Non-Cooperative Distributed MPC with Iterative Learning

Haimin Hu[1], Konstantinos Gatsis[2], Manfred Morari[1], George J. Pappas[1]

[1]University of Pennsylvania
[2]University of Oxford

Acknowledgement: Work supported by NSF and Intel ISTC-WAS

Overview

Background

- Large-scale distributed systems
- Modeling of distributed systems
- Distributed model predictive control (DMPC)

Contribution

- Distributed learning model predictive control (DLMPC)
- Algorithms and properties of DLMPC
- Numerical Example

Control of Large-scale Interconnected Systems

Large-scale distributed systems

- Centralized solution not possible
- Interconnected plants with complex interactions
- Big data available

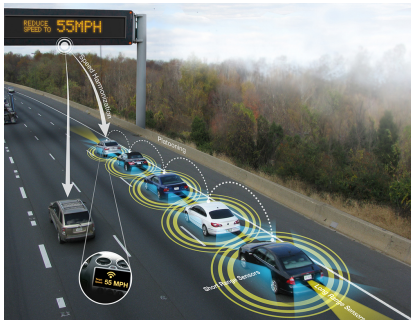

Control of Large-scale Interconnected Systems

Large-scale distributed systems

- Centralized solution not possible
- Interconnected plants with complex interactions
- Big data available

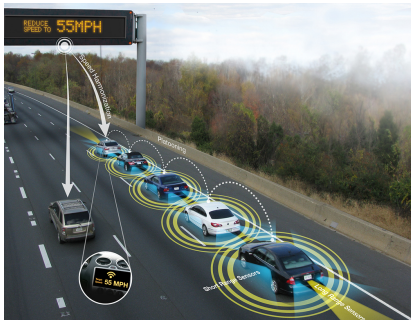

Goal of our work

- Use **data** to improve performance
- Design **distributed** control law to maintain operational scalability

Modeling of Distributed Systems

Global system: $\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t \quad \text{s.t. } \mathbf{x}_t \in \mathbb{X},\ \mathbf{u}_t \in \mathbb{U}$

Partitioned subsystem: $\Sigma_i : x_{t+1}^{[i]} = A_{ii}x_t^{[i]} + B_i u_t^{[i]} + \sum_{j \in \mathcal{N}_i} A_{ij} x_t^{[j]}$

$$\text{s.t. } x_t^{[i]} \in \mathbb{X}^{[i]},\ u_t^{[i]} \in \mathbb{U}^{[i]}$$

Example:

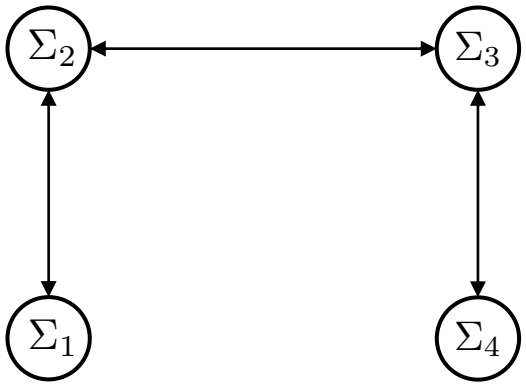

$$\mathbf{A} = \begin{bmatrix} A_{11} & A_{12} & 0 & 0 \\ A_{21} & A_{22} & A_{23} & 0 \\ 0 & A_{32} & A_{33} & A_{34} \\ 0 & 0 & A_{43} & A_{44} \end{bmatrix}$$

$$\mathbf{B} = \begin{bmatrix} B_1 & 0 & 0 & 0 \\ 0 & B_2 & 0 & 0 \\ 0 & 0 & B_3 & 0 \\ 0 & 0 & 0 & B_4 \end{bmatrix}$$

Reference Trajectory

Partitioned subsystem: $\Sigma_i : x_{t+1}^{[i]} = A_{ii} x_t^{[i]} + B_i u_t^{[i]} + \sum_{j \in \mathcal{N}_i} A_{ij} x_t^{[j]}$

- We want to get rid of neighbors' states $x_t^{[j]}$

Reference Trajectory

Partitioned subsystem: $\Sigma_i : x_{t+1}^{[i]} = A_{ii} x_t^{[i]} + B_i u_t^{[i]} + \sum_{j \in \mathcal{N}_i} A_{ij} x_t^{[j]}$

- We want to get rid of neighbors' states $x_t^{[j]}$

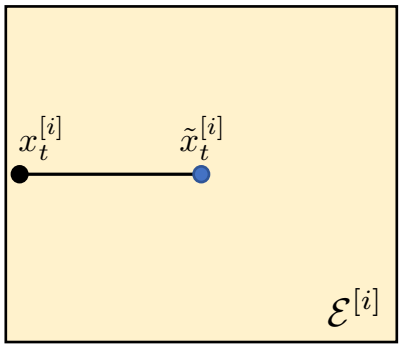

Each subsystem Σ_i

- Maintains a reference trajectory: $\tilde{x}_t^{[i]},\ \tilde{x}_{t+1}^{[i]}, \ldots$
- Ensures that $x_t^{[i]} - \tilde{x}_t^{[i]} \in \mathcal{E}^{[i]}$
- Transmits the reference trajectory to neighbors at each time

Reference Trajectory

Partitioned subsystem: $\Sigma_i : x_{t+1}^{[i]} = A_{ii}x_t^{[i]} + B_i u_t^{[i]} + \sum_{j\in\mathcal{N}_i} A_{ij}x_t^{[j]}$

- We want to get rid of neighbors' states $x_t^{[j]}$

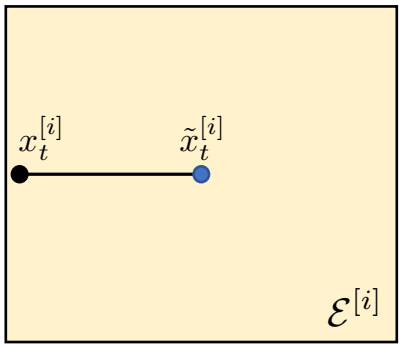

Each subsystem Σ_i

- Maintains a reference trajectory: $\tilde{x}_t^{[i]},\ \tilde{x}_{t+1}^{[i]}, \ldots$
- Ensures that $x_t^{[i]} - \tilde{x}_t^{[i]} \in \mathcal{E}^{[i]}$
- Transmits the reference trajectory to neighbors at each time

"Perturbed" subsystem:

$$x_{t+1}^{[i]} = A_{ii}x_t^{[i]} + B_i u_t^{[i]} + \sum_{j\in\mathcal{N}_i} A_{ij}\tilde{x}_t^{[j]} + \sum_{j\in\mathcal{N}_i} A_{ij}\underbrace{\left(x_t^{[j]} - \tilde{x}_t^{[j]}\right)}_{\in\mathcal{E}_j}$$

$$= A_{ii}x_t^{[i]} + B_i u_t^{[i]} + c_t^{[i]} + \underbrace{w_t^{[i]}}_{\in\mathbb{W}^{[i]} = \bigoplus_{j\in\mathcal{N}_i} A_{ij}\mathcal{E}^{[j]}}$$

known input

Nominal Subsystem

"Perturbed" subsystem: $x_{t+1}^{[i]} = A_{ii} x_t^{[i]} + B_i u_t^{[i]} + c_t^{[i]} + w_t^{[i]}$

Nominal subsystem: $z_{t+1}^{[i]} = A_{ii} z_t^{[i]} + B_i v_t^{[i]} + c_t^{[i]}$

Nominal Subsystem

"Perturbed" subsystem: $x_{t+1}^{[i]} = A_{ii}x_t^{[i]} + B_i u_t^{[i]} + c_t^{[i]} + w_t^{[i]}$

Nominal subsystem: $z_{t+1}^{[i]} = A_{ii}z_t^{[i]} + B_i v_t^{[i]} + c_t^{[i]}$

- Input parameterization: $u_t^{[i]} = v_t^{[i]} + K_i\left(x_t^{[i]} - z_t^{[i]}\right)$
- Define error states: $e_t^{[i]} = x_t^{[i]} - z_t^{[i]}$

