AMPyC

Chapter 6: Stochastic Model Predictive Control I

30. May 2025

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

ETH Zurich

MPC for additive disturbances - Robust setting

Uncertain constrained system
$x(k+1) = f(x(k), u(k)) + w(k) \qquad x, u \in \mathcal{X}, \mathcal{U} \qquad w \in \mathcal{W}$

Design control law $u(k) = \pi(x(k))$ such that the system:

1. Satisfies constraints : $\{x(k)\} \subset \mathcal{X}$, $\{u(k)\} \subset \mathcal{U}$ for **all** disturbance realizations
2. Is stable: Converges to a neighborhood of the origin
3. Optimizes (nominal/worst-case) "performance"
4. Maximizes the set $\{x(0) \,|\, \text{Conditions 1-3 are met}\}$

MPC for additive disturbances - Stochastic setting

Uncertain constrained system
$x(k+1) = f(x(k), u(k)) + w(k) \quad \Pr(x(k) \in \mathcal{X}) \geq p, \Pr(u(k) \in \mathcal{U}) \geq p, \quad w(k) \sim \mathcal{Q}^w,$ i.i.d.

Design control law $u(k) = \pi(x(k))$ such that the system:

1. Satisfies constraints : $x(k) \in \mathcal{X}$, $u(k) \in \mathcal{U}$ **with given probability p**
2. Is 'stable': Converges to the origin **in a suitable sense**
3. Optimizes (nominal/**expected**) "performance"
4. Maximizes the set $\{x(0) \,|\, \text{Conditions 1-3 are met}\}$

Receding Horizon Control for Stochastic Systems

In the last chapter we formulated 'open-loop' stochastic optimal control problems

This (& next) chapter: **Receding horizon control** of stochastic systems

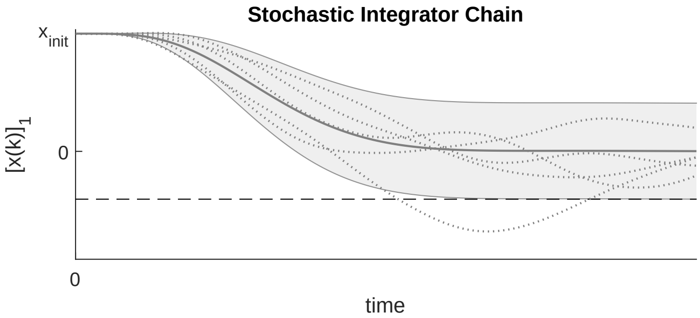

solid line:	mean trajectory
dotted lines:	trajectory samples
shaded area:	confidence region

1. Stability & Performance
2. Feasibility and (chance) constraint satisfaction

Receding Horizon Control for Stochastic Systems

In the last chapter we formulated 'open-loop' stochastic optimal control problems

This (& next) chapter: **Receding horizon control** of stochastic systems

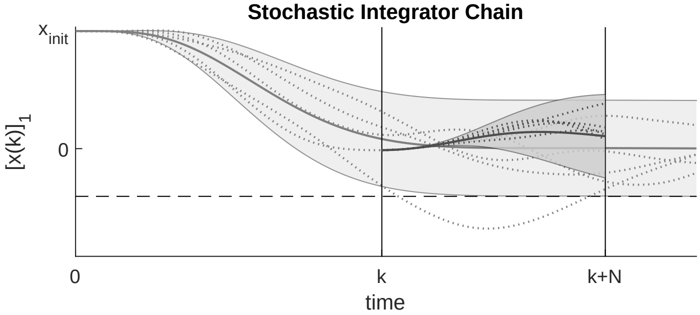

solid line: mean trajectory
dotted lines: trajectory samples
shaded area: confidence region

1. Stability & Performance
2. Feasibility and (chance) constraint satisfaction

Receding Horizon Control for Stochastic Systems

In the last chapter we formulated 'open-loop' stochastic optimal control problems

This (& next) chapter: **Receding horizon control** of stochastic systems

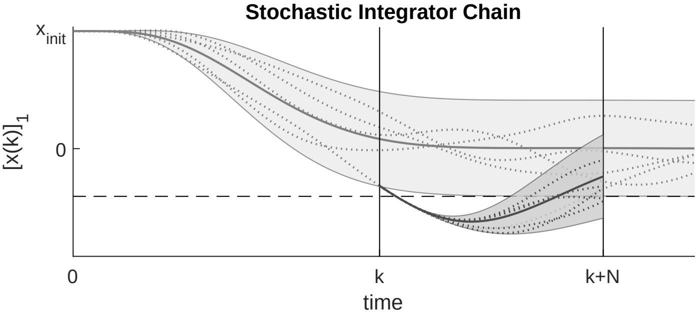

solid line:	mean trajectory
dotted lines:	trajectory samples
shaded area:	confidence region

1. Stability & Performance
2. Feasibility and (chance) constraint satisfaction

Contents – 6: Stochastic MPC I

- Understand stability concepts in stochastic model predictive control
  - Asymptotic average performance as a typical criterion for additive disturbances
  - Derive asymptotic average performance bounds for (linear) stochastic MPC

- Understand feasibility issues arising in stochastic MPC
  - Derive recursively feasible stochastic MPC for bounded random disturbances
  - Show closed-loop chance constraint satisfaction for recursively feasible stochastic MPC

Outline

1. Stability Criteria in Stochastic Receding Horizon Control

2. Chance Constraints in SMPC and Feasibility Issues

3. "Constraint Tightening" SMPC for Bounded Disturbances

Outline

1. Stability Criteria in Stochastic Receding Horizon Control

2. Chance Constraints in SMPC and Feasibility Issues

3. "Constraint Tightening" SMPC for Bounded Disturbances

Outline

1. Stability Criteria in Stochastic Receding Horizon Control

   Asymptotic Average Performance under Additive Disturbances

   Asymptotic Average Performance Bounds in Stochastic MPC

Asymptotic Average Performance bounds

- Additive noise requires adjusted stability concept (cf. Input-to-State Stability)
- Stochastic generalization with expected Lyapunov-like decrease:

$$\mathbb{E}(V(x(k+1)) \mid x(k)) - V(x(k)) \leq -l(x(k), \pi(x(k))) + C$$

- We focus on asymptotic average performance as most common stability/convergence criterion in stochastic MPC.

