Two-Mode Switching PID Control of a Nonlinear SISO System

TargetDeLab Technical Paper 01

TargetDeLab Group

2026

Abstract

This study investigates trajectory tracking for the nonlinear SISO system $\dot{x} = f(x) + u$, with a dedicated numerical study of f(x)=x^n. The proposed architecture combines compensation of the known nonlinearity, feedforward of $\dot{r}(t)$, and PID feedback. The implemented architecture expresses the ON/OFF concept as a clearly defined two-state PID with an error hysteresis band and a minimum dwell time to avoid chattering. A causal reference filter and conditional-integration anti-windup are also added when actuator saturation is active. The theory shows that, under exact compensation of f(x), the error dynamics are second order and independent of the trajectory shape apart from the initial conditions. The reproducible simulations show that the switching controller provides an accuracy-energy trade-off: in the baseline scenario it reduces control energy from 117.65 to 82.06 (about 30.2%) relative to permanently high gains, with IAE 0.02163 versus 0.00746. The high-gain state is used for only about 1.28% of the time. The results also clarify that increasing the order n primarily affects compensation energy rather than the ideal tracking error, while saturation and noise introduce the expected trade-offs.

Keywords: PID, switched control, gain scheduling, nonlinear control, feedback linearization, dwell time, saturation, anti-windup, tracking.

1 Introduction and Positioning of the Study

PID controllers remain a core industrial control technology, but their practical application to nonlinear processes requires careful treatment of saturation, noise, operating-point changes, and actuator constraints. This study does not propose a new universal PID law; it examines a specific hybrid architecture in which two predefined gain sets are switched according to the magnitude of the error.

The idea is related to gain scheduling and switched/hybrid control. The relevant theory warns that stability of the individual modes alone is not sufficient under arbitrary switching. For this reason, the implementation uses hysteresis and a minimum

dwell time. In parallel, cancellation of f(x) is a form of feedback linearization for this particular relative-degree-one system.

A principal contribution of this chapter is to clarify what is actually supported by theory and what is supported by the numerical tests: switching does not by itself guarantee minimum energy, but it provides a tunable trade-off between error and control effort. The energy required to compensate the nonlinearity may dominate the total u(t), and therefore it should be reported separately from the fraction of time for which the high-gain mode is active.

1.1 Nonlinear Tracking and the Role of Feedback Linearization

For a nonlinear state-space system $\dot{x}$=f(x)+g(x)u with output y=h(x), tracking design begins by asking how the input enters the derivatives of the regulated output. The number of differentiations required before u appears explicitly is the relative degree. When the relative degree is well defined and the associated decoupling term does not vanish, a state-feedback transformation can cancel known nonlinear terms and impose a desired input-output relationship. This idea is the basis of input-output feedback linearization [7,11-13]. The scalar plant used in this chapter is an especially transparent relative-degree-one case because the control input enters $\dot{x}$ directly with unit gain.

Feedback linearization is exact only with respect to the model used in the cancellation. If the implemented compensation is g_c(x) while the true drift is f(x), the residual $\Delta f(x) = f(x) - g_c(x)$ remains in the tracking-error dynamics as a disturbance. The resulting closed loop must then be assessed using robustness concepts such as input-to-state stability, bounded disturbance response, or an explicit uncertainty model [5,11]. This distinction is central to the interpretation of the present simulations: exact cancellation explains the nominal error equation, whereas the noise and saturation scenarios illustrate mechanisms by which the ideal reduction is lost.

For the particular system $\dot{x} = f(x) + u$ and control u=-f(x)+$\dot{r}$+v, exact cancellation gives $\dot{e}$=-v. The nonlinear plant has therefore not become globally linear in every physical sense; rather, its nominal input-output tracking-error dynamics have been linearized by state feedback. The actual actuator signal still contains the nonlinear cancellation term -f(x), so control effort can remain strongly dependent on the nonlinear plant even when the ideal error dynamics do not.

1.2 PID Feedback on a Linearized Error Channel

A PID controller combines proportional action, integral action, and derivative action. Proportional feedback reacts to the instantaneous error, integral feedback accumulates persistent bias and is therefore important for offset rejection, and derivative action adds phase lead by responding to the rate of change of the error. In a conventional LTI setting these roles can be studied through loop transfer functions and sensitivity functions; in the present chapter the same gain structure is applied to

the feedback-linearized error channel rather than directly interpreted as an LTI controller around the nonlinear plant [2,11].

