Optimal Decentralized Control of System with Partially Exchangeable Agents

Aditya Mahajan
McGill University

Joint work with Jalal Arabneydi

Allerton Conference on Communication, Control, and Computing
28 Sep, 2016

Optimal decentralized control: Applications

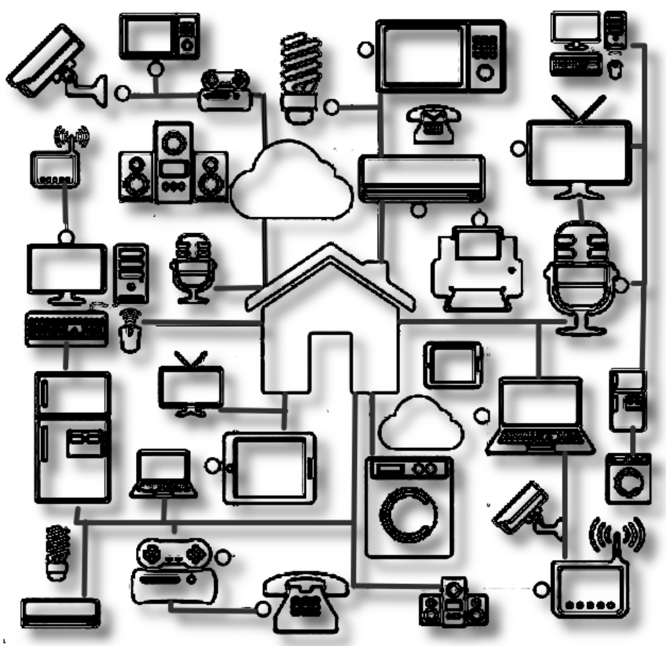

Internet of Things

Optimal decentralized control: Applications

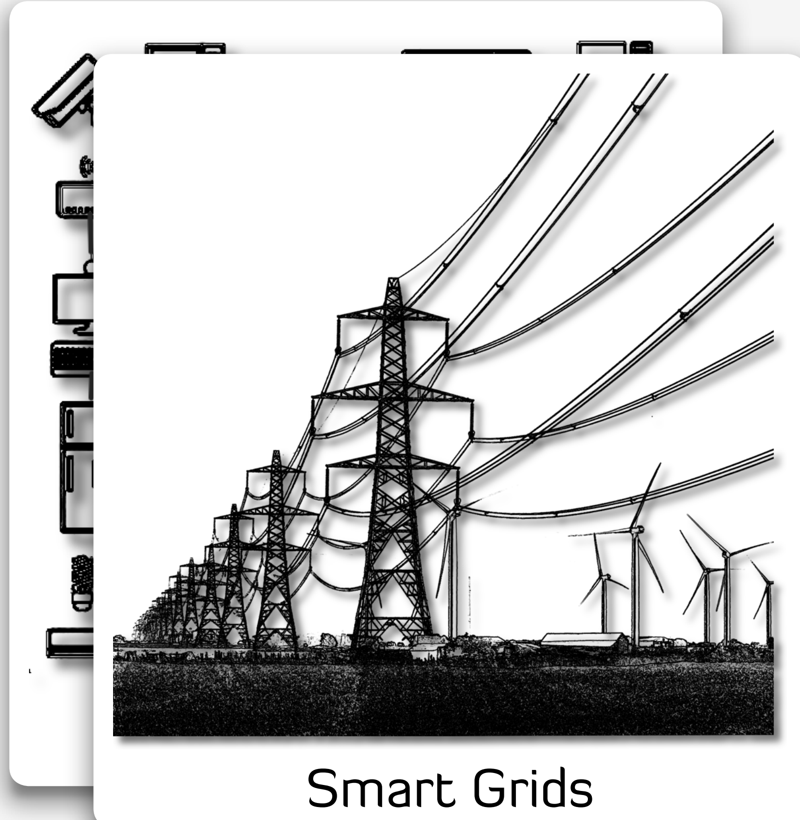

Smart Grids

Optimal decentralized control: Applications

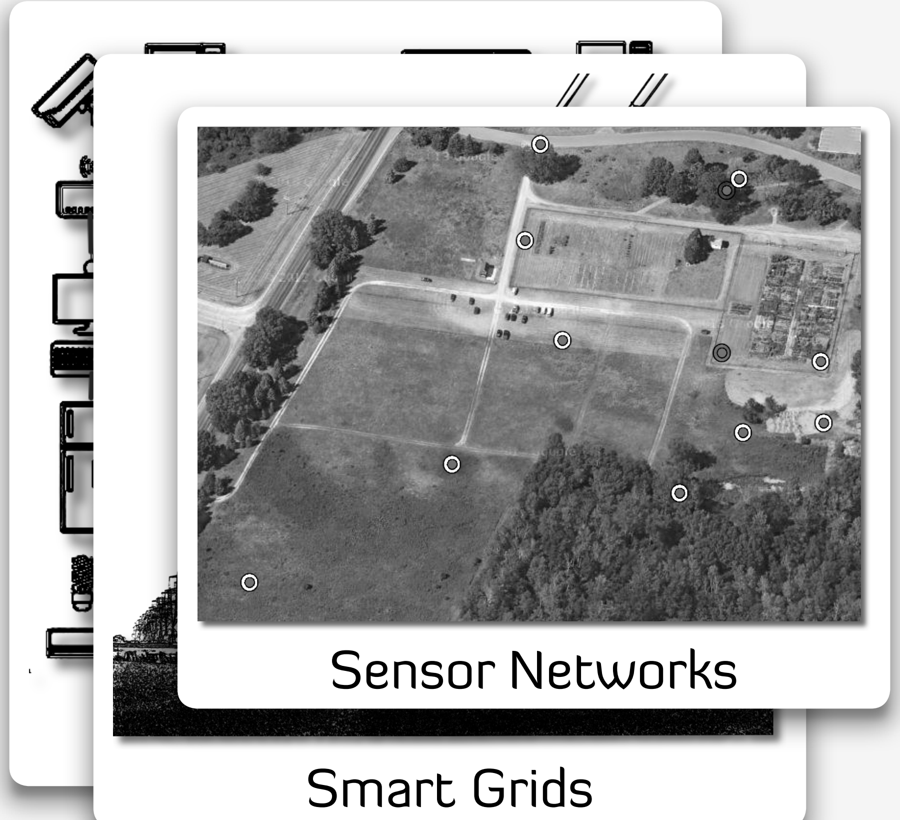

Optimal decentralized control: Applications