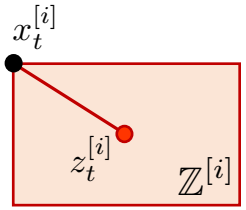

Error subsystem: $e_{t+1}^{[i]} = \underbrace{(A_{ii} + B_i K_i)}_{\Phi_i} e_t^{[i]} + w_t^{[i]}$

If Φ_i is stable: $\in \mathbb{Z}^{[i]}$ (the RPI set) $\quad \forall t \geq 0,\ \forall w_t^{[i]} \in \mathbb{W}^{[i]}$

Nominal Subsystem

"Perturbed" subsystem: $x_{t+1}^{[i]} = A_{ii}x_t^{[i]} + B_i u_t^{[i]} + c_t^{[i]} + w_t^{[i]}$

Nominal subsystem: $z_{t+1}^{[i]} = A_{ii}z_t^{[i]} + B_i v_t^{[i]} + c_t^{[i]}$

- Input parameterization: $u_t^{[i]} = v_t^{[i]} + K_i\left(x_t^{[i]} - z_t^{[i]}\right)$
- Define error states: $e_t^{[i]} = x_t^{[i]} - z_t^{[i]}$

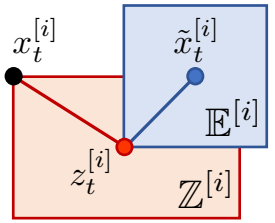

Error subsystem: $e_{t+1}^{[i]} = \underbrace{(A_{ii} + B_i K_i)}_{\Phi_i} e_t^{[i]} + w_t^{[i]}$

If Φ_i is stable: $\in \mathbb{Z}^{[i]}$ (the RPI set) $\quad \forall t \geq 0,\ \forall w_t^{[i]} \in \mathbb{W}^{[i]}$

$z_t^{[i]} - \tilde{x}_t^{[i]} \in \mathbb{E}^{[i]}$ "Bounded incompliance"
(will appear in MPC constraints)

Nominal Subsystem

"Perturbed" subsystem: $x_{t+1}^{[i]} = A_{ii}x_t^{[i]} + B_i u_t^{[i]} + c_t^{[i]} + w_t^{[i]}$

Nominal subsystem: $z_{t+1}^{[i]} = A_{ii}z_t^{[i]} + B_i v_t^{[i]} + c_t^{[i]}$

- Input parameterization: $u_t^{[i]} = v_t^{[i]} + K_i\left(x_t^{[i]} - z_t^{[i]}\right)$
- Define error states: $e_t^{[i]} = x_t^{[i]} - z_t^{[i]}$

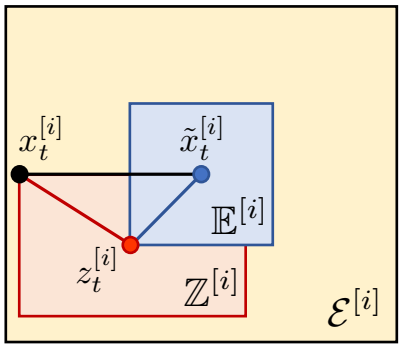

Error subsystem: $e_{t+1}^{[i]} = \underbrace{(A_{ii} + B_i K_i)}_{\Phi_i} e_t^{[i]} + w_t^{[i]}$

If Φ_i is stable: $\in \mathbb{Z}^{[i]}$ (the RPI set) $\quad \forall t \geq 0,\ \forall w_t^{[i]} \in \mathbb{W}^{[i]}$

Designing the sets:

$$\mathbb{E}^{[i]} \oplus \mathbb{Z}^{[i]} \subseteq \mathcal{E}^{[i]}$$

$$\left.\begin{array}{l} z_t^{[i]} - \tilde{x}_t^{[i]} \in \mathbb{E}^{[i]} \\ x_t^{[i]} - z_t^{[i]} \in \mathbb{Z}^{[i]} \end{array}\right\} \Rightarrow\ x_t^{[i]} - \tilde{x}_t^{[i]} \in \mathcal{E}^{[i]}$$

Robust constraint satisfaction (for each subsystem Σ_i):

$$\left.\begin{array}{l} z_t^{[i]} \in \mathbb{X}^{[i]} \ominus \mathbb{Z}^{[i]},\ v_t^{[i]} \in \mathbb{U}^{[i]} \ominus K_i\mathbb{Z}^{[i]} \\ z_t^{[i]} - \tilde{x}_t^{[i]} \in \mathbb{E}^{[i]} \end{array}\right. \Rightarrow\ x_t^{[i]} \in \mathbb{X}^{[i]}, u_t^{[i]} \in \mathbb{U}^{[i]}, \quad \forall w_t^{[i]} \in \mathbb{W}^{[i]}$$

Nominal Subsystem

"Perturbed" subsystem: $x_{t+1}^{[i]} = A_{ii}x_t^{[i]} + B_i u_t^{[i]} + c_t^{[i]} + w_t^{[i]}$

Nominal subsystem: $z_{t+1}^{[i]} = A_{ii}z_t^{[i]} + B_i v_t^{[i]} + c_t^{[i]}$

Reference on robust MPC:

L. Chisci, J.A. Rossiter and G. Zappa, "Systems with persistent disturbances: predictive control with restricted constraints," in *Automatica*, 37(7), pp. 1019–1028, 2001.

D.Q. Mayne, M.M. Seron and S.V. Raković, "Robust model predictive control of constrained linear systems with bounded disturbances," in *Automatica*, 41(2), pp. 219-224, 2005.

Reference on distributed MPC:

W.B. Dunbar and R.M. Murray, "Distributed receding horizon control for multi-vehicle formation stabilization," in *Automatica*, 42(4), pp. 549-558, 2006.

C. Conte, N.R. Voellmy, M.N. Zeilinger, M. Morari and C.N. Jones, "Distributed synthesis and control of constrained linear systems," in *American Control Conference (ACC)*, 2012, pp. 6017-6022. IEEE.

M. Farina and R. Scattolini, "Distributed predictive control: A non-cooperative algorithm with neighbor-to-neighbor communication for linear systems," in *Automatica*, 48(6), pp. 1088-1096, 2012.

Problem Setup

Infinite horizon non-cooperative control sub-problem:

$$
\begin{aligned}
\min_{\bar{v}^{[i]}_{[0:\infty]}} \quad & \sum_{k=0}^{\infty} h_i(\bar{z}^{[i]}_k, \bar{v}^{[i]}_k) \\
\text{s.t.} \quad & \bar{z}^{[i]}_{k+1} = A_{ii}\bar{z}^{[i]}_k + B_i\bar{v}^{[i]}_k + c^{[i]}_k \\
& \bar{z}^{[i]}_k \in \bar{\mathbb{X}}^{[i]}, \ \bar{v}^{[i]}_k \in \bar{\mathbb{U}}^{[i]} \\
& \bar{z}^{[i]}_k - \tilde{x}^{[i]}_k \in \mathbb{E}^{[i]} \\
& \bar{z}^{[i]}_0 = z^{[i]}_0, \ k = 0, 1, \ldots,
\end{aligned}
$$

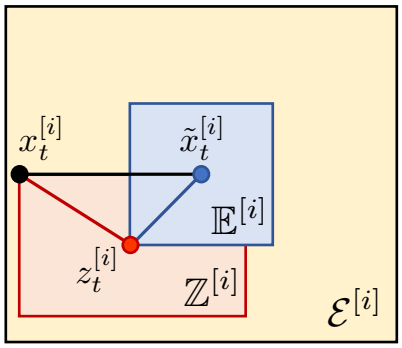

- The solution is optimal for the nominal model of Σ_i given $\tilde{x}^{[i]}_{[0:\infty]}$ and $c^{[i]}_{[0:\infty]}$
- Needs approximation to avoid the infinite horizon

