Team optimal control of coupled major-minor subsystems with mean-field sharing

Jalal Arabneydi and Aditya Mahajan

McGill University

Indian Control Conference

6 Jan, 2015

Motivation

Optimal multi-agent control:

- Multiple controllers with a common optimization objective
- Key feature: information decentralization

Motivation

Smart Grids

Optimal multi-agent control:

- Multiple controllers with a common optimization objective
- Key feature: information decentralization

Motivation

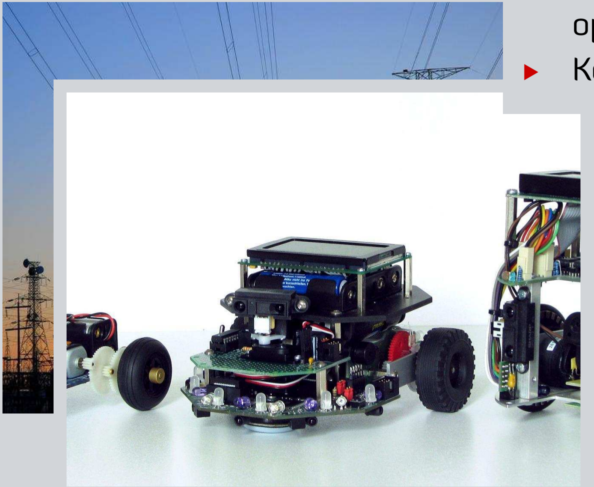

Robotics

Optimal multi-agent control:

- Multiple controllers with a common optimization objective
- Key feature: **information decentralization**

Motivation

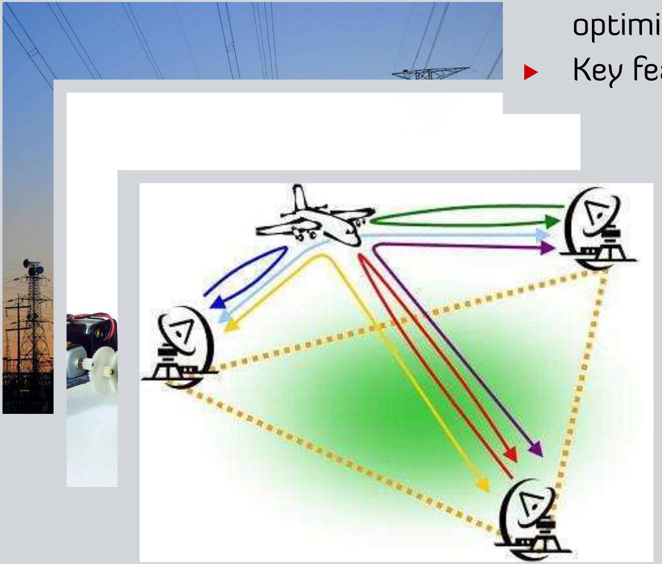

Communication Networks

Optimal multi-agent control:

- Multiple controllers with a common optimization objective
- Key feature: **information decentralization**

Motivation

Optimal multi-agent control:

- Multiple controllers with a common optimization objective
- Key feature: **information decentralization**

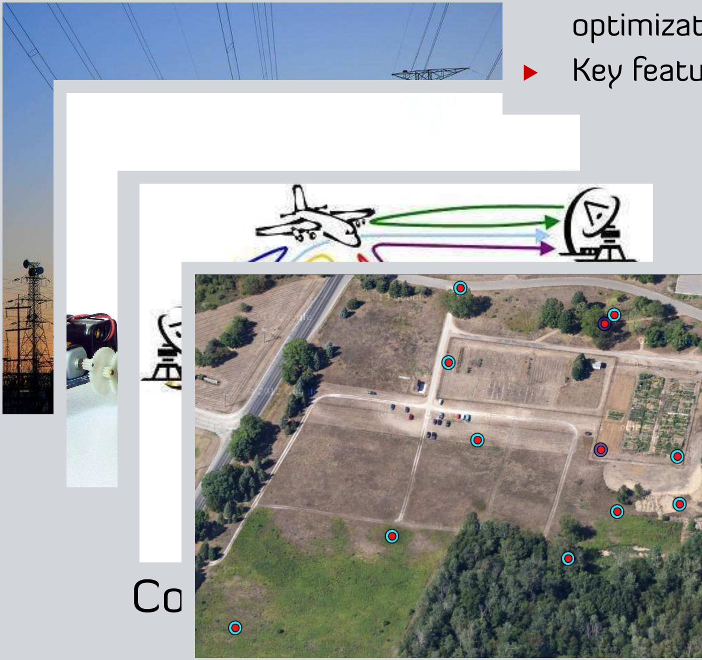

Sensor Networks

Motivation

Optimal multi-agent control:

- Multiple controllers with a common optimization objective
- Key feature: information decentralization

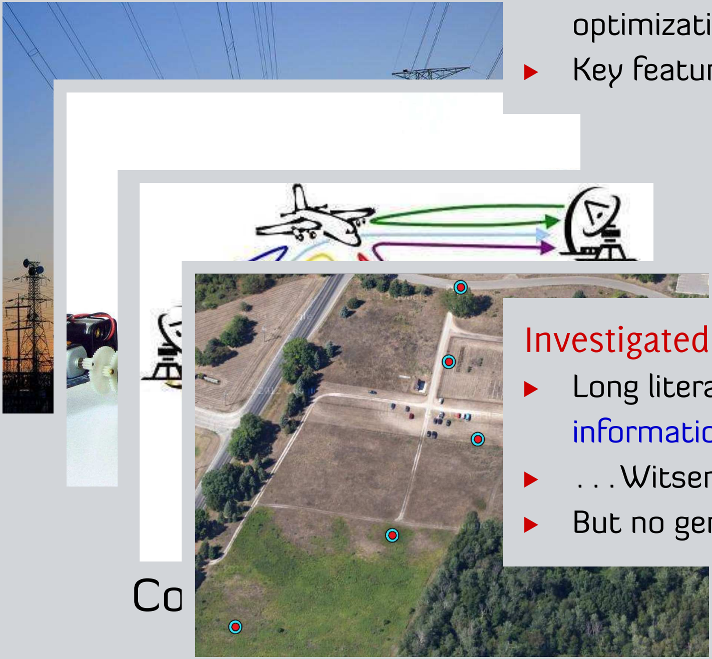

Sensor Networks

Investigated using team theory

- Long literature on solution for specific information structures
- . . . Witsenhausen, Ho, Varaiya, and others.
- But no generic solution approach