The formal differentiation that produces the second-order error equation in the Model and Control Law section is valid on intervals where the selected gain set is fixed and the signals possess the required regularity. Switching introduces piecewise-smooth dynamics, while saturation and conditional integration further modify the equations. Consequently, the characteristic polynomial associated with each fixed mode establishes the stability of that mode under the idealized assumptions, but it is not by itself a proof of stability under arbitrary switching.

1.3 Switched and Hybrid Systems, Hysteresis, and Dwell Time

A switched system consists of a family of continuous-time subsystems together with a rule that selects the active subsystem. Even when every subsystem is asymptotically stable, arbitrary switching can produce an unstable overall trajectory unless additional structure is available. Standard sufficient approaches include a common Lyapunov function, multiple-Lyapunov-function conditions, or restrictions on the frequency of switching [4-6]. Average dwell-time theory formalizes the latter idea by requiring the number of switches over any interval to grow no faster than a prescribed affine bound; sufficiently slow switching can preserve exponential stability for broad classes of switched systems [14].

Hysteresis and dwell time serve different purposes. Hysteresis creates a state-dependent dead band between the turn-on and turn-off thresholds, preventing repeated mode changes caused by small perturbations near a single threshold. Dwell time imposes a temporal constraint after a switch and therefore limits how rapidly the supervisor may switch again. Their combination is a standard practical mechanism for suppressing chattering and avoiding numerically or physically unrealistic high-frequency switching. In this chapter these devices are used as implementable design rules rather than as a substitute for a general switched-system theorem.

1.4 Saturation, Windup, and Noise in Practical PID Loops

Actuator saturation is a dynamic nonlinearity because the commanded input and the applied input differ whenever the command exceeds the admissible range. If an integral controller continues integrating while the actuator is saturated, the integrator state can grow to a value that delays recovery after the command re-enters the admissible region. This windup phenomenon motivates anti-windup architectures [9]. Conditional integration, used here, is among the simplest remedies: integration is suspended only when the integral contribution would drive the saturated actuator further into saturation.

Derivative action and numerical differentiation are also sensitive to high-frequency measurement components. A causal low-pass filter reduces this sensitivity at the

price of additional phase lag. The first-order recursive filter used in Section 1.2.3 is therefore not merely a numerical convenience; it represents the causal filtering that would be required in an online implementation. The resulting trade-off between noise attenuation, tracking bandwidth, and control effort is a recurring theme throughout the book.

2 Model and Control Law

Control-loop topology used in this chapter. The reference r(t) is compared with y(t)=x(t). The selected low/high PID mode generates v(t); the reference feedforward $\dot{r}(t)$ and nonlinear compensation -g(x) are combined to form u(t), which drives the nonlinear plant. The measured state/output is fed back to the controller.

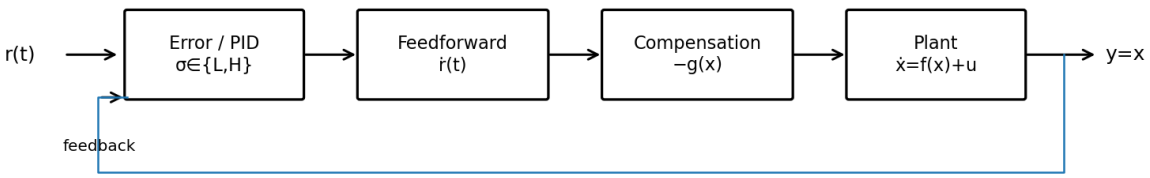


Figure 1. Control topology used in Chapter 1.

Consider the scalar nonlinear system:

$$\dot{x}(t) = f(x(t)) + u(t),\ y(t) = x(t) \tag{1}$$

For the numerical tests, f(x)=x^n, $n \in \mathbb{N}$, is used. The tracking error is e(t)=r(t)-x(t).

The control law is written as:

$$u(t) = -g(x(t)) + \dot{r}(t) + v(t) \tag{2}$$

$$v(t) = K_{p,\sigma}\, e(t) + K_{d,\sigma}\, \dot{e}(t) + K_{i,\sigma} \int_0^t e(\tau)\, d\tau \tag{3}$$

where $\sigma(t)$ in {L,H} is the discrete low- or high-gain state. For exact compensation g(x)=f(x), $\dot{x}=\dot{r}+v$ and therefore $\dot{e}=-v$.

Differentiating the relation $\dot{e}$=-Kp e-Kd $\dot{e}$-Ki integral e dt gives:

$$(1 + K_{d,\sigma})\ddot{e} + K_{p,\sigma}\dot{e} + K_{i,\sigma}e = 0 \tag{4}$$

This relation establishes an important point: the error equation is second order, and cancellation of the nonlinearity follows from g(x)=f(x), not from the definition of u(t) itself. For positive Kp, Ki, and 1+Kd>0, the characteristic polynomial has positive coefficients and its roots lie in the left half-plane.

