AMPyC

Chapter 5: Linear Stochastic Optimal Control

30. May 2025

Authors: Prof. Melanie Zeilinger
Dr. Lukas Hewing
Dr. Andrea Carron

ETH Zurich

Stochastic Optimal Control

Optimal control problem that we ideally would like to solve

$$
\begin{aligned}
J^*(x) = \min_{\{\pi_k\}} \quad & \mathbb{E}\left(L(X, U)\right) \\
\text{s.t.} \quad & x(k+1) = f_k(x(k), u(k), w(k), \theta), \\
& u(k) = \pi_k(x(0), \dots, x(k), u(0) \dots, u(k-1)), \\
& W \sim \mathcal{Q}^W, \ \theta \sim \mathcal{Q}^\theta, \\
& Pr\left(x(k) \in \mathcal{X}, u(k) \in \mathcal{U}\right) \geq p, \\
& x(0) = x
\end{aligned}
$$

- State sequence $X = [x(0)^\mathsf{T}, \dots, x(\bar{N}-1)^\mathsf{T}]^\mathsf{T}$
- Input sequence $U = [u(0)^\mathsf{T}, \dots, u(\bar{N}-1)^\mathsf{T}]^\mathsf{T}$
- Disturbance sequence $W = [w(0)^\mathsf{T}, \dots, w(\bar{N}-1)^\mathsf{T}]^\mathsf{T}$
- Objective function $L(X, U)$
- Constraints $\mathcal{X}$, $\mathcal{U}$
- Uncertainties $\theta \sim \mathcal{Q}^\theta$, $W \sim \mathcal{Q}^W$,

Graphical Outline

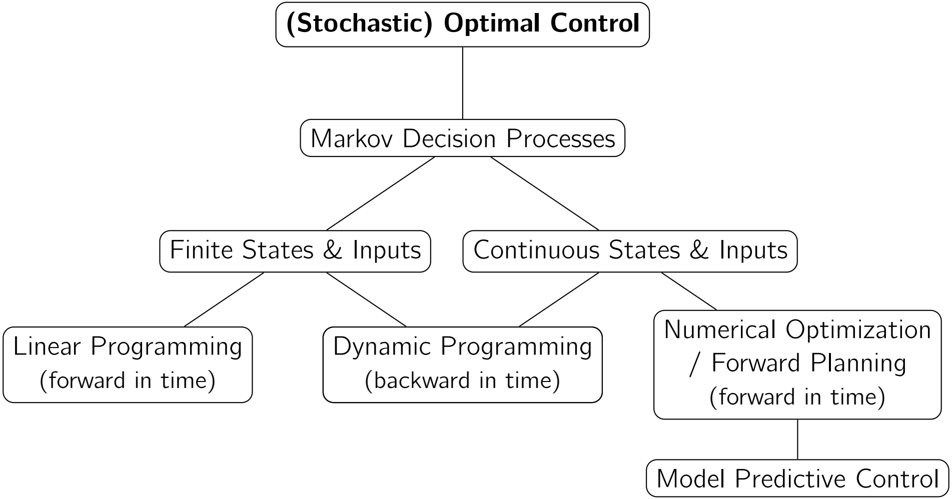

Graphical Outline

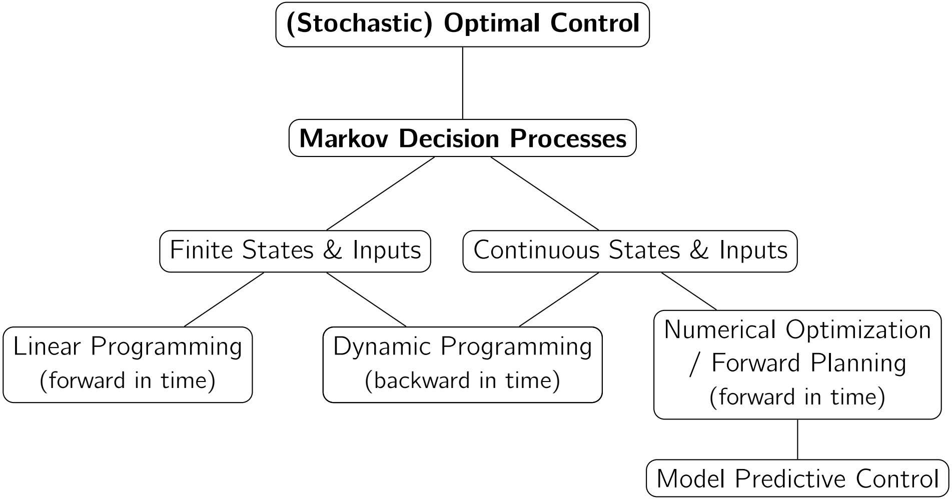

Assumptions for Markov Decision Process

To be a Markov Decision Process (MDP), the problem needs to decompose into different **stages**

- Independent disturbances $w(k) \sim Q^{w_k}$, no parametric uncertainty θ
  (Ensures that distribution of $x(k+1)$ is fully determined given $u(k)$ and $x(k)$ (Markov Property))
- Stage cost $L(X, U) = \sum l_k(x, u)$

$$
\begin{aligned}
J^*(x) = \quad \min_{\{\pi_k\}} \quad & \mathbb{E}\left(\sum_{k=0}^{\bar{N}} l_k(x, u)\right) \\
\text{s.t.} \quad & x(k+1) = f_k(x(k), u(k), w(k)), \\
& u(k) = \pi_k(x(k)), \\
& w(k) \sim \mathcal{Q}^{w_k}, \text{ independent}, \\
& x(0) = x
\end{aligned}
$$

Dynamic Programming: Principle of Optimality

Let $\Pi^* = \{\pi_0, \dots, \pi_{\bar{N}}\}$ be an optimal policy for the stochastic optimal control problem. Consider the subproblem whereby we are at $x(k)$ at time k and wish to minimize the **"cost-to-go"** from time k to time $\bar{N}$

$$\mathbb{E}\left(\sum_{i=k}^{\bar{N}} l_i(x(i), u(i)) \,\middle|\, x(k)\right)$$

Then the truncated policy $\{\pi_k, \dots, \pi_{\bar{N}}\}$ is optimal for this subproblem[1].

Interpretation

Suppose that the fastest route from Los Angeles to Boston passes through Chicago. Then the principle of optimality formalizes the fact that the Chicago to Boston portion of the route is also the fastest route for a trip that starts from Chicago and ends in Boston.