Motivation

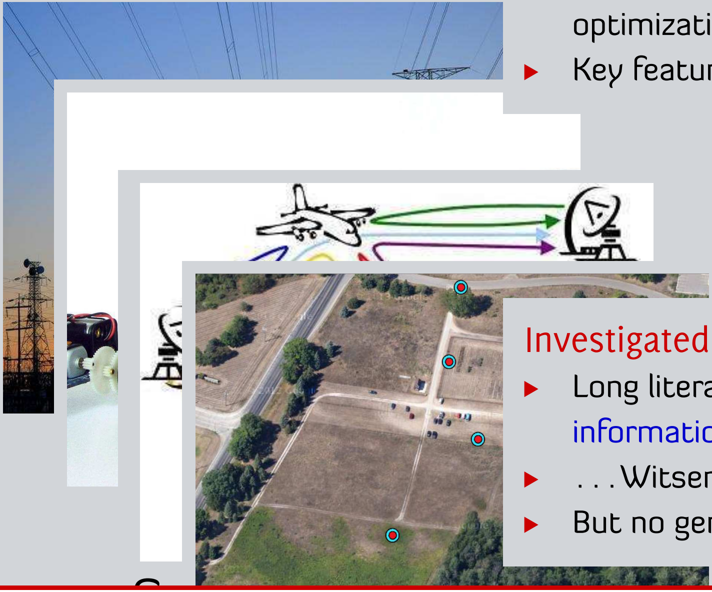

Optimal multi-agent control:

- Multiple controllers with a common optimization objective
- Key feature: information decentralization

Investigated using team theory

- Long literature on solution for specific information structures
- . . . Witsenhausen, Ho, Varaiya, and others.
- But no generic solution approach

Analyze and solve a stylized model for large-scale systems

Motivating setup: System with major and minor subsystems

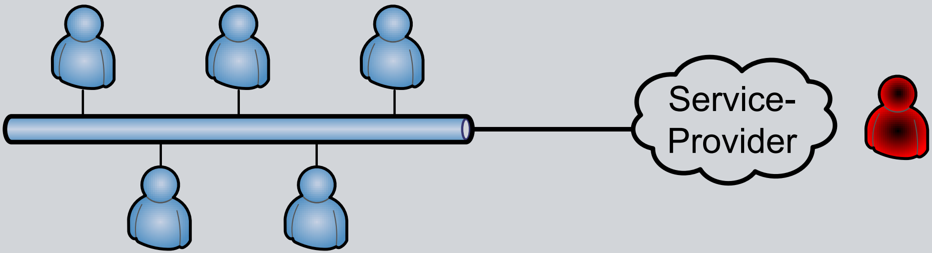

Major-subsystem (e.g., a service provider)

- Controls operating conditions of the system e.g., price, capacity, etc.
- The dynamics of the major-subsystem's state depend on the minor-subsystem's state through their mean-field (or empirical distribution).

Minor homogeneous subsystems

- Dynamics are affected by the state of the major-subsystem.
- Influence each other only though their mean-field (equivalent to a interacting particle model).

(MF-MM) Model and Problem Formulation

Major subsystem

- State $X_t^0 \in \mathcal{X}^0$ Indexed by 0.
- Action $U_t^0 \in \mathcal{U}^0$

(MF-MM) Model and Problem Formulation

Major subsystem
- State $X_t^0 \in \mathcal{X}^0$
- Action $U_t^0 \in \mathcal{U}^0$

Indexed by 0.

Minor subsystems
- State $X_t^i \in \mathcal{X}$
- Action $U_t^i \in \mathcal{U}$

Indexed by $i \in \{1, \dots, n\}$

(MF-MM) Model and Problem Formulation

Major subsystem
- State $X_t^0 \in \mathcal{X}^0$ — Indexed by 0.
- Action $U_t^0 \in \mathcal{U}^0$

Minor subsystems
- State $X_t^i \in \mathcal{X}$ — Indexed by $i \in \{1, \dots, n\}$
- Action $U_t^i \in \mathcal{U}$

- Mean-field of minor subsystems

$$Z_t(x) = \frac{1}{n} \sum_{i=1}^{n} \mathbb{1}\{X_t^i = x\} \quad \text{or} \quad Z_t = \frac{1}{n} \sum_{i=1}^{n} \delta_{X_t^i}$$

(MF-MM) Model and Problem Formulation

Major subsystem
- State $X_t^0 \in \mathcal{X}^0$ — Indexed by 0.
- Action $U_t^0 \in \mathcal{U}^0$

Minor subsystems
- State $X_t^i \in \mathcal{X}$ — Indexed by $i \in \{1, \dots, n\}$
- Action $U_t^i \in \mathcal{U}$

- Mean-field of minor subsystems

$$Z_t(x) = \frac{1}{n}\sum_{i=1}^{n} \mathbb{1}\{X_t^i = x\} \quad \text{or} \quad Z_t = \frac{1}{n}\sum_{i=1}^{n} \delta_{X_t^i}$$

Major subsystem

Dynamics $\quad X_{t+1}^0 = f_t^0(Z_t, X_t^0, U_t^0, W_t^0)$

(MF-MM) Model and Problem Formulation

Major subsystem
- State $X_t^0 \in \mathcal{X}^0$ Indexed by 0.
- Action $U_t^0 \in \mathcal{U}^0$

Minor subsystems
- State $X_t^i \in \mathcal{X}$ Indexed by $i \in \{1, \dots, n\}$
- Action $U_t^i \in \mathcal{U}$
- Mean-field of minor subsystems

$$Z_t(x) = \frac{1}{n} \sum_{i=1}^{n} \mathbb{1}\{X_t^i = x\} \quad \text{or} \quad Z_t = \frac{1}{n} \sum_{i=1}^{n} \delta_{X_t^i}$$

	Major subsystem	**Minor subsystems**
Dynamics	$X_{t+1}^0 = f_t^0(Z_t, X_t^0, U_t^0, W_t^0)$	$X_{t+1}^i = f_t(Z_t, X_t^0, X_t^i, U_t^i, W_t^i)$

(MF-MM) Model and Problem Formulation

Major subsystem
- State $X_t^0 \in \mathcal{X}^0$ — Indexed by 0.
- Action $U_t^0 \in \mathcal{U}^0$

