Adaptive Model Predictive Control: Robustness and Parameter Estimation

Mark Cannon

Joint work with Matthias Lorenzen, University of Stuttgart
and Xiaonan Lu, University of Oxford

Motivation

Robust MPC paradigm:

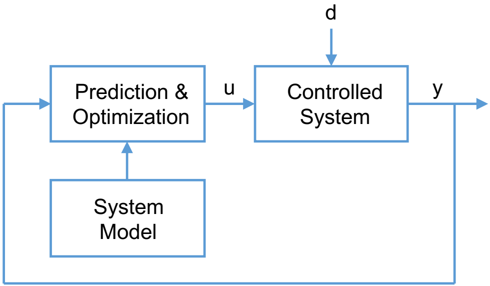

- ▷ MPC requires adequate models of the system, uncertainty, disturbances
- ▷ Amount of uncertainty in the model crucially affects performance
- ▷ Large effort (time & money) spent on model identification offline

Motivation

Adaptive MPC paradigm:

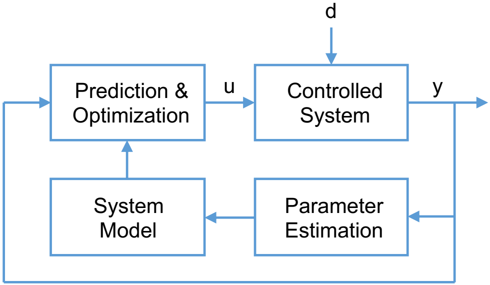

- ▷ Identify model online
- ▷ Require: robust constraint satisfaction
  closed loop stability & performance guarantees
  parameter convergence

Motivation

An idea with a long history: e.g. self-tuning control, DMC, GPC . . .
[Clarke, Tuffs, Mohtadi, 1987]

Revisited with new tools:

- Set membership estimation
  [Bai, Cho, Tempo, 1998]
- Robust tube MPC
  [Langsson, Chryssochoos, Rakovic, Mayne, 2004]
- Dual adaptive/predictive control
  [Lee & Lee, 2009]

Motivation

Recent work on MPC with model adaptation

- Focus on online learning & identification:
  - Persistency of Excitation constraints
    [Marafioti, Bitmead, Hovd, 2014]
  - Kalman filter-based parameter estimation with covariance matrix in cost
    [Heirung, Ydstie, Foss, 2017]
  - Gaussian process regression, particle filtering
    [Klenske, Zeilinger, Scholkopf, Hennig, 2016]
    [Bayard & Schumitzky, 2010]
- Focus on robust constraint satisfaction and performance:
  - Constraints based on prior uncertainty set, online update of cost only
    [Aswani, Gonzalez, Sastry, Tomlin, 2013]
  - Set-based identification, stable FIR plant model
    [Tanaskovic, Fagiano, Smith, Morari, 2014]

Motivation

This talk considers how to

- ensure robust constraint satisfaction;
- update constraints & costs online via set-membership & point estimates;
- enforce parameter convergence via persistency of excitation conditions.

Outline:

1. Set membership parameter estimation
2. Polytopic tube robust MPC
3. Parameter convergence and time-varying parameters

Parameter set estimate

Plant model with unknown parameter vector $\theta^\star$ and disturbance w:

$$x_{k+1} = A(\theta^\star)x_k + B(\theta^\star)u_k + w_k$$

Assume the model is affine in $\theta^\star$ (assumed constant)

$$x_{k+1} = D_k\theta^\star + d_k + w_k \qquad \begin{cases} D_k = D(x_k, u_k) \\ d_k = A_0 x_k + B_0 u_k \end{cases}$$

with stochastic disturbance $w_k \in \mathcal{W}$ a.s., $\mathcal{W}$ known, compact, polytopic

If x_k, x_{k-1}, u_{k-1} are known, then $\theta^\star$ must lie in the "unfalsified set":

$$\Delta_k = \{\theta : x_k = D_{k-1}\theta + d_{k-1} + w,\ w \in \mathcal{W}\}$$

Hence update the parameter set estimate Θ_k via

$$\Theta_k = \Theta_{k-1} \cap \Delta_k$$

Parameter set estimate

- If Θ_0 is a compact polytope, then Θ_k is a compact polytope for all $k > 0$
  But the update $\Theta_k = \Theta_{k-1} \cap \Delta_k$ has potentially unbounded complexity!