Asymptotic average performance

$$l_{\text{avg}} = \lim_{\bar{N} \to \infty} \frac{1}{\bar{N}} \sum_{k=0}^{\bar{N}-1} \mathbb{E}\left(l(x(k), \pi(x(k)))\right)$$

Example: Overhead Crane

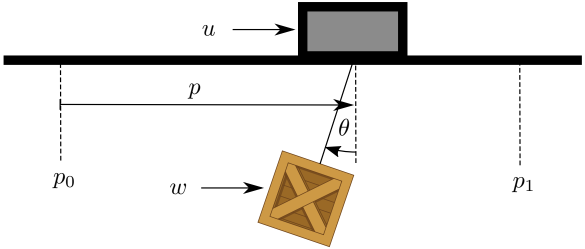

Damped cart pole system

- States: Slider position p, velocity v, load angle Θ, ang. velocity $\dot{\Theta}$
- Input: Slider acceleration u
- For now: Deterministic, $w = 0$

$$\text{Regularization task: } l(x,u) = \|x\|_Q^2 + \|u\|_R^2,\ Q = I,\ R = 10^{-3}$$

Example: Overhead Crane (LQR)

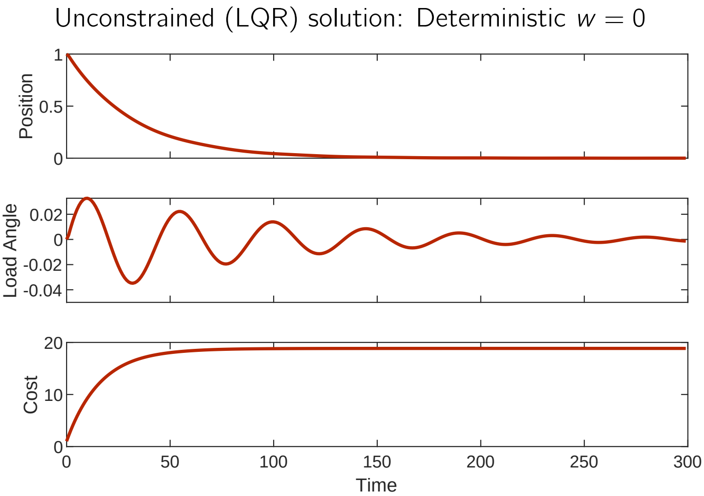

Infinite Horizon Linear Quadratic Cost with Disturbances I

Consider now the same system subject to additive zero mean disturbances

$$
\begin{aligned}
x(k+1) &= (A + BK)x(k) + w(k), \\
\mathbb{E}(w(k)) &= 0, \ \mathrm{var}(w(k)) = \Sigma_w \\
\mathbb{E}(l(x, Kx)) &= \|\mathbb{E}(x)\|^2_{\tilde{Q}} + \mathrm{tr}(\tilde{Q}\,\mathrm{var}(x))
\end{aligned}
$$

with $\tilde{Q} = Q + K^\mathsf{T} RK$

What happens to the expected cost as horizon $\bar{N}$ is going towards infinity?

$$
\lim_{\bar{N}\to\infty} \sum_{k=0}^{\bar{N}-1} \mathbb{E}(l(x(k), Kx(k)))
$$

Example: Overhead Crane (LQR)

Unconstrained (LQR) solution: Stochastic i.i.d. $w \sim \mathcal{N}(0, 1)$

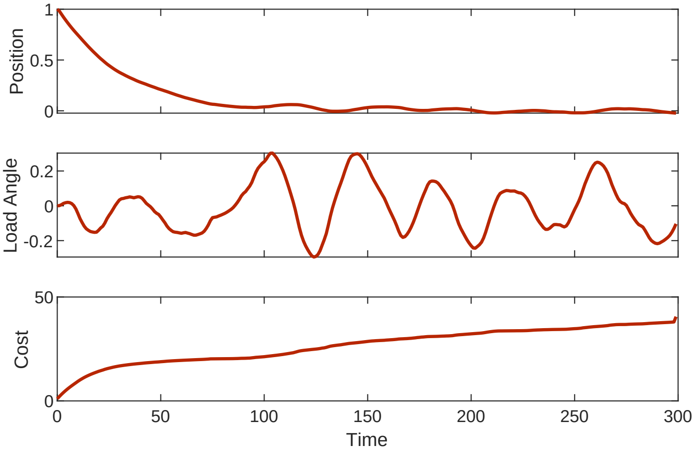

Infinite Horizon Linear Quadratic Cost with Disturbances II

$$\lim_{\bar N \to \infty} \sum_{k=0}^{\bar N - 1} \mathbb{E}(\|x(k)\|_{\tilde Q}^2) = x(0)^\mathsf{T} P x(0) + \lim_{\bar N \to \infty} \sum_{k=0}^{\bar N - 1} \mathrm{tr}(\tilde Q \operatorname{var} x(k))$$

But $\operatorname{var}(x(k))$ converges to non-zero value as $k \to \infty$

$$\lim_{k \to \infty} (\operatorname{var}(x(k))) = \Sigma_\infty = (A + BK)\Sigma_\infty (A + BK)^\mathsf{T} + \Sigma_w \text{ and } \lim_{k \to \infty} \mathrm{tr}(\tilde Q \operatorname{var}(x(k))) = \mathrm{tr}(\tilde Q \Sigma_\infty)$$

- Stage cost keeps accumulating: $\lim_{\bar N \to \infty} \sum_{k=0}^{\bar N - 1} \mathrm{tr}(\tilde Q \operatorname{var}(x(k)))$ diverges,
- **But** "speed of accumulation" $\lim_{\bar N \to \infty} \mathrm{tr}(\tilde Q \operatorname{var} x(k)) = \mathrm{tr}(\tilde Q \Sigma_\infty)$ converges

$$l_{\mathrm{avg}} = \lim_{\bar N \to \infty} \frac{1}{\bar N} \mathbb{E}\left(\sum_{k=0}^{\bar N - 1} l(x(k), Kx(k)) \,\middle|\, x(0) \right) = \mathrm{tr}(\tilde Q \Sigma_\infty) = \mathrm{tr}(P \Sigma_w)$$