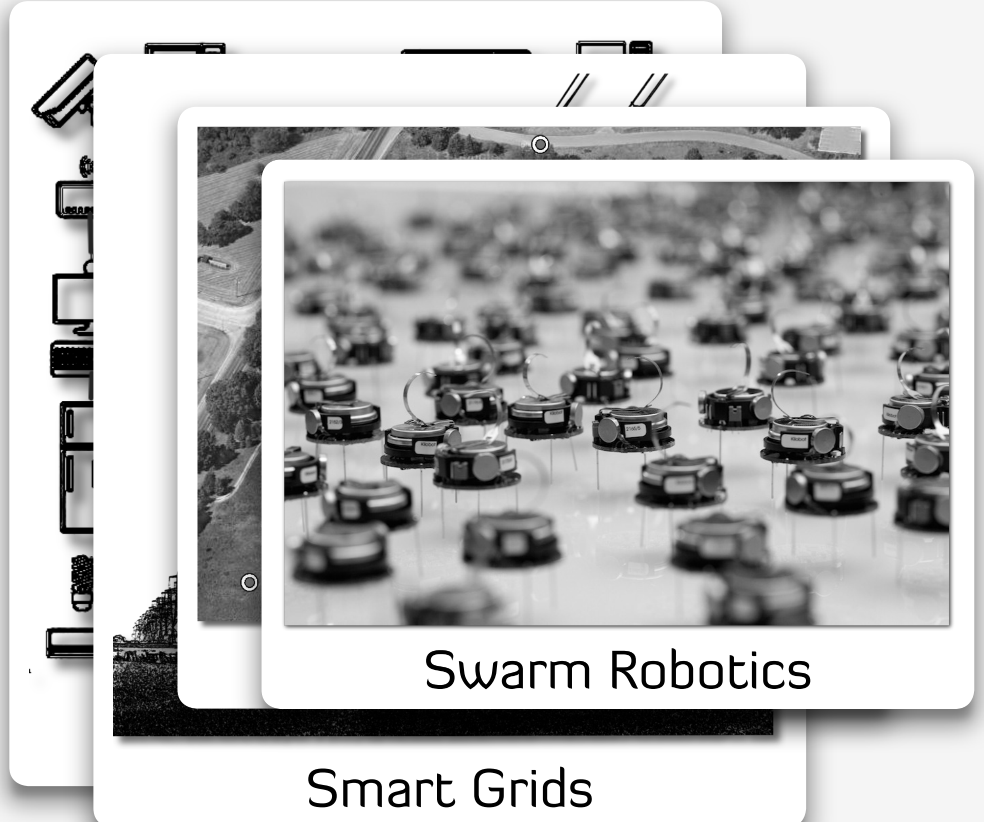

Optimal decentralized control: Applications and Theory

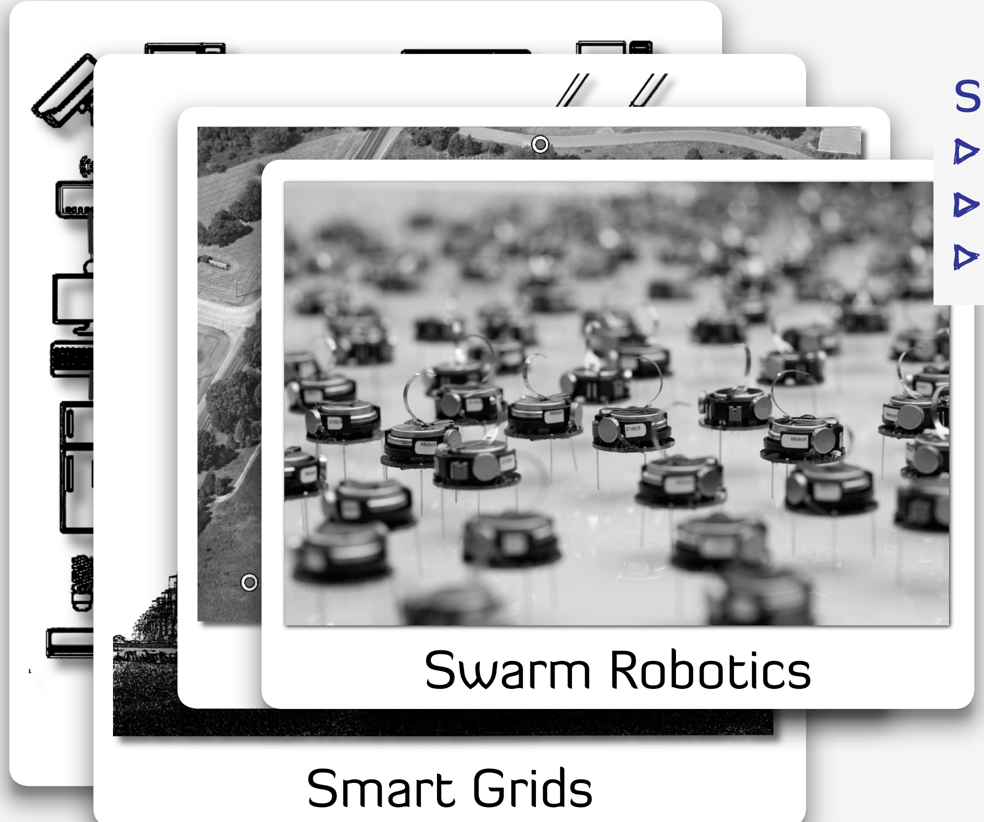

Salient features

- Multiple decision makers
- Access to different information
- Cooperate towards a common objective

Optimal decentralized control: Applications and Theory

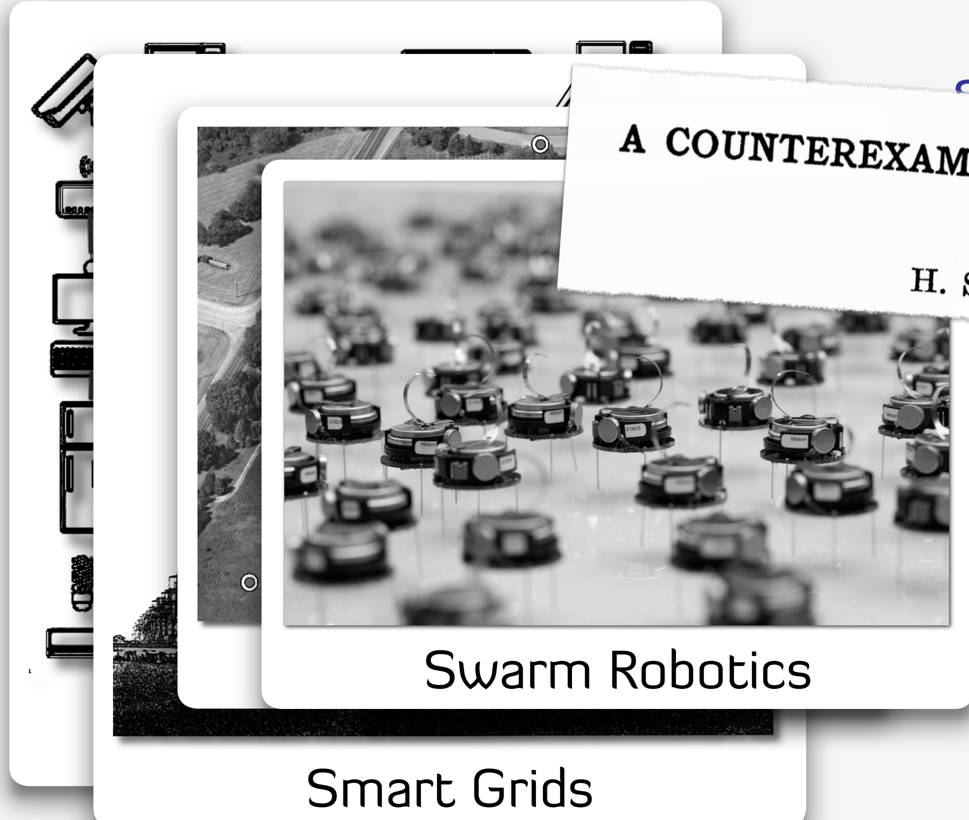

Salient features

A COUNTEREXAMPLE IN STOCHASTIC OPTIMUM CONTROL*

H. S. WITSENHAUSEN†

Decentralized control with exchangeable agents–(Arabneydi and Mahajan)