- Instead, use fixed complexity sets defined for given H_Θ by

$$\Theta_k = \{\theta : H_\Theta \theta \le h_k\}$$

  and update Θ_k by solving a set of linear programs:

$$[h_k]_i = \max_{w \in \mathcal{W},\, \theta \in \Theta_{k-1}} [H_\Theta]_i \theta \quad \text{s.t.} \quad x_k = D_{k-1}\theta + d_{k-1} + w$$

  Then

$$\Theta_{k-1} \cap \Delta_k \subseteq \Theta_k \subseteq \Theta_{k-1}$$

  since
  - $[H_\Theta]_i \theta \le [h_k]_i$ for all $\theta \in \Delta_k \cap \Theta_{k-1}$ implies $\Theta_{k-1} \cap \Delta_k \subseteq \Theta_k$
  - $[h_k]_i \le \max_{\theta \in \Theta_{k-1}} [H_\Theta]_i \theta = [h_{k-1}]_i$ implies $\Theta_k \subseteq \Theta_{k-1}$

Parameter point estimate

To ensure closed loop l^2 stability, we define the MPC cost in terms of a point estimate $\hat{\theta}_k$ of $\theta^\star$, computed using a LMS filter

Given a parameter estimate $\hat{\theta}_k$, let $\hat{x}_{1|k} = A(\hat{\theta}_k)x_k + B(\hat{\theta}_k)u_k$
Then for a given parameter update gain $\mu > 0$ satisfying

$$1/\mu > \sup_{(x,u)\in\mathcal{Z}} \|D(x,u)\|^2$$

the point estimate $\hat{\theta}_k$ is defined

$$\begin{aligned}\tilde{\theta}_k &= \hat{\theta}_{k-1} + \mu D^\top(x_{k-1}, u_{k-1})(x_k - \hat{x}_{1|k-1}) \\ \hat{\theta}_k &= \Pi_{\Theta_k}(\tilde{\theta}_k)\end{aligned}$$

where Π_{Θ_k} is the Euclidean projection onto Θ_k

Here $\mathcal{Z}$ is the joint state and control constraint set (assumed bounded) and the point estimate update becomes simply a projection onto Θ_k if $\mu \to 0$

Parameter point estimate

The closed loop l^2 gain property is based on the following result

Lemma (Point estimate)

If $\sup_{k\in\mathbb{N}} \|x_k\| < \infty$ *and* $\sup_{k\in\mathbb{N}} \|u_k\| < \infty$*, then* $\theta_k \in \Theta_k$ *for all* k *and*

$$\sup_{T\in\mathbb{N}, w_k\in\mathcal{W}, \hat{\theta}_0\in\Theta_0} \frac{\sum_{k=0}^{T} \|\tilde{x}_{1|k}\|^2}{\frac{1}{\mu}\|\hat{\theta}_0 - \theta^\star\|^2 + \sum_{k=0}^{T} \|w_k\|^2} \leq 1$$

where $\tilde{x}_{1|k} = A(\theta^\star)x_k + B(\theta^\star)u_k - \hat{x}_{1|k}$ *is the 1-step prediction error*

Control Problem

Consider robust regulation of the system

$$x_{k+1} = A(\theta)x_k + B(\theta)u_k + w_k$$

with $\theta \in \Theta_k$, $w_k \in \mathcal{W}$, subject to the state and control constraints

$$Fx_k + Gu_k \leq \mathbf{1} = [1 \quad \cdots \quad 1]^\top$$

Assumption (Robust stabilizabilty)

There exists a set $\mathcal{X} = \{x : Vx \leq \mathbf{1}\}$ and feedback gain K such that $\mathcal{X}$ is λ-contractive for some $\lambda \in [0, 1)$, i.e.

$$V\Phi(\theta)x \leq \lambda\mathbf{1}, \quad \text{for all } x \in \mathcal{X}, \theta \in \Theta_0.$$

where $\Phi(\theta) = A(\theta) + B(\theta)K$.

