INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 3 Feedforward, Cascade, and Ratio Control

CHAPTER 9

An Inner Current Loop Against the Right-Half-Plane Zero of a Boost Converter

Pages 196–224

Edited by John D. Hedengren

How to cite this chapter. "An Inner Current Loop Against the Right-Half-Plane Zero of a Boost Converter," in *Feedforward, Cascade, and Ratio Control*, John D. Hedengren, Ed., vol. 3 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 9, pp. 196–224.

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

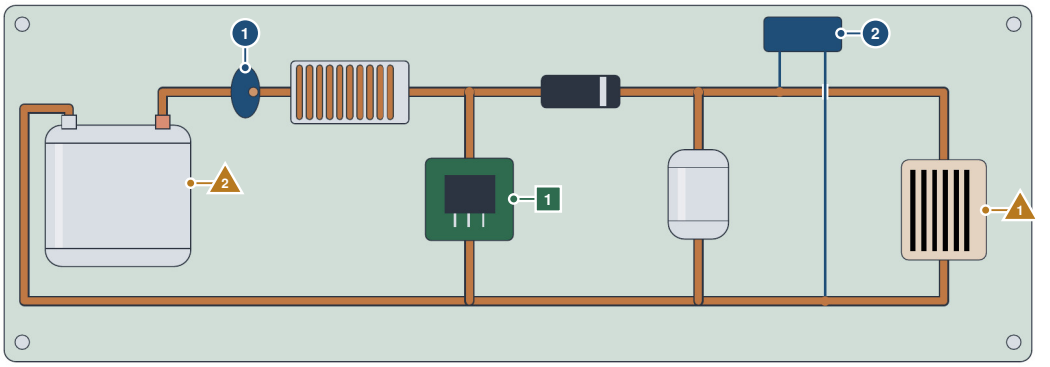
The app for this chapter
https://apmonitor.com/apps/boost-converter/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v03-feedforward-cascade/#boost-converter

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

CHAPTER 9

An Inner Current Loop Against the Right-Half-Plane Zero of a Boost Converter



The controller moves 1 the transistor switch, through the share of each cycle it stays on

The controller reads 1 current in the coil (fast inner loop) 2 output voltage (slow outer loop)

What disturbs it 1 load changes 2 supply voltage sags

Notice that the current is read on the coil ahead of the switch and the voltage across the output, so the inner loop watches the fast store of energy and the outer loop the slow one. Illustration, not to scale.

IN PLAIN TERMS A Power Box That Dips Before It Climbs

The main idea. A boost converter makes a higher voltage from a lower one by opening and closing a switch very fast. The longer the switch stays closed each time, the higher the output ends up.

The challenge. A longer closed time first cuts a coil of wire off from the output, so the output dips before it rises. A second, quicker loop that watches the current in the coil cannot remove the dip, but it keeps that current safe.

Why it matters. These converters sit between solar panels, batteries and the circuits they feed, and those circuits feel every dip.

THREE WAYS TO READ THIS CHAPTER

⊘	**Explore**	12 min	the first page, the fact on switching speed (Section 9.1), the history of averaged models (Section 9.2), the history of predictive control in converters (Section 9.5.1), the two try-it segments and the fact from the simulator (Section 9.6) and the key points
🕮	**Study**	60 min	the whole chapter with its three pause points and the twelve exercises at its end
⛰	**Practice**	25 min	the first page, the results (Section 9.6), Technology in Practice (Section 9.7), Limitations (Section 9.8) and the two open challenges

The control question Raising the duty of a boost converter starves the output before it feeds it, and that right-half-plane zero belongs to the power stage. What does it cost a voltage loop, and what does an inner current loop buy that more gain cannot?

Learning objectives

- Derive the right-half-plane zero from the averaged equations and compute it at an operating point.
- Explain why duty to inductor current is minimum phase and what that gives the outer loop.
- Measure the trade between dip and recovery time under gain, feedforward and component changes.
- Compute both feedforward terms and the output a current limit holds.
- Judge what an explicit MPC buys over the cascade and what the zero still charges it.

Prerequisites Sections 1.4.5 (inner-loop speed rule), 1.4.3 (static feedforward), 1.6.2 (velocity form and back-calculation) and 1.6.4 (scenarios and exported files). Section 1.6.3 for the MPC only.

The app https://apmonitor.com/apps/boost-converter/

See also Volume 7, Chapter 6 for an explicit MPC whose gain over a PID is state feedback, Volume 7, Chapter 7 for the terminal cost of a short horizon. Volume 9, Chapter 4 for offset-free MPC on a drum unit whose level swells.

9.1 Introduction

Almost every electrical load in a modern vehicle, laptop, photovoltaic string or battery pack is fed through a switching converter rather than a linear regulator, because chopping a voltage and filtering the result wastes a fraction of a percent where dissipating the difference would waste tens of percent [1, 2]. The boost converter is the member of that family that produces an output above its input, and it is also the member whose dynamics place the strongest limits on feedback design. Raising the duty cycle keeps the transistor closed for a longer share of each switching period, which is exactly the interval during which the inductor is disconnected from the output and the capacitor supplies the load alone. The immediate consequence of a duty increase is therefore a reduction in delivered charge, and the output falls. Only after the inductor current has grown does the larger current through a shorter delivery window overtake the loss, and the output climbs to a value above where it started.

That behavior is inverse response, and in transfer function terms it is a zero in the right half plane. A right-half-plane zero cannot be canceled by a controller, because canceling it would require an unstable pole in the controller. It cannot be moved by tuning, because it belongs to the power stage. The zero bounds the achievable crossover frequency and imposes a trade between the speed of a response and the size of its initial excursion in the wrong direction, and the classical sensitivity integrals make that trade quantitative rather than qualitative [3–5]. The same response shape appears in a boiler drum whose level rises when cold feedwater is added and in a distillation column whose bottom composition initially moves away from its final value after a reflux increase [6, 7]. The drum of Volume 18, Chapter 4 has a right-half-plane zero that, like the one of this chapter, moves with load.

Textbook treatments of inverse response usually begin from a second-order transfer function with a numerator zero placed by hand, followed by a step response that dips [6]. That presentation ties the phenomenon to an equation rather than to a mechanism. The physical statement is that a fixed quantity of charge per switching period is reallocated between two destinations, and the reallocation is instantaneous while its benefit is delayed. However, a simulation in which the zero is a coefficient supplies no mechanism to examine, and a hardware converter fast enough to be realistic reverses direction in tens of microseconds, well below what an oscilloscope laboratory in a first control course resolves.

This chapter works with a boost converter simulator, built on the framework of Section 1.6, whose right-half-plane zero is produced by the averaged circuit equations rather than inserted by hand. It quantifies the bandwidth cost of the zero and the effect of an inner current loop. The averaged model and current-mode control are both settled engineering. What the simulator adds is a single power

stage on which a single voltage controller and a voltage-to-current cascade share one actuator, one sample period and one set of limits, so the two architectures differ in loop structure and in nothing else. An explicit model predictive controller (MPC) on the same actuator, sample period and limits is a third structure, so the comparison can separate what a hard constraint on the inductor current and a model that contains the zero add to the cascade from what the zero still takes. The scenarios separate what the zero fixes from what the tuning chooses, including one that moves the zero by changing the inductor and one that deepens the dip by changing the capacitor while leaving the zero exactly where it was. Together they cover inverse response, the bandwidth limit a right-half-plane zero imposes, cascade structure, feedforward, anti-windup, actuator saturation and constrained predictive control, principles treated in the process control textbooks [5–7]. The model rests on two conservation statements, a voltage balance on the inductor and a charge balance on the capacitor, and the cascade rests on the time-scale separation between the two storage elements.

> ✦ DID YOU KNOW? **Ten Times the Frequency, a Tenth of the Inductor**
>
> Inductance and capacitance scale with the switching frequency. A national technology assessment of 2015 states that a tenfold increase in frequency gives a tenfold decrease in the required capacitance and inductance [8]. Section 9.4 shows that a smaller inductor also moves the right-half-plane zero farther from the origin.

Section 9.2 traces averaged models and current-mode control. Section 9.3 states the equations, the parameters and the numerics closely enough that the power stage can be reconstructed from the text. Section 9.4 derives the right-half-plane zero and its dependence on load, and Section 9.5 describes the cascade, the single loop, the feedforward and the explicit MPC. Section 9.6 measures the fourteen scenarios, Section 9.7 describes the hardware that carries the cascade, Section 9.8 states the limitations, Section 9.9 discusses the design trade the zero sets, Section 9.10 concludes, and the exercises of Section 9.11 close the chapter.

9.2 Background

A switching converter has no single set of differential equations. During each period the topology changes at least once, so the circuit is described by one linear system while the transistor conducts and by a different one while the diode conducts. Averaging over one switching period answered that difficulty, as the history segment below tells. The result is a continuous model that is accurate well below the switching frequency and that carries the operating point inside its coefficients. Later work extended averaging to quantities that are not slowly varying, by carrying complex Fourier coefficients of the state rather than its period average, which is what resonant converters require [9].

> ⧖ HISTORY **Averaging the Switch Away, 1973 to 1976**
>
> A converter's circuit changes every time its transistor opens or closes, and by 1976 two ways of modeling it had grown up side by side. One wrote the exact state equations of each switched circuit and relied on numerical computation. The other averaged by manipulating equivalent circuits, which gave the designer physical insight [10]. A 1973 paper had characterized the low-frequency behavior of switched dc-dc converters [11]. A 1976 conference paper joined the two traditions: average the two exact state-space descriptions over one switching cycle, then perturb and linearize. Its canonical circuit model predicted that switching puts zeros as well as poles into the response from duty to output, and it pointed to a real zero in the right half-plane of the buck-boost converter that would need proper compensation. Averaged models of this kind agreed with measured responses up to close to half the switching frequency [10].

The small-signal consequences of averaging a boost converter were catalogued soon afterwards. A canonical circuit model collects the duty-cycle dependence into controlled sources and shows that the boost and buck-boost topologies both acquire a numerator zero in the right half plane while the buck

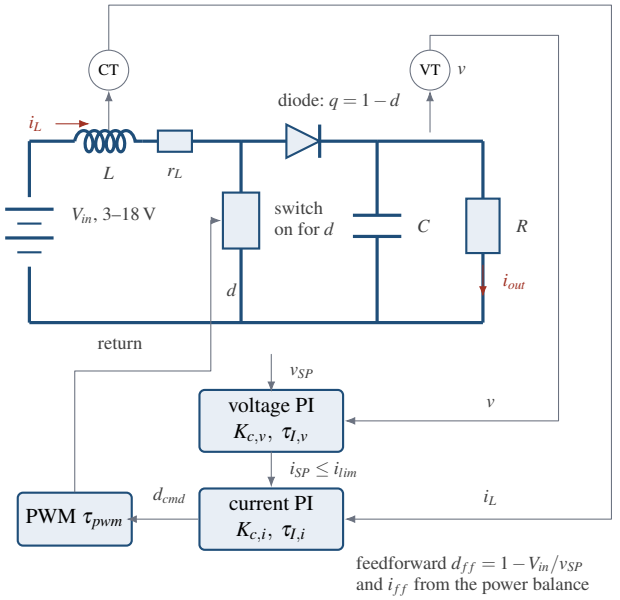


Figure 9.1 Boost power stage and the cascade that drives it. The transistor shorts the switching node for the fraction d of each switching period and the diode delivers inductor charge to C and R for the complement $q = 1 - d$. A current transformer reports i_L to the inner loop and a voltage transmitter reports v to the outer loop, whose output is the inductor-current setpoint i_{SP}. The commanded duty is clamped to $[0, d_{max}]$ and the current setpoint to $[0, i_{lim}]$, and the load current is $i_{out} = v/R$.

does not [12]. An alternative construction replaces the transistor and diode pair with a three-terminal averaged element, which allows the same zero to be read from a small-signal circuit without repeating the matrix algebra for every topology [13]. Modern treatments derive the control-to-output transfer function of the boost stage directly and give the zero as a function of load resistance, duty cycle and inductance [14].

Current-mode control developed as the practical remedy. Regulating the inductor current with a fast inner loop and letting a slower voltage loop set the current reference removes the right-half-plane zero from the loop that has to be tuned, because duty cycle to inductor current is minimum phase. Small-signal models of the arrangement followed its adoption in practice. A first small-signal treatment of current-programmed buck and boost regulators established the low-frequency behavior and the effect of the ramp added for stability [15], and a continuous-time model that also predicts the high-frequency sampling behavior of the current loop followed [16]. The engineering conclusion is long established, that the remedy for a right-half-plane zero is a different loop structure rather than a higher gain, and the simulator of this chapter makes that conclusion measurable. Section 1.4.5 gives the inner-loop speed rule that the same structure obeys on a process plant.

9.3 Converter Model

The simulated hardware is a single boost stage. Figure 9.1 places each symbol of the equations that follow onto the element it belongs to. A stiff source V_{in} drives an inductor L with series resistance r_L. The switching node is pulled to the negative rail by a transistor for a fraction d of each period and released for the complement $q = 1 - d$, during which the diode conducts and the inductor delivers charge to the output capacitor C and the load resistance R. Both states are instrumented, the inductor current i_L and the output voltage v, and both the source voltage and the load current are available to the feedforward terms.

