Autonomous Robotics

Control Theory

Josiah Hanna
University of Wisconsin — Madison

Reading Responses

- Overall, great responses.
- If you submitted, then you likely received most of the 10 possible points.
  - Sometimes, points were lost when responses lacked detail demonstrating the text had been read.
- Reading on Bayes filter is now available on the course website.

Learning Outcomes

After today's lecture, you will:

- Be able to define and give examples of set points and system error.
- Be able to define and implement common control laws: bang-bang control, P-, PI-, PID control
- Be able to compare and contrast open- vs closed-loop control.

Deterministic Interaction Model

$$u_t \leftarrow H(y_t)$$

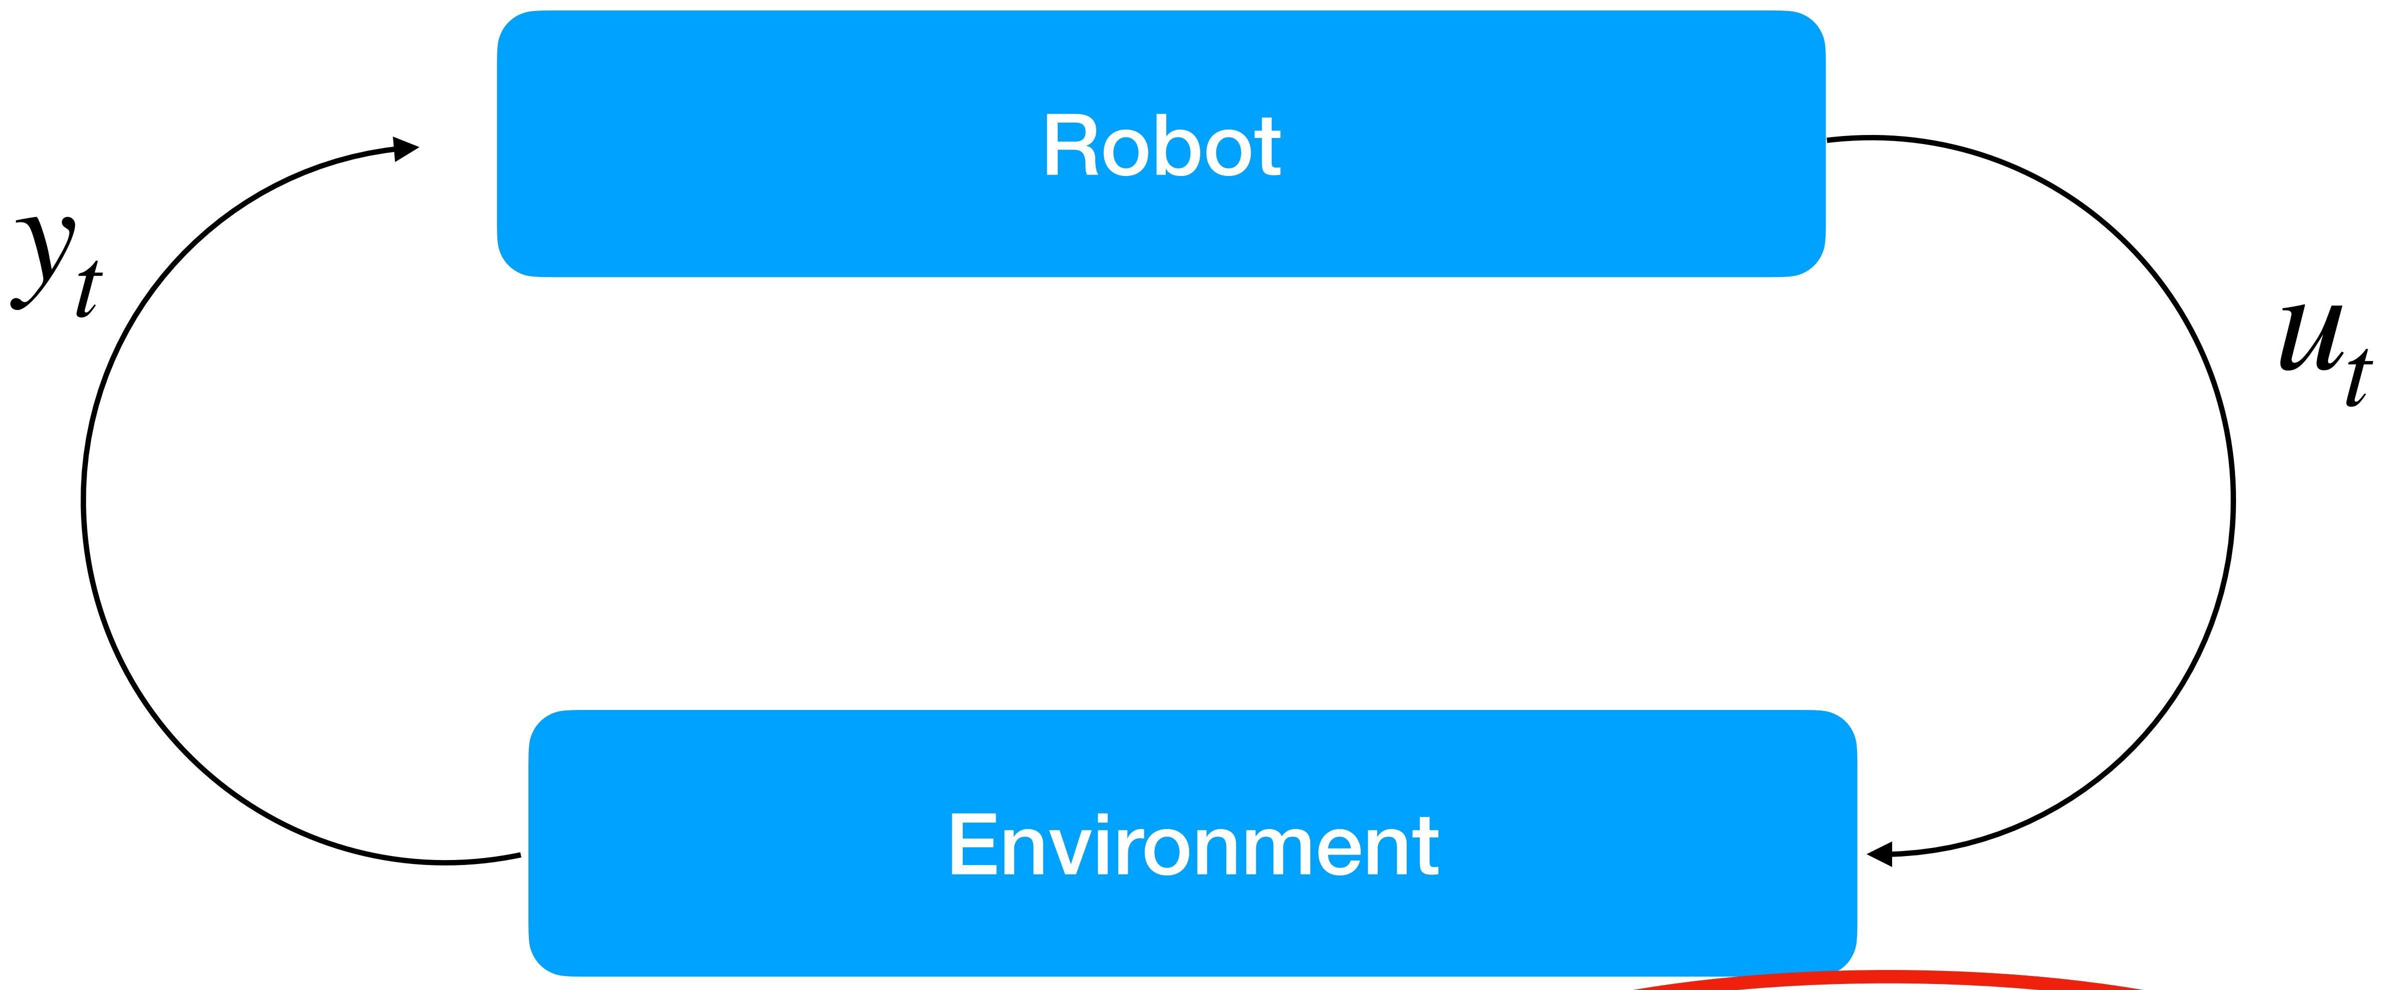

$$x_{t+1} \leftarrow F(x_t, u_t)$$

Or

$$\dot{x}_t \leftarrow F(x_t, u_t)$$
$$x_{t+1} \leftarrow x_t + \dot{x}\Delta t$$

$$y_{t+1} \leftarrow G(x_{t+1})$$

Control Objective

- Goal: bring the robot's state, x, to a desired state, $x_{\texttt{set}}$.
- Use error, $e = x - x_{\texttt{set}}$, to measure how close the robot is to achieving this.
- Assumptions:
  - The state x is observable.
  - Increasing u will increase x.
  - Simplification: everything is 1-dimensional