Independence of the ideal error dynamics from n does not imply independence of the required energy. The term -x^n remains in the actual control signal u and can substantially increase the integral $\int u^2(t)\,dt$.

2.1 Switching Logic with Hysteresis and Dwell Time

The implementation employs a hysteresis band on |e|, which is clearer, reproducible, and less sensitive to noise:

$\sigma(t)$ = H if |e| $\geq e_{\text{ON}}$; $\sigma(t)$ = L if |e| $\leq e_{\text{OFF}}$; otherwise retain the previous state (5)

with e_{ON}=0.05, e_{OFF}=0.025, and minimum dwell time T_d=0.01 s. The high gains are (Kp,Kd,Ki)H=(250,0.1,375), and the low gains are one tenth of these values.

Hysteresis prevents rapid switching near the threshold, while dwell time imposes an explicit temporal constraint on the switching frequency. The numerical study documents the performance of this specific tuning; no universal stability theorem is claimed for arbitrary switching.

2.2 Saturation and Anti-Windup

With actuator limit A, u=sat(u0,-A,A) is applied. Because saturation breaks the exact cancellation of f(x), the ideal linear error equation does not hold during saturation intervals. The implementation freezes integration when the integral term would drive the actuator deeper into saturation (conditional integration), thereby avoiding windup.

2.3 Noise and Causal Filtering

Numerical differentiation of a noisy r(t) amplifies high-frequency content. A causal first-order filter is used:

r_f[k] = r_f[k-1] + α(r_n[k]-r_f[k-1]), $\alpha = \Delta t/(\tau_f + \Delta t)$ (6)

with τ_f=0.002 s. The simulation therefore corresponds to an implementable online filter.

3 Numerical Methodology

Scenario settings used for the results below. The baseline reference is a smooth accelerated command: $r(t) = 5t^2$ from 0 to 0.5 s, then it is held at 1.25 until T=1 s.

Unless a subsection states otherwise, x(0)=1, n=2, Δt=10^{-4} s, the high gains are (250,0.1,375), the low gains are one tenth of these values, e_{ON}=0.05, e_{OFF}=0.025, T_d=0.01 s, and the noise seed is 123. The first comparison uses permanently high, permanently low, and switching gains. The nonlinearity-order subsection then sweeps n=1,...,10 with the other baseline settings fixed; the saturation subsection applies actuator limits $|u| \leq 20, 10, 5$; and the noise subsection applies clean, 40, 30, and 20 dB AWGN cases with the causal reference filter τ_f=0.002 s.

Baseline scenario: T=1 s, Δt=10^{-4} s, x(0)=1, $r(t) = 5t^2$ for $t < 0.5$ s and $r(t) = 1.25$ for $t \geq 0.5$ s, n=2. Euler integration is used for integral calculations. All noisy runs use the fixed seed 123.

Four groups of tests are considered: (i) permanently high, permanently low, and switching gains; (ii) n=1,...,10; (iii) limits |u|<=20, 10, 5; and (iv) a clean signal and SNR=40, 30, 20 dB.

Table 1. Comparison of the three operating modes.

Controller	final_error	IAE	RMSE	energy	high_pct
high	0.0014538	0.0074591	0.046583	117.65	100
low	0.014936	0.061392	0.14179	17.198	0
switch	0.004655	0.02163	0.096206	82.061	1.2799

The switching controller reduces Eu by about 30.2% relative to the permanently high-gain mode (82.06 versus 117.65), while IAE increases from 0.00746 to 0.02163. The high-gain state is active for only 1.28% of the time. The principal documented property is therefore an accuracy-energy trade-off, not simultaneous minimization of both quantities.

Interpretation of the baseline result. The numerical comparison shows that the switching law is functioning primarily as a resource-allocation mechanism: the high-gain controller is invoked only during the small fraction of the trajectory for which the error requires aggressive correction, whereas the low-gain controller carries most of the motion. Consequently, the reduction in control energy is obtained by avoiding continuous high-gain action rather than by changing the compensated plant dynamics. The accompanying increase in IAE is therefore the expected price of spending less time in the aggressive mode. This makes the result most useful as a quantitative accuracy-versus-effort trade-off rather than as evidence that switching dominates both fixed-gain controllers in every metric.

