INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 19 Control of Biological Systems and Human Performance

CHAPTER 2

Artificial Pancreas Simulator for Direct Action and One-Sided Insulin Delivery

Pages 23–53

Edited by John D. Hedengren

How to cite this chapter. "Artificial Pancreas Simulator for Direct Action and One-Sided Insulin Delivery," in *Control of Biological Systems and Human Performance*, John D. Hedengren, Ed., vol. 19 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 2, pp. 23–53.

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

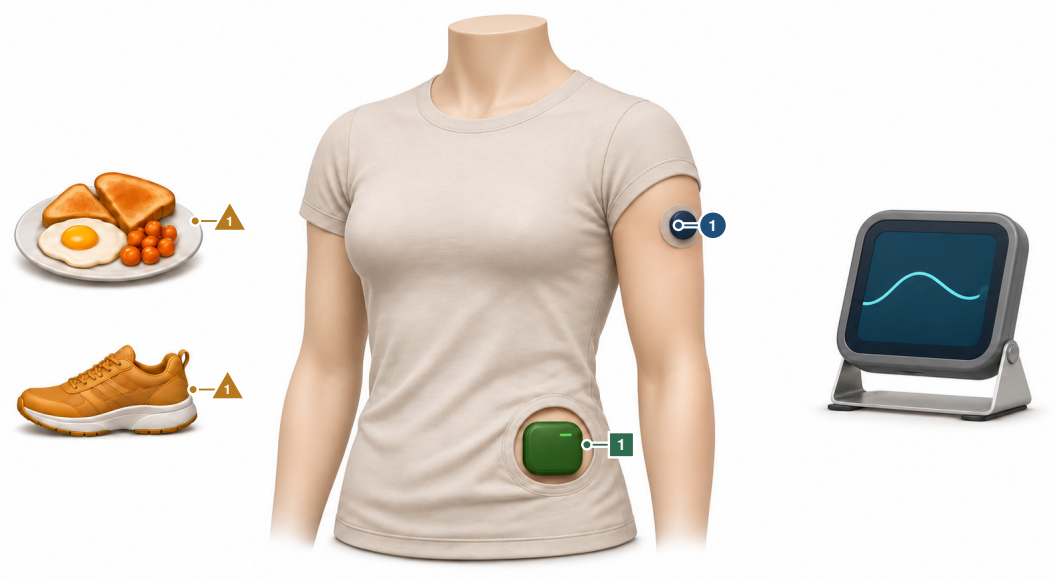
The app for this chapter
https://apmonitor.com/apps/diabetic/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v19-biological-human/#diabetic

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

CHAPTER 2

Artificial Pancreas Simulator for Direct Action and One-Sided Insulin Delivery



The controller moves 1 insulin pump worn on the abdomen

The controller reads 1 continuous glucose monitor on the arm

What disturbs it 1 a meal, or exercise

Notice that the monitor on the arm and the pump on the abdomen are separate devices, and that the pump can add insulin or stop but can never take it back. Illustration, not to scale.

IN PLAIN TERMS **Insulin That Can Only Push One Way**

The main idea. People with type 1 diabetes no longer make their own insulin, so a small pump worn on the body adds it. An automatic system reads the blood sugar from a sensor and sets the speed of the pump.

The challenge. Insulin only lowers blood sugar, and the pump cannot take back what it has given, so stopping it does not stop a fall for more than an hour. A low sugar is also more dangerous than a high one of the same size.

Why it matters. A score that treats highs and lows alike can pick a setting that is worse on the side that matters most.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history boxes of Section 2.2, the fact boxes of Sections 2.6 and 2.7, the try-it boxes of Section 2.8, and the key points
Study	65 min	The whole chapter with its pause points and the twelve exercises
Practice	15 min	The first page, Section 2.7, Section 2.10, Section 2.11 and the open challenge

The control question An insulin pump can only lower glucose, it stops at zero, and a low glucose costs more than a high one of the same size. How should a direct-acting PI loop be tuned when its actuator pushes one way and its error is not symmetric, and what does a predictive controller gain when it is told of a meal?

Learning objectives

- Derive the fasting steady-state map of the minimal model and the pump rate that holds a chosen glucose.
- Identify a first-order plus dead-time model from a pump step and explain why its dead time is the plasma insulin lag.
- Explain why the loop is direct acting and why no retuning shortens a low once the pump is at zero.
- Compare tunings on one-sided excursion areas, and show where the symmetric integrated error misranks them.
- Separate what a meal announcement, a soft glucose band and a bounded meal estimate each change in a predictive controller.

Prerequisites Sections 1.3 (loop action), 1.4.2 (compartments), 1.4.6 (one-sided actuators and asymmetric cost), 1.6.2 (velocity form) and 1.6.4 (scenarios and data). Section 1.6.3 is needed only for the predictive controller and the estimators.

The app https://apmonitor.com/apps/diabetic/

See also Volume 2, Chapter 2 for velocity-form PID with back-calculation and a relay test placed by its switching band, and Volume 2, Chapter 3 for anti-windup at a bound the supply sets. Volume 9, Chapter 9 for control against a band rather than a setpoint, and Volume 9, Chapter 5 for a predictive controller weighed against releases ahead of published demand.

2.1 Introduction

Most regulatory loops taught in a first course have an actuator that works in both directions. A heater can be turned up or down, a valve can be opened or closed, and the error signal that drives the actuator is scored with a metric that treats an excursion above setpoint and an excursion below setpoint as equally bad. Subcutaneous insulin delivery satisfies neither condition. An infusion pump adds insulin, and insulin removes glucose, so the process gain from the manipulated variable to the controlled variable is negative and the controller must be direct acting in the ISA sense, with an output that rises as the measurement rises [1]. In the convention of Section 1.3 the loop has the action $a = -1$. The pump has no negative setting, because there is no way to withdraw insulin that has already reached the tissue, so the manipulated variable is bounded below at zero and the loop has authority in one direction only. The cost of the two error signs is also different. A sustained high glucose is a long-term risk, while a sufficiently deep low is an immediate one.

The dynamics compound the actuator constraint. A meal delivers glucose to the blood within an hour, while the insulin path from the pump runs through a plasma compartment and then through a remote action compartment whose time constant in this model is $1/p_2 = 81.3\,\text{min}$. The disturbance is therefore faster than the actuator, and the same lag that delays the effect of switching the pump on also delays the effect of switching it off. Once glucose is falling, shutting the pump down does not stop the fall, because the insulin already in the action compartment continues to clear glucose for more than an hour.

Teaching simulators for glucose control frequently present the plant as a first-order lag with a fixed gain and an unbounded actuator. However, a symmetric linear plant with a two-sided actuator removes exactly the three properties that make this loop worth studying, because it hides the sign reversal, it hides the shutdown floor, and it makes a symmetric error integral a legitimate score. This chapter studies a simulation of Type 1 diabetic glucose regulation that retains the sign reversal, the zero floor on the pump and the asymmetric cost of error, and it quantifies what each of them does to the tuning of a proportional-integral controller. Neither the model nor the control algorithms are new.

Section 2.2 traces automated insulin delivery from bedside instruments to hybrid closed loop. Section 2.3 sets out the six-state model, its disturbances, instrument, actuator and timing, with

every relation and constant a reimplementation needs. Section 2.4 derives the closed-form fasting steady-state map of the model, shows that its slope falls by a factor of 1.75 between infusion rates of 3.0 and 4.0 µU/min, and shows that the apparent dead time of the identified first-order plus dead-time model is the plasma insulin clearance time constant $1/k_e$. Section 2.5 describes the direct-acting PID, a linear model predictive controller that can replace it, two moving horizon estimators and a relay autotune. Section 2.6 defines the one-sided scores, and Section 2.7 measures the loop over a full simulated day. In a sweep of integral gain over that day a fivefold change in integral gain moves the integrated absolute error by 8.7 % and the glucose nadir by 6.4 mg/dL, which quantifies the claim that a symmetric metric is not a suitable objective for this loop. Section 2.8 compares the predictive controller with the PID to separate what a meal announcement and a soft glucose band each change, and Section 2.9 compares the estimators and the constants a relay test writes. Section 2.10 describes the sensor, pump and controller platform of real systems, and Sections 2.11 to 2.13 state the limitations of the model and what the results mean beyond this plant.

The chapter rests on standard process control fundamentals. The model is a set of compartment mass balances with first-order exchange between pools, and the loop behavior follows from the time-scale separation between a fast glucose balance and two slower insulin lags [1, 2]. The principles it applies are controller action against a negative process gain, actuator saturation and anti-windup, identification of a first-order plus dead-time model from a step test, steady-state offset and its removal by bias or integral action, and performance evaluation under sampled and noisy measurements [3]. Loops with one-sided actuators appear industrially wherever the actuator works only in the safe direction, such as pressure relief and emergency cooling, so the tuning question posed here is not specific to glucose. Section 1.4.6 sets out the general case, and Chapter 3 meets the same one-way pump with an anesthetic.

2.2 Background

Automated insulin delivery began as bedside instrumentation rather than as a wearable device [4, 5], as the first history box below tells.

> ⧖ HISTORY **Bedside Machines Close the Loop (1964 to 1979)**
>
> Automatic control of blood glucose began at the hospital bedside. After pioneering work in 1964, five teams reported closed-loop results between 1974 and 1978. Their machines measured glucose in blood from a vein and infused insulin, and glucose as well, into the vein. Their control laws were of the proportional-derivative kind, which computed the insulin dose from the glucose value and its rate of change. One of these systems became the first commercial device in 1977. The intravenous route could not leave the hospital, but the machines proved that external control of glucose was feasible. In 1979 two studies showed that insulin could be delivered continuously under the skin, the route a wearable pump uses today [6].

The move from bedside to wearable removed the glucose infusion channel and replaced intravenous measurement with a subcutaneous sensor, which added transport lag on both sides of the loop. Feasibility of a fully automatic subcutaneous system was established with a proportional integral derivative algorithm operating on a continuous glucose monitor, and the reported result was glucose inside 70 mg/dL to 180 mg/dL for 75 % of the time under closed loop against 63 % open loop [7]. Model predictive formulations followed, first as simulation studies [8] and then in clinical protocols [9]. A review of the algorithms available at that stage set out why the problem resists a textbook solution, naming the sensor lag, the actuator lag, the unmeasured meal and the impossibility of negative insulin as the four difficulties any candidate must address [10].

Safety layers rather than better regulators account for much of the progress since. Constraining the controller by the insulin already delivered, an approach usually called insulin on board, prevents the algorithm from stacking correction doses whose effect has not yet appeared, and it does so as an explicit constraint inside the optimization rather than as a detuning [11]. Testing moved to virtual

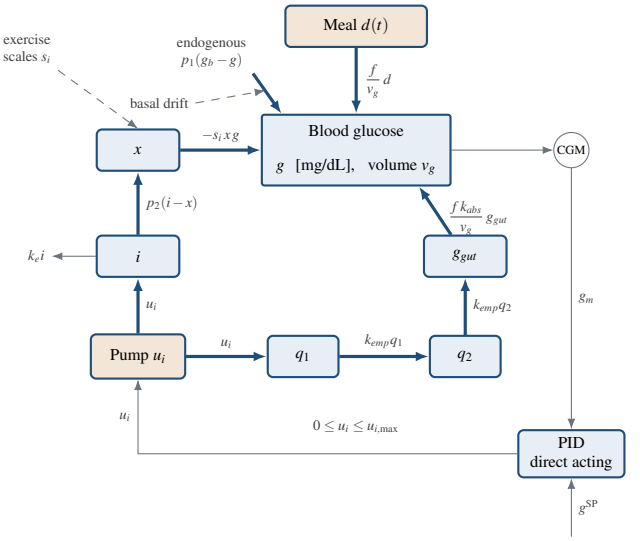


Figure 2.1 Compartments and flows of the expanded Bergman minimal model as implemented. The infusion rate u_i enters both the plasma insulin pool and the gastric transit chain, and the meal enters the blood directly through $f d / v_g$, which is how the published assignment equations are written.

patients, as the second history box tells. Systems that automate basal delivery while leaving the meal bolus to the user, usually called hybrid closed loop, reached clinical practice on that foundation [12], and a six-month randomized multicenter trial reported time in the 70 mg/dL to 180 mg/dL range of 71 % under closed-loop control against 59 % in the control group [13].

> ⧖ HISTORY **Virtual Patients in Place of Animal Trials (2009)**
>
> A report of 2009 described a simulator for testing closed-loop algorithms before any clinical study. It held 300 virtual subjects, 100 adults, 100 adolescents and 100 children, built from the data of real people to span the variability seen across the population with type 1 diabetes, together with models of sensor error and of insulin delivery by pump [14]. The United States Food and Drug Administration accepted it as a substitute for animal trials in the preclinical testing of closed-loop control strategies. Because an experiment in the computer costs a fraction of the time and money of an animal trial, it became a primary test bed for new control algorithms [6]. The simulator of this chapter holds one virtual subject, as Section 2.3 states.

2.3 Process Model

The simulator implements the expanded Bergman minimal model exactly as it is written in the APMonitor diabetic blood glucose assignment. The original minimal model was built to estimate insulin sensitivity from an intravenous glucose tolerance test using three states, glucose, remote insulin action and plasma insulin [15]. The expansion adds a three-compartment chain that delays material on its way into the blood, which is the same structural device used by larger meal models to reproduce the shape of a postprandial glucose curve [16]. Figure 2.1 sets out the compartments, the flows between them, the pump and meal inputs, and the point of measurement.

2.3.1 Roles of the Signals

The controlled variable is blood glucose as reported by the continuous glucose monitor, the manipulated variable is the pump infusion rate, and the disturbances are the meal appearance rate, the exercise term and the basal drift. The complete classification is given in Table 2.1. The distinction

Table 2.1 Classification of Signals in the Glucose Loop

Symbol	Role	Unit	Range
g	state, CV source	mg/dL	≥ 0
g_m	measurement (CV)	mg/dL	≥ 0
x	state, remote action	µU/mL	≥ 0
i	state, plasma insulin	µU/mL	≥ 0
q_1	state, transit pool 1	model units	≥ 0
q_2	state, transit pool 2	model units	≥ 0
g_{gut}	state, gut glucose	model units	≥ 0
u_i	MV, infusion rate	µU/min	0 to $u_{i,\max}$
d	DV, glucose appearance	model units	≥ 0
e_x	DV, exercise	%	0 to 100
b_d	DV, basal drift	%	-30 to $+30$
g^{SP}	setpoint	mg/dL	40 to 250

between the true state g and the reported measurement g_m matters throughout, because the controller only ever sees g_m while every performance number quoted in Section 2.7 is computed from g.

2.3.2 Governing Equations

The six state equations are written with every parameter in reciprocal minutes and every right-hand side multiplied by 60, so the independent variable of the simulation is time in hours while the parameter values stay in the units in which they are published. The set is given in Equation (2.1).

$$\frac{\mathrm{d}g}{\mathrm{d}t} = 60\left[-p_1(g-g_b) - s_i x g + \frac{f k_{abs}}{v_g} g_{gut} + \frac{f}{v_g} d\right] \tag{2.1a}$$

$$\frac{\mathrm{d}x}{\mathrm{d}t} = 60\, p_2\,(i-x) \tag{2.1b}$$

$$\frac{\mathrm{d}i}{\mathrm{d}t} = 60\,(-k_e i + u_i) \tag{2.1c}$$

$$\frac{\mathrm{d}q_1}{\mathrm{d}t} = 60\,(u_i - k_{emp} q_1) \tag{2.1d}$$

$$\frac{\mathrm{d}q_2}{\mathrm{d}t} = -60\,k_{emp}(q_2 - q_1) \tag{2.1e}$$

$$\frac{\mathrm{d}g_{gut}}{\mathrm{d}t} = 60\,(k_{emp} q_2 - k_{abs} g_{gut}) \tag{2.1f}$$

where g is blood glucose, x is remote insulin action, i is plasma insulin, q_1 and q_2 are the two transit pools, g_{gut} is the gut glucose pool, u_i is the infusion rate, d is the glucose appearance rate, g_b is the basal glucose parameter, p_1 is the glucose clearance rate, p_2 is the insulin action rate, s_i is insulin sensitivity, k_e is plasma insulin clearance, k_{abs} is gut absorption, k_{emp} is gastric emptying, f is the fraction of appearing glucose that reaches the blood and v_g is the glucose distribution volume. Every state is clamped at zero after each integration step, because g, i and the three pools are concentrations and masses that cannot go negative.