Control Examples

Bang-Bang Control

- Simplest control law: toggle between choosing one of two values for u.

$$\begin{array}{llll} \text{if} & e < -\epsilon & \text{then} & u := \mathit{on} \\ \text{if} & e > +\epsilon & \text{then} & u := \mathit{off} \end{array}$$

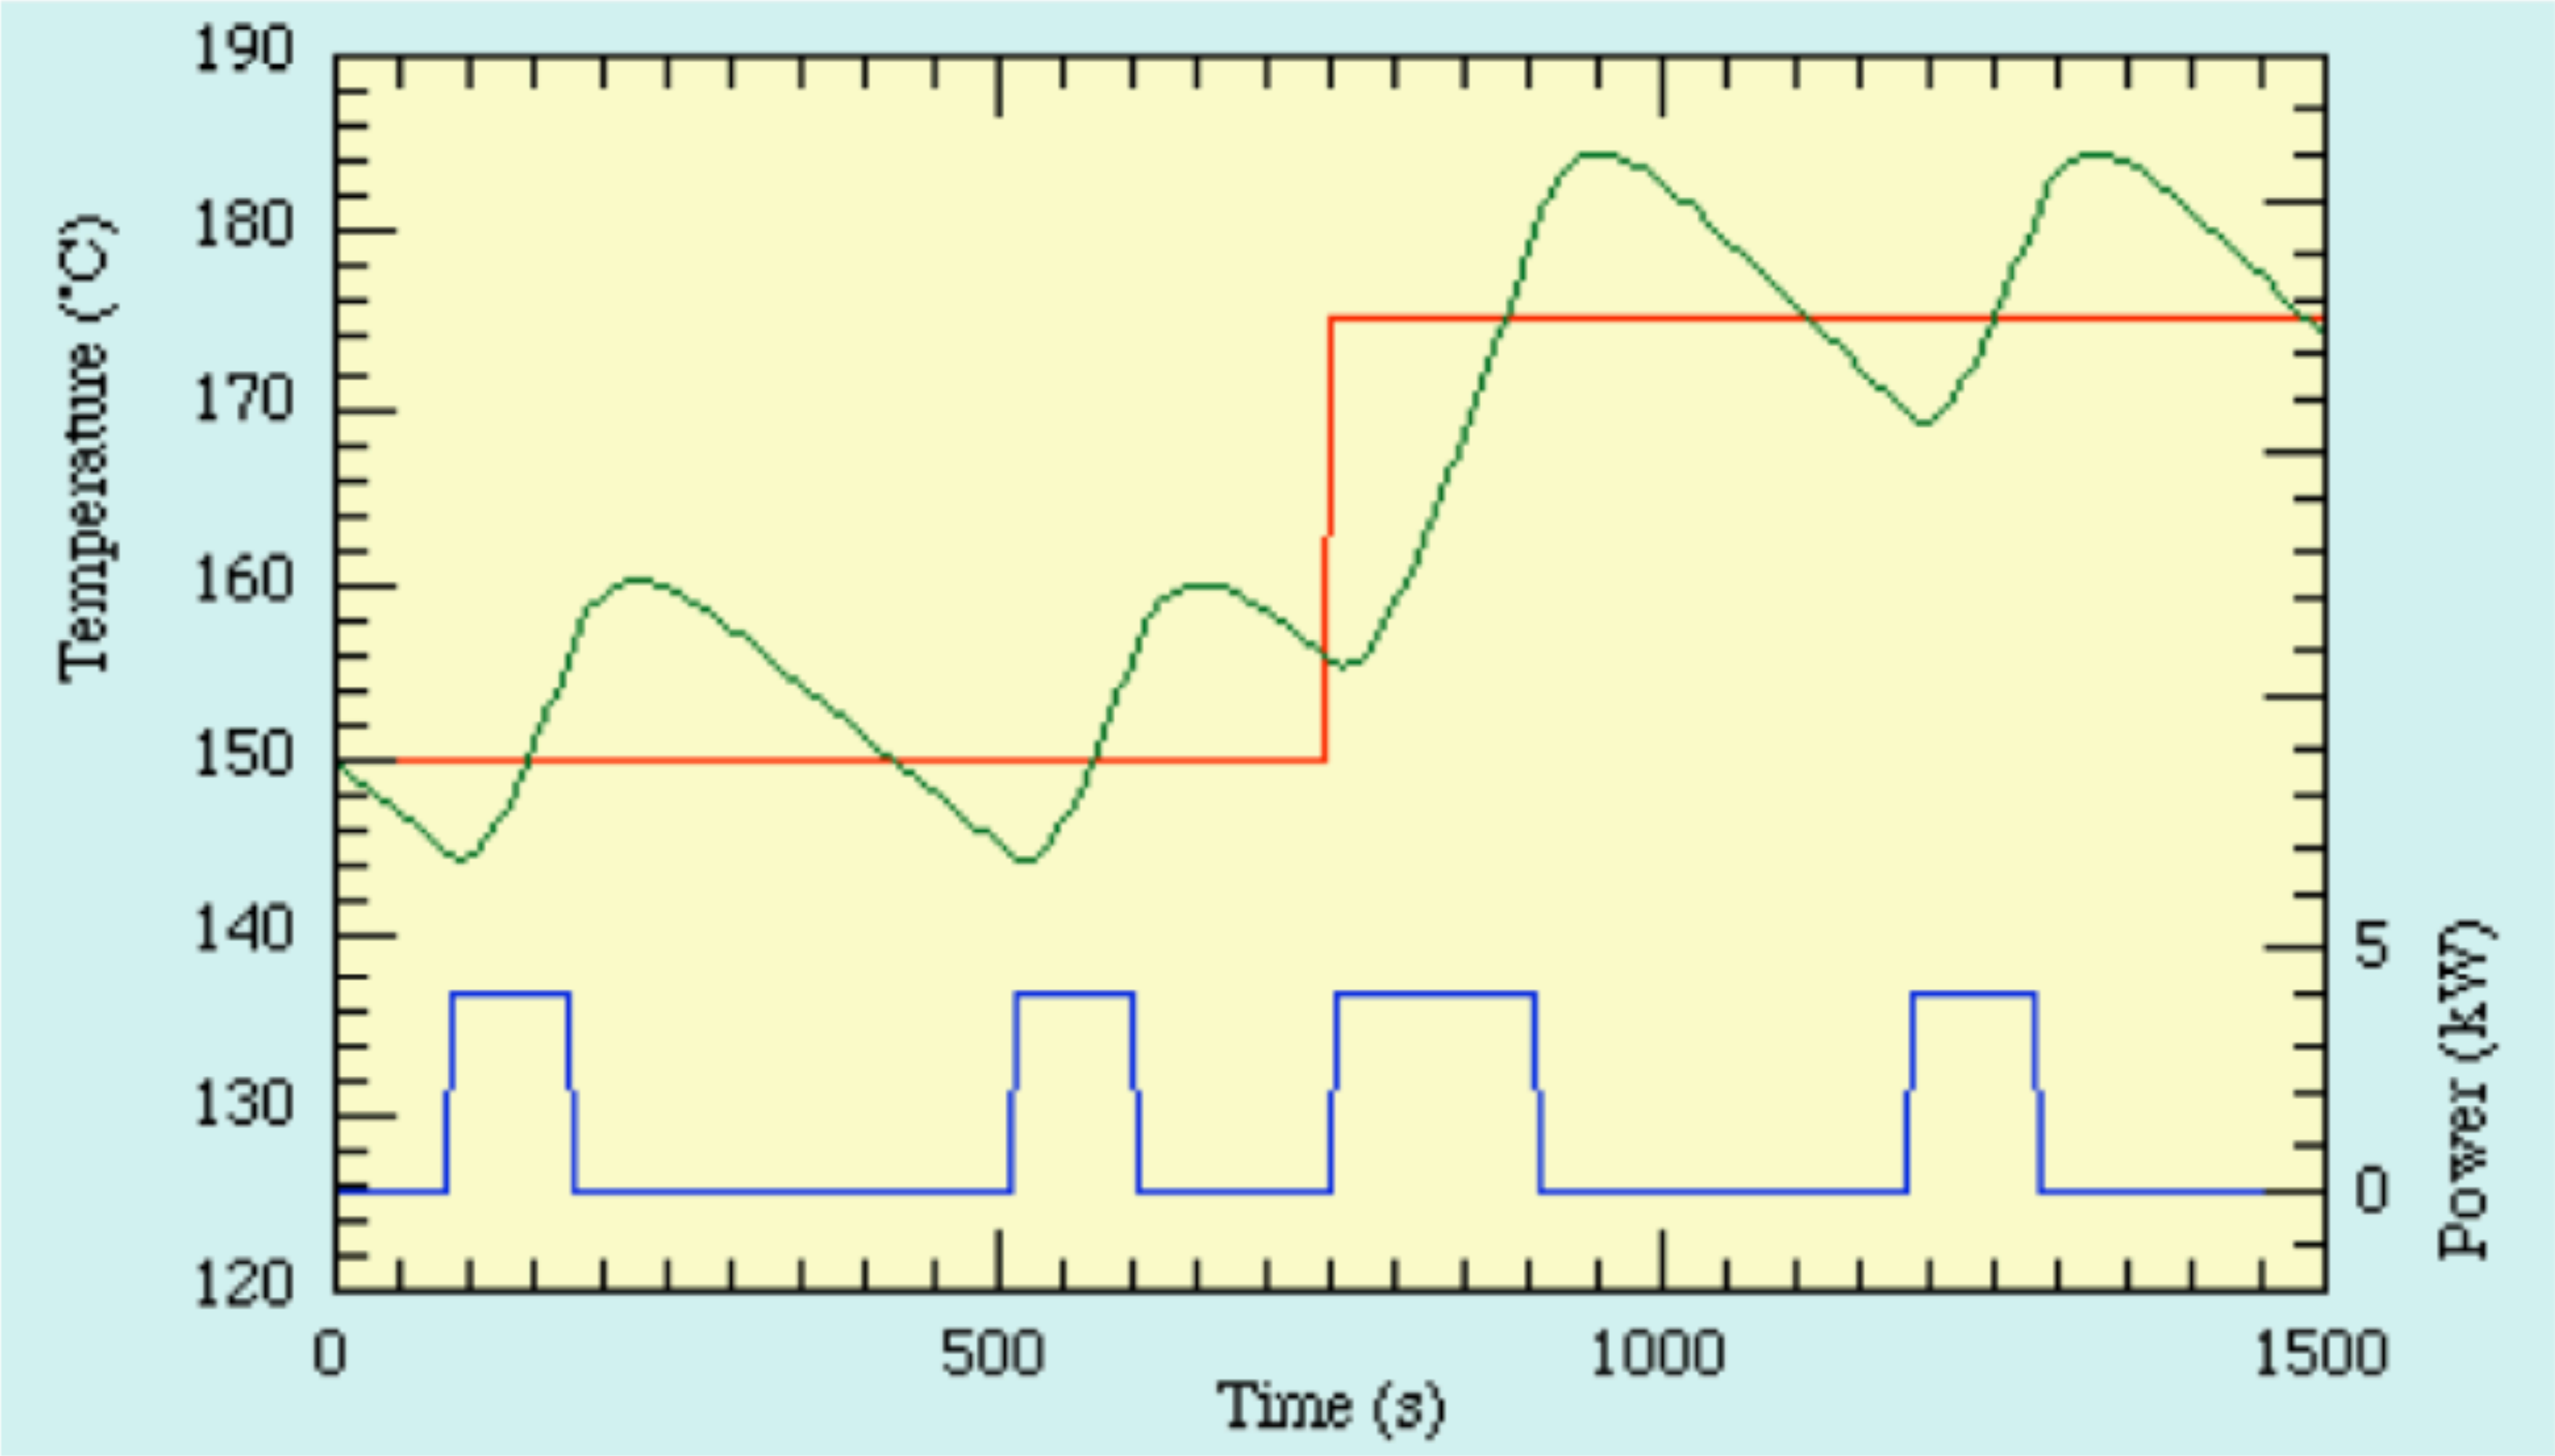

$$e = x - x_{\texttt{set}}$$

P-Control

- Improve upon bang-bang control by acting proportionally to error.

$$u = -k_P * e + u_b$$

- Constant offset is necessary to counteract external forces.
  - Example: a robot arm falls in the absence of any control.
  - Must set u_b so that $\dot{x} = F(x_{\texttt{set}}, u) = 0$