Example: Overhead Crane (LQR)

Unconstrained (LQR) solution: Stochastic $w \sim \mathcal{N}(0, 1)$

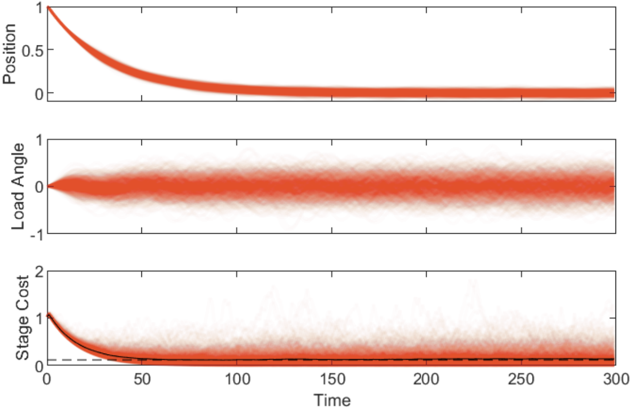

- Dashed: Theoretical asymptotic expected stage cost
- Full line: Empirical average stage cost over 1000 trajectories

Outline

1. Stability Criteria in Stochastic Receding Horizon Control

    Asymptotic Average Performance under Additive Disturbances

    Asymptotic Average Performance Bounds in Stochastic MPC

MPC under Additive Disturbances: Robust vs. Stochastic

$$x(k+1) = f(x(k), \pi(x(k))) + w(k)$$

with MPC control law π and cost function $J = l_f(x_N) + \sum_{i=0}^{N-1} l(x_i, u_i)$

Recall ISS-Lyapunov result:

$$\boxed{J^*(x(k+1)) - J^*(x(k) \leq -l(x(k), \pi(x(k)) + C}$$

where C depends on the "size" of the disturbance. $\rightarrow$ ISS-stability (convergence to a set)

Similarly, we will show in the stochastic case that

$$\boxed{\mathbb{E}(J^*(x(k+1)) \,|\, x(k)) - J^*(x(k)) \leq -l(x(k), \pi(x(k)) + C}$$

where C depends on the "size" of the disturbance. $\rightarrow$ Bounds on average expected closed-loop cost.

Asymptotic Average Performance in Stochastic MPC

Asymptotic average performance bound

$$l_{\text{avg}} = \lim_{\bar{N}\to\infty} \frac{1}{\bar{N}} \sum_{k=0}^{\bar{N}-1} \mathbb{E}\left(l(x(k), \pi(x(k)))\right) \leq C$$

We will derive such a performance bound for stochastic MPC in three steps:

1. Show that Lyapunov-like decrease implies asymptotic average performance bound

$$\mathbb{E}(V(x(k+1)) \,|\, x(k)) - V(x(k)) \leq -l(x(k), \pi(x(k))) + C \Rightarrow l_{\text{avg}} \leq C$$

2. Apply this to the cost function decrease in MPC

$$\mathbb{E}(J^*(x(k+1)) \,|\, x(k)) - J^*(x(k)) \leq -l(x(k), \pi(x(k))) + C$$

3. Derive tractable solution for (unconstrained/recursively feasible) linear stochastic MPC

Lyapunov-like Decrease implies Peformance Bound I

$$\mathbb{E}(V(x(k+1))\,|\,x(k)) - V(x(k)) \leq -l(x(k),\pi(x(k))) + C \Rightarrow l_{\text{avg}} \leq C$$

Apply **law of iterated expectation:** $\mathbb{E}_{x,y}(f(x,y)) = \mathbb{E}_y(\mathbb{E}_x(f(x,y)|y))$

$$\begin{aligned}
&\mathbb{E}(V(x(k+1))\,|\,x(0)) - V(x(0)) \\
=&\mathbb{E}\Big(\mathbb{E}\big(V(x(k+1))\,\big|\,x(k)\big)\,\Big|\,x(0)\Big) - V(x(0)) \\
=&\mathbb{E}\left(\underbrace{\mathbb{E}\left(V(x(k+1))\,|\,x(k)\right) - V(x(k))}_{\leq -l(x(k),\pi(x(k)))+C} + V(x(k))\,\middle|\,x(0)\right) - V(x(0)) \\
\leq&\mathbb{E}\big(-l(x(k),\pi(x(k))) + C + V(x(k))\,\big|\,x(0)\big) - V(x(0)) \\
\leq&\ldots \leq \mathbb{E}\left(\sum_{i=0}^{k} -l(x(i),\pi(x(i))) + C\,\middle|\,x(0)\right) = (k+1)C + \sum_{i=0}^{k}\mathbb{E}\left(-l(x(i),\pi(x(i)))\,|\,x(0)\right)
\end{aligned}$$

Lyapunov-like Decrease implies Peformance Bound II

$$\sum_{i=0}^{k-1} \mathbb{E}\left(l(x(i), \pi(x(i))) \,|\, x(0)\right) - kC \leq V(x(0)) - \mathbb{E}(V(x(k)) \,|\, x(0))$$

We then have that

$$\begin{aligned}
&\lim_{\bar{N}\to\infty} \frac{1}{\bar{N}} V(x(0)) = 0 \quad \mathbb{E}(V(x(k)) \,|\, x(0)) \geq 0 \\
\Rightarrow &\lim_{\bar{N}\to\infty} \frac{1}{\bar{N}} \left(\sum_{k=0}^{\bar{N}-1} \mathbb{E}\left(l(x(k), \pi(x(k))) \,|\, x(0)\right) - \bar{N}C \right) \leq 0 \\
\Leftrightarrow &\lim_{\bar{N}\to\infty} \frac{1}{\bar{N}} \sum_{k=0}^{\bar{N}-1} \mathbb{E}\left(l(x(k), \pi(x(k))) \,|\, x(0)\right) \leq C
\end{aligned}$$

Asymptoptic Average Performance: Unconstrained MPC

$$
\begin{aligned}
\min_{\{u_i\}} \quad & \mathbb{E}\left(l_f(x_N) + \sum_{i=0}^{N-1} l(x_i, u_i) \right) \\
\text{s.t.} \quad & x_{i+1} = f(x_i, u_i, w_i), \\
& w_i \sim \mathcal{Q}^w \text{ i.i.d.}, \\
& x_0 = x(k)
\end{aligned}
$$