[1] Adapted from **Bertsekas, "Dynamic Programming & Optimal Control"**

Graphical Outline

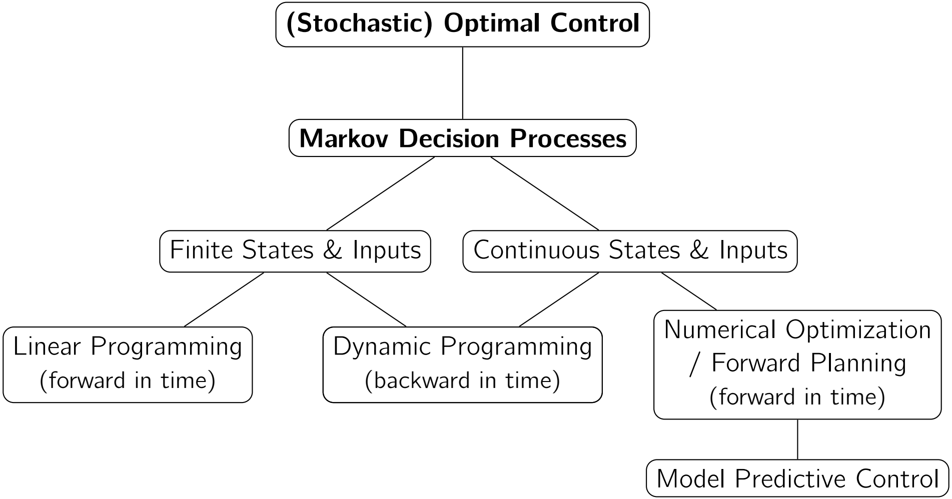

Graphical Outline

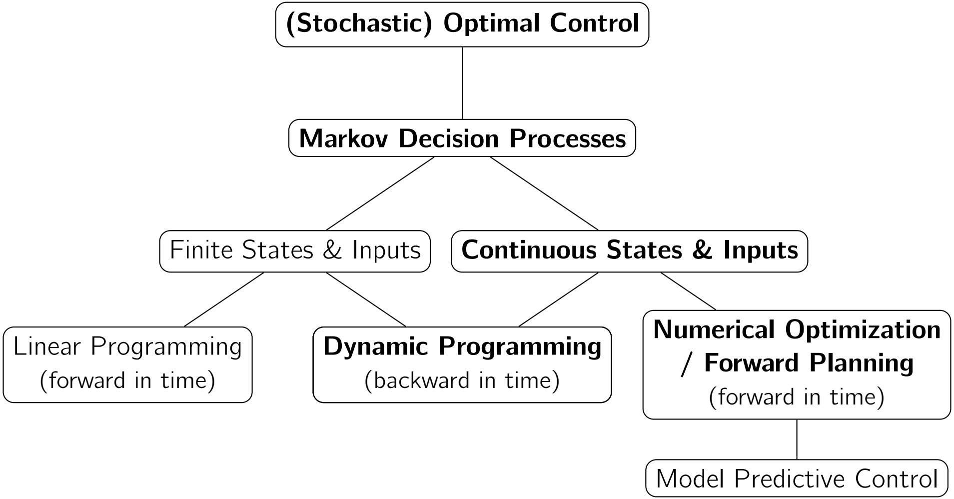

Contents – 5: Linear Stochastic Optimal Control

- Derive solutions for continuous state-input space MDPs

$\rightarrow$ backward in time using dynamic programming
(linear-quadratic & unconstrained, optimal)

$\rightarrow$ forward in time using forward planning
(potentially constrained, approximate)

- Understand strenghts and weaknesses of the different solution approachens
  - Dynamic Programming: More narrow problem class, optimal and computationally cheap
  - Forward Planning: Wider problem class, often approximate and computationally more expensive

- Appreciate the potential for receding horizon control, i.e. MPC

Outline

1. Problem Formulation: Linear Quadratic Control
2. Dynamic Programming Solution (backwards in time)
3. Forward Planning Solution (forwards in time)
4. Constrained Stochastic Optimal Control Problems
5. Outlook Model Predictive Control

Outline

1. Problem Formulation: Linear Quadratic Control

2. Dynamic Programming Solution (backwards in time)

3. Forward Planning Solution (forwards in time)

4. Constrained Stochastic Optimal Control Problems

5. Outlook Model Predictive Control

Linear Quadratic Control: Dynamics

$$x(k+1) = A_k x(k) + B_k u(k) + w(k),\ \ k = 0, \dots, \bar{N} - 1$$

- $w(k)$ are random disturbances (process noise) at time step k
- all $w(k)$ are independent and distributed with

$$\mathbb{E}(w(k)) = \mu_w(k),\ \operatorname{var}(w(k)) = \Sigma_w(k)$$

Many systems can be (approximately) represented in this form, enabling high-performance control

- Random walk
- (Wind) Turbines
- Thermal Building Control
- Quadrotors, airplanes, rockets, ...

Linear Quadratic Control: Dynamics

$$x(k+1) = \mathbf{A}x(k) + \mathbf{B}u(k) + w(k),\ \ k = 0, \dots, \bar{N} - 1$$

- $w(k)$ are random disturbances (process noise) at time step k
- all $w(k)$ are **zero-mean independent and identically distributed (i.i.d.)** with

$$\mathbb{E}(w(k)) = \mathbf{0},\ \ \mathrm{var}(w(k)) = \mathbf{\Sigma_w}$$

Many systems can be (approximately) represented in this form, enabling high-performance control

- Random walk
- (Wind) Turbines
- Thermal Building Control
- Quadrotors, airplanes, rockets, ...

Linear Quadratic Control: Dynamics

$$x(k+1) = \mathbf{A}x(k) + \mathbf{B}u(k) + w(k),\ \ k = 0, \ldots, \bar{N} - 1$$

- $w(k)$ are random disturbances (process noise) at time step k
- all $w(k)$ are **zero-mean independent and identically distributed (i.i.d.)** with

$$\mathbb{E}(w(k)) = \mathbf{0},\ \ \mathrm{var}(w(k)) = \mathbf{\Sigma_w}$$

Many systems can be (approximately) represented in this form, enabling high-performance control

- Random walk
- (Wind) Turbines
- Thermal Building Control
- Quadrotors, airplanes, rockets, ...