Table 9.1 Variable Classification

Symbol	Meaning	Unit	Range
Controlled variables			
v	Output voltage, outer loop CV	V	13–48
i_L	Inductor current, inner loop CV	A	0–9
Setpoints			
v_{SP}	Output voltage setpoint	V	13–48
i_{SP}	Current setpoint, outer loop MV	A	0–9
Manipulated variable			
d	Duty cycle, inner loop MV	—	0–0.85
Disturbances			
V_{in}	Source voltage	V	3–18
R	Load resistance	Ω	4–48
States			
i_L	Inductor current	A	≥ 0
v	Capacitor voltage	V	≥ 0
d	Applied duty, lagged command	—	0–0.85
Derived outputs			
q	Delivery fraction $1-d$	—	0.15–1
i_{out}	Load current v/R	A	—
d_{ff}	Duty feedforward	—	0–0.85
i_{ff}	Current feedforward	A	0–9
η	Instantaneous P_{out}/P_{in}	%	—
E_L, E_C	Stored energy in L and in C	J	—
z	Right-half-plane zero	rad/s	—
ω_n	Natural frequency of the LC pair	rad/s	—

9.3.1 Variables and Signal Classification

The model quantities are listed in Table 9.1, separated into what is controlled, what is manipulated, what is disturbed and what is integrated. The classification is stated because the cascade assigns one signal three roles. The inductor current is a state of the power stage, the controlled variable of the inner loop, and the manipulated variable of the outer loop.

9.3.2 Averaged State Equations

Nothing in the simulator resolves an individual switching cycle. Both circuit equations are averaged over one period, with the transistor conducting for the fraction d and the diode for the complement $q = 1 - d$, which is the standard continuous-conduction averaged model [10, 14]. The inductor sees the full source voltage while the switch is closed and must also overcome the output voltage while the diode conducts, so the average opposing voltage is $q\,v$, as shown in (9.1).

$$L\frac{\mathrm{d}i_L}{\mathrm{d}t} = V_{in} - r_L\,i_L - q\,v \tag{9.1}$$

The capacitor is charged only during the diode interval and is discharged by the load at all times, as shown in (9.2).

$$C\frac{\mathrm{d}v}{\mathrm{d}t} = q\,i_L - \frac{v}{R} \tag{9.2}$$

The applied duty follows its command through a single lag that stands for the digital pulse-width modulator update and the gate driver, as shown in (9.3).

$$\tau_{pwm}\frac{\mathrm{d}d}{\mathrm{d}t} = d_{cmd} - d \tag{9.3}$$

where L is the inductance in H, C the output capacitance in F, r_L the inductor series resistance in Ω, R the load resistance in Ω, V_{in} the source voltage in V, τ_{pwm} the modulator lag in seconds and d_{cmd} the duty the controller asks for. Every term in (9.1) is a voltage and every term in (9.2) is a current, so both balance dimensionally.

The inverse response originates in the product $q\,i_L$ in (9.2). A step increase in d reduces q immediately while i_L is a state and cannot move at all in the same instant, so the delivered current $q\,i_L$ drops by exactly $i_L\,\Delta d$ and the capacitor is starved. At the default operating point a duty step of $\Delta d = 0.10$ gives an initial slope of $-i_L\,\Delta d/C = -207\,\mathrm{V/s}$. The recovery has to wait for i_L to build through (9.1), which takes on the order of $L\,i_L/(V_{in} - r_L i_L) = 7.2\,\mathrm{ms}$.

Pause. Try Exercise 9.1 before reading on.

A useful accounting of the same event is the energy balance obtained by multiplying (9.1) by i_L and (9.2) by v and adding, as shown in (9.4).

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\tfrac{1}{2}L i_L^2 + \tfrac{1}{2}C v^2\right) = i_L\,(V_{in} - r_L i_L) - \frac{v^2}{R} \tag{9.4}$$

The switch appears nowhere in (9.4), because an ideal switch stores and dissipates nothing and only decides which reservoir is being filled. During the initial dip the capacitor energy falls while the inductor energy rises faster, so the total stored energy is already increasing while the output voltage is still decreasing. The charge withheld from the output is stored in the inductor as current.

9.3.3 Steady State and the Effect of Inductor Resistance

Setting both derivatives to zero in (9.1) and (9.2) gives $i_L = v/(qR)$ and the steady-state output in (9.5).

$$v = \frac{V_{in}\,q\,R}{q^2 R + r_L}, \qquad q = 1 - d \tag{9.5}$$

With $r_L = 0$ this collapses to the familiar ideal relation $v = V_{in}/(1-d)$, and the application never uses that relation as a model. At the default point the difference matters. A duty of exactly 0.5 gives 23.226 V from (9.5) rather than the ideal 24.000 V, 0.774 V short, and the duty that actually holds 24.000 V is 0.517263. The inductor drops $r_L i_L = 0.414\,\mathrm{V}$, which the duty must make up.

Inverting (9.5) for a target voltage requires the larger root of $Rq^2 - (V_{in}R/v)\,q + r_L = 0$, as shown in (9.6).

$$q = \frac{1}{2R}\left(\frac{V_{in}R}{v} + \sqrt{\left(\frac{V_{in}R}{v}\right)^2 - 4R\,r_L}\right) \tag{9.6}$$

The smaller root is the high-current branch beyond the peak of the voltage-versus-duty curve and is not the physical operating point of a regulator. A negative discriminant means the target voltage is unreachable at any duty, in which case the initialization falls back on $q = 1 - d_{max}$. Equation (9.6) places every automatic scenario at an exact equilibrium at $t = 0$, so the only transient in a run is the one the scenario introduces.

9.3.4 Parameters, Limits and Numerics

Table 9.2 collects the seven parameters, their default values and their adjustable intervals. Three of them are hardware limits rather than tuning choices and are fixed, so the saturation statistics and the displayed limits derive from one set of declared values.

Three hard events are handled once per step by a post-step routine and never by the right-hand side itself, which keeps every intermediate integration stage on a smooth vector field. The applied duty is clamped to $[0, d_{max}]$. The output voltage is clamped at zero from below. If the averaged inductor current reaches zero the converter has left continuous conduction, the state is clamped at zero and a fault is annunciated that clears only once the current exceeds 0.05 A again, because the averaged

Table 9.2 Model Parameters and Their Adjustable Ranges

Symbol	Meaning	Value	Unit	Range
L	Inductance	20	mH	2–60
C	Output capacitance	2.0	mF	0.2–6.0
r_L	Inductor resistance	0.10	Ω	0–0.60
τ_{pwm}	PWM and driver lag	0.2	ms	0–2.0
d_{max}	Maximum duty (hidden)	0.85	—	fixed
i_{lim}	Current setpoint limit (hidden)	9.0	A	fixed
i_{trip}	Overcurrent trip (hidden)	12.0	A	fixed
Default operating point, $V_{in} = 12\,\text{V}$, $R = 12\,\Omega$				
v	Output voltage	24.000	V	—
d	Duty cycle	0.517263	—	—
i_L	Inductor current	4.1430	A	—
P_{in}	Input power	49.72	W	—
P_{out}	Output power	48.00	W	—
η	Efficiency	96.55	%	—
E_L	Inductor energy	171.6	mJ	—
E_C	Capacitor energy	576.0	mJ	—
z	Right-half-plane zero	134.8	rad/s	—
ω_n	*LC* natural frequency	77.68	rad/s	—
ζ	*LC* damping ratio	0.300	—	—

model no longer describes the circuit past that boundary. A separate overcurrent flag with hysteresis annunciates when i_L exceeds $i_{trip} = 12\,\text{A}$ and clears below 10.8 A.

Three clocks run at once, as in every framework app (Section 1.6.1). A fixed-step fourth-order Runge-Kutta integrator advances the states every $\Delta t = 0.1\,\text{ms}$, the two loops execute on a $\Delta t_c = 0.2\,\text{ms}$ cycle that stands for one period of a 5 kHz digital controller, and 500 samples of every logged quantity are recorded per simulated second. The logging rate is set by the phenomenon, because the open-loop inverse response bottoms out 7.5 ms after a duty step and is over by 15.8 ms, and a 5 mH inductor compresses the closed-loop dip into 2.5 ms. Each run exports as comma-separated data, including both states, the applied and commanded duty, both feedforward terms, the operating-point zero and natural frequency, the stored energy in each element, the saturation and fault flags, all seven parameters and both disturbance channels.

9.3.5 Assumptions

The model is switching-averaged and assumes continuous conduction throughout. It says nothing about output ripple, nothing about subharmonic oscillation, and nothing about the peak-current-mode instability that appears above 50 % duty in hardware and that a compensating ramp is added to suppress [16]. Averaging is accurate for behavior well below the switching frequency and breaks down as the closed-loop bandwidth approaches it. Nothing in the equations resolves an individual switching cycle.

The switch and the diode are ideal, with no forward drop, no on-resistance, no reverse-recovery charge and no dead time. A real diode loses roughly a volt, which matters most at low output voltage and would depress the achievable boost ratio in the brownout scenario more than the model shows. The only dissipation is the inductor series resistance r_L. Capacitor equivalent series resistance, core loss and switching loss are absent, so the reported efficiency is an upper bound and the step in output voltage that equivalent series resistance produces at the instant of a load change is missing entirely.

The load is a fixed resistor and the source is ideal. A constant-power load such as a downstream regulator presents a negative incremental resistance and would move the closed-loop poles, and a source with series inductance would add a resonance this model cannot show. Sensing is perfect and instantaneous, with no current-sense filter, no divider dynamics, no quantization and no noise, so the

derivative-free tuning used here is not challenged by measurement noise the way a real current loop is. The only actuator dynamic is the first-order modulator lag of 0.2 ms.

The parameter values are exaggerated on purpose. A 5 kHz converter with a 20 mH inductor and a 2 mF capacitor is slow by the standards of commercial designs, and no product would be built with an inverse response that lasts 16 ms. The exaggeration makes the phenomenon legible on a trend display at reduced playback speed. Every qualitative statement of Section 9.4 survives the scaling, because (9.9) is homogeneous in the component values.

9.4 Origin and Load Dependence of the Right-Half-Plane Zero

Linearizing (9.1) and (9.2) about an operating point (i_0, v_0, q_0) and writing δ for a deviation gives (9.7).

$$\begin{aligned} L\,\frac{\mathrm{d}\delta i}{\mathrm{d}t} &= -r_L\,\delta i - q_0\,\delta v + v_0\,\delta d \\ C\,\frac{\mathrm{d}\delta v}{\mathrm{d}t} &= \quad q_0\,\delta i - \frac{\delta v}{R} - i_0\,\delta d \end{aligned} \tag{9.7}$$

The two duty terms carry opposite signs, and that sign difference is the origin of the inverse response. Additional duty raises the average inductor voltage, which acts through the inductor dynamics, and simultaneously removes $i_0\,\delta d$ of delivered current, which acts immediately. Eliminating δi gives the control-to-output transfer function in (9.8).

$$\frac{\delta v(s)}{\delta d(s)} = \frac{(q_0 v_0 - r_L i_0)\,(1 - s/z)}{LC\,s^2 + \left(\dfrac{L}{R} + r_L C\right) s + \left(\dfrac{r_L}{R} + q_0^2\right)} \tag{9.8}$$

where the numerator zero is given by (9.9).

$$z = \frac{q_0 v_0 - r_L i_0}{L\,i_0} \;\approx\; \frac{q_0^2 R}{L} \;>\; 0 \tag{9.9}$$

The zero is positive for every physically reachable operating point, because $q_0 v_0$ is the average voltage the inductor pushes against and $r_L i_0$ is a small fraction of it. At the default point (9.9) evaluates to $+134.8\,\mathrm{rad/s}$ against the approximation $q_0^2 R/L = 139.8\,\mathrm{rad/s}$, and the denominator of (9.8) gives $\omega_n = 77.68\,\mathrm{rad/s}$ with $\zeta = 0.300$. The steady-state gain of (9.8) is 46.3 V per unit of duty, so half a percent more duty is worth 0.23 V of output.

Two features of (9.9) are tested by separate scenarios. The capacitance does not appear in the zero at all, so a smaller capacitor deepens the initial dip through the $1/C$ in (9.2) and leaves the zero exactly where it was. The inductance appears in the denominator, so a smaller inductor moves the zero away from the origin and relaxes the bandwidth limit. The load appears through $q_0^2 R$, which means the zero moves toward the origin as the converter is loaded more heavily, and the operating point with the most restrictive bandwidth limit is the one at which the converter delivers the most power.

🕮 **Pause.** Try Exercise 9.2 before reading on.

Figure 9.2 plots (9.9) and the natural frequency of (9.8) across the load range the application permits. Between 48 Ω and 4 Ω the zero sweeps from 585 rad/s down to 34.4 rad/s, a factor of 17, while ω_n moves only from 78.7 to 74.5 rad/s. The two curves cross at a load of 7.36 Ω. Above that load the zero lies below the resonance, the most restrictive arrangement, because the crossover must then be placed below a resonant peak the loop cannot attenuate. The default 12 Ω load lies on the favorable side of that crossing, with $z = 1.74\,\omega_n$.

The bandwidth consequence is quantitative. For a stable plant with a single right-half-plane zero at z the Poisson integral forces the sensitivity function to exceed unity over some frequency band whenever it is pushed below unity elsewhere, and the area of that penalty is weighted toward z [3, 5].