$$e = x - x_{\texttt{set}}$$

P-Control

$$u = -k_P * e + u_b$$

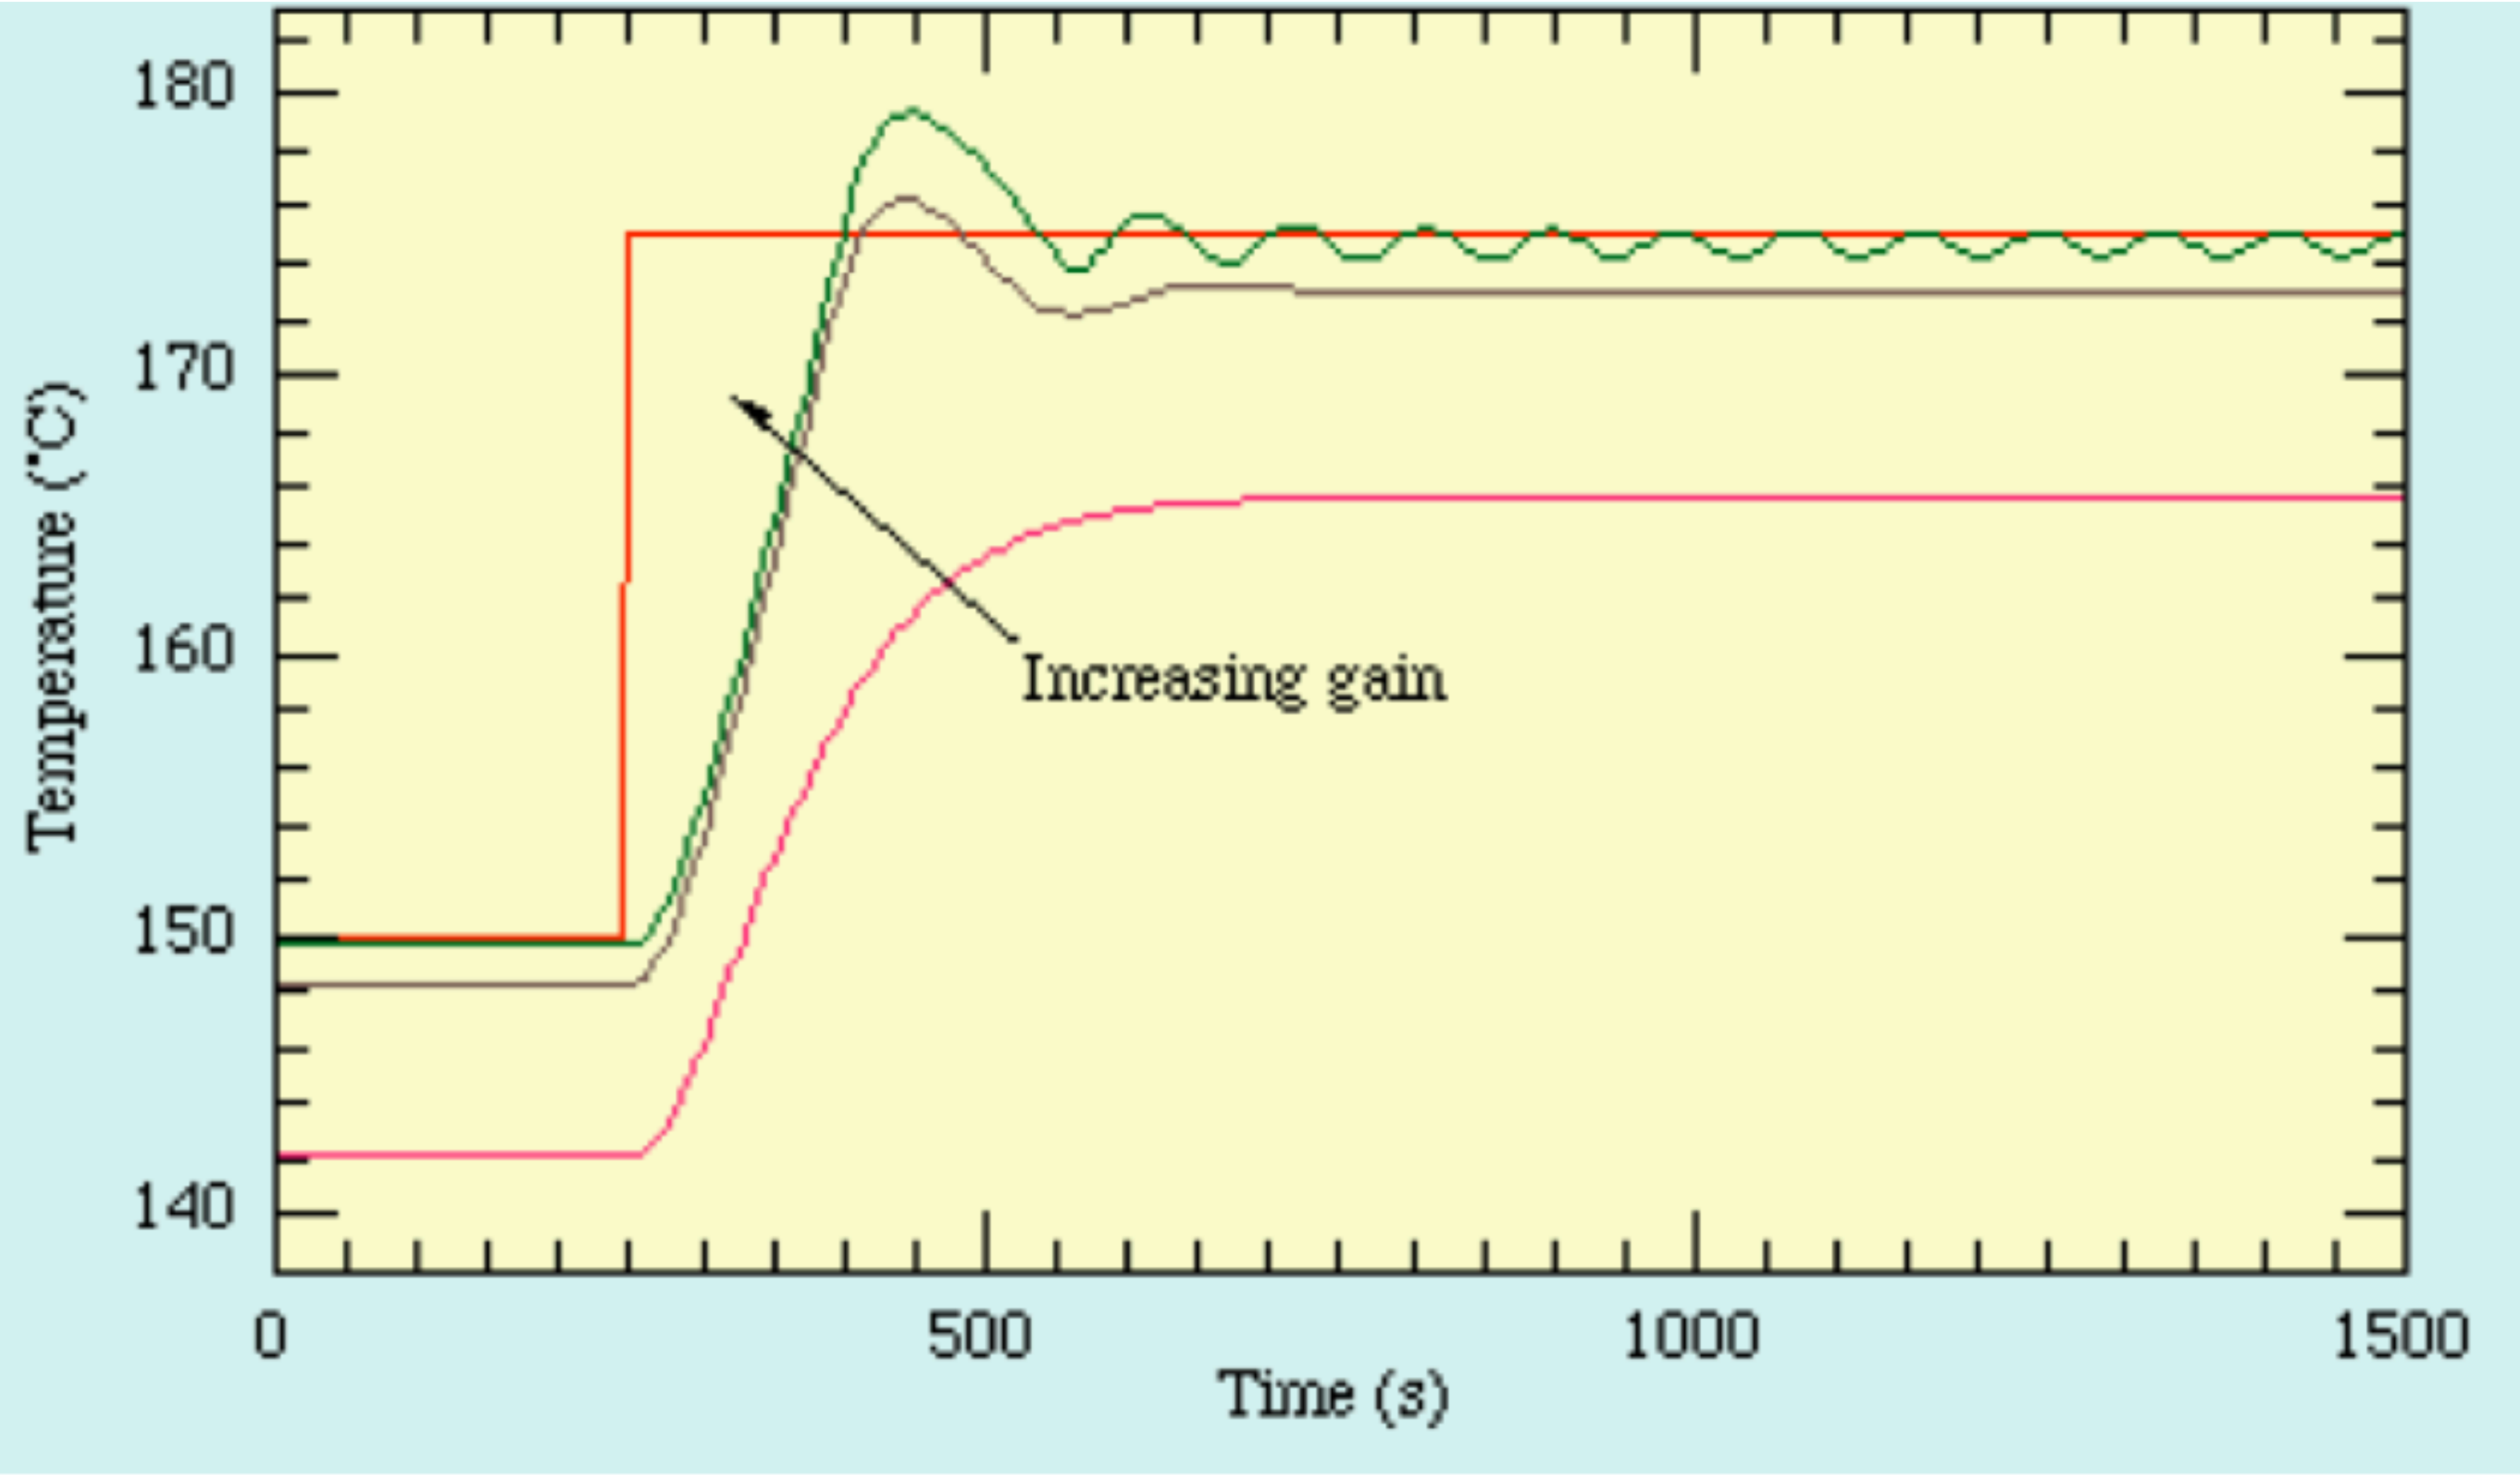

$$e = x - x_{\mathtt{set}}$$

Adaptive Control

- Update controller parameters to decrease persistent errors.
  - Can be thought of as a form of reinforcement learning.
- Habituation: change the set point to decrease the error.

$$x_{\texttt{set}} = k_H e$$

- Why do we do this?
  - Avoid excessively strong controls.

$$e = x - x_{\texttt{set}}$$

PI-Control

- P-control might result in convergence to a steady-state offset.
- Solution: "push" harder when error is not decreasing.

$$u_t = -k_P e(t) - k_I \int_{i=0}^{t} e(i)di$$

- In discrete-time systems, implement integral with a sum over time.
- Beware of wind-up!

$$u_t = -k_P e(t) - k_I \sum_{i=0}^{t} e(i)$$

  - Important to reset cumulative error if $x_{\texttt{set}}$ changes.
  - May need to cap cumulative error or limit use of I-term.

$$e = x - x_{\texttt{set}}$$

PD-Control

- P-control might result in overshoot and oscillation
- Solution: dampen push when the error derivative is large.

$$u_t = -k_P e - k_D \dot{e}$$

- In discrete-time systems, must approximate $\dot{e} \approx e(t) - e(t-1)$.

