

Script generated by TTT

Title: Sturm: Visual Navigation (22.05.2012)

Date: Tue May 22 10:17:32 CEST 2012

Duration: 89:20 min

Pages: 79

Visual Navigation for Flying Robots

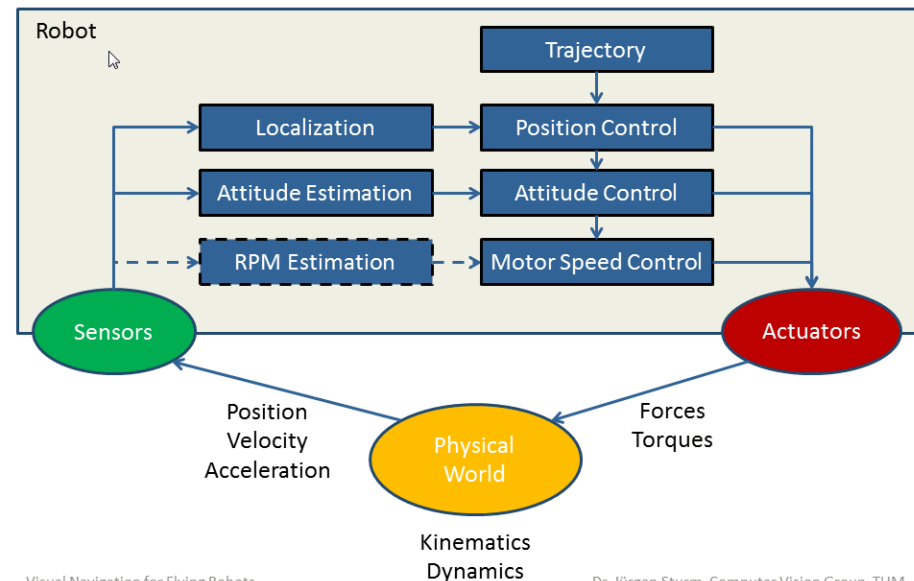
Robot Control

Dr. Jürgen Sturm

Organization - Exam



- Oral exams **in teams** (2-3 students)
- At least 15 minutes per student
→ individual grades
- Questions will address
 - Material from the lecture
 - Material from the exercise sheets
 - Your mini-project

Control Architecture



DC Motors



- Maybe you built one in school
- Stationary permanent magnet
- Electromagnet induces torque
- Split ring switches direction of current



Visual Navigation for Flying Robots

Brushless Motors


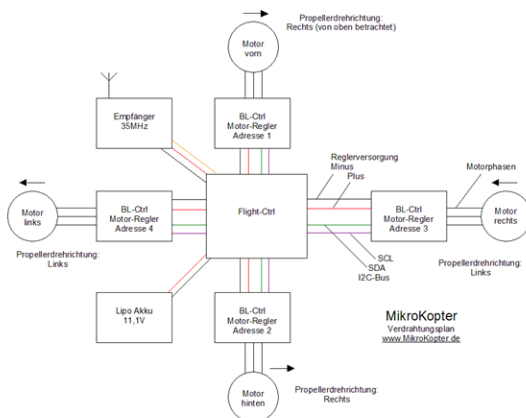
- Used in most quadcopters
 - Permanent magnets on the axis
 - Electromagnets on the outside
 - Requires motor controller to switch currents
- Does not require brushes (less maintenance)



Visual Navigation for Flying Robots

Attitude + Motor Controller Boards


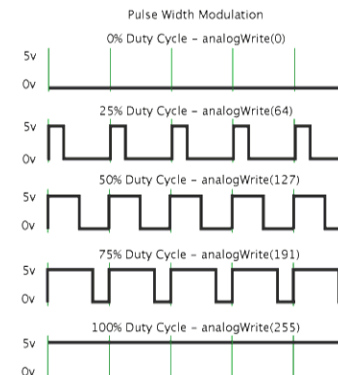
- Example: MikroKopter Platform



Visual Navigation for Flying Robots

Pulse Width Modulation (PWM)