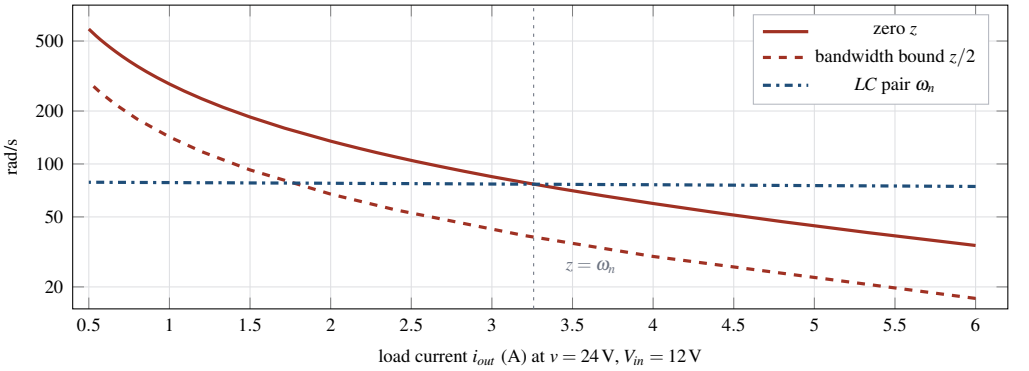


Figure 9.2 Right-half-plane zero (solid) and LC natural frequency (dash-dotted) against load current, with the bandwidth bound $z/2$ (dashed), computed from (9.5) and (9.9) over the load range the application allows. The dotted vertical line marks $z = \omega_n$. The zero sweeps by a factor of 17 between 48 Ω and 4 Ω while ω_n barely moves, and the two cross at $i_{out} = 3.26\,\mathrm{A}$, which is a 7.36 Ω load.

The practical rule that follows is that the crossover frequency of a loop closed around such a plant should stay comfortably below z, and $z/2$ is the usual working bound. Figure 9.2 draws that bound. At the default point it is $67\,\mathrm{rad/s} = 10.7\,\mathrm{Hz}$, which is already below the LC natural frequency, and at a 4 Ω load it falls to 17.2 rad/s. No choice of K_c, τ_I or controller structure changes those numbers, because they are properties of L, C, R and the operating duty.

The reason a cascade helps is that the inner loop closes around a different transfer function. Eliminating δv from (9.7) instead gives the control-to-current response in (9.10).

$$\frac{\delta i(s)}{\delta d(s)} = \frac{v_0\,(Cs + 1/R) + q_0 i_0}{LC\,s^2 + \left(\dfrac{L}{R} + r_L C\right) s + \left(\dfrac{r_L}{R} + q_0^2\right)} \tag{9.10}$$

Every coefficient in the numerator of (9.10) is positive, so the zero of the current response sits at $-(v_0/R + q_0 i_0)/(C v_0)$, in the left half plane. The physical statement is the one (9.1) makes directly, that a longer on-time raises $\mathrm{d}i_L/\mathrm{d}t$ the instant q falls with nothing competing to reverse the sign. The inner loop can therefore be closed as fast as the modulator and the sampling allow. With the current effectively an input, the output node obeys the power balance in (9.11), obtained from (9.4) with the inductor term removed.

$$C v \frac{\mathrm{d}v}{\mathrm{d}t} = i_L\,(V_{in} - r_L i_L) - \frac{v^2}{R} \tag{9.11}$$

Equation (9.11) is first order and its right-hand side is increasing in i_L for all $i_L < V_{in}/2r_L$, which at the default supply is 60 A and therefore never binding. The plant seen by the outer loop retains no inverse response.

9.5 Controller Design

Both loops are proportional-integral with the derivative time set to zero, in the velocity form of Section 1.6.2, executing every $\Delta t_c = 0.2\,\mathrm{ms}$. The outer loop converts volts of error into amps of

current demand, as shown in (9.12).

$$\begin{aligned} e_{v,k} &= v_{SP,k} - v_k \\ \Delta u^I_{v,k} &= \frac{K_{c,v}}{\tau_{I,v}} e_{v,k} \Delta t_c \\ s_{v,k} &= s_{v,k-1} + K_{c,v}\left(e_{v,k} - e_{v,k-1}\right) \\ c_{v,k} &= i_{ff,k} + s_{v,k} + I_{v,k-1} \\ i_{SP,k} &= \text{clamp}\left(c_{v,k}, 0, i_{lim}\right) \\ I_{v,k} &= I_{v,k-1} + \Delta u^I_{v,k} + \frac{\Delta t_c}{\tau_{I,v}}\left(i_{SP,k} - c_{v,k}\right) \end{aligned} \tag{9.12}$$

The inner loop converts amps of error into duty, as shown in (9.13).

$$\begin{aligned} e_{i,k} &= i_{SP,k} - i_{L,k} \\ \Delta u^I_{i,k} &= \frac{K_{c,i}}{\tau_{I,i}} e_{i,k} \Delta t_c \\ s_{i,k} &= s_{i,k-1} + K_{c,i}\left(e_{i,k} - e_{i,k-1}\right) \\ c_{i,k} &= d_{ff,k} + s_{i,k} + I_{i,k-1} \\ d_{cmd,k} &= \text{clamp}\left(c_{i,k}, 0, d_{max}\right) \\ I_{i,k} &= I_{i,k-1} + \Delta u^I_{i,k} + \frac{\Delta t_c}{\tau_{I,i}}\left(d_{cmd,k} - c_{i,k}\right) \end{aligned} \tag{9.13}$$

where I is the integral accumulator of each loop, s accumulates the proportional increments without limit and c is the output the loop demands before clipping. In the symbols of Equations (1.9) to (1.12), $s = S$, $c = v$, the feedforward term is the bias u_b and the tracking time is $T_t = \tau_I$. Anti-windup is back-calculation, so while an output is clipped the final term of each update moves the integral by $\Delta t_c/\tau_I$ of the clipped excess (1/250 per scan for the outer loop at $\tau_{I,v} = 0.05\,\text{s}$ and 1/40 for the inner loop at $\tau_{I,i} = 8\,\text{ms}$) [17]. With the tracking time equal to τ_I the error terms cancel in the update of a clipped loop, so its integral settles at the limit less the feedforward bias, $I_v = i_{lim} - i_{ff}$ for an outer loop held at i_{lim}. At a limit the tracking time shapes the release, as Section 9.6.5 shows. Both loops are reverse acting in the ISA sense, because a low output voltage calls for more current and a low current calls for more duty. Neither loop has a derivative term, and neither output is slew limited beyond what τ_{pwm} imposes on the actuator.

The alternative architecture replaces (9.12) and (9.13) by a single proportional-integral controller of identical form with gain $K_{c,dv}$ and integral time $\tau_{I,dv}$, taking e_v and producing d_{cmd} clamped to $[0, d_{max}]$. In that configuration nothing measures the inductor current for control purposes and nothing enforces i_{lim}. The inductor current and its feedforward reference remain recorded together, so the size of the unregulated current excursion is measurable.

Both feedforward terms come from the measured input voltage and load current, with the declared $r_L = 0.1\,\Omega$ in the current term. The duty feedforward is the ideal duty for the requested boost ratio, and the current feedforward is the inductor current the power balance $V_{in} i_L = v_{SP}\, i_{out} + r_L i_L^2$ demands at the present load, as shown in (9.14).

$$\begin{aligned} d_{ff} &= \text{clamp}\left(1 - \frac{V_{in}}{v_{SP}}, 0, d_{max}\right) \\ i_{ff} &= \text{clamp}\left(\frac{V_{in} - \sqrt{V_{in}^2 - 4 r_L v_{SP}\, i_{out}}}{2 r_L}, 0, i_{lim}\right) \end{aligned} \tag{9.14}$$

Table 9.3 Controller Settings and Their Adjustable Ranges

Symbol	Loop	Unit	Default	Range
$K_{c,v}$	Voltage, outer	A/V	0.12	0–1.00
$\tau_{I,v}$	Voltage, outer	s	0.05	0.01–3
$K_{c,i}$	Current, inner	%/A	50	0–200
$\tau_{I,i}$	Current, inner	ms	8	1–100
$K_{c,dv}$	Voltage only	%/V	0.1	0–3
$\tau_{I,dv}$	Voltage only	s	0.20	0.01–3
v_{SP}	Setpoint	V	24.0	13–48

The smaller root is the operating one. A negative discriminant means the demanded power cannot be pushed through r_L at any current, and i_{ff} is set to i_{lim} in that case, the largest deliverable request. The r_L in (9.14), in the integrator preload and in the bias shift described below is the declared value, because no instrument measures the inductor resistance, so a resistance changed in the simulated hardware reaches the cascade only through the measured current and voltage. Both terms enter as the bias of their loop, so the integral accumulators carry only what the feedforward misses. At the default point $d_{ff} = 0.5000$ against a true steady-state duty of 0.517263, so the inner integrator holds 0.017263 permanently and that offset is precisely the inductor drop. Both terms are static feedforward computed from a balance, as in Section 1.4.3. A dynamic compensator built as in Equation (1.4) around the duty-to-voltage path would have to invert its right-half-plane zero, which takes an unstable pole.

Three mechanisms keep transitions free of actuator jumps. In MANUAL the operator holds the duty and all three controllers track it, as Section 1.6.2 describes. In AUTO the controller that is not in use tracks the applied duty every step, so switching architecture mid-run moves nothing. Switching the feedforward on or off would step every bias by the whole feedforward amount, so each integral is shifted by the same quantity in the opposite direction at the instant of the switch, which makes the comparison in the load scenario a comparison of two responses rather than of two different startup transients. The manual duty command tracks the controller output while in AUTO, so the reverse transition is bumpless as well.

Default tuning is listed in Table 9.3. The inner gain places the current loop crossover near $K_{c,i}\, v_0/L = 600\,\mathrm{rad/s}$, using the high-frequency approximation $\delta i/\delta d \approx v_0/(Ls)$ from (9.10), which is about four times the zero and far faster than the outer loop, whose recovery to within 0.5 V of a 6 V setpoint change takes 196 ms. That separation is the inner-loop speed rule of Section 1.4.5, which Section 9.6.4 tests on this plant.

9.5.1 Explicit Model Predictive Control

A third architecture moves the duty directly in automatic operation. It is a model predictive controller computed in advance rather than at run time, where the MPC option of Section 1.6.3 solves its program online at every decision. A quadratic program whose cost gradient and constraint bounds vary linearly with a parameter vector has an optimal solution that is a continuous piecewise affine function of that vector, defined over a finite set of polyhedra, so the constrained linear quadratic regulator (LQR) reduces to a stored list of polyhedra, each paired with its own affine control law, and a decision becomes a point-location search [18]. Predictive control of the boost stage has also been formulated online, by enumerating switch sequences for direct voltage control [19]. The cascade remains the default architecture, and only scenario 13 selects the MPC.

HISTORY Predictive Control Reaches the Converter, 1980s to 2000s

Model predictive control emerged in the 1960s and over the next decade became an established strategy for constrained plants, mostly in the process industry. The power electronics community began to investigate it in the 1980s, but the meager computing resources of the time, together with new power semiconductors that allowed higher switching frequencies, limited its use. After a hiatus of about two decades, more powerful microprocessors revived the interest in the 2000s [20]. One route that followed solves the optimization offline for all possible states and stores the resulting control law in look-up tables, the explicit form this chapter's MPC takes. The same review warns that such tables suit only small problems, because their memory grows exponentially with the size of the problem [20].

The prediction model is (9.7) with the modulator lag (9.3) appended, linearized at the declared default point of Table 9.2 ($L = 20\,\text{mH}$, $C = 2\,\text{mF}$, $r_L = 0.1\,\Omega$, $\tau_{pwm} = 0.2\,\text{ms}$, $V_{in} = 12\,\text{V}$, $R = 12\,\Omega$, $v_0 = 24\,\text{V}$, $d_0 = 0.517263$, $i_0 = 4.1430\,\text{A}$). Its state $x = (i_L - i_{ss},\, v - v_{ss},\, d - d_{ss})$ and its input $\delta u = d_{cmd} - d_{ss}$ are deviations from a target (i_{ss}, v_{ss}, d_{ss}), and an exact zero-order hold over $T_s = \Delta t_c = 0.2\,\text{ms}$ gives $x_{k+1} = \Phi x_k + \Gamma\,\delta u_k$. No step test was used to identify it. Equation (9.7) is already that linearization, and a finite-difference Jacobian taken from the simulator's right-hand side matches its coefficients within a relative 1.3×10^{-9}. Its open-loop eigenvalues are $-23.3 \pm 74.1\text{j}$ and –5000/s, and duty to voltage keeps the zero of (9.9) at +134.8 rad/s, so the prediction contains the inverse response. The same matrices serve every setpoint.

Every 0.2 ms the controller plans three duty moves, one per period, and applies the first, as shown in (9.15).

$$\begin{aligned} \min_{\delta u_0, \delta u_1, \delta u_2} \; & \sum_{k=1}^{2} Q_v x_{2,k}^2 + x_3^T P x_3 + R \sum_{k=0}^{2} \delta u_k^2 \\ \text{subject to } & x_{k+1} = \Phi x_k + \Gamma\,\delta u_k, \quad x_0 = \hat{x} \\ & 0 \le d_{ss} + \delta u_k \le d_{max} \\ & i_{floor} \le i_{ss} + x_{1,k+1} \le i_{lim} \end{aligned} \tag{9.15}$$

where $x_{1,k}$ and $x_{2,k}$ are the predicted current and voltage deviations, $\hat{x}$ is the estimated state measured from the target, and P solves the discrete algebraic Riccati equation for Φ, Γ, $\text{diag}(0, Q_v, 0)$ and R. Charging the Riccati cost-to-go after the last move is the usual route to closed-loop stability in receding-horizon control [21], and no terminal set is imposed here. Volume 7, Chapter 7 shows a predictive controller stripped of that terminal cost becoming unstable at short horizons. Bryson's rule on the 0.5 V band of the voltage loop and on the duty headroom $d_{max} - d_0 = 0.3327$ gives $Q_v = 4/\text{V}^2$ and $R = 9.032$, with no weight on the current and no move penalty. With no row active the first move is the LQR's, $\delta u = -(1.2688/\text{A}\,x_1 + 0.103\,81/\text{V}\,x_2 + 0.24164\,x_3)$, which places the closed-loop poles at -135.9, -1434 and –4772/s. The slowest of them lies near the mirror image of the zero, at $-1.008\,z$.

