

Speed Control Lab

Three controllers, one motor and a bounded comparison

The lab compares filtered PID, pole-placement state feedback and discrete LQR on the same ideal DC motor. The parameters and implementation are newly selected for this portfolio in 2026. Results are numerical predictions, with no identified hardware, historical measurements or experimental validation.

Plant, inputs and boundaries

The states are speed ω [rad/s] and current i [A], with applied voltage u [V] and prescribed resisting torque τ [N m]:

$$J\dot{\omega} = K_t i - b\omega - \tau, \quad L\dot{i} = u - Ri - K_e \omega.$$

Parameter	Value	Parameter	Value
J	0.02 kg m ²	b	0.04 N m s/rad
K_t	0.10 N m/A	K_e	0.10 V s/rad
R	2 Ω	L	0.05 H
Voltage	± 12 V	Sample period	0.001 s

Each run starts at $[\omega, i] = [0, 0]$, steps the target to 5 rad/s, applies 0.15 N m from 2.5 through just before 4.5 s, and ends at 7 s. Inputs hold between controller samples. Exact zero-order-hold propagation uses the matrix exponential, splitting intervals at load edges. The endpoint is recorded without an extra integration interval.


Controller definitions

PID: $u^* = K_p(r - \omega) + z - K_d \dot{\omega}$, with $K_p = 2.4$, $K_i = 6$, $K_d = 0.04$ in V s/rad, V/rad and V s²/rad. The measurement derivative has a 0.02 s exponential filter. An explicit integral update adds $K_i(r - \omega)\Delta t$ in volts, except when conditional anti-windup rejects an increment pushing farther into saturation. The initial derivative is zero.

State feedback: $u^* = Nr - k_\omega \omega - k_i i$. Placing nominal continuous poles at $-5, -50$ s⁻¹ gives $K = [1.34, 0.65]$ and $N = 2.5$ in physical units. The actual sampled poles are 0.99500723 and 0.95096493. Both states are known ideally; there is no integral disturbance rejection.

LQR: deviations from $[r, 0.4r]$ and $u = 0.9r$ are normalized by 10 rad/s, 1 A and 12 V. The discrete cost is $\sum \Delta t [10x_{n,1}^2 + x_{n,2}^2 + u_n^2]$. The DARE gives $K = [1.08094974, 9.01535024]$ and $N = 5.58708983$. This is an unconstrained nominal optimum. Applied clipping is outside that optimization. The reference equilibrium assumes zero load.

Default response and tradeoffs




Shading marks the prescribed load pulse. All controllers receive the same experiment; PID rejects the sustained offset through its integral state, while the two state-feedback designs have no integral state. Comparing the final error alone would hide the error accumulated during the pulse.

Controller	Error at 7 s rad/s	IAE rad	$\int u^2 dt$ V ² s	Peak $ i $ A
PID	-0.001934	1.886	231.066	4.985
State feedback	0.000006	4.281	187.817	4.211
LQR	0.005104	7.916	151.419	2.507

IAE and voltage-squared effort use left-held sample quadrature over actual intervals. The effort metric is **not electrical energy**. Current is calculated and displayed without any current-protection model. Peak current in the broader 394-case range test reaches 6.345 A. A numerical range check does not establish hardware safety.

Parameter assumptions and a useful failure



The study holds nominal controller gains fixed while perturbing each of J, b, K_t, K_e, R, L by $\pm 20\%$, then adds two combined corners: 15 plants and 45 runs. These are exploratory coefficient choices, not measured manufacturing tolerances. Separate K_t and K_e perturbations isolate coefficients and need not represent a reciprocal physical motor; the two combined corners retain $K_t = K_e$. Full parameters and metrics are saved for every run. No architecture is asserted to win every objective.

The baseline has equal SI torque and back-EMF constants. Its stored mechanical plus magnetic energy satisfies

$$\frac{d}{dt} \left(\frac{1}{2} J \omega^2 + \frac{1}{2} L i^2 \right) = u i - R i^2 - b \omega^2 - \tau \omega.$$

The independent check evaluates this identity at 100 seeded states and inputs. The maximum discrepancy is 4.27×10^{-14} W.

At a 5 rad/s target and 0.8 N m load, equilibrium requires 10 A and 20.5 V. The available 12 V can sustain only 0.375 N m at that speed. The model therefore retains the resulting failure, including load-driven reverse rotation. The imposed torque retains its positive sign through zero speed; it is an external torque, not dry friction.

Conditional anti-windup lowers the recovery peak from 13.218 to 5.554 rad/s in that experiment. The minimum speed remains about -4.355 rad/s; anti-windup cannot supply missing voltage. Saturation duration is 2.107 s with anti-windup and 4.150 s without it. These predictions explain the controller mechanism and actuator boundary; they are not motor test measurements.

Scope limits

States are ideal and known. There is no observer, sensor noise, nonlinear friction, heating, backlash, current protection, inverter switching or physical calibration. A fabricated motor/controller, CAD, physical BOM, harness and fabrication drawings are outside this mathematical boundary. The source package includes the applicable software dependency list, SI configuration, equations, test inputs, complete model and reproducible results.

Verification and reproduction

The existing range suite covers 394 complete runs and 27 LQR designs. It checks equilibrium, held-input semigroup behavior, independent P-only and load-edge analytic responses, pole placement, DARE residual, accumulated optimal cost, weight scaling, bounds and endpoint metric accounting. The largest analytic load-edge discrepancy is 4.56×10^{-14} rad/s.

A separate Python reference constructs the augmented continuous matrices from SI equations and uses SciPy `expm`. One-step error is below 1.60×10^{-15} SI; replay of each nominal held-voltage trajectory differs by less than 6.0×10^{-13} SI. SciPy's Schur DARE solution differs from the iterative JavaScript physical gain by less than 2.33×10^{-10} . These are independent numerical methods, not experimental evidence.

Sample-period refinement compares 2, 1 and 0.5 ms. The table shows the maximum difference between aligned 1 and 0.5 ms samples. The plant itself propagates exactly for a held command; these differences arise from sampled control and PID update timing. Refinement changes the controller, including the DARE design.

Controller	Speed [rad/s]	Current [A]	Effort [V ² s]
PID	0.001464	0.002456	0.01544
State feedback	0.001534	0.005652	0.01508
LQR	0.000390	0.000470	0.00781

All satisfy the frozen gates: speed < 0.08 rad/s, current < 0.12 A, IAE change < 0.03 rad and effort change < 0.5 V²s. Numerical agreement does not bound errors in the selected physical assumptions.

The browser, static default plots and CSV use the same saved model. The shared CSV serializer records settings and units, exports every tenth sample for all three controllers, and includes the final sample even when the interval does not divide the run. The default export contains 2,103 data rows. Full 1 ms reference traces are regenerated by the study runner.

From an extracted source package, run with Node.js 20 or newer:

```
node test-model.cjs
node test-speed-package.cjs
node study.cjs
```

For independent verification and figures, install the Python dependencies in `requirements.txt`, then run `python verify_independent.py` and `python make_figures.py`. Serve `index.html` with a static server to inspect the portable interface. Reset restores the complete default experiment. The voltage and current plots follow the selected controller, while all three speed responses and metrics remain comparable.

Theory reference: Russ Tedrake, *Linear Quadratic Regulators*, MIT Underactuated Robotics, discrete-time Riccati formulation: <https://underactuated.mit.edu/lqr.html>. Motor values and new implementation are illustrative selections.