. 2, pp. 93–111, 2004.

[4] D. P. Loucks and E. van Beek, *Water Resource Systems Planning and Management*. Springer International Publishing, 2017.

[5] A. T. Schwarm and M. Nikolaou, "Chance-constrained model predictive control," *AIChE Journal*, vol. 45, no. 8, pp. 1743–1752, 1999.

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

[7] J. R. Birge and F. Louveaux, *Introduction to Stochastic Programming*, ser. Springer Series in Operations Research and Financial Engineering. Springer New York, 2011.

[8] J. R. Stedinger, B. F. Sule, and D. P. Loucks, "Stochastic dynamic programming models for reservoir operation optimization," *Water Resources Research*, vol. 20, no. 11, pp. 1499–1505, 1984.

[9] H. L. Cloke and F. Pappenberger, "Ensemble flood forecasting: A review," *Journal of Hydrology*, vol. 375, no. 3–4, pp. 613–626, 2009.

[10] M. C. Campi and S. Garatti, "A sampling-and-discarding approach to chance-constrained optimization: Feasibility and optimality," *Journal of Optimization Theory and Applications*, vol. 148, no. 2, pp. 257–280, 2011.

[11] L. Raso, D. Schwanenberg, N. C. van de Giesen, and P. J. van Overloop, "Short-term optimal operation of water systems using ensemble forecasts," *Advances in Water Resources*, vol. 71, pp. 200–208, 2014.

[12] G. Uysal, R. Alvarado-Montero, D. Schwanenberg, and A. Şensoy, "Real-time flood control by tree-based model predictive control including forecast uncertainty: A case study reservoir in Turkey," *Water*, vol. 10, no. 3, p. 340, 2018.

[13] Lake Mendocino FIRO Steering Committee, "Lake Mendocino forecast informed reservoir operations final viability assessment," UC San Diego, Tech. Rep., Dec. 2020. [Online]. Available: https://cw3e.ucsd.edu/FIRO_docs/LakeMendocino_FIRO_FVA.pdf

[14] A. Castelletti, F. Pianosi, and R. Soncini-Sessa, "Water reservoir control under economic, social and environmental constraints," *Automatica*, vol. 44, no. 6, pp. 1595–1607, 2008.

[15] B. A. Faber and J. R. Stedinger, "Reservoir optimization using sampling SDP with ensemble streamflow prediction (ESP) forecasts," *Journal of Hydrology*, vol. 249, no. 1–4, pp. 113–133, 2001.

[16] J. E. Nash, "The form of the instantaneous unit hydrograph," in *IAHS Red Book Series, Publication 45*. International Association of Hydrological Sciences, 1958, pp. 114–121. [Online]. Available: https://iahs.info/uploads/dms/045011.pdf

[17] J. A. Cunge, "On the subject of a flood propagation computation method (Muskingum method)," *Journal of Hydraulic Research*, vol. 7, no. 2, pp. 205–230, 1969.

[18] V. B. Sauer and D. P. Turnipseed, "Stage measurement at gaging stations," U.S. Geological Survey, Reston, VA, Techniques and Methods 3-A7, 2010.

[19] U.S. Bureau of Reclamation, *Design of Small Dams*, 3rd ed. Denver, CO: U.S. Department of the Interior, Bureau of Reclamation, 1987. [Online]. Available: https://www.usbr.gov/tsc/techreferences/mands/mands-pdfs/SmallDams.pdf

7.10 Exercises

7.1 Concept check · 2 min

In case 2 the rule curve meets a second storm with 92 hm^3 already in its 300 hm^3 flood pool. The town's river first exceeds the channel limit on day 17.4, with the lake at 866 hm^3, far below its 1100 hm^3 crest. Why does the town flood while the lake still has room?

(a) The tributary below the dam rises first, and a cut in the release reaches the town a day after it is made

(b) The flood pool is too small, so the rule has to release at the channel's limit through the whole storm

(c) The lake spills over its crest, and the crest flow adds to what the gates already release

(d) The diversion stops taking the supply in a storm, so its share of the release goes down the river

7.2 Concept check · 2 min

The scenario tree plans ten days on seven forecast members, and for the first $b = 2$ days the release at each day boundary is one variable shared by every member. Why are those releases shared?