Optimal decentralized control: Applications and Theory

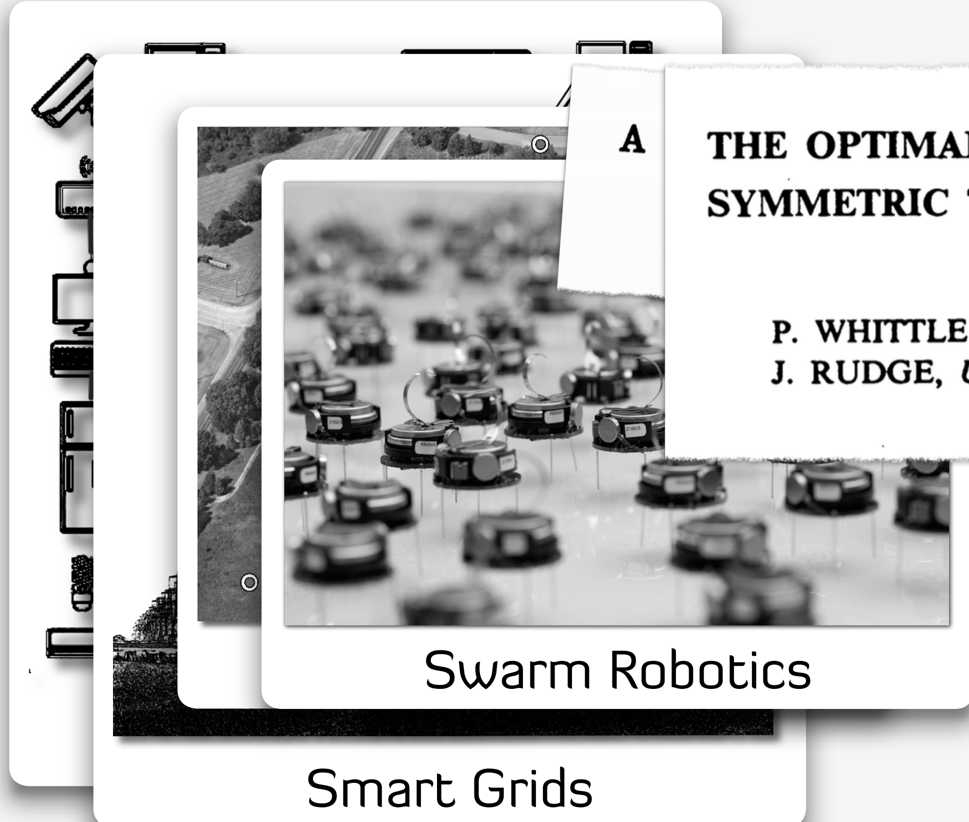

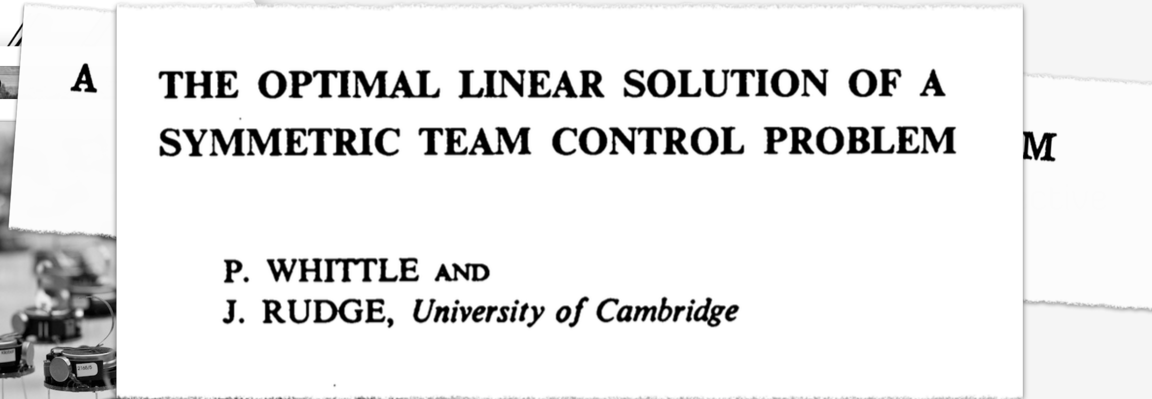
THE OPTIMAL LINEAR SOLUTION OF A SYMMETRIC TEAM CONTROL PROBLEM

P. WHITTLE AND
J. RUDGE, *University of Cambridge*

Optimal decentralized control: Applications and Theory

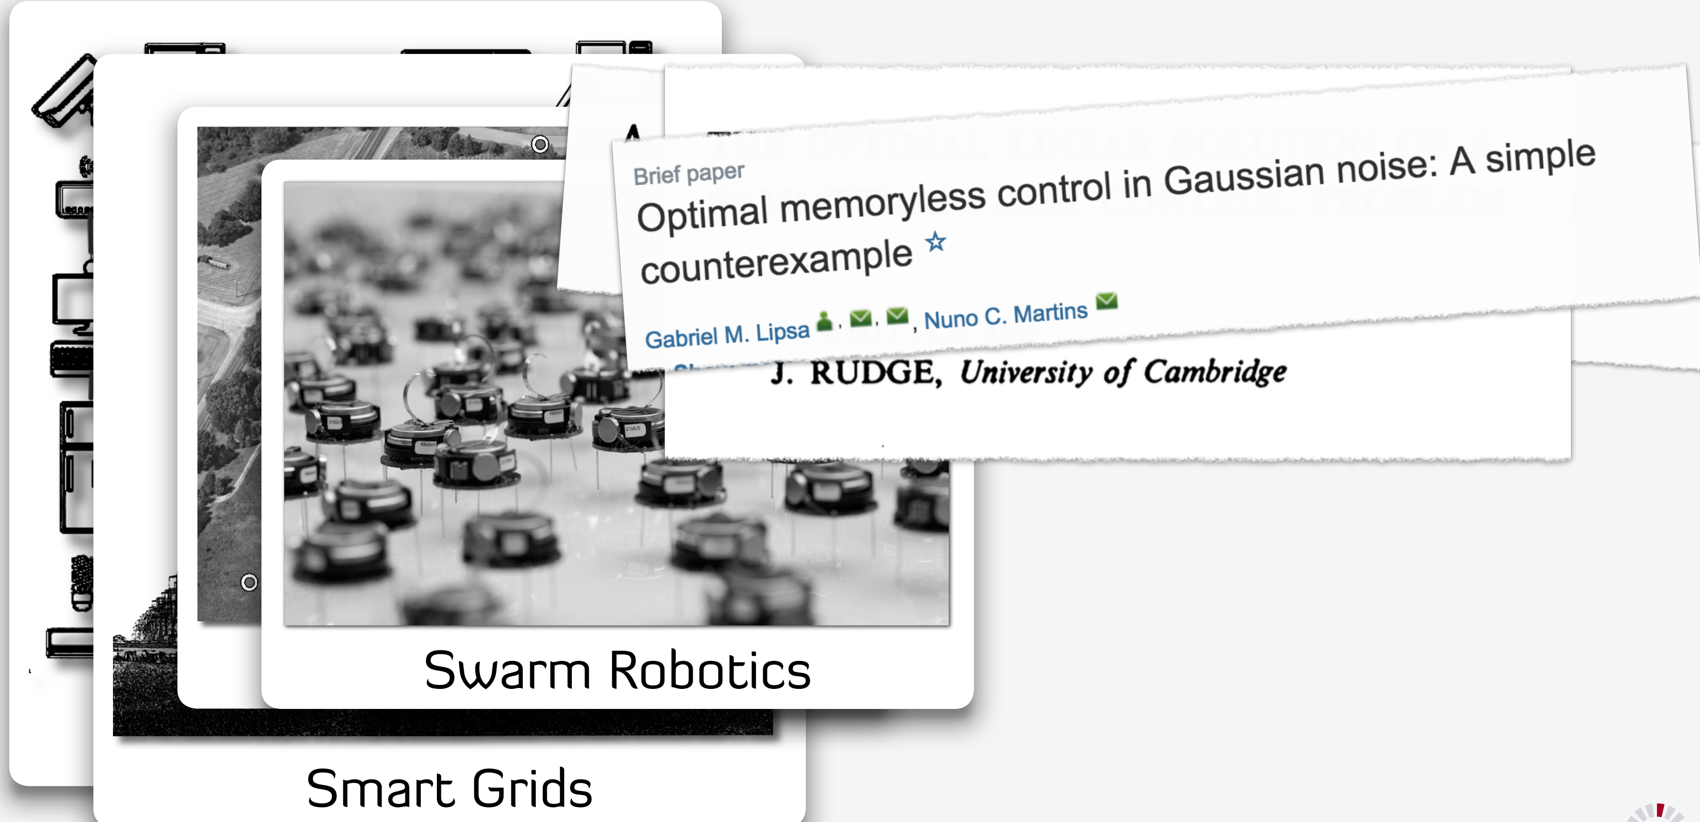

Decentralized control with exchangeable agents–(Arabneydi and Mahajan)