Two features of Equation (2.1) are consequences of following the published assignment rather than of a modeling choice made here. The transit chain (2.1d) to (2.1f) is driven by the infusion rate u_i and not by the ingested carbohydrate, so q_1, q_2 and g_{gut} track the pump. The meal reaches the blood only through the direct term $f d / v_g$ in (2.1a). A reimplementation that routes the meal into q_1 instead is a different model and will not reproduce the numbers reported here.

The nonlinearity of the plant lives in the single product $s_i x g$ in (2.1a). Glucose clearance by insulin is proportional to both the insulin action and the glucose concentration, so the incremental effect of

insulin falls as glucose falls, and a loop that recovers from a postprandial peak can still overshoot below the setpoint during the descent.

2.3.3 Meals, Exercise and Basal Drift

The glucose appearance rate is the sum of a fasting baseline and the contribution of every meal that has started, as shown in Equation (2.2).

$$d(t) = d_0 \left(1 + \frac{b_d}{100}\right) + \sum_k m(t - t_k, A) \tag{2.2a}$$

$$m(\tau, A) = \begin{cases} 0, & \tau < 0 \\ A\,\tau, & 0 \le \tau < 1\,\mathrm{h} \\ A\,\mathrm{e}^{-(\tau-1)}, & \tau \ge 1\,\mathrm{h} \end{cases} \tag{2.2b}$$

where $d_0 = 1000$ is the fasting baseline in model units, b_d is the basal drift in percent, t_k is the start time of meal k, and A is the meal amplitude. Each meal ramps linearly to its full amplitude over one hour and then decays with a one-hour time constant, so the appearance rate peaks exactly one hour after the meal begins. The scheduled plan places meals at 08:00, 12:00 and 18:00 on the simulated clock. The schedule is off in every case as the application loads it and is switched on as a disturbance setting, and an additional unscheduled meal can be injected at any simulated time with the current amplitude. Three of the shipped cases script one such meal of amplitude 1800 at $t = 2\,\mathrm{h}$.

The baseline term d_0 is not part of the published assignment code. It is required here because the published initial condition is a steady state only when the appearance rate carries that value. With $d = 1000$ and $u_i = 3.0\,\mu\mathrm{U/min}$ the steady state of Equation (2.1) lies within 0.005 of every published value, and the reference case starts at that steady state, where no state moves, which is what makes a fasting step test meaningful.

Exercise scales insulin sensitivity by $s_i(1 + e_x/100)$, so the insulin already present clears glucose faster without any change in the pump rate. Basal drift scales only the baseline term of d and stands in for a change in endogenous glucose production. Neither disturbance is available to the controller, which sees only g_m.

2.3.4 Instrument and Actuator

The continuous glucose monitor adds a zero-mean Gaussian deviate of standard deviation σ to the true glucose and clamps the result at zero, as shown in Equation (2.3).

$$g_m = \max(0,\ g + \sigma\,\mathcal{N}(0,1)) \tag{2.3}$$

The deviate is drawn once per control update by the Box-Muller transform from a generator that is reseeded with a fixed seed whenever a case starts, so a run with noise repeats exactly. There is no sensor lag and no calibration error in the instrument model, so the subcutaneous transport delay of a real monitor is absent. The pump applies the commanded rate without lag, without a minimum delivery increment and without a delivery quantum, and it is clamped to $0 \le u_i \le u_{i,\max}$ with $u_{i,\max} = 10\,\mu\mathrm{U/min}$ by default. The manual pump setting follows the same maximum.

2.3.5 Parameters, Limits and Initial Conditions

The constants of Equation (2.1) and the settings that the application exposes are listed in Table 2.2. Five model parameters, the pump maximum and the sensor noise are adjustable inputs, and the remainder are fixed. The insulin distribution volume $v_i = 12\,\mathrm{L}$ appears in the published parameter list but in no equation of the implemented model, so it has no effect on any trajectory.

Equation (2.1b) equates x with i, so the remote action state carries the units of plasma insulin and s_i carries mL per µU per minute in order that $s_i x g$ have the units of g divided by time. The three

Table 2.2 Model Parameters, Defaults and Adjustable Ranges

Symbol	Meaning	Value	Range in settings
g_b	basal glucose parameter	291 mg/dL	fixed
p_1	glucose clearance rate	0.0317 min^{-1}	fixed
p_2	insulin action rate	0.0123 min^{-1}	fixed
f	fraction reaching blood	0.8	fixed
v_i	insulin volume (unused)	12 L	fixed
s_i	insulin sensitivity	0.029	0.01 to 0.06
k_e	plasma insulin clearance	0.09 min^{-1}	0.03 to 0.2
k_{abs}	gut absorption rate	0.012 min^{-1}	0.004 to 0.03
k_{emp}	gastric emptying rate	0.18 min^{-1}	0.05 to 0.35
v_g	glucose volume	12 L	5 to 25
$u_{i,\max}$	pump upper limit	10 µU/min	5 to 20
σ	CGM noise standard dev.	1.5 mg/dL	0 to 5
d_0	fasting appearance baseline	1000	fixed
A	meal amplitude	1200	0 to 2400
e_x	exercise	0 %	0 to 100
b_d	basal drift	0 %	−30 to +30
g^{SP}	glucose setpoint	80 mg/dL	40 to 250
K_c	proportional gain	0.3	0 to 1.5
τ_I	integral time	3 h	0.002 to 48[a]
τ_D	derivative time	0 h	0 to 16.8
T_p	MPC prediction horizon	6 h	1 to 12
λ	MPC move suppression	3	0 to 30
u_{bias}	controller bias	u_i at start	not set directly

[a] The top of the range, 48 h, switches integral action off.

transit states carry the infusion rate integrated over time, and the coefficient $f k_{abs}/v_g$ that returns g_{gut} to the blood is applied as a numerical factor rather than as a dimensionally consistent conversion, which is how the published assignment writes it.

The published initial condition is $[g, x, i, q_1, q_2, g_{gut}] = [76.22, 33.33, 33.33, 16.67, 16.67, 250.0]$ with the pump at the basal rate 3.0 µU/min. Setting the derivatives of Equation (2.1) to zero at that infusion rate gives $i = x = u_i/k_e = 33.33$, $q_1 = q_2 = u_i/k_{emp} = 16.67$, $g_{gut} = u_i/k_{abs} = 250$ and $g = 76.216\,\text{mg/dL}$ from Equation (2.4), so the published vector is the steady state to the digits it is written with. The application starts every case at a steady state computed this way. In manual mode it is the steady state of the pump rate in force, and in automatic mode it is the steady state of the setpoint in force, with the pump at the rate of Equation (2.5) that holds it. The parameters and disturbances in force enter the calculation, so a setting chosen before a run starts changes the starting state and is not a step. The reference case therefore starts at $[76.216, 33.333, 33.333, 16.667, 16.667, 250.0]$, and every automatic case starts at 80 mg/dL with the pump at 2.853 µU/min.

2.3.6 Integration and Timing

The engine of Section 1.6.1 advances the state by explicit Euler on a fixed step of $1/500$ hours, or 7.2 sec. The controller and the continuous glucose monitor execute on every 84th step, so the control period is $\Delta t = 0.168\,\text{h}$ (10.08 min), which is the period the application declares. A nominal ten-minute period would not be a whole number of steps. When a control period has elapsed, the measurement is read and the controller executes first, the state is then advanced by one Euler step, and each state is clamped at zero. A change the operator makes to a disturbance or a model parameter reaches the plant at the next controller execution, while an unscheduled meal enters at once. The speed setting spans 0.5 to 10 simulated hours per real second and leaves the trajectory unchanged. At the default $1\times$ a twenty-four hour day takes 24 sec.

The run is logged once per control period, and the exported time column is the model clock that the meal schedule uses, on which the 08:00 meal begins at $t = 8\,\mathrm{h}$. A row stamped t holds the state at t, the glucose reading on which the controller acts at t and the infusion rate held over the control period that ends at t. Every time quoted in this chapter is model time.

2.3.7 Assumptions

The model is a lumped compartment description with a small number of pools, and each of the following simplifications has a regime in which it fails.

One well-mixed glucose pool. Blood glucose is a single concentration in a fixed volume v_g, so there is no distinction between plasma and interstitial glucose. A real continuous glucose monitor reads interstitial fluid and lags plasma by five to fifteen minutes, and that lag is absent here.

Insulin action through one remote compartment. The path from plasma insulin to glucose clearance is a single first-order lag with rate p_2. Subcutaneous depot kinetics, which dominate the delay in a real wearable pump, are not represented at all, so the simulated actuator is faster than a real one.

Linear endogenous balance. The term $-p_1(g - g_b)$ is linear and restores glucose toward $g_b = 291\,\mathrm{mg/dL}$. That value is a fitting parameter of the minimal model and not a physiological setpoint. The model carries no counter-regulatory response, so there is no glucagon action, no hepatic glucose release and no floor under a low excursion. With the pump at zero the fasting steady state is 2394 mg/dL, which is not a physiological number and shows that the model is calibrated for a band around the basal point and not across its full mathematical range.

Stylized disturbances. The amplitude A is in the units of d and has no gram equivalent, and the ramp and decay shape is fixed, so a meal of the same carbohydrate content but a different fat content cannot be expressed. A single scaling of s_i captures the direction of the exercise effect and not its mechanism, and it applies instantly rather than building and decaying around the activity.

Fluxes that are not physiologically scaled. At the basal operating point the four terms of (2.1a) evaluate to $+6.81$, -73.67, $+0.20$ and $+66.67$ in mg/dL/min, and they sum to within 0.02 mg/dL/min of zero. Individual terms of that size are one to two orders of magnitude above physiological glucose turnover, so the parameter set reproduces the net dynamics around its own operating point rather than a flux balance whose individual terms can be read as rates of appearance and disappearance.

One patient. The parameters describe a single virtual subject. Between subject variability and within subject variability across a day are what in silico protocols exist to test [14], and neither is represented by a single parameter set.

2.4 Steady-State and Dynamic Analysis

2.4.1 Steady-State Map

Holding the infusion rate constant and setting the derivatives of Equation (2.1) to zero gives $i = x = u_i/k_e$, $q_1 = q_2 = u_i/k_{emp}$ and $g_{gut} = u_i/k_{abs}$, and substituting those into (2.1a) collapses the balance to a single algebraic relation, shown in Equation (2.4).

$$g_{ss}(u_i) = \frac{p_1 g_b + \dfrac{f}{v_g} u_i + \dfrac{f}{v_g} d}{p_1 + \dfrac{s_i}{k_e} u_i} \tag{2.4}$$

Equation (2.4) is a hyperbola in u_i, not a line, and its slope is $(BC - AD)/(C + Du_i)^2$ with $A = p_1 g_b + fd/v_g$, $B = f/v_g$, $C = p_1$ and $D = s_i/k_e$. At the fasting appearance rate $d = 1000$ the numerator of that slope is -24.45, so the tangent gain is $-24.5\,\mathrm{mg/dL}$ per $\mu\mathrm{U/min}$ at $u_i = 3.0\,\mu\mathrm{U/min}$ and $-14.0\,\mathrm{mg/dL}$ per $\mu\mathrm{U/min}$ at $u_i = 4.0\,\mu\mathrm{U/min}$, a factor of 1.75 across one unit of infusion rate. The map and the effect of a 30 % basal drift on it are shown in Figure 2.2.

Inverting Equation (2.4) for the infusion rate that holds a chosen glucose gives Equation (2.5), which is the bias a positional controller would need to hold that glucose without integral action.

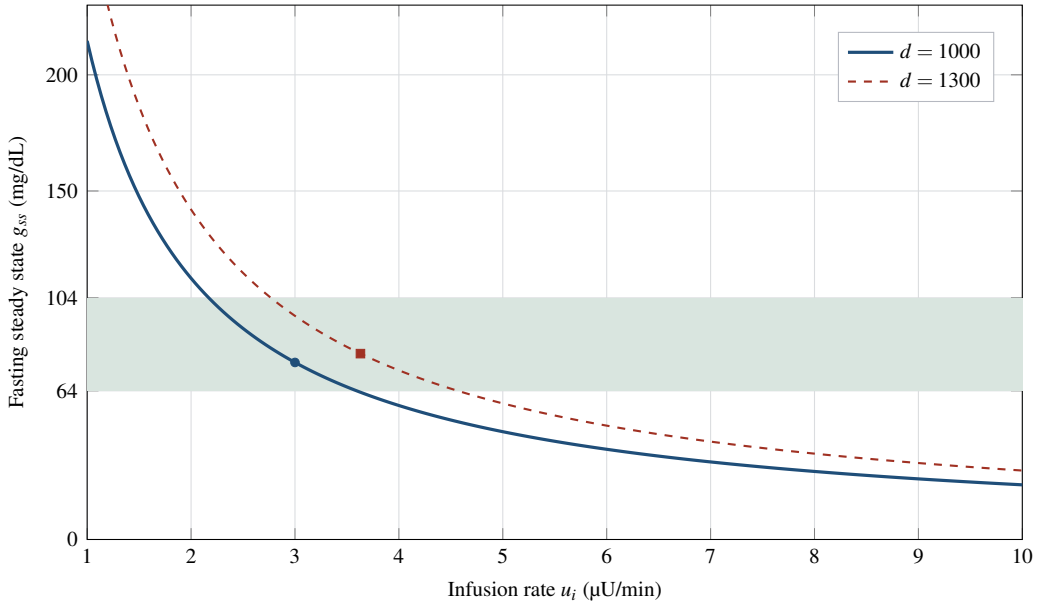


Figure 2.2 Fasting steady-state glucose against infusion rate from Equation (2.4). The shaded strip is the 64 mg/dL to 104 mg/dL teaching band of Section 2.6, the circle is the basal operating point and the square is the infusion rate that holds 80 mg/dL after a 30 % basal drift.

$$u_{bias}(g) = \frac{p_1(g_b - g) + \dfrac{fd}{v_g}}{\dfrac{s_i}{k_e}g - \dfrac{f}{v_g}} \tag{2.5}$$

At $g = 76.2\,\mathrm{mg/dL}$ and $d = 1000$ this returns 3.00 µU/min, at the default setpoint of 80 mg/dL it returns 2.85 µU/min, and after a 30 % basal drift it returns 3.63 µU/min. The application starts the patient at the pump rate of Equation (2.5) for the setpoint and disturbances in force when a run starts, and the controller takes that rate as its bias, 3.00 µU/min at 76.22 mg/dL and 2.85 µU/min at 80 mg/dL. A drift chosen before the run starts enters that calculation and returns 3.63 µU/min, but a drift that arrives during a run is unmeasured, so its value cannot be computed in advance and must be recovered by feedback.