Minor subsystems
- State $X_t^i \in \mathcal{X}$ — Indexed by $i \in \{1, \dots, n\}$
- Action $U_t^i \in \mathcal{U}$
- Mean-field of minor subsystems

$$Z_t(x) = \frac{1}{n} \sum_{i=1}^{n} \mathbb{1}\{X_t^i = x\} \quad \text{or} \quad Z_t = \frac{1}{n} \sum_{i=1}^{n} \delta_{X_t^i}$$

	Major subsystem	**Minor subsystems**
Dynamics	$X_{t+1}^0 = f_t^0(Z_t, X_t^0, U_t^0, W_t^0)$	$X_{t+1}^i = f_t(Z_t, X_t^0, X_t^i, U_t^i, W_t^i)$
Control	$U_t^0 = g_t^0(Z_{1:t}, X_{1:t}^0)$	

(MF-MM) Model and Problem Formulation

Major subsystem
- State $X_t^0 \in \mathcal{X}^0$ — Indexed by 0.
- Action $U_t^0 \in \mathcal{U}^0$

Minor subsystems
- State $X_t^i \in \mathcal{X}$ — Indexed by $i \in \{1, \dots, n\}$
- Action $U_t^i \in \mathcal{U}$
- Mean-field of minor subsystems

$$Z_t(x) = \frac{1}{n} \sum_{i=1}^{n} \mathbb{1}\{X_t^i = x\} \quad \text{or} \quad Z_t = \frac{1}{n} \sum_{i=1}^{n} \delta_{X_t^i}$$

	Major subsystem	**Minor subsystems**
Dynamics	$X_{t+1}^0 = f_t^0(Z_t, X_t^0, U_t^0, W_t^0)$	$X_{t+1}^i = f_t(Z_t, X_t^0, X_t^i, U_t^i, W_t^i)$
Control	$U_t^0 = g_t^0(Z_{1:t}, X_{1:t}^0)$	$U_t^i = g_t(Z_{1:t}, X_{1:t}^0, X_t^i)$

(MF-MM) Model and Problem Formulation

Major subsystem

- State $X_t^0 \in \mathcal{X}^0$ Indexed by 0.
- Action $U_t^0 \in \mathcal{U}^0$

Minor subsystems

- State $X_t^i \in \mathcal{X}$ Indexed by $i \in \{1, \dots, n\}$
- Action $U_t^i \in \mathcal{U}$
- Mean-field of minor subsystems

$$Z_t(x) = \frac{1}{n} \sum_{i=1}^{n} \mathbb{1}\{X_t^i = x\} \quad \text{or} \quad Z_t = \frac{1}{n} \sum_{i=1}^{n} \delta_{X_t^i}$$

	Major subsystem	**Minor subsystems**
Dynamics	$X_{t+1}^0 = f_t^0(Z_t, X_t^0, U_t^0, W_t^0)$	$X_{t+1}^i = f_t(Z_t, X_t^0, X_t^i, U_t^i, W_t^i)$
Control	$U_t^0 = g_t^0(Z_{1:t}, X_{1:t}^0)$	$U_t^i = g_t(Z_{1:t}, X_{1:t}^0, X_t^i)$
Objective	$\min \mathbb{E}\left[\sum_{t=1}^{T} \ell_t(X_t^0, X_t, U_t^0, U_t)\right]$	Arbitrary cost coupling

Assumptions on the model

Assumption (A1) The primitive random variables:

- initial state X_1^0 of the major subsystem
- initial states $(X_1^1, \ldots, X_1^n)$ of the minor subsystems
- process noises $\{(W_t^0, \ldots, W_t^n)\}_{t=1}^T$

are indepedent

Furthermore the initial states $(X_1^1, \ldots, X_1^n)$ and the process noise $(W_t^1, \ldots, W_t^n)$ of the minor subsystem are identically distributed

Assumptions on the model

Assumption (A1) The primitive random variables:

- initial state X_1^0 of the major subsystem
- initial states $(X_1^1, \ldots, X_1^n)$ of the minor subsystems
- process noises $\{(W_t^0, \ldots, W_t^n)\}_{t=1}^T$

are indepedent

Furthermore the initial states $(X_1^1, \ldots, X_1^n)$ and the process noise $(W_t^1, \ldots, W_t^n)$ of the minor subsystem are identically distributed

Assumption (A2) All minor subsystems use identical control laws

- Standard assumption to ensure simplicity, fairness, and robustness.
- Leads to loss in performance

Salient features and main results

Features of the model

- Decentralized control system with non-classical information structure
- Mean-field coupled dynamics and arbitrarily coupled cost.
- Seek globally optimal solution for arbitrary # of minor controllers

Salient features and main results

Features of the model

- Decentralized control system with **non-classical information structure**
- Mean-field coupled dynamics and arbitrarily coupled cost.
- Seek globally optimal solution for arbitrary # of minor controllers

Main results

- Indetify the **structure** of optimal control strategies.
- Obtain a dynamic program that determines optimal control strategies at all controllers.

Salient features and main results

Features of the model

- Decentralized control system with non-classical information structure
- Mean-field coupled dynamics and arbitrarily coupled cost.
- Seek globally optimal solution for arbitrary # of minor controllers

Main results

- Indetify the structure of optimal control strategies.
- Obtain a dynamic program that determines optimal control strategies at all controllers.

Features of the solution

- State space of the DP increases polynomially (rather than exponentially) with the number of minor subsystems.
- Action space of DP does not depend on the # of minor subsystems.
- State and action spaces do not depend on time; hence, the results extend naturally to infinite horizon

Proof outline

First analyze basic MF model [Arabneydi Mahajan, CDC 2014]

Multiple types of minor subsystems but no major subsystem.

Proof outline

First analyze basic MF model [Arabneydi Mahajan, CDC 2014]

Multiple types of minor subsystems but no major subsystem.

Step 1 Follow the common information approach [NMT13] to convert the decentralized control problem into a centralized control problem

- Nayyar, Mahajan, Teneketzis, "Decentralized stochastic control with partial history sharing: A common information approach," IEEE TAC 2013.

Step 2 Exploit symmetry of the system (with respect to the controllers) to show that the mean-field is an information state.

Proof outline

First analyze basic MF model [Arabneydi Mahajan, CDC 2014]

Multiple types of minor subsystems but no major subsystem.

Step 1 Follow the common information approach [NMT13] to convert the decentralized control problem into a centralized control problem

- Nayyar, Mahajan, Teneketzis, "Decentralized stochastic control with partial history sharing: A common information approach," IEEE TAC 2013.