(a) The forecast's first two days carry no rain error, so every member agrees on them

(b) They are decided before anyone can tell which member will come true, so one release must serve all seven

(c) The ramp limit stops the release from changing within two days, so the members cannot differ

(d) The dynamic program's water value fixes the release for two days, and the tree adds the branches after it

7.3 Concept check · 2 min

The chance-constrained tree holds the channel on $k = 6$ of the $M = 7$ members. What does this setting promise?

(a) The town stays under its limit with a probability of at least 6/7, as the scenario approach guarantees

(b) Nothing about the wettest member: its branch may release anything, since its excess carries no cost

(c) The limit holds on the six driest members, the wettest pays a token price, and no flood probability is stated

(d) All seven members are held, each to six sevenths of the channel's limit, to leave a margin

7.4 Theory and calculation · 6 min

The rule curve's flood rule releases $R_{ch} = (1-\delta)\,Q_{flood} - \hat{L} + D_s$ (Equation (7.8)) with $\delta = 0.07$ and $Q_{flood} = 450\,\mathrm{m}^3/\mathrm{s}$. Take the winter supply $D_s = 6\,\mathrm{m}^3/\mathrm{s}$, and recall that 1 m^3/s held for a day is 0.0864 hm^3.

(a) With no tributary flow ($\hat{L} = 0$), compute R_{ch} and the volume it removes from the lake in a day. *Checkpoint:* 424.5 m^3/s and 36.7 hm^3 a day.

(b) How many days does that release take to empty a full 300 hm^3 flood pool with no inflow, and with a steady inflow of 100 m^3/s? *Checkpoint:* 8.2 d and 10.7 d.

(c) The tributary runs at 200 m^3/s. What is R_{ch} now, and how much longer does the same pool take to empty with no inflow? Use the result to say why the rule of case 3 needs days of forecast lead and not hours. *Checkpoint:* 224.5 m^3/s, 19.4 hm^3 a day, 15.5 d, nearly twice as long.

7.5 Theory and calculation · 6 min

The reach is one Muskingum element with $K = 1$ d and $X = 0.2$ (Equation (7.5)).

(a) Compute the coefficients $C_0 = (\Delta t - 2KX)/D$, $C_1 = (\Delta t + 2KX)/D$ and $C_2 = (2K(1-X) - \Delta t)/D$, with $D = 2K(1-X) + \Delta t$, for a daily step $\Delta t = 1$ d. Check that they are positive and sum to one. *Checkpoint:* 0.231, 0.538, 0.231.

(b) The continuous law answers a unit step in I_r from rest with $O(t) = 1 - e^{-t/K(1-X)}/(1-X)$. Find O at $t = 0$ and at 1 d, and the time at which O crosses zero. *Checkpoint:* −0.25, 0.642, 0.18 d.

(c) The dam raises its release by 100 m^3/s at once. By how much does the flow leaving the reach first fall, and why does the plant's ramped release keep this dip small? *Checkpoint:* 25 m^3/s.

7.6 Theory and calculation · 5 min

The lake's depth at the dam is $h = 6.69\,S^{0.4}$ with S in hm^3, the head is $H = h - 10$, and the powerhouse makes $P_e = 9.81 \times 10^{-3}\,\eta\,Q_t H$ in MW with $\eta = 0.88$ (Equation (7.3)).

(a) Compute the head and the full-powerhouse power ($Q_t = 180\,m^3/s$) at 800 hm^3 and at the 1100 hm^3 crest. *Checkpoint:* 87.0 m and 135 MW; 100.1 m and 155.6 MW.

(b) How much energy does 1 hm^3 make through the turbines at 800 hm^3, and at 305 hm^3, the storage at which both trees end the challenge year? *Checkpoint:* 0.209 GWh and 0.134 GWh.

(c) By what fraction does a lake held at 305 hm^3 lower the energy of each later 1 hm^3 against one at 800 hm^3? Say why the challenge's score compares each run with perfect foresight ending at the run's own storage. *Checkpoint:* About 36 %.

7.7 App challenge · 8 min

Load scenario 3, "Pre-release on a forecast: how much lead the channel needs", default seed, Rule + forecast with the Pre-release lead at 2 d as loaded. Run at 7 d/s: the case pauses itself at day 45.

(a) Read the Town peak tile, the days over the limit on the Energy, supply and exceedance tile, and the Flood pool tile. *Checkpoint:* 552 m^3/s, 4 days, 322 hm^3.

(b) Set the Pre-release lead to 3, then 4, then 5 d, each time before Run or with Reset. Record the town peak and the days over the limit, and give the shortest lead that keeps the town under 450 m^3/s. *Checkpoint:* 525 and 4 days, 489 and 3 days, 425 and 0 days: 5 d.