$e = x - x_{\texttt{set}}$

PD-Control

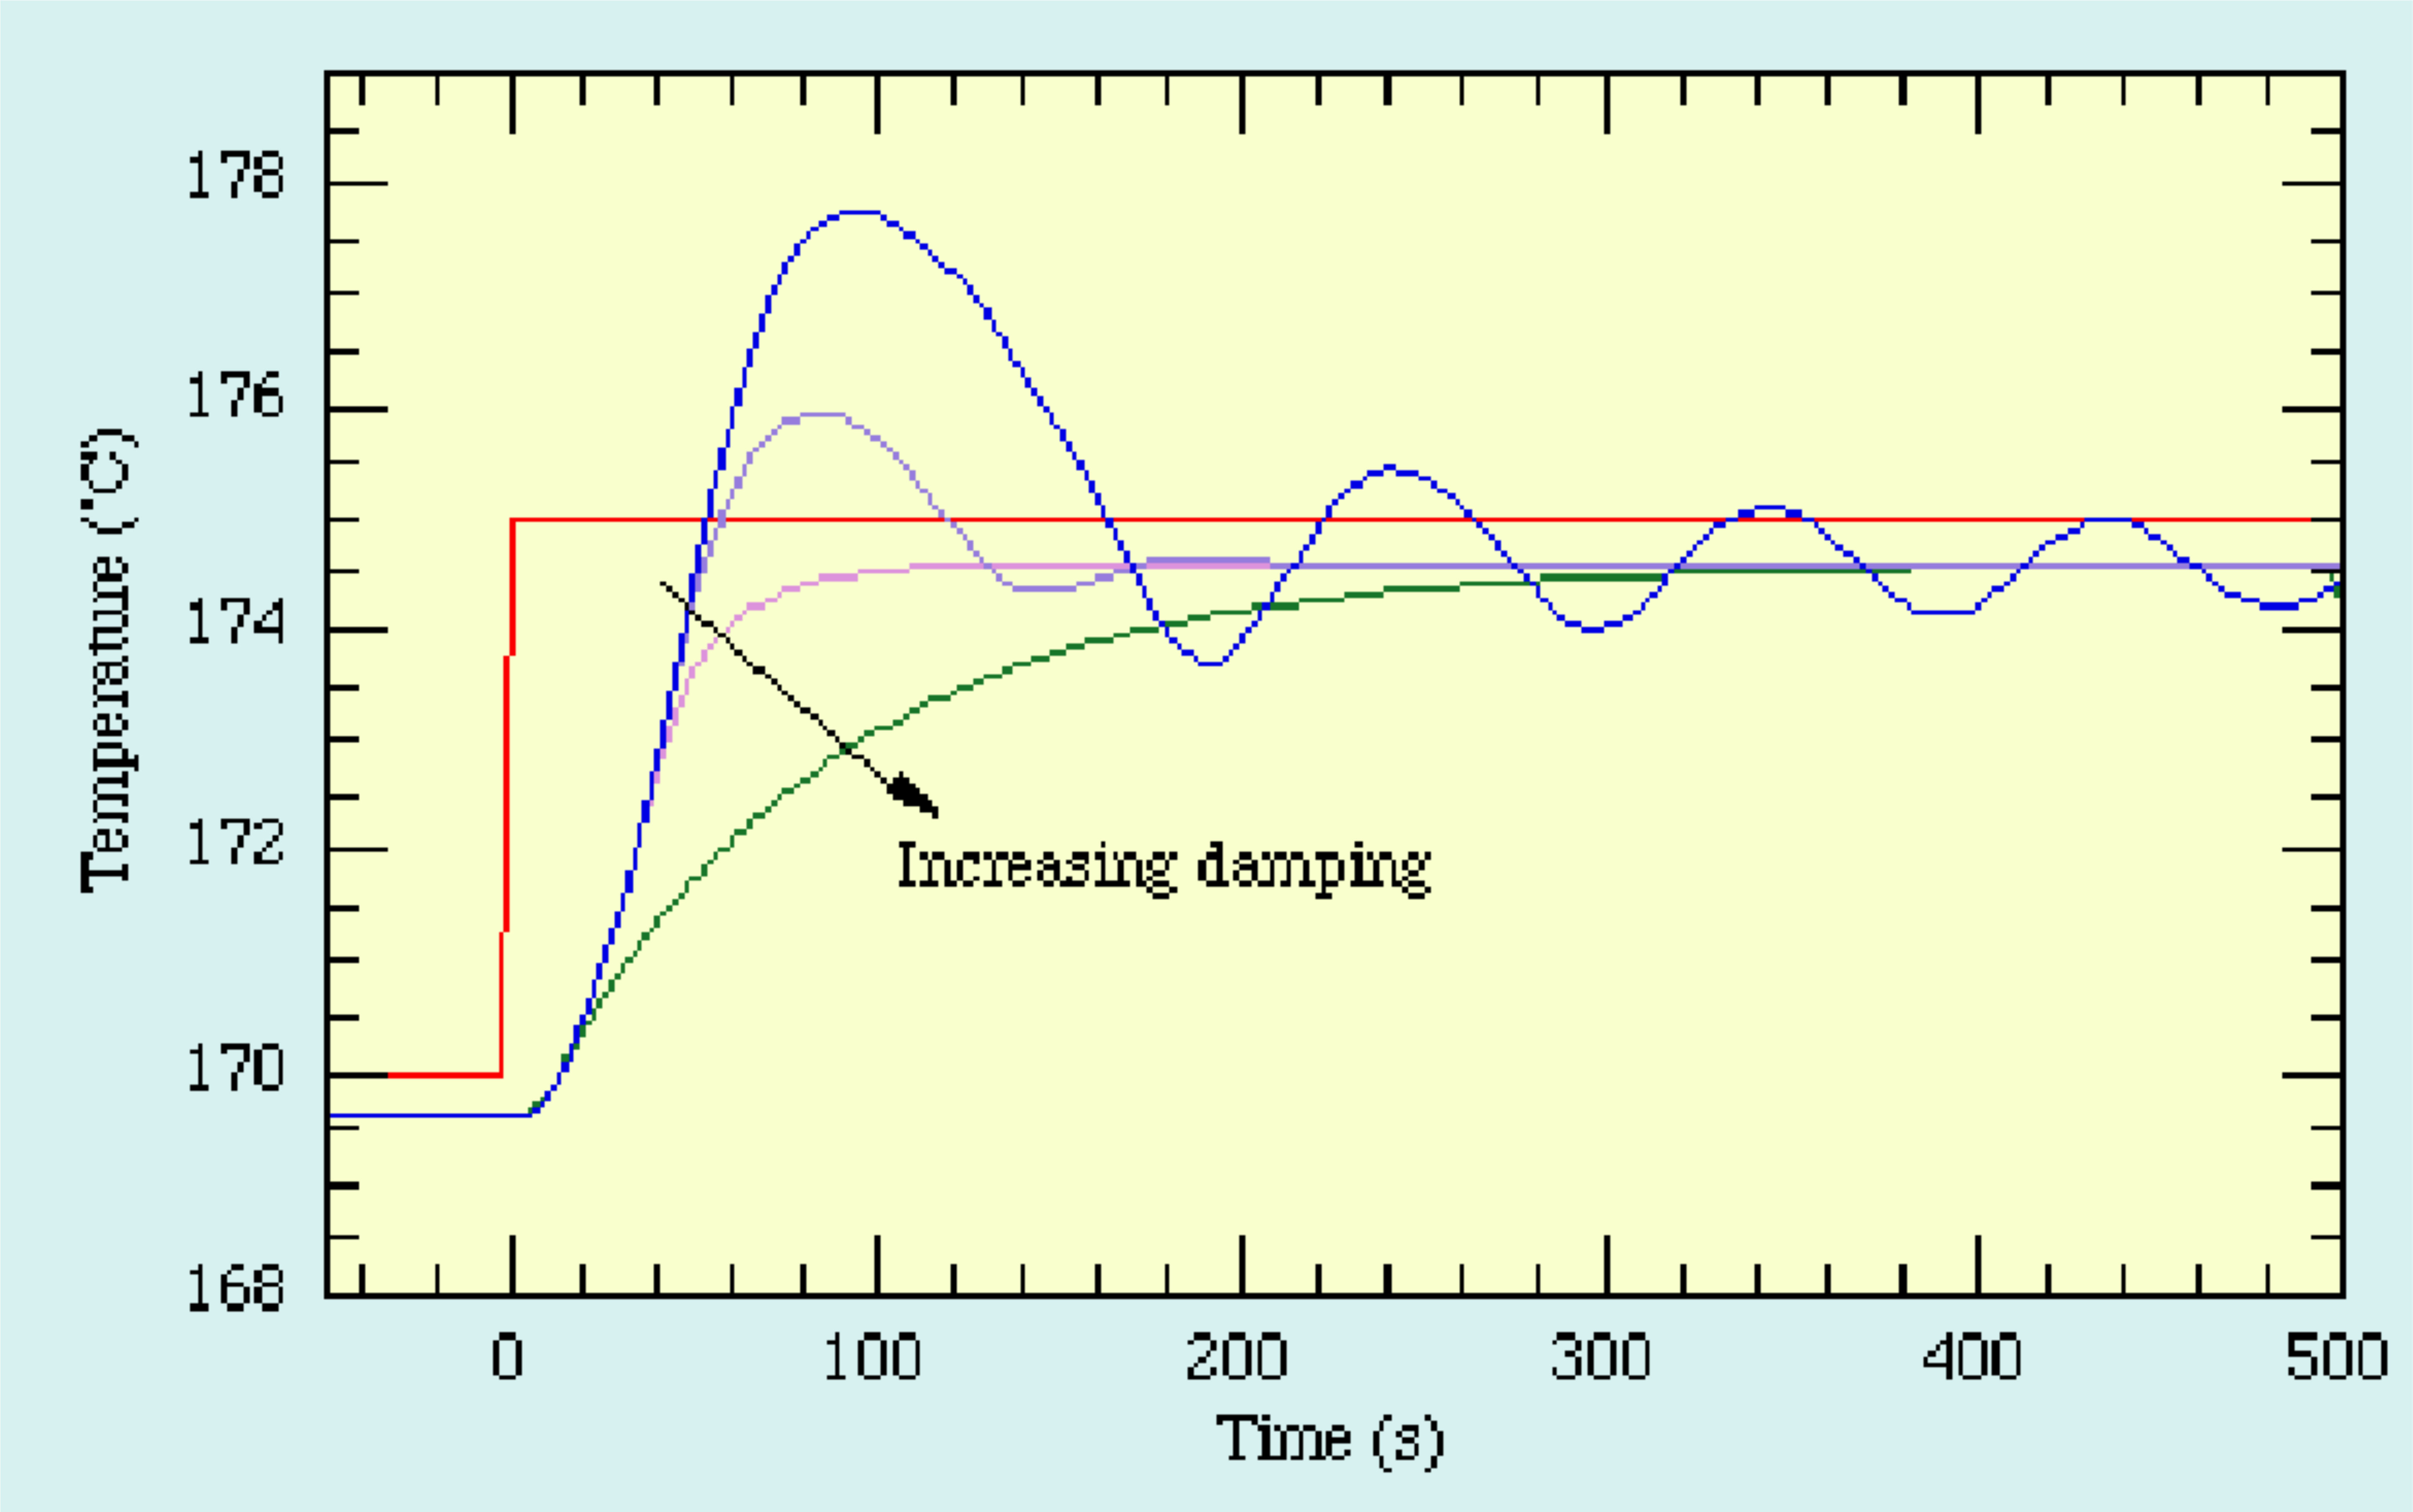

$e = x - x_{\text{set}}$

PID-Control

$$u_t = -k_P e(t) - k_I \int_{i=0}^{t} e(i)di - k_D \dot{e}$$

- Combine PI and PD controllers:
  - P term pushes back in proportion to error.
  - I term increases the push if error does not decay to zero.
  - D term dampens the push if error is decreasing too fast.

$$e = x - x_{\texttt{set}}$$

Closed Loop Control

- Choose action in response to state / sensor observation.
- Also known as feedback control.
- Advantages: most up-to-date information to base decision on, robust to stochastic disturbances.
- Disadvantages: need to wait for sensor reading, potentially more computationally complex.

Open Loop Control

- If we know x_0 and we know $F(x, u)$, then we don't need to know any future x_t to compute an optimal control sequence.