Optimal decentralized control: Applications and Theory

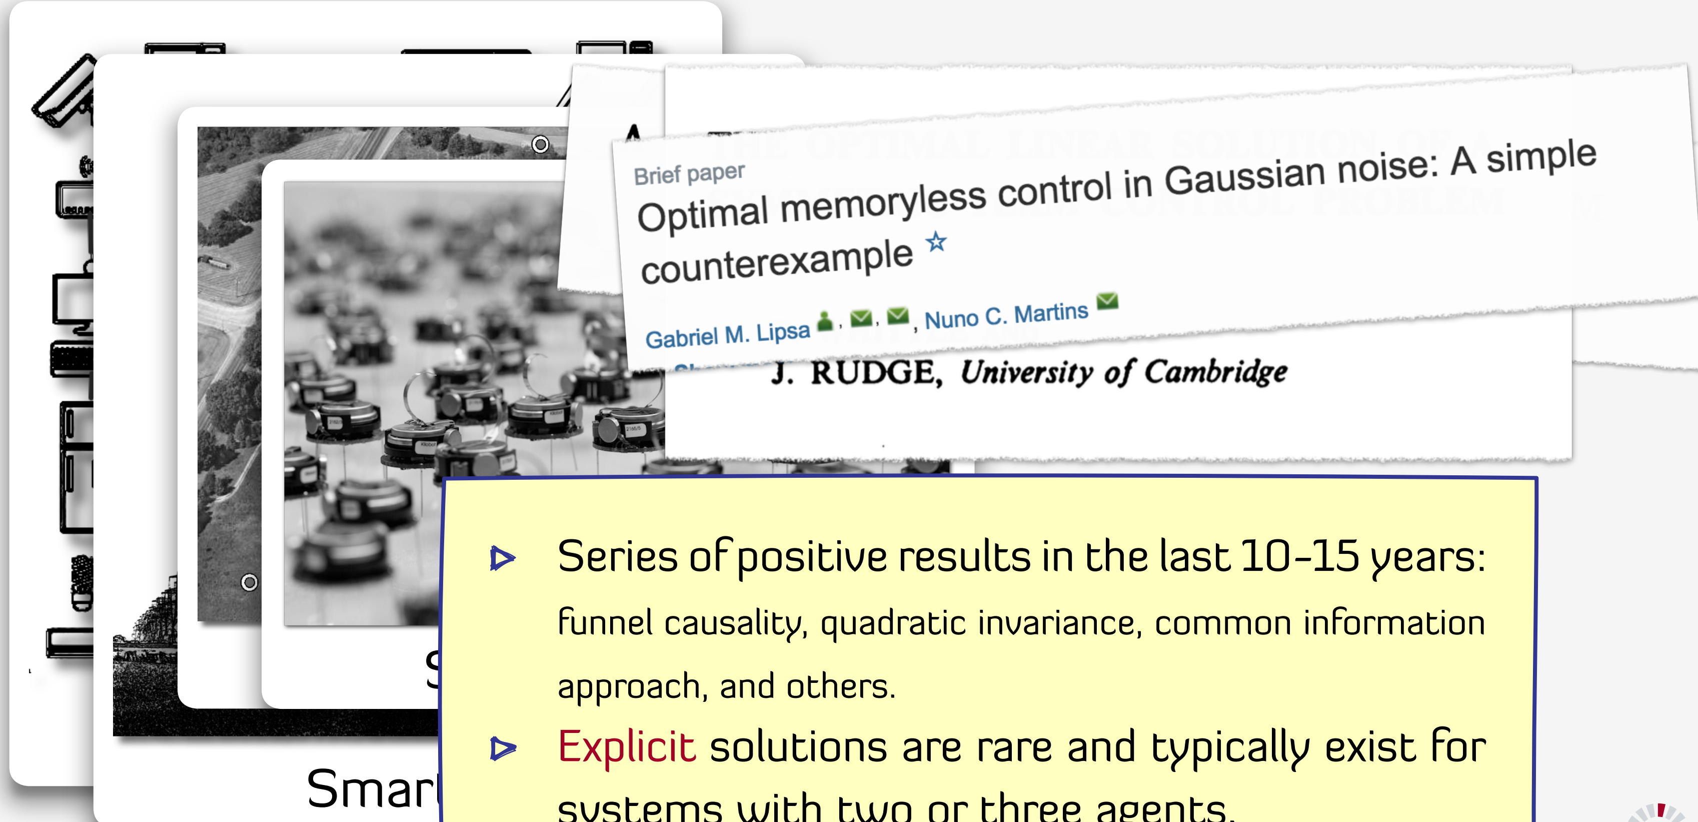

- Series of positive results in the last 10-15 years: funnel causality, quadratic invariance, common information approach, and others.
- Explicit solutions are rare and typically exist for systems with two or three agents.

Decentralized control with ex

Are there features that are present in the applications but are missing from the theory?

System with exchangeable agents

Dynamics $\mathbf{x}_{t+1} = f_t(\mathbf{x}_t, \mathbf{u}_t, \mathbf{w}_t)$ with per-step cost $c_t(\mathbf{x}_t, \mathbf{u}_t)$.

System with exchangeable agents

Dynamics $x_{t+1} = f_t(x_t, u_t, w_t)$ with per-step cost $c_t(x_t, u_t)$.

Pair of exchangeable agents

Agents i and j are exchangeable if

- $\mathcal{X}^i = \mathcal{X}^j$, $\mathcal{U}^i = \mathcal{U}^j$, $\mathcal{W}^i = \mathcal{W}^j$.
- $f_t(\sigma_{ij} \mathbf{x}_t, \sigma_{ij} \mathbf{u}_t, \sigma_{ij} \mathbf{w}_t) = \sigma_{ij} \big(f_t(\mathbf{x}_t, \mathbf{u}_t, \mathbf{w}_t) \big)$
- $c_t(\sigma_{ij} \mathbf{x}_t, \sigma_{ij} \mathbf{u}_t) = c_t(\mathbf{x}_t, \mathbf{u}_t)$.

System with exchangeable agents

Dynamics $x_{t+1} = f_t(x_t, u_t, w_t)$ with per-step cost $c_t(x_t, u_t)$.

Pair of exchangeable agents

Agents i and j are exchangeable if

- $\mathcal{X}^i = \mathcal{X}^j$, $\mathcal{U}^i = \mathcal{U}^j$, $\mathcal{W}^i = \mathcal{W}^j$.
- $f_t(\sigma_{ij} x_t, \sigma_{ij} u_t, \sigma_{ij} w_t) = \sigma_{ij}\big(f_t(x_t, u_t, w_t)\big)$
- $c_t(\sigma_{ij} x_t, \sigma_{ij} u_t) = c_t(x_t, u_t)$.

Set of exchangeable agents

A set of agents is exchangeable if every pair in that set is exchangeable

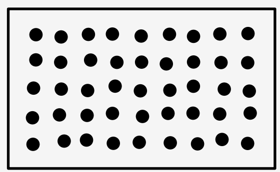

System with exchangeable agents

Dynamics $x_{t+1} = f_t(x_t, u_t, w_t)$ with per-step cost $c_t(x_t, u_t)$.

Pair of exchangeable agents

Agents i and j are exchangeable if

- $\mathcal{X}^i = \mathcal{X}^j$, $\mathcal{U}^i = \mathcal{U}^j$, $\mathcal{W}^i = \mathcal{W}^j$.
- $f_t(\sigma_{ij} x_t, \sigma_{ij} u_t, \sigma_{ij} w_t) = \sigma_{ij}\big(f_t(x_t, u_t, w_t)\big)$
- $c_t(\sigma_{ij} x_t, \sigma_{ij} u_t) = c_t(x_t, u_t)$.

Set of exchangeable agents

A set of agents is exchangeable if every pair in that set is exchangeable

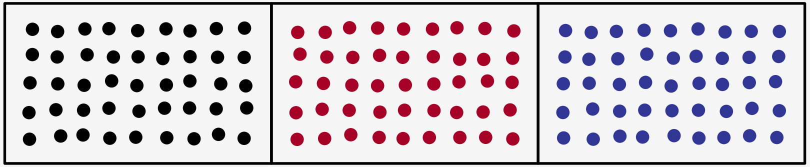

System with partially exchangeable agents

... is a multi-agent system where the set of agents can be partitioned into disjoint sets of exchangeable agents.

Notation

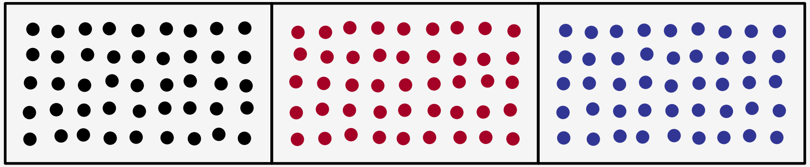

- N : number of heterogeneous agents
- K : number of subpopulations

Notation

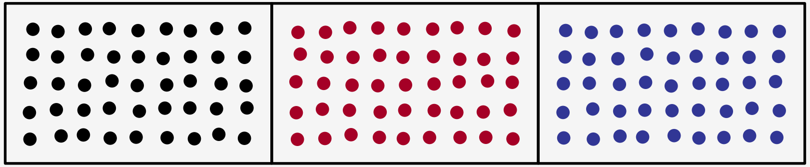

- N : number of heterogeneous agents
- K : number of subpopulations

For agent i of sub-population k

- $x_t^i \in \mathbb{R}^{d_x^k}$: state of agent i
- $u_t^i \in \mathbb{R}^{d_u^k}$: control action of agent i

Notation

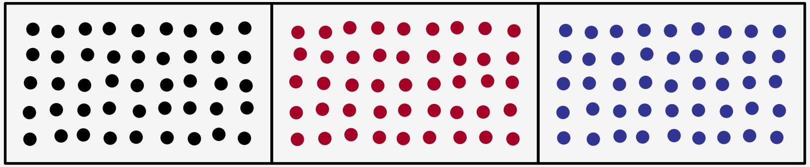

- N : number of heterogeneous agents
- K : number of subpopulations

For agent i of sub-population k

- $x_t^i \in \mathbb{R}^{d_x^k}$: state of agent i
- $u_t^i \in \mathbb{R}^{d_u^k}$: control action of agent i

For sub-population k

- $\mathcal{N}^k$: set of agents in sub-popln k
- $\bar{x}_t^k = \dfrac{1}{|\mathcal{N}^k|} \sum_{i \in \mathcal{N}^k} x_t^i$: mean-field of states
- $\bar{u}_t^k = \dfrac{1}{|\mathcal{N}^k|} \sum_{i \in \mathcal{N}^k} u_t^i$: mean-field of actions