We focus on linear **time-invariant** (LTI) systems; extension to time-varying possible

Segway: Illustrative Example

States: Position $p \in \mathbb{R}$
Velocity $v \in \mathbb{R}$

Input: Acceleration $a \in \mathbb{R}$

Disturbances $w \sim \mathcal{N}(0, 0.1)$

Dynamics
$$p(k+1) = p(k) + v(k)$$
$$v(k+1) = v(k) + a(k) + w(k)$$

[freepik.com]

Linear time invariant (LTI) system: $x(k+1) = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} x(k) + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u(k) + \begin{bmatrix} 0 \\ 1 \end{bmatrix} w(k)$

Segway: Heuristic Policies

Policy:

1. Do nothing: $\pi_k(x) = 0$

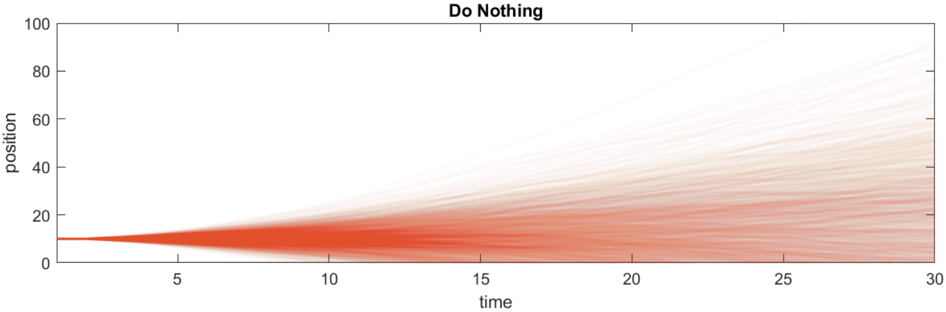

Segway: Heuristic Policies

Policy:

1. Do nothing: $\pi_k(x) = 0$
2. Full throttle: $\pi_k(x) = 1$

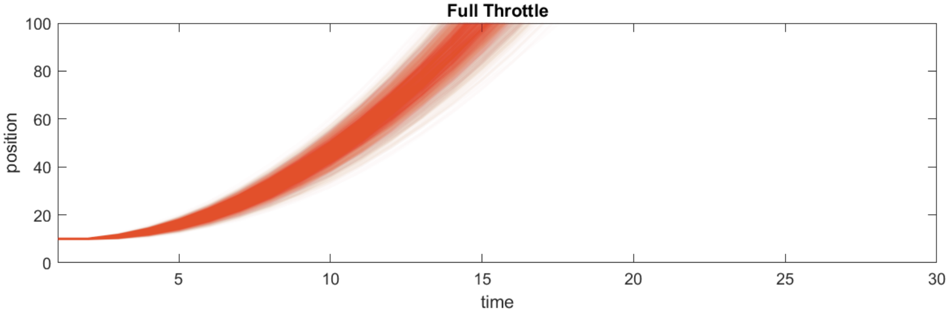

Segway: Heuristic Policies

Policy:

1. Do nothing: $\pi_k(x) = 0$
2. Full throttle: $\pi_k(x) = 1$
3. Hold velocity: $\pi_k(x) = \begin{cases} 1 & \text{if } v < 3 \\ 0 & \text{if } v = 3 \\ -1 & \text{if } v > 3 \end{cases}$

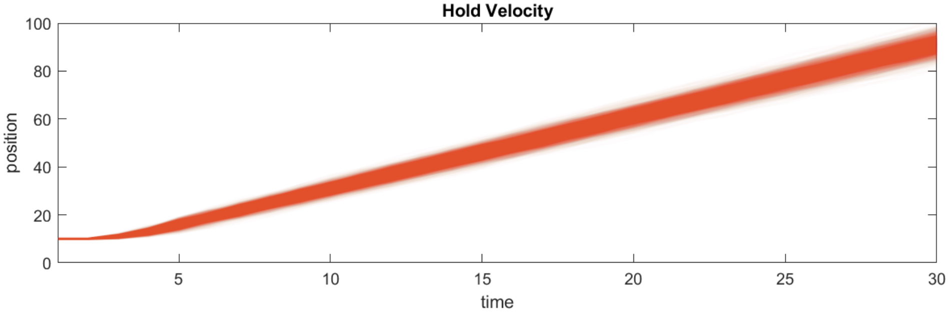

Segway: Heuristic Policies

Policy:

1. Do nothing: $\pi_k(x) = 0$
2. Full throttle: $\pi_k(x) = 1$
3. Hold velocity: $\pi_k(x) = \begin{cases} 1 & \text{if } v < 3 \\ 0 & \text{if } v = 3 \\ -1 & \text{if } v > 3 \end{cases}$
4. Hold position: $\pi_k(x) = \begin{cases} 1 & \text{if } p < 50 \\ 0 & \text{if } p = 50 \\ -1 & \text{if } p > 50 \end{cases}$

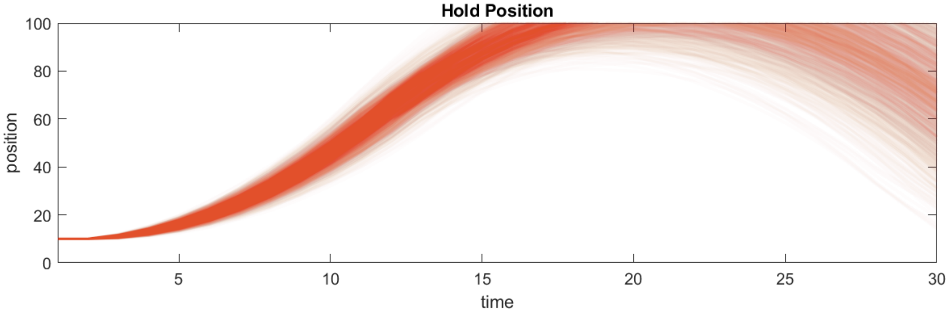

Linear Quadratic Control: Cost

$$J = \mathbb{E}\left(\|x(\bar{N})\|_P^2 + \sum_{k=0}^{\bar{N}-1} \|x(k)\|_Q^2 + \|u(k)\|_R^2 \right)$$

- Positive (semi-)definite cost matrices P, $Q \geq 0$, $R > 0$
- Can be generalized to time-varying cost matrices and additional linear terms (e.g. reference tracking)

Objective

Find control policies $u(k) = \pi_k(x(k))$ which minimize the expected cost J subject to dynamics $x(k+1) = Ax(k) + Bu(k) + w(k)$ and independent and identically distributed disturbances $w(k)$ with zero mean and known variance.