Control Problem

State and control input sequences predicted at time k: $u_{i|k}, x_{i|k}$, $i = 0, 1, \ldots$ are expressed in terms of decision variables $\mathbf{v} = (v_{0|k}, \ldots, v_{N|k})$:

$$u_{i|k} = \begin{cases} Kx_{i|k} + v_{i|k} & i = 0, 1, \ldots \\ Kx_{i|k} & \end{cases}$$

The regulation cost is defined in terms of point estimate $\hat{\theta}_k$:

$$J_N(x_k, \hat{\theta}_k, \mathbf{v}_k) = \sum_{i=0}^{N-1} \left(\|\hat{x}_{i|k}\|_Q^2 + \|\hat{u}_{i|k}\|_R^2 \right) + \|\hat{x}_{N|k}\|_P^2$$

where $\hat{x}_{i|k}$, $\hat{u}_{i|k}$ are defined by

$$\begin{aligned} \hat{x}_{i+1|k} &= A(\hat{\theta}_k)\hat{x}_{i|k} + B(\hat{\theta}_k)\hat{u}_{i|k} \\ \hat{u}_{i|k} &= K\hat{x}_{i|k} + v_{i|k} \end{aligned}$$

and $P \succeq \Phi^\top(\theta) P \Phi(\theta) + Q + K^\top R K$ for all $\theta \in \Theta_0$

Tube MPC

A sequence of sets (a "tube") is constructed to bound the predicted state $x_{i|k}$, with ith cross section, $\mathcal{X}_{i|k}$:

$$\mathcal{X}_{i|k} = \{x : Vx \le \alpha_{i|k}\}$$

where V is determined offline and $\alpha_{i|k}$ are online decision variables

- (A) For robust satisfaction of $x_{i|k} \in \mathcal{X}_{i|k}$, we require

  $$V\Phi(\theta)x + VB(\theta)v_{i|k} + \bar{w} \le \alpha_{i+1|k} \quad \text{for all } x \in \mathcal{X}_{i|k},\ \theta \in \Theta_k$$

  where $[\bar{w}]_i = \max_{w \in \mathcal{W}} [V]_i w$

- (B) For robust satisfaction of $Fx_{i|k} + Gu_{i|k} \le \mathbf{1}$, we require

  $$(F + GK)x + Gv_{i|k} \le \mathbf{1} \quad \text{for all } x \in \mathcal{X}_{i|k}$$

Condition (A) is bilinear in x and θ, but it can be expressed in terms of linear inequalities using a vertex representation of either $\mathcal{X}_{i|k}$ or Θ_k

Tube MPC

We generate the vertex representation:

$$\mathcal{X}_{i|k} = \text{co}\{x^1_{i|k}, \dots x^m_{i|k}\}$$

using the property that $\{x : [V]_r x \le [\alpha_{i|k}]_r\}$ is a supporting hyperplane of $\mathcal{X}_{i|k}$ for each r:

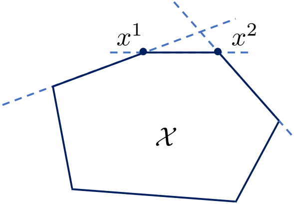

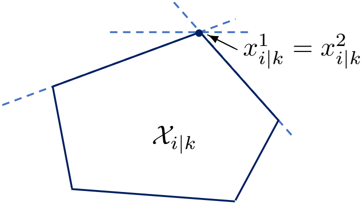

Hence each vertex $x^j_{i|k}$ is given by the intersection of hyperplanes corresponding to a fixed set of rows of V, and

$$x^j_{i|k} = U^j \alpha_{i|k}$$

for some U^j, determined offline from the vertices of $\mathcal{X} = \{x : Vx \le \mathbf{1}\}$

Tube MPC

In terms of both hyperplane and the vertex descriptions of $\mathcal{X}_{i|k}$, the robust tube constraints become

- (A) $V\Phi(\theta)U^j\alpha_{i|k} + VB(\theta)v_{i|k} + \bar{w} \leq \alpha_{i+1|k}$ for all $\theta \in \Theta_k$, $j = 1, \ldots, m$
- (B) $(F + GK)U^j\alpha_{i|k} + Gv_{i|k} \leq \mathbf{1}$, $j = 1, \ldots, m$

Now condition (B) is linear and (A) can be equivalently written as linear constraints using

Lemma (Polyhedral set inclusion)

Let $\mathcal{P}_i = \{x : F_i x \leq f_i\} \subset \mathbb{R}^n$ *for* $i = 1, 2$. *Then* $\mathcal{P}_1 \subseteq \mathcal{P}_2$ *iff*

$$\exists \Lambda \geq 0 \text{ such that } \Lambda F_1 = F_2 \text{ and } \Lambda f_1 \leq f_2$$