Non-cooperative Distributed MPC

Non-cooperative DMPC[1]:

$$
\begin{aligned}
\min_{\bar{v}^{[i]}_{[t:t+N-1]}} \quad & \sum_{k=t}^{t+N-1} h_i(\bar{z}^{[i]}_k, \bar{v}^{[i]}_k) + P_i(\bar{z}^{[i]}_{t+N}) \\
\text{s.t.} \quad & \bar{z}^{[i]}_{k+1} = A_{ii}\bar{z}^{[i]}_k + B_i\bar{v}^{[i]}_k + c^{[i]}_k \\
& \bar{z}^{[i]}_k \in \bar{\mathbb{X}}^{[i]}, \ \bar{v}^{[i]}_k \in \bar{\mathbb{U}}^{[i]} \\
& \bar{z}^{[i]}_k - \tilde{x}^{[i]}_k \in \mathbb{E}^{[i]} \\
& \bar{z}^{[i]}_t = z^{[i]}_t, \ \bar{z}^{[i]}_{t+N} \in \mathbb{X}^{[i]}_F \\
& k = t, \ldots, t+N-1
\end{aligned}
$$

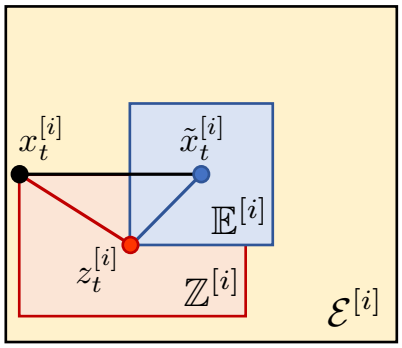

- **Advantages**: Scalability (low-dimensional sub-problems, efficient communication)

[1] M. Farina and R. Scattolini, "Distributed predictive control: A non-cooperative algorithm with neighbor-to-neighbor communication for linear systems," in *Automatica*, 48(6), pp. 1088-1096, 2012.

Non-cooperative Distributed MPC

Non-cooperative DMPC[1]:

$$
\begin{aligned}
\min_{\bar{v}^{[i]}_{[t:t+N-1]}} \quad & \sum_{k=t}^{t+N-1} h_i(\bar{z}^{[i]}_k, \bar{v}^{[i]}_k) + \boxed{P_i(\bar{z}^{[i]}_{t+N})} \\
\text{s.t.} \quad & \bar{z}^{[i]}_{k+1} = A_{ii}\bar{z}^{[i]}_k + B_i\bar{v}^{[i]}_k + c^{[i]}_k \\
& \bar{z}^{[i]}_k \in \bar{\mathbb{X}}^{[i]}, \ \bar{v}^{[i]}_k \in \bar{\mathbb{U}}^{[i]} \\
& \bar{z}^{[i]}_k - \tilde{x}^{[i]}_k \in \mathbb{E}^{[i]} \\
& \bar{z}^{[i]}_t = z^{[i]}_t, \ \bar{z}^{[i]}_{t+N} \in \boxed{\mathbb{X}^{[i]}_F} \\
& k = t, \ldots, t+N-1
\end{aligned}
$$

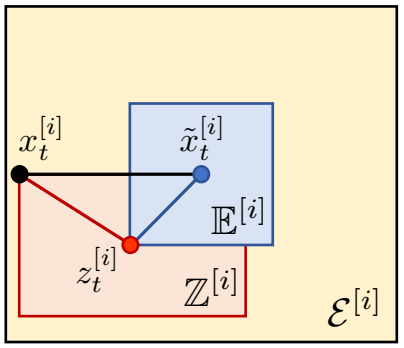

- **Advantages**: Scalability (low-dimensional sub-problems, efficient communication)
- **Limitations**:
  - Terminal sets and costs can be difficult to construct
  - Nominal trajectory is suboptimal for the infinite-horizon control problem

Non-cooperative Distributed MPC

Non-cooperative DMPC[1]:

$$
\begin{aligned}
\min_{\bar{v}^{[i]}_{[t:t+N-1]}} \quad & \sum_{k=t}^{t+N-1} h_i(\bar{z}^{[i]}_k, \bar{v}^{[i]}_k) + \boxed{P_i(\bar{z}^{[i]}_{t+N})} \\
\text{s.t.} \quad & \bar{z}^{[i]}_{k+1} = A_{ii}\bar{z}^{[i]}_k + B_i\bar{v}^{[i]}_k + c^{[i]}_k \\
& \bar{z}^{[i]}_k \in \bar{\mathbb{X}}^{[i]}, \ \bar{v}^{[i]}_k \in \bar{\mathbb{U}}^{[i]} \\
& \bar{z}^{[i]}_k - \tilde{x}^{[i]}_k \in \mathbb{E}^{[i]} \\
& \bar{z}^{[i]}_t = z^{[i]}_t, \ \bar{z}^{[i]}_{t+N} \in \boxed{\mathbb{X}^{[i]}_F} \\
& k = t, \ldots, t+N-1
\end{aligned}
$$

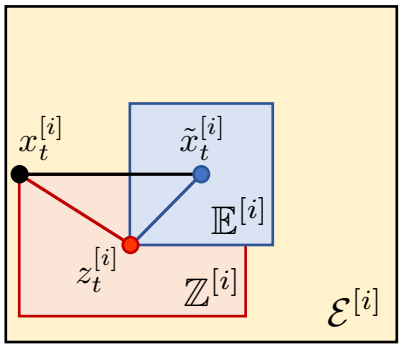

- **Advantages**: Scalability (low-dimensional sub-problems, efficient communication)
- **Limitations**:
  - Terminal sets and costs can be difficult to construct
  - Nominal trajectory is suboptimal for the infinite horizon control problem
- **This work**: Use **historical data** to overcome the above limitations, i.e. to **ensure safety** and **improve performance** for non-cooperative DMPC

Overview

Background

- Large-scale distributed systems
- Modeling of distributed systems
- Distributed model predictive control (DMPC)

Contribution

- Distributed learning model predictive control (DLMPC)
- Algorithms and properties of DLMPC
- Numerical Example

Iterative Leaning for DMPC: Setup

Each subsystem Σ_i :

- Same initial state $x_0^{[i]}$ and terminal state (the origin) at each iteration
- Same reference trajectory $\tilde{x}_{[0:\infty]}^{[i]}$ at each iteration
- Closed-loop trajectories at the l-th iteration:

  - Nominal subsystem: $\left(z_{[0:\infty]}^{l,[i]}, v_{[0:\infty]}^{l,[i]}\right)$

  - Actual subsystem: $\left(x_{[0:\infty]}^{l,[i]}, u_{[0:\infty]}^{l,[i]}\right)$

Assumption on the reference trajectory:

There exists a **finite time** t' such that $\tilde{x}_t^{[i]} = 0,\ \forall i,\ \forall t \geq t'$. Can be generated by [1].

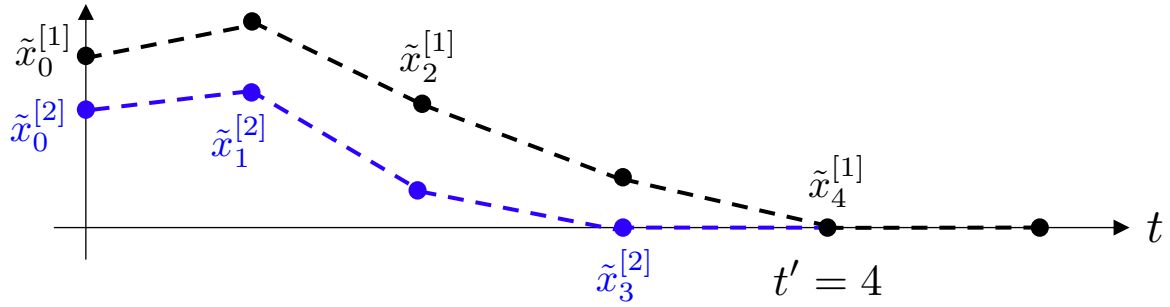

Iterative Learning MPC

LMPC for LTI systems[2]**:** the convex safe set $\mathcal{CS}^l = \text{conv}\left(\left\{\bigcup_{h=0}^{l}\bigcup_{t=0}^{\infty} x_t^h\right\}\right)$

is control invariant due to
1. Linearity of dynamics
2. Convexity and time-invariance of the MPC problem

[2] U. Rosolia, X. Zhang, and F. Borrelli, "Learning model predictive control for iterative tasks: A computationally efficient approach for linear system," in *CDC*. IEEE, 2017, pp. 1157–1162.