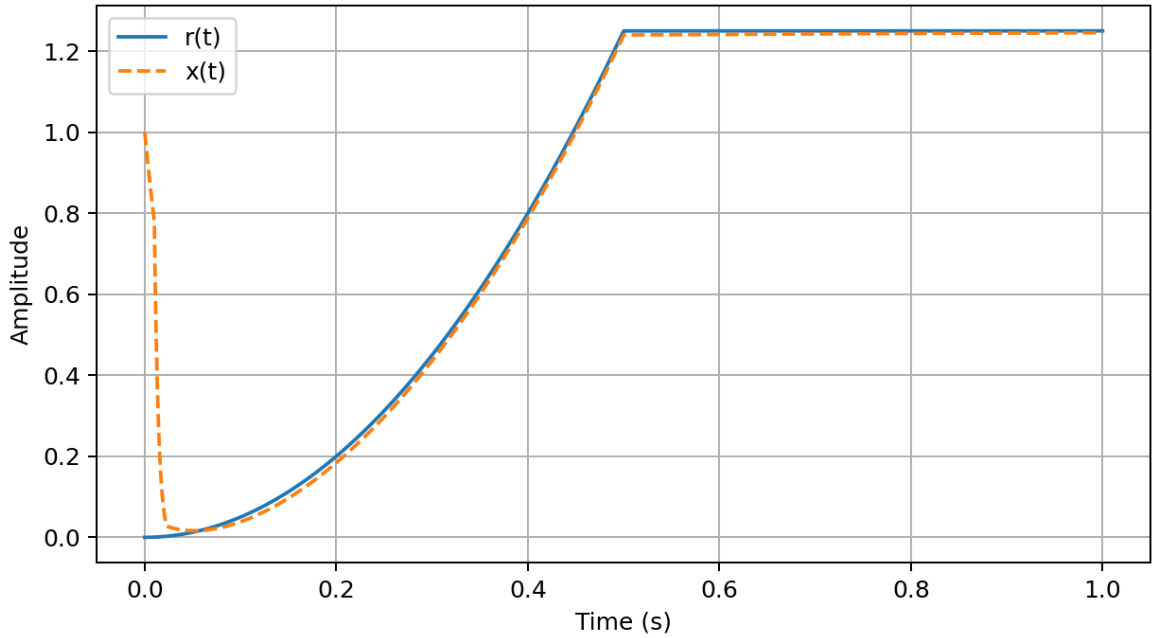


Figure 2. Trajectory tracking with switching gains.

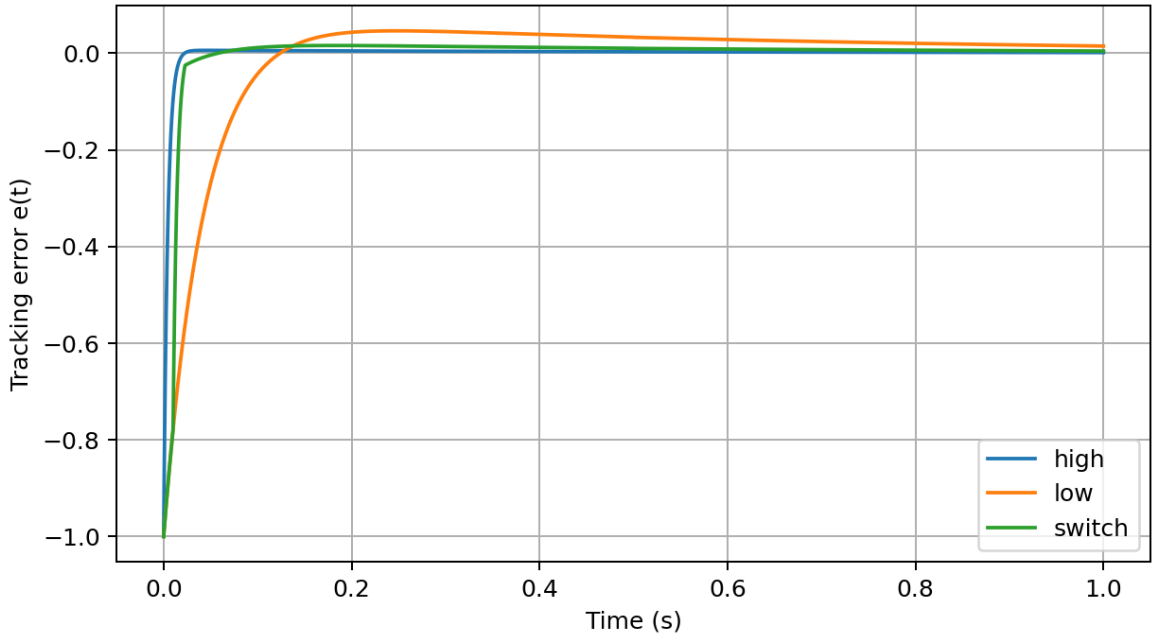


Figure 3. Error comparison for high, low, and switching gains.