- Protocol used to control motor speed
- Remote controls typically output PWM



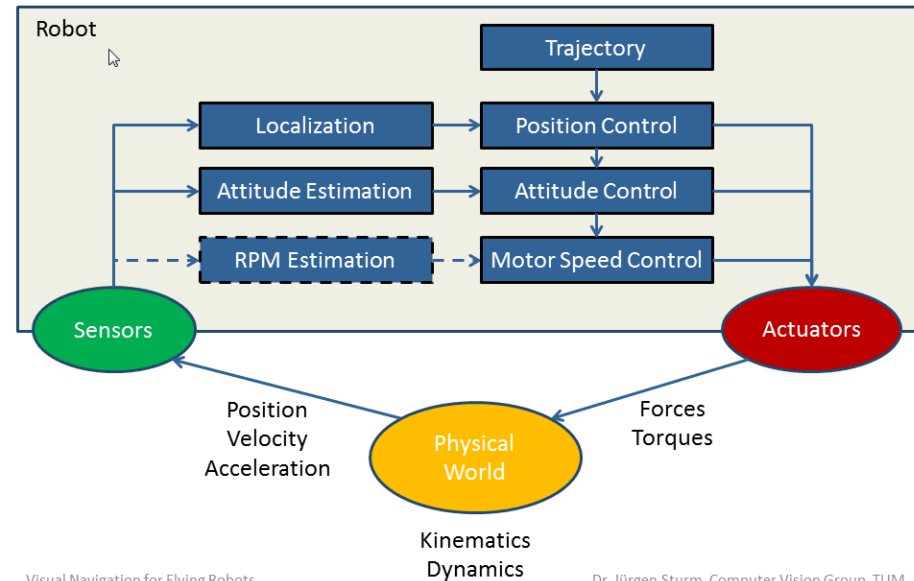
Visual Navigation for Flying Robots

I2C Protocol

- Serial data line (SDA) + serial clock line (SCL)
- All devices connected in parallel
- 7-10 bit address, 100-3400 kbit/s speed
- Used by Mikrocopter for motor control



Control Architecture



Kinematics and Dynamics

- Kinematics**
 - Integrate acceleration to get velocity
 - Integrate velocity to get position
- Dynamics**
 - Actuators induce forces and torques
 - Forces induce linear acceleration
 - Torques induce angular acceleration
- What types of forces do you know?
- What types of torques do you know?

Example: 1D Kinematics

- State** $\mathbf{x} = (x \ \dot{x} \ \ddot{x})^T \in \mathbb{R}^3$
- Action** $u \in \mathbb{R}$
- Process model**

$$\mathbf{x}_t = \begin{pmatrix} 1 & \Delta t & 0 \\ 0 & 1 & \Delta t \\ 0 & 0 & 1 \end{pmatrix} \mathbf{x}_{t-1} + \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} u_t$$

- Kalman filter
- How many states do we need for 3D?



Dynamics - Essential Equations

- Force (Kraft)

$$m\ddot{\mathbf{x}} = \sum_i F_i$$



- Torque (Drehmoment)

$$J\alpha = \sum_i \tau_i$$



Example: Spring-Damper System

- Combination of spring and damper
- Forces $F = F_{\text{damping}} + F_{\text{spring}}$
- Resulting dynamics $m\ddot{x} = D\dot{x} + K(x - x_{\text{eq}})$




Forces

- Gravity $F_{\text{grav}} = mg$
- Friction
 - Stiction (static friction) $F_{\text{stiction}} = c_s \text{sign } \dot{x}$
 - Damping (viscous friction) $F_{\text{damping}} = D\dot{x}$
- Spring $F_{\text{spring}} = K(x - x_{\text{eq}})$
- Magnetic force
- ...



Torques


- Definition $\tau = F \times \mathbf{r}$
- Torques sum up $\tau_{\text{net}} = \sum \tau_i$
- Torque results in angular acceleration $\tau = J\alpha$ (with $\alpha = \frac{d\omega}{dt}$, J moment of inertia)
- Friction same as before...