Step 2 Exploit symmetry of the system (with respect to the controllers) to show that the mean-field is an information state.

Note that MF-MM model is a special case of basic MF model

The major-minor model corresponds to a basic MF system with 1-subsystem of type 0 and n-subsystems of type 1.

Proof outline

First analyze basic MF model [Arabneydi Mahajan, CDC 2014]

Multiple types of minor subsystems but no major subsystem.

Step 1 Follow the common information approach [NMT13] to convert the decentralized control problem into a centralized control problem

- Nayyar, Mahajan, Teneketzis, "Decentralized stochastic control with partial history sharing: A common information approach," IEEE TAC 2013.

Step 2 Exploit symmetry of the system (with respect to the controllers) to show that the mean-field is an information state.

Note that MF-MM model is a special case of basic MF model

The major-minor model corresponds to a basic MF system with 1-subsystem of type 0 and n-subsystems of type 1.

Translate the results of basic MF model to the MF-MM model

Basic MF-model [Arabneydi Mahajan, CDC 2014]

Minor subsystems ▸ Type $k \in \{1, \ldots, m\}$. $\mathcal{N}^k = \{$ subsystems of type-$k\}$. $|\mathcal{N}^k| = n^k$.

Basic MF-model [Arabneydi Mahajan, CDC 2014]

Minor subsystems

- Type $k \in \{1, \dots, m\}$. $\mathcal{N}^k = \{$ subsystems of type-$k\}$. $|\mathcal{N}^k| = n^k$.
- State $X_t^i \in \mathcal{X}^k$ Indexed by $i \in \mathcal{N}^k$
- Action $U_t^i \in \mathcal{U}^k$

Basic MF-model [Arabneydi Mahajan, CDC 2014]

Minor subsystems

- Type $k \in \{1, \ldots, m\}$. $\mathcal{N}^k = \{$ subsystems of type-$k\}$. $|\mathcal{N}^k| = n^k$.

- State $X_t^i \in \mathcal{X}^k$ Indexed by $i \in \mathcal{N}^k$
- Action $U_t^i \in \mathcal{U}^k$

- Mean-field of minor subsystems

$$Z_t^k(x) = \frac{1}{n^k} \sum_{i \in \mathcal{N}^k} \mathbb{1}\{X_t^i = x\} \quad \text{and} \quad Z_t = (Z_t^1, \ldots, Z_t^m)$$

Basic MF-model [Arabneydi Mahajan, CDC 2014]

Minor subsystems
- Type $k \in \{1, \dots, m\}$. $\mathcal{N}^k = \{$ subsystems of type-$k\}$. $|\mathcal{N}^k| = n^k$.
- State $X^i_t \in \mathcal{X}^k$ Indexed by $i \in \mathcal{N}^k$
- Action $U^i_t \in \mathcal{U}^k$
- Mean-field of minor subsystems

$$Z^k_t(x) = \frac{1}{n^k} \sum_{i \in \mathcal{N}^k} \mathbb{1}\{X^i_t = x\} \quad \text{and} \quad Z_t = (Z^1_t, \dots, Z^m_t)$$

Dynamics $X^i_{t+1} = f^k_t(Z_t, X^i_t, U^i_t, W^i_t), \quad i \in \mathcal{N}^k.$

Basic MF-model [Arabneydi Mahajan, CDC 2014]

Minor subsystems

- Type $k \in \{1, \ldots, m\}$. $\mathcal{N}^k = \{$ subsystems of type-$k\}$. $|\mathcal{N}^k| = n^k$.
- State $X_t^i \in \mathcal{X}^k$ Indexed by $i \in \mathcal{N}^k$
- Action $U_t^i \in \mathcal{U}^k$
- Mean-field of minor subsystems

$$Z_t^k(x) = \frac{1}{n^k} \sum_{i \in \mathcal{N}^k} \mathbb{1}\{X_t^i = x\} \quad \text{and} \quad Z_t = (Z_t^1, \ldots, Z_t^m)$$

Dynamics $X_{t+1}^i = f_t^k(Z_t, X_t^i, U_t^i, W_t^i), \quad i \in \mathcal{N}^k.$

Controls $U_t^i = g_t^k(Z_{1:t}, X_t^i), \quad i \in \mathcal{N}^k.$

Equiv. to (A2) All subsystems of the same type use identical control laws

Basic MF-model [Arabneydi Mahajan, CDC 2014]

Minor subsystems

- Type $k \in \{1, \dots, m\}$. $\mathcal{N}^k = \{$ subsystems of type-$k\}$. $|\mathcal{N}^k| = n^k$.

- State $X^i_t \in \mathcal{X}^k$ Indexed by $i \in \mathcal{N}^k$
- Action $U^i_t \in \mathcal{U}^k$

- Mean-field of minor subsystems

$$Z^k_t(x) = \frac{1}{n^k} \sum_{i \in \mathcal{N}^k} \mathbb{1}\{X^i_t = x\} \quad \text{and} \quad Z_t = (Z^1_t, \dots, Z^m_t)$$

Dynamics $X^i_{t+1} = f^k_t(Z_t, X^i_t, U^i_t, W^i_t), \quad i \in \mathcal{N}^k.$

Controls $U^i_t = g^k_t(Z_{1:t}, X^i_t), \quad i \in \mathcal{N}^k.$

Equiv. to (A2) All subsystems of the same type use identical control laws

Objective $\min \mathbb{E}\Big[\sum_{t=1}^{T} \ell_t(X_t, U_t)\Big]$, where $X_t = (X^1_t, \dots, X^n_t)$; $U_t = (U^1_t, \dots, U^n_t)$

From decentralized to centralized control: the common information approach

▸ Nayyar, Mahajan, Teneketzis, "Decentralized stochastic control with partial history sharing: A common information approach," IEEE TAC 2013.

From decentralized to centralized control: the common information approach

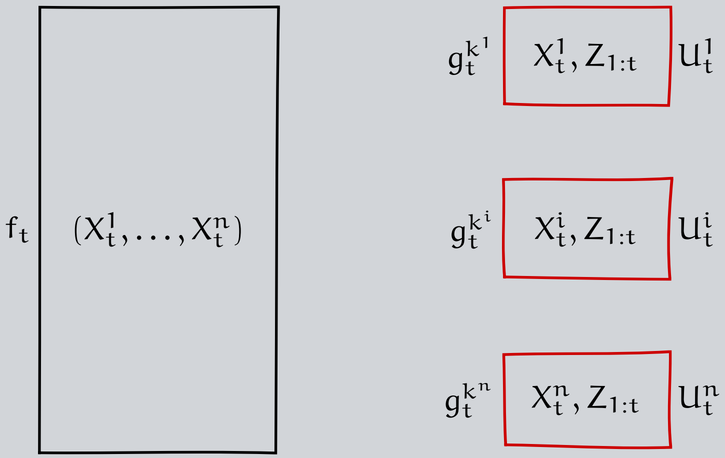

- Nayyar, Mahajan, Teneketzis, "Decentralized stochastic control with partial history sharing: A common information approach," IEEE TAC 2013.