Robust MPC online optimization problem

Summary of constraints in the online MPC optimization at time k:

$$V x_k \leq \alpha_{0|k}$$

$$\Lambda^j_{i|k} H_\Theta = V D(U^j \alpha_{i|k}, K U^j \alpha_{i|k} + v_{i|k})$$

$$\Lambda^j_{i|k} h_k \leq \alpha_{i+1|k} - V d(u^j \alpha_{i|k}, K U^j \alpha_{i|k} + v_{i|k}) - \bar{w}$$

$$\Lambda^j_{i|k} \geq 0$$

$$(F + GK) U^j \alpha_{i|k} + G v_{i|k} \leq \mathbf{1}$$

$$\Lambda^j_{N|k} H_\Theta = V D(U^j \alpha_{N|k}, K U^j \alpha_{N|k})$$

$$\Lambda^j_{N|k} h_k \leq \alpha_{N|k} - V d(u^j \alpha_{N|k}, K U^j \alpha_{N|k}) - \bar{w}$$

$$\Lambda^j_{N|k} \geq 0$$

$$(F + GK) U^j \alpha_{N|k} \leq \mathbf{1}$$

for $i = 0, \ldots, N-1$, $j = 1, \ldots, m$

Let $\mathcal{D}(x_k, \Theta_k)$ be the feasible set for the decision variables $\mathbf{v}_k, \boldsymbol{\alpha}_k, \boldsymbol{\Lambda}_k$

Robust adaptive MPC algorithm

Offline: Choose Θ_0, $\mathcal{X}$, feedback gain K, and compute P

Online, at each time $k = 1, 2, \ldots$:

1. Given x_k, update the set (Θ_k) and point ($\hat{\theta}_k$) parameter estimates
2. Compute the solution $(\mathbf{v}_k^*, \boldsymbol{\alpha}_k^*, \boldsymbol{\Lambda}_k^*)$ of the QP

$$\min_{\mathbf{v}_k, \boldsymbol{\alpha}_k, \boldsymbol{\Lambda}_k} J(x_k, \hat{\theta}_k, \mathbf{v}_k)$$
$$\text{subject to } (\mathbf{v}_k, \boldsymbol{\alpha}_k, \boldsymbol{\Lambda}_k) \in \mathcal{D}(x_k, \Theta_k)$$

3. Apply the control law $u_k^* = Kx_k + v_{0|k}^*$

Robust adaptive MPC algorithm

Theorem (Closed loop properties)

If $\theta^\star \in \Theta_0$ *and* $\mathcal{D}(x_0, \Theta_0) \neq \emptyset$*, then for all* $k > 0$*:*

1. $\theta^\star \in \Theta_k$
2. $D(x_k, \Theta_k) \neq \emptyset$
3. $F x_k + G u_k \leq \mathbf{1}$

and the closed loop system is finite-gain l^2*-stable, i.e. there exist constants* $c_0, c_1, c_2 > 0$ *such that for all* T*:*

$$\sum_{k=0}^{T} \|x_k\|^2 \leq c_0 \|x_0\|^2 + c_1 \|\hat{\theta}_0 - \theta^\star\|^2 + c_2 \sum_{k=0}^{T} \|w_k\|^2$$

A numerical example

Second-order linear system with

$$(A(\theta), B(\theta)) = (A_0, B_0) + \sum_{i=1}^{3} (A_i, B_i)\theta_i$$

$$A_0 = \begin{bmatrix} 0.5 & 0.2 \\ -0.1 & 0.6 \end{bmatrix}, \quad B_0 = \begin{bmatrix} 0 \\ 0.5 \end{bmatrix}, \quad A_1 = \begin{bmatrix} 0.042 & 0 \\ 0.072 & 0.03 \end{bmatrix}, \quad B_1 = \begin{bmatrix} 0 \\ 0 \end{bmatrix},$$

$$A_2 = \begin{bmatrix} 0.015 & 0.019 \\ 0.009 & 0.035 \end{bmatrix}, \quad B_2 = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \quad A_3 = \begin{bmatrix} 0 & 0 \\ -0 & 0 \end{bmatrix}, \quad B_3 = \begin{bmatrix} 0.0397 \\ 0.059 \end{bmatrix}.$$

- ▷ true parameter $\theta^\star = [0.8 \;\; 0.2 \;\; -0.5]^\top$, initial set $\Theta_0 = \{\theta : \|\theta\|_\infty \leq 1\}$.
- ▷ disturbance uniformly distributed on $\mathcal{W} = \{w \in \mathbb{R}^2 : \; \|w\|_\infty \leq 0.1\}$, w_k
- ▷ state and input constraints: $[x]_2 \geq -0.3$ and $u_k \leq 1$.