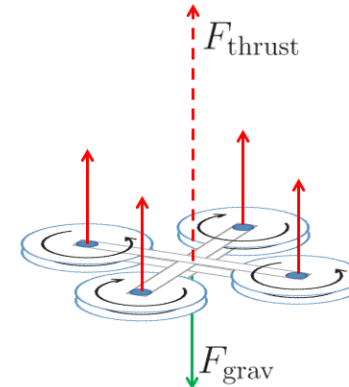
Dynamics of a Quadcopter

- Each propeller induces force and torque by accelerating air
- Gravity pulls quadcopter downwards




Vertical Acceleration

- Thrust $F_{\text{thrust}} = F_1 + F_2 + F_3 + F_4$




Vertical and Horizontal Acceleration

- Thrust $F_{\text{thrust}} = F_1 + F_2 + F_3 + F_4$
- Acceleration $\ddot{\mathbf{x}}_{\text{global}} = R_{RPY} F_{\text{thrust}} - F_{\text{grav}}$




Pitch (and Roll)

- Attitude changes when opposite motors generate unequal thrust
- Induced torque $\tau = (F_1 - F_3) \times \mathbf{r}$
- Induced angular acceleration




Torques

- Definition $\tau = F \times \mathbf{r}$
- Torques sum up $\tau_{\text{net}} = \sum \tau_i$
- Torque results in angular acceleration $\tau = J\alpha$ (with $\alpha = \frac{d\omega}{dt}$, J moment of inertia)
- Friction same as before...



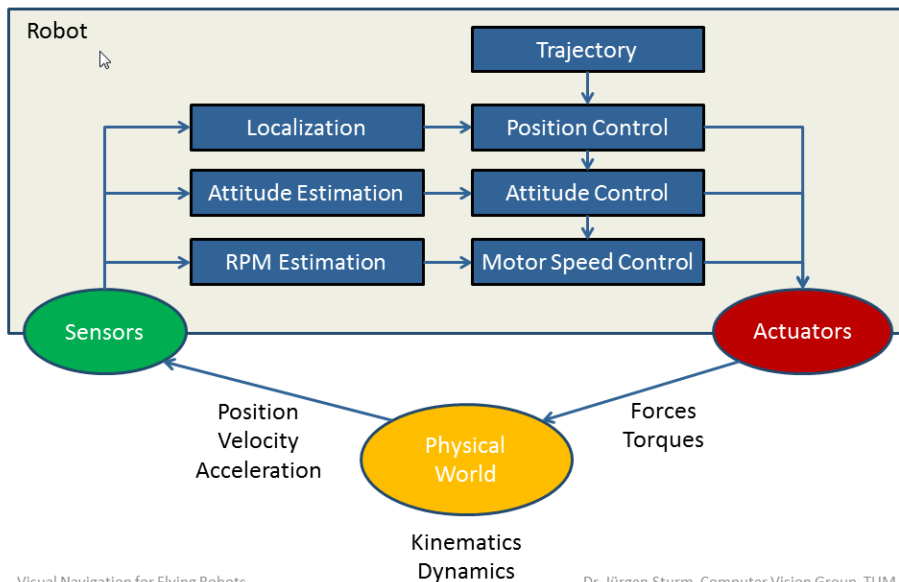
Pitch (and Roll)

- Attitude changes when opposite motors generate unequal thrust
- Induced torque $\tau = (F_1 - F_3) \times \mathbf{r}$
- Induced angular acceleration