Resulting Formulation

$$
\begin{aligned}
J^*(x) = \min_{\{\pi_k\}} \quad & \mathbb{E}\left(\|x(\bar{N})\|_P^2 + \sum_{k=0}^{\bar{N}-1} \|x(k)\|_Q^2 + \|u(k)\|_R^2 \right) \\
\text{s.t.} \quad & x(k+1) = Ax(k) + Bu(k) + w(k), \\
& u(k) = \pi_k(x(k)), \\
& \mathbb{E}(w(k)) = 0, \ \mathrm{var}(w(k)) = \Sigma_w, \ \text{i.i.d.}, \\
& x(0) = x
\end{aligned}
$$

Exact solution possible by Dynamic Programming (LQR)

Outline

DP Solution

Let $J_k^*(x)$ be the optimal cost to go from time k and state $x(k) = x$

$$J_k^*(x) = \min_{\{\pi_i(\cdot)\}} \mathbb{E}\left(\|x(\bar{N})\|_P^2 + \sum_{i=k}^{\bar{N}-1} \|x(i)\|_Q^2 + \|u(i)\|_R^2 \,\middle|\, x(k) = x \right)$$

then we have

- optimal cost at time step 0: $J_0^*(x) = J^*(x)$ (desired)
- optimal cost at time step $\bar{N}$: $J_{\bar{N}}^*(x) = \mathbb{E}(\|x(\bar{N})\|_P^2 \,|\, x(\bar{N}) = x) = \|x\|_P^2$ (tractable initialization)

$\boxed{\rightarrow \text{ Backwards recursion: successive solution of 1-step problems}}$

DP Solution: Backwards Recursion I

- 1-step cost to go:

$$\begin{aligned} J^*_{\bar{N}-1}(x) &= \min_u \mathbb{E}\left(J^*_{\bar{N}}(x(\bar{N})) + \|x(\bar{N}-1)\|^2_Q + \|u\|^2_R \,\middle|\, x(\bar{N}-1) = x\right) \\ &= \min_u \mathbb{E}_w\left(\|Ax + Bu + w\|^2_P + \|x\|^2_Q + \|u\|^2_R\right) \end{aligned}$$

- Expectation over quadratic function:

$$\mathbb{E}_w\left(\|Ax + Bu + w\|^2_P\right) = \|Ax + Bu\|^2_P + \mathrm{tr}(P\Sigma_w), \text{ where } \Sigma_w := \mathrm{var}(w)$$

- Equivalent to deterministic problem (only offset by constant $C = \mathrm{tr}(P\Sigma_w)$):

$$J^*_{\bar{N}-1}(x) = \min_u (\underbrace{\|Ax + Bu\|^2_P + C}_{J^*_{\bar{N}}(x^+)} + \|x\|^2_Q + \|u\|^2_R)$$

- Unconstrained QP, minimized by $u = \underbrace{-(B^\top PB + R)^{-1}B^\top PA}_{:=K_{\bar{N}-1}}\, x$

DP Solution: Backwards Recursion II

Observation: Cost-to-go is quadratic, only depends on x

$$J^*_{\bar{N}-1}(x) = \min_u (\underbrace{\|Ax + Bu\|_P^2 + C}_{J^*_{\bar{N}}(x^+)} + \|x\|_Q^2 + \|u\|_R^2)$$

$$\text{inserting } u^* = K_{\bar{N}-1}x: \quad = \|(A + BK_{\bar{N}-1})x\|_P^2 + C + \|x\|_Q^2 + \|K_{\bar{N}-1}x\|_R^2 := x^\top P_{\bar{N}-1}x + c_{\bar{N}-1},$$

where $P_{\bar{N}-1} = (A + BK_{\bar{N}-1})^\top P(A + BK_{\bar{N}-1}) + Q + (K_{\bar{N}-1})^\top RK_{\bar{N}-1}$, $c_{\bar{N}-1} = C = \mathrm{tr}(P\Sigma_w)$

After substituting K and rearranging terms, this leads to Riccati backwards recursion

$$\begin{aligned}
P_{i-1} &= A^\top P_i A - A^\top P_i B(B^\top P_i B + R)^{-1} B^\top P_i A + Q \\
c_{i-1} &= c_i + \mathrm{tr}(\Sigma_w P_i) \\
K_i &= -(B^\top P_i B + R)^{-1} B^\top P_i A
\end{aligned}$$

with $P_{\bar{N}} = P$, $c_{\bar{N}} = 0$.

DP Solution: Optimal Cost

Following this recursion leads to solution for optimal cost

$$J^*(x) = J_0^*(x) = x^\top P_0 x + c_0$$

- Quadratic function in x
- Note that $x^\top P_0 x$ is independent of Σ_w (cost due to initial condition) while c_0 is independent of x (cost due to disturbances)
- (This decomposition is unique for linear systems under quadratic cost)

Equivalency to Deterministic Case

The resulting control law is identical to the deterministic LQR, i.e. the control law is independent of the disturbance intensity Σ_w

Segway: Illustrative Example

States: Position $p \in \mathbb{R}$
Velocity $v \in \mathbb{R}$

Input: Acceleration $a \in \mathbb{R}$

Disturbance: $w \sim \mathcal{N}(0, 0.1)$

Dynamic: $p(k+1) = p(k) + v(k)$
$v(k+1) = v(k) + a(k) + w(k)$

[freepik.com]

Cost
- Define $x := [p - 50,\ v]^\mathsf{T}$ (shift does not affect dynamics)
- $l(x, u) = \|x\|^2 + 1000\|u\|^2$