Notation

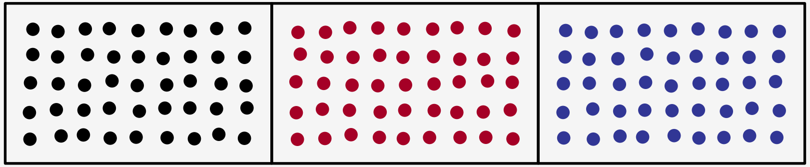

- N : number of heterogeneous agents
- K : number of subpopulations

For the entire population

- $\mathcal{N} = \mathcal{N}^1 \cup \cdots \cup \mathcal{N}^K$: set of all agents
- $\mathcal{K} = \{1, \dots, K\}$: set of all sub-populations

- $\mathbf{x}_t = (x_t^i)_{i \in \mathcal{N}}$: global state of the system
- $\mathbf{u}_t = (u_t^i)_{i \in \mathcal{N}}$: joint actions of all agents

- $\bar{\mathbf{x}}_t = \mathrm{vec}(\bar{x}_t^1, \dots, \bar{x}_t^K)$: global mean-field of states
- $\bar{\mathbf{u}}_t = \mathrm{vec}(\bar{u}_t^1, \dots, \bar{u}_t^K)$: global mean-field of actions

Linear quadratic system with partially exchangeable agents

Dynamics $\quad \mathbf{x}_{t+1} = A_t \mathbf{x}_t + B_t \mathbf{u}_t + \mathbf{w}_t$

Cost $\quad \displaystyle\sum_{t=1}^{T} \Big[\mathbf{x}_t^\top Q_t \mathbf{x}_t + \mathbf{u}_t^\top R_t \mathbf{u}_t \Big]$

Linear quadratic system with partially exchangeable agents

Dynamics $\quad x_{t+1} = A_t x_t + B_t u_t + w_t$

Cost $\quad \sum_{t=1}^{T} \left[x_t^\top Q_t x_t + u_t^\top R_t u_t \right]$

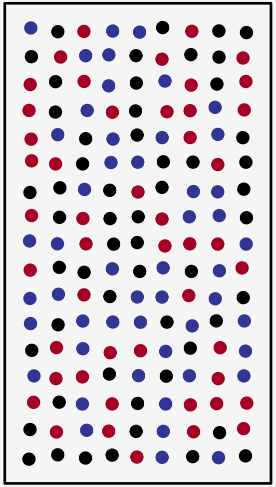

Linear quadratic system with partially exchangeable agents

Dynamics $$x_{t+1} = A_t x_t + B_t u_t + w_t$$

Cost $$\sum_{t=1}^{T} \left[x_t^\top Q_t x_t + u_t^\top R_t u_t \right]$$

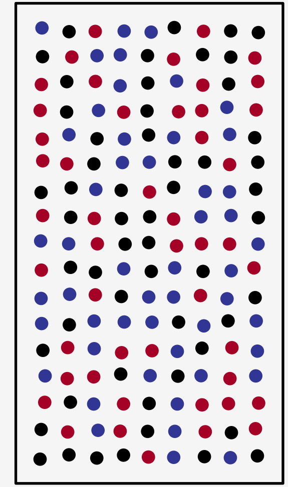

Irrespective of the information structure
such a system is equivalent to a mean-field coupled system

Linear quadratic system with partially exchangeable agents

Dynamics $$x_{t+1} = A_t x_t + B_t u_t + w_t$$

Cost $$\sum_{t=1}^{T} \left[x_t^\top Q_t x_t + u_t^\top R_t u_t \right]$$

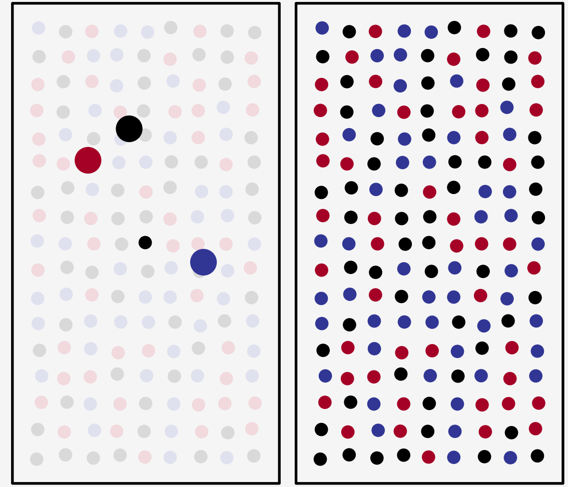

Irrespective of the information structure
such a system is equivalent to a mean-field coupled system

Agent dynamics in sub-population k $$x_{t+1}^i = A_t^k x_t^i + B_t^k u_t^i + D_t^k \bar{x}_t + E_t^k \bar{u}_t + w_t^i$$

Cost $$\sum_{t=1}^{T} \left[\sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{N}^k} \frac{1}{|\mathcal{N}^k|} \left[(x_t^i)^\top Q_t^k x_t^i + (u_t^i)^\top R_t^k u_t^i \right] + \bar{x}_t^\top P_t^x \bar{x}_t + \bar{u}_t^\top P_t^u \bar{u}_t \right]$$

There is a long history of mean-field approximations

Mean-field approximation in statistical physics (Weiss 1907; Landau 1937)

There is a long history of mean-field approximations

Mean-field approximation in statistical physics (Weiss 1907; Landau 1937)

> It is a well-known phenomenon in many branches of the exact and physical sciences that very great numbers are often easier to handle than those of medium size. An almost exact theory of a gas, containing about 10^{25} freely moving particles, is incomparably easier than that of the solar system, made up of 9 major bodies... This is, of course, due to the excellent possibility of applying the laws of statistics and probabilities in the first case.
>
> — von Neumann and Morgenstern,
> Theory of Games and Economic Behavior (1944) §2.4.2

There is a long history of mean-field approximations

Mean-field approximation in statistical physics (Weiss 1907; Landau 1937)

- . . .

Mean-field approximations in Game Theory

- Jovanovic Rosenthal 1988
- Bergin Bernhardt 1995
- Weintraub Benkard Van Roy 2008
- . . .

Anonymous games

There is a long history of mean-field approximations

Mean-field approximation in statistical physics (Weiss 1907; Landau 1937)

- . . .

Mean-field approximations in Game Theory

- Jovanovic Rosenthal 1988
- Bergin Bernhardt 1995
- Weintraub Benkard Van Roy 2008
- . . .

Anonymous games

Mean-field approximations in Systems and Control (Mean-field games)

- Huang Caines Malhalmé 2003, . . .
- Larsy Lions 2006, . . .
- . . .

Our results are different

There is no approximation!

Results are applicable to systems with arbitrary (not necessarily large) number of agents

Main idea: What happens if mean-field is observed?

Mean-field sharing
information structure

$$I^i_t = \{x^i_{1:t}, u^i_{1:t-1}, \bar{x}_{1:t}\}$$

Main idea: What happens if mean-field is observed?

Mean-field sharing information structure

$$I^i_t = \{x^i_{1:t}, u^i_{1:t-1}, \bar{x}_{1:t}\}$$

Is it a restrictive assumption?

- Not really. Mean-field can be shared using small communication overhead (using consensus algorithms)
- We later provide approx. results when mean-field is not shared.

Main idea: What happens if mean-field is observed?

Mean-field sharing information structure

$$I_t^i = \{x_{1:t}^i, u_{1:t-1}^i, \bar{x}_{1:t}\}$$

Is it a restrictive assumption?

- Not really. Mean-field can be shared using small communication overhead (using consensus algorithms)
- We later provide approx. results when mean-field is not shared.

Not one of the known tractable information structures

- Not partially nested (or stochastically nested)
- Not quadratic invariant
- Not partial history sharing

A surprisingly simple solution . . .

A surprisingly simple solution . . .

Parallel axis Theorem
$$\frac{1}{|\mathcal{N}^k|}\sum_{i\in\mathcal{N}^k}(x_t^i)^\top Q_t^k x_t^i = \frac{1}{|\mathcal{N}^k|}\sum_{i\in\mathcal{N}^k}(\breve{x}_t^i)^\top Q_t^k \breve{x}_t^i + (\bar{x}_t^k)^\top Q_t^k \bar{x}_t^k,$$

where $\breve{x}_t^i = x_t^i - \bar{x}_t^k$.

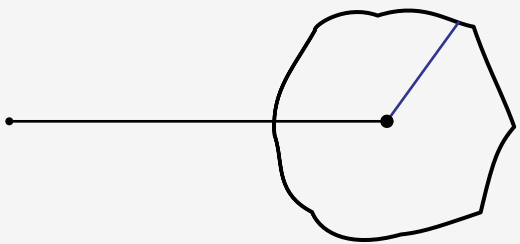

A surprisingly simple solution . . .