Side view of quadcopter

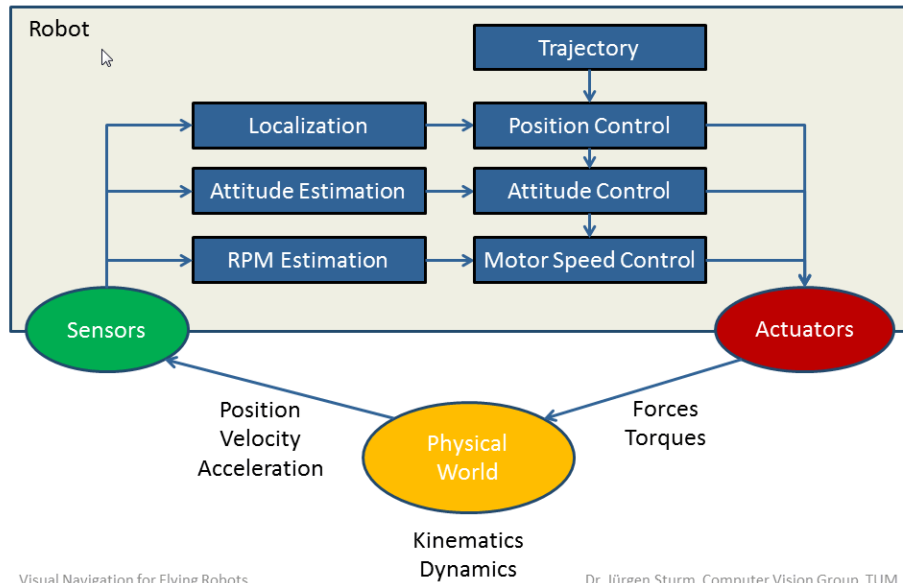
Cascaded Control



Assumptions of Cascaded Control

- Dynamics of inner loops is so fast that it is not visible from outer loops
- Dynamics of outer loops is so slow that it appears as static to the inner loops