All twelve rows of (9.15) are hard. The duty range and i_{lim} are read as the cascade reads them, but the current rows bound the predicted inductor current itself, where the cascade clamps only the setpoint of its inner loop. The floor $i_{floor} = \min(0.3\,\text{A}, i_{ss}/2)$ keeps the plan out of discontinuous conduction. With a floor of zero a step from 24 to 13 V reaches zero current, because the linear model's one-period current prediction errs by up to 125 mA that far from 24 V. Unlike a box on the input alone, the current rows can make the program infeasible, for example when the current already sits at the limit and the three moves cannot keep the prediction below it.

The target and the limits enter only the right-hand sides, through the parameter $\theta = (\hat{x}, -d_{ss}, d_{max} - d_{ss}, i_{lim} - i_{ss}, i_{floor} - i_{ss}, 1)$, so one table serves every setpoint, load, supply and limit. Writing (9.15) as $\min \frac{1}{2} z^T G z + z^T F \hat{x}$ subject to $C_q z \ge b(\theta)$ with z the three moves, an active set $\mathscr{A}$ gives the moves in (9.16).

$$\begin{aligned} z_u &= -G^{-1} F \hat{x}, \qquad \lambda_{\mathscr{A}} = W_{\mathscr{A}} \left(b_{\mathscr{A}}(\theta) - C_{\mathscr{A}} z_u \right) \\ z &= z_u + G^{-1} C_{\mathscr{A}}^T \lambda_{\mathscr{A}}, \qquad W_{\mathscr{A}} = \left(C_{\mathscr{A}} G^{-1} C_{\mathscr{A}}^T \right)^{-1} \end{aligned} \tag{9.16}$$

The law holds on the critical region where $\lambda_{\mathscr{A}} \geq 0$ and every row of $C_q z \geq b(\theta)$ is met, so a region that accepts θ returns the exact optimum. Every candidate set of up to three rows was tested, and a set enters the table only if its region contains a ball of nonzero radius somewhere in a stated domain: setpoints 13 to 48 V, supplies 3 to 18 V, loads 4 to 48 Ω, duty ceilings 0.5 to 0.95, current limits 3 to 20 A, outputs within 30 V of the target and currents from zero to 3 A above the limit. The table holds 109 regions. An independent enumeration that ties each target coordinate to its physical relation (a target current never above the limit, a floor never above the target current, a target duty inside its range) finds 85 reachable regions, all of them in the table. The other 24 exist only where the box lets a target coordinate pass its limit, and none was chosen in 359 976 closed-loop decisions. Outputs up to 60 V from the target would add six regions that are not stored.

Each decision scans the stored regions in a fixed order that begins with the unconstrained law. Over 135 000 test states, 15 000 of them within 10^{-6} of a region boundary and 105 000 taken from closed-loop runs of every scenario, the table's duty differs by at most 8.8×10^{-11} from that of the framework's dual active-set solver applied online to the identical program, and both agree with an independent solution of the optimality conditions to 3.2×10^{-9}. Of those states 4408 have no feasible plan, and the lookup misses every one of them. In Node.js a lookup averages 0.18 µs over 209 986 closed-loop states and 1.86 µs in the slowest region, against a median of 6.0 µs for the online solve. Played in Chrome by its animation loop, which held 60 frames per second in every case, the whole controller execution under the lookup takes 0.22 to 0.24 % of one core at the default playback speed of 0.1 times real time and 1.20 to 1.25 % at real time (5000 decisions per second), against 0.13 and 0.43 % for the cascade, whose execution includes the MPC's estimator. The online solve takes 1.06 to 1.11 % and 5.96 to 6.95 %. Every timing was taken on a host shared with other computations.

The MPC reads the measured inductor current, output voltage, input voltage and load current (the two signals the feedforward reads), the setpoint, its own commands, d_{max} and i_{lim}. A constant-gain extended Kalman filter estimates (i_L, v, d, w_C, w_L), where w_L is a voltage added to (9.1) and w_C a current drawn from (9.2), both random walks. It predicts with the declared nonlinear equations (9.1) to (9.3) at $R = 12\,\Omega$, integrated by a fourth-order Runge-Kutta method in two half-period substeps, and corrects with the steady Kalman gain of the design linearization augmented with w_C and w_L, designed for 0.01 A and 0.01 V of measurement noise and random walks of 0.05 A and 0.05 V per period, which leaves error poles at -792 ± 789j and -2548 ± 2446j 1/s. The target is the declared steady state at v_{SP} for the measured supply and the measured load resistance $\hat{R} = v / i_{out}$, with w_L passed through a first-order filter of 0.1 s to give $\bar{w}_L$, as shown in (9.17).

$$
\begin{aligned}
(V_{in} + \bar{w}_L)\, i_{ss} - r_L i_{ss}^2 &= \frac{v_{SP}^2}{\hat{R}}, \qquad v_{ss} = v_{SP} \\
d_{ss} &= 1 - \frac{v_{SP}}{\hat{R}\, i_{ss}}
\end{aligned}
\tag{9.17}
$$

The smaller root of the power balance is taken, as in (9.14). The filtered estimate shifts the target, the target-shift form of offset-free predictive control [22], so the output returns to v_{SP} with r_L or V_{in} in error. The offset-free MPC of Volume 9, Chapter 4 carries a fuel bias in its estimator to the same end. Where the setpoint cannot be held, in an overload past i_{lim} or a brownout past d_{max}, the duty and current targets become those of the nearest reachable steady state while v_{ss} stays at v_{SP}. An earlier design that estimated the load current instead of measuring it did not settle with $C = 0.5$ mF (16.82 to 45.68 V over 4 s of scenario 8), because the estimate absorbed the capacitor's model error, while the adopted design reaches the band in 18 ms on the same case.

When no stored region contains θ the decision is rejected as a miss. The duty then passes without a step to the cascade, which has tracked the applied duty at every execution (its outer loop toward the measured current), the historian logs the fallback, and the table regains the duty once 20 consecutive decisions have found a region, ramped in over another 20 decisions (4 ms). A miss forced in the middle of scenario 13's setpoint step moves the duty by 0.021 at the hand-over. With the current held at the limit in an overload the program has no solution, so a miss there leaves the cascade in charge

Table 9.4 Scenario Definitions

#	Scenario	Mode	Event
1	`steady`	MANUAL	none, $d = 0.517263$
2	`inverse`	MANUAL	d 0.5173→0.6173 at 0.2 s
3	`sp-step`	AUTO	v_{SP} 24→30 V at 0.3 s
4	`load-step`	AUTO	R 12→6→12 Ω
5	`vin-sag`	AUTO	V_{in} 12→8 V at 0.4 s
6	`duty-sat`	AUTO	$V_{in} \to 3.5$ V, $R = 24\,\Omega$
7	`current-limit`	AUTO	R 12→4→12 Ω
8	`small-cap`	AUTO	$C = 0.5$ mF, 24→28 V
9	`small-ind`	AUTO	$L = 5$ mH, 24→28 V
10	`voltage-only`	AUTO	single PI, 24→30 V
11	`cascade`	AUTO	cascade, 24→30 V
12	`aggressive`	AUTO	$K_{c,v} = 0.5$, $\tau_{I,v} = 0.02$ s
13	`explicit-mpc`	AUTO	MPC, v_{SP} 24→30 V, R 12→4→12 Ω
14	`challenge`	AUTO	R 12→7 Ω at 1 s

until the load is released. Forced at 1.2 s in scenario 13, it returns the duty to the MPC at 2.004 s, and meanwhile the cascade reaches 9.03 A where the MPC alone holds 8.84 A.

Three models stay separate. The simulated plant evaluates (9.1) to (9.3) with whatever component values the run carries, while the MPC's prediction and estimator use the declared values frozen when the application loads and never read the plant's. A changed L, C, r_L or τ_{pwm} is therefore a mismatch that the MPC absorbs through its disturbance estimate, much as the cascade absorbs a feedforward error through its integrators. The estimator executes under every architecture but feeds nothing to the proportional-integral loops, so each scenario under the cascade and the single loop reproduces its earlier run bit for bit on every previously logged column.

9.6 Results

Table 9.4 lists the fourteen scenarios the application supplies. Each scenario fixes only the quantities it requires and leaves everything else at the defaults. Summary metrics for all of them, with the shipped settings and from the scorecard of Section 1.6.4, are given in Table 9.5, read at the stated simulated time. Overshoot and settling are reported only after a setpoint change, so the disturbance cases carry a dash. In scenario 10 the single loop commands the duty directly, and in scenario 13 the MPC does. The displayed current reference is then not a commanded actuator, so its MV travel and saturation entries are dashes. The duty-cycle scorecard still reports actual actuator travel and saturation. If architectures are switched during a run, current-setpoint travel and saturation accumulate only while that setpoint commands the inner loop; saturation is normalized by that active time. Depths, peaks and crossing times quoted below are computed from the state at every 0.1 ms integration step, because the 2 ms log reads the bottom of a brief dip late and shallow (0.21 V against 0.50 V for the 5 mH inductor of scenario 9). Recovery is the first instant at which the output enters ±0.5 V of setpoint and stays inside it, which agrees with the settling column of Table 9.5 to within one 2 ms sample.

9.6.1 Open-Loop Inverse Response

Scenario 2 records the inverse response with no controller acting and is shown in Figure 9.3. The converter starts in exact equilibrium at 24.000 V with the duty held constant at 0.517263, and at $t = 0.2$ s a timed event steps the duty by ten percentage points to 0.617263. The new steady state is 29.67 V, so the eventual motion is upward by more than 5 V. The output falls 0.723 V first, bottoms out 7.5 ms after the step, crosses back through its starting value at 15.8 ms and reaches 31.46 V at the first peak of the underdamped recovery, 62 ms after the step. No tuning is implicated, because no controller is active. The output is recorded together with $V_{in}/(1-d)$, the steady-state voltage the applied duty requests, and the two signals move in opposite directions for the first 15.8 ms.

Table 9.5 Measured Voltage-Loop Performance at the Stated Simulated Time

#	t (s)	IAE (V s)	OS (%)	settling (s)	MV travel (A)	sat (%)	peak i_L (A)
1	1.03	0.000	—	—	0.00	0.0	4.14
2	0.24	0.078	—	—	3.26	0.0	7.40
3	1.08	0.305	14.8	0.194	3.49	0.0	7.04
4	4.80	0.556	—	—	13.26	1.0	9.00
5	1.97	0.208	—	—	4.04	0.0	6.84
6	2.00	7.267	—	—	11.76	74.2	6.05
7	4.78	8.376	—	—	14.01	41.6	9.47
8	1.02	0.123	10.7	0.016	1.94	0.0	5.90
9	1.04	0.115	8.8	0.028	1.90	0.0	5.88
10	1.95	0.557	30.5	0.134	—	—	7.59
11	2.00	0.305	14.8	0.194	3.49	0.0	7.04
12	2.00	0.218	177.9	0.056	8.92	0.5	9.50
13	3.00	10.319	0.0	1.706	—	—	8.85
14	3.00	0.179	—	—	5.04	0.0	7.58

TRY IT IN THE APP **2 · Inverse Response · Step the Duty by Hand**

Load the scenario and press Run. Watch the output voltage trend after the duty step at 0.2 s. The output falls first, bottoms out 7.5 ms after the step and crosses back through 24 V at 15.8 ms before it climbs toward its new steady state of 29.67 V. No controller is acting.

The stored-energy accounting makes the mechanism explicit. At the instant shown in Figure 9.3 the capacitor holds 865 mJ and the inductor 548 mJ, against 576 mJ and 172 mJ at the start. During the first 7.5 ms the capacitor energy is falling and the inductor energy is rising faster, so (9.4) is positive throughout an interval in which the measured output is negative-going. The reported zero location is $+71\,\mathrm{rad/s}$ at that instant rather than the $+135\,\mathrm{rad/s}$ reported at rest, because (9.9) moves with the operating point and the converter is delivering more current.

9.6.2 Undershoot Against Recovery Time

Scenario 3 closes the cascade around the same power stage and steps the setpoint from 24 to 30 V. Raising the outer gain does not remove the dip, it deepens it, and the exchange holds over a factor of seven in gain. At $K_{c,v}$ of 0.06, 0.12, 0.24 and 0.45 A/V the initial undershoot below 24 V is 1.80, 2.33, 3.50 and 4.62 V and the time to recover to within 0.5 V of the new setpoint is 225, 196, 146 and 114 ms. The fastest tuning produces an undershoot of 4.62 V on a 6 V setpoint change. The trade is a property of the plant rather than of the tuning, because a faster output response requires a larger immediate duty excursion, and (9.2) falls the instant the duty rises regardless of which term raised it. At every gain in the sweep the commanded duty reaches d_{max} in the first controller cycle, because the jump of both feedforward terms with the setpoint adds to the proportional kick, and the gain decides how long it stays there, from 2.7 ms at 0.06 A/V to 12.1 ms at 0.45 A/V.