2.4.2 Time Constants of the Linearized Model

Linearizing Equation (2.1) about the basal point separates the insulin path into three first-order lags in series. The remote action follows plasma insulin as an effect site follows plasma in Section 1.4.2. The plasma pool (2.1c) contributes $1/k_e = 11.1\,\mathrm{min}$, the remote action pool (2.1b) contributes $1/p_2 = 81.3\,\mathrm{min}$, and the glucose balance itself contributes $1/(p_1 + s_i\bar{x}) = 1.0\,\mathrm{min}$ at $\bar{x} = 33.33$. The dominant time constant is therefore an insulin compartment and not glucose, and the glucose pool is fast enough to be invisible in the response. An identification that fits a first-order plus dead-time form to this cascade should return a time constant close to 81 min and an apparent dead time close to 11 min, and Section 2.4.3 confirms both.

2.4.3 Step-Test Identification

A fasting step test runs from the reference operating point, which the application loads in manual mode with the scheduled meals off and the sensor noise at zero. The pump is held at 3.0 µU/min

with glucose flat at 76.22 mg/dL, and the operator moves it to 4.0 µU/min at $t = 2$ h, which the plant receives at the controller execution at 2.016 h. Glucose falls from 76.22 mg/dL to 57.67 mg/dL over the following twelve hours, which is a secant process gain of $K_p = -18.55$ mg/dL per µU/min. The analytic value from Equation (2.4) is $57.669 - 76.215 = -18.546$, so the integration reproduces the steady-state map to four significant figures. The case starts at the exact steady state, so the unstepped run does not move and no drift correction is needed. The application's own open-loop step case starts at 80 mg/dL with the pump at 2.85 µU/min and moves it by 2.0 µU/min at 2 h, and glucose settles at 47.8 mg/dL. Its secant gain of −16.1 mg/dL per µU/min is 60 % of the tangent gain of −27.0 mg/dL per µU/min at its starting rate, which is the curvature of Figure 2.2 measured directly.

The response reaches 63.2 % of its change at $t_{63} = 1.34$ h after the step. A two-point fit through the 28.3 % and 63.2 % crossings, which occur at 0.553 h and 1.340 h, returns $\tau_p = 1.18$ h and $\theta_p = 0.16$ h [1]. The fitted dead time of 0.16 h is 9.6 min against a plasma clearance time constant $1/k_e$ of 11.1 min, and the fitted time constant of 1.18 h is 70.8 min against a remote action time constant $1/p_2$ of 81.3 min. The apparent dead time of the identified model is not a transport delay. It is the second lag of the insulin chain reported as a delay by the first-order fit.

Internal model control correlations for a plant of this form give $K_c = \tau_p/(|K_p|(\tau_c + \theta_p))$ with $\tau_I = \tau_p$ [17]. A conservative closed-loop time constant $\tau_c = \tau_p$ returns $K_c = 0.048$, and an aggressive $\tau_c = \theta_p$ returns $K_c = 0.20$. The shipped default of $K_c = 0.3$ is more aggressive than either, and the shipped integral time $\tau_I = 3.0$ h, which is an integral gain $K_I = K_c/\tau_I = 0.1$ in µU/min per mg/dL per hour, is 2.5 times the identified time constant rather than equal to it. Section 2.7 shows why the integral time is deliberately slow.

Four models appear in the analysis and are kept distinct. The six-state set of Equation (2.1) is the simulation model. The first-order plus dead-time fit is the identified control model and is used only to produce tuning constants. The PID itself carries no internal model. Its starting bias is the pump rate in force when the run starts, read from the pump, which on a steady start is the value of Equation (2.5) at the setpoint in force. The predictive controller of Section 2.5.4 plans on the fourth, a linear model fitted separately at 80 mg/dL.

2.4.4 Actuator Authority

The manipulated variable is bounded on both sides, but the two bounds are not alike. The upper bound $u_{i,\max}$ is a device setting, and raising it provides faster rejection of a meal. The lower bound is zero and is a property of the physics, because a pump can add insulin and cannot remove it. Stopping delivery does not remove the insulin already in the plasma and action compartments, and that stored quantity decays with the time constants of (2.1c) and (2.1b). The slowest of those is $1/p_2 = 81.3$ min, the time constant that sets the slow way down of Equation (1.9), and it is the quantitative statement of why a low excursion is not recoverable by the actuator and why deployed systems suspend delivery on a prediction rather than on a measurement [10]. Constraining the controller by the insulin already on board addresses the same limit from the opposite side [11].

2.5 Controller Design

2.5.1 Velocity Form and Direct Action

The pump is driven by the velocity-form PID of Section 1.6.2, with the proportional action on the error and the derivative taken on the measurement. It executes every $\Delta t = 0.168$ h, and each execution computes an increment rather than the pump rate itself. The algorithm as coded for this loop is given by Equations (2.6) to (2.10).

$$e_k = a\left(g^{\mathrm{SP}} - g_{m,k}\right), \qquad a = -1 \tag{2.6}$$

$$\Delta u_k^P = K_c\left(e_k - e_{k-1}\right) \tag{2.7}$$

$$\Delta u_k^I = \frac{K_c}{\tau_I}\, e_k\, \Delta t \tag{2.8}$$

$$\Delta u_k^D = -K_c \tau_D\, a\, \frac{g_{m,k} - 2g_{m,k-1} + g_{m,k-2}}{\Delta t} \tag{2.9}$$

$$\begin{aligned} s_k &= s_{k-1} + \Delta u_k^P + \Delta u_k^D \\ c_k &= u_{bias} + s_k + I_{k-1} \\ u_k &= \min(u_{i,\max}, \max(0, c_k)) \\ I_k &= I_{k-1} + \Delta u_k^I + \frac{\Delta t}{\tau_I}\left(u_k - c_k\right) \end{aligned} \tag{2.10}$$

where a is the controller direction, s_k accumulates the proportional and derivative increments without limit, c_k is the infusion rate the algorithm demands, u_k is that demand clipped to the pump range and applied as u_i until the next execution, and I is the integral state. The integral used at execution k was accumulated through execution $k-1$, and u_{bias} is the infusion rate when the run starts. In the notation of Equations (1.13) to (1.15), s_k is S_k, c_k is v_k, u_{bias} is u_b and the tracking time T_t equals τ_I. The glucose loop is constructed with $a = -1$, so a measurement above setpoint raises the pump, which is direct action in the ISA sense, and every tuning constant entered in the application is a positive magnitude. Section 1.6.2 explains why a setpoint change kicks the demand by K_c times the change and why a change of K_c acts at once on the whole error.

Pause. Try Exercise 2.1 before reading on.

2.5.2 Initialization and Mode Transfer

A velocity-form controller has no absolute output of its own, so the value at which the recursion starts is the bias of the equivalent positional law, which the application sets to the infusion rate at the start of the run. A case that opens in automatic mode starts at the steady state of its setpoint with $I = 0$, $s = 0$ and a zero previous proportional term, so its first execution contributes the whole proportional term, as shown in Equation (2.11).

$$u_0 = \min(u_{i,\max}, \max(0, u_{bias} + K_c\, e_0)) \tag{2.11}$$

Away from the pump limits the demand then equals the positional law $u_{bias} + K_c e_k$ plus the accumulated integral and derivative contributions. Manual tracking makes transfer bumpless in both directions, as Section 1.6.2 describes. A mode or setpoint chosen before the run starts is part of the starting state rather than a step, and the same choice carried through Reset gives the same run. The closed-loop runs of Section 2.7 choose automatic mode and the 80 mg/dL setpoint in the reference case before the run starts, so they start at 80 mg/dL with the pump and the bias at the 2.85 µU/min that Equation (2.5) gives for that setpoint. With no sensor noise the starting error is zero, and nothing moves until the first meal.

2.5.3 Anti-Windup at the Actuator Limits

Anti-windup is the back-calculation term of Equation (2.10), which corrects the integral alone, as in Section 1.6.2 [3]. Volume 2, Chapter 3 studies the same scheme at a bound set by the supply rather than by a valve. At the shipped $\tau_I = 3$ h each execution removes 5.6 % of the clipped excess, so a demand that has run well past a limit takes several executions to return into range. With the integral switched off at the top of the τ_I range, I keeps its value and no back-calculation is applied.

Against a one-sided actuator this behavior is not symmetric in its consequences. At the upper limit the clamp is protective, because the pump is already delivering as fast as the device allows and the excess demand is discarded rather than stored. At the lower limit the clamp marks the point at which the loop has no remaining authority. The controller output is zero, the error is still telling it to reduce insulin delivery further, and there is no further reduction available. Around the deepest point of each noise-free meal day in Section 2.7 the pump has already been cut below 0.9 µU/min, less than a third of the basal rate, and after the exercise step it sits at zero for 0.84 h while glucose recovers. Once the command is at zero, no retuning of K_c, τ_I or τ_D can shorten the excursion. Tuning can only change how much insulin is already stored when the fall begins, and the quantities that set how quickly the stored insulin clears are plant parameters.

2.5.4 Model Predictive Control

In automatic mode the pump can be driven instead by a linear model predictive controller (MPC), chosen with the Controller option, which Section 1.6.3 describes in general, while the PID stays in charge by default in all cases but cases 9 and 10 of Table 2.3. At every execution, so with a decision period $T_s = \Delta t = 0.168\,\mathrm{h}$, the MPC plans the infusion rate over a horizon, keeps only the first move and replans at the next execution, the receding-horizon practice of industrial predictive controllers [18]. It plans on Equation (2.12), in deviations from the design point $\bar{g} = 80\,\mathrm{mg/dL}$, $\bar{u} = 2.853\,\mu\mathrm{U/min}$.

$$\begin{aligned} \tau_1 \frac{\mathrm{d}\xi_1}{\mathrm{d}t} &= -\xi_1 + K_u\,(u_i - \bar{u}) \\ \tau_2 \frac{\mathrm{d}\xi_2}{\mathrm{d}t} &= -\xi_2 + \xi_1 + \delta \\ \tau_m \frac{\mathrm{d}\xi_3}{\mathrm{d}t} &= -\xi_3 + K_m\,\hat{d}(t), \qquad \hat{g} = \bar{g} + \xi_2 + \xi_3 \end{aligned} \tag{2.12}$$

Here ξ_1 is the slow insulin state, ξ_2 the glucose deviation from insulin and from δ, an unmeasured disturbance in mg/dL that enters at glucose, where a meal acts, ξ_3 the glucose from announced meals of appearance rate $\hat{d}$, and $\hat{g}$ the predicted glucose. Pump steps of $\pm 0.2\,\mu\mathrm{U/min}$ about $\bar{u}$, meals and noise off, gave $K_u = -27.16\,\mathrm{mg/dL}$ per µU/min, $\tau_1 = 1.358\,\mathrm{h}$ and $\tau_2 = 0.1876\,\mathrm{h}$, which lie 0.5, 0.2 and 1.3 % above the tangent gain of Section 2.4.3, $1/p_2$ and $1/k_e$, and the model misses their averaged response by 0.012 mg/dL root mean square, 0.21 % of its span. One meal of amplitude 1800 with the pump held gave $K_m = 0.0703$ mg/dL per model unit and $\tau_m = 0.0155\,\mathrm{h}$, with a residual of 0.24 mg/dL on a 124.7 mg/dL rise. Both paths are frozen at 80 mg/dL, whereas the plant's insulin acts harder at higher glucose, Equation (2.4) giving tangent gains of −43.1 and −60.9 mg/dL per µU/min at 101 and 120 mg/dL, and a meal raises glucose less as insulin action grows.

A steady-state Kalman filter of Equation (2.12), sampled with a zero-order hold at Δt and extended with δ as a random walk, the disturbance arrangement that leaves a predictive controller without steady-state offset [19], is updated at each execution under either controller, as Equation (2.13) gives.

$$\begin{aligned} \bar{\xi}_k &= \Phi\hat{\xi}_{k-1} + \gamma\,(u_{k-1} - \bar{u}) + \gamma_\delta\,\hat{\delta}_{k-1} + \gamma_m \hat{d}_{k-1/2} \\ \begin{bmatrix} \hat{\xi}_k \\ \hat{\delta}_k \end{bmatrix} &= \begin{bmatrix} \bar{\xi}_k \\ \hat{\delta}_{k-1} \end{bmatrix} + \ell\,\left(g_{m,k} - \bar{g} - \bar{\xi}_{2,k} - \bar{\xi}_{3,k}\right) \end{aligned} \tag{2.13}$$

Here Φ, γ, γ_δ and γ_m are the sampled model, u_{k-1} the rate over the period just ended, $\hat{d}_{k-1/2}$ the announced appearance at its middle and $g_{m,k}$ the reading the PID also receives. For a measurement variance $r = \sigma^2 = 2.25\,(\mathrm{mg/dL})^2$, a state noise variance $10^{-6}r$ and a disturbance noise variance $q_\delta = 100\,r$ the gain ℓ has entries 0.974 for ξ_2, 1.62 for δ and below 10^{-7} otherwise, so each innovation is booked almost wholly as glucose and as a lasting offset.

Each plan minimizes Equation (2.14) over $N_p = 36$ steps, 6.05 h for the 6 h horizon setting T_p, with eleven infusion rates μ_j held over blocks of $(1,1,1,1,2,2,3,4,6,9,6)$ steps, subject to $0 \le \mu_j \le u_{i,\max}$

and $70-\varepsilon_- \le \hat{g}_k \le 180+\varepsilon_+$ with $\varepsilon_\pm \ge 0$.

$$J = \frac{T_s}{\tau_n}\sum_{k=1}^{N_p}\left(\frac{\hat{g}_k - g^{\mathrm{SP}}}{y_s}\right)^2 + \lambda\sum_{j=1}^{11}\left(\frac{\mu_j-\mu_{j-1}}{u_s}\right)^2 + \rho_+\varepsilon_+^2 + \rho_-\varepsilon_-^2 \tag{2.14}$$

Here μ_0 is the rate applied now, $\tau_n = 1.55\,\mathrm{h}$, $y_s = 20\,\mathrm{mg/dL}$, $u_s = 10\,\mu\mathrm{U/min}$, $\lambda = 3$ a dimensionless move suppression, and ε_+ and ε_- one slack for each side of the band over the horizon, priced $\rho_+ = 100$ and $\rho_- = 1000$ per $(\mathrm{mg/dL})^2$. The infusion limits are hard, with $u_{i,\max}$ read as the PID reads it, and the soft band cannot leave the program without a solution. Against a tracking weight of 2.7×10^{-4} per $(\mathrm{mg/dL})^2$ and step, these prices make the band nearly a limit. The tenfold price of the low side is a teaching choice after Section 2.6 and the zone controllers with asymmetric costs proposed for the artificial pancreas [20], not a clinical calibration. The decisions are departures from $\bar{u}$, a dual active-set method solves the program, and λ and q_δ were fixed before any comparison by the tuning rule shared by the seven single-loop studios of the series.

With Meals announced on, $\hat{d}(t) = \sum_k m(t-t_k, A_k)$ from Equation (2.2) over the meals eaten and those the case has planned, so a planned meal is known from the start, and each plan reads $\hat{d}$ at the middle of every step. The announcement is exact, a best case that a person's estimate would not reach, and its glucose comes from the fitted meal path, never from the plant. Hybrid closed-loop systems ask the user for such an announcement, and in a pediatric study small manual premeal boluses added to closed-loop control gave average postprandial peaks of 194 against 226 mg/dL without them [21].

A plan without an optimum, or with a first rate that is not finite or lies outside the pump range, is set aside, and the PID, which follows the applied rate, takes the pump without a bump, moving it by 0.002 µU/min when a failure is forced at 1.5 h of case 9 with the noise off. The MPC reads the CGM, its own rate, the setpoint, the infusion limits and the announced meals, and no state of the plant, so Equations (2.1), (2.3) and (2.12) remain three distinct models.