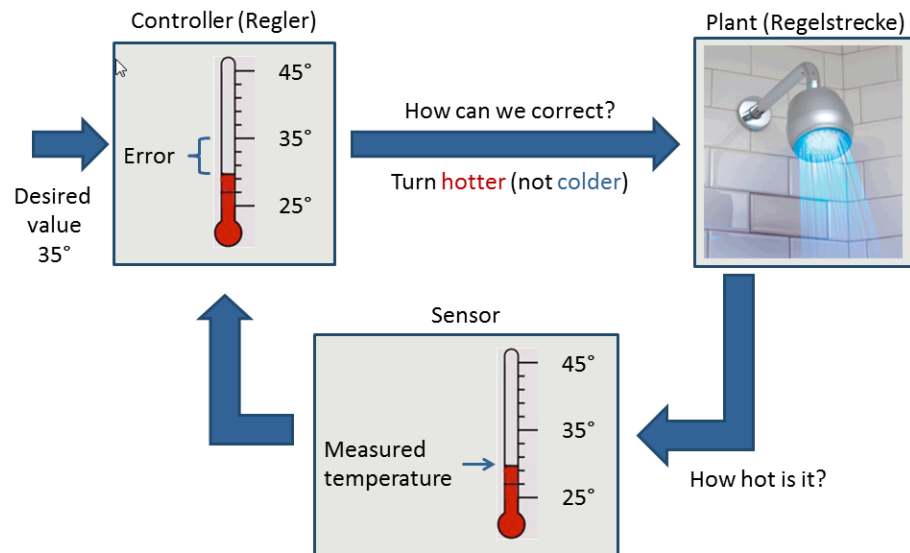
Cascaded Control



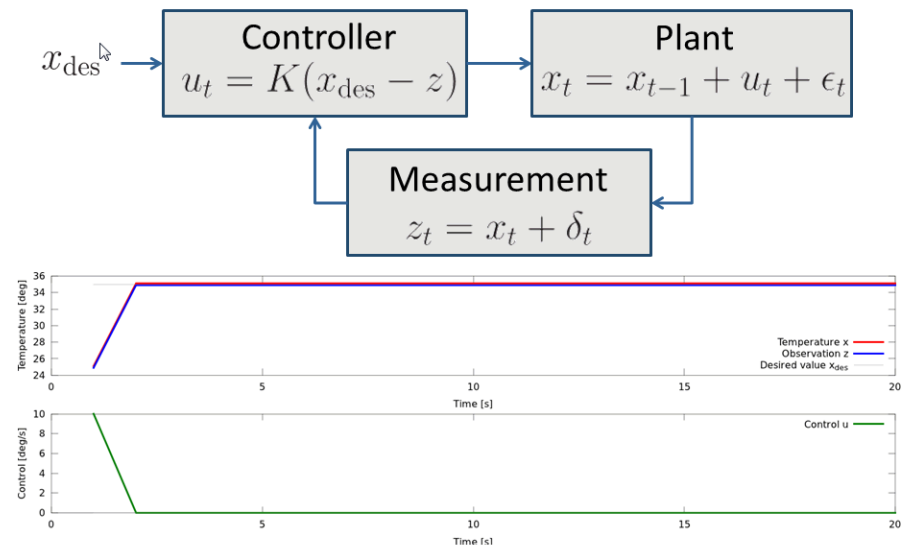
Cascaded Control Example

- Motor control happens on motor boards (controls every motor tick)
- Attitude control implemented on micro-controller with hard real-time (at 1000 Hz)
- Position control (at 10 – 250 Hz)
- Trajectory (waypoint) control (at 0.1 – 1 Hz)

Feedback Control - Generic Idea




Feedback Control - Example



Measurement Noise


- What effect has noise in the measurements?



- Poor performance for $K=1$
- How can we fix this?


Proper Control with Measurement Noise

- Lower the gain... ($K=0.15$)




What do High Gains do?

- High gains are always problematic ($K=2.15$)




What happens if sign is messed up?

- Check $K=-0.5$

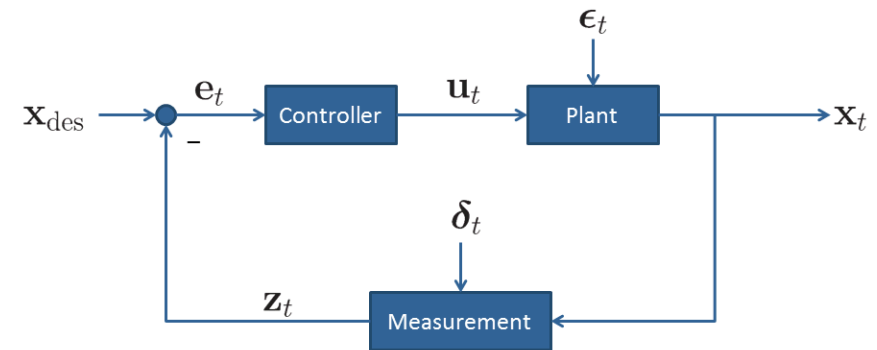


Saturation

- In practice, often the set of admissible controls u is bounded
- This is called (control) saturation




Block Diagram



Delays

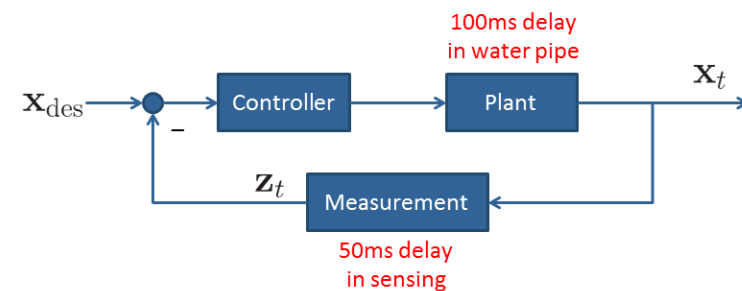
- In practice most systems have delays
- Can lead to overshoots/oscillations/de-stabilization



- One solution: lower gains (why is this bad?)

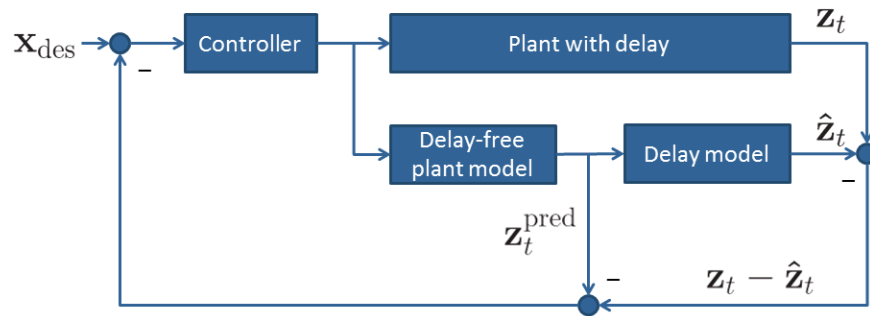
Delays

- What is the total dead time of this system?