Scenario 12 extends the same sweep past the useful range. With $K_{c,v} = 0.5$ and $\tau_{I,v} = 0.02\,\mathrm{s}$ on a 4 V setpoint change the undershoot reaches 4.62 V, which is larger than the step itself, the output overshoots to 35.16 V for a peak excursion of 179 % of the step, and the inductor current is driven to 9.50 A against a 9 A setpoint clamp. The undershoot is the same 4.62 V that scenario 3 produces at $K_{c,v} = 0.45\,\mathrm{A/V}$ on a step half again as large. In both runs the commanded duty sits at d_{max} from the first controller cycle until about 12 ms after the step, past the bottom of the dip, so the excursion is set by the limit rather than by the gain. At $K_{c,v} = 0.24\,\mathrm{A/V}$ in scenario 3 the duty leaves the limit after 6.9 ms. Raising the gain to 0.75 A/V produces a sustained cycle between 20.7 and 37.0 V that persists 4 s after the step. That behavior is the practical consequence of the bound in Figure 9.2. A controller

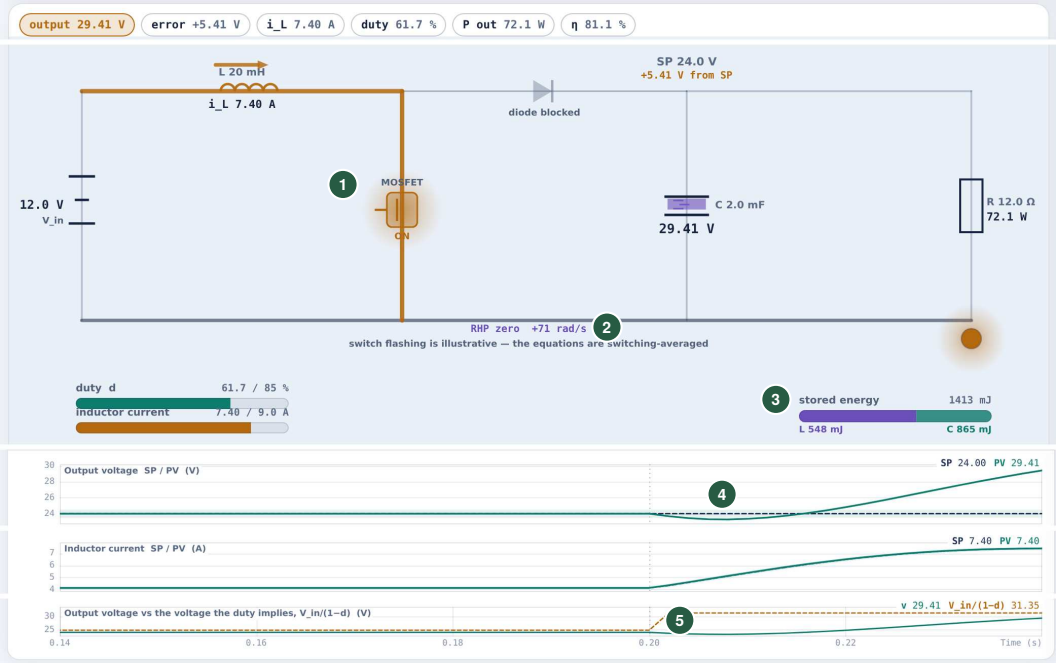


Figure 9.3 Inverse response to a duty step of ten percentage points applied in MANUAL at $t = 0.2$ s (scenario 2), shown at $t = 0.24$ s on a 0.1 s window. The picture is a composite of the application's window: the header values, the circuit view with a blank band of its canvas removed, and three of the five trend strips scaled to the circuit's width. The Control and Learn panels and the scorecard are cut. Callouts: (1) the transistor, on for the larger share of each period, (2) the right-half-plane zero at the operating point of that instant, (3) the stored energy split between L and C, (4) the output dipping below its starting value before it climbs, (5) the requested steady-state voltage $V_{in}/(1-d)$ stepping up while the output still falls. In the logged run the output falls to 23.28 V and is back through 24 V 16 ms after the step while the inductor current builds from 4.14 to 7.40 A, and the requested steady-state voltage steps to 31.35 V while the output is still falling. At the instant shown the output is 29.41 V and total stored energy has risen from 748 mJ at the start to the 1413 mJ the energy bar reads.

that inverts the plant to cross above roughly $z/2$ must generate an ever larger excursion in the wrong direction first, and the appropriate design response is to change L, C or the load rather than the gain.

The feedforward makes the same trade, as scenario 4 shows. A load step from 12 to 6 Ω with feedforward on produces a 7.82 V dip and recovery in 43 ms, and the same step with feedforward off produces a smaller dip of 5.69 V and a recovery of 804 ms, nearly nineteen times longer. The current feedforward is computed from the measured load current, so it jumps the instant the load changes and drives the duty hard immediately, which is what makes the recovery fast and what makes the dip deep. Feedforward on a nonminimum-phase plant trades undershoot for recovery speed exactly as a higher gain does. Whether 7.82 V for 43 ms is preferable to 5.69 V for 804 ms depends on what the downstream circuit cares about, and that is a specification decision rather than a tuning one.

9.6.3 Effect of Capacitance and Inductance Changes

Scenarios 8 and 9 change one component each and separate two effects that a single step response mixes together. Both apply a 24 to 28 V setpoint step with the default tuning, and the reference case at $L = 20$ mH and $C = 2$ mF gives an undershoot of 1.62 V and recovery in 127 ms.

Quartering the capacitance to 0.5 mF leaves the zero exactly where it was, at $+135\,\mathrm{rad/s}$ at 24 V and $+95\,\mathrm{rad/s}$ at the 28 V endpoint in both cases, because C does not appear in (9.9). The undershoot nevertheless grows from 1.62 to 4.27 V (3.97 V on the 2 ms log), because the same missing charge is being taken out of a quarter of the capacitance. What the smaller capacitor does move is the pole pair, from $\omega_n = 77.7$ to 155.4 rad/s, which puts the plant on the unfavorable side of the crossing in Figure 9.2, with the resonance above the bandwidth the zero permits.

> **TRY IT IN THE APP** **8 · Smaller Output Capacitor (0.5 mF)**
>
> Run the scenario and watch the RHP zero tile and the output voltage. The tile reads +135 rad/s at 24 V, as it does with the 2 mF capacitor, yet the first dip after the step to 28 V is 4.27 V, against 1.62 V with the default capacitor. The 2 ms log reads that dip as 3.97 V.

Quartering the inductance to 5 mH moves the zero by the factor of four that (9.9) predicts, from 134.8 to $+539\,\mathrm{rad/s}$ at 24 V, and the undershoot collapses from 1.62 to 0.50 V, bottoming out 1.2 ms after the step and over within 2.5 ms, too brief for the 2 ms log, which reads 0.21 V. A small inductor shortens the current rise time and the interval during which the capacitor is undersupplied. The cost appears outside this model, because peak-to-average current ripple in real hardware goes as $1/L$, so a small inductor buys loop bandwidth with switch rating, capacitor size, conduction loss and radiated emissions. The zero is a design variable that trades against all of those rather than a fixed property of the topology.

9.6.4 Single-Loop and Cascade Comparison

Scenarios 10 and 11 apply an identical 24 to 30 V setpoint step to the same power stage, differing only in whether an inner current loop exists. The single voltage controller is not the slower of the two. It recovers to within 0.5 V in 135 ms against the cascade's 196 ms, and its initial dip is the smaller of the two at 0.78 V against 2.33 V, for the same reason it is quicker to the band and for the same reason it then overshoots to 31.83 V against the cascade's 30.89 V.

The advantages of the cascade lie elsewhere. The first is that the inductor current becomes a regulated and limited variable. In scenario 7 a 12 to 4 Ω overload demands 13.5 A at steady state from a 12 V supply, the outer loop's request is clamped at 9.00 A for 99.5 % of the overload interval, the converter becomes a 9 A current source and the output settles at 19.99 V with the voltage error simply persisting. The peak current reached during that event is 9.47 A with the default inner gain and 11.32 A with $K_{c,i}$ reduced to 5 %/A, so a current limit is only as good as the loop enforcing it. The second is how far the outer loop can be pushed. At $K_{c,v} = 0.45\,\mathrm{A/V}$ the peak excursion on the same setpoint step is 58.6 % of the step with $K_{c,i} = 50\,\%/\mathrm{A}$ and 161.9 % with $K_{c,i} = 5\,\%/\mathrm{A}$, and the peak current is 9.51 against 11.09 A. The third is that load-current feedforward becomes possible at all, because there is a current loop to send it to.

Pause. Try Exercise 9.3 before reading on.

Adequate loop separation has an operational definition in this application. At $K_{c,v} = 0.12\,\mathrm{A/V}$ the recovery barely depends on the inner gain, 196 against 202 ms for inner gains of 50 and 5 %/A, although the slower inner loop moves the duty more gently and dips less (1.44 against 2.33 V). For inner gains from 15 to 200 %/A the recovery stays between 193 and 197 ms and the overshoot between 14.0 and 15.1 %, which means the inner loop has left the outer loop's dynamics. Below roughly 15 %/A the outer loop inherits the inner loop's lag, and the difference between 58.6 and 161.9 % at $K_{c,v} = 0.45$ measures the cost of that inheritance.

Raising the single-loop gain instead of changing the structure produces the failure mode this chapter quantifies. At $K_{c,dv} = 1.2\,\%/\mathrm{V}$ the voltage-only architecture rings, peaking at 37.88 V on a 30 V setpoint and drawing 10.16 A through an unregulated inductor, and it is not inside the band for good until 1.67 s after the step. At 1.5 %/V it never settles. The output cycles between 16.4 and 58.5 V with inductor current peaks of 18.25 A, past the 12 A trip. That gain is past the point where

the right-half-plane zero has consumed the available phase, and no value of $\tau_{I,dv}$ between 0.01 and 3 s recovers a settled response.

9.6.5 Saturation and Anti-Windup

Two scenarios test the actuator boundaries. In scenario 6 the supply collapses to 3.5 V into a 24 Ω load, the ideal ceiling is $V_{in}/(1-d_{max}) = 23.3$ V and the inductor drop pulls the achievable output down to 19.69 V, more than 4 V short of setpoint. The duty is pinned at d_{max} for the whole interval from the event to $t = 2.0$ s and the current setpoint is at its clamp for 92.7 % of it. No tuning fixes that, because no duty in the allowed range holds 24 V from a 3.5 V supply. Back-calculation keeps both integrals bounded while they are clipped [17], but it does not make the release free. The outer integral settles at $i_{lim} - i_{ff} = 1.96$ A, and when the nominal supply and load return at $t = 2.0$ s that stored current carries the output to 29.07 V, 5.07 V above setpoint, and keeps it outside the band for 326 ms. In a counterfactual run that zeroes both integrals at the same instant the peak is 24.11 V, and with the back-calculation term removed it is 102 V.

In scenario 7 the overload is released at $t = 2.4$ s after two seconds against the current clamp. The current feedforward is itself at the 9 A limit during the overload, so the outer integral is held at zero and the current setpoint leaves the clamp within one controller cycle of the release. The inductor current still flowing has nowhere to go but the capacitor, and the output spikes to 33.09 V 6.1 ms after the release. The default integrator then over-corrects into an undershoot to 23.12 V at 86 ms, and the output is inside the ±0.5 V band for good at 163 ms. With $\tau_{I,v} = 3$ s it is inside at 50 ms, because the slow integrator does not chase the spike. With the back-calculation term removed the outer integral would reach 19.5 A during the overload, the current setpoint would stay on its clamp for 614 ms after the release, and the output would not be inside the band for good until 1.01 s. A recursion that absorbed the whole clipped excess at every scan, instead of 1/250 of it, would settle the default case at 229 ms rather than 163 ms, so the tracking time is part of the result.

The graded challenge, scenario 14, applies a 12 to 7 Ω load step at $t = 1$ s to the cascade at its default tuning. It asks for the output to be inside ±0.5 V of 24 V for good within 30 ms, with no more than 0.3 V of overshoot above 24 V and a peak inductor current below 8 A, judged on the logged run at $t = 3$ s, on the case as given, so a changed model parameter or disturbance voids the attempt. The default tuning rides the step out with a 5.78 V dip and a 7.58 A peak, but it is not inside the band for good until 34 ms after the step and then coasts 0.43 V past 24 V, because an integral time of 0.05 s winds in the error from the dip and spends it on the other side. The suggested tuning of $K_{c,v} = 0.25$ A/V and $\tau_{I,v} = 0.15$ s passes with 24 ms, 0.21 V and 7.44 A, which leaves 6 ms, 0.09 V and 0.56 A to spare. The load-current feedforward already supplies the new 7.30 A, so the integrator has little to correct, and lengthening its time to 0.15 s at the default gain cuts the overshoot from 0.43 to 0.16 V while slowing the recovery to 38 ms. The higher gain supplies the speed. The extra speed costs 1.1 V of additional dip (6.88 against 5.78 V), which is the right-half-plane zero at work. At $K_{c,v} = 0.35$ A/V with the same integral time the recovery takes 16 ms but the current peaks at 8.22 A, and at 0.40 A/V the overshoot is 1.27 V, so the passing region is bounded on both sides. The challenge assigns the two terms separate jobs, with the proportional gain buying recovery speed that is paid for in dip depth and current and the integral time setting the overshoot.