A numerical example: constraint satisfaction

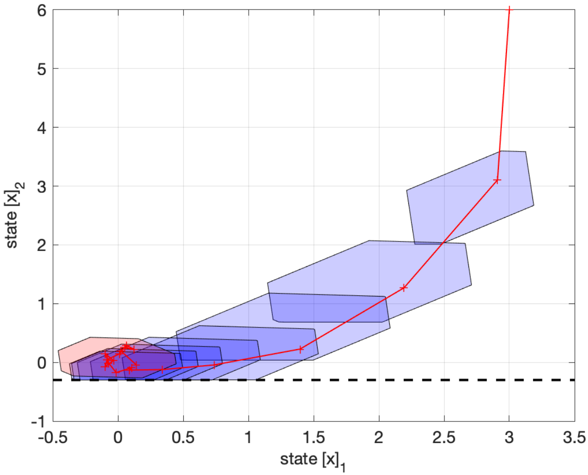

Figure: Realized closed-loop trajectory from initial condition $x_0 = [3\ 6]^\top$ (red line), predicted state tube at time $k = 0$ (tube cross-sections: blue, terminal set: pink)

A numerical example: constraint satisfaction

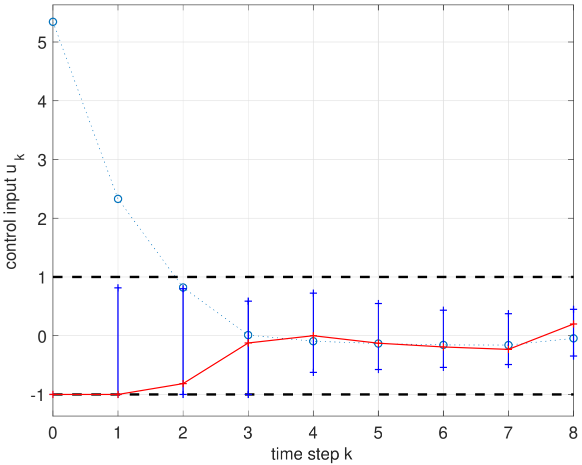

Figure: Realized closed-loop trajectory from initial condition $x_0 = [3\ 6]^\top$ (red line), predicted control tube at time $k = 0$ (tube cross-sections: blue)

Persistent excitation

- A regressor Ψ_k is persistently exciting (PE) if

$$\beta_1{}^2\mathbb{I} \preceq \frac{1}{l}\sum_{k=k_0+1}^{k_0+l} \Psi_k\Psi_k^\top \preceq \beta_2{}^2\mathbb{I}$$

  for some $\beta_1, \beta_2, l > 0$ and all k_0 (Narendra,1987).

- Define the diameter of Θ as $\mathrm{dia}(\Theta) = \sup_{\theta_1,\theta_2\in\Theta} \|\theta_1 - \theta_2\|$

Convergence of set membership parameter estimate

If the noise bound $w \in \mathcal{W}$ is tight and the regressor D_k is persistently exciting, then $\mathrm{dia}(\Theta_k) \to 0$ with probability one [Bai, Cho, Tempo, 1998].

- $\mathcal{W}$ is a tight noise bound if the support of the probability distribution of w is equal to $\mathcal{W}$

Persistent excitation

Regressor: $\Psi_k = D_k^\top = \begin{bmatrix} A_1 x_k + B_1 u_k & \cdots & A_p x_k + B_p u_k \end{bmatrix}^\top$