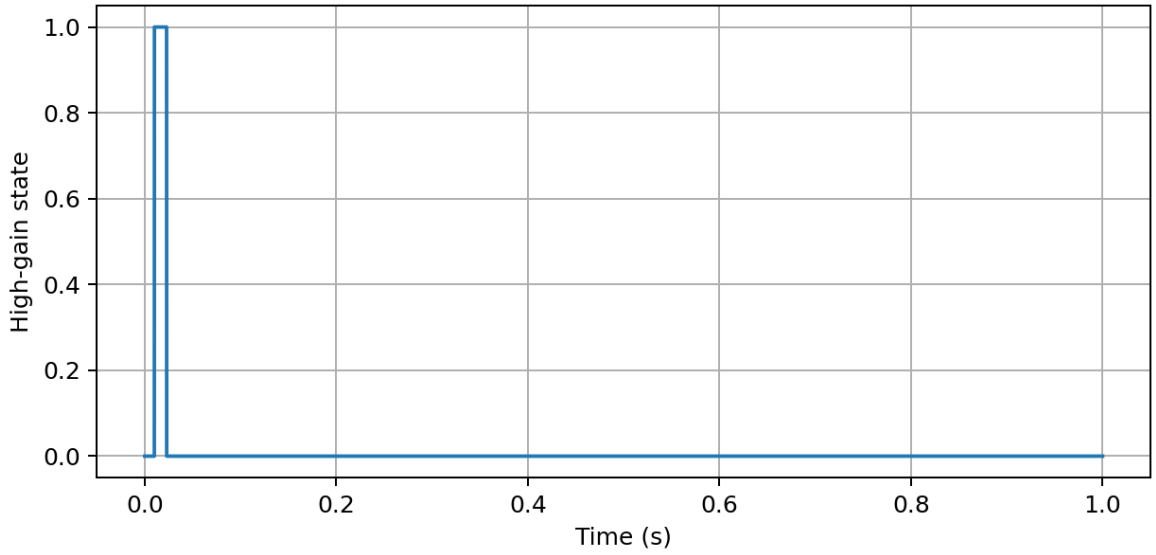


Figure 4. High-gain state of the switching PID controller.

3.1 Effect of Nonlinearity Order

Table 2. Sweep of nonlinearity order n=1,…,10.

n	final_error	IAE	energy	high_pct
1	0.00465499	0.0216301	81.7333	1.27987
2	0.00465499	0.0216301	82.0607	1.27987
3	0.00465499	0.0216301	82.6357	1.27987
4	0.00465499	0.0216301	83.5669	1.27987
5	0.00465499	0.0216301	85.0445	1.27987
6	0.00465499	0.0216301	87.3671	1.27987
7	0.00465499	0.0216301	90.9982	1.27987
8	0.00465499	0.0216301	96.6551	1.27987
9	0.00465499	0.0216301	105.447	1.27987
10	0.00465499	0.0216301	119.088	1.27987

The quantities |e(T)|, IAE, and duty cycle remain practically unchanged, as predicted by exact compensation. In contrast, Eu changes with n because the term -x^n belongs to the actual control input u. This result establishes that the nonlinearity may disappear from the ideal error dynamics but not from the actuator energy cost.

The order sweep also clarifies where the nonlinear term still matters after ideal cancellation. Because the compensated tracking-error equation is independent of x^n, similar error metrics across n are consistent with the theory. The physical actuator, however, must still generate the cancellation term -x^n. The changing Eu values therefore measure a real implementation cost that is hidden if only the ideal error dynamics are examined. In practical terms, two plants can exhibit nearly identical tracking while demanding very different actuator effort.