$\rightarrow Q = P = I,\ R = 1000$

Segway: DP Solution

Policy: Dynamic programming with stage cost

$$l_k(x, u) = \|x\|^2 + 1000\|u\|^2$$

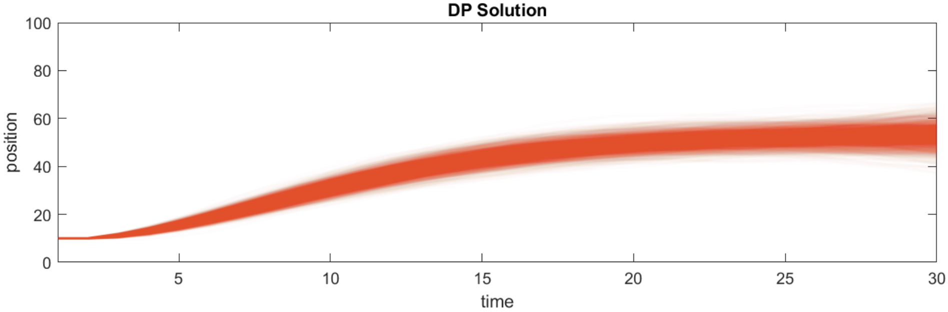

Segway: DP Solution

Policy: Dynamic programming with stage cost

$$l_k(x, u) = \|x\|^2 + 1000\|u\|^2$$

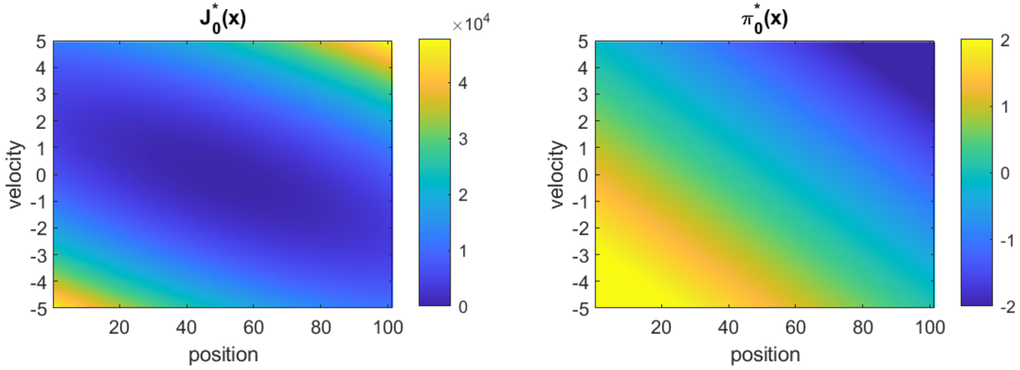

Outline

1. Problem Formulation: Linear Quadratic Control
2. Dynamic Programming Solution (backwards in time)
3. Forward Planning Solution (forwards in time)
4. Constrained Stochastic Optimal Control Problems
5. Outlook Model Predictive Control

Motivation & Idea

Motivation: (Approximately) solve more challenging problem classes

- Constraints
- Nonlinear dynamics
- Non-quadratic objectives

Idea: Construct state-input distribution forward in time and optimize directly.

Problem: Finite dimensional description of policies and distributions not available.

Approach: Approximate solution over restrained policy class

Most common in MPC:

$$\text{(Open-loop) input sequences}: \{\pi_k(x)\} = \{u_k\}$$
$$\text{Affine policies}: \{\pi_k(x)\} = \{K_k x + v_k\}$$

Motivation & Idea

Motivation: (Approximately) solve more challenging problem classes

- **Constraints**
- Nonlinear dynamics
- Non-quadratic objectives

Idea: Construct state-input distribution forward in time and optimize directly.

Problem: Finite dimensional description of policies and distributions not available.

Approach: Approximate solution over restrained policy class

Most common in MPC:

$$\begin{aligned} \text{(Open-loop) input sequences} : \{\pi_k(x)\} &= \{u_k\} \\ \text{Affine policies} : \{\pi_k(x)\} &= \{K_k x + v_k\} \end{aligned}$$

Mean-Variance Prediction Dynamics

$$x(k+1) = Ax(k) + Bu(k) + w(k)$$

All $w(k)$ are independent identically distributed (i.i.d.) with $\mathbb{E}(w(k)) = 0$, $\text{var}(w(k)) = \Sigma_w$.

Under affine control policy $u(k) = \pi_k(x) = K_k x + v_k$ predicted mean and variance is

$$\mathbb{E}(x(k+1)) = (A + BK_k)\mathbb{E}(x(k)) + Bv_k$$
$$\text{var}(x(k+1)) = (A + BK_k)\,\text{var}(x(k))(A + BK_k)^T + \Sigma^w$$

- in other cases (nonlinear, non-i.i.d.) propagation very challenging
- if $w(k)$ is Gaussian, so is $x(k)$ and $u(k)$

$\rightarrow$ finite and very compact description of resulting state distributions

Forward Optimization Problem: Open-loop Sequence

Simplest case: Open-loop input sequences $\{\pi_k(x)\} = \{u_k\}$

$$
\begin{aligned}
J^*_{\text{seq}}(x) = \min_{\{u_k\}} \quad & \mathbb{E}\left(\|x(\bar{N})\|_P^2 + \sum_{k=0}^{\bar{N}-1} \|x(k)\|_Q^2 + \|u(k)\|_R^2 \right) \\
\text{s.t.} \quad & x(k+1) = Ax(k) + Bu(k) + w(k), \\
& u(k) = u_k, \\
& w(k) \sim \mathcal{N}(0, \Sigma_w), \text{ i.i.d.}, \\
& x(0) = x
\end{aligned}
$$

Forward Optimization Problem: Open-loop Sequence

Simplest case: Open-loop input sequences $\{\pi_k(x)\} = \{u_k\}$

$$
\begin{aligned}
\tilde{J}^*_{\text{seq}}(x) = \min_{\{u_k\}} \quad & \left\| \mathbb{E}(x(\bar{N})) \right\|_P^2 + \sum_{k=0}^{\bar{N}-1} \|\mathbb{E}(x(k))\|_Q^2 + \|u_k\|_R^2 \\
\text{s.t.} \quad & \mathbb{E}(x(k+1)) = A\mathbb{E}(x(k)) + Bu_k, \\
& \mathbb{E}(x(0)) = x
\end{aligned}
$$

- Equivalent to deterministic optimization problem over $\mathbb{E}(x(k))$, independent of Σ_W
- Quadratic Program (convex, efficiently solvable)

Segway: Open-loop Sequence

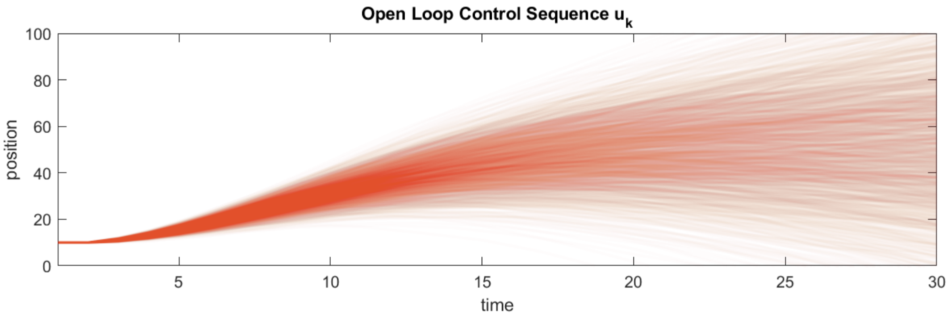

Problem: Without state feedback, trajectories diverge rapidly over time
Idea: 'Stabilize' trajectories using linear control gain $\pi_k(x)$
But... optimization over K_k typically non-convex[a] and computationally challenging