2.5.5 Moving Horizon Estimation

The state that opens each plan can come instead from a moving horizon estimator (MHE), a least-squares fit over a sliding window of readings with hard bounds on the estimated states [22]. The filter of Equation (2.13) runs under either choice and still supplies the logged $\hat{\delta}$. Writing $\zeta = [\xi_1\ \xi_2\ \xi_3\ \delta]^{\mathsf{T}}$ for the filter's augmented state, execution k refits the readings $j = s,\dots,k$ with $s = k - N_e$ by solving Equation (2.15).

$$\begin{aligned}\min_{\zeta_s,\,\omega}\ & \|\zeta_s - \bar{\zeta}_s\|^2_{\bar{P}^{-1}} + \sum_{j=s}^{k-1}\omega_j^{\mathsf{T}}Q^{-1}\omega_j \\ & + \sum_{j=s}^{k}\frac{\left(g_{m,j} - \bar{g} - \xi_{2,j} - \xi_{3,j}\right)^2}{r} \\ \text{s.t. }\ & \zeta_{j+1} = \Phi_e\zeta_j + \Gamma_e\left(u_j - \bar{u}\right) + \Gamma_m\hat{d}_{j+1/2} + \omega_j \\ & K_u\left(u_{i,\max,0} - \bar{u}\right) \le \xi_{1,j} \le -K_u\,\bar{u}, \quad \xi_{3,j},\, a_j \ge 0\end{aligned} \tag{2.15}$$

The matrices $\Phi_e = \begin{bmatrix}\Phi & \gamma_\delta \\ 0 & 1\end{bmatrix}$, $\Gamma_e = [\gamma^{\mathsf{T}}\ 0]^{\mathsf{T}}$ and $\Gamma_m = [\gamma_m^{\mathsf{T}}\ 0]^{\mathsf{T}}$ restate the prediction of Equation (2.13), u_j is the rate held after reading j, the decisions are ζ_s and a noise ω_j on every state, and $Q = \mathrm{diag}(10^{-6}r,\, 10^{-6}r,\, 10^{-6}r,\, q_\delta)$ and r are the variances from which the filter was designed. The prior $\bar{\zeta}_s$ is the estimate the MHE itself made at execution $s-1$, propagated one period, and its weight $\bar{P}$ is the steady prior covariance of the filter, held constant. With $N_e = 10$ the window spans 1.68 h and eleven readings, which gives the dual active-set method of Section 2.5.4 44 decisions and 33 bound rows. The bounds come from declared values only. The insulin state stays between −194.1 and 77.5 mg/dL, the steady values of ξ_1 for the declared pump maximum $u_{i,\max,0} = 10\,\mu\mathrm{U/min}$ and for a

stopped pump, and the announced-meal glucose ξ_3 cannot fall below zero because a meal only adds glucose. The glucose deviation ξ_2 and the offset δ are left free, since glucose is the sum $\bar{g}+\xi_2+\xi_3$ and δ may stand for a meal, exercise or drift of either sign. Without an active bound the window problem has the filter's model, weights and prior, so it returns the filter's estimate. If a window's solution is suboptimal, contains a non-finite value or leaves the bounds, the estimate becomes a Kalman correction of the MHE's own one-period prediction, projected onto the bounds, and the history records the substitution. A first window, also one opened by a change of estimator during a run, starts from the filter's own prediction.

The third option, MHE with a meal state, gives ζ an unknown input of known shape, the appearance a of a meal nobody announced [23], which reaches glucose as $\tau_m\dot{\xi}_3 = -\xi_3 + K_m(\hat{d}+a)$ and fades as $\tau_a\dot{a} = -a+\omega_a$ with the 1 h decay of Equation (2.2). The noise ω_a has a standard deviation of 201.6 per period, what the default meal adds in one period of its ramp. A meal only adds glucose, so $a_j \geq 0$, a bound no Kalman filter can hold because its estimate is linear in the readings. With $q_\delta = 3.11\,(\mathrm{mg/dL})^2$, which lets δ drift by the 21.1 mg/dL of the largest basal drift in a day, and the rest unchanged, the window has 55 decisions and 44 bound rows, and its arrival cost is the augmented filter's steady prior, so without bounds it returns that filter's estimate, to 1.2×10^{-9} of a standard deviation, and as ω_a vanishes at $q_\delta = 225$ the MHE above, to 3.3×10^{-6}. The plan predicts on the same model, so it expects $\hat{a}$ to fade. Estimates and plans used 0.62 % of a core in Chrome at 1×.

2.5.6 Relay Autotuning

PI constants can also be measured from a limit cycle by the relay feedback experiment of [24], which the reader starts from the PID in automatic and which no shipped case runs. In place of the PID, a relay holds the pump at $u_0+\delta_r$ until the error e_k of Equation (2.6) falls below $-\varepsilon_r$, then at $u_0-\delta_r$ until it rises above ε_r, and so on. The center u_0 is the mean rate applied over the last 14 executions (2.35 h), the amplitude δ_r is 5 % of the pump span, 0.5 µU/min, or 0.9 of the distance from u_0 to the nearer stop when that is smaller, and the band $\varepsilon_r = 3\sigma_0 = 4.5\,\mathrm{mg/dL}$ is three times the declared CGM noise σ_0. A cycle in which either level lasts fewer than four executions is discarded. Once the period and the swing $\pm\alpha$ of the reading repeat within 5 % over two consecutive accepted cycles, the first harmonic of the applied levels, corrected for the band, gives the ultimate gain K_{cu}, and the Tyreus–Luyben rule turns it and the period P_u into PI constants, as Equation (2.16) shows.

$$K_{cu} = \frac{4\delta_r \sin(\pi D)}{\pi\sqrt{\alpha^2-\varepsilon_r^2}}, \qquad K_c = \frac{K_{cu}}{3.2}, \qquad \tau_I = 2.2\,P_u \tag{2.16}$$

Here D is the share of the cycle during which the relay sits at its upper level, one half only when u_0 is the rate that holds the setpoint. The constants are entered to three significant figures with $\tau_D = 0$ and remain editable. An unconverged test stops at 30 cycles or 48 h, or earlier at the reader's request, and in every outcome the PID, which has tracked the applied rate, takes the pump back without a bump. The band delays each switch, so the cycle settles where the loop's phase is $\psi = \arcsin(\varepsilon_r/\alpha)$ short of -180°, and a result with $\psi \geq 10^\circ$ is reported as placed by the band rather than by the insulin path.

2.6 Asymmetric Performance Metrics

A symmetric error integral scores an excursion of a given size the same way whichever side of setpoint it falls on. That score is not appropriate for this loop, and the reason is clinical rather than mathematical. Glucose measurements are usually summarized by the fraction of time spent inside a target band, and the international consensus recommendation reports that fraction together with separate thresholds for time below 70 mg/dL and time below 54 mg/dL, because the two low thresholds carry different consequences and neither is the mirror of the corresponding high threshold [25]. Time in range has been validated against microvascular outcomes and is now used as a trial endpoint [26]. The asymmetry was recognized in the measurement scale itself long before it was recognized in the outcome measure, and a transformation that symmetrizes the glucose axis about its clinical center was proposed for exactly that purpose [27].

> ✦ DID YOU KNOW? **A Small Allowance Below the Range**
>
> An international consensus of 2019 set targets for most adults with type 1 or type 2 diabetes: more than 70 % of sensor readings between 70 and 180 mg/dL, under 25 % above 180 mg/dL, under 4 % below 70 mg/dL, and under 1 %, less than 15 min a day, below 54 mg/dL [25].

The application shades a band of $\pm 20\,\mathrm{mg/dL}$ about the setpoint on the glucose trend, 60 mg/dL to 100 mg/dL at the default setpoint, and flags glucose below 64 mg/dL as low and above 180 mg/dL as high. The analysis here uses a fixed teaching band from 64 mg/dL to 104 mg/dL, whose lower edge is the application's low flag. That band is narrower and lower than the 70 mg/dL to 180 mg/dL consensus target range, so it is a tighter test than clinical practice and not a clinical target. Both bands are reported in Section 2.7 to keep the two readings apart.

The results tables therefore carry four numbers rather than one. The integrated absolute error against setpoint is the symmetric score. The excursion areas $\int \max(0, g-104)\,\mathrm{d}t$ and $\int \max(0, 64-g)\,\mathrm{d}t$ are the two one-sided scores, and they are what an asymmetric objective would weight unequally. Any weighted combination $J = \int \max(0, g-104)\,\mathrm{d}t + w \int \max(0, 64-g)\,\mathrm{d}t$, the objective of Equation (1.10) with the band edges of this chapter, ranks the two sides differently when $w > 1$, and the results show that the ranking of two runs reverses at $w = 7.9$. A specification written as a band rather than a setpoint recurs in Volume 9, Chapter 9, where a comfort band leaves the zone setpoint free.

2.7 Results

2.7.1 Protocol

Every run in this section starts from the application's reference operating point, case 1 of Table 2.3, which loads in manual mode with the pump at 3.0 µU/min, the setpoint at 76.22 mg/dL, the scheduled meals off and the sensor noise at zero. The settings listed in Table 2.4 are chosen before the run starts, so they set the starting state, and every step and upset is made after the run has started. The runs were driven headless through the application's own simulation engine, one 7.2 sec step at a time, exactly as the browser drives it. Time fractions, excursion areas, integrated absolute error and pump travel are computed from the logged noise-free glucose and pump columns of the exported file, which carries one row per control period of 0.168 h. Glucose extremes are taken at the integration step, because the logged rows miss the sharp meal peaks by up to 5.5 mg/dL. The application's own scorecard computes the integrated absolute error on the measured glucose, which without sensor noise equals the true glucose, and for the noise-free closed-loop day it reports the same 95.2 mg/dL h, pump travel and time at a limit as computed here. The noisy run uses the default noise with the case's fixed seed and repeats exactly. The configurations are listed in Table 2.4 and the outcomes in Table 2.5.

The app ships the eleven scenarios of Table 2.3, each credited by a completion rule of the kind described in Section 1.6.4, and the runs below use several of them by number.

2.7.2 Open-Loop Reference Day

With the pump held at the basal 3.0 µU/min through three meals of amplitude 1200, glucose reaches 155.5 mg/dL after breakfast, 157.0 mg/dL after lunch and 155.7 mg/dL after dinner, each about one hour after the meal begins. It is still at 80.3 mg/dL when lunch begins four hours after breakfast, and it is back within 0.6 mg/dL of the basal value before dinner and at the end of the day. The day spends 78.9 % of its samples inside the 64 mg/dL to 104 mg/dL band and none of them below it, and the area above 104 mg/dL is 122.5 mg/dL h. The repeatability of the three peaks to within 1.5 mg/dL is a property of the model rather than of a patient, and it is what makes the open-loop day a usable reference case.

Table 2.3 Scenarios of the App and What Each Requires

#	Scenario	Kind	Required to complete
1	Reference operating point	Experiment	Run 6 h in Manual at the basal steady state (pump 3.0 µU/min, glucose 76.22 mg/dL, meals off, no sensor noise) and note which states stay stationary.
2	Open-loop actuator step	Experiment	Starting in Manual at 80 mg/dL with the pump at 2.85 µU/min, run through the pump step to 4.85 µU/min at 2 h and on to 14 h.
3	Closed-loop setpoint step	Experiment	Starting in Auto at 80 mg/dL, run through the setpoint step to 101 mg/dL at 2 h until glucose has stayed within $\pm 5\,\mathrm{mg/dL}$ of the new setpoint for 1.2 h.
4	Disturbance rejection	Experiment	Starting in Auto at 80 mg/dL, run through a meal of amplitude 1800 at 2 h and on to 14 h.
5	Recover by hand	Exploration	Starting in Manual at 80 mg/dL with the pump at 2.85 µU/min, move the pump by hand after the same meal and bring glucose back within $\pm 20\,\mathrm{mg/dL}$ of 80 mg/dL without letting it pass 180 mg/dL, with the pump at a rate that holds glucose inside that band once settled, and keep it there for 1.2 h.
6	Conservative tuning	Experiment	In Auto with $K_c = 0.12$, run through the setpoint step from 80 to 96.8 mg/dL at 2 h and on to 14 h.
7	Aggressive tuning	Experiment	As case 6 with $K_c = 0.6$.
8	Actuator limit	Demonstration	Starting in Auto at 80 mg/dL, run through the setpoint step to 120 mg/dL at 2 h until the pump reaches a limit, and on to 8 h.
9	MPC with an announced meal	Experiment	In Auto at 80 mg/dL under the MPC with meals announced, run from the start through the planned meal of amplitude 1800 at 2 h to 8 h, then switch to PID, press Reset and run to 8 h again.
10	Unannounced meal, MHE + meal or Kalman	Experiment	As case 9 with the meal not announced and the MHE + meal estimator, then with the estimator switched to Kalman before Reset.
11	Challenge	Challenge	After the meal of amplitude 1800 at 2 h, keep measured glucose within $\pm 40\,\mathrm{mg/dL}$ of 80 mg/dL from 2.64 h to 4 h, the last third of a 4 h run, with the model parameters and disturbances as given.[a]

[a] With the shipped tuning ($K_c = 0.06$, $\tau_I = 3\,\mathrm{h}$) glucose peaks at 150.3 mg/dL and the worst measured error over the judged window is 70.6 mg/dL, 30.6 mg/dL outside the band. The suggested tuning changes only the gain, to $K_c = 0.3$, and the worst error is then 33.4 mg/dL, inside the band with 6.6 mg/dL to spare, in one run whether the tuning is chosen before Run or carried through Reset.