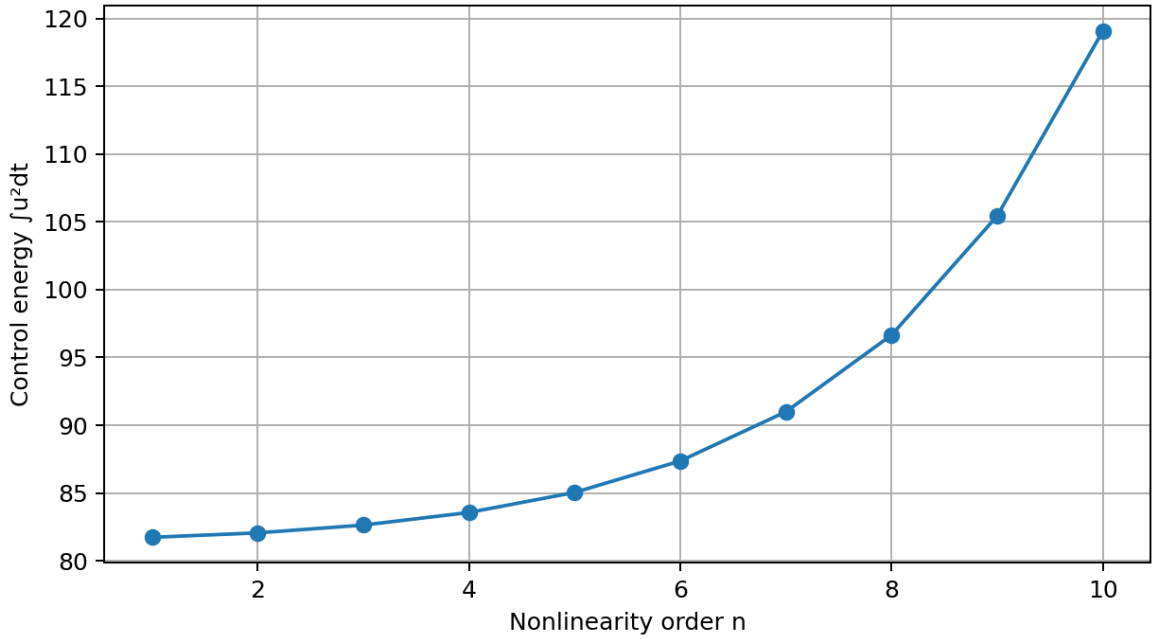


Figure 5. Control energy versus nonlinearity order n.

3.2 Effect of Saturation

Table 3. Effect of the actuator limit.

limit	final_error	IAE	RMSE	energy	high_pct	sat_pct
none	0.00465499	0.0216301	0.0962064	82.0607	1.27987	0
20	0.000513069	0.0271722	0.130744	23.2933	4.13959	4.58954
10	0.000476439	0.0524833	0.186029	13.7072	8.71913	9.42906
5	0.000509386	0.100189	0.262598	8.38485	16.8383	17.7682

More restrictive saturation reduces energy but increases the integral error and the time for which high gain is required. Final error alone is misleading because the trajectory becomes constant during the second half of the run and the system has time to recover. For this reason, IAE and RMSE should always accompany |e(T)|.

The saturation experiment should consequently be read as a constraint study. A lower actuator limit caps the instantaneous command and necessarily suppresses part of the cancellation and PID action. The controller can recover after the demanding transient, which explains why the final error can look satisfactory even when the accumulated error is larger. The result is a reminder that final-value accuracy alone is insufficient for constrained nonlinear tracking; integral metrics and saturation activity are needed to reveal what occurred during the transient.