- Receding horizon: $\pi(x(k)) = u_0^* :=$ "first element of" $\operatorname{argmin} J(x(k))$
- Assumption: Prediction process is identical to "real" process

$$x_{i+1} = f(x_i, u_i, w_i),\ w_i \sim \mathcal{Q}^w \text{ i.i.d.} \quad \leftrightarrow \quad x(k+1) = f(x(k), u(k), w(k)),\ w(k) \sim \mathcal{Q}^w \text{ i.i.d.}$$

Asymptoptic Average Performance: Unconstrained MPC

Use $J^*(x)$ as Lyapunov-like decrease function: Candidate sequence approach

$$\mathbb{E}(J^*(x(k+1))\,|\,x(k)) - J^*(x(k)) \leq \mathbb{E}(\bar{J}(x(k+1))\,|\,x(k)) - J^*(x(k))$$

where $\mathbb{E}(\bar{J}(x(k+1)\,|\,x(k))$ is cost of candidate solution with terminal control law $\pi_f(\cdot)$[1]

Candidate solution: $\bar{U} = \{u_0, \ldots, u_{N-1}\} = \{u_1^*, \ldots, u_{N-1}^*, \pi_f(\cdot)\}$

resulting in $\bar{X} = \{x_0, \ldots, x_N\} \overset{d}{=} \{x_1^*, \ldots, x_{N-1}^*, f(x_N^*, \pi_f(\cdot), w)\}$ (equal in distribution)

and $\mathbb{E}_{w(k)}(\bar{J}(x(k+1))\,|\,x(k)) = \mathbb{E}_{w(k)}\left(\mathbb{E}\left(l_f(x_N) + \sum_{i=0}^{N-1} l(x_i, u_i)\right)\right)$

$$= \mathbb{E}\left(l_f(f(x_N^*, \pi_f(\cdot), w)) + l(x_N^*, \pi_f(\cdot)) + \sum_{i=1}^{N-1} l(x_i^*, u_i^*)\right)$$

[1] mapping previous solution $\{u_i^*\}_0^{N-1}$ and measured state $x(k)$ to a candidate u_{N-1}

Asymptoptic Average Performance: Unconstrained MPC

From

$$\mathbb{E}_{w(k)}(\bar{J}(x(k+1))\,|\,x(k)) = \mathbb{E}\left(l_f(f(x_N^*, \pi_f(\cdot), w)) + l(x_N^*, \pi_f(\cdot)) + \sum_{i=1}^{N-1} l(x_i^*, u_i^*) \right)$$

we then have

$$\begin{aligned}
&\mathbb{E}(\bar{J}(x(k+1))\,|\,x(k)) - J^*(x(k)) \\
=&\mathbb{E}\left(l_f(f(x_N^*, \pi_f(\cdot), w)) + l(x_N^*, \pi_f(\cdot)) + \sum_{i=1}^{N-1} l(x_i^*, u_i^*) \right) - \mathbb{E}\left(l_f(x_N^*) + \sum_{i=0}^{N-1} l(x_i^*, u_i^*) \right) \\
=&\mathbb{E}(l_f(f(x_N^*, \pi_f(\cdot), w))) - \mathbb{E}(l_f(x_N^*)) + \mathbb{E}(l(x_N^*, \pi_f(\cdot))) - l(x(k), u_0^*)
\end{aligned}$$

Cost Decrease: Nominal MPC vs. Stochastic

Standard MPC Lyapunov decrease

$$l_f(f(x_N^*, \pi_f(x_N^*))) - l_f(x_N^*)$$
$$+l(x_N^*, \pi_f(x_N^*)) - l(x(k), u_0^*)$$

Typical assumption:

$$l_f(f(x_N^*, \pi_f(x_N^*))) - l_f(x_N^*)$$
$$\leq - l(x_N, \pi_f(x_N^*))$$

implies

$$J^*(x(k+1) - J^*(x(k)) \leq -l(x(k), u_0^*)$$

Result: asymptotic stability
$\lim_{\bar{N}\to\infty} \frac{1}{\bar{N}} \sum l(x(k), u(k)) = 0$

Stochastic MPC Lyapunov decrease

$$\mathbb{E}(l_f(f(x_N^*, \pi_f(\cdot), w))) - \mathbb{E}(l_f(x_N^*))$$
$$+\mathbb{E}(l(x_N^*, \pi_f(\cdot))) - l(x(k), u_0^*)$$

Corresponding assumption:

$$\mathbb{E}(l_f(f(x_N^*, \pi_f(\cdot), w)))) - \mathbb{E}(l_f(x_N^*))$$
$$\leq - \mathbb{E}(l(x_N^*, \pi_f(\cdot))) + C$$

implies

$$\mathbb{E}(J^*(x(k+1)) \,|\, x(k)) - J^*(x(k)) \leq -l(x(k), u_0^*) + C$$

Result: average performance bound
$\lim_{\bar{N}\to\infty} \frac{1}{\bar{N}} \sum \mathbb{E}\left(l(x(k), u(k))\right) \leq C$

Tractable Case: Linear Dynamics, Quadratic Cost I

Typical assumption nominal and robust MPC:

$$\begin{aligned} &l_f(f(x_N^*, \pi_f(x_N^*))) - l_f(x_N^*) \\ \leq &- l(x_N, \pi_f(x_N)) \end{aligned}$$

in a **terminal invariant set** $\mathcal{X}_f \subseteq \mathcal{X}$

Corresponding assumption in Stochastic MPC:

$$\begin{aligned} &\mathbb{E}(l_f(f(x_N^*, \pi_f(\cdot), w)))) - \mathbb{E}(l_f(x_N^*)) \\ \leq &- \mathbb{E}(l(x_N^*, \pi_f(\cdot))) + C \end{aligned}$$

hard to enforce locally (further reading [3])

For the special case of LTI systems with quadratic cost, this can be ensured globally!

$$\begin{aligned} x(k+1) &= Ax(k) + Bu(k) + w(k),\ \mathbb{E}(w(k)) = 0,\ \mathrm{var}(w(k)) = \Sigma_w, \text{ i.i.d.} \\ l(x,u) &= \|x(k)\|_Q^2 + \|u(k)\|_R^2 \end{aligned}$$