Table 2.4 Configurations of the Reported Runs

Label	Configuration
Basal	Manual, pump held at 3.0 µU/min, meal schedule on with $A = 1200$, 24 h
PI day	Automatic, $K_c = 0.3$, $\tau_I = 3\,\mathrm{h}$, $\tau_D = 0$, $g^{\mathrm{SP}} = 80\,\mathrm{mg/dL}$, meal schedule on, 24 h, starting at 80 mg/dL
Noisy	As PI day with the default CGM noise $\sigma = 1.5\,\mathrm{mg/dL}$
Large meal	As PI day with $A = 2400$
Exercise	As PI day, exercise raised to 60 % at $t = 20\,\mathrm{h}$ (in effect at 20.16 h), run to 26 h
Step test	Manual, meals off, pump moved 3.0 µU/min to 4.0 µU/min at $t = 2\,\mathrm{h}$ (in effect at 2.016 h), run to 14.1 h
Drift, $\tau_I = 3\,\mathrm{h}$	Automatic at $g^{\mathrm{SP}} = 76.22\,\mathrm{mg/dL}$, meals off, basal drift +30 % at $t = 2\,\mathrm{h}$ (in effect at 2.016 h), run to 14.1 h
Drift, P only	As above with τ_I at its 48 h top, which switches integral action off

2.7.3 Closed-Loop Day

The shipped tuning of $K_c = 0.3$, $\tau_I = 3\,\mathrm{h}$ and $\tau_D = 0$ holds glucose between 67.2 mg/dL and 103.1 mg/dL across the same day, which is 100.0 % of the samples inside the 64 mg/dL to 104 mg/dL band against 78.9 % open loop, and it reduces the integrated absolute error against the 80 mg/dL setpoint from 345.5 to 95.2 mg/dL h. The three peaks fall from the open-loop range 155.5 mg/dL to

Table 2.5 Measured Outcomes of the Reported Runs. Glucose Extremes Are Taken at the Integration Step and the Other Statistics From the Logged Noise-Free Glucose. Excursion Areas and IAE Are in mg/dL h, With IAE Against 80 mg/dL Except in the Drift Runs, Which Use Their 76.22 mg/dL Setpoint. Pump Travel Is $\sum |\Delta u_i|$ in µU/min, and the Last Column Is the Percentage of Samples at Either Pump Stop

	Glucose (mg/dL)		Time (% of run)				Excursion area		Pump		
Run	min	max	64–104	70–180	< 64	< 54	above 104	below 64	IAE	travel	at limit
Basal	76.2	157.0	78.9	100.0	0.0	0.0	122.5	0.0	345.5	0.0	0.0
PI day	67.2	103.1	100.0	94.4	0.0	0.0	0.0	0.0	95.2	59.7	1.4
Noisy	67.0	103.4	100.0	93.7	0.0	0.0	0.0	0.0	104.5	124.9	2.8
Large meal	64.4	130.9	87.3	90.8	0.0	0.0	56.2	0.0	210.3	58.6	14.1
Exercise	47.8	103.1	97.4	91.6	2.6	1.3	0.0	7.1	118.5	73.0	4.5
Step test	57.7	76.2	23.8	17.9	76.2	0.0	0.0	59.9	—	1.0	0.0
Drift, $\tau_I = 3$ h	71.7	96.3	100.0	100.0	0.0	0.0	0.0	0.0	11.8	15.7	0.0
Drift, P only	73.1	96.3	100.0	100.0	0.0	0.0	0.0	0.0	33.7	16.1	0.0

157.0 mg/dL to 101.0 mg/dL to 103.1 mg/dL, and they occur 0.42 h to 0.45 h after each meal begins rather than an hour after. On the clinical 70 mg/dL to 180 mg/dL band the ordering reverses, 94.4 % closed loop against 100.0 % open loop, because each postprandial descent carries glucose below 70 mg/dL. The pump reaches its 10 µU/min limit on 1.4 % of the samples, one after breakfast and one after dinner, and its lowest commanded value is 0.63 µU/min.

The upper excursion at this tuning is shaped jointly by the pump limit and the proportional gain. The breakfast peak occurs with the pump at its limit and the lunch and dinner peaks with the pump at 9.93 and 9.31 µU/min, so the controller is close to asking for what the device can deliver. Raising K_c to 0.45 still lowers the highest peak from 103.1 mg/dL to 99.8 mg/dL and the integrated absolute error to 74.7 mg/dL h, at the cost of 4.9 % of the samples with the pump at zero and 6.3 % at its maximum.

> ✦ **DID YOU KNOW?** **Sensor Noise Doubles the Pump's Work**
>
> From the simulator: a sensor noise with a standard deviation of only 1.5 mg/dL leaves the closed-loop day 100.0 % inside the 64 mg/dL to 104 mg/dL band, yet the pump's travel over the day rises from 59.7 to 124.9 µU/min (Table 2.5).

2.7.4 Large Meal and Saturation at the Upper Limit

Doubling the meal amplitude to 2400 with the same tuning raises the peaks to 128.5 mg/dL to 130.9 mg/dL and leaves 12.7 % of the day above 104 mg/dL, with an area above the band of 56.2 mg/dL h. The pump now sits on its upper limit for seven consecutive samples after breakfast and lunch and six after dinner, which is 1.0 h to 1.2 h each time, and it falls to a minimum of 0.61 µU/min without reaching zero. In this case the loop is authority-limited rather than tuning-limited. Doubling K_c lowers the highest peak by only 3.3 mg/dL, to 127.6 mg/dL, while doubling the pump maximum to 20 µU/min lowers it by 8.3 mg/dL, to 122.6 mg/dL, and with the shipped settings the manipulated variable spends 14.1 % of the day at its upper stop.

2.7.5 Exercise Disturbance

Raising the exercise disturbance to 60 % at $t = 20$ h, in effect from the controller execution at 20.16 h, multiplies s_i by 1.6 with no change in any measured quantity. Glucose falls from 74.9 mg/dL to 47.8 mg/dL within 0.04 h, before the controller next executes, because the glucose balance has a time constant of about one minute. The next execution drives the pump to zero, where it stays for 0.84 h, and glucose returns above 64 mg/dL 0.80 h after the change. The run spends 2.6 % of its samples below 64 mg/dL and 1.3 % below 54 mg/dL, and the area below the band is 7.1 mg/dL h.

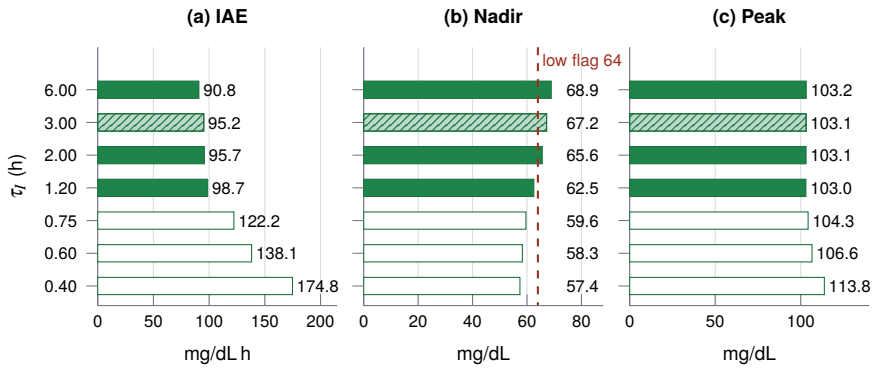


Figure 2.3 The integral gain sweep of Table 2.6: (a) integrated absolute error, (b) glucose nadir with the 64 mg/dL low flag (dashed), and (c) postprandial peak, for each integral time at $K_c = 0.3$. Hatched: the shipped 3.00 h. Solid: the other rows of the fivefold range from 6.00 to 1.20 h. White: the three faster rows. Notice that across the solid and hatched rows the IAE bars barely change and the peaks do not move, while the nadir walks down across the low flag.

Table 2.6 Integral Gain Sweep Over a Full Day at $K_c = 0.3$

K_I	τ_I (h)	nadir	peak	IAE	travel	% at zero
0.05	6.00	68.9	103.2	90.8	59.2	0.0
0.10	3.00	67.2	103.1	95.2	59.7	0.0
0.15	2.00	65.6	103.1	95.7	61.5	0.0
0.25	1.20	62.5	103.0	98.7	68.5	2.1
0.40	0.75	59.6	104.3	122.2	79.3	8.5
0.50	0.60	58.3	106.6	138.1	88.4	9.9
0.75	0.40	57.4	113.8	174.8	121.0	14.1

Pause. Try Exercise 2.2 before reading on.

The symmetric score misranks this run. Its integrated absolute error is 118.5 mg/dL h against 210.3 mg/dL h for the large-meal case, so the symmetric metric ranks it as the better of the two runs. The large-meal run never falls below 64 mg/dL, while the exercise run reaches 47.8 mg/dL and then holds the pump at zero for 0.84 h. Any asymmetric objective of the form given in Section 2.6 reverses the ranking once the low-side weight exceeds $w = 7.9$, which is the ratio of the two excursion areas 56.2 and 7.1.

2.7.6 Integral Gain Sweep

Repeating the closed-loop day at seven values of the integral time with $K_c = 0.3$ fixed separates the two costs, as listed in Table 2.6. Figure 2.3 draws the same rows as bars. Over a fivefold change in integral gain, from $\tau_I = 6$ h to 1.2 h, the integrated absolute error rises by only 8.7 %, from 90.8 to 98.7 mg/dL h, while the glucose nadir falls by 6.4 mg/dL, from 68.9 to 62.5 mg/dL, and the day begins to spend time below the 64 mg/dL flag. The postprandial peak moves by 0.15 mg/dL over the same range, which confirms that the high side is insensitive to the integral and the low side is not. Below $\tau_I = 1.2$ h both sides deteriorate, and at $\tau_I = 0.4$ h the peak reaches 113.8 mg/dL, the nadir 57.4 mg/dL, and the pump spends 14.1 % of the day at zero.

Within the sweep both objectives select the slowest integral, and switching integral action off at the 48 h top of the τ_I range lowers the integrated absolute error to 70.9 mg/dL h and raises the nadir to 72.1 mg/dL. On a day of meals with no unmeasured drift the integral action therefore carries no benefit, and its cost falls on the low side, where a faster integral keeps pushing insulin while glucose

is already falling and pulls the nadir below the band in the last four rows of Table 2.6. The integral is still needed, because Section 2.7.7 shows that only integral action returns glucose to setpoint after an unmeasured change in endogenous appearance. The shipped integral time of 3.0 h, 2.5 times the identified time constant, is a compromise between the two cases. It settles that drift within twelve hours while the day's nadir stays 3.2 mg/dL above the low flag.

Pause. Try Exercise 2.3 before reading on.

2.7.7 Basal Drift and Steady-State Offset

A 30 % step in the baseline appearance term with the meals switched off and the loop closed at the reference setpoint of 76.22 mg/dL raises the required infusion rate from 3.00 µU/min to 3.82 µU/min by Equation (2.5). Glucose rises to 96.3 mg/dL 0.17 h after the step, because the fast glucose balance reaches its new quasi-steady value before the insulin path can respond. With the shipped $\tau_I = 3$ h the loop returns to 76.25 mg/dL and 3.82 µU/min within the 12.1 h that follow, and to 76.22 mg/dL and 3.816 µU/min after 24 h, which matches Equation (2.5) to four significant figures. At the 48 h top of its range the integral-time setting switches integral action off, and the proportional-only loop settles at 78.55 mg/dL with the pump at 3.70 µU/min and stays there. The integrated absolute error over the 12.1 h after the step is 11.8 mg/dL h with $\tau_I = 3$ h and 33.7 mg/dL h without integral action.

The offset is small because the loop gain $K_c|K_p| = 5.6$ is well above unity. Equation (2.4) with the starting bias of 3.00 µU/min and $K_c = 0.3$ places the proportional-only equilibrium at 78.55 mg/dL and 3.70 µU/min, an offset of 2.33 mg/dL, which the run reproduces to four significant figures. The bias is set only when a run starts, so for a drift that arrives during the run the integral term is the only mechanism in the application that removes the offset.

2.7.8 Sensor Noise and the Pump

Turning the continuous glucose monitor noise back on at its default 1.5 mg/dL standard deviation changes the glucose statistics very little, because the noise is a small fraction of the postprandial excursion. The total variation of the pump rate roughly doubles, from 59.7 to 124.9 µU/min summed over the day, and the pump reaches zero on one sample, 0.7 % of the day, where the noise-free run never does. Eight other noise seeds give a pump travel of 109 to 124 µU/min and zero or one sample at zero, so the doubling is robust and the single sample at zero is within the spread between seeds. A proportional gain of 0.3 converts a 1.5 mg/dL measurement deviation into a 0.45 µU/min command deviation, which is 15 % of the basal rate. Adding derivative action with $\tau_D > 0$ magnifies the effect, because the second difference of the measurement is divided by $\Delta t = 0.168$ h.

2.7.9 Application

The app offers the manual and automatic modes, scorecard and data export of Section 1.6.4, and its scorecard reports the integrated absolute error, pump travel and time at a limit. The interface is shown in Figure 2.4, with a schematic beside the glucose reading and trends of glucose, infusion rate and insulin action.

2.8 Predictive Control and the Announced Meal

Table 2.7 runs cases 3, 4, 8 and 9 over 24 h and the challenge over 4 h, the controller chosen before each run, so the PID and the MPC face identical events and CGM noise within a case. Unlike Section 2.7, every value is read on the 0.168 h log, as the scorecard reads it. The PID keeps $K_c = 0.3$, $\tau_I = 3$ h and $\tau_D = 0$, the challenge shipping $K_c = 0.06$ and suggesting 0.3, while the MPC holds the values given in Section 2.5.4, with meals announced only in case 9. Cases 4 and 9 schedule the same meal of amplitude 1800 at 2 h, so case 4 is the like-for-like comparison and case 9 adds only the announcement.

At the 101 mg/dL of case 3 the tangent gain of Equation (2.4) is 1.6 times its value at 80 mg/dL, where every case starts and the MPC's model was fitted. The PID holds the pump at zero for 0.84 h,

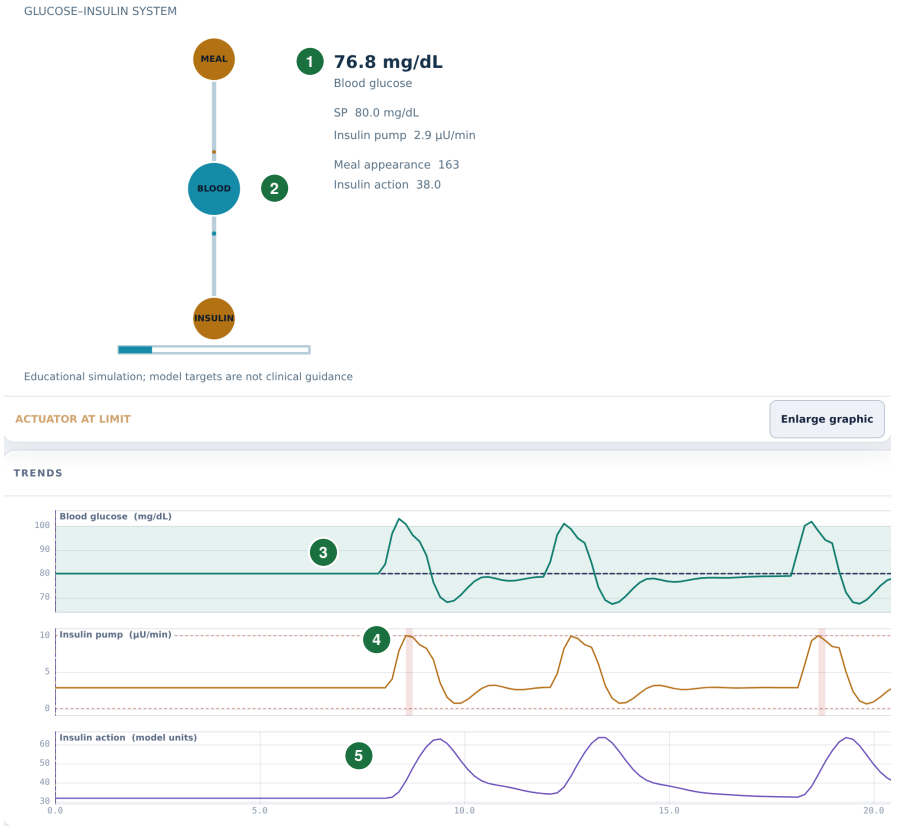


Figure 2.4 Closed-loop day under the shipped tuning, 21 h into the run with the 24 h trend window, from the reference case with the meal schedule on, the setpoint at 80 mg/dL and automatic control chosen before the run, so glucose holds 80 mg/dL until breakfast. The plant view and the trends are shown, cropped just before the end of the trends. (1) The glucose reading, with the setpoint and the pump rate. (2) The schematic of meal, blood and insulin. (3) The glucose trend with its shaded band. (4) The pump trend at its limit. (5) The insulin action. The three postprandial peaks rise just above the shaded ± 20 mg/dL band, the pump touches its 10 µU/min limit after breakfast and dinner, and the insulin action peaks about 0.9 h after the pump. At 21 h glucose is 76.8 mg/dL and the pump delivers 2.9 µU/min.

overshoots to 109.4 mg/dL and stays within 5 mg/dL only from 11.6 h after the step, while the MPC keeps the pump above 0.75 µU/min, peaks at 105.8 mg/dL and stays within 5 mg/dL from 0.69 h. At the 120 mg/dL of case 8 the gain is 2.25 times that value and the PID never settles, its pump at zero for 6.9 h and glucose still between 105.7 and 133.3 mg/dL after 8 h. The first three plans there bind the pump's floor, and the MPC spends 0.34 h at zero, peaks at 129.5 mg/dL, stays within 5 mg/dL from 1.36 h and over the last 12 h averages the 1.87 µU/min that Equation (2.5) requires, with $\hat{\delta}$ averaging 13.3 mg/dL.