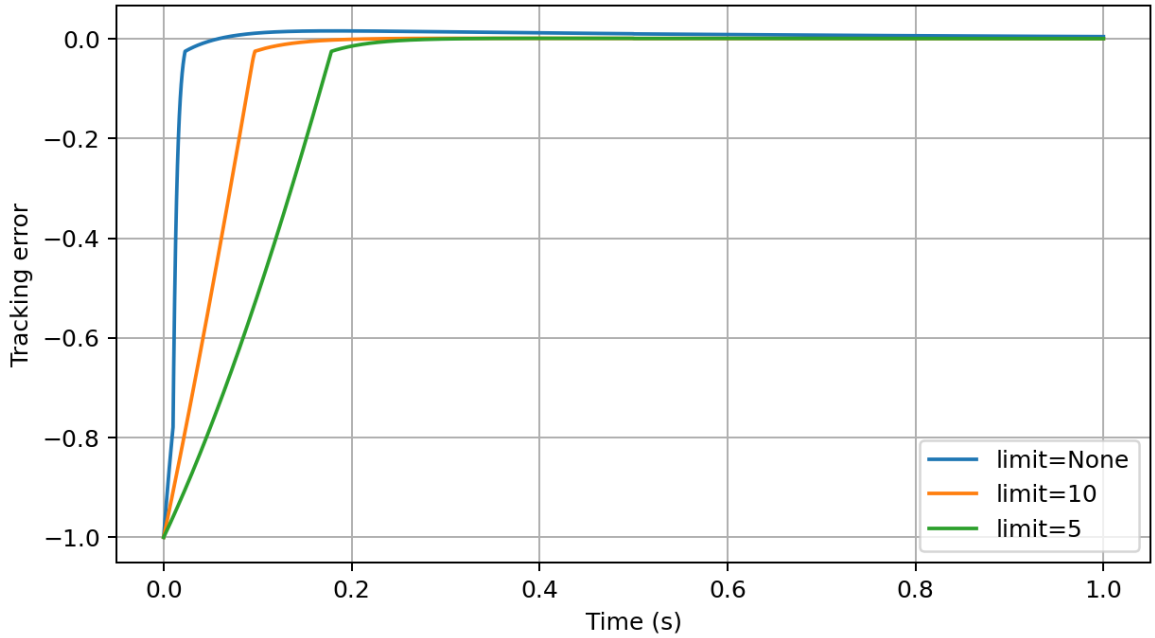


Figure 6. Tracking-saturation trade-off.

3.3 Effect of Noise

Table 4. Effect of AWGN after causal filtering.

SNR_dB	final_error	IAE	RMSE	energy	high_pct
clean	0.00465499	0.0216301	0.0962064	82.0607	1.27987
40	0.00462007	0.021808	0.0971331	93.7797	1.28987
30	0.00457354	0.0222615	0.0991476	193.377	1.28987
20	0.00442636	0.0237034	0.105533	1177.01	1.28987

With the causal filter and error hysteresis, high-gain utilization remains about 1.29% even at 20 dB. IAE increases smoothly from 0.02163 to 0.02370, while Eu increases because of the effort required to track the residual noise. This is clearer and practically implementable behavior.

The noise sweep indicates that the combination of causal filtering and hysteretic switching prevents measurement fluctuations from producing excessive mode changes. The modest increase in IAE with decreasing SNR reflects residual tracking disturbance, while the rise in Eu reflects the additional corrective action required against the filtered noise. Thus, the principal benefit observed here is not immunity to noise, but controlled degradation: the switching duty cycle remains essentially unchanged and the performance metrics vary smoothly rather than becoming erratic.

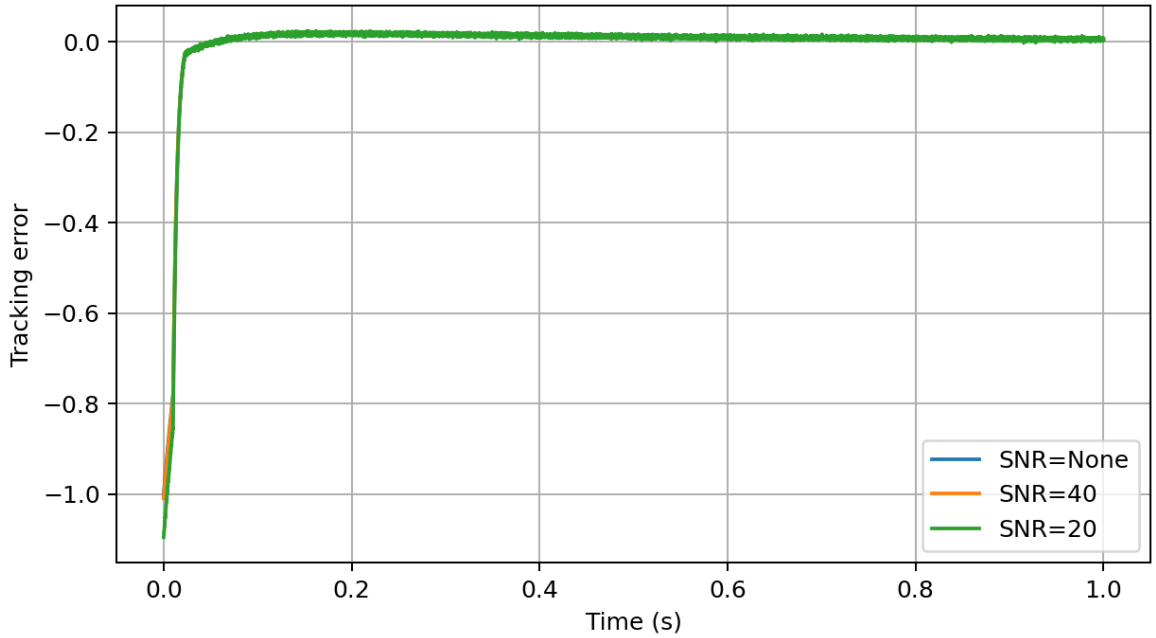


Figure 7. Tracking error for different SNR values.