9.6.6 Explicit MPC Against the Cascade

Table 9.6 runs scenarios 3, 4, 5 and 7 and the challenge under every architecture on one engine, with the same scenario script, plant, events and duration in every row and no noise or random draw anywhere. The cascade and the single loop run at their shipped tuning (the cascade also at the challenge's suggested tuning), and the MPC runs at scenario 13's settings with nothing retuned per case. The fourth row is the cascade retuned once by the MPC's own rule, the LQR's gains mapped onto the two loops ($K_{c,i} = 127$ %/A from the current gain, $K_{c,v} = 0.08$ A/V from the ratio of the voltage gain to the current gain and $\tau_{I,v} = 0.10$ s from the target filter), with the shipped feedforward, clamp and anti-windup. The shipped feedforward already carries the measured supply voltage and

Table 9.6 Explicit MPC Against the Shipped Architectures and the Cascade Retuned With the LQR's Gains

Structure	dip (V)	t_{band} (ms)	OS (V)	i_{pk} (A)	$t_{>9\,A}$ (ms)	IAE (V s)
Setpoint 24 to 30 V (scenario 3)						
Cascade	2.33	196	0.89	7.044	0	0.305
Voltage only	0.78	135	1.83	7.586	0	0.456
MPC	3.74	29	0.00	7.525	0	0.128
Cascade, LQR gains	2.00	56	0.42	6.812	0	0.346
Load 12 to 6 Ω (scenario 4)						
Cascade	7.82	43	0.49	9.003	5.1	0.263
Voltage only	6.95	1948	0.15	8.614	0	1.372
MPC	10.58	21	0.00	9.019	4.8	0.116
Cascade, LQR gains	7.39	54	0.24	8.808	0	0.303
Supply 12 to 8 V (scenario 5)						
Cascade	4.44	145	0.51	6.840	0	0.208
Voltage only	4.24	449	0.34	6.719	0	0.858
MPC	5.81	30	0.00	7.064	0	0.087
Cascade, LQR gains	4.15	62	0.23	6.670	0	0.237
Overload 12 to 4 Ω (scenario 7)[a]						
Cascade	12.87	—	—	9.475	145.5	8.113
Voltage only	11.10	—	—	12.281[b]	1976.9	2.710
MPC	14.26	—	—	9.005	0.7	8.245
Cascade, LQR gains	12.31	—	—	9.184	31.0	8.115
Challenge, load 12 to 7 Ω						
Cascade	5.78	35	0.43	7.578	0	0.179
Cascade, suggested	6.88	24	0.21	7.441	0	0.156
Voltage only	5.34	322	0.43	7.523	0	0.995
MPC	8.14	18	0.00	8.088	0	0.074
Cascade, LQR gains	5.46	45	0.19	7.424	0	0.205

Dip below 24 V, the time t_{band} after the event at which the output last leaves ±0.5 V of setpoint, the overshoot OS above setpoint, the peak inductor current i_{pk} and the time $t_{>9\,A}$ above the limit come from the state at every 0.1 ms step, so the challenge's 35 and 24 ms read 34 and 24 ms on the logged trajectory. IAE is the scorecard's. [a]No row holds 24 V with 9 A, and the window ends at 2.4 s. [b]Past the 12 A trip.

load current that the MPC's target reads, so both structures receive the same disturbance information. The online solve reproduces every MPC row to the printed digits, and no case produces a miss.

The MPC is inside the band for good 29, 21, 30 and 18 ms after the step, the load step, the sag and the challenge's load step, against 196, 43, 145 and 35 ms for the cascade, which makes it 6.8, 2.0, 4.8 and 1.9 times sooner (1.3 times against the suggested tuning's 24 ms), with no overshoot and with the IAE cut from 0.305 to 0.128 V s on the step. It dips deeper in every case, 3.74 against 2.33 V on the step, 10.58 against 7.82 V on the load step, 5.81 against 4.44 V on the sag, 14.26 against 12.87 V in the overload and 8.14 against 5.78 V on the challenge, and its peak current is higher on every step, 7.525 against 7.044 A on scenario 3. On the challenge the MPC meets the time and overshoot clauses with 18 of 30 ms and 0.00 V, but its 8.088 A peak breaks the 8 A clause, so the MPC as configured fails the graded challenge that the suggested cascade passes. Figure 9.4 draws the dips and the times of Table 9.6 side by side.

A counterfactual run locates the speed. The MPC's own law with its current rows removed, with the same estimator, target and gains and the duty clipped to its range, reproduces the MPC's rows on scenarios 3 and 5 and the challenge to the last digit. The faster recovery is therefore the LQR's state feedback on the estimated current, voltage and applied duty, not the look-ahead and not the constraint. The explicit MPC of Volume 7, Chapter 6 owes its advantage over a PID to state feedback in the same way. Mapping the same gains onto the two proportional-integral loops recovers only part of it, 56 ms against 29 ms on the step, with a smaller dip of 2.00 V. The current rows matter near the limit. Without them the load step peaks at 9.696 A and the overload at 10.461 A with 1989.2 ms above 9 A, while with them the MPC holds the overload to 9.005 A and 0.7 ms above the limit. The cascade, which clamps only the setpoint of its inner loop, reaches 9.475 A and spends 145.5 ms above it, and the retuned cascade 9.184 A and 31.0 ms. The rows bound a linear prediction, so the plant's current

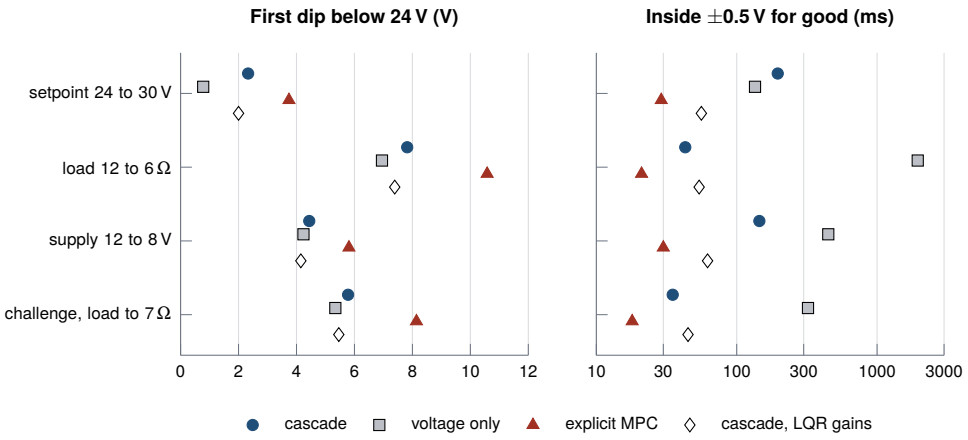


Figure 9.4 Table 9.6 redrawn for the four cases in which the output returns to its setpoint (the overload rows and the challenge's suggested tuning are left out). Each case carries four structures on the same plant: cascade (circles), single voltage loop (squares), explicit MPC (triangles) and the cascade with the LQR's gains (diamonds). The time axis is logarithmic. Notice that the triangles sit farthest left in time and farthest right in dip in every case: the fastest structure always dips deepest, which is the trade the right-half-plane zero sets.

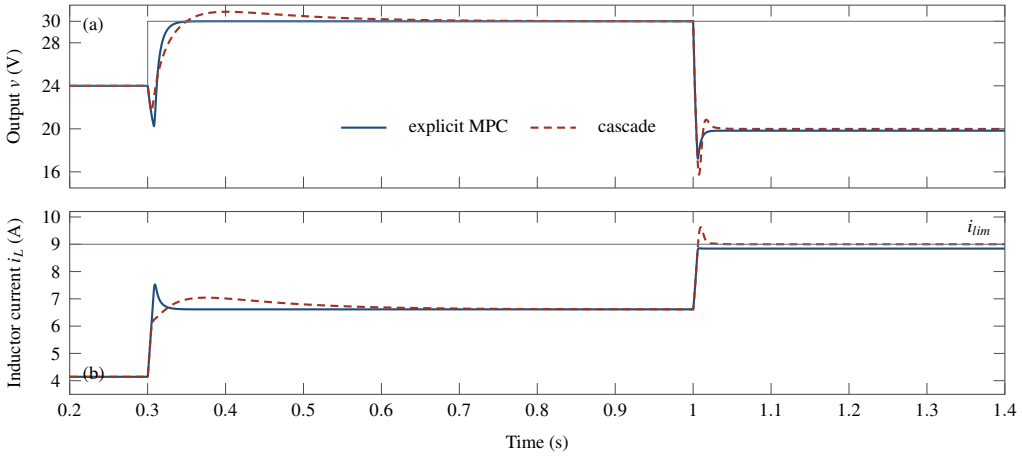


Figure 9.5 Scenario 13 run twice, first with the explicit MPC as loaded and then with the cascade selected before the step, from 0.2 to 1.40 s. MPC solid, cascade densely dashed. (a) Output voltage against the setpoint (thin gray line). (b) Inductor current against the 9 A limit (thin gray line). After the setpoint step at 0.3 s the MPC dips 3.74 V and is inside ±0.5 V in 29 ms, against 2.33 V and 196 ms. After the 4 Ω overload at 1.0 s the cascade's current peaks at 9.623 A and stays above the limit for 169.6 ms, while the MPC's peaks at 8.852 A.

still reaches 9.019 A for 4.8 ms on the load step. The price of keeping v_{SP} as the voltage reference while the limit holds the target is a held output of 19.92 V in scenario 7 against the cascade's 19.99 V.

Scenario 13 puts the step and the overload on one run, as shown in Figure 9.5. The setpoint steps from 24 to 30 V at 0.3 s, and at 1.0 s the load drops to 4 Ω, more than 9 A can supply at 30 V, until it returns at 2.0 s. The MPC's output bottoms out 8.1 ms after the step at 20.26 V and the cascade's 4.7 ms after it at 21.67 V. From 1.0063 s the cascade's current is above 9 A for 169.6 ms, with a peak

Table 9.7 Worst Case of Each Safety-Relevant Output Over 39 Cases

Output	Cascade	Voltage only	MPC
Peak i_L (A)	9.749	19.794	9.222
Time above 9 A (ms)	173.5	2385.8	4.8
Overcurrent trips	0	9	0
Discontinuous conduction	0	3	0
Highest v (V)	37.83	53.99	53.91
Lowest v (V)	7.84	10.11	5.79

Every scenario for 3 s, the challenge at its suggested tuning, every model parameter at both ends of its range on scenario 11's step, the supply and load at their extremes, the disturbance events in combination, the setpoint range and two component changes with an overload, each under the three architectures.

of 9.623 A, and it then holds 9.000 A and 19.99 V. The MPC's current peaks at 8.852 A and at 1.40 s holds 8.842 A and 19.83 V. During the hold the binding rows are the current limit and the duty ceiling on the third move only, the applied duty is 0.4394 against the target's 0.4447, and the plan predicts 8.889 A one period ahead, 47.8 mA above the current then measured. The scorecard of the whole run is row 13 of Table 9.5, whose settling of 1.706 s counts from the step because the overload keeps the output out of the band until the load is released.

The right-half-plane zero survives the change of controller. The prediction contains it, so every plan accounts for the output falling while the current builds, but three moves span 0.6 ms, a small fraction of the 7.5 ms the open-loop dip of scenario 2 takes to reach its bottom, and the terminal cost carries the rest of the response. The plan cannot remove the dip or move the zero. It places the closed loop at a different point of the trade, 3.74 V for 29 ms on the 24 to 30 V step, where the cascade sweep of Section 9.6 gives 3.50 V for 146 ms at $K_{c,v} = 0.24\,\mathrm{A/V}$ and 4.62 V for 114 ms at 0.45 A/V. Feedback on the estimated current and applied duty buys recovery time that the proportional-integral cascade does not reach, and it pays for that time in dip depth exactly as a higher outer gain does.

Table 9.7 gives the worst case of each safety-relevant output over 39 cases. The MPC's peak current is at most 9.222 A, with the 5 mH inductor and an overload, against 9.749 A for the cascade and 19.794 A for the single loop, and the MPC never trips the 12 A flag or enters discontinuous conduction. It is safer on current in 15 of the 39 cases and worse than the cascade on some output in 20. In 12 of those 20 only its peak current is higher, the price of its speed, and it stays under the limit in all of them except the 9.019 A of the load step. Seven of the other eight are model mismatch with a component changed. With $C = 0.5\,\mathrm{mF}$ the output overshoots to 33.96 V against 28.43 V, and with an overload added it falls to 5.79 V against 7.84 V. With $L = 60\,\mathrm{mH}$ it overshoots to 40.74 V against 32.16 V. With $C \le 0.3\,\mathrm{mF}$ or $L \le 3\,\mathrm{mH}$ the MPC does not settle on scenario 11's step, ending 0.660 V short at $C = 0.3\,\mathrm{mF}$ and 0.529 V short at $L = 3\,\mathrm{mH}$ after 3 s, where the cascade ends on setpoint. The one loss at the declared components is a supply sagged three times to 3 V, where the MPC's output reaches 25.77 V against 24.29 V. The highest output of 53.91 V occurs at $C = 0.2\,\mathrm{mF}$, and the application has no overvoltage limit, so no protection is crossed.

DID YOU KNOW? Nine Trips Without an Inner Loop

From the simulator: over 39 cases the single voltage loop drives the inductor to 19.794 A, trips the 12 A overcurrent flag 9 times and leaves continuous conduction 3 times. The cascade and the MPC never trip (Table 9.7). The inner loop is first of all a guard on the inductor current.

The single linearization at 24 V serves the whole setpoint range. Over steps from 24 to targets between 13 and 30 V the one-period current prediction errs by 0.9 mA or less on average and by at most 125 mA, and steps to 13 and 18 V settle in 28 and 25 ms against 180 and 128 ms for the cascade. Targets of 36 to 48 V are held by the current limit, where the target is not an equilibrium of the model and the mean error grows to 5.6 to 52 mA. No step from 24 V produced a miss. A step from 48 to 13 V leaves the table for 4.2 ms and still reaches 13 V within 0.5 V in 41 ms, where the cascade enters discontinuous conduction and takes 418 ms.