Smith Predictor

- Allows for higher gains
- Requires (accurate) model of plant




Smith Predictor

- Plant model is available
- 5 seconds delay
- Results in perfect compensation
- Why is this unrealistic in practice?




Smith Predictor

- Time delay (and plant model) is often not known accurately (or changes over time)
- What happens if time delay is **over**estimated?

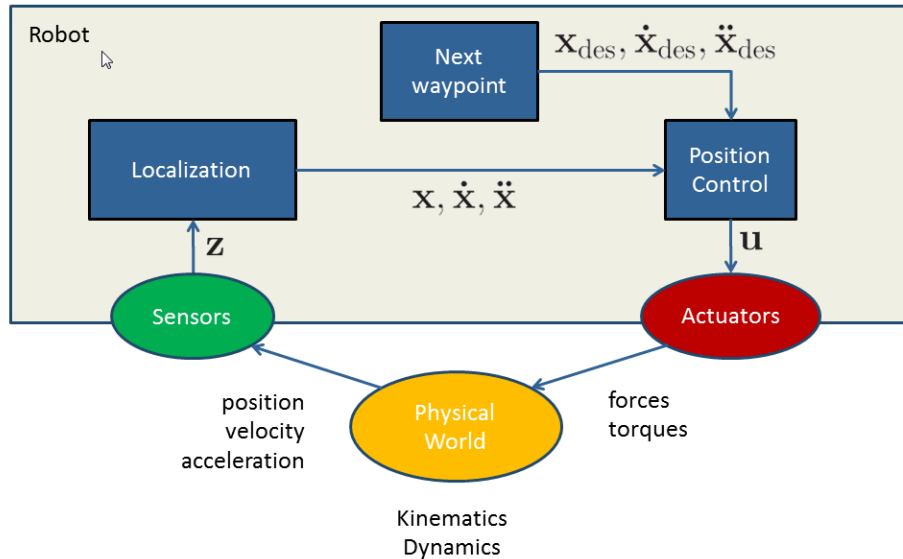


Smith Predictor

- Time delay (and plant model) is often not known accurately (or changes over time)
- What happens if time delay is **under**estimated?



Position Control




Rigid Body Kinematics

- Consider a rigid body
- Free floating in 1D space, no gravity
- How does this system evolve over time?


Rigid Body Kinematics

- Consider a rigid body
- Free floating in 1D space, no gravity
- How does this system evolve over time?
- Example: $x_0 = 0, \dot{x}_0 = 0$



Rigid Body Kinematics

- Consider a rigid body
- Free floating in 1D space, no gravity
- How does this system evolve over time?
- Example: $x_0 = 0, \dot{x}_0 = 0.1$



Rigid Body Kinematics


- Consider a rigid body
- Free floating in 1D space, no gravity
- In each time instant, we can apply a force F
- Results in acceleration $\ddot{x} = F/m$
- Desired position $x_{des} = 1$

P Control

- What happens for this control law?

$$u_t = K(x_{des} - x_{t-1})$$

- This is called proportional control




PD Control

- What happens for this control law?

$$u_t = K_P(x_{des} - x_{t-1}) + K_D(\dot{x}_{des} - \dot{x}_{t-1})$$

- Proportional-Derivative control




PD Control

- What happens for this control law?

$$u_t = K_P(x_{des} - x_{t-1}) + K_D(\dot{x}_{des} - \dot{x}_{t-1})$$

- What if we set **higher** gains?




PD Control

- What happens for this control law?


$$u_t = K_P(x_{\text{des}} - x_{t-1}) + K_D(\dot{x}_{\text{des}} - \dot{x}_{t-1})$$

- What if we set **lower** gains?



PD Control

- What happens when we add gravity?