4 Discussion

The architecture is effective when f(x) is known accurately enough for the compensation to be reliable. In the presence of model error, the residual nonlinearity appears as a disturbance in the error dynamics and requires robust/ISS analysis. Accordingly, exact cancellation is not presented as a general solution for unknown nonlinearities.

The use of two PID gain sets is closely related in practice to gain scheduling. The addition of hysteresis and dwell time is essential because switched-systems theory shows that fast or arbitrary switching may destroy stability properties even when the individual modes are stable. In the present scalar example, the common structure of the error equation and the positive coefficients in each mode provide strong evidence of stable behavior, but a general theorem under model uncertainty remains future work.

Saturation must be treated as part of the dynamics rather than merely as amplitude protection. When it becomes active, cancellation of f(x) is no longer exact and integrator windup can worsen the transition. For this reason, the implementation includes anti-windup.

5 Conclusions

The controller is consistently described as a switched-gain PID applied to a feedback-linearized scalar plant, with hysteresis, dwell time, causal filtering, and anti-windup.

In the baseline scenario, the switching law reduces energy by about 30.2% relative to the permanently high-gain PID, at the cost of increased IAE but with very limited use of the high-gain state (about 1.28%). The permanently low-gain PID requires still less energy but exhibits clearly worse tracking. The switching PID therefore occupies an intermediate Pareto point between accuracy and control effort.

The study of order n confirms that exact cancellation decouples the ideal error dynamics from the nonlinearity, but not the actual control energy. Saturation reduces Eu at the cost of increased transient error, while the causal filter and hysteresis maintain controlled switching under noise. These results provide a basis for future extensions to uncertain f(x), robust switching criteria, and multi-objective optimization of the gain sets and thresholds.

Taken together, the simulations support a bounded conclusion: for the exactly modelled scalar plant, switched gains can reduce actuator effort relative to permanent high-gain operation while retaining substantially better tracking than permanent low-gain operation. The order, saturation, and noise studies then identify the conditions under which this advantage is preserved or weakened. The results do not establish a universal optimal switching controller; instead, they provide a reproducible baseline showing how cancellation accuracy, actuator limits, noise treatment, hysteresis, and dwell time determine the practical value of the architecture.

Relationship to the Complete Book

This paper is a self-contained study extracted and reformatted from the corresponding research chapter of *Nonlinear Control: From Model to Evidence*. The complete book provides the unified progression across all seven studies, cross-study synthesis, extended methodological context, and the MATLAB/Octave and Python software-companion architecture used for replication and further experimentation.

5.1 References

[1] J. G. Ziegler and N. B. Nichols, “Optimum Settings for Automatic Controllers,” Trans. ASME, 1942.

[2] K. J. Åström and T. Hägglund, Advanced PID Control, ISA, 2006; Wiley online edition, 2026.

[3] W. J. Rugh and J. S. Shamma, “Research on gain scheduling,” Automatica, 36(10), 1401–1425, 2000.

[4] D. Liberzon, Switching in Systems and Control, Birkhäuser, 2003.

[5] L. Vu, D. Chatterjee and D. Liberzon, “Input-to-state stability of switched systems and switching adaptive control,” Automatica, 43(4), 639–646, 2007.

[6] D. Chatterjee and D. Liberzon, “Stability analysis of deterministic and stochastic switched systems via a comparison principle and multiple Lyapunov functions,” SIAM

J. Control Optim., 2006.

[7] C. Kravaris, “Nonlinear state feedback synthesis by global input/output linearization,” AIChE Journal, 33(4), 592–603, 1987.

[8] W. Kang, “Approximate linearization of nonlinear control systems,” Systems & Control Letters, 23(1), 43–52, 1994.

[9] L. Zaccarian and A. R. Teel, Modern Anti-windup Synthesis: Control Augmentation for Actuator Saturation, Princeton Univ. Press, 2011.

[10] M. Caparroz, K. Soltesz, T. Hägglund and J. L. Guzmán, “Anti-Windup in PID Control: Review, Analysis, and New Tuning Directions,” arXiv:2606.01959, 2026.

[11] H. K. Khalil, Nonlinear Systems, 3rd ed., Prentice Hall, 2002.

[12] J.-J. E. Slotine and W. Li, Applied Nonlinear Control, Prentice Hall, 1991.

[13] A. Isidori, Nonlinear Control Systems, 3rd ed., Springer, 1995. DOI: 10.1007/978-1-84628-615-5.

[14] J. P. Hespanha and A. S. Morse, “Stability of switched systems with average dwell-time,” Proc. 38th IEEE Conf. Decision and Control, Phoenix, AZ, 1999, pp. 2655-2660.