The MPC loses case 4. The meal shows clearly in the reading of 2.184 h, from which the proportional term of the PID holds the pump at its maximum. The logged $\hat{\delta}$ climbs from 38 to 101 mg/dL by 3.19 h, but the plans, which expect it to persist and pay for each move, raise the pump through 5.4, 7.4, 9.1 and 9.6 µU/min to the maximum only from 2.856 h, and glucose peaks at 130.9 mg/dL against 116.8. On the descent $\hat{\delta}$ is still 85 mg/dL at 3.36 h, the MPC holds 6.7 and 5.7 µU/min where the PID is down to 4.4 and 2.3, and glucose falls to 64.9 mg/dL against 66.2, until two plans stop the pump on the low side of the band.

Table 2.7 Five Cases Under the Shipped PID and the Predictive Controller, With the Same CGM Noise in Every Run of a Case

Case	Controller	IAE (mg/dL h)	IAE_g	g_{max} (mg/dL)	g_{min}	$t_{<70}$ (h)	Travel	Zero (%)	Max (%)
3, setpoint	PID	67.65	58.31	109.4	79.1	0	125.1	3.5	0
	MPC	46.93	29.26	105.8	79.8	0	30.8	0	0
4, meal	PID	83.67	68.97	116.8	66.2	0.67	101.9	0.7	4.2
	MPC	98.98	82.54	130.9	64.9	0.84	51.1	1.4	1.4
8, limit	PID	164.17	160.28	133.3	79.1	0	219.2	28.9	0
	MPC	64.42	48.84	129.5	79.8	0	34.1	1.4	0
9, announced[a]	MPC	50.94	34.42	98.6	75.0	0	41.9	0	2.1
11, challenge	PID, $K_c = 0.06$	70.89	68.00	148.6	77.5	0	10.0	0	0
	PID, $K_c = 0.3$[b]	45.79	42.70	116.8	66.2	0.50	28.3	4.3	26.1
	MPC	56.50	53.12	130.9	64.9	0.34	20.2	4.3	8.7

IAE uses the CGM reading, as the scorecard does, and IAE_g the true glucose. g_{max} and g_{min} are the highest and lowest logged true glucose from the event at 2 h, $t_{<70}$ the logged time below 70 mg/dL, Travel the summed absolute change of the infusion rate in µU/min, and Zero and Max the time with the pump at zero and at its 10 µU/min maximum as shares of the run. No run exceeds 180 mg/dL. [a]PID as in case 4. [b]Suggested, $\tau_I = 3$ h.

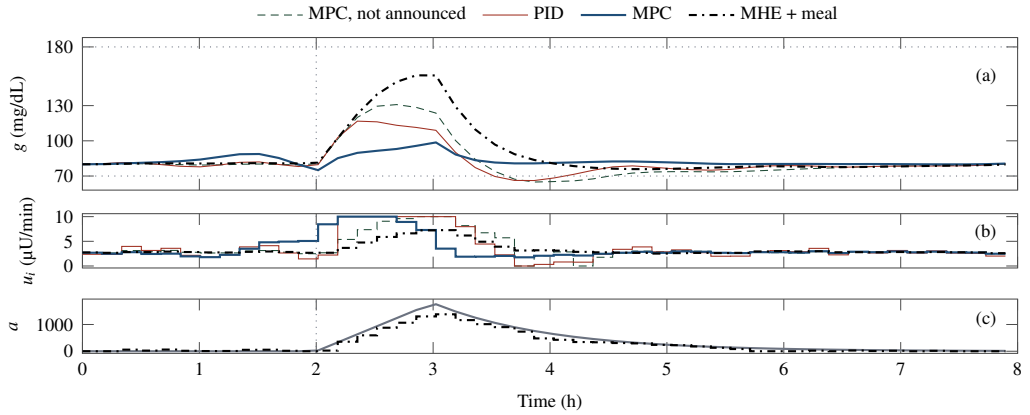


Figure 2.5 Case 9 run three ways from the same CGM noise, 8 h of the 0.168 h log, with the meal beginning at 2 h (dotted), and case 10, the unannounced run with the meal-state MHE in place of the Kalman filter. (a) True glucose and the 70 and 180 mg/dL sides of the band. (b) Infusion rate, each value drawn over the period it was held. (c) True meal appearance (gray) and $\hat{a}$, each drawn over the period after it was made.

> **TRY IT IN THE APP** **4 · Disturbance Rejection**
>
> Load scenario 4 and press Run. A meal of amplitude 1800 begins at 2 h. Watch the glucose trend: under the shipped PID it peaks at 116.8 mg/dL, and on the way down it falls to 66.2 mg/dL, below 70 mg/dL for 0.67 h (Table 2.7). Then watch the pump trend reach its maximum after the meal.

Told from the start of case 9 that the meal begins at 2 h, the twelve plans made before it each predict a lowest glucose of exactly 70 mg/dL, so the band decides how much insulin can go in early (Figure 2.5). They first lower the pump to 1.79 µU/min, letting glucose rise to 88.8 mg/dL, then raise it to 5.07 µU/min at 1.848 h, so glucose is 75.0 mg/dL in the first sample of the meal, which meets insulin already acting. Glucose peaks at 98.6 mg/dL, or 99.8 at the integration step, and stays above 78.7 mg/dL afterward. Over the first 8 h the MPC delivers 28.1 against 28.8 µU/min h of insulin, nearly the same amount given earlier, which costs 6.7 against 2.3 mg/dL h of true IAE before the meal and returns 18.7 against 52.2 over the next 6 h.

Table 2.8 The Meal of Cases 4 and 9 Under Each Estimator Over 24 h

	Case 4, not announced						Case 9	
Controller	g_{max}	g_{min}	$t_{<70}$	IAE	IAE$_g$	Area	IAE	IAE$_g$
PID	117.4	65.8	0.61	83.7	69.4	—	83.7	69.4
MPC, Kalman or MHE	131.1	64.9	0.89	99.0	82.6	—	50.9	34.1
MPC, MHE + meal	157.9	75.8	0	114.2	103.9	2730	78.6	72.4
$q_\delta = 12.4$	155.4	71.7	0	118.4	107.1	2292	67.4	57.7
$q_\delta = 44.5$	149.7	67.6	1.41	121.7	108.0	1668	60.8	48.6
$q_\delta = 225$	140.5	65.5	1.05	110.6	95.1	812	54.3	39.7
no sign bounds	153.4	72.7	0	138.5	131.6	298	49.1	31.9
MPC, Kalman[a]	142.1	61.6	2.20	123.3	106.8	—	48.4	31.9
PID, $\hat{a}$ fed forward	117.5	60.6	0.94	88.2	74.3	2951	88.8	75.0

One noise seed. Extremes (mg/dL) and $t_{<70}$ (h) at the integration step from 2 h, IAE as in Table 2.7, Area of $\hat{a}$ in units h (true 2700). [a] $q_\delta = 3.11$. No run leaves 54 mg/dL to 180 mg/dL, and in case 9 only the PID runs fall below 70.

> **TRY IT IN THE APP** **9 · MPC With an Announced Meal**
>
> Load scenario 9 and press Run. The predictive controller knows from the start that a meal begins at 2 h. Watch the pump before the meal: it first drops to 1.79 µU/min, then rises to 5.07 µU/min at 1.848 h, so insulin is already acting when the meal arrives. Glucose then peaks at 98.6 mg/dL.

With only the band removed, in a run outside the application, the announced plans give up to 7.3 µU/min before the meal without lowering the pump first, glucose is 67.6 mg/dL at the meal, below 70 for 0.17 h, and the IAE falls to 43.8 mg/dL h, while case 4 falls to 58.7 mg/dL with 2.18 h below 70. Pricing both sides alike, at 100 or 1000, reproduces both cases to every digit shown, since no plan finds the 180 mg/dL side active and every active low side has a slack of at most 1.1×10^{-4} mg/dL or the pump at zero.

Unaffected by K_c, the MPC repeats in the challenge its first 4 h of case 4, and its largest reading error from 2.64 to 4 h, 50.5 mg/dL, fails the $\pm 40\,\mathrm{mg/dL}$ test as the shipped gain does with 70.6, while the suggested $K_c = 0.3$ passes with 33.4. All plans were optimal, 143 per 24 h case, and their mean solve time, a wall-clock figure from one desktop computer, was 0.05 to 0.2 ms under Node.js and 216 and 206 µs in two Chrome runs of case 9, 0.13 % of one core at the default speed. The MPC wins cases 3, 8 and 9 and loses the two with a meal it is not told of.

2.9 Estimator Choice and Relay-Written Constants

The MHE without a meal state repeats the filter's run in every case, within 4.8×10^{-8} µU/min over 72 h, while the meal state changes it. Case 10 is case 4 with that estimator, and Table 2.8 compares the estimators on its meal, unannounced as in case 4 and announced as in case 9.

Unannounced, the meal-state MPC never falls below 70 mg/dL and pays at the peak, 157.9 against 131.1 mg/dL under the filter. Its $\hat{a}$ peaks at 1381 against a true 1800, a root-mean-square error of 77, and its plans, expecting the meal to pass, stop at 7.3 µU/min, while the filter books the rise as a lasting $\hat{\delta}$ and its plans keep insulin flowing as the meal fades (Figure 2.5). Over seeds 1001 to 1010 it stays above 74.8 mg/dL, against 0.90 to 1.38 h below 70 for the filter. Neither half suffices alone, since the meal state beside the quick $q_\delta = 225$ and the filter given the slow one both fall below 70, as does the filter's MPC at every λ and horizon tried, at best 0.67 h at $\lambda = 30$. The sign bounds buy the estimate, not the absence of the low, as unbounded $\hat{a}$ is negative at 111 of 142 estimates. Announced, it loses on every seed, a true IAE of 56.7 to 71.9 against 26.8 to 31.7, because $\hat{a}$ cannot fall below zero where the meal path overpredicts, and freeing a alone gives 55.3. The slow q_δ came after 225 was measured, and its 24 h scale is the trend window, so the table adds the rule over the 6 h horizon, which keeps the result, and the 1.68 h window, which loses it.

Over 74 runs at extreme settings of the controls, 24 after the meal of case 4 or 9, the meal-state MPC is worse on the low side than the filter's in 13, longer below 54 in one (0.43 against 0.39 h) and

Table 2.9 Shipped PID, Relay-Written PI and Two Counterfactual PI Loops on Cases 3 and 4 Over 72 h With One Noise Seed

Case	Controller	IAE	IAE_g	IAE_g^{24+}	$g_{\max}$	$g_{\min}$	$t_{<70}$	Travel
		(mg/dL h)			(mg/dL)		(h)	
3, setpoint	PID, shipped	192.7	172.8	114.5	110.3	78.8	0	379.9
	PI, relay	129.4	72.2	21.0	102.5	79.9	0	46.7
	PI[a]	186.3	165.3	109.1	109.8	78.9	0	374.0
	PI[b]	116.6	54.5	21.7	102.6	79.9	0	46.2
4, meal	PID, shipped	164.7	122.7	53.3	117.4	65.8	0.61	279.3
	PI, relay	205.0	150.1	16.2	152.5	70.2	0	52.0
	PI[a]	161.2	119.5	52.2	117.5	67.6	0.45	281.7
	PI[b]	210.4	151.4	15.6	151.5	66.2	2.28	51.8

PID, shipped: $K_c = 0.3$, $\tau_I = 3$ h. PI, relay: $K_c = 0.0578$, $\tau_I = 5.14$ h. [a]$K_c = 0.3$, $\tau_I = 5.14$ h. [b]$K_c = 0.0578$, $\tau_I = 3$ h. Neither counterfactual is shipped or written. IAE_g^{24+} is the true IAE after the first 24 h. Glucose extremes and $t_{<70}$ are taken at the integration step from the event at 2 h. Other columns as in Table 2.7.

above 180 in seven (up to 2.30 against 0.84 h), since a slow $\hat{\delta}$ and a nonnegative $\hat{a}$ correct a fall that is no meal late. In all twelve added runs after the unannounced meal it is safer, exercise at 40 % at 4 h leaving 56.1 against 47.0 mg/dL.

With the setpoint step of case 3 deleted, the relay test began at 2.4 h, glucose steady at 80 mg/dL and u_0 at 2.87 µU/min, against the 2.853 µU/min that holds that glucose, so δ_r took its full 0.5 µU/min. It converged after three cycles and 5.83 h with $\alpha = 5.64\,\text{mg/dL}$, $D = 0.457$, $P_u = 2.34$ h and $K_{cu} = 0.185$ in µU/min per mg/dL, and it wrote $K_c = 0.0578$ and $\tau_I = 5.14$ h, a fifth of the shipped gain and 1.7 times its integral time, flagged as band-limited at $\psi = 52.9°$. The ten seeds 1001 to 1010 converged as well, with K_{cu} from 0.114 to 0.176, P_u from 2.22 to 3.03 h and ψ from 38.8° to 52.1°.

That result reflects the switching band and the CGM noise rather than the insulin path, as the relay test of Volume 2, Chapter 2 does on a vehicle. The two lags of Equation (2.12) approach a phase of −180° without crossing it, and only the zero-order hold at Δt gives the sampled model an ultimate point, at a period of 0.77 h and a gain of 0.821. The measured cycle is three times slower, and its K_{cu} is 23 % of that gain. The same relay returns 0.279 at 2.69 h on the noise-free model and 0.278 at 2.69 h on the simulator with the CGM noise off, where $D = 0.527$ and $\psi = 63.1°$. Noise lowers K_{cu} because α lies close to ε_r. A swing 12 % wider, 5.64 against 5.05 mg/dL, makes $\sqrt{\alpha^2 - \varepsilon_r^2}$ in Equation (2.16) half as large again.

The written constants exchange the meal peak for the low side, as listed in Table 2.9. On the setpoint step the scorecard IAE falls from 193 to 129 mg/dL h and the true IAE from 173 to 72.2, mostly after the first day, where the shipped gain passes the CGM noise to the pump (travel 380 against 46.7) and the true glucose departs from 101 mg/dL by 2.4 mg/dL on average over the last 48 h against 0.44. On the meal the highest glucose rises from 117.4 to 152.5 mg/dL and the scorecard IAE from 165 to 205, while the lowest glucose rises from 65.8 to 70.2 mg/dL and the 0.61 h below 70 mg/dL vanishes. The counterfactual rows separate the two effects. The shipped K_c with the written τ_I keeps the peak at 117.5 mg/dL with 0.45 h below 70, and the written K_c with the shipped τ_I repeats the high peak, 151.5 mg/dL, with 2.28 h below 70. The peak follows the proportional gain, and the descent stays above 70 mg/dL only when the small gain is paired with the slow integral, an integral gain K_c/τ_I of 0.011 against the shipped 0.1, which is the low-side effect of Table 2.6 at a different gain. Over 21 relay tests (the shipped seed and 1001 to 1020) the written constants won the setpoint step every time, with scorecard IAE from 119 to 192, and lost the meal every time, with 186 to 235 mg/dL h against 165.