(c) With the 5 d lead the Flood pool tile reads only about 15 hm^3 less than with 2 d. What did the extra lead buy, if not room in the lake? *Checkpoint:* 307 hm^3: the lead buys timing.

Read every value to $\pm 1\,m^3/s$ or $\pm 1\,hm^3$.

7.8 App challenge · 8 min

Load scenario 7, "The chance constraint: letting the wettest member go", default seed, with Scenario MPC as loaded. Run at 7 d/s: the case pauses itself at day 30.

(a) Read the energy on the Energy, supply and exceedance tile, the Town peak tile and the storage badge. *Checkpoint:* 67.6 GWh, 412 m^3/s, 1082 hm^3.

(b) Choose Chance-constrained MPC with Members held (k) at 6, press Reset and run to the pause again. Read the same three values. What did letting the wettest member go buy, and what did it cost the town? *Checkpoint:* 78.0 GWh, 340 m^3/s, 1082 hm^3.

(c) Set Members held to 5, press Reset and run again. Compare the energy with part (b), and say why the second member let go buys so little here. *Checkpoint:* 78.3 GWh, 345 m^3/s.

Read the energy to ± 0.1 GWh and the flows to $\pm 1\,\mathrm{m^3/s}$.

7.9 App challenge · 10 min

Load the Challenge, "a water year: two storms, a bust and a dry summer", default seed, with the Rule curve as loaded. Run at 30 d/s. The run pauses if the levee fails, press Run again, and it pauses at 30 September with the verdict on the Challenge target tile.

(a) Read the verdict, the days over the limit and the share of the perfect-foresight energy. *Checkpoint:* Not met, 8 days over, 75.6 % of 529.6 GWh, levee failed.

(b) Choose SDP policy, press Reset and run the year again. Read the same values and the end storage the tile names. *Checkpoint:* Not met, 9 days over, 89.5 % of 586.2 GWh, ending at 341 $\mathrm{hm^3}$.

(c) Choose Chance-constrained MPC with Members held at 6 (the suggested settings), press Reset and run the year. Read the verdict, the share, the end storage and the Town peak tile. *Checkpoint:* Met, 0 days over, 89.5 % of 590.8 GWh, ending at 305 $\mathrm{hm^3}$, town peak 425 $\mathrm{m^3/s}$.

(d) The SDP policy and the tree reach the same share. Why does only one of them pass, and why is the perfect-foresight energy different in each run? *Checkpoint:* The SDP floods the town; each bound ends at its own run's storage.

Read the shares to ± 0.1 %.

7.10 Application area · 3 min

In case 8 the tributary below the dam peaks at 394 $\mathrm{m^3/s}$. The scenario tree, with the ramp at 40 $\mathrm{m^3/s}$ a day, starts its cut before the rain and holds the town at 429 $\mathrm{m^3/s}$. With the environmental flow raised from 6 to 60 $\mathrm{m^3/s}$ the town reaches 459 $\mathrm{m^3/s}$ and is over the limit on 1 day.

(a) Using $Q_{town} = O + L$ and a river of at least Q_{env} below the diversion, find the lowest town flow the dam can give at the tributary's peak with each environmental flow, once the reach has passed the cut. *Checkpoint:* 400 $\mathrm{m^3/s}$ and 454 $\mathrm{m^3/s}$.

(b) What should an engineer tell the holder of the dam's permit?

7.11 Application area · 4 min

The stochastic dynamic program learns its statistics from a 100-year synthetic record whose storm days are one-day bursts of about 71 mm, 1.66 a year, while the cases declare three-day atmospheric rivers of 260 to 440 mm.

(a) What does this mismatch do to the policy in case 5 and in the challenge?

(b) Trained on full three-day storms, the program's gain falls from \$15.73 M to $-$\$4.20 M a year, and the tree that prices water with it draws the challenge's lake to 171 $\mathrm{hm^3}$ by 1 November and makes 66.3 % of its bound. Why was that record not adopted, and what would a utility check before trusting a value function?

7.12 Review the agent · 5 min

An AI assistant reviews the challenge results and writes: "The suggested chance-constrained tree makes 528.9 GWh, more than the 515.5 GWh perfect-foresight energy for the year from 760 $\mathrm{hm^3}$, which is 102.6 %. A real controller beat perfect foresight, so the bound is loose and the 80 % target is too easy." Find the flaw and say how you would check it in the app.