Gravity compensation

- Add as an additional term in the control law


$$u_t = K_P(x_{\text{des}} - x_{t-1}) + K_D(\dot{x}_{\text{des}} - \dot{x}_{t-1}) + F_{\text{grav}}$$

- Any known (inverse) dynamics can be included



PD Control

- What happens when we have systematic errors? (noise with non-zero mean)
- Example: unbalanced quadcopter, wind, ...
- Does the robot ever reach its desired location?




PID Control

- Idea: Estimate the system error (bias) by integrating the error

$$u_t = K_P(x_{des} - x_t) + K_D(\dot{x}_{des} - \dot{x}_t) + K_I \int_{-\infty}^t x_{des} - x_t dt$$

- Proportional+Derivative+Integral Control



PID Control

- Idea: Estimate the system error (bias) by integrating the error

$$u_t = K_P(x_{des} - x_t) + K_D(\dot{x}_{des} - \dot{x}_t) + K_I \int_{-\infty}^t x_{des} - x_t dt$$

- Proportional+Derivative+Integral Control

- For steady state systems, this can be reasonable

- Otherwise, it may create havoc or even disaster (wind-up effect)

PID Control

- Idea: Estimate the system error (bias) by integrating the error

$$u_t = K_P(x_{des} - x_t) + K_D(\dot{x}_{des} - \dot{x}_t) + K_I \int_{-\infty}^t x_{des} - x_t dt$$

- Proportional+Derivative+Integral Control


- For steady state systems, this can be reasonable

- Otherwise, it may create havoc or even disaster (wind-up effect)

Example: Wind-up effect


- Quadrocopter gets stuck in a tree → does not reach steady state

- What is the effect on the I-term?




De-coupled Control

- So far, we considered only single-input, single-output systems (SISO)
- Real systems have multiple inputs + outputs
- MIMO (multiple-input, multiple-output)
- In practice, control is often de-coupled



How to Choose the Coefficients?


- Gains too large: overshooting, oscillations
- Gains too small: long time to converge
- Heuristic methods exist
- In practice, often tuned manually



Example: Ardrone

Cascaded control

- Inner loop runs on embedded PC and stabilizes flight
- Outer loop runs externally and implements position control



Ardrone: Inner Control Loop

- Plant input: motor torques

$$\mathbf{u}_{\text{inner}} = (\tau_1 \quad \tau_2 \quad \tau_3 \quad \tau_4)^\top$$

- Plant output: roll, pitch, yaw rate, z velocity

$$\mathbf{x}_{\text{inner}} = (\omega_x \quad \omega_y \quad \dot{\omega}_z \quad \dot{z})^\top$$

attitude
(measured using gyro +
accelerometer)
z velocity
(measured using ultrasonic
distance sensor + attitude)



Ardrone: Outer Control Loop

- Outer loop sees inner loop as a plant (black box)
- Plant input: roll, pitch, yaw rate, z velocity

$$\mathbf{u}_{\text{outer}} = (\omega_x \ \omega_y \ \dot{\omega}_z \ \dot{z})^T$$


- Plant output:

$$\mathbf{x}_{\text{outer}} = (x \ y \ z \ \psi)^T$$



Mechanical Equivalent

- PD Control is equivalent to adding spring-dampers between the desired values and the current position



PID Control – Summary

PID is the most used control technique in practice

- P control → simple proportional control, often enough
- PI control → can compensate for bias (e.g., wind)
- PD control → can be used to reduce overshoot (e.g., when acceleration is controlled)
- PID control → all of the above



Optimal Control

What other control techniques do exist?

- Linear-quadratic regulator (LQR)
- Reinforcement learning
- Inverse reinforcement learning
- ... and many more



Optimal Control

- Find the controller that provides the best performance
- Need to define a measure of performance
- What would be a good performance measure?
 - Minimize the error?
 - Minimize the controls?
 - Combination of both?