[a] Convex reformulations exist, e.g. disturbance feedback

Predefined Tube Controllers

Common Solution: Use **predefined** feedback gain K and only optimize over $\bar{u}_k$ of policy

$$\pi_k(x) = K(x - \bar{x}_k) + \bar{u}_k$$

with $\bar{x} := \mathbb{E}(x(k))$ resulting in $\bar{u}_k = \mathbb{E}(u(k))$

Resulting expectation dynamics:

$$\begin{aligned}\mathbb{E}(x(k+1)) &= \mathbb{E}(Ax(k) + B(K(x(k) - \bar{x}_k) + \bar{u}_k) + w(k)) \\ &= A\mathbb{E}(x(k)) + B(K(\underbrace{\mathbb{E}(x(k)) - \bar{x}_k}_{=0}) + \bar{u}_k) \\ &= A\mathbb{E}(x(k)) + B\bar{u}_k\end{aligned}$$

$$\boxed{\text{Open-loop optimization problem remains unchanged}}$$

Segway: Open-Loop Sequence with Predefined Tube Controller

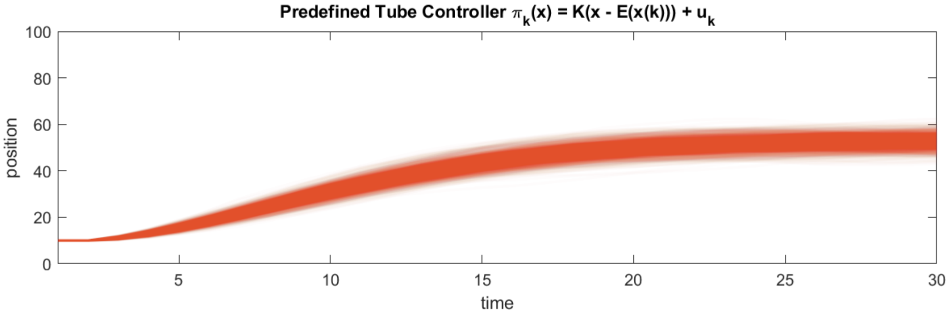

- Related to typical split into 'trajectory generation' and 'tracking'
- Can have very good performance for 'good' choice of K
- Here (infinite horizon LQR as tube controller):

$$J^*(x(0)) = 15081 \text{ vs. } J^*_{\text{pre}}(x(0)) = 15195 \text{ vs. } J^*_{\text{seq}}(x(0)) = 21246$$

Outline

1. Problem Formulation: Linear Quadratic Control
2. Dynamic Programming Solution (backwards in time)
3. Forward Planning Solution (forwards in time)
4. Constrained Stochastic Optimal Control Problems
5. Outlook Model Predictive Control

Forward Optimization Problem

$$\begin{aligned}
\tilde{J}^*_{\text{seq}}(x) = \min_{\{\bar{u}_k\}} \quad & \left\|\mathbb{E}(x(\bar{N}))\right\|_P^2 + \sum_{k=0}^{\bar{N}} \|\mathbb{E}(x(k))\|_Q^2 + \|\bar{u}_k\|_R^2 \\
\text{s.t.} \quad & \mathbb{E}(x(k+1)) = A\mathbb{E}(x(k)) + B\bar{u}_k, \\
& \mathbb{E}(x(0)) = x
\end{aligned}$$

- Deterministic (unconstrained) QP in $\bar{u}_k$ and $\mathbb{E}(x(k))$
- Straightforward to add (convex/polytopic) constraints to any of these variables

Segway Example

"Avoid velocities > 2.5": enforce expected value constraint: $\begin{bmatrix} 0 & 1 \end{bmatrix} \mathbb{E}(x(k)) \leq 2$

Segway: Open-Loop Sequence with Predefined Tube Controller

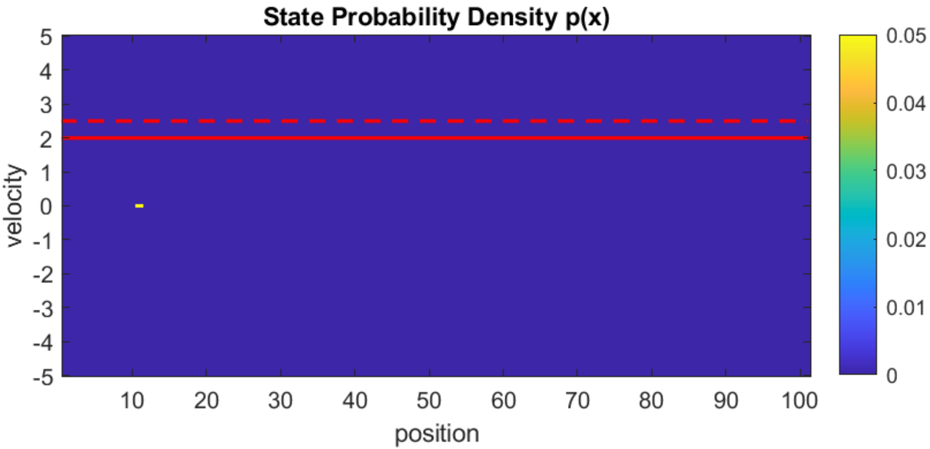

unconstrained

Segway: Open-Loop Sequence with Predefined Tube Controller

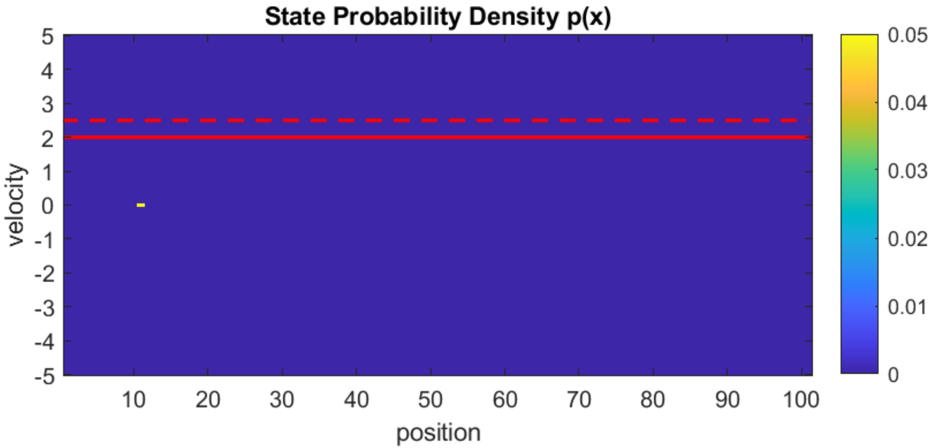

Added constraint: $\begin{bmatrix} 0 & 1 \end{bmatrix} \mathbb{E}(x(k)) \leq 2$ (red line)