2.10 Technology in Practice

The symbols of Figure 2.1 stand, on a person, for a glucose sensor worn on the skin, a controller that runs in software, and an insulin pump that infuses under the skin. Figure 2.6 draws the loop with that hardware in place of the schematic's symbols.

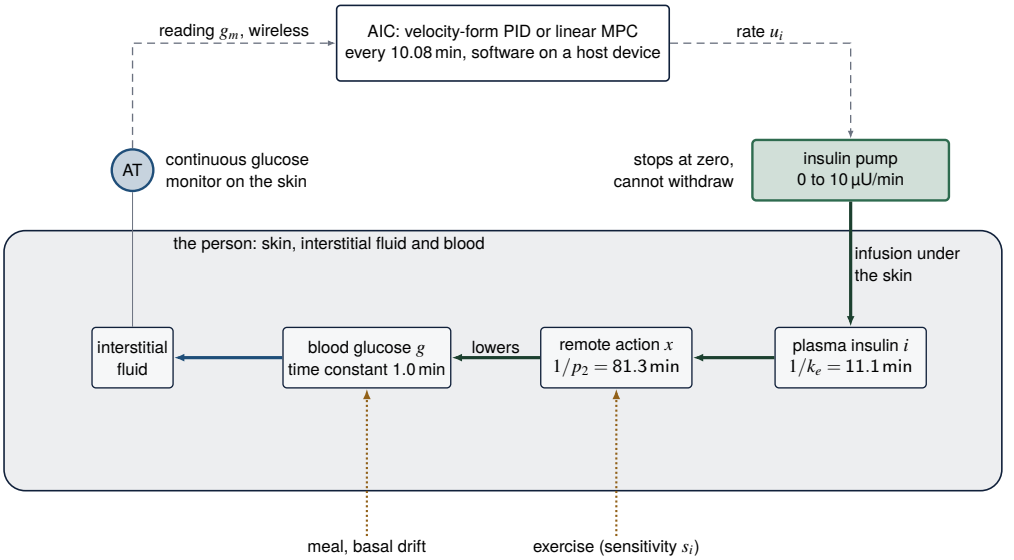


Figure 2.6 The loop of Figure 2.1 with the hardware of a wearable system. A continuous glucose monitor (AT) reads the interstitial fluid and sends its reading to the controller AIC, which every 10.08 min sets the rate of the insulin pump. Insulin reaches glucose through the plasma and remote action pools, whose time constants set the lag the loop must work through. Thick solid arrows carry insulin and glucose, thin dashed arrows signals, dotted arrows disturbances. In this model the pump feeds the plasma pool directly, and the sensor has no lag.

What Measures It

The controller reads one signal, the glucose that a continuous glucose monitor reports. A real monitor reads the interstitial fluid, and its data need filtering, denoising and artifact rejection before they can drive a controller [6]. A monitor meant to feed an automated system is held to accuracy limits written into regulation. Across its measuring range, the lower 95 % confidence bound on the share of readings within 20 % of blood glucose must exceed 87 %, and no reading above 180 mg/dL may correspond to a blood glucose below 70 mg/dL [28]. The chapter's monitor of Equation (2.3) keeps only the noise, with no lag and no calibration error.

> **TECHNOLOGY Sensor: The Continuous Glucose Monitor**
>
> A continuous glucose monitor places a small electrochemical sensor under the skin, where the enzyme glucose oxidase turns the glucose of the interstitial fluid into an electrical signal. It gives a reading every 1 to 5 min, and recent sensors reach a mean absolute relative difference below 10 % against a laboratory reference. The reading lags blood glucose by 5 to 10 min, the time glucose takes to pass from the blood into the interstitial fluid. Earlier sensors also suffered from imperfect calibration, from compression artifacts and from interference by drugs such as acetaminophen [29]. The chapter's monitor adds noise of 1.5 mg/dL standard deviation and has no lag.

What Moves It

The final element is a pump worn on the body that delivers insulin at set or variable basal rates and in boluses, and whose labeling states its minimum and maximum basal rates and its delivery increments [30]. Its fail position is the one the chapter studies: the pump can stop, and it cannot withdraw what it has given. The delay that a real pump adds sits in the subcutaneous depot, where insulin waits to be absorbed [6]. The chapter's pump has no depot, no delivery increment and no lag, as Section 2.11 states, so its insulin reaches the plasma pool of Equation (2.1c) at once.

TECHNOLOGY Final Element: The Insulin Pump

Continuous infusion of insulin under the skin has been an accepted therapy since the late 1990s. Pumps allow fine adjustment of the basal rate, and patch pumps are worn directly on the skin. The weak link is absorption. Regular insulin given under the skin peaks up to 120 min after a bolus, and altered absorption at the infusion site can underdeliver insulin, which may lead to ketoacidosis if it is not detected and corrected in time [6]. A pump that takes commands from an outside controller must detect occlusions and air in the line within a clinically relevant time and state its delivery accuracy over the first 1, 6 and 12 h [30]. The chapter's pump applies its rate at once between zero and 10 µU/min.

Where the Controller Runs

In current hybrid closed-loop systems the algorithm runs on one device of the system, and in one of them it is an app on a smartphone that talks to the sensor and the pump. The algorithms in use include model predictive control, PID control and fuzzy logic [31], and the regulation of an interoperable controller requires specifications for the device that hosts it [32]. A controller cannot act more often than its readings arrive, every 1 to 5 min from current sensors [29]. The chapter's controller acts every 10.08 min. Its structure is one block of software: the velocity-form PID with back-calculation of Equation (2.10), or the linear predictive controller with its Kalman filter, which hands the pump back to the PID when a plan fails.

IN PRACTICE How an Automated Insulin System Is Cleared

In the United States an automated insulin system can be assembled from three separately cleared devices: an integrated continuous glucose monitor (21 CFR 862.1355), an interoperable automated glycemic controller (21 CFR 862.1356) and an alternate controller enabled infusion pump (21 CFR 880.5730). The controller may take glucose only from such a monitor and command only such a pump. Each device must keep therapy safe when communication is interrupted or lost, for example by reverting to a preprogrammed safe delivery rate, and must log the commands and their confirmations, malfunctions such as a pump occlusion, alarms, and every loss of connection [30, 32].

What Is Modernizing It

Three developments are changing this loop. Interoperable devices let users combine a sensor, a pump and a controller into a closed-loop system of their own, and the three device classes of the in practice box above make that possible in regulation [31, 32]. Sensors now reach an accuracy close to that of fingerstick meters, and some can be used for ten days without calibration by the user [29]. Do-it-yourself closed-loop systems, built by user communities on open software without regulatory review, are used by several thousand people, and observational studies report gains in time in range, although no randomized trial had yet tested them at the time of the review [31].

2.11 Limitations

The application is an educational simulator. It is not a medical device, it is not validated against patient data, and no tuning obtained from it should be transferred to any clinical setting. The parameter set describes one virtual subject, and every result in Section 2.7 is a property of that parameter set.

The model omits the two largest sources of delay in a real system. There is no subcutaneous insulin depot between the pump and the plasma compartment, and there is no interstitial transport lag between plasma glucose and the reported measurement. Both delays are on the order of tens of minutes in a wearable device, and adding either would make the loop substantially harder to stabilize than the one measured here. The identified apparent dead time of 0.16 h is a consequence of the insulin cascade and not of any transport delay in the simulator.

The transit chain of Equation (2.1d) to Equation (2.1f) is driven by the infusion rate rather than by the meal, which follows the published assignment code. A model in which the meal drives the chain would place the meal disturbance behind two additional lags and would change both the shape and

the timing of the postprandial peak. The application's own model notes state the same routing, with the meal entering the blood directly.

The teaching band of 64 mg/dL to 104 mg/dL and the band of $\pm 20\,\mathrm{mg/dL}$ about the setpoint that the application shades are not the consensus clinical target range of 70 mg/dL to 180 mg/dL, and the low-glucose flag at 64 mg/dL is not either of the two clinical thresholds at 70 mg/dL and 54 mg/dL. Table 2.5 reports both bands so that the two readings cannot be confused.

The sensor noise comes from a generator reseeded with a fixed seed whenever a case starts, so a noisy run repeats exactly but represents one noise realization, and eight other seeds move the pump travel of the noisy day between 109 and 124 µU/min. The PID has no feedforward path and no meal announcement, and neither controller has an insulin-on-board constraint, all three of which are standard in deployed systems, so the closed-loop performance reported in Section 2.7 is that of feedback alone.

The predictive controller plans on a linear model frozen at 80 mg/dL, without the plant's insulin gain, 1.6 and 2.25 times larger at 101 and 120 mg/dL, or the smaller rise of a meal against insulin already acting, so its first plan in case 9 predicted a peak of 130.4 mg/dL where the plant, fed that plan open loop, peaks at 111.4 mg/dL. Its announcement is exact, its tenfold low-side price is never exercised by the compared cases, it has no insulin-on-board constraint, and it loses both cases with an unannounced meal.

The MHE constrains the states of the controller's linear model, never the plant's, so where the model is wrong the bounded estimate is wrong in the same way, and on this loop the filter crosses the one bound ever active by at most 6.2×10^{-7} mg/dL. The meal state knows one meal shape, a one-hour fade, so a longer meal is dosed as one about to pass, and its gain rests on a slow δ chosen after a first rule failed. The relay test reads a cycle placed by its 4.5 mg/dL band and by the CGM noise, at ψ of 38.8° to 52.9°, and converts it by one conservative rule, so its constants are a starting point that the meal case shows must still be judged against the low side.

> ⚠ **OPEN CHALLENGE** **Meals and Exercise Without Being Told**
>
> Rises in glucose after meals remain a challenge for closed-loop systems, because insulin given under the skin is absorbed slowly. Good control still needs the user to count carbohydrate and give a bolus before the meal, and attempts to lighten that burden with simplified boluses or fully closed loops have compromised glucose control. Exercise raises the risk of low glucose and changes insulin sensitivity, so users usually announce it to the algorithm in advance [31]. Cases 4 and 9 of this chapter measure the same gap on one virtual patient. Faster-acting insulins, and extra signals such as heart rate or body motion that reveal activity sooner than glucose alone, are the candidates for progress [31].

2.12 Discussion

Faster insulin, a shorter sensor lag and a larger pump all attack the delay between a decision and its effect, and none of them changes the sign structure of the problem. Insulin will still only lower glucose, the pump will still stop at zero, and a low will still cost more than a high of the same size. That is why deployed systems treat the insulin already delivered as a state that limits what the controller may ask for next [11], and why the outcome measure the field settled on reports the two sides of a band separately [25]. The lesson generalizes past this plant. Any loop with a one-sided actuator and an unequal cost of the two error signs, such as emergency cooling, pressure relief or an override acting only in the safe direction, shows the pattern of Table 2.6, a symmetric index nearly flat across tunings over which the constrained side moves a great deal, so tuning by that index selects a controller materially worse on the constraint.

2.13 Conclusion

This chapter studied Type 1 diabetic glucose regulation on the six-state expanded Bergman minimal model, three meals on a twenty-four hour clock, an exercise disturbance that scales insulin sensitivity

and a basal drift that shifts endogenous appearance, all stated in full for an independent rebuild. A fasting step test from 3.0 µU/min to 4.0 µU/min returns a process gain of −18.55 mg/dL per µU/min and a first-order plus dead-time fit of $\tau_p = 1.18$ h and $\theta_p = 0.16$ h, in which the apparent dead time is the plasma insulin clearance time constant rather than a transport delay. The shipped direct-acting velocity-form controller raises time inside the 64 mg/dL to 104 mg/dL band from 78.9 % to 100.0 % over a full day while reaching the pump limit on 1.4 % of the samples. A fivefold change in integral gain moves the integrated absolute error by 8.7 % and the glucose nadir by 6.4 mg/dL, which is the measured form of the claim that a symmetric metric cannot express the objective of this loop. Eleven shipped cases, from a reference operating point to a postprandial challenge, carry explicit completion requirements. Told of a meal, the predictive controller beside the PID scores an IAE of 50.9 against the PID's 83.7 mg/dL h, and 99.0 when not told. In place of its Kalman filter a moving horizon estimator moves no case by more than 4.8×10^{-8} µU/min, and one with a nonnegative meal state removes the 0.89 h below 70 mg/dL after the unannounced meal while the peak rises from 131.1 to 157.9 mg/dL. The relay-written constants, a fifth of the shipped gain, keep the meal case above 70 mg/dL while its peak climbs from 117.4 to 152.5 mg/dL.

Several things the simulator does not carry would change the numbers, including a subcutaneous insulin depot, an interstitial sensor lag, feedforward of an announced meal to the PID, an insulin-on-board constraint and more than one virtual subject. The app still shows direct action against a negative process gain, a one-sided actuator, a disturbance faster than the manipulated path, and a performance cost that is not symmetric, in a single run of a simulated day that takes 24 sec at the default speed.

KEY POINTS

- Insulin only lowers blood sugar and a pump cannot take it back, so the controller pushes one way only.
- The apparent dead time of the identified model is the plasma insulin lag, reported as a delay by the first-order fit.
- Once the pump is at zero, no retuning shortens a low.
- A fivefold change in integral gain moves the integrated absolute error by 8.7 % and the nadir by 6.4 mg/dL, so a symmetric score cannot express the objective of this loop.
- Told of a meal, the predictive controller scores an IAE of 50.9 against 83.7 for the PID.

Where else this applies. The one-way pump returns in Chapter 3, and anti-windup at a bound set by the supply in Volume 2, Chapter 3.

References

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

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

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

[4] E. F. Pfeiffer, C. Thum, and A. H. Clemens, "The artificial beta cell - a continuous control of blood sugar by external regulation of insulin infusion (glucose controlled insulin infusion system)," *Hormone and Metabolic Research*, vol. 6, no. 5, pp. 339–342, 1974.

[5] A. M. Albisser, B. S. Leibel, T. G. Ewart, Z. Davidovac, C. K. Botz, and W. Zingg, "An artificial endocrine pancreas," *Diabetes*, vol. 23, no. 5, pp. 389–396, 1974.

[6] C. Cobelli, E. Renard, and B. Kovatchev, "Artificial pancreas: Past, present, future," *Diabetes*, vol. 60, no. 11, pp. 2672–2682, 2011.

[7] G. M. Steil, K. Rebrin, C. Darwin, F. Hariri, and M. F. Saad, "Feasibility of automating insulin delivery for the treatment of type 1 diabetes," *Diabetes*, vol. 55, no. 12, pp. 3344–3350, 2006.

[8] R. S. Parker, F. J. Doyle, III, and N. A. Peppas, "A model-based algorithm for blood glucose control in type i diabetic patients," *IEEE Transactions on Biomedical Engineering*, vol. 46, no. 2, pp. 148–157, 1999.

[9] R. Hovorka, V. Canonico, L. J. Chassin, U. Haueter, M. Massi-Benedetti, M. O. Federici, T. R. Pieber, H. C. Schaller, L. Schaupp, T. Vering, and M. E. Wilinska, "Nonlinear model predictive control of glucose concentration in subjects with type 1 diabetes," *Physiological Measurement*, vol. 25, no. 4, pp. 905–920, 2004.

[10] B. W. Bequette, "A critical assessment of algorithms and challenges in the development of a closed-loop artificial pancreas," *Diabetes Technology & Therapeutics*, vol. 7, no. 1, pp. 28–47, 2005.

[11] C. Ellingsen, E. Dassau, H. Zisser, B. Grosman, M. W. Percival, L. Jovanovič, and F. J. Doyle, III, "Safety constraints in an artificial pancreatic β cell: An implementation of model predictive control with insulin on board," *Journal of Diabetes Science and Technology*, vol. 3, no. 3, pp. 536–544, 2009.