Consider the PE condition evaluated over a window that includes n past time-steps plus current time:

$$\sum_{k=-n}^{k=0} D_k^\top D_k \succeq \beta_1^2 \mathbb{I}$$

This is nonconvex in $u_0 = K x_k + v_{0|k}$, but we can linearise to obtain a convex condition. Thus, let $u_0 = u_0^* + \delta u$, so that

$$\begin{aligned} D_0^\top D_0 \succeq\; & D(x_0, u_0^*)^\top D(x_0, u_0^*) + D^\top(x_0, u_0^*) \begin{bmatrix} B_1 \delta u & \cdots & B_p \delta u \end{bmatrix} \\ & + \begin{bmatrix} B_1 \delta u & \cdots & B_p \delta u \end{bmatrix}^\top D(x_0, u_0^*) \end{aligned}$$

Therefore a sufficient condition for $\sum_{k=-n}^{k=0} D_k^\top D_k \succeq \beta_1^2 \mathbb{I}$ is an LMI in δu:

$$\begin{aligned} \text{LMI}(\delta u) : & \sum_{k=-n}^{k=-1} D_k^\top D_k + D(x_0, u_0^*)^\top D(x_0, u_0^*) \\ & + D(x_0, u_0^*)^\top \begin{bmatrix} B_1 \delta u & \cdots & B_p \delta u \end{bmatrix} + \begin{bmatrix} B_1 \delta u & \cdots & B_p \delta u \end{bmatrix}^\top D(x_0, u_0^*) \succeq \beta_1^2 \mathbb{I} \end{aligned}$$

Robust adaptive MPC algorithm with PE constraint

Offline: Choose Θ_0, $\mathcal{X}$, β_1, feedback gain K, and compute P

Online, at each time $k = 1, 2, \ldots$:

1. Given x_k, update the set (Θ_k) and point ($\hat{\theta}_k$) parameter estimates

2. Compute the solution $(\mathbf{v}_k^*, \boldsymbol{\alpha}_k^*, \boldsymbol{\Lambda}_k^*)$ of the QP

$$\min_{\mathbf{v}_k, \boldsymbol{\alpha}_k, \boldsymbol{\Lambda}_k} J(x_k, \hat{\theta}_k, \mathbf{v}_k)$$

$$\text{subject to } (\mathbf{v}_k, \boldsymbol{\alpha}_k, \boldsymbol{\Lambda}_k) \in \mathcal{D}(x_k, \Theta_k)$$

3. If $\displaystyle\sum_{k=-n}^{k=-1} D_k^\top D_k + D(x_0, u_0^*)^\top D(x_0, u_0^*) \not\succeq \beta_1 \mathbb{I}$:
   (a) Re-run the MPC optimization with $v_{0|k} = v_{0|k}^* + \delta u$ and LMI(δu) as additional constraints
   (b) If a feasible solution exists, set $v_{0|k}^* \leftarrow v_{0|k}^* + \delta u^*$

4. Apply the control law $u_k^* = K x_k + v_{0|k}^*$

A numerical example: parameter set

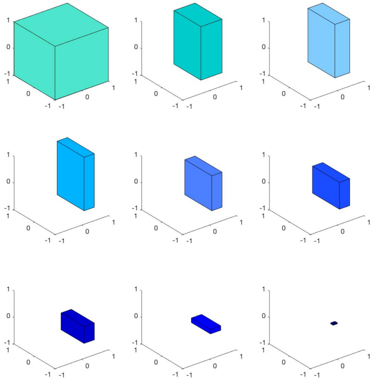

Figure: Parameter set Θ_k at time steps $k \in \{0, 1, 2; 10, 25, 50; 100, 500, 5000\}$

Θ set	Volume (%)	Cost*
Θ_0	100	62.22
Θ_1	26.1	61.13
Θ_2	18.3	61.03
Θ_{10}	12.7	60.96
Θ_{25}	8.3	60.93
Θ_{50}	6.3	60.77
Θ_{100}	3.4	59.45
Θ_{500}	0.7	57.94
Θ_{5000}	0.0089	53.95
$\theta^\star$	-	52.70

Table: Volume of Θ_k as $\Theta_k/\Theta_0 \times 100\%$; Cost* with same initial x_0 and constraints

Time-varying parameters

Assumption (time-varying parameters)

There exists a constant r_θ *such that the parameter vector* $\theta_k^\star$ *satisfies* $\theta_k^\star \in \Theta_0$ *for all* k *and* $\|\theta_{k+1}^\star - \theta_k^\star\| \leq r_\theta$