Iterative Learning MPC

LMPC for LTI systems[2]**:** the convex safe set $\mathcal{CS}^l = \text{conv}\left(\left\{\bigcup_{h=0}^{l}\bigcup_{t=0}^{\infty} x_t^h\right\}\right)$

is control invariant due to
1. Linearity of dynamics
2. Convexity and time-invariance of the MPC problem

In our case:

- The prediction model is nonlinear: $z_{t+1}^{[i]} = A_{ii} z_t^{[i]} + B_i v_t^{[i]} + c_t^{[i]}$
- Time-varying constraint in MPC: $z_t^{[i]} \in \mathbb{E}_t^{[i]} = \tilde{x}_t^{[i]} \oplus \mathbb{E}^{[i]}$

[2] U. Rosolia, X. Zhang, and F. Borrelli, "Learning model predictive control for iterative tasks: A computationally efficient approach for linear system," in *CDC*. IEEE, 2017, pp. 1157–1162.

Iterative Learning MPC

LMPC for LTI systems[2]**:** the convex safe set $\mathcal{CS}^l = \text{conv}\left(\left\{\bigcup_{h=0}^{l}\bigcup_{t=0}^{\infty} x_t^h\right\}\right)$

is control invariant due to
1. Linearity of dynamics
2. Convexity and time-invariance of the MPC problem

In our case:

- The prediction model is nonlinear: $z_{t+1}^{[i]} = A_{ii} z_t^{[i]} + B_i v_t^{[i]} + c_t^{[i]}$
- Time-varying constraint in MPC: $z_t^{[i]} \in \mathbb{E}_t^{[i]} = \tilde{x}_t^{[i]} \oplus \mathbb{E}^{[i]}$

Solution:

- Introduce **time-varying** convex safe sets: $\mathcal{CS}_t^{l,[i]} = \text{conv}\left(\left\{\bigcup_{h=0}^{l} z_t^{h,[i]}\right\}\right)$

[2] U. Rosolia, X. Zhang, and F. Borrelli, "Learning model predictive control for iterative tasks: A computationally efficient approach for linear system," in *CDC*. IEEE, 2017, pp. 1157–1162.

Time-varying Safe Sets and Cost Functions

Time-varying convex safe set:

$$\mathcal{CS}_t^{l,[i]} := \begin{cases} \text{conv}\left(\left\{\bigcup_{h=0}^{l} z_t^{h,[i]}\right\}\right) & \text{if } t < t' \\ \mathcal{CS}_\infty^{l,[i]} = \text{conv}\left(\left\{\bigcup_{h=0}^{l}\bigcup_{t\geq t'} z_t^{h,[i]}\right\}\right) & \text{otherwise} \end{cases}$$

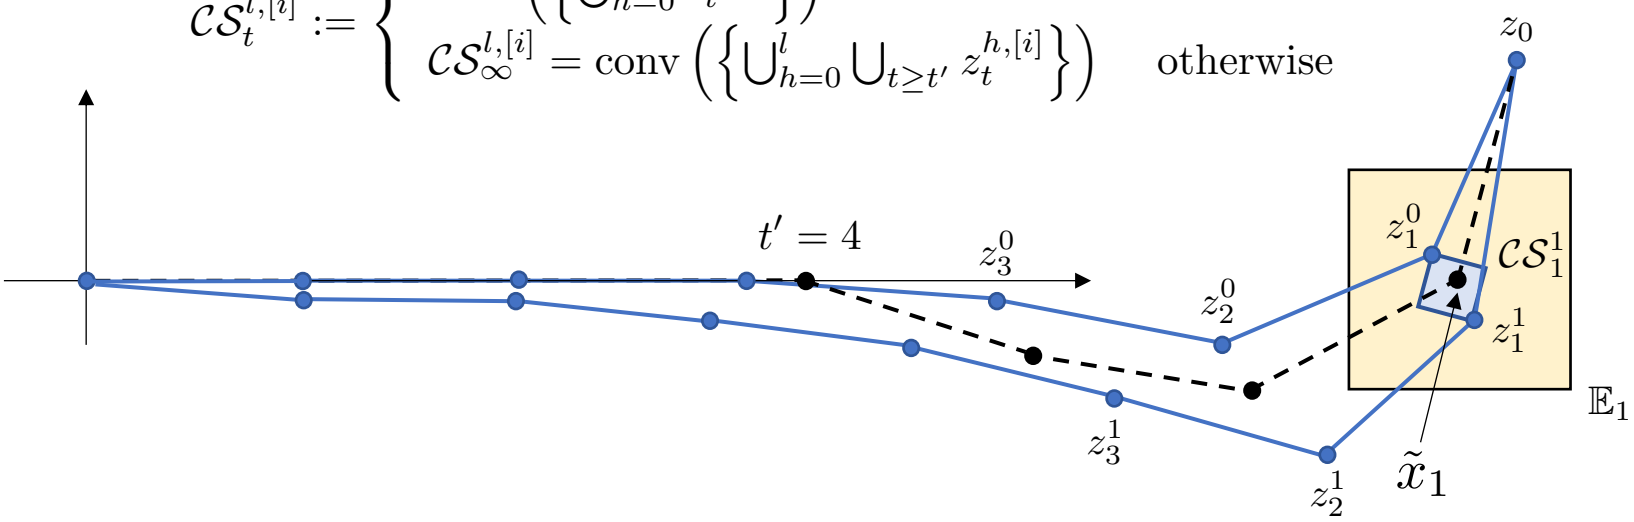

Convex combination of points in $\mathcal{CS}_1^1$ satisfies:

- Nonlinear dynamics: $z_{t+1}^{h,[i]} = A_{ii} z_t^{h,[i]} + B_i v_t^{h,[i]} + c_t^{[i]}$
- Time-varying constraint: $z_t^{h,[i]} \in \mathbb{E}_t^{[i]}$

Time-varying Safe Sets and Cost Functions

Time-varying convex safe set:

$$\mathcal{CS}_t^{l,[i]} := \begin{cases} \text{conv}\left(\left\{\bigcup_{h=0}^{l} z_t^{h,[i]}\right\}\right) & \text{if } t < t' \\ \mathcal{CS}_\infty^{l,[i]} = \text{conv}\left(\left\{\bigcup_{h=0}^{l}\bigcup_{t\geq t'} z_t^{h,[i]}\right\}\right) & \text{otherwise} \end{cases}$$