Reinforcement Learning

- In principle, any measure can be used
- Define reward for each state-action pair

$$R(x_t, u_t)$$
- Find the policy (controller) that maximizes the expected future reward
- Compute the expected future reward based on
 - Known process model
 - Learned process model (from demonstrations)



Linear Quadratic Regulator

Given:

- Discrete-time **linear** system

$$x_{k+1} = Ax_k + Bu_k$$

- **Quadratic** cost function

$$J = \sum_{k=0}^{\infty} (x_k^T Q x_k + u_k^T R u_k)$$

Goal: Find the controller with the lowest cost →
LQR control



Inverse Reinforcement Learning

- Parameterized reward function
- Learn these parameters from expert demonstrations and refine
- Example: [Abbeel and Ng, ICML 2010]





Reinforcement Learning

- In principle, any measure can be used
- Define reward for each state-action pair

$$R(x_t, u_t)$$
- Find the policy (controller) that maximizes the expected future reward
- Compute the expected future reward based on
 - Known process model
 - Learned process model (from demonstrations)



Interesting Papers at ICRA 2012

- Flying robots are a hot topic in the robotics community
- 4 out of 27 sessions on flying robots
- Robots: quadcopters, nano quadcopters, fixed-wing airplanes
- Sensors: monocular cameras, Kinect, motion capture, laser-scanners



Inverse Reinforcement Learning




- Parameterized reward function
- Learn these parameters from expert demonstrations and refine
- Example: [Abbeel and Ng, ICML 2010]



Autonomous Indoor 3D Exploration with a Micro-Aerial Vehicle

Shaojie Shen, Nathan Michael, and Vijay Kumar

- Map a previously unknown building
- Find good exploration frontiers in partial map






Decentralized Formation Control with Variable Shapes for Aerial Robots


Matthew Turpin, Nathan Michael, and Vijay Kumar

- Move in formation (e.g., to traverse a window)
- Avoid collisions
- Dynamic role switching



Visual Navigation for Flying Robots

71




Dr. Jürgen Sturm, Computer Vision Group, TUM



Versatile Distributed Pose Estimation and Sensor Self-Calibration for an Autonomous MAV


Stephan Weiss, Markus W. Achtelik, Margarita Chli, Roland Siegwart

- IMU, camera
- EKF for pose, velocity, sensor bias, scale, inter-sensor calibration



Visual Navigation for Flying Robots

72



Dr. Jürgen Sturm, Computer Vision Group, TUM



On-board Velocity Estimation and Closed-loop Control of a Quadrotor UAV based on Optical Flow


Volker Grabe, Heinrich H. Bühlhoff, and Paolo Robuffo Giordano

- Ego-motion from optical flow using homography constraint
- Use for velocity control



Visual Navigation for Flying Robots

73




Dr. Jürgen Sturm, Computer Vision Group, TUM



Autonomous Landing of a VTOL UAV on a Moving Platform Using Image-based Visual Servoing

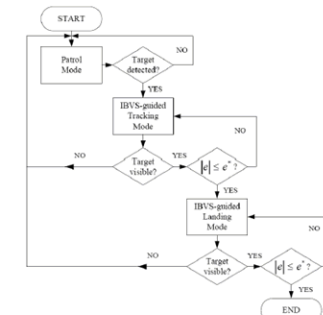
Daewon Lee, Tyler Ryan and H. Jin. Kim

- Tracking and landing on a moving platform
- Switch between tracking and landing behavior



Visual Navigation for Flying Robots

74




Dr. Jürgen Sturm, Computer Vision Group, TUM


Resonant Wireless Power Transfer to Ground Sensors from a UAV

Brent Griffin and Carrick Detweiler

- Quadcopter transfers power to light a LED



Visual Navigation for Flying Robots




75

Dr. Jürgen Sturm, Computer Vision Group, TUM

Using Depth in Visual Simultaneous Localisation and Mapping

Sebastian A. Scherer, Daniel Dube and Andreas Zell

- Combine PTAM with Kinect
- Monocular SLAM: scale drift
- Kinect: has small maximum range



Visual Navigation for Flying Robots

76

Dr. Jürgen Sturm, Computer Vision Group, TUM



ICRA Papers

- Will put them in our paper repository
- Remember password (or ask by mail)
- See course website

77

Visual Navigation for Flying Robots

77

Dr. Jürgen Sturm, Computer Vision Group, TUM



Ende der Bildschirmpräsentation. Zum Beenden klicken.