Parallel axis Theorem $$\frac{1}{|\mathcal{N}^k|}\sum_{i\in\mathcal{N}^k}(x_t^i)^\top Q_t^k x_t^i = \frac{1}{|\mathcal{N}^k|}\sum_{i\in\mathcal{N}^k}(\breve{x}_t^i)^\top Q_t^k \breve{x}_t^i + (\bar{x}_t^k)^\top Q_t^k \bar{x}_t^k,$$

where $\breve{x}_t^i = x_t^i - \bar{x}_t^k$.

Decoupled Per-step cost $$\sum_{k\in\mathcal{K}}\sum_{i\in\mathcal{N}^k}\frac{1}{|\mathcal{N}^k|}[(\breve{x}_t^i)^\top Q_t^k \breve{x}_t^i] + \bar{\mathbf{x}}_t^\top(\bar{Q}_t + P_t^x)\bar{\mathbf{x}}_t$$

$+$ similar u-terms

A surprisingly simple solution . . .

Parallel axis Theorem

$$\frac{1}{|\mathcal{N}^k|}\sum_{i\in\mathcal{N}^k}(x_t^i)^\top Q_t^k x_t^i = \frac{1}{|\mathcal{N}^k|}\sum_{i\in\mathcal{N}^k}(\breve{x}_t^i)^\top Q_t^k \breve{x}_t^i + (\bar{x}_t^k)^\top Q_t^k \bar{x}_t^k,$$

where $\breve{x}_t^i = x_t^i - \bar{x}_t^k$.

Decoupled Per-step cost

$$\sum_{k\in\mathcal{K}}\sum_{i\in\mathcal{N}^k}\frac{1}{|\mathcal{N}^k|}[(\breve{x}_t^i)^\top Q_t^k \breve{x}_t^i] + \bar{\mathbf{x}}_t^\top(\bar{Q}_t + P_t^x)\bar{\mathbf{x}}_t$$

$+$ similar u-terms

Noise coupled Dynamics

$$\breve{x}_{t+1}^i = A_t^k \breve{x}_t^i + B_t^k \breve{u}_t^i + \breve{w}_t^i, \quad \bar{\mathbf{x}}_{t+1} = A_t^k \bar{\mathbf{x}}_t + B_t^k \bar{\mathbf{u}}_t + \bar{\boldsymbol{w}}_t$$

A surprisingly simple solution . . .

Parallel axis Theorem

$$\frac{1}{|\mathcal{N}^k|}\sum_{i\in\mathcal{N}^k}(x_t^i)^\top Q_t^k x_t^i = \frac{1}{|\mathcal{N}^k|}\sum_{i\in\mathcal{N}^k}(\breve{x}_t^i)^\top Q_t^k \breve{x}_t^i + (\bar{x}_t^k)^\top Q_t^k \bar{x}_t^k,$$

where $\breve{x}_t^i = x_t^i - \bar{x}_t^k$.

Decoupled Per-step cost

$$\sum_{k\in\mathcal{K}}\sum_{i\in\mathcal{N}^k}\frac{1}{|\mathcal{N}^k|}[(\breve{x}_t^i)^\top Q_t^k \breve{x}_t^i] + \bar{\mathbf{x}}_t^\top(\bar{Q}_t + P_t^x)\bar{\mathbf{x}}_t$$

$+$ similar u-terms

Noise coupled Dynamics

$$\breve{x}_{t+1}^i = A_t^k \breve{x}_t^i + B_t^k \breve{u}_t^i + \breve{w}_t^i, \quad \bar{\mathbf{x}}_{t+1} = A_t^k \bar{\mathbf{x}}_t + B_t^k \bar{\mathbf{u}}_t + \bar{\mathbf{w}}_t$$

We still have a non-classical information structure

Assume centralized information and use certainty equivalence

	Local States	Mean-field state
Dynamics	$\breve{x}^i_{t+1} = A^k_t \breve{x}^i_t + B^k_t \breve{u}^i_t + \breve{w}^i_t$	$\bar{\mathbf{x}}_{t+1} = A_t \bar{\mathbf{x}}_t + B_t \bar{\mathbf{u}}_t + \bar{\mathbf{w}}_t$
Cost	$(\breve{x}^i_t)^\top Q^k_t \breve{x}^i_t + (\breve{u}^i_t)^\top R^k_t \breve{u}^i_t$	$(\bar{\mathbf{x}}_t)^\top (P^x_t + Q_t) \bar{\mathbf{x}}_t + (\bar{\mathbf{u}}_t)^\top (P^u_t + R_t) \bar{\mathbf{u}}_t$

Assume centralized information and use certainty equivalence

	Local States	Mean-field state
Dynamics	$\breve{x}^i_{t+1} = A^k_t \breve{x}^i_t + B^k_t \breve{u}^i_t + \breve{w}^i_t$	$\bar{\mathbf{x}}_{t+1} = A_t \bar{\mathbf{x}}_t + B_t \bar{\mathbf{u}}_t + \bar{\mathbf{w}}_t$
Cost	$(\breve{x}^i_t)^\top Q^k_t \breve{x}^i_t + (\breve{u}^i_t)^\top R^k_t \breve{u}^i_t$	$(\bar{\mathbf{x}}_t)^\top (P^x_t + Q_t) \bar{\mathbf{x}}_t + (\bar{\mathbf{u}}_t)^\top (P^u_t + R_t) \bar{\mathbf{u}}_t$
Control Law	$\breve{u}^i_t = \breve{L}^k_t \breve{x}^i_t$	$\bar{\mathbf{u}}_t = \bar{L}_t \bar{\mathbf{x}}_t$

Assume centralized information and use certainty equivalence

	Local States	Mean-field state
Dynamics	$\breve{x}^i_{t+1} = A^k_t \breve{x}^i_t + B^k_t \breve{u}^i_t + \breve{w}^i_t$	$\bar{\boldsymbol{x}}_{t+1} = A_t \bar{\boldsymbol{x}}_t + B_t \bar{\boldsymbol{u}}_t + \bar{\boldsymbol{w}}_t$
Cost	$(\breve{x}^i_t)^\top Q^k_t \breve{x}^i_t + (\breve{u}^i_t)^\top R^k_t \breve{u}^i_t$	$(\bar{\boldsymbol{x}}_t)^\top (P^x_t + Q_t)\bar{\boldsymbol{x}}_t + (\bar{\boldsymbol{u}}_t)^\top (P^u_t + R_t)\bar{\boldsymbol{u}}_t$
Control Law	$\breve{u}^i_t = \breve{L}^k_t \breve{x}^i_t$	$\bar{\boldsymbol{u}}_t = \bar{\boldsymbol{L}}_t \bar{\boldsymbol{x}}_t$
Gains	$\breve{L}^k_t = -(\cdots)^{-1}(B^k_t)^\top \breve{M}^k_{t+1} A^k_t$	$\bar{L}_t = -(\cdots)^{-1}(\bar{B}_t)^\top \bar{\boldsymbol{M}}_{t+1} \bar{A}_t$
Riccati Equation	$\breve{M}^k_{1:T} = \mathrm{DRE}(A^k_{1:T}, B^k_{1:T}, Q^k_{1:T}, R^k_{1:T})$	$\bar{M}_{1:T} = \mathrm{DRE}(\bar{A}_{1:T}, \bar{B}_{1:T}, \bar{Q}_{1:T} + P^x_{1:T}, \bar{R}_{1:T} + P^u_{1:T})$

Assume centralized information and use certainty equivalence

	Local States	Mean-field state
Dynamics	$\breve{x}^i_{t+1} = A^k_t \breve{x}^i_t + B^k_t \breve{u}^i_t + \breve{w}^i_t$	$\bar{\mathbf{x}}_{t+1} = A_t \bar{\mathbf{x}}_t + B_t \bar{\mathbf{u}}_t + \bar{\mathbf{w}}_t$
Cost	$(\breve{x}^i_t)^\top Q^k_t \breve{x}^i_t + (\breve{u}^i_t)^\top R^k_t \breve{u}^i_t$	$(\bar{\mathbf{x}}_t)^\top (P^x_t + Q_t)\bar{\mathbf{x}}_t + (\bar{\mathbf{u}}_t)^\top (P^u_t + R_t)\bar{\mathbf{u}}_t$
Control Law	$\breve{u}^i_t = \breve{L}^k_t \breve{x}^i_t$	$\bar{\mathbf{u}}_t = \bar{\mathbf{L}}_t \bar{\mathbf{x}}_t$
Gains	$\breve{L}^k_t = -(\cdots)^{-1} (B^k_t)^\top \breve{M}^k_{t+1} A^k_t$	$\bar{L}_t = -(\cdots)^{-1} (\bar{B}_t)^\top \bar{\mathbf{M}}_{t+1} \bar{A}_t$
Riccati Equation	$\breve{M}^k_{1:T} = \mathsf{DRE}(A^k_{1:T}, B^k_{1:T}, Q^k_{1:T}, R^k_{1:T})$	$\bar{M}_{1:T} = \mathsf{DRE}(\bar{A}_{1:T}, \bar{B}_{1:T}, \bar{Q}_{1:T} + P^x_{1:T}, \bar{R}_{1:T} + P^u_{1:T})$

K equations, one for each sub-population

1 equation for all mean-fields

$$u_t^i = \breve{u}_t^i + \bar{u}_t^k = \breve{L}_t^k (x_t^i - \bar{x}_t^k) + \bar{L}_t^k \bar{\mathbf{x}}_t$$

Optimal centralized solution can be implemented with mean-field sharing information structure.