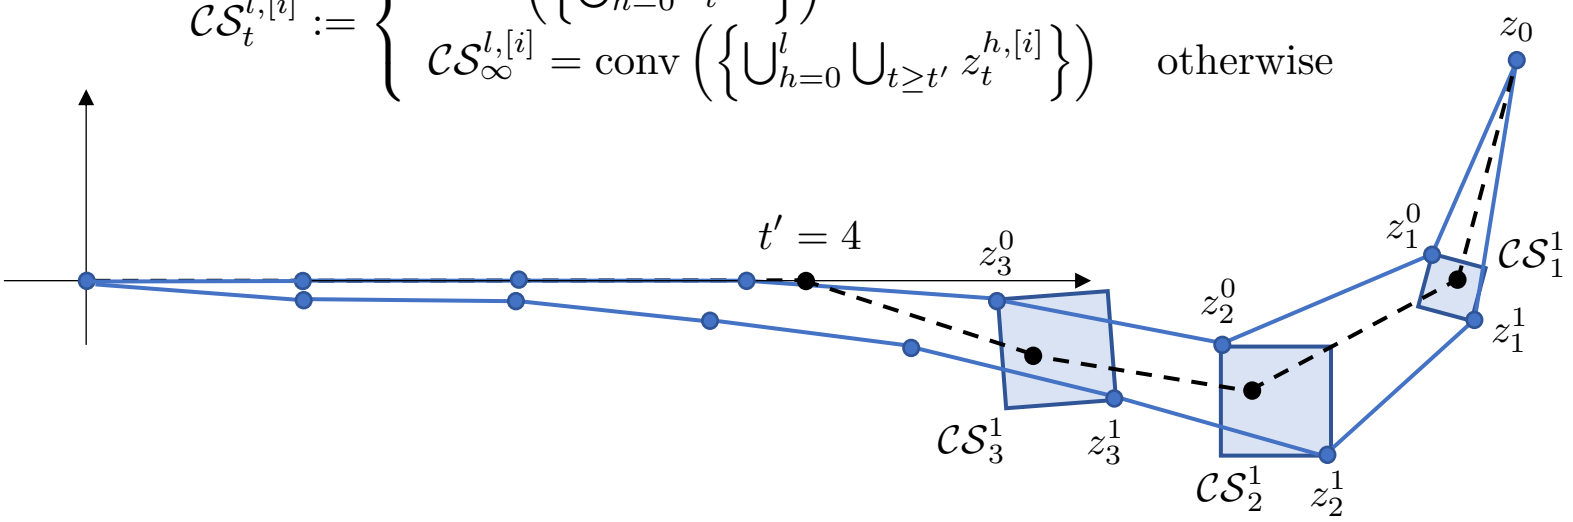

Convex combination of points in $\mathcal{CS}_1^1$ satisfies:

- Nonlinear dynamics: $z_{t+1}^{h,[i]} = A_{ii} z_t^{h,[i]} + B_i v_t^{h,[i]} + c_t^{[i]}$
- Time-varying constraint: $z_t^{h,[i]} \in \mathbb{E}_t^{[i]}$

Time-varying Safe Sets and Cost Functions

Time-varying convex safe set:

$$\mathcal{CS}_t^{l,[i]} := \begin{cases} \text{conv}\left(\left\{\bigcup_{h=0}^{l} z_t^{h,[i]}\right\}\right) & \text{if } t < t' \\ \mathcal{CS}_\infty^{l,[i]} = \text{conv}\left(\left\{\bigcup_{h=0}^{l}\bigcup_{t\geq t'} z_t^{h,[i]}\right\}\right) & \text{otherwise} \end{cases}$$

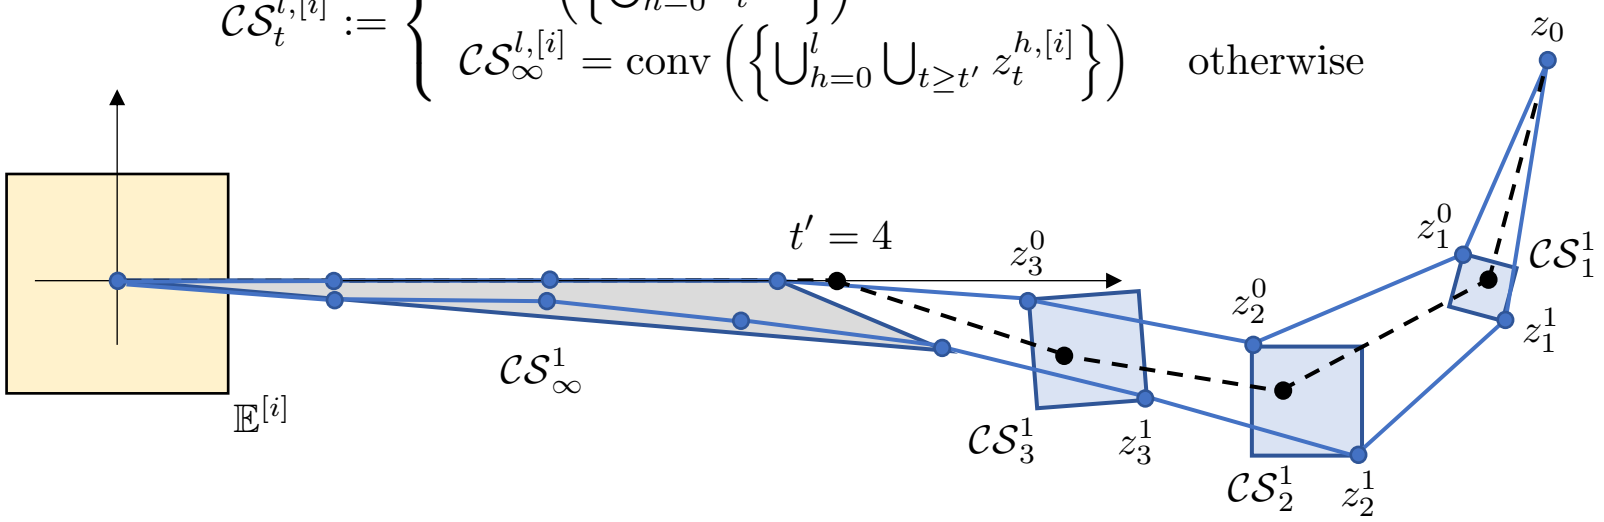

Key observations for $t \geq t'$ **:**

- LTI dynamics: $z_{t+1}^{h,[i]} = A_{ii} z_t^{h,[i]} + B_i v_t^{h,[i]}$
- Time-invariant constraint: $z_t^{h,[i]} \in \mathbb{E}^{[i]}$

➢ The same as: [2] Rosolia et al., "Learning model predictive control for iterative tasks: A computationally efficient approach for linear system."

Time-varying Safe Sets and Cost Functions

Time-varying convex safe set:

$$\mathcal{CS}_t^{l,[i]} := \begin{cases} \operatorname{conv}\left(\left\{\bigcup_{h=0}^{l} z_t^{h,[i]}\right\}\right) & \text{if } t < t' \\ \mathcal{CS}_\infty^{l,[i]} = \operatorname{conv}\left(\left\{\bigcup_{h=0}^{l} \bigcup_{t \geq t'} z_t^{h,[i]}\right\}\right) & \text{otherwise} \end{cases}$$

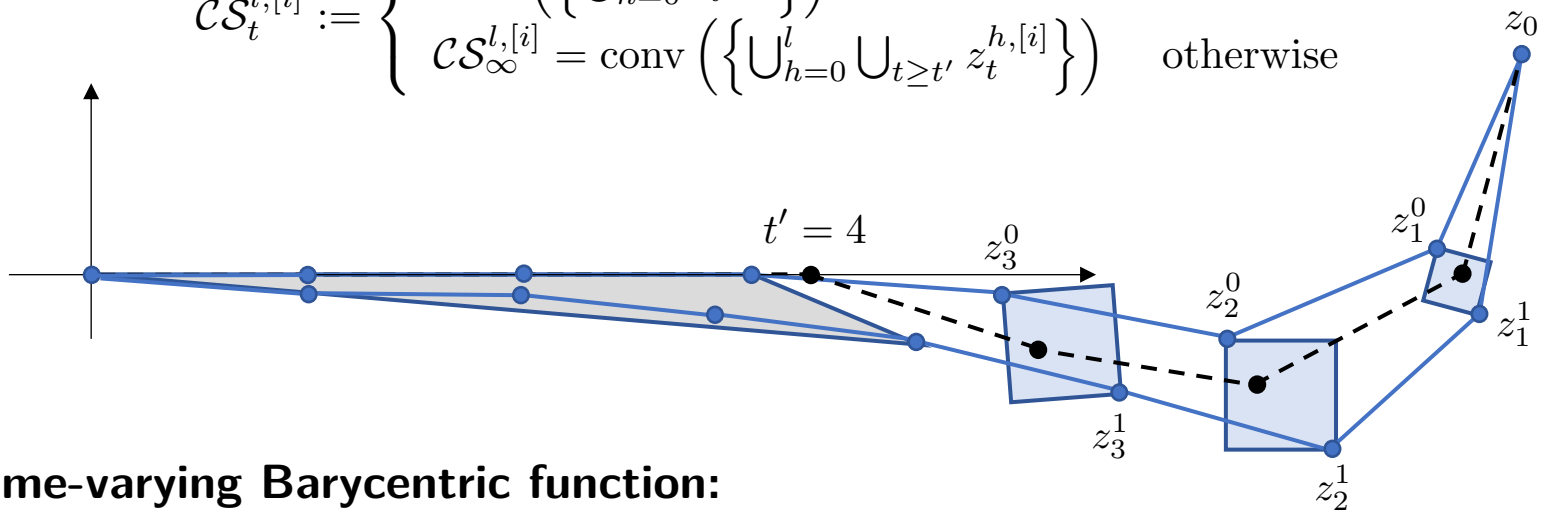

Time-varying Barycentric function:

$$P_t^{l,[i]}(z_t) := \begin{cases} p_t^{l,\star}(z_t) & \text{if } z_t \in \mathcal{CS}_t^{l,[i]} \\ +\infty & \text{if } z_t \notin \mathcal{CS}_t^{l,[i]} \end{cases}$$

where $p_t^{l,\star}(z_t) := \min_{\lambda^h \geq 0} \sum_{h=0}^{l} \lambda^h J_{t\to\infty}^{h,[i]}(z_t^{h,[i]})$ ← empirical cost-to-go

$$\text{s.t. } \sum_{h=0}^{l} \lambda^h = 1, \ \sum_{h=0}^{l} \lambda^h z_t^{h,[i]} = z_t$$

Distributed Learning MPC (DLMPC)

Each subsystem $\sum_i$ solves:

$$
\begin{aligned}
&J^{l,[i],\star}_{t\to t+N}(z^{l,[i]}_t) = \\
&\min_{\bar{v}^{l,[i]}_{[t:t+N-1]}} \quad \sum_{k=t}^{t+N-1} h_i(\bar{z}^{[i]}_k, \bar{v}^{[i]}_k) + P^{l-1,[i]}_{t+N}(\bar{z}^{[i]}_{t+N}) \\
&\text{s.t.} \quad \bar{z}^{[i]}_{k+1} = A_{ii}\bar{z}^{[i]}_k + B_i\bar{v}^{[i]}_k + c^{[i]}_k \\
&\qquad \bar{z}^{[i]}_k \in \bar{\mathbb{X}}^{[i]},\ \bar{v}^{[i]}_k \in \bar{\mathbb{U}}^{[i]} \\
&\qquad \bar{z}^{[i]}_k - \tilde{x}^{[i]}_k \in \mathbb{E}^{[i]} \\
&\qquad \bar{z}^{[i]}_t = z^{l,[i]}_t,\ \bar{z}^{[i]}_{t+N} \in \mathcal{CS}^{l-1,[i]}_{t+N} \\
&\qquad k = t, \dots, t+N-1
\end{aligned}
$$

Input for the nominal subsystem: $v^{l,[i]}_t = \bar{v}^{l,[i],\star}_t$

Input for the actual subsystem: $u^{l,[i]}_t = \bar{v}^{l,[i],\star}_t + K_i\left(x^{l,[i]}_t - z^{l,[i]}_t\right)$

Distributed Learning MPC (DLMPC)

Each subsystem $\sum_i$ solves:

$$
\begin{aligned}
J^{l,[i],\star}_{t\to t+N}(z^{l,[i]}_t) = \\
\min_{\bar{v}^{l,[i]}_{[t:t+N-1]}} \quad & \sum_{k=t}^{t+N-1} h_i(\bar{z}^{[i]}_k, \bar{v}^{[i]}_k) + P^{l-1,[i]}_{t+N}(\bar{z}^{[i]}_{t+N}) \\
\text{s.t.} \quad & \bar{z}^{[i]}_{k+1} = A_{ii}\bar{z}^{[i]}_k + B_i\bar{v}^{[i]}_k + c^{[i]}_k \\
& \bar{z}^{[i]}_k \in \bar{\mathbb{X}}^{[i]},\ \bar{v}^{[i]}_k \in \bar{\mathbb{U}}^{[i]} \\
& \bar{z}^{[i]}_k - \tilde{x}^{[i]}_k \in \mathbb{E}^{[i]} \\
& \bar{z}^{[i]}_t = z^{l,[i]}_t,\ \bar{z}^{[i]}_{t+N} \in \mathcal{CS}^{l-1,[i]}_{t+N} \\
& k = t, \dots, t+N-1
\end{aligned}
$$

Input for the nominal subsystem: $v^{l,[i]}_t = \bar{v}^{l,[i],\star}_t$

Input for the actual subsystem: $u^{l,[i]}_t = \bar{v}^{l,[i],\star}_t + K_i\left(x^{l,[i]}_t - z^{l,[i]}_t\right)$

- Low-dimensional QP that can be solved in real-time
- No communication among agents required during online control

DLMPC: Theoretical Results

For the nominal subsystem $z_{t+1}^{[i]} = A_{ii} z_t^{[i]} + B_i v_t^{[i]} + c_t^{[i]}$

- Recursive feasibility and asymptotic stability
- Non-increasing cost over iterations: $J_{0\to\infty}^{l,[i]}(z_0^{l,[i]}) \leq J_{0\to\infty}^{l-1,[i]}(z_0^{l-1,[i]})$
- Convergence to the globally optimal solution of the infinite horizon problem

DLMPC: Theoretical Results

For the nominal subsystem $z_{t+1}^{[i]} = A_{ii} z_t^{[i]} + B_i v_t^{[i]} + c_t^{[i]}$

- Recursive feasibility and asymptotic stability
- Non-increasing cost over iterations: $J_{0\to\infty}^{l,[i]}(z_0^{l,[i]}) \leq J_{0\to\infty}^{l-1,[i]}(z_0^{l-1,[i]})$
- Convergence to the globally optimal solution of the infinite horizon problem

For the actual subsystem $x_{t+1}^{[i]} = A_{ii} x_t^{[i]} + B_i u_i^{[i]} + c_t^{[i]} + w_t^{[i]}$

- Robust constraint satisfaction: $x_t^{[i]} \in \mathbb{X}^{[i]}, u_t^{[i]} \in \mathbb{U}^{[i]}, \quad \forall w_t^{[i]} \in \mathbb{W}^{[i]}$
- Asymptotic stability (**to the origin**)

DLMPC: Theoretical Results

For the nominal subsystem $z_{t+1}^{[i]} = A_{ii} z_t^{[i]} + B_i v_t^{[i]} + c_t^{[i]}$

- Recursive feasibility and asymptotic stability
- Non-increasing cost over iterations: $J_{0\to\infty}^{l,[i]}(z_0^{l,[i]}) \leq J_{0\to\infty}^{l-1,[i]}(z_0^{l-1,[i]})$
- Convergence to the globally optimal solution of the infinite horizon problem

For the actual subsystem $x_{t+1}^{[i]} = A_{ii} x_t^{[i]} + B_i u_i^{[i]} + c_t^{[i]} + w_t^{[i]}$

- Robust constraint satisfaction: $x_t^{[i]} \in \mathbb{X}^{[i]}, u_t^{[i]} \in \mathbb{U}^{[i]}, \quad \forall w_t^{[i]} \in \mathbb{W}^{[i]}$
- Asymptotic stability (**to the origin**)
- ➢ **How about control performance?**

Evaluation of Performance

Metric 1: Individual subsystem $\sum_i$

- Empirical cost-to-go of $\left(x^{l,[i]}_{[0:\infty]}, u^{l,[i]}_{[0:\infty]}\right)$: $\bar{J}^{l,[i]}_{0\to\infty}\left(x^{[i]}_0\right) = \sum_{k=0}^{\infty} h\left(x^{l,[i]}_k, u^{l,[i]}_k\right)$
- Select DLMPC output according to: $l^\star = \text{argmin}_{l=0,\ldots,\texttt{MaxIter}}\, \bar{J}^{l,[i]}_{0\to\infty}\left(x^{[i]}_0\right)$