Segway: Open-Loop Sequence with Predefined Tube Controller

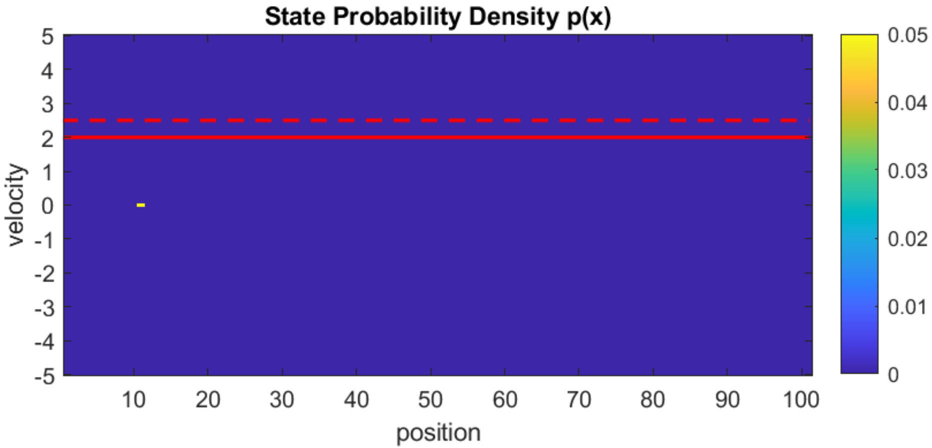

What is resulting probability $\Pr(\begin{bmatrix} 0 & 1 \end{bmatrix} x(k) \leq 2.5)$? (dashed red line)

Segway: Open-Loop Sequence with Predefined Tube Controller

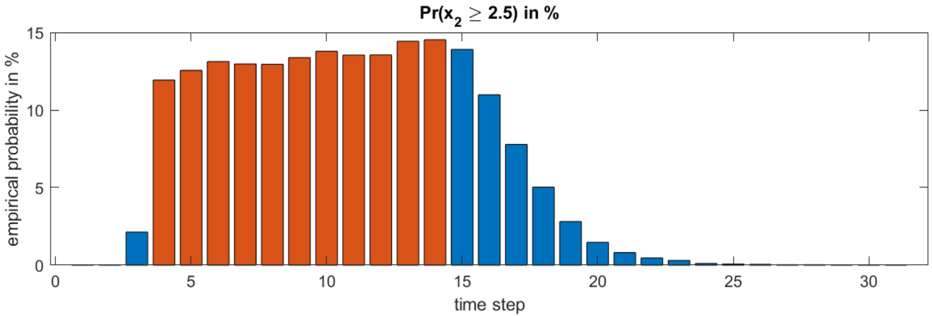

- 20.000 runs
- red: time step with active constraint

No direct handle on violation probability

Chance Constrained Linear Quadratic Optimal Control

$$
\begin{aligned}
J^*(x) = \min_{\{\pi_k\}} \quad & \mathbb{E}\left(\|x(\bar{N})\|_P^2 + \sum_{k=0}^{\bar{N}} \|x(k)\|_Q^2 + \|u(k)\|_R^2 \right) \\
\text{s.t.} \quad & x(k+1) = Ax(k) + Bu(k) + w(k), \\
& u(k) = \pi_k(x(k)), \\
& w(k) \sim \mathcal{N}(0, \Sigma_w), \\
& \mathsf{Pr}(x(k) \in \mathcal{X}) \geq p, \\
& x(0) = x
\end{aligned}
$$

Objective

Find control policies $u(k) = \pi_k(x(k))$ that minimize J^* subject to the chance constraints

Affine Policies & Gaussian Disturbances

- Affine feedback policies $\pi_k(x) = K_k x + v_k$, (or $\pi_k(x) = K_k(x - \bar{x}_k) + \bar{u}_k$)

$\rightarrow$ The closed-loop dynamics is linear/affine.

- Deterministic initial condition[2] & Gaussian i.i.d. disturbances $w(k)$

$\rightarrow$ State distribution is **Gaussian** in each time step

$$x(k) \sim \mathcal{N}(\mathbb{E}(x(k)), \mathrm{var}(x(k)))$$

Evaluating $\Pr(x(k) \in \mathcal{X})$ requires integrating Gaussian distributions over $\mathcal{X}$

- Typically very challenging (even for convex/polytopic $\mathcal{X}$)
- Analytical solution possible for half-spaces $\mathcal{X} := \{x \,|\, h^\mathsf{T} x \leq b\}$

[2] or normally distributed

Gaussian Half-Space Chance Constraint I

Given half-space chance constraint and Gaussian state distribution

$$x(k) \sim \mathcal{N}(\mathbb{E}(x(k)), \mathrm{var}(x(k)))$$
$$\mathcal{X} = \{x \mid h^\top x \leq b\}$$

we can construct the marginal distribution in direction of the constraint

$$h^\top x(k) \sim \mathcal{N}(h^\top \mathbb{E}(x(k)), h^\top \mathrm{var}(x(k)) h) \textbf{ (scalar!)}$$

$$\Pr(x(k) \in \mathcal{X}) = \Pr(h^\top x(k) \leq b) = \phi\left(\frac{b - h^\top \mathbb{E}(x(k))}{\sqrt{h^\top \mathrm{var}(x(k)) h}}\right)$$

- ϕ is the cumulative distribution function of the standard normal distribution (available)

$$\phi(\bar{x}) := \Pr(x \leq \bar{x}) \text{ with } x \sim \mathcal{N}(0,1)$$

- depends only on $\mathbb{E}(x(k))$ and $\mathrm{var}(x(k))$ (available)

Gaussian Half-Space Chance Constraint II

$$
\begin{aligned}
\Pr(x(k) \in \mathcal{X}) \geq p \Leftrightarrow \phi\left(\frac{b - h^\mathsf{T}\mathbb{E}(x(k))}{\sqrt{h^\mathsf{T}\mathrm{var}(x(k))h}}\right) \geq p \\
\Leftrightarrow \frac{b - h^\mathsf{T}\mathbb{E}(x(k))}{\sqrt{h^\mathsf{T}\mathrm{var}(x(k))h}} \geq \phi^{-1}(p) \\
\Leftrightarrow -h^\mathsf{T}\mathbb{E}(x(k)) \geq -b + \sqrt{h^\mathsf{T}\mathrm{var}(x(k))h}\,\phi^{-1}(p) \\
\Leftrightarrow h^\mathsf{T}\mathbb{E}(x(k)) \leq b \underbrace{- \sqrt{h^\mathsf{T}\mathrm{var}(x(k))h}\,\phi^{-1}(p)}_{\text{tightening/back-off term}}
\end{aligned}
$$

with ϕ^{-1} the inverse cumulative distribution function of the standard normal distribution (available)

$$\boxed{\text{Tightened half-space constraint when optimizing over } \mathbb{E}(x(k))}$$