From decentralized to centralized control: the common information approach

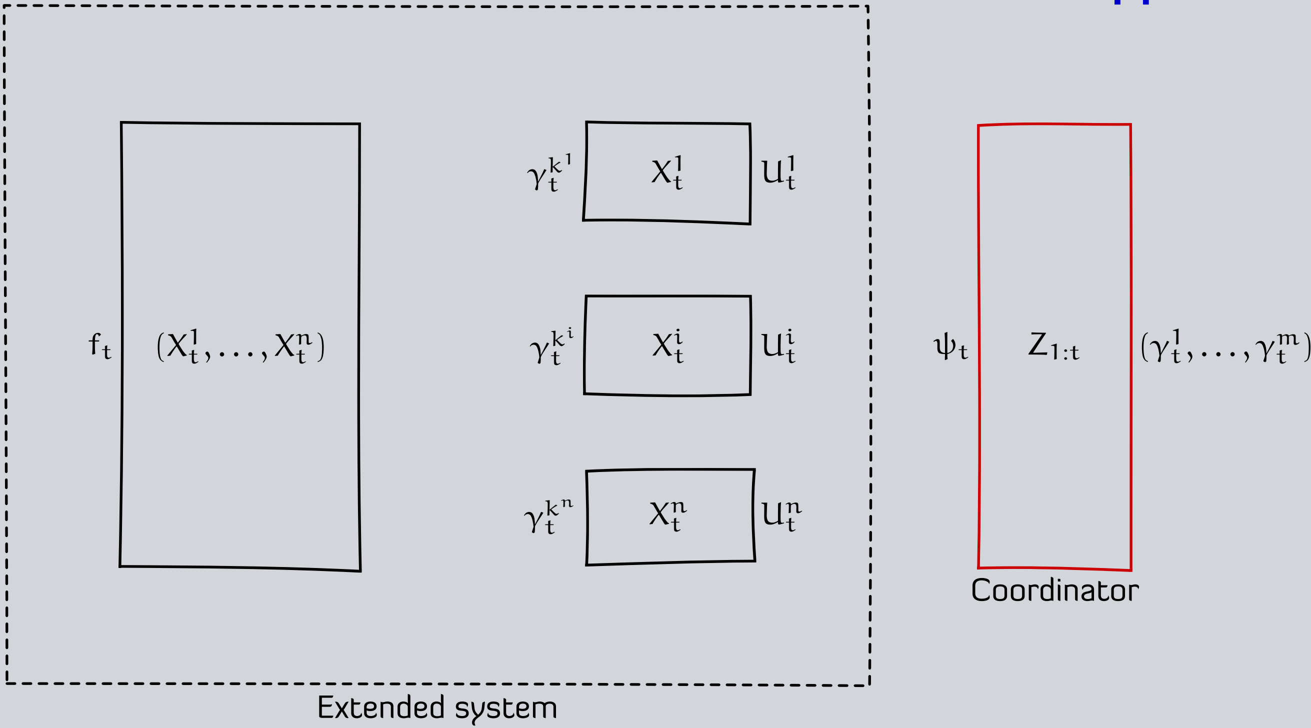

- Nayyar, Mahajan, Teneketzis, "Decentralized stochastic control with partial history sharing: A common information approach," IEEE TAC 2013.

From decentralized to centralized control: the common information approach

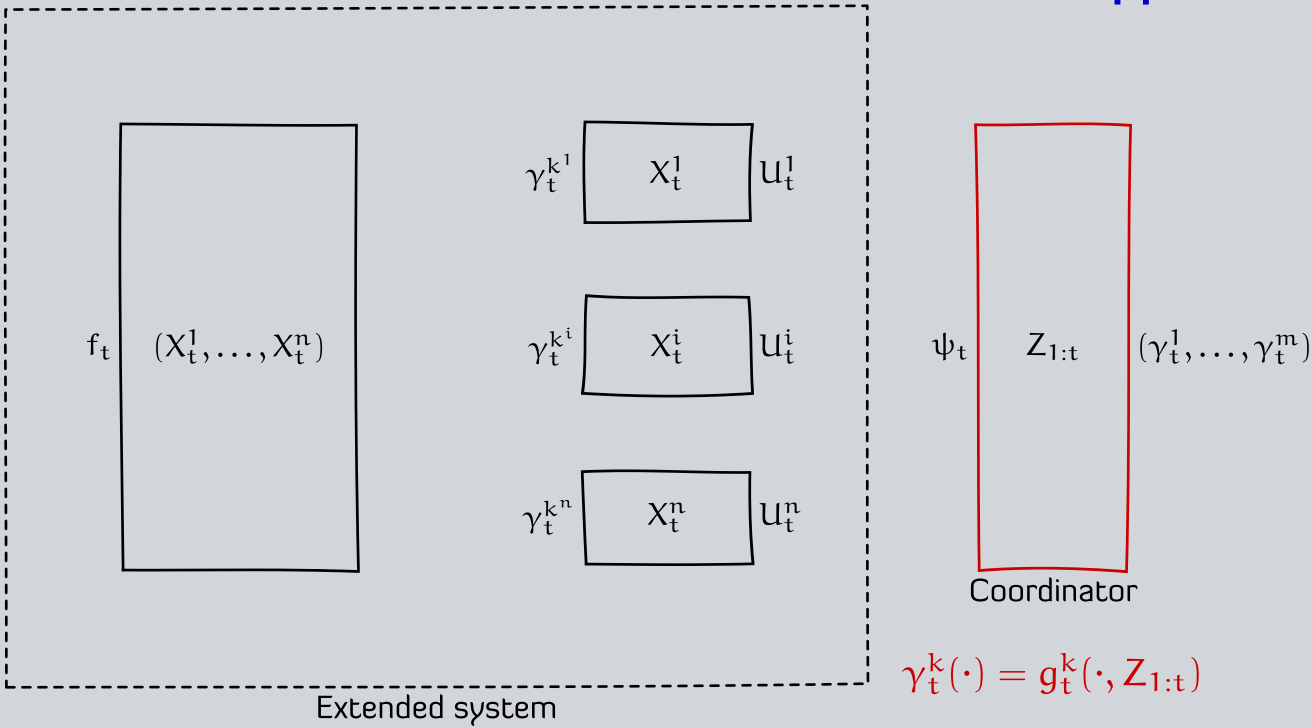

$$\gamma^k_t(\cdot) = g^k_t(\cdot, Z_{1:t})$$

- Nayyar, Mahajan, Teneketzis, "Decentralized stochastic control with partial history sharing: A common information approach," IEEE TAC 2013.