Terminal cost as infinite horizon cost $l_f(x) = \|x\|_P^2$ of autonomous system[2]

$$P = A^\mathsf{T} P A + Q$$

[2] Assumes asympt. stable A

Tractable Case: Linear Dynamics, Quadratic Cost II

$$
\begin{aligned}
&\mathbb{E}(l_f(f(x_N^*, \pi_f(\cdot), w_N)) && - \mathbb{E}(l_f(x_N^*)) && + \mathbb{E}(l(x_N^*, \pi_f(\cdot))) && - l(x(k), u_0^*) \\
=&\mathbb{E}(\|Ax_N^* + w\|_P^2 && - \|x_N^*\|_P^2 && + \|x_N^*\|_Q^2 + \|0\|_R^2) && - \|x(k)\|_Q^2 - \|u_0^*\|_R^2 \\
=&\mathbb{E}(\|x_N^*\|_{A^\top PA}^2 + \|w\|_P^2 + && - \|x_N^*\|_P^2 && + \|x_N^*\|_Q^2) && - \|x(k)\|_Q^2 - \|u_0^*\|_R^2 \\
=&\mathbb{E}(\|x_N^*\|_{A^\top PA - P + Q}^2) + \mathrm{tr}(P\Sigma_w) && && && - \|x(k)\|_Q^2 - \|u_0^*\|_R^2
\end{aligned}
$$

Where due to the choice of P we have $A^\top PA - P + Q = 0$ (!)

$$\mathbb{E}(J^*(x(k+1)) \,|\, x(k)) - J^*(x(k))) \leq -\|x(k)\|_Q^2 - \|u(k)\|_R^2 + \mathrm{tr}(P\Sigma_w)$$

Same argument holds when optimizing over affine part $\bar{u}_i$ of tube policy $u_i = K(x_i - \mathbb{E}(x_i)) + \bar{u}_i$

(Unconstrained) linear quadratic MPC has better (or equal) performance as tube controller K when choosing corresponding terminal cost $P = (A + BK)^\top P(A + BK) + Q + K^\top RK$, i.e.,

$$\lim_{\bar{N} \to \infty} \frac{1}{\bar{N}} \sum_{k=0}^{\bar{N}-1} \mathbb{E}\left(l(x(k), u(k)) \,|\, x(0)\right) \leq \mathrm{tr}(P\Sigma_w)$$

Example: Overhead Crane (MPC vs. Linear Controller)

Linear Controller $\pi(x) = Kx$ (hand-tuned)

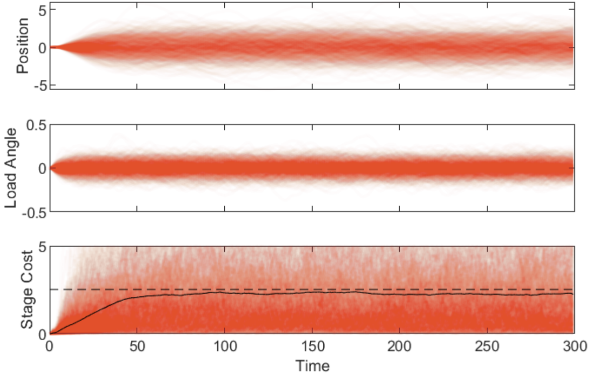

Dashed line: Performance bound $\mathrm{tr}(P\Sigma_w)$
Solied line: Mean performance $\mathbb{E}(l(x, u))$

with K: $\approx \mathrm{tr}(P\Sigma_w)$

with MPC: $\ll \mathrm{tr}(P\Sigma_w)$
(P as terminal weight)

Example: Overhead Crane (MPC vs. Linear Controller)

Unconstrained Model Predictive Controller

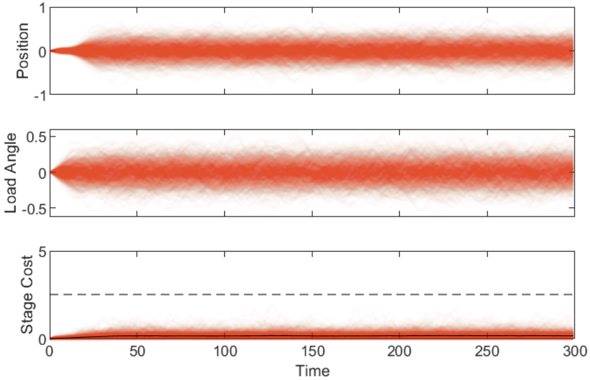

Dashed line: Performance bound $\mathrm{tr}(P\Sigma_w)$

Solied line: Mean performance $\mathbb{E}(l(x, u))$

with K: $\approx \mathrm{tr}(P\Sigma_w)$

with MPC: $\ll \mathrm{tr}(P\Sigma_w)$

(P as terminal weight)

Outline

1. Stability Criteria in Stochastic Receding Horizon Control

2. Chance Constraints in SMPC and Feasibility Issues

3. "Constraint Tightening" SMPC for Bounded Disturbances

Recall: Forward Planning (Chapter 5)

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

- optimize over affine part of tube control policy $u(k) = K(x(k) - \mathbb{E}(x(k)) + \overline{u}_k$

Recall: Forward Planning (Chapter 5)

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

- optimize over affine part of tube control policy $u(k) = K(x(k) - \mathbb{E}(x(k)) + \overline{u}_k$
- tractable (approximate) solution to chance constrained problem via deterministic optimization

Recall: Forward Planning (Chapter 5)

$$
\begin{aligned}
\tilde{J}^*(x(k)) = \min_{\{\overline{u}_i\}} \quad & \|\mathbb{E}(x_N)\|_P^2 + \sum_{i=0}^{N-1} \|\mathbb{E}(x_i)\|_Q^2 + \|\overline{u}_i\|_R^2 \\
\text{s.t.} \quad & \mathbb{E}(x_{i+1}) = A\mathbb{E}(x_i) + B\overline{u}_i, \\
& h^\mathsf{T}\mathbb{E}(x_i) \le b - \sqrt{h^\mathsf{T}\mathrm{var}(x_i)h}\,\phi^{-1}(p), \\
& \mathbb{E}(x_0) = x(k)
\end{aligned}
$$

- optimize over affine part of tube control policy $u_i = K(x_i - \mathbb{E}(x_i)) + \overline{u}_i$
- tractable (approximate) solution to chance constrained problem via deterministic optimization
- straightforward conversion to receding horizon controller (in principle...)