Solution generalizes to . . .

Major-minor setup — One major agent and a population of minor agents.

Tracking cost function

$$\sum_{k\in\mathcal{K}}\sum_{i\in\mathcal{N}^k}\frac{1}{|\mathcal{N}^k|}\Big[(x_t^i-\mathring{x}_t^i)^\top Q_t^k(x_t^i-\mathring{x}_t^i)+(u_t^i)^\top R_t^k u_t^i\Big]$$
$$+(\bar{\mathbf{x}}_t-r_t)^\top P_t^x(\bar{\mathbf{x}}_t-r_t)+\bar{\mathbf{u}}_t^\top P_t^u\bar{\mathbf{u}}_t$$

Systems coupled through weighted mean-field

$$\bar{x}_t^k=\frac{1}{|\mathcal{N}^k|}\sum_{i\in\mathcal{N}^k}\lambda^i x_t^i,\qquad \bar{u}_t^k=\frac{1}{|\mathcal{N}^k|}\sum_{i\in\mathcal{N}^k}\lambda^i u_t^i.$$

But what if the mean-field is not observed?

Partial mean-field sharing information structure

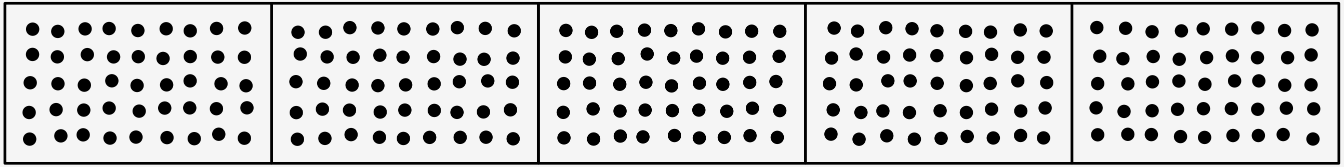

Partial mean-field sharing information structure

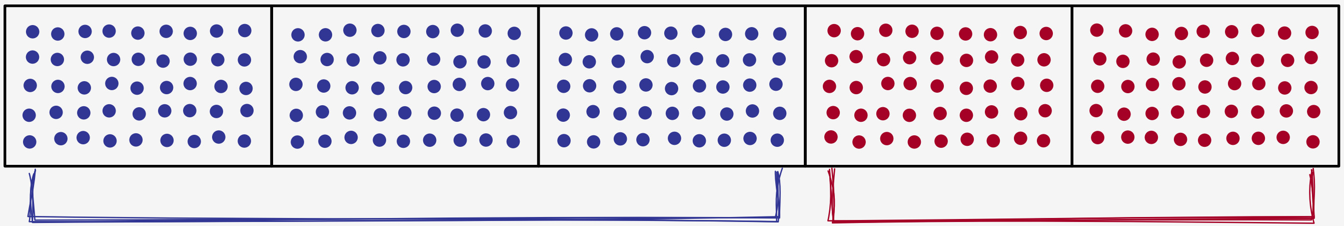

Set $\mathcal{S}$: MF observed   Set $\mathcal{S}^c$: MF not observed

Notation

We will compare performance with system where mean-field is completely observed. To avoid confusion, use

State: s^i_t;   Actions: ν^i_t.

and similar notation for mean-field $\bar{s}^k_t$, etc.

Partial mean-field sharing information structure

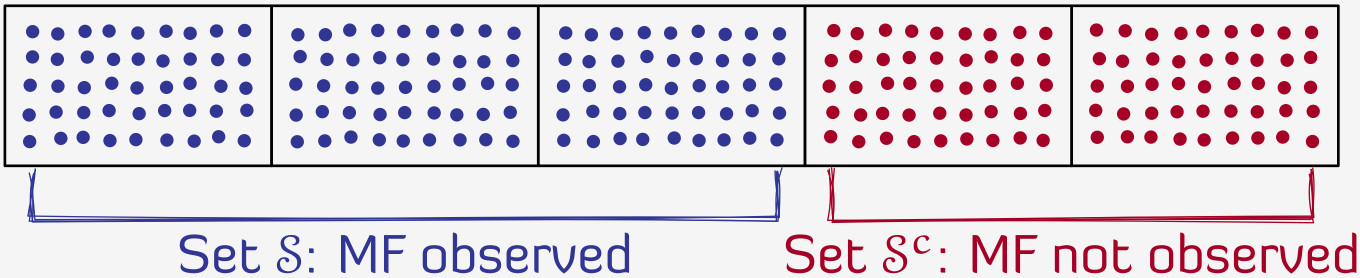

Estimated mean-field $\quad \mathbf{z}_t = (z_t^1, \dots, z_t^K) = \mathbb{E}[\bar{\mathbf{s}}_t \mid \{\bar{s}_t^k\}_{k \in \mathcal{N}}],$

where $z_{t+1}^k = \begin{cases} \bar{s}_{t+1}^k, & k \in \mathcal{S} \\ A_t^k z_t^k + (B_t^k \bar{L}_t^k + D_t^k + E_t^k \bar{L}_t)\mathbf{z}_t, & k \notin \mathcal{S} \end{cases}$

Certainty equivalence controller and its performance

Certainty equivalence controller

$$u_t^i = \breve{L}_t^k (s_t^i - z_t^k) + \bar{L}_t^k z_t$$

Certainty equivalence controller and its performance

Certainty equivalence controller $u^i_t = \breve{L}^k_t (s^i_t - z^k_t) + \bar{L}^k_t z_t$

Key Lemma Under the certainty equivalence control: $\breve{s}^i_t = \breve{x}^i_t$.

Certainty equivalence controller and its performance

Certainty equivalence controller $\quad u_t^i = \breve{L}_t^k(s_t^i - z_t^k) + \bar{L}_t^k z_t$

Key Lemma $\quad$ Under the certainty equivalence control: $\breve{s}_t^i = \breve{x}_t^i$. Thus,

$$\hat{J} - J^* = \mathbb{E}\Big[\sum_{t=1}^{T}[\bar{\mathbf{s}}_t^\top \hat{Q}_t \bar{\mathbf{s}}_t + \bar{\boldsymbol{v}}_t^\top \hat{R}_t \bar{\boldsymbol{v}}_t - \bar{\mathbf{x}}_t^\top \hat{Q}_t \bar{\mathbf{x}}_t - \bar{\mathbf{u}}_t^\top \hat{R}_t \bar{\mathbf{u}}_t]\Big]$$

Certainty equivalence controller and its performance

Certainty equivalence controller

$$u_t^i = \breve{L}_t^k (s_t^i - z_t^k) + \bar{L}_t^k z_t$$

Key Lemma

Under the certainty equivalence control: $\breve{s}_t^i = \breve{x}_t^i$. Thus,

$$\begin{aligned}\hat{J} - J^* &= \mathbb{E}\Big[\sum_{t=1}^{T} [\bar{\mathbf{s}}_t^\top \hat{Q}_t \bar{\mathbf{s}}_t + \bar{\mathbf{v}}_t^\top \hat{R}_t \bar{\mathbf{v}}_t - \bar{\mathbf{x}}_t^\top \hat{Q}_t \bar{\mathbf{x}}_t - \bar{\mathbf{u}}_t^\top \hat{R}_t \bar{\mathbf{u}}_t]\Big] \\ &= \mathbb{E}\Big[\sum_{t=1}^{T} \begin{bmatrix} \zeta_t & \xi_t \end{bmatrix} \tilde{Q} \begin{bmatrix} \zeta_t \\ \xi_t \end{bmatrix}\Big], \text{ where } \zeta_t^k = \bar{x}_t^k - z_t^k \text{ and } \xi_t^k = \bar{s}_t^k - z_t^k\end{aligned}$$