Equivalent centralized problem

Dynamical system	State	: $(X_t^1, \ldots, X_t^n)$
	Observations	: Z_t
	Control actions	: $(\gamma_t^1, \ldots, \gamma_t^m)$, where $\gamma_t^k : \mathcal{X}^k \mapsto \mathcal{U}^k$.
	Control law	:

$$(\gamma_t^1, \ldots, \gamma_t^m) = \psi_t(Z_{1:t})$$

Equivalent centralized problem

Dynamical system

- State: $(X_t^1, \ldots, X_t^n)$
- Observations: Z_t
- Control actions: $(\gamma_t^1, \ldots, \gamma_t^m)$, where $\gamma_t^k : \mathcal{X}^k \mapsto \mathcal{U}^k$.
- Control law:

$$(\gamma_t^1, \ldots, \gamma_t^m) = \psi_t(Z_{1:t})$$

"Standard" centralized POMDP

Equivalent centralized problem

Dynamical system

- State : $(X_t^1, \ldots, X_t^n)$
- Observations : Z_t
- Control actions: $(\gamma_t^1, \ldots, \gamma_t^m)$, where $\gamma_t^k : \mathcal{X}^k \mapsto \mathcal{U}^k$.
- Control law :

$$(\gamma_t^1, \ldots, \gamma_t^m) = \psi_t(Z_{1:t})$$

"Standard" centralized POMDP

Information state

Belief state: $\mathbb{P}(\text{state} \mid \text{observations}) = \mathbb{P}(X_t^1, \ldots, X_t^n \mid Z_{1:t})$

Equivalent centralized problem

Dynamical system	State	: $(X_t^1, \dots, X_t^n)$
	Observations	: Z_t
	Control actions	: $(\gamma_t^1, \dots, \gamma_t^m)$, where $\gamma_t^k : \mathcal{X}^k \mapsto \mathcal{U}^k$.
	Control law	:

$$(\gamma_t^1, \dots, \gamma_t^m) = \psi_t(Z_{1:t})$$

"Standard" centralized POMDP

Information state

Belief state: $\mathbb{P}(\text{state} \mid \text{observations}) = \mathbb{P}(X_t^1, \dots, X_t^n \mid Z_{1:t})$

Because of the symmetry in the problem, Z_t is also an information state.

Identifying the information state

Controlled Markov property

$$\mathbb{P}(Z_{t+1} = z \mid Z_{1:t} = z_{1:t}, \Gamma_{1:t} = \gamma_{1:t}) = \mathbb{P}(Z_{t+1} = z \mid Z_t = z_t)$$

Sufficient for performance evaluation

$$\mathbb{E}[\ell_t(X_t, U_t) \mid Z_{1:t}, \Gamma_{1:t}] = \hat{\ell}_t(Z_t, \Gamma_t).$$

Identifying the information state

Key Lemma

$$\mathbb{P}(X_t = x \mid Z_{1:t} = z_{1:t}, \Gamma_{1:t} = \gamma_{1:t}) = \mathbb{P}(X_t = x \mid Z_t = z_t, \Gamma_t = \gamma_t)$$

$$= \frac{\mathbb{1}\{x \in H(z_t)\}}{|H(z_t)|}$$

where $H(z) = \{(x^1, \dots, x^n) \in \mathcal{X}^n : \text{emperical dist}(x^1, \dots, x^n) = z\}$

Controlled Markov property

$$\mathbb{P}(Z_{t+1} = z \mid Z_{1:t} = z_{1:t}, \Gamma_{1:t} = \gamma_{1:t}) = \mathbb{P}(Z_{t+1} = z \mid Z_t = z_t)$$

Sufficient for performance evaluation

$$\mathbb{E}[\ell_t(X_t, U_t) \mid Z_{1:t}, \Gamma_{1:t}] = \hat{\ell}_t(Z_t, \Gamma_t).$$

Identifying the information state

Key Lemma

$$\mathbb{P}(X_t = x \mid Z_{1:t} = z_{1:t}, \Gamma_{1:t} = \gamma_{1:t}) = \mathbb{P}(X_t = x \mid Z_t = z_t, \Gamma_t = \gamma_t)$$

$$= \frac{\mathbb{1}\{x \in H(z_t)\}}{|H(z_t)|}$$

where $H(z) = \{(x^1, \dots, x^n) \in \mathcal{X}^n : \text{emperical dist}(x^1, \dots, x^n) = z\}$

Controlled Markov property

$$\mathbb{P}(Z_{t+1} = z \mid Z_{1:t} = z_{1:t}, \Gamma_{1:t} = \gamma_{1:t}) = \mathbb{P}(Z_{t+1} = z \mid Z_t = z_t)$$

Follows from Lemma and (A1)

Sufficient for performance evaluation

$$\mathbb{E}[\ell_t(X_t, U_t) \mid Z_{1:t}, \Gamma_{1:t}] = \hat{\ell}_t(Z_t, \Gamma_t).$$

Identifying the information state

Key Lemma

$$\mathbb{P}(X_t = x \mid Z_{1:t} = z_{1:t}, \Gamma_{1:t} = \gamma_{1:t}) = \mathbb{P}(X_t = x \mid Z_t = z_t, \Gamma_t = \gamma_t)$$

$$= \frac{\mathbb{1}\{x \in H(z_t)\}}{|H(z_t)|}$$

where $H(z) = \{(x^1, \dots, x^n) \in \mathcal{X}^n : \text{emperical dist}(x^1, \dots, x^n) = z\}$

Controlled Markov property

$$\mathbb{P}(Z_{t+1} = z \mid Z_{1:t} = z_{1:t}, \Gamma_{1:t} = \gamma_{1:t}) = \mathbb{P}(Z_{t+1} = z \mid Z_t = z_t)$$