Metric 2: Plant-wide performance

- Empirical cost-to-go of $\left(\mathbf{x}^{l}_{[0:\infty]}, \mathbf{u}^{l}_{[0:\infty]}\right)$: $\bar{J}^{l,\text{sum}}_{0\to\infty}\left(\mathbf{x}_0\right) = \sum_{i\in\mathcal{I}_{1:M}} \bar{J}^{l,[i]}_{0\to\infty}\left(x^{[i]}_0\right)$
- Select DLMPC output according to: $l^\star = \text{argmin}_{l=0,\ldots,\texttt{MaxIter}}\, \bar{J}^{l,\text{sum}}_{0\to\infty}\left(\mathbf{x}_0\right)$

Output of DLMPC: $\left(\mathbf{x}^{l^*}_{[0:\infty]}, \mathbf{u}^{l^*}_{[0:\infty]}\right)$

Numerical Example: Power Network System[3],[4]

Scenario 1 in [4]:

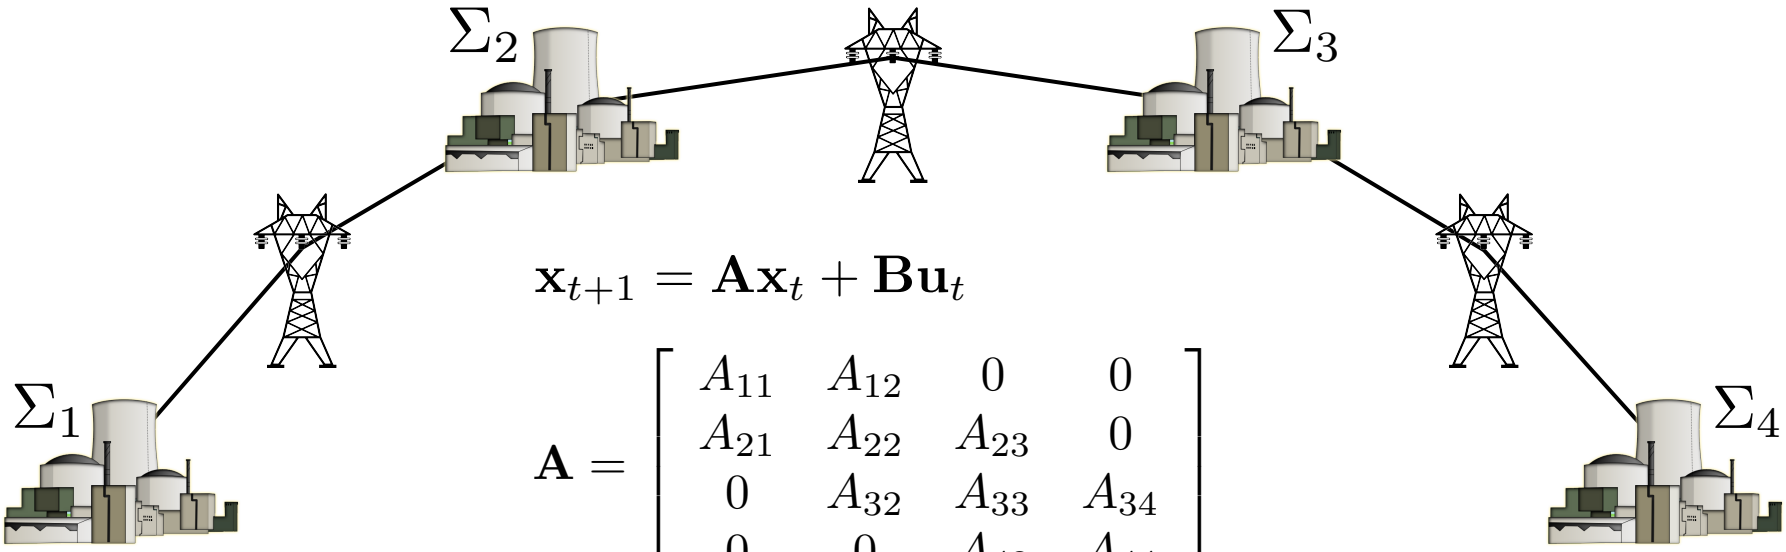

$$\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t$$

$$\mathbf{A} = \begin{bmatrix} A_{11} & A_{12} & 0 & 0 \\ A_{21} & A_{22} & A_{23} & 0 \\ 0 & A_{32} & A_{33} & A_{34} \\ 0 & 0 & A_{43} & A_{44} \end{bmatrix}$$

$$A_{ij} \in \mathbb{R}^{4\times 4}$$

$$\mathbf{B} = \begin{bmatrix} B_1 & 0 & 0 & 0 \\ 0 & B_2 & 0 & 0 \\ 0 & 0 & B_3 & 0 \\ 0 & 0 & 0 & B_4 \end{bmatrix}$$

[3] H. Saadat et al., "Power system analysis."

[4] S. Riverso and G. Ferrari-Trecate, "Hycon2 benchmark: Power network system."

DLMPC Setup

- Initial state: $x_0^{[i]} = \left[0, 0, -\Delta P_L^{[i]}, -\Delta P_L^{[i]}\right]^\top$

  $\Delta P_L^{[i]}$ represents a step change in power load

- Target state: $x_F^{[i]} = [0, 0, 0, 0]^\top$
- Horizon: $N = 5$

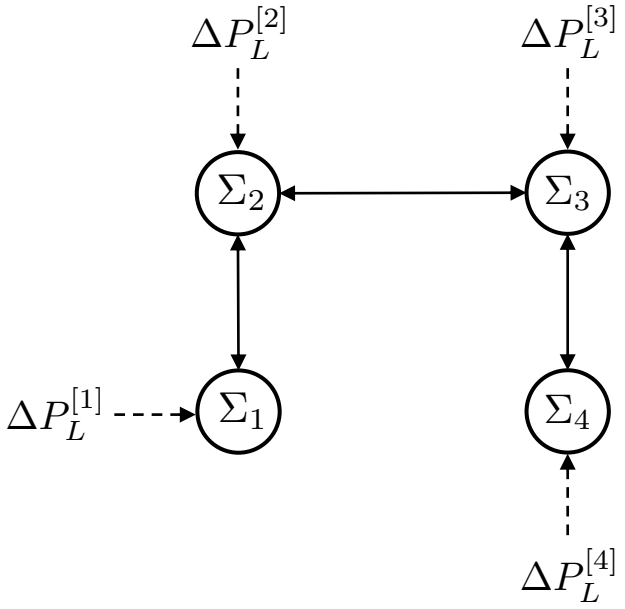

$$
\begin{aligned}
\min_{\bar{v}^{l,[i]}_{[t:t+N-1]}} \quad & \sum_{k=t}^{t+N-1} h_i(\bar{z}_k^{[i]}, \bar{v}_k^{[i]}) + P_{t+N}^{l-1,[i]}(\bar{z}_{t+N}^{[i]}) \\
\text{s.t.} \quad & \bar{z}_{k+1}^{[i]} = A_{ii}\bar{z}_k^{[i]} + B_i\bar{v}_k^{[i]} + c_k^{[i]} \\
& \bar{z}_k^{[i]} \in \bar{\mathbb{X}}^{[i]},\ \bar{v}_k^{[i]} \in \bar{\mathbb{U}}^{[i]} \\
& \bar{z}_k^{[i]} - \tilde{x}_k^{[i]} \in \mathbb{E}^{[i]} \\
& \bar{z}_t^{[i]} = z_t^{l,[i]},\ \bar{z}_{t+N}^{[i]} \in \mathcal{CS}_{t+N}^{l-1,[i]} \\
& k = t, \ldots, t+N-1
\end{aligned}
$$

Deviation constraints (on the first state):

$$
\begin{aligned}
\mathbb{E}^{[1]} &= [-0.0451, 0.0451] \\
\mathbb{E}^{[2]} &= [-0.0241, 0.0241] \\
\mathbb{E}^{[3]} &= [-0.0380, 0.0380] \\
\mathbb{E}^{[4]} &= [-0.060, 0.060]
\end{aligned}
$$