- Open-loop control: choose and then execute control actions without adjusting control based on new sensor observations.

- Advantages: no need to wait for sensors to respond, less computation

- Disadvantages: cannot take advantage of new information; not robust to stochastic disturbances or error in the <u>model of F.</u>

Model Predictive Control

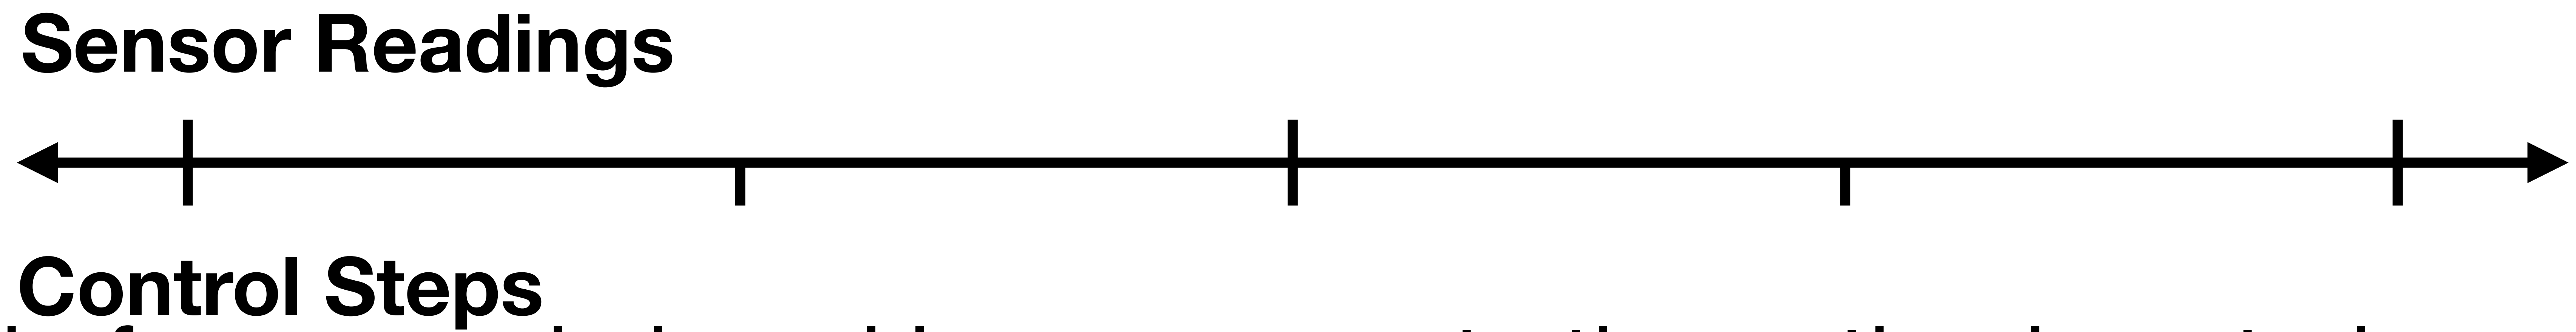

- Hybrid of open and closed loop: compute the optimal control sequence, $u^*_{t:T}$

- Follow control sequence until a new state is observed. Then re-compute control sequence.

- Advantages: Best of both open and closed loop

- Disadvantages: Potential computation cost (mitigate with models that enable fast planning).

Summary

- Covered basic control laws: bang-bang, proportional or P, PI, PD, and PID.
- Introduced open-loop and closed-loop control and discussed blending them with model-predictive control.