[12] R. Hovorka, "Closed-loop insulin delivery: from bench to clinical practice," *Nature Reviews Endocrinology*, vol. 7, no. 7, pp. 385–395, 2011.

[13] S. A. Brown, B. P. Kovatchev, D. Raghinaru, J. W. Lum, B. A. Buckingham, Y. C. Kudva, L. M. Laffel, C. J. Levy, J. E. Pinsker, R. P. Wadwa, E. Dassau, F. J. Doyle, III, S. M. Anderson, M. M. Church, V. Dadlani, L. Ekhlaspour, G. P. Forlenza, E. Isganaitis, D. W. Lam, C. Kollman, and R. W. Beck, "Six-month randomized, multicenter trial of closed-loop control in type 1 diabetes," *New England Journal of Medicine*, vol. 381, no. 18, pp. 1707–1717, 2019.

[14] B. P. Kovatchev, M. Breton, C. Dalla Man, and C. Cobelli, "In silico preclinical trials: A proof of concept in closed-loop control of type 1 diabetes," *Journal of Diabetes Science and Technology*, vol. 3, no. 1, pp. 44–55, 2009.

[15] R. N. Bergman, Y. Z. Ider, C. R. Bowden, and C. Cobelli, "Quantitative estimation of insulin sensitivity," *American Journal of Physiology-Endocrinology and Metabolism*, vol. 236, no. 6, p. E667, 1979.

[16] C. Dalla Man, R. A. Rizza, and C. Cobelli, "Meal simulation model of the glucose-insulin system," *IEEE Transactions on Biomedical Engineering*, vol. 54, no. 10, pp. 1740–1749, 2007.

[17] D. E. Rivera, M. Morari, and S. Skogestad, "Internal model control: PID controller design," *Industrial & Engineering Chemistry Process Design and Development*, vol. 25, no. 1, pp. 252–265, 1986.

[18] S. J. Qin and T. A. Badgwell, "A survey of industrial model predictive control technology," *Control Engineering Practice*, vol. 11, no. 7, pp. 733–764, 2003.

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

[20] R. Gondhalekar, E. Dassau, and F. J. Doyle, "Periodic zone-MPC with asymmetric costs for outpatient-ready safety of an artificial pancreas to treat type 1 diabetes," *Automatica*, vol. 71, pp. 237–246, 2016.

[21] S. A. Weinzimer, G. M. Steil, K. L. Swan, J. Dziura, N. Kurtz, and W. V. Tamborlane, "Fully automated closed-loop insulin delivery versus semiautomated hybrid control in pediatric patients with type 1 diabetes using an artificial pancreas," *Diabetes Care*, vol. 31, no. 5, pp. 934–939, 2008.

[22] C. V. Rao, J. B. Rawlings, and D. Q. Mayne, "Constrained state estimation for nonlinear discrete-time systems: stability and moving horizon approximations," *IEEE Transactions on Automatic Control*, vol. 48, no. 2, pp. 246–258, 2003.

[23] H. Lee and B. W. Bequette, "A closed-loop artificial pancreas based on model predictive control: Human-friendly identification and automatic meal disturbance rejection," *Biomedical Signal Processing and Control*, vol. 4, no. 4, pp. 347–354, 2009.

[24] K. J. Åström and T. Hägglund, "Automatic tuning of simple regulators with specifications on phase and amplitude margins," *Automatica*, vol. 20, no. 5, pp. 645–651, 1984.

[25] T. Battelino, T. Danne, R. M. Bergenstal, S. A. Amiel, R. Beck, T. Biester, E. Bosi, B. A. Buckingham, W. T. Cefalu, K. L. Close, C. Cobelli, E. Dassau, J. H. DeVries, K. C. Donaghue, K. Dovc, F. J. Doyle, III, S. Garg, G. Grunberger, S. Heller, L. Heinemann, I. B. Hirsch, R. Hovorka, W. Jia, O. Kordonouri, B. Kovatchev, A. Kowalski, L. Laffel, B. Levine, A. Mayorov, C. Mathieu, H. R. Murphy, R. Nimri, K. Nørgaard, C. G. Parkin, E. Renard, D. Rodbard, B. Saboo, D. Schatz, K. Stoner, T. Urakami, S. A. Weinzimer, and M. Phillip, "Clinical targets for continuous glucose monitoring data interpretation: Recommendations from the international consensus on time in range," *Diabetes Care*, vol. 42, no. 8, pp. 1593–1603, 2019.

[26] R. W. Beck, R. M. Bergenstal, T. D. Riddlesworth, C. Kollman, Z. Li, A. S. Brown, and K. L. Close, "Validation of time in range as an outcome measure for diabetes clinical trials," *Diabetes Care*, vol. 42, no. 3, pp. 400–405, 2019.

[27] B. P. Kovatchev, D. J. Cox, L. A. Gonder-Frederick, and W. Clarke, "Symmetrization of the blood glucose measurement scale and its applications," *Diabetes Care*, vol. 20, no. 11, pp. 1655–1658, 1997.

[28] U.S. Food and Drug Administration, "Integrated continuous glucose monitoring system," Code of Federal Regulations, Title 21, Section 862.1355, revised as of April 1, 2024, U.S. Government Publishing Office, govinfo.gov, 2024.

[29] G. Cappon, M. Vettoretti, G. Sparacino, and A. Facchinetti, "Continuous glucose monitoring sensors for diabetes management: A review of technologies and applications," *Diabetes & Metabolism Journal*, vol. 43, no. 4, pp. 383–397, 2019.

[30] U.S. Food and Drug Administration, "Alternate controller enabled infusion pump," Code of Federal Regulations, Title 21, Section 880.5730, revised as of April 1, 2024, U.S. Government Publishing Office, govinfo.gov, 2024.

[31] C. K. Boughton and R. Hovorka, "New closed-loop insulin systems," *Diabetologia*, vol. 64, no. 5, pp. 1007–1015, 2021.

[32] U.S. Food and Drug Administration, "Interoperable automated glycemic controller," Code of Federal Regulations, Title 21, Section 862.1356, revised as of April 1, 2024, U.S. Government Publishing Office, govinfo.gov, 2024.

2.14 Exercises

2.1 Concept check · 2 min

A glucose loop is built on the insulin pump of this chapter. In the convention of Section 1.3, which controller action does the loop need, and why?

(a) Reverse acting, $a = +1$, because the pump must rise when the glucose rises above setpoint

(b) Reverse acting, $a = +1$, because the process gain from the pump to the glucose is positive

(c) Direct acting, $a = -1$, because insulin lowers glucose and the process gain is negative

(d) Either action, because a negative proportional gain entered in the app reverses the loop

2.2 Concept check · 1 min

In the noise-free closed-loop day, exercise raised to 60 % at 20 h drives glucose from 74.9 to 47.8 mg/dL within 0.04 h. The next execution sets the pump to zero, where it stays for 0.84 h. Which change to the controller would shorten the low from that moment on?

(a) A larger proportional gain, which cuts the pump harder for each mg/dL of fall

(b) A shorter integral time, which empties the stored integral faster at the floor

(c) Derivative action, which answers the rate of the fall before its size

(d) None of them, because the pump is at zero and the insulin stored clears at plant rates

2.3 Concept check · 2 min

Table 2.6 repeats the closed-loop day at $K_c = 0.3$ for integral times from 6 to 0.4 h. What does it show between $\tau_I = 6\,\text{h}$ and $\tau_I = 1.2\,\text{h}$, a fivefold change in integral gain?

(a) The integrated error rises by under 9 % while the nadir falls 6.4 mg/dL and crosses the low flag

(b) The integrated error has a minimum near the shipped 3 h, which is why that value is shipped

(c) The meal peak falls by several mg/dL as the integral gets faster, at the price of pump travel

(d) The nadir rises as the integral gets faster, because glucose returns to its setpoint sooner

2.4 Theory and calculation · 6 min

Use Equation (2.4) with $g_b = 291\,\text{mg/dL}$, $p_1 = 0.0317/\text{min}$, $f = 0.8$, $v_g = 12\,\text{L}$, $s_i = 0.029$, $k_e = 0.09/\text{min}$ and the fasting appearance $d = 1000$.

(a) Find the fasting glucose at infusion rates of 3.0 and 4.0 µU/min, and the secant gain between them. *Checkpoint:* 76.22 and 57.67 mg/dL, and −18.55 mg/dL per µU/min.

(b) Find the tangent gain at each of the two rates and the ratio between them. *Checkpoint:* −24.5 and −14.0 in the same units, a ratio of 1.75.

(c) Find from Equation (2.5) the pump rate that holds 80 mg/dL at $d = 1000$ and after a 30 % basal drift, $d = 1300$. *Checkpoint:* 2.85 and 3.63 µU/min.

2.5 Theory and calculation · 6 min

The fasting step test of Section 2.4.3 moves the pump from 3.0 to 4.0 µU/min. The glucose response crosses 28.3 % and 63.2 % of its change at 0.553 and 1.340 h after the step, and the process gain is $K_p = -18.55$ mg/dL per µU/min.

(a) Fit a first-order plus dead-time model by the two-point rule $\tau_p = 1.5\,(t_{63} - t_{28})$ and $\theta_p = t_{63} - \tau_p$. *Checkpoint:* $\tau_p = 1.18$ h and $\theta_p = 0.16$ h.

(b) Compare both, in minutes, with $1/k_e$ and $1/p_2$ for $k_e = 0.09$/min and $p_2 = 0.0123$/min, and find the time constant of the glucose balance, $1/(p_1 + s_i\bar{x})$, at $\bar{x} = 33.33$. *Checkpoint:* 9.6 against 11.1 min, 70.8 against 81.3 min, and 1.0 min.

(c) Find the internal model control gain $K_c = \tau_p/(|K_p|(\tau_c + \theta_p))$ for $\tau_c = \tau_p$ and for $\tau_c = \theta_p$. *Checkpoint:* 0.048 and 0.20.

2.6 Theory and calculation · 8 min

The loop holds the reference setpoint of 76.22 mg/dL in Auto with $K_c = 0.3$, integral action off, and the bias at the 3.00 µU/min that holds that glucose with $d = 1000$. The basal drift then steps to 30 %, so $d = 1300$. Use the constants of the first theory exercise.

(a) Find the glucose at which the drift would settle with the pump held at 3.00 µU/min, and the pump rate that an integral would settle at to return glucose to 76.22 mg/dL. *Checkpoint:* 96.25 mg/dL and 3.82 µU/min.

(b) The proportional law is $u_i = 3.00 + 0.3\,(g - 76.22)$. Solve it together with Equation (2.4) for the equilibrium glucose, pump rate and offset. *Checkpoint:* 78.55 mg/dL, 3.70 µU/min and 2.33 mg/dL.

(c) Find the share of the open-loop rise of (a) that the proportional loop removes. *Checkpoint:* 88 %.

2.7 App challenge · 8 min

Load scenario 2, Open-loop actuator step, and change no setting. It starts in Manual at 80 mg/dL with the pump at 2.85 µU/min, and the pump steps by 2.0 µU/min at 2 h, in effect from the execution at 2.016 h. Use the default seed, run to at least 14 h at speed 5 and export the CSV file.

(a) Read the true glucose, column `true_pv`, in the row at 13.944 h. Compute the secant gain of the step and compare it with the tangent gain of −27.0 mg/dL per µU/min at the starting rate. *Checkpoint:* 47.77 mg/dL within 0.05, a secant gain of −16.1, 60 % of the tangent gain.

(b) Find the time after 2.016 h at which `true_pv` has covered 63.2 % of the change from 80 to the value of (a), interpolating between rows, and compare it with the 1.34 h of the fasting one-unit step. *Checkpoint:* 1.17 h, within 0.03 h.

2.8 App challenge · 8 min

Load scenario 1, Reference operating point. It starts in Manual with the pump at 3.0 µU/min, the glucose and its setpoint at 76.22 mg/dL and no sensor noise. Before Run, set the mode to Auto. Run at speed 1, pause between 1.5 and 2.5 h, set Basal drift to 30 %, run on past 14.2 h and export the CSV file.

(a) Read `true_pv` and `mv` in the row at 14.112 h, and the IAE tile of the scorecard. *Checkpoint:* 76.25 mg/dL, 3.82 µU/min and 11.8 mg/dL h, within 0.01, 0.01 and 0.1.

(b) Press the reload button (not Reset), set Integral time to its top, 48 h, which switches integral action off, set the mode to Auto and repeat the run. Read the same row. *Checkpoint:* 78.55 mg/dL and 3.70 µU/min, within 0.01.

Compare both rows with Equation (2.5) and with the proportional-only equilibrium of Equation (2.4), and say why the reload button is needed in (b).

2.9 App challenge · 8 min

Load scenario 9, MPC with an announced meal. It starts at 80 mg/dL in Auto under the MPC with Meals announced on, and a meal of amplitude 1800 arrives at 2 h. Use the default seed and run each case to 24 h at speed 10. For each run read the IAE tile of the scorecard, and find in the exported file the highest and lowest `true_pv` from 2 h on and the time below 70 mg/dL (the number of rows below 70 times 0.168 h).

(a) Run the scenario as shipped. *Checkpoint:* 50.9 mg/dL h, 98.6 and 75.0 mg/dL, 0 h.

(b) Set Meals announced to Off, press Reset and run again. *Checkpoint:* 99.0 mg/dL h, 130.9 and 64.9 mg/dL, 0.84 h.

(c) Set the Controller to PID, press Reset and run again. *Checkpoint:* 83.7 mg/dL h, 116.8 and 66.2 mg/dL, 0.67 h.

Accept an IAE within 0.1 and glucose within 0.1 mg/dL. Explain the order of the three runs.

2.10 Application area · 4 min

A real continuous glucose monitor reads interstitial glucose 5 to 15 min behind plasma, and the model has no such lag. Take the identified model of Section 2.4.3, $K_p = -18.55\,\mathrm{mg/dL}$ per µU/min, $\tau_p = 1.18\,\mathrm{h}$ and $\theta_p = 0.16\,\mathrm{h}$, and treat a 15 min sensor lag as added dead time.

(a) Find the new dead time and the internal model control gain for $\tau_c = \theta_p$. *Checkpoint:* 0.41 h and 0.078.

(b) Compare it with the shipped $K_c = 0.3$, and say what the Limitations section allows an engineer to take from the app to a device. *Checkpoint:* The shipped gain is 3.9 times larger.

2.11 Application area · 3 min

(a) The six-month trial cited in Section 2.2 reported time in the 70 mg/dL to 180 mg/dL range of 71 % under closed loop against 59 % in the control group. Express both in hours per day and give the difference. *Checkpoint:* 17.0 and 14.2 h, 2.9 h a day.

(b) The noise-free closed-loop day of Table 2.5 spends 94.4 % of its samples in the same range, none above 180 mg/dL and none below 64 mg/dL. Find the time it spends below 70 mg/dL and explain why the teaching band of 64 to 104 mg/dL reads 100 % for the same run. *Checkpoint:* 1.34 h.

2.12 Review the agent · 5 min

An AI assistant is asked to retune the glucose loop and replies as follows. “The step test gives $K_p = -18.55\,\mathrm{mg/dL}$ per µU/min, $\tau_p = 1.18\,\mathrm{h}$ and $\theta_p = 0.16\,\mathrm{h}$. The internal model control rule with $\tau_c = \theta_p$ gives $K_c = 0.198$ on the slider grid and $\tau_I = \tau_p = 1.18\,\mathrm{h}$. That is a third less proportional gain than the shipped 0.3, so the loop is more conservative and will lower the risk of low glucose after meals.” Find the flaw in the proposal and say how it could be checked.