Gaussian Linear Quadratic Control with Half-Space Constraint

$$
\begin{aligned}
J^*(x) = \min_{\{\bar{u}_k\}} \quad & \mathbb{E}\left(\|x(\bar{N})\|_P^2 + \sum_{k=0}^{\bar{N}} \|x(k)\|_Q^2 + \|u(k)\|_R^2 \right) \\
\text{s.t.} \quad & x(k+1) = Ax(k) + Bu(k) + w(k), \\
& u(k) = K(x(k) - \mathbb{E}(x(k)) + \bar{u}_k, \\
& w(k) \sim \mathcal{N}(0, \Sigma_w), \\
& \Pr(h^\top x(k) \leq b) \geq p, \\
& x(k) = x
\end{aligned}
$$

Gaussian Linear Quadratic Control with Half-Space Constraint

$$
\begin{aligned}
\tilde{J}^*_{\text{pre,cc}}(x) = \min_{\{\bar{u}_k\}} \quad & \left\|\mathbb{E}(x(\bar{N}))\right\|_P^2 + \sum_{k=0}^{\bar{N}} \|\mathbb{E}(x(k))\|_Q^2 + \|\bar{u}_k\|_R^2 \\
\text{s.t.} \quad & \mathbb{E}(x(k+1)) = A\mathbb{E}(x(k)) + B\bar{u}_k, \\
& h^\mathsf{T}\mathbb{E}(x(k)) \leq b - \sqrt{h^\mathsf{T}\text{var}(x(k))h}\,\phi^{-1}(p), \\
& \mathbb{E}(x(0)) = x
\end{aligned}
$$

- $\text{var}(x(k))$ is iteratively precomputed from $\text{var}(x(k+1)) = (A+BK)\,\text{var}(x(k))(A+BK)^T + \Sigma_w$

$\rightarrow$ solution for $\{\bar{u}_k\}$ now depends on K

Segway: Open-Loop Sequence with Predefined Tube Controller

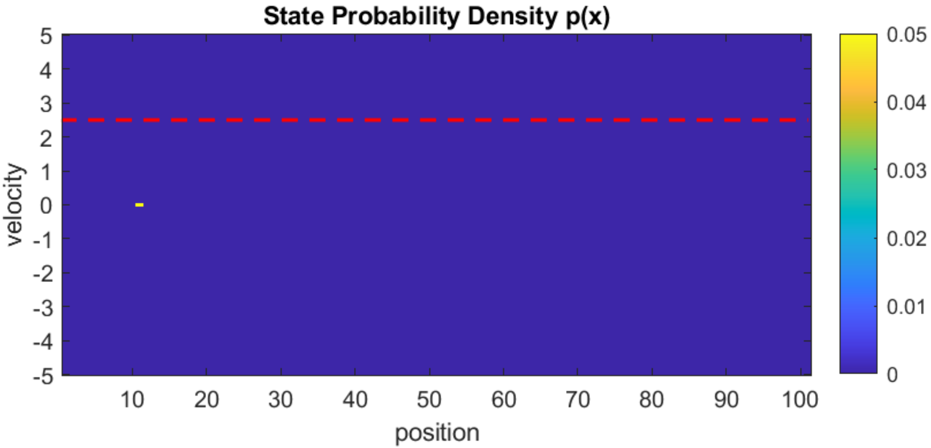

Chance Constraint with $p = 0.9$: $\begin{bmatrix} 0 & 1 \end{bmatrix} \mathbb{E}(x(k)) \leq 2.5 - \sqrt{\mathrm{var}(x_2(k))}\,\phi^{-1}(p)$

Segway: Open-Loop Sequence with Predefined Tube Controller

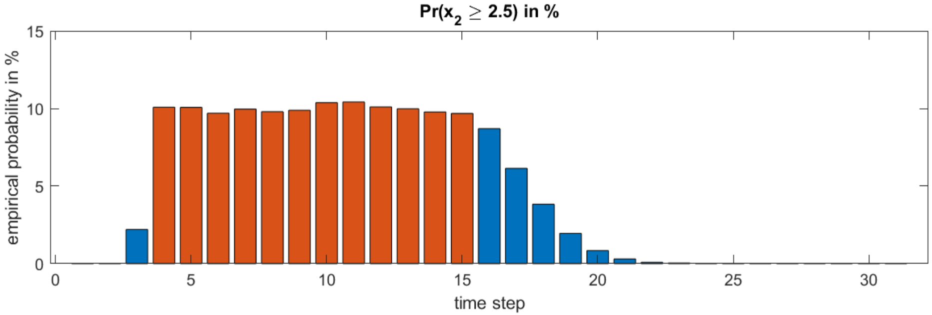

direct handle on violation probability

Outline

1. Problem Formulation: Linear Quadratic Control
2. Dynamic Programming Solution (backwards in time)
3. Forward Planning Solution (forwards in time)
4. Constrained Stochastic Optimal Control Problems
5. Outlook Model Predictive Control

Receding Horizon Control

So far: Optimization carried out offline, directly returning feedback control law

$$\{\pi_k(x)\}_{k=0}^{\bar{N}-1} = \operatorname{argmin} J(x_{\text{init}})$$

and feedback typically reduced to linear feedback.

In MPC (receding/shrinking horizon control) the optimization is instead repeated at each time step, with only the first computed control input applied to the system

$$\pi_k(x) = \text{"first element of"} \ \operatorname{argmin} J_k(x)$$

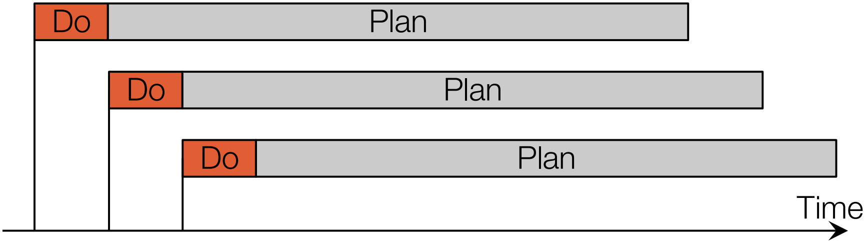

Unique Challenges in MPC

- Control law defined indirectly via repeated optimization
- Introduces (nonlinear) feedback
- Optimization typically over a shortened horizon (typically task horizon assumed infinite)

Unique challenges resulting from the receding horizon nature and mismatch between prediction and closed-loop.

- Feasibility
- Stability
- Resulting closed-loop trajectory distributions
- (Closed-loop) chance-constraint satisfaction

References and further reading

- Lecture Slides: Linear Dynamical Systems (EE363), Prof. Boyd, Stanford
- D. Bertsekas, "Dynamic Programming and Optimal Control I", 1995 (Chapter 4.1)
- Farina et al., "Stochastic linear Model Predictive Control with chance constraints – A review", Journal of Process Control, 2016