Braitenberg: Vehicles 1 and 2

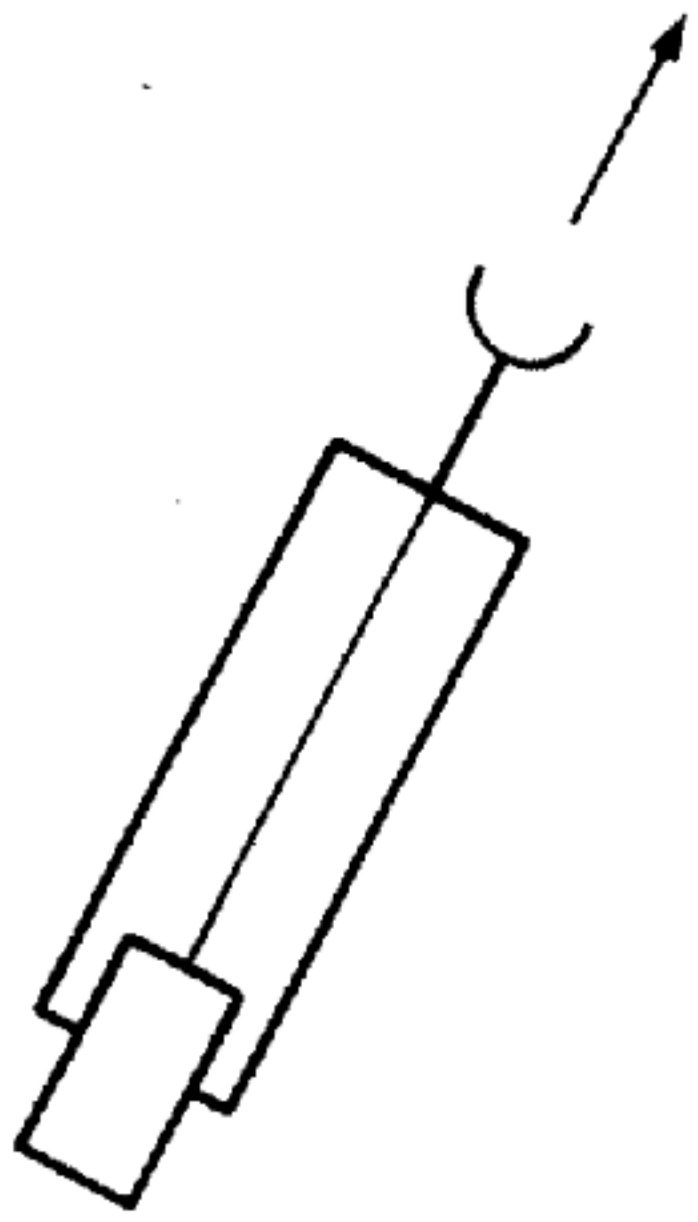

Vehicle 1

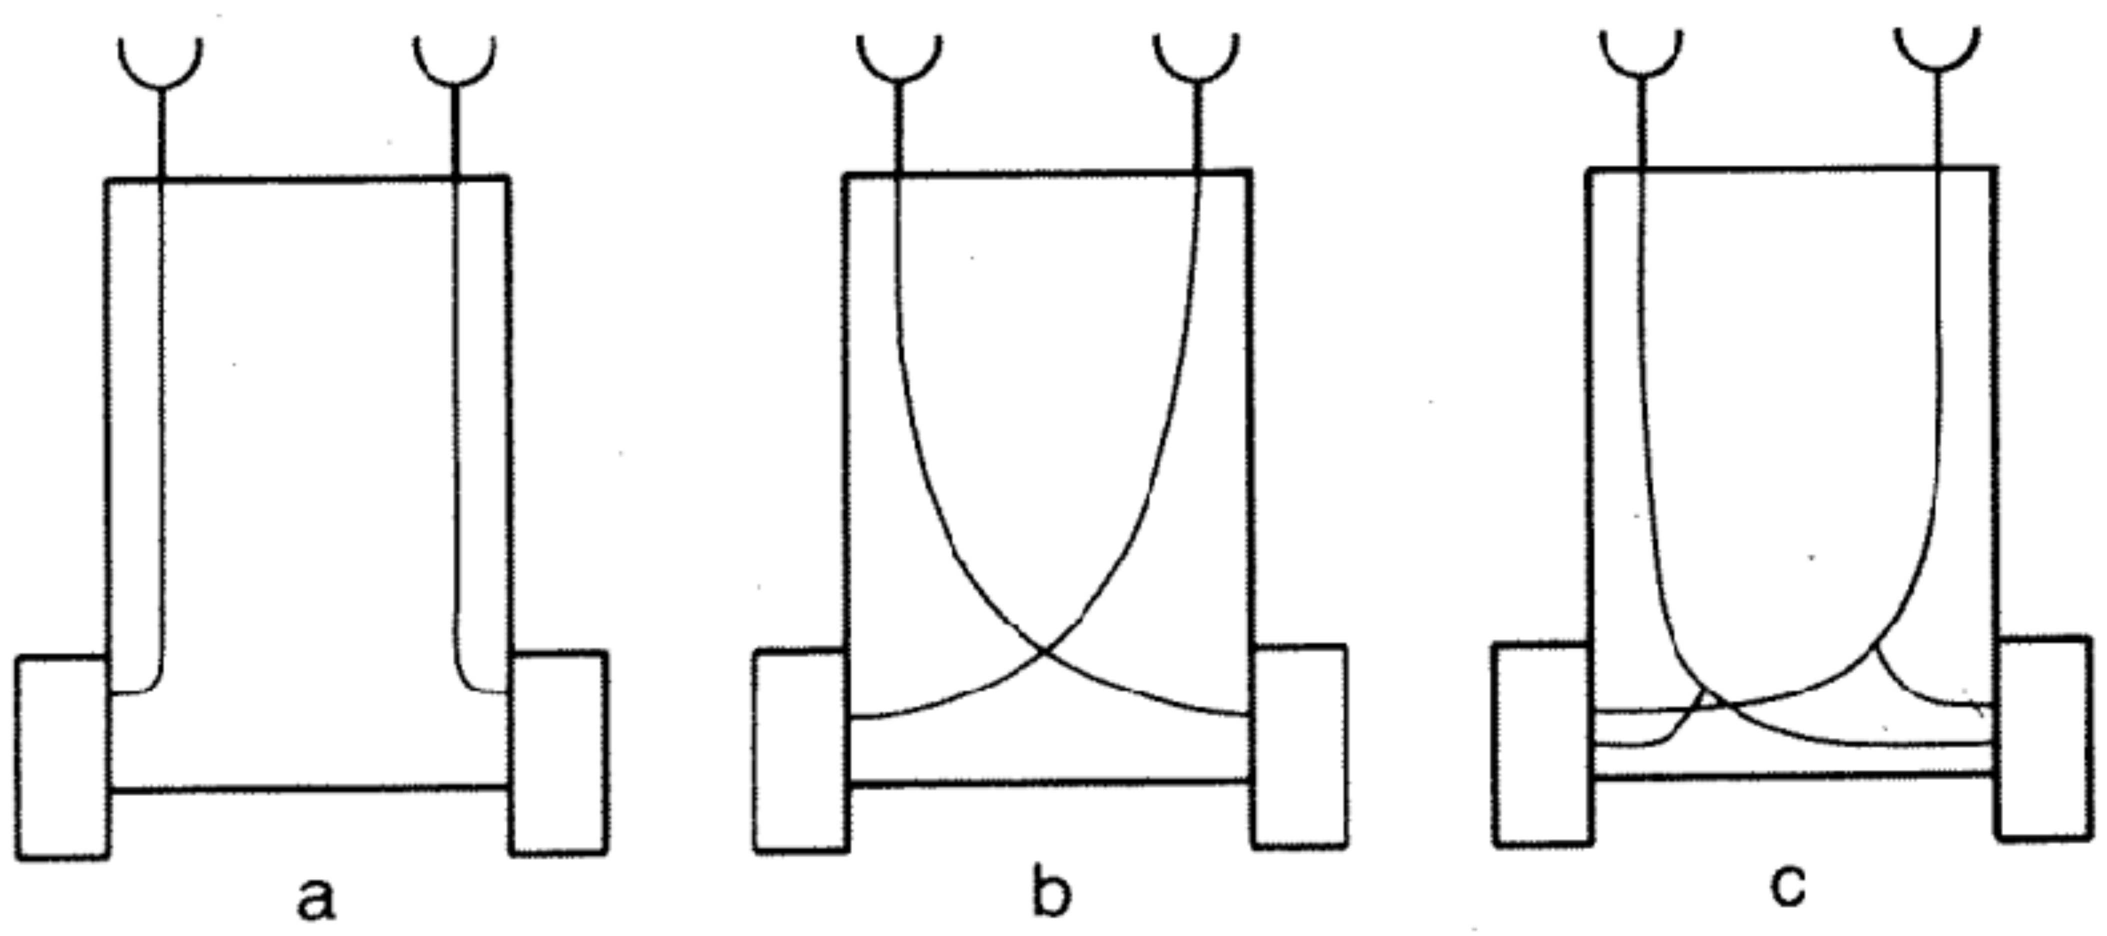

Vehicle 2

Braitenberg: Vehicles 1 and 2

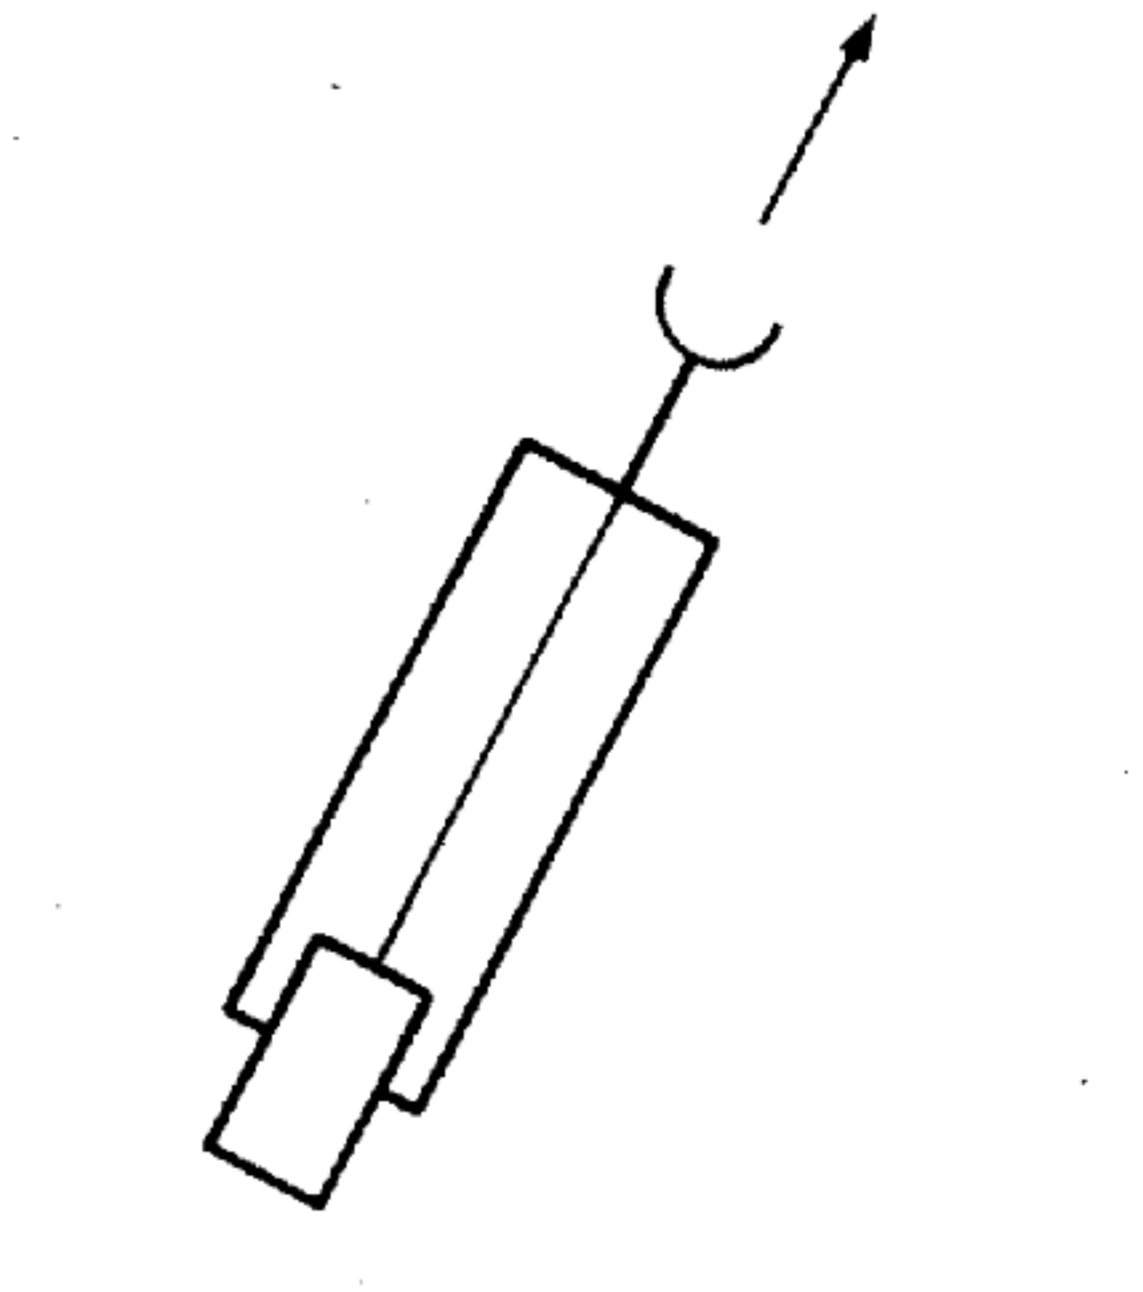

Vehicle 1

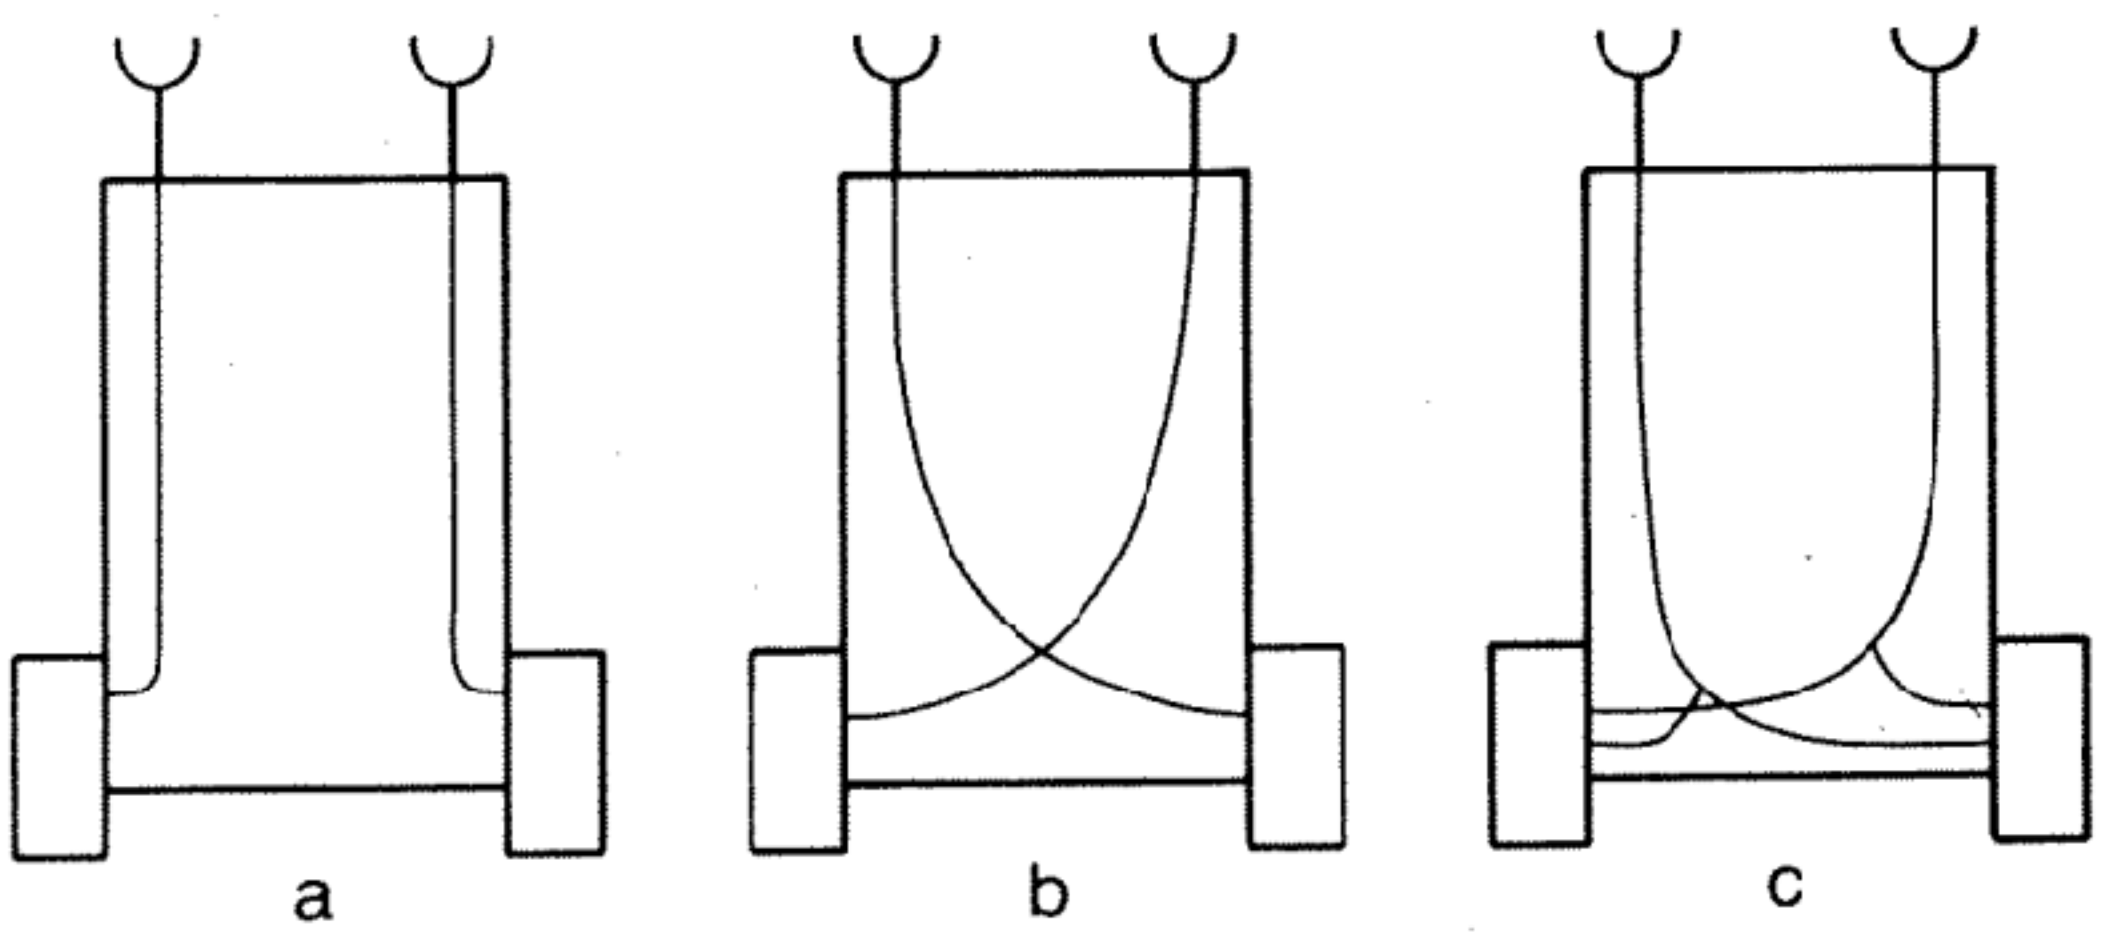

Vehicle 2

Action Items

- Join Piazza and Gradescope.
- Complete the background survey: https://forms.gle/d8hmnQGWQc9SMVcN6
- Begin the first programming assignment on control.
- Read on Bayes filter for next week; send a reading response by 12 pm on Monday.