Illustration: Feasibility Issues in SMPC

Under general stochastic disturbances, feasibility of this formulation **cannot** be guaranteed

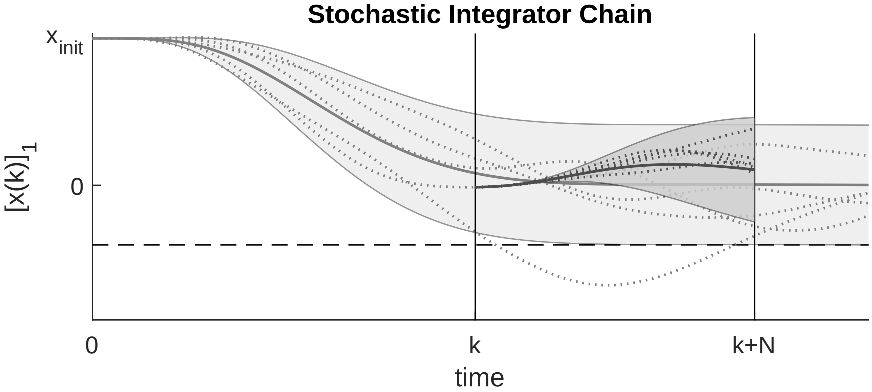

Illustration: Feasibility Issues in SMPC

Under general stochastic disturbances, feasibility of this formulation **cannot** be guaranteed

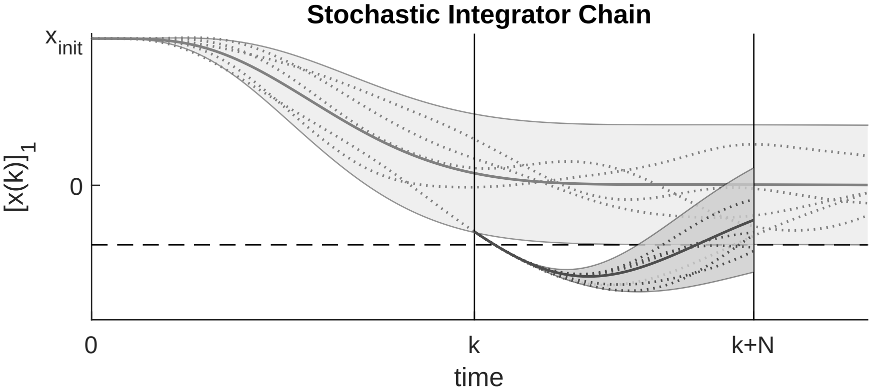

Feasibility Issues in SMPC

Problem: (Potentially unbounded) stochastic disturbances can drive state initial state $x_0 = x(k)$ into infeasible region.

Several strategies to handle this problem

- Assume bounded disturbances
  $\rightarrow$ use robust techniques
- Make use of recovery mechanisms
- Alternative forms of feedback

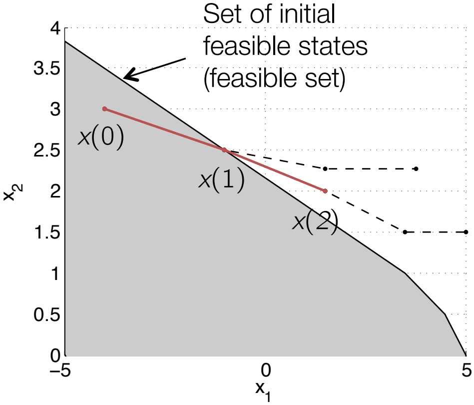

Feasibility Issues in SMPC

Problem: (Potentially unbounded) stochastic disturbances can drive state initial state $x_0 = x(k)$ into infeasible region.

Several strategies to handle this problem

- **Assume bounded disturbances**
  $\rightarrow$ use robust techniques
- Make use of recovery mechanisms
- Alternative forms of feedback
  (next lecture)

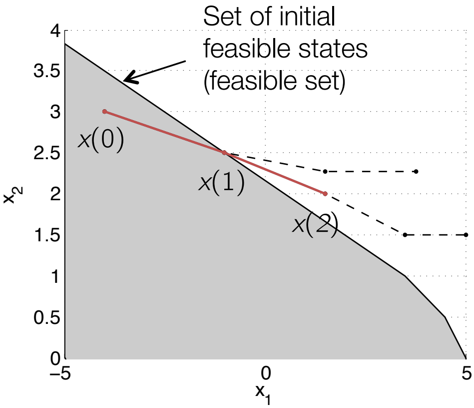

Outline

1. Stability Criteria in Stochastic Receding Horizon Control
2. Chance Constraints in SMPC and Feasibility Issues
3. "Constraint Tightening" SMPC for Bounded Disturbances

Stochastic MPC with Bounded Disturbances

Setup:

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

where $w(k) \sim \mathcal{Q}^w$ i.i.d. **and** all $w(k) \in \mathcal{W}$ with $\mathcal{W}$ a compact set.