9.6.7 Overshoot Measurement Under Inverse Response

The overshoot in Table 9.5 normalizes the peak against the span from the process value before the setpoint change to the new target, and on this plant the choice of the starting value matters. The first logged sample that carries the new setpoint is taken 2 ms after the step. On a minimum-phase plant the process value has not moved by then. On this plant it has moved the wrong way. At $K_{c,v} = 0.45$ that sample already reads 22.855 V, so a normalization against it would use a span of 7.145 V instead of 6 V and report 49.0 % where the logged trajectory gives 58.3 %, and scenario 12 would read 138.3 % instead of 177.9 %. The application normalizes against the last sample before the change, so the overshoot column agrees with the logged trajectory. A metric that is exact on a minimum-phase plant can understate the overshoot by a fifth or more on a nonminimum-phase one.

9.7 Technology in Practice

Figure 9.6 puts the cascade of Section 9.5 on real hardware: a current sensor and a voltage divider feed a digital controller, whose two loops set the duty of a power transistor through its gate driver.

What Measures It

Two readings carry the cascade, and they are not equally hard to get. The output voltage reaches the controller through a divider and a conversion to digital form, whose dynamics and quantization the model omits, as its assumptions state. The inductor current is harder, because every current-sensing method trades size, bandwidth, linearity, isolation, accuracy and cost against one another [23]. The CT of Figure 9.1 stands for a perfect and instantaneous reading of the averaged current. A plain current transformer cannot pass the direct part of that current, so a real inner loop reads a shunt or a magnetic sensor that responds down to direct current [23], and the noise, offset and bandwidth of that sensor are what Section 9.8 says the model leaves out.

TECHNOLOGY Sensor: Reading the Inductor Current

A cascade's inner loop is only as good as its reading of the inductor current, and three families of sensor serve power converters. A shunt resistor turns current into a voltage by Ohm's law and can follow it from direct current up to megahertz, but it sits in the power path without isolation, adds loss and parasitic inductance, and its reading drifts as its resistance rises with temperature. Hall-effect and other magnetic sensors are non-intrusive but need temperature compensation, and the offset drift of a Hall device produces readings unrelated to the actual current. A current transformer saturates on large currents and, like a Rogowski coil, cannot detect direct current at all [23]. On a trend an offset drift reads as current the inductor does not carry, and the inner loop then holds that false value.

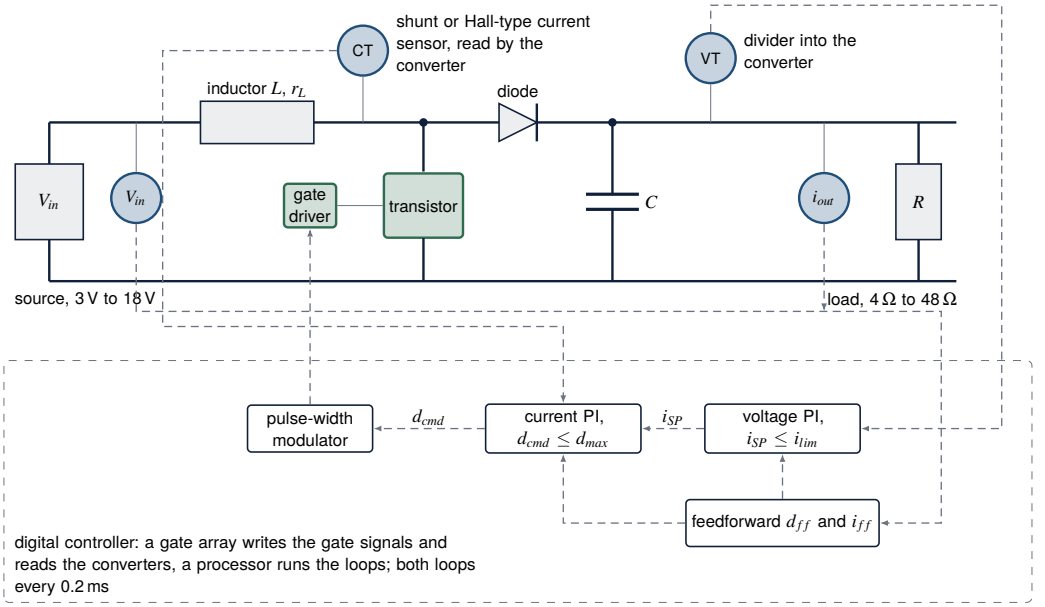


Figure 9.6 The cascade of Figure 9.1 on real hardware. The voltage loop reads the output through a divider and sets the current setpoint, the current loop reads a shunt or Hall-type sensor on the inductor path and sets the duty, and a pulse-width modulator and a gate driver turn that duty into on and off times of the transistor. The source voltage and the load current feed both feedforward terms. Solid lines are the power path and dashed gray arrows are signals. Notice that the inner loop's sensor sits in the power path beside the switch, which is why Section 9.8 calls current sensing the hardest part of building this cascade.

What Moves It

The final element is the power transistor and its gate driver, and the modulator turns the commanded duty into on and off times once per switching period. The model replaces that chain by the first-order lag $\tau_{pwm} = 0.2\,\text{ms}$ of (9.3) and by the clamp at $d_{max} = 0.85$. The clamp to $[0, d_{max}]$ is the only nonlinearity of the actuator that the chapter carries, and scenario 6 shows it binding.

> **TECHNOLOGY** **Final Element: The Power Transistor**
>
> The duty cycle acts through a power transistor, such as a metal-oxide-semiconductor field-effect transistor, that a gate driver turns on and off at the modulator's command. Silicon devices have matured to the point of pushing the fundamental limits of the material. Wide bandgap devices of silicon carbide and gallium nitride operate at higher voltages and temperatures and switch at higher frequencies with greater efficiency, with silicon carbide expected to prevail above 600 V and gallium nitride below it [8]. Faster switching brings its own trouble, because parasitic inductance and resistance in the layout cause losses and overshoots of voltage and current that raise electromagnetic interference [8]. Aging of the switches and the capacitors shows mostly in current transients and high-frequency content, so watching for it takes a fast current sensor [23].

Where the Controller Runs

A converter controller is an embedded system. Typically a field-programmable gate array manages the inputs and outputs, generating the gate signals and reading the analog-to-digital converters, while one or more microcontrollers or digital signal processors compute all or part of the control law [20]. In one arrangement the fast controller runs in the gate array while an outer loop on a processor core supplies its reference at a lower rate [20], which is the cascade of this chapter split across two layers

of hardware. Sampling intervals of 150 µs or less are common in predictive control of converters, a few tens of microseconds on modern platforms [20]. The application's control period of 0.2 ms is one period of a 5 kHz digital controller, deliberately slow so that the inverse response can be watched.

> ⏢ **IN PRACTICE** **Above Half Duty, a Peak-Current Loop Needs a Ramp**
>
> A widely used form of current-mode control compares the peak inductor current with its reference once per clock period, which suits supplies regulated over a wide range of loads at a constant frequency. Such a peak loop is unstable at a duty of 0.5 or greater: a small error grows from one cycle to the next and appears as a cycle-to-cycle oscillation of the inductor current at half the switching frequency [24]. The cure is a compensating slope added to the sensed current, and a very commonly recommended slope is half the falling slope of the inductor current, which guarantees suppression without saying how fast [24]. The default duty of 0.517263 already lies past that boundary.

What Is Modernizing It

Three developments bear on this power stage. The first is the switch itself. Wide bandgap transistors allow higher switching frequencies and therefore smaller inductors and capacitors [8], which by (9.9) also moves the right-half-plane zero away from the origin, as scenario 9 shows with a quarter of the inductance. The second is the sensor. Current sensors embedded in the converter can reduce the usual problems of sensors used for control, and monitoring the condition of switches and capacitors through their current can extend the life of a converter [23]. The third is predictive control. Platforms that combine a gate array with processor cores on one chip, and solvers that keep the computing burden down, are where research on predictive control of converters is heading [20], and the explicit table of Section 9.5.1 is one answer to that burden.

9.8 Limitations

The averaged model is the largest limitation and it is deliberate. Peak-current mode control in hardware becomes subharmonically unstable above 50 % duty unless a compensating ramp is added, and that instability lives inside a switching period and is invisible to (9.1) and (9.2) [15, 16]. The inner loop in this application is an average-current loop sampled at 5 kHz, which does not have that failure mode, so the application teaches why an inner current loop is the right architecture and not how a peak-current comparator is stabilized.

The continuous-conduction assumption is enforced rather than modeled. When the averaged inductor current reaches zero the state is clamped, a fault is raised and the run is declared out of range, so the application says nothing about the discontinuous mode that a real converter enters at light load. With the shipped disturbance ranges the equilibrium current never reaches zero, and the lightest available load of 48 Ω at 24 V still draws 1.01 A through the inductor. A full load shed from 12 to 48 Ω in AUTO does not trigger the flag either, because the inductor current bottoms out at 0.15 A.

Instantaneous efficiency is reported as the ratio of output to input power and therefore exceeds 100 % for a few milliseconds whenever stored energy is leaving the converter. That behavior is arithmetic rather than a fault, but the number invites misreading. The reported efficiency is in any case an upper bound, because switching loss, core loss and diode conduction loss are all absent and those three dominate the loss budget of a real converter at 50 W.

The measurement chain is perfect. Real current sensing is the hardest part of building the architecture this application recommends, because a shunt or a current transformer contributes noise, offset and bandwidth limits that interact directly with the inner loop gain. Nothing here penalizes a large $K_{c,i}$, so the application cannot show why real designs stop increasing it.

The explicit MPC's model leaves out the bilinear terms of (9.1) and (9.2) away from 24 V, any component a learner changes, discontinuous conduction and the trip. Its prediction errs by up to 125 mA per period at a 13 V target, which is why the current floor carries a margin, and with the capacitance or inductance a quarter of its declared value or less the plan stops settling. Its current

rows bound the predicted current, so the plant reaches 9.019 A on the load step. Because those rows are hard, a miss while an overload holds the current at the limit leaves the cascade in charge until the load is released, and a softened current row, which would change the program and its partition, is not implemented. The MPC loses the dip on every case and the challenge's current clause, and its three moves cover 0.6 ms of a response that lasts tens of milliseconds, so the look-ahead decides only near the limits.

> ⚠ **OPEN CHALLENGE** **Optimizing Within a Few Microseconds**
>
> Converter controllers act within a few tens of microseconds, so solving the optimization in real time is the main obstacle to predictive control of power electronics, and it calls for powerful platforms and efficient solvers [20]. Many algorithms also assume a linear model while many converters are not linear, so the model is linearized and loses accuracy [20], as the single linearization of this chapter does away from 24 V. Progress would be nonlinear formulations and solvers that fit inside the sampling interval. Cutting the computational burden is itself a main industrial research focus [20].

9.9 Discussion

Boost stages sit between a photovoltaic string and an inverter, between a fuel cell and a traction inverter, and between a battery pack and a motor drive, and in each of those the conversion efficiency is what decides how much of a scarce resource reaches the load [1, 2]. The design choices that set efficiency are also the choices that set the zero. A larger inductor lowers current ripple and conduction loss and moves the zero toward the origin, which (9.9) states directly and Figure 9.2 quantifies over the load range. A converter sized for the lightest ripple is a converter whose voltage loop must be the slowest, and a converter sized for loop bandwidth pays in conduction loss, switch area and radiated emissions.

The load dependence is the part that has operational consequences. Because z falls as $q_0^2 R$ falls, the zero migrates toward the origin exactly when the converter is delivering the most power and the load transients are largest. A loop tuned at light load for a crossover near $z/2$ can lie above the bound at full load, and Figure 9.2 puts numbers on the range, with the bound falling from 292 rad/s at 48 Ω to 17.2 rad/s at 4 Ω. Current-mode control provides more than a current limit and a feedforward path. It gives the outer loop a plant, (9.11), whose bandwidth does not collapse with load in the same way.

> ⚠ **OPEN CHALLENGE** **No Current Sensor Does Everything**
>
> As switches get faster, the current sensor becomes the weak link of the inner loop. Almost no current-sensing method provides the necessary size, bandwidth, linearity, isolation, accuracy and cost in one place, and almost none of the inexpensive options suits the control of high-frequency wide bandgap converters [23]. Progress would be sensors embedded in the converter, such as a Rogowski coil built into the circuit board or magnetoresistive and fluxgate elements, which can reduce or remove the usual problems of sensors used for control [23], fast enough for the inner loop and small enough to sit beside the switch.

9.10 Conclusion

This chapter studied a boost converter simulator with three states, a switching-averaged continuous-conduction model, two feedforward terms computed from measurements the hardware really has, a voltage-to-current cascade, a single-loop alternative and an explicit MPC on the same power stage, and fourteen scenarios. The control-to-output transfer function carries a right-half-plane zero at $(qv - r_L i_L)/(L i_L)$, which is +134.8 rad/s at the default operating point and sweeps by a factor of 17 across the allowed load range while the *LC* natural frequency stays near 77 rad/s. That zero fixes a trade the controller cannot alter, measured here as an undershoot growing from 1.80 to 4.62 V as the recovery time falls from 225 to 114 ms over a factor of seven in outer gain, and reproduced by the feedforward, which buys a 43 ms load recovery in place of 804 ms at the cost of a 7.82 V dip in place of 5.69 V. The inner current loop does not shrink the first dip, and this chapter measured the three

things it does instead, a current limit that holds the peak to 9.47 A instead of 11.32 A, an outer loop that can be driven to 0.45 A/V with 58.6 % rather than 161.9 % excursion, and a place to send load feedforward. The explicit MPC, a 109-region table of a three-move program with the inductor current as a hard constraint, reaches the band 1.9 to 6.8 times sooner than the cascade and holds the overload to 9.005 A against 9.475 A, while the same zero deepens its dip on every case.