Learning Nominal Trajectories

Iteration	Iteration cost $J^{l,[i]}_{0\to\infty}(z_0^{[i]})$			
	Σ_1	Σ_2	Σ_3	Σ_4
$l=0$	0.48156	0.32609	0.18844	1.52845
$l=1$	0.40011	0.29701	0.17509	1.36560
$l=2$	0.35381	0.28352	0.17219	1.26562
$l=3$	0.32999	0.27668	0.17080	1.21154
$\cdots$			$\cdots$	
$l=25$	0.30121	0.26827	0.16923	1.15885
$l=26$	**0.30120**	**0.26826**	0.16923	1.15884
$l=27$	0.30120	0.26826	**0.16922**	1.15884
$l=28$	0.30120	0.26826	0.16922	1.15884
$l=29$	0.30120	0.26826	0.16922	**1.15883**
$J^{\star,[i]}_{0\to\infty}(z_0^{[i]})$	0.30119	0.26824	0.16922	1.15881

- Recursive feasibility and asymptotic stability ✅
- Non-increasing cost over iterations ✅
- Convergence to the globally optimal solution of the infinite horizon problem ✅

Control Performance of the Actual Systems

Iter. l	Iteration cost $\bar{J}^{l,[i]}_{0\to\infty}(x_0^{[i]})$ Σ_1	Σ_2	Σ_3	Σ_4	$\bar{J}^{l,\text{sum}}_{0\to\infty}(\cdot)$	
0	0.47871	0.31987	0.18944	1.53063	2.51864	← Baseline
1	0.39814	0.29101	0.17355	1.36619	2.22888	
2	0.35275	0.27864	0.17005	1.26567	2.06712	
3	0.32970	0.27332	**0.16923**	1.21172	1.98397	
4	0.31728	0.27118	0.16953	1.18564	1.94364	
5	0.31056	0.27051	0.17006	1.17283	1.92396	
6	0.30701	**0.27050**	0.17055	1.16675	1.91480	
7	0.30512	0.27073	0.17097	1.16381	1.91063	
8	0.30414	0.27105	0.17130	1.16235	1.90885	
9	0.30364	0.27138	0.17156	1.16163	1.90820	
10	0.30337	0.27168	0.17176	1.16126	**1.90806**	← Metric 2: $\bar{J}^{10,\text{sum}}_{0\to\infty}(\mathbf{x}_0)$
11	**0.30324**	0.27193	0.17191	**1.16108**	1.90815	
CMPC	0.30261	0.26915	0.17231	1.15890	1.90297	

- Robust constraint satisfaction and asymptotic stability ✅
- **24.24%** lower cost than state-of-the-art non-cooperative DMPC ☺
- **0.268%** higher cost than centralized MPC ☺

Learned Trajectories

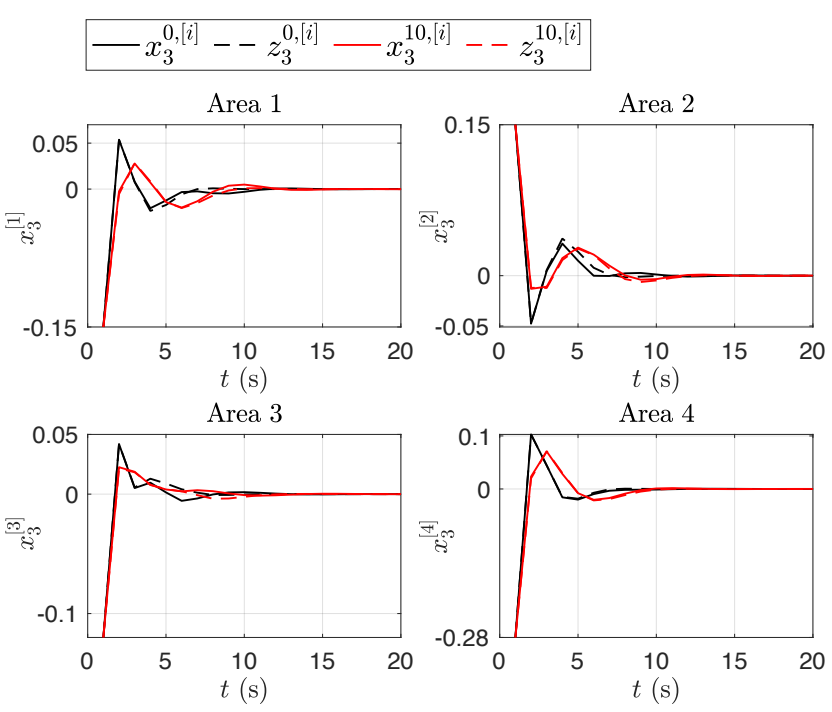

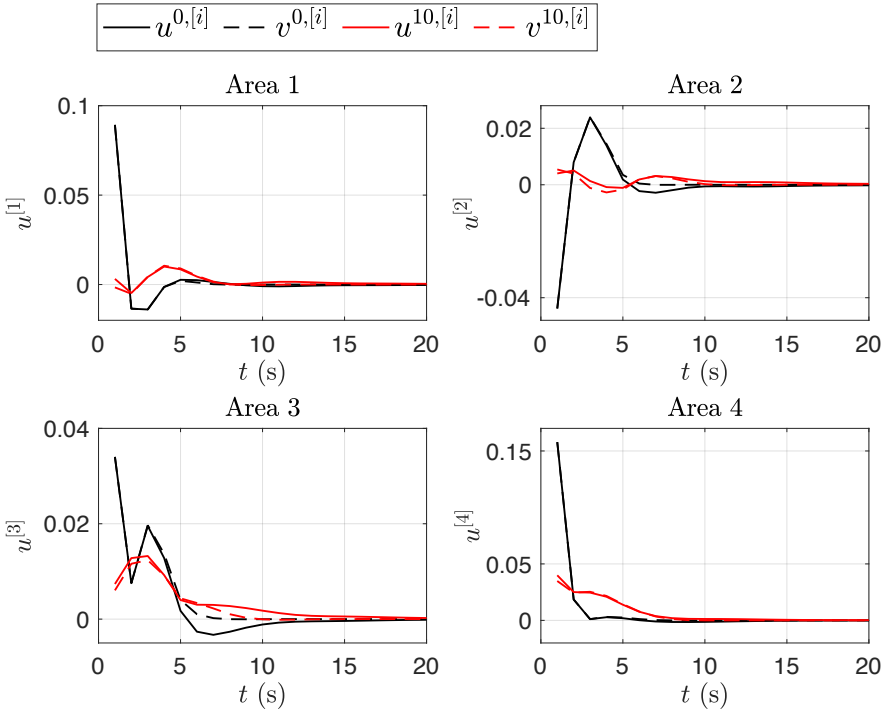

- Less oscillations in state trajectories
- Much less control efforts needed for convergence

DLMPC Summary

Conclusions:

- Proposed an iterative learning control framework for distributed systems
- Strong safety guarantees; empirically providing good performance improvements
- Scalable and easy to deploy:
  - terminal costs and sets constructed from local trajectories
  - no communication required during online control

DLMPC Summary

Conclusions:

- Proposed an iterative learning control framework for distributed systems
- Strong safety guarantees; empirically providing good performance improvements
- Scalable and easy to deploy:
  - terminal costs and sets constructed from local trajectories
  - no communication required during online control

Future work:

- Co-learning the reference trajectory
- Iterative learning for cooperative DMPC
- Connection with multi-agent R.L.

DLMPC Summary

Conclusions:

- Proposed an iterative learning control framework for distributed systems
- Strong safety guarantees; empirically providing good performance improvements
- Scalable and easy to deploy:
  - terminal costs and sets constructed from local trajectories
  - no communication required during online control

Future work:

- Co-learning the reference trajectory
- Iterative learning for cooperative DMPC
- Connection with multi-agent R.L.

Contact: haiminh@princeton.edu

Slides available on my website: https://haiminhu.org/research

Thank you! Questions?