Certainty equivalence controller and its performance

Certainty equivalence controller $\quad u_t^i = \breve{L}_t^k(s_t^i - z_t^k) + \bar{L}_t^k z_t$

Key Lemma $\quad$ Under the certainty equivalence control: $\breve{s}_t^i = \breve{x}_t^i$. Thus,

$$\hat{J} - J^* = \mathbb{E}\Big[\sum_{t=1}^{T}[\bar{\mathbf{s}}_t^\top \hat{Q}_t \bar{\mathbf{s}}_t + \bar{\boldsymbol{v}}_t^\top \hat{R}_t \bar{\boldsymbol{v}}_t - \bar{\mathbf{x}}_t^\top \hat{Q}_t \bar{\mathbf{x}}_t - \bar{\mathbf{u}}_t^\top \hat{R}_t \bar{\mathbf{u}}_t]\Big]$$

$$= \mathbb{E}\Big[\sum_{t=1}^{T}\begin{bmatrix}\zeta_t & \xi_t\end{bmatrix}\tilde{Q}\begin{bmatrix}\zeta_t \\ \xi_t\end{bmatrix}\Big],$$ where $\zeta_t^k = \bar{x}_t^k - z_t^k$ and $\xi_t^k = \bar{s}_t^k - z_t^k$

Moreover, $\begin{bmatrix}\zeta_{t+1} \\ \xi_{t+1}\end{bmatrix} = \tilde{A}_t\begin{bmatrix}\zeta_t \\ \xi_t\end{bmatrix} + \begin{bmatrix}h \circ \bar{\boldsymbol{w}}_t \\ h \circ \bar{\boldsymbol{w}}_t\end{bmatrix}$

Certainty equivalence controller and its performance

Certainty equivalence controller

$$u_t^i = \breve{L}_t^k (s_t^i - z_t^k) + \bar{L}_t^k z_t$$

Key Lemma

Under the certainty equivalence control: $\breve{s}_t^i = \breve{x}_t^i$. Thus,

$$\hat{J} - J^* = \mathbb{E}\Big[\sum_{t=1}^{T} [\bar{\mathbf{s}}_t^\top \hat{Q}_t \bar{\mathbf{s}}_t + \bar{\boldsymbol{v}}_t^\top \hat{R}_t \bar{\boldsymbol{v}}_t - \bar{\mathbf{x}}_t^\top \hat{Q}_t \bar{\mathbf{x}}_t - \bar{\mathbf{u}}_t^\top \hat{R}_t \bar{\mathbf{u}}_t]\Big]$$

Quadratic Cost

$$= \mathbb{E}\Big[\sum_{t=1}^{T} \begin{bmatrix} \zeta_t & \xi_t \end{bmatrix} \tilde{Q} \begin{bmatrix} \zeta_t \\ \xi_t \end{bmatrix}\Big], \text{ where } \zeta_t^k = \bar{x}_t^k - z_t^k \text{ and } \xi_t^k = \bar{s}_t^k - z_t^k$$

Moreover, $\begin{bmatrix} \zeta_{t+1} \\ \xi_{t+1} \end{bmatrix} = \tilde{A}_t \begin{bmatrix} \zeta_t \\ \xi_t \end{bmatrix} + \begin{bmatrix} h \circ \bar{\boldsymbol{w}}_t \\ h \circ \bar{\boldsymbol{w}}_t \end{bmatrix}$

Linear Dynamics

Certainty equivalence controller and its performance

Exact Performance $\quad \hat{J} - J^* = \mathrm{Tr}(\tilde{X}_1 \tilde{M}_1) + \sum_{t=1}^{T-1} \mathrm{Tr}(\tilde{W}_t \tilde{M}_{t+1})$ where $\tilde{M}_{1:T} = \mathrm{DLE}(\tilde{A}_{1:T}, \tilde{Q}_{1:T})$

Certainty equivalence controller and its performance

Exact Performance

$$\hat{J} - J^* = \mathrm{Tr}(\tilde{X}_1 \tilde{M}_1) + \sum_{t=1}^{T-1} \mathrm{Tr}(\tilde{W}_t \tilde{M}_{t+1}) \text{ where } \tilde{M}_{1:T} = \mathrm{DLE}(\tilde{A}_{1:T}, \tilde{Q}_{1:T})$$

Performance bound

Let $n = \min_{k \notin S}\{|\mathcal{N}^k|\}$. Suppose all noises are independent. Then, there exists a matrix C such that $\tilde{X}_1 \leqslant C/n$ and $\tilde{W}_t \leqslant C/n$. Thus,

$$\hat{J} - J^* \in \mathcal{O}\left(\frac{T}{n}\right),$$

Certainty equivalence controller and its performance

Exact Performance

$$\hat{J} - J^* = \mathrm{Tr}(\tilde{X}_1 \tilde{M}_1) + \sum_{t=1}^{T-1} \mathrm{Tr}(\tilde{W}_t \tilde{M}_{t+1}) \text{ where } \tilde{M}_{1:T} = \mathrm{DLE}(\tilde{A}_{1:T}, \tilde{Q}_{1:T})$$

Performance bound

Let $n = \min_{k \notin S}\{|\mathcal{N}^k|\}$. Suppose all noises are independent. Then, there exists a matrix C such that $\tilde{X}_1 \leqslant C/n$ and $\tilde{W}_t \leqslant C/n$. Thus,

$$\hat{J} - J^* \in \mathcal{O}\left(\frac{T}{n}\right),$$

Infinite horizon

Results extend to infinite horizon setup under standard assumptions. For both discounted and average cost setup:

$$\hat{J} - J^* \in \mathcal{O}\left(\frac{1}{n}\right)$$

Certainty equivalence controller and its performance

Exact Performance

$$\hat{J} - J^* = \mathrm{Tr}(\tilde{X}_1 \tilde{M}_1) + \sum_{t=1}^{T-1} \mathrm{Tr}(\tilde{W}_t \tilde{M}_{t+1}) \text{ where } \tilde{M}_{1:T} = \mathrm{DLE}(\tilde{A}_{1:T}, \tilde{Q}_{1:T})$$

Performance bound

Let $n = \min_{k \notin S}\{|\mathcal{N}^k|\}$. Suppose all noises are independent. Then, there exists a matrix C such that $\tilde{X}_1 \leqslant C/n$ and $\tilde{W}_t \leqslant C/n$. Thus,

$$\hat{J} - J^* \in \mathcal{O}\left(\frac{T}{n}\right),$$

Infinite horizon

Results extend to infinite horizon setup under standard assumptions.

For both discounted and average cost setup:

$$\hat{J} - J^* \in \mathcal{O}\left(\frac{1}{n}\right), \qquad \text{c.f. } \mathcal{O}\left(\frac{1}{\sqrt{n}}\right) \text{ in MFG}$$

An example: Demand response with minimum discomfort to users

Demand response of space heaters

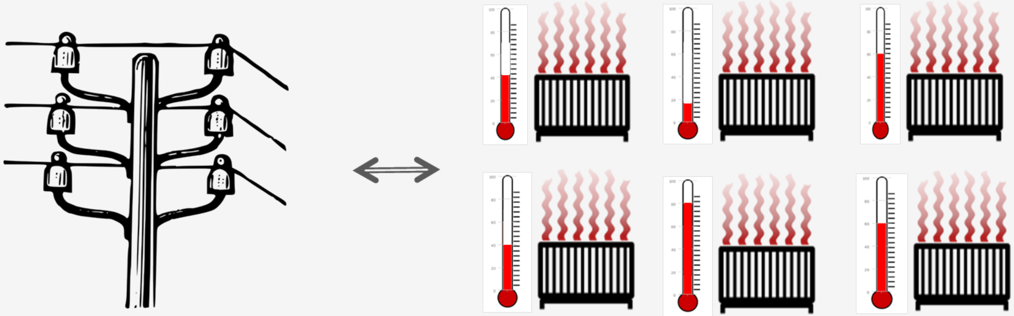

Dynamics of space heater $$x^i_{t+1} = a(x^i_t - x_{\text{nom}}) + b(u^i_t + u_{\text{nom}}) + w^i_t$$

Objective $$\mathbb{E}\left[\frac{1}{n}\sum_{t=1}^{T}\sum_{i=1}^{n}\left[q_t(x^i_t - x^i_{\text{des}})^2 + r_t(u^i_t)^2\right] + p_t(\bar{x}_t - \bar{x}^{\text{ref}}_t)^2\right]$$

Everyone follows the mean-field