Two of the scenarios are component changes rather than tuning changes, and they separate the capacitor, which deepens the dip from 1.62 to 4.27 V without moving the zero, from the inductor, which moves the zero from 135 to 539 rad/s and collapses the dip to 0.50 V. The model omits diode drop, switching loss, capacitor equivalent series resistance, current-sense dynamics and everything that happens inside a switching period, so the efficiency figures are upper bounds and the peak-current-mode instability that motivates slope compensation cannot appear. Extending the simulator with a constant-power load, a sensed-current filter and a discontinuous-conduction branch would each add a distinct failure mode to the same power stage. Within those limits the simulator turns the cost of a right-half-plane zero and the return on rewiring a loop into measured numbers.

KEY POINTS

- Raising the duty first starves the output and only later feeds it, so the output dips before it climbs.
- The zero belongs to the power stage: no gain, integral time or controller structure moves it, while the inductor and the load do.
- A faster recovery costs a deeper first dip whether the speed comes from a higher gain, from feedforward or from an MPC.
- The inner current loop does not shrink the first dip. It limits the current, tolerates a harder outer gain and gives load feedforward a place to enter.

Where else this applies. Volume 2, Chapter 12 closes an armature-current loop inside a speed loop, and the drum level of Volume 18, Chapter 4 carries a right-half-plane zero that moves with load.

References

[1] F. Blaabjerg, Z. Chen, and S. B. Kjaer, “Power electronics as efficient interface in dispersed power generation systems,” *IEEE Transactions on Power Electronics*, vol. 19, no. 5, pp. 1184–1194, 2004.

[2] A. Emadi, Y. J. Lee, and K. Rajashekara, “Power electronics and motor drives in electric, hybrid electric, and plug-in hybrid electric vehicles,” *IEEE Transactions on Industrial Electronics*, vol. 55, no. 6, pp. 2237–2245, 2008.

[3] J. Freudenberg and D. Looze, “Right half plane poles and zeros and design tradeoffs in feedback systems,” *IEEE Transactions on Automatic Control*, vol. 30, no. 6, pp. 555–565, 1985.

[4] J. B. Hoagg and D. S. Bernstein, “Nonminimum-phase zeros - much to do about nothing - classical control - revisited part II,” *IEEE Control Systems*, vol. 27, no. 3, pp. 45–57, 2007.

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

[6] 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.

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

[8] U.S. Department of Energy, “Wide bandgap semiconductors for power electronics,” U.S. Department of Energy, Quadrennial Technology Review 2015, Chapter 6, Technology Assessment 6N, 2015. [Online]. Available: https://www.energy.gov/sites/prod/files/2016/02/f29/QTR2015-6N-Wide-Bandgap-Semiconductors-for-Power-Electronics.pdf

[9] S. R. Sanders, J. M. Noworolski, X. Z. Liu, and G. C. Verghese, “Generalized averaging method for power conversion circuits,” *IEEE Transactions on Power Electronics*, vol. 6, no. 2, pp. 251–259, 1991.

[10] R. D. Middlebrook and S. Cuk, “A general unified approach to modelling switching-converter power stages,” in *1976 IEEE Power Electronics Specialists Conference*, 1976, pp. 18–34.

[11] G. W. Wester and R. D. Middlebrook, “Low-frequency characterization of switched dc-dc converters,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-9, no. 3, pp. 376–385, 1973.

[12] R. D. Middlebrook, "Small-signal modeling of pulse-width modulated switched-mode power converters," *Proceedings of the IEEE*, vol. 76, no. 4, pp. 343–354, 1988.

[13] V. Vorperian, "Simplified analysis of PWM converters using model of PWM switch. continuous conduction mode," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 26, no. 3, pp. 490–496, 1990.

[14] R. W. Erickson and D. Maksimović, *Fundamentals of Power Electronics*. Cham, Switzerland: Springer International Publishing, 2020.

[15] R. D. Middlebrook, "Modeling current-programmed buck and boost regulators," *IEEE Transactions on Power Electronics*, vol. 4, no. 1, pp. 36–52, 1989.

[16] R. B. Ridley, "A new, continuous-time model for current-mode control (power convertors)," *IEEE Transactions on Power Electronics*, vol. 6, no. 2, pp. 271–280, 1991.

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

[18] A. Bemporad, M. Morari, V. Dua, and E. N. Pistikopoulos, "The explicit linear quadratic regulator for constrained systems," *Automatica*, vol. 38, no. 1, pp. 3–20, 2002.

[19] P. Karamanakos, T. Geyer, and S. Manias, "Direct voltage control of DC–DC boost converters using enumeration-based model predictive control," *IEEE Transactions on Power Electronics*, vol. 29, no. 2, pp. 968–978, 2014.

[20] P. Karamanakos, E. Liegmann, T. Geyer, and R. Kennel, "Model predictive control of power electronic systems: Methods, results, and challenges," *IEEE Open Journal of Industry Applications*, vol. 1, pp. 95–114, 2020.

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

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

[23] A. Parsa Sirat and B. Parkhideh, "Current sensor integration issues with wide-bandgap power converters," *Sensors*, vol. 23, no. 14, p. 6481, 2023.

[24] A. Tiwary and G. A. Rincón-Mora, "Sub-harmonic stability and slope compensation boundary in peak/valley DC–DC converters," *Electronics*, vol. 15, no. 8, p. 1609, 2026.

9.11 Exercises

9.1 ⌕ Concept check · 2 min

In scenario 2 the duty steps from 0.517 to 0.617 in MANUAL, and the output voltage starts to fall although its new steady state is near 29.7 V. What sets the sign of the initial slope?

(a) The modulator lag holds the old duty, so the output follows it at first

(b) The delivered current $q\,i_L$ drops at once, while i_L cannot change in that instant

(c) The inductor drop $r_L i_L$ rises before the output can respond

(d) The load draws more current as soon as the duty rises

9.2 ⌕ Concept check · 1 min

Scenario 8 quarters the output capacitor, from 2 to 0.5 mF, and repeats the 24 to 28 V step. What happens to the right-half-plane zero and to the first dip?

(a) The zero moves four times farther out and the dip shrinks

(b) The zero moves toward the origin and the dip deepens

(c) The zero stays where it was and the dip deepens

(d) The zero stays where it was and the dip is unchanged

9.3 ⌕ Concept check · 2 min

On the 24 to 30 V step of scenarios 10 and 11, the single voltage loop dips 0.78 V and is inside the band in 135 ms, and the cascade dips 2.33 V and takes 196 ms. What does the inner current loop buy?

(a) A smaller first dip, because the inner loop removes the zero from the plant
(b) A faster recovery at the shipped tuning, because the inner loop crosses near 600 rad/s
(c) A controller pole that cancels the zero, so the outer loop can cross above z
(d) A limited inductor current, room for a harder outer gain and a path for load feedforward

9.4 Theory and calculation · 6 min

The default point is $V_{in} = 12\,\text{V}$, $R = 12\,\Omega$, $r_L = 0.1\,\Omega$ and $v = 24\,\text{V}$.

(a) The ideal relation $v = V_{in}/(1-d)$ asks for $d = 0.5$. Put that duty into Equation (9.5) and find how far short of 24 V the output settles. *Checkpoint:* 23.226 V, 0.774 V short.

(b) Use Equation (9.6) to find the duty that holds 24 V, the inductor current and the inductor drop. *Checkpoint:* $q = 0.482737$, $d = 0.517263$, $i_L = 4.1430\,\text{A}$, $r_L i_L = 0.414\,\text{V}$.

(c) Compute d_{ff} and i_{ff} of Equation (9.14) at this point. What does each integral accumulator hold at steady state, and what is the efficiency? *Checkpoint:* $d_{ff} = 0.5000$, $i_{ff} = 4.1430\,\text{A}$. Inner integral 0.017263, outer integral zero, 96.55 %.

9.5 Theory and calculation · 6 min

Use the operating point of the previous exercise and $L = 20\,\text{mH}$, $C = 2\,\text{mF}$.

(a) Evaluate the zero of Equation (9.9) exactly and with the approximation $q_0^2 R/L$. *Checkpoint:* 134.8 rad/s and 139.8 rad/s.

(b) From the denominator of Equation (9.8) find ω_n and ζ, and the ratio z/ω_n. *Checkpoint:* 77.68 rad/s, 0.300, 1.74.

(c) Repeat the exact zero for a $4\,\Omega$ load at the same 24 V, and give the bandwidth bound $z/2$ at both loads. *Checkpoint:* $q_0 = 0.44365$, $i_0 = 13.52\,\text{A}$, $z = 34.4\,\text{rad/s}$. Bounds 67 and 17.2 rad/s.

9.6 Theory and calculation · 6 min

The current feedforward of Equation (9.14) uses the measured load current at the setpoint $v_{SP} = 24\,\text{V}$, with $V_{in} = 12\,\text{V}$, $r_L = 0.1\,\Omega$ and $i_{lim} = 9\,\text{A}$.

(a) The challenge drops the load from 12 to $7\,\Omega$. Compute i_{ff} at the instant of the step, while the output is still 24 V, and the jump from its value at $12\,\Omega$. *Checkpoint:* 7.30 A, a jump of 3.16 A.

(b) Scenario 7 drops the load to $4\,\Omega$. Compute the current the power balance asks for at 24 V, with and without r_L, and what the clamp passes. *Checkpoint:* 13.52 A, 12 A lossless, clamped to 9 A.

(c) With i_L held at 9 A, use Equation (9.11) at steady state to find the output the overload holds. *Checkpoint:* 19.99 V.

9.7 App challenge · 8 min

Load scenario 2, Inverse response, step the duty by hand, and use the default seed. Change nothing, run to $t = 0.24\,\text{s}$ and export the CSV file. The scenario steps the duty from 0.517 to 0.617 at 0.2 s.

(a) Report the lowest `output_voltage_V` after 0.2 s, the row where it occurs and the first later row at or above 24 V. *Checkpoint:* 23.28 V in the 0.208 s row, back at 0.216 s.

(b) Report `output_voltage_V` and `duty_implied_V` at 0.24 s, and `stored_energy_J` at 0.2 and 0.24 s. *Checkpoint:* 29.41 and 31.35 V. 748 and 1413 mJ.

(c) Report `capacitor_energy_J` and `inductor_energy_J` at 0.2 and 0.206 s, and say where the charge withheld from the output has gone. *Checkpoint:* 576 to 543 mJ and 172 to 238 mJ.

9.8 App challenge · 8 min

Load scenario 3, Output setpoint step, 24 to 30 V, and use the default seed. The setpoint steps at 0.3 s.

(a) Change nothing, run to 0.6 s and export the CSV file. Report the lowest `output_voltage_V` after the step as a dip below 24 V, and the Settling, Overshoot and Peak i_L readings of the scorecard. *Checkpoint:* Dip 2.20 V (within 0.05 V), settling 0.194 s, overshoot 14.8 %, peak 7.04 A.

(b) Press Reset, set Voltage gain (outer) to 0.24 A/V before Run, and repeat. Then Reset, set 0.45 A/V and repeat. *Checkpoint:* 3.43 V, 0.144 s, 19.9 %, 7.50 A. 4.61 V, 0.112 s, 58.3 %, 9.51 A.

(c) From 0.12 to 0.45 A/V, how many volts of dip does each millisecond of settling cost? Why can no outer gain avoid the exchange? *Checkpoint:* About 0.03 V per millisecond.

9.9 App challenge · 9 min

Load the challenge, Ride out a load step, and use the default seed. The load falls from 12 to 7 Ω at 1 s, and each run lasts 3 s.

(a) Change nothing and run. Report the three readings of the challenge verdict and name the clauses that fail. *Checkpoint:* 34 ms, 0.43 V, 7.58 A. The time and overshoot clauses fail.

(b) Press Reset, set Voltage gain (outer) to 0.25 A/V and Voltage integral time to 0.15 s before Run, and run. Then Reset, raise the gain to 0.35 A/V and run again. *Checkpoint:* 24 ms, 0.21 V, 7.44 A, met. 16 ms, 0.28 V, 8.22 A, the current clause fails.

(c) Press the reload button, set Control architecture to Explicit MPC before Run, and run. Report the verdict. *Checkpoint:* 16 ms, 0.00 V, 8.09 A, the current clause fails.

Say what the gain and the integral time each did in (b), and why the MPC fails.

9.10 Application area · 4 min

A designer wants the bandwidth bound $z/2$ at the 4 Ω full load to equal the bound at the 12 Ω default point with $L = 20\,\text{mH}$, by changing only the inductor. The operating point at a given load does not depend on L.

(a) What inductance is needed? Use $z = 134.8\,\text{rad/s}$ at 12 Ω and 34.4 rad/s at 4 Ω. *Checkpoint:* About 5.1 mH.

(b) By what factor does the current ripple grow at the same switching frequency, and what pays for it? *Checkpoint:* About 3.9.

9.11 Application area · 3 min

The app reports 96.55 % efficiency at the default point, with $P_{in} = 49.72\,\text{W}$ and $P_{out} = 48\,\text{W}$, and its diode is ideal.

(a) A real diode drops roughly 1 V and carries the output current of 2 A on average. Estimate the efficiency with that loss added. *Checkpoint:* About 92.8 %.

(b) The efficiency badge can read above 100 % for a few milliseconds. When, and is it a fault?

9.12 Review the agent · 5 min

On scenario 4, the 12 to 6 Ω load step, the cascade with feedforward on dips 7.82 V. An engineer asks an AI assistant how to make the dip smaller. It replies: “Load-current feedforward acts before an error appears, so it can only reduce the dip. The 7.82 V must come from the outer loop. Keep feedforward on and halve the outer gain to 0.06 A/V.” Find the flaw, and say how you would check the claim in the app.