- Stochastic disturbance $w(k)$ has **bounded support** $\mathcal{W}$
- Enables use of robust techniques for recursive feasibility

Outline:

1. Constraint-tightening SMPC with recursive feasibility
2. Recursive feasibility $\Rightarrow$ closed-loop satisfaction of chance constraints

Recursive feasibility in SMPC for bounded disturbances

Different techniques exist to ensure recursive feasibility based on robust arguments

Essential difference to robust MPC:

$$\underbrace{\text{Robustly ensuring } \textbf{chance} \text{ constraints}}_{\text{Stochastic MPC with bounded disturbances}} \not\Leftrightarrow \underbrace{\text{Robustly ensuring } \textbf{deterministic} \text{ constraints}}_{\text{Robust MPC}}$$

In the following, we discuss one possible approach related to "constraint-tightening" robust MPC

- Enforce **chance** constraints w.r.t. **all** possible previous disturbances ($i-1$-steps robust, 1-step stochastic)
- Enforce terminal robust invariant set (within constraints) robustly
- For simplicity, we neglect input constraints for now, extension is straightforward
- References: [4, 5]

Robust "Constraint-Tightening" MPC (Chapter 3)

$$
\begin{aligned}
\min_{\{v_i\}} \quad & \|z_N\|_P^2 + \sum_{i=0}^{N-1} \|z_i\|_Q^2 + \|v_i\|_R^2 \\
\text{s.t.} \quad & z_{i+1} = Az_i + Bv_i, \quad i \in [0, N-1], \\
& \{z_i\} \oplus \mathcal{W} \oplus (A+BK)\mathcal{W} \oplus \ldots \oplus (A+BK)^{i-1}\mathcal{W} \subseteq \mathcal{X}, \quad i \in [0, N-1], \\
& \{z_N\} \oplus \mathcal{W} \oplus (A+BK)\mathcal{W} \oplus \ldots \oplus (A+BK)^{N-1}\mathcal{W} \subseteq \mathcal{X}_f, \\
& z_0 = x(k)
\end{aligned}
$$

- Applied control: $u(k) = v_0^*$
- Tightening: $\{z_i\} \oplus \mathcal{W} \oplus \ldots \subseteq \mathcal{X} \Leftrightarrow z_i \in \mathcal{X} \ominus (\mathcal{W} \oplus \ldots)$
- Robustly ensure satisfaction of constraints at each time step
- Terminal robust invariant set under tube controller $\mathcal{X}_f$

Now: Stochastic "Constraint-Tightening" MPC

$$
\begin{aligned}
\min_{\{\overline{u}_i\}} \quad & \|\overline{x}_N\|_P^2 + \sum_{i=0}^{N-1} \|\overline{x}_i\|_Q^2 + \|\overline{u}_i\|_R^2 \\
\text{s.t.} \quad & \overline{x}_{i+1} = A\overline{x}_i + B\overline{u}_i, \quad i \in [0, N-1], \\
& \Pr(\{\overline{x}_i + w_{i-1}\} \oplus (A+BK)\mathcal{W} \oplus \ldots \oplus (A+BK)^{i-1}\mathcal{W} \subseteq \mathcal{X}) \geq p, \quad i \in [1, N-1], \\
& \{\overline{x}_N\} \oplus \mathcal{W} \oplus (A+BK)\mathcal{W} \oplus \ldots \oplus (A+BK)^{N-1}\mathcal{W} \subseteq \mathcal{X}_f, \\
& \overline{x}_0 = x(k)
\end{aligned}
$$

- Applied control: $u(k) = \overline{u}_0^*$
- Expected state $\mathbb{E}(x_i|x(k)) = \overline{x}_i$
- Robustly ensure satisfaction of chance constraints at each time step
- Terminal robust invariant set under tube controller $\mathcal{X}_f$

Now: Stochastic "Constraint-Tightening" MPC

We can deterministically reformulate this chance constraint as

$$\Pr(\{\bar{x}_i + w_{i-1}\} \oplus (A+BK)\mathcal{W} \oplus \ldots \oplus (A+BK)^{i-1}\mathcal{W} \subseteq \mathcal{X}) \geq p$$
$$\Leftarrow \{\bar{x}_i\} \oplus \mathcal{F}_w(p) \oplus (A+BK)\mathcal{W} \oplus \ldots \oplus (A+BK)^{i-1}\mathcal{W} \subseteq \mathcal{X}$$

where $\mathcal{F}_w(p)$ is a set containing w with probability p.

For half-space constraints $\mathcal{X} = \{x \,|\, h^\mathsf{T} x \leq b\}$, least conservative set $\mathcal{F}_w$ given by

$$\mathcal{F}_w(p) = \{w \,|\, h^\mathsf{T} w \leq F_w(p)\}$$

where $F_w(p)$ is the cumulative distribution function of $h^\mathsf{T} w$.

Closed-loop chance constraint satisfaction

Closed-loop chance constraints:

$$(*)\ \Pr(x(k) \in \mathcal{X} \mid x(0)) \geq p,\ \forall k \geq 0$$

But MPC formulation successively enforces

$$(**)\ \Pr(x(k+1) \in \mathcal{X} \mid x(k)) \geq p,\ \forall k \geq 0$$

It can be easily seen that $(**) \Rightarrow (*)$ since

$$\Pr(x(k+1) \in \mathcal{X} \mid x(0)) = \int \underbrace{\Pr(x(k+1) \in \mathcal{X} \mid x(k))}_{\geq p} p(x(k) \mid x(0))\, \mathrm{d}x(k) \geq p$$

This analysis relies on the fact that the MPC formulation is feasible $\forall k \geq 0$.

Illustration: SMPC with Bounded Uncertainties

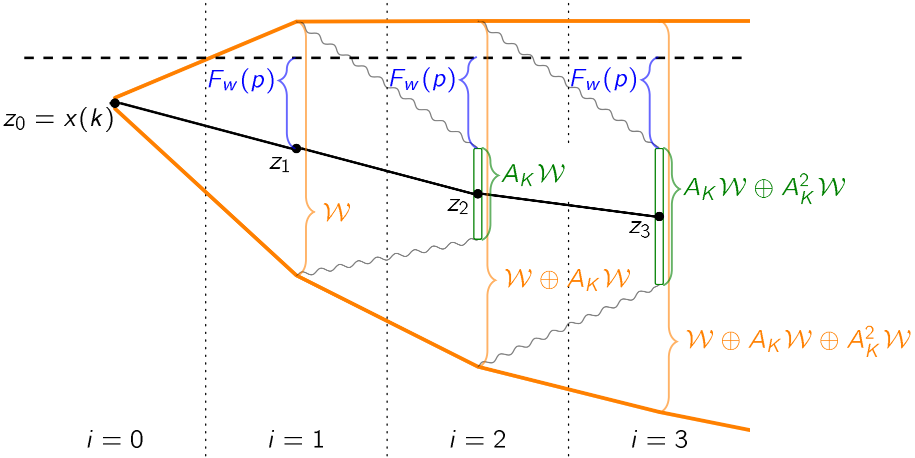

Figure adapted from H. Schlüter, F. Allgöwer, "A Constraint-Tightening Approach to Nonlinear Stochastic Model Predictive Control for Systems under General Disturbances", 2020