Follows from Lemma and (A1)

Sufficient for performance evaluation

$$\mathbb{E}[\ell_t(X_t, U_t) \mid Z_{1:t}, \Gamma_{1:t}] = \hat{\ell}_t(Z_t, \Gamma_t).$$

Follows from Lemma and the coordinated system: $U_t = \gamma_t(X_t)$

Basic MF model: Main results

Theorem 1 In the equivalent centralized system, there is no loss of optimality in restricting attention to coordination strategies of the form

$$(\gamma_t^1, \ldots, \gamma_t^m) = \psi_t(z_t).$$

Equivalently, in the original decentralized system, there is no loss of optimality in restricting attention to control strategies of the form

$$U_t^i = g_t^k(X_t^i, Z_t), \quad i \in \mathcal{N}^k.$$

Basic MF model: Main results

Theorem 1 In the equivalent centralized system, there is no loss of optimality in restricting attention to coordination strategies of the form

$$(\gamma_t^1, \dots, \gamma_t^m) = \psi_t(z_t).$$

Equivalently, in the original decentralized system, there is no loss of optimality in restricting attention to control strategies of the form

$$U_t^i = g_t^k(X_t^i, Z_t), \quad i \in \mathcal{N}^k.$$

Theorem 2 Let (V_t, ψ_t^*), where $V_t : \mathcal{Z} \to \mathbb{R}$, $\psi_t^* : \mathcal{Z} \mapsto (\gamma^1, \dots, \gamma^m)$, and $\gamma^k : \mathcal{X}^k \to \mathcal{U}^k$, be the solution to the following dynamic program:

$$V_t(z) = \min_{(\gamma^1, \dots, \gamma^m)} \mathbb{E}[\ell_t(X_t, U_t) + V_{t+1}(Z_{t+1}) \mid Z_t = z, \Gamma_t = \gamma]$$

Then, $g_t^{*,k}(z, x) = \psi^{*,k}(z)(x)$, is an optimal strategy for controller of type k.

Back to the MF-MM Model: Main results

Note that MF-MM model is a special case of basic MF model

The major-minor model corresponds to a basic MF system with 1-subsystem of type 0 and n-subsystems of type 1.

Back to the MF-MM Model: Main results

Note that MF-MM model is a special case of basic MF model

The major-minor model corresponds to a basic MF system with 1-subsystem of type 0 and n-subsystems of type 1.

Translate the results of basic MF model to the MF-MM model

Back to the MF-MM Model: Main results

Note that MF-MM model is a special case of basic MF model

The major-minor model corresponds to a basic MF system with 1-subsystem of type 0 and n-subsystems of type 1.

Translate the results of basic MF model to the MF-MM model

Theorem 1a There is no loss of optimality in restricting attention to control laws of the form

$$U_t^0 = g_t^0(X_t^0, Z_t) \quad \text{and} \quad U_t^i = g_t^k(X_t^i, X_t^0, Z_t), \ \forall i \in \mathcal{N}^k.$$

Back to the MF-MM Model: Main results

Note that MF-MM model is a special case of basic MF model

The major-minor model corresponds to a basic MF system with 1-subsystem of type 0 and n-subsystems of type 1.

Translate the results of basic MF model to the MF-MM model

Theorem 1a There is no loss of optimality in restricting attention to control laws of the form

$$U_t^0 = g_t^0(X_t^0, Z_t) \quad \text{and} \quad U_t^i = g_t^k(X_t^i, X_t^0, Z_t), \ \forall i \in \mathcal{N}^k.$$

Theorem 2a Let (V_t, ψ_t^*) be the solution to the following dynamic program:

$$V_t(z, x^0) = \min_{u^0, \gamma} \mathbb{E}[\ell_t(X_t^0, X_t, U_t^0, U_t) + V_{t+1}(Z_{t+1}, X_{t+1}^0) \mid Z_t = z, X_t^0 = x_t^0,$$
$$\Gamma_t = \gamma, U_t^0 = u^0]$$

Then, $g_t^{*,0}(z, x^0) = \psi^{*,1}(z, x^0)$, and $g_t^*(z, x^0, x) = \psi^{*,2}(z, x^0)(x)$ is an optimal strategy.

Summary

(MF-MM) Model and Problem Formulation

Major subsystem
- State $X^0_t \in \mathcal{X}^0$ — Indexed by 0.
- Action $U^0_t \in \mathcal{U}^0$

Minor subsystems
- State $X^i_t \in \mathcal{X}$ — Indexed by $i \in \{1, \dots, n\}$
- Action $U^i_t \in \mathcal{U}$

- Mean-field of minor subsystems

$$Z_t(x) = \frac{1}{n}\sum_{i=1}^{n} \mathbb{1}\{X^i_t = x\} \quad \text{or} \quad Z_t = \frac{1}{n}\sum_{i=1}^{n} \delta_{X^i_t}$$

	Major subsystem	**Minor subsystems**
Dynamics	$X^0_{t+1} = f^0_t(Z_t, X^0_t, U^0_t, W^0_t)$	$X^i_{t+1} = f_t(Z_t, X^0_t, X^i_t, U^i_t, W^i_t)$
Control	$U^0_t = g^0_t(Z_{1:t}, X^0_{1:t})$	$U^i_t = g_t(Z_{1:t}, X^0_{1:t}, X^i_t)$
Objective	$\min \mathbb{E}\left[\sum_{t=1}^{T} \ell_t(X^0_t, X_t, U^0_t, U_t)\right]$	Arbitrary cost coupling

Summary

Proof outline

First analyze basic MF model [Arabneydi Mahajan, CDC 2014]

Multiple types of minor subsystems but no major subsystem.

Step 1 Follow the common information approach [NMT13] to convert the decentralized control problem into a centralized control problem

- Nayyar, Mahajan, Teneketzis, "Decentralized stochastic control with partial history sharing: A common information approach," IEEE TAC 2013.

Step 2 Exploit symmetry of the system (with respect to the controllers) to show that the mean-field is an information state.

Note that MF-MM model is a special case of basic MF model

The major-minor model corresponds to a basic MF system with 1-subsystem of type 0 and n-subsystems of type 1.