Define the dilation operator:

$$R_i(\Theta) = \{\theta : H_\Theta \theta \leq h + i r_\theta \mathbf{1}\}$$

Then the parameter set update can be expressed

$$\Theta_k = R_1(\Theta_{k-1} \cap \Delta_k) \cap \Theta_0$$

and Θ_k is replaced in the tube MPC constraints by

$$\Theta_{i|k} = R_i(\Theta_k) \cap \Theta_0$$

Robust adaptive MPC algorithm with time-varying parameters

Parameter estimate bounds and recursive feasibility properties are unchanged:

Theorem (Closed loop properties)

If $\theta^\star \in \Theta_0$ and $\mathcal{D}(x_0, \Theta_0) \neq \emptyset$, then for all $k > 0$:

1. $\theta^\star \in \Theta_k$
2. $D(x_k, \Theta_k) \neq \emptyset$
3. $F x_k + G u_k \leq \mathbf{1}$

But the LMS filter has an additional tracking error, which invalidates the l^2 properties, i.e. "certainty equivalence" no longer applies

However other performance measures can be used in this context, such as the min-max approach of [Lorenzen, Allgöwer, Cannon, 2017]

A numerical example: time-varying parameters

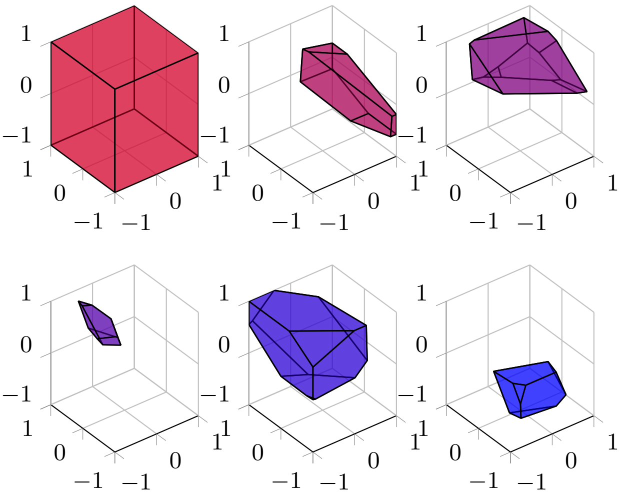

Figure: Parameter set Θ_k at times $k \in \{0, 100, 200, 300, 400, 500\}$ for the time-varying system with $r_\theta = 0.01$

A numerical example: time-varying parameters

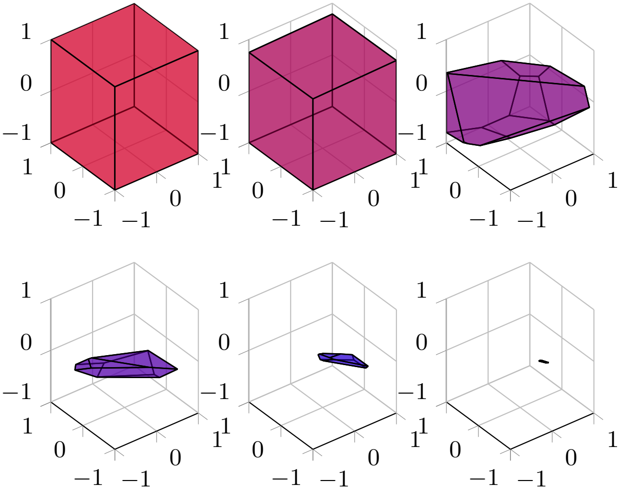

Figure: Parameter set Θ_k at times $k \in \{0, 5, 25, 70, 120, 500\}$ for the non-time-varying case for comparison

Conclusions & Outlook

Conclusions:

- Adaptive robust MPC with closed loop guarantees is computationally tractable
- Set-membership parameter estimation and LMS point estimates are obvious choices for MPC cost functions and robust constraints
- Nonconvex PE conditions can be relaxed to convex sufficient conditions

Future work

- How to ensure recursive feasibility with PE constraints?
- Are PE conditions better handled by adding terms to the MPC cost (similar to MPC-based dual control)?
- How can we relax the requirement of prior knowledge of a robustly stabilizing local feedback law?

References:

– M. Lorenzen, M. Cannon, & F. Allgöwer, "Robust MPC with recursive model update" Automatica (to appear)

– X. Lu, M. Cannon, "Robust adaptive tube model predictive control" ACC19 (submitted)