Stochastic vs. Robust MPC: Constraint-Tightening

Robust constraint tightening:

$$z_i \in \mathcal{X}_i^{\mathrm{robust}} = \mathcal{X} \ominus (\mathcal{W} \oplus (A+BK)\mathcal{W} \oplus \ldots \oplus (A+BK)^{i-1}\mathcal{W})$$

Stochastic constraint tightening:

$$\bar{x}_i \in \mathcal{X}_i^{\mathrm{stochastic}} = \mathcal{X} \ominus (\mathcal{F}_w(p) \oplus (A+BK)\mathcal{W} \oplus \ldots \oplus (A+BK)^{i-1}\mathcal{W})$$

Stochastic vs. Robust MPC: Constraint-Tightening

Robust constraint tightening:

$$z_i \in \mathcal{X}_i^{\text{robust}} = \mathcal{X} \ominus \left(\mathcal{W} \oplus (A+BK)\mathcal{W} \oplus \ldots \oplus (A+BK)^{i-1}\mathcal{W}\right)$$

Stochastic constraint tightening:

$$\bar{x}_i \in \mathcal{X}_i^{\text{stochastic}} = \mathcal{X} \ominus \left(\mathcal{F}_w(p) \oplus (A+BK)\mathcal{W} \oplus \ldots \oplus (A+BK)^{i-1}\mathcal{W}\right)$$

- Simply replace worst-case $w \in \mathcal{W}$ by $\Pr(w \in \mathcal{F}_w(p)) \geq p$.
- Reduction in conservatism can be small.
- Resulting tightening can be quite conservative compared to exact cumulative distribution.

Stochastic vs. Robust MPC: Constraint-Tightening

Robust constraint tightening:

$$z_i \in \mathcal{X}_i^{\text{robust}} = \mathcal{X} \ominus (\mathcal{W} \oplus (A+BK)\mathcal{W} \oplus \ldots \oplus (A+BK)^{i-1}\mathcal{W})$$

Stochastic constraint tightening:

$$\bar{x}_i \in \mathcal{X}_i^{\text{stochastic}} = \mathcal{X} \ominus (\mathcal{F}_w(p) \oplus (A+BK)\mathcal{W} \oplus \ldots \oplus (A+BK)^{i-1}\mathcal{W})$$

- Simply replace worst-case $w \in \mathcal{W}$ by $\Pr(w \in \mathcal{F}_w(p)) \geq p$.
- Reduction in conservatism can be small.
- Resulting tightening can be quite conservative compared to exact cumulative distribution.

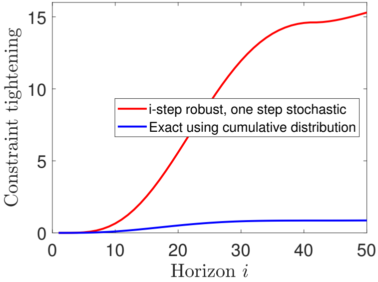

Figure adapted from J. Köhler*, F. Geuss*, M. Zeilinger, "On stochastic MPC formulations with closed-loop guarantees: Analysis and a unifying framework", 2023.

Stochastic vs. Robust MPC: Feasible Region

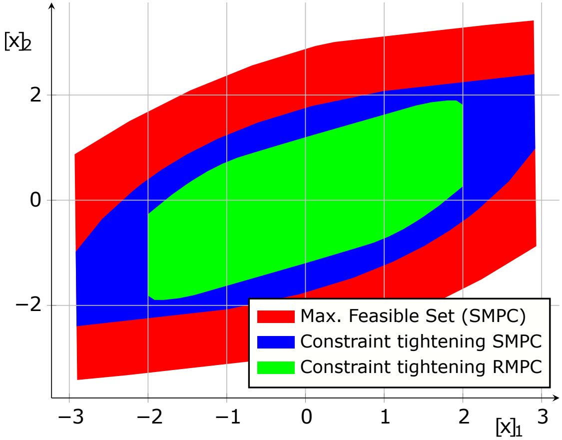

- SMPC reduces conservatism compared to RMPC.
- Constraint tightening SMPC introduces conservatism.

Figure adapted from M. Lorenzen et al, "Constraint-Tightening and Stability in Stochastic Model Predictive Control", Trans. Automatic Control, 2017

Stochastic MPC with Bounded Support

Summary: Stochastic "Constraint Tightening" MPC

- Recursive feasibility properties follow directly from robust case:
  - computed sequence robustly fulfills all chance constraints
  - terminal set $\mathcal{X}_f$ is robust invariant (and inside constraints)

  $\rightarrow$ shifted sequence remains feasible
  - In stochastic setting $\mathcal{X}_f$ can be slightly enlarged (not necessarily completely inside constraints)
- Alternatives: "Brute force" computation of feasible set [4,5]
- Asymptotic average performance analysis applicable (candidate solution feasible).
- Robust-stochastic treatment can be subject to considerable conservatism
  $\Rightarrow$ Study alternatives in the next chapter.

References and further reading

[1] H. Kushner, "Introduction to Stochastic Control", 1971

[2] S.P. Meyn and R.L Tweedie, "Markov Chains and Stochastic Stability", 1993

[3] D. Chatterjee and J. Lygeros, "On Stability and Performance of Stochastic Predictive Control Techniques", Trans. Automatic Control, 2015

[4] M. Korda et al, "Strongly feasible stochastic model predictive control", Conf. Decision Control, 2011

[5] M. Lorenzen et al, "Constraint-Tightening and Stability in Stochastic Model Predictive Control", Trans. Automatic Control, 2017

- M. Cannon et al, "Mean-variance receding horizon control for discrete time linear stochastic systems", IFAC World Congress, 2008
- B. Kouvaritakis, M. Cannon, "Model Predictive Control: Classical, Robust and Stochastic", 2016, (Chapter 7.1 & 7.2)
- M. Farina et al., "Stochastic linear Model Predictive Control with chance constraints – A review", J. Process Control, 2016
- A. Mesbah, "Stochastic model predictive control: An overview and perspectives for future research", IEEE Control Systems Magazine, 2016