Translate the results of basic MF model to the MF-MM model

Summary

From decentralized to centralized control: the common information approach

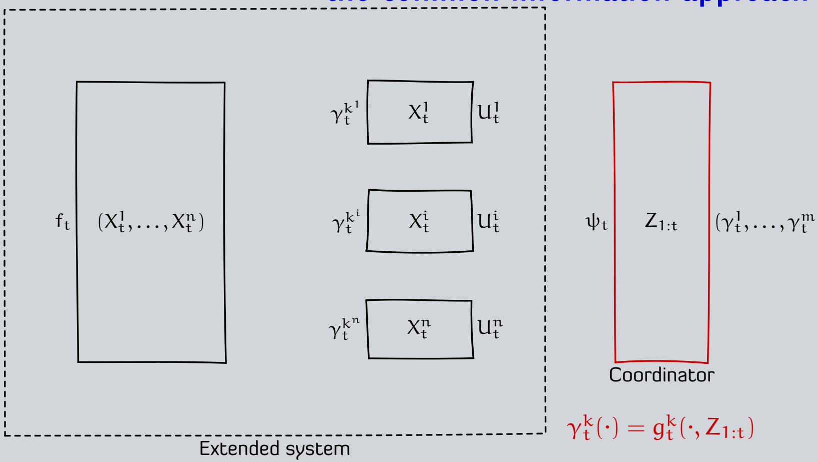

▸ Nayyar, Mahajan, Teneketzis, "Decentralized stochastic control with partial history sharing: A common information approach," IEEE TAC 2013.

Summary

Identifying the information state

Key Lemma

$$\mathbb{P}(X_t = x \mid Z_{1:t} = z_{1:t}, \Gamma_{1:t} = \gamma_{1:t}) = \mathbb{P}(X_t = x \mid Z_t = z_t, \Gamma_t = \gamma_t) = \frac{\mathbb{1}\{x \in H(z_t)\}}{|H(z_t)|}$$

where $H(z) = \{(x^1, \dots, x^n) \in \mathcal{X}^n : \text{emperical dist}(x^1, \dots, x^n) = z\}$

Controlled Markov property

$$\mathbb{P}(Z_{t+1} = z \mid Z_{1:t} = z_{1:t}, \Gamma_{1:t} = \gamma_{1:t}) = \mathbb{P}(Z_{t+1} = z \mid Z_t = z_t)$$

Follows from Lemma and (A1)

Sufficient for performance evaluation

$$\mathbb{E}[\ell_t(X_t, U_t) \mid Z_{1:t}, \Gamma_{1:t}] = \hat{\ell}_t(Z_t, \Gamma_t).$$

Follows from Lemma and the coordinated system: $U_t = \gamma_t(X_t)$

Summary

Basic MF model: Main results

Theorem 1 In the equivalent centralized system, there is no loss of optimality in restricting attention to coordination strategies of the form

$$(\gamma_t^1, \dots, \gamma_t^m) = \psi_t(z_t).$$

Equivalently, in the original decentralized system, there is no loss of optimality in restricting attention to control strategies of the form

$$U_t^i = g_t^k(X_t^i, Z_t), \quad i \in \mathcal{N}^k.$$

Theorem 2 Let (V_t, ψ_t^*), where $V_t : \mathcal{Z} \to \mathbb{R}$, $\psi_t^* : \mathcal{Z} \mapsto (\gamma^1, \dots, \gamma^m)$, and $\gamma^k : \mathcal{X}^k \to \mathcal{U}^k$, be the solution to the following dynamic program:

$$V_t(z) = \min_{(\gamma^1, \dots, \gamma^m)} \mathbb{E}[\ell_t(X_t, U_t) + V_{t+1}(Z_{t+1}) \mid Z_t = z, \Gamma_t = \gamma]$$

Then, $g_t^{*,k}(z, x) = \psi^{*,k}(z)(x)$, is an optimal strategy for controller of type k.

Team optimal control of major-minor subsystems— (Arabneydi and Mahajan)

Summary

Back to the MF-MM Model: Main results

Note that MF-MM model is a special case of basic MF model

The major-minor model corresponds to a basic MF system with 1-subsystem of type 0 and n-subsystems of type 1.

Translate the results of basic MF model to the MF-MM model

Theorem 1a There is no loss of optimality in restricting attention to control laws of the form

$$U^0_t = g^0_t(X^0_t, Z_t) \quad \text{and} \quad U^i_t = g^k_t(X^i_t, X^0_t, Z_t), \ \forall i \in \mathcal{N}^k.$$

Theorem 2a Let (V_t, ψ^*_t) be the solution to the following dynamic program:

$$V_t(z, x^0) = \min_{u^0, \gamma} \mathbb{E}[\ell_t(X^0_t, X_t, U^0_t, U_t) + V_{t+1}(Z_{t+1}, X^0_{t+1}) \mid Z_t = z, X^0_t = x^0_t, \Gamma_t = \gamma, U^0_t = u^0]$$

Then, $g^{*,0}_t(z, x^0) = \psi^{*,1}(z, x^0)$, and $g^*_t(z, x^0, x) = \psi^{*,2}(z, x^0)(x)$ is an optimal strategy.

Conclusion

Features of the solution

- State space of the DP increases polynomially (rather than exponentially) with the number of minor subsystems.
- Action space of DP does not depend on the # of minor subsystems.
- State and action spaces do not depend on time; hence, the results extend naturally to infinite horizon

$$V_t(z, x^0) = \min_{u^0, \gamma} \mathbb{E}[\ell_t(X_t^0, X_t, U_t^0, U_t) + V_{t+1}(Z_{t+1}, X_{t+1}^0) \mid Z_t = z, X_t^0 = x_t^0, \Gamma_t = \gamma, U_t^0 = u^0]$$

Conclusion

Features of the solution

- State space of the DP increases polynomially (rather than exponentially) with the number of minor subsystems.
- Action space of DP does not depend on the # of minor subsystems.
- State and action spaces do not depend on time; hence, the results extend naturally to infinite horizon

Appropriateness of the model

- Assume that the mean-field is observed by all users.
- Happens naturally in some applications (e.g., EV charging, comm. nets)
- Can be computed in a distributed manner using consensus protocols.

Conclusion

Features of the solution

- State space of the DP increases polynomially (rather than exponentially) with the number of minor subsystems.
- Action space of DP does not depend on the # of minor subsystems.
- State and action spaces do not depend on time; hence, the results extend naturally to infinite horizon

Appropriateness of the model

- Assume that the mean-field is observed by all users.
- Happens naturally in some applications (e.g., EV charging, comm. nets)
- Can be computed in a distributed manner using consensus protocols.

Future directions

- Simplification for LQG setups.
- Comparsion with results in mean-field games.
- Asymptotic properties as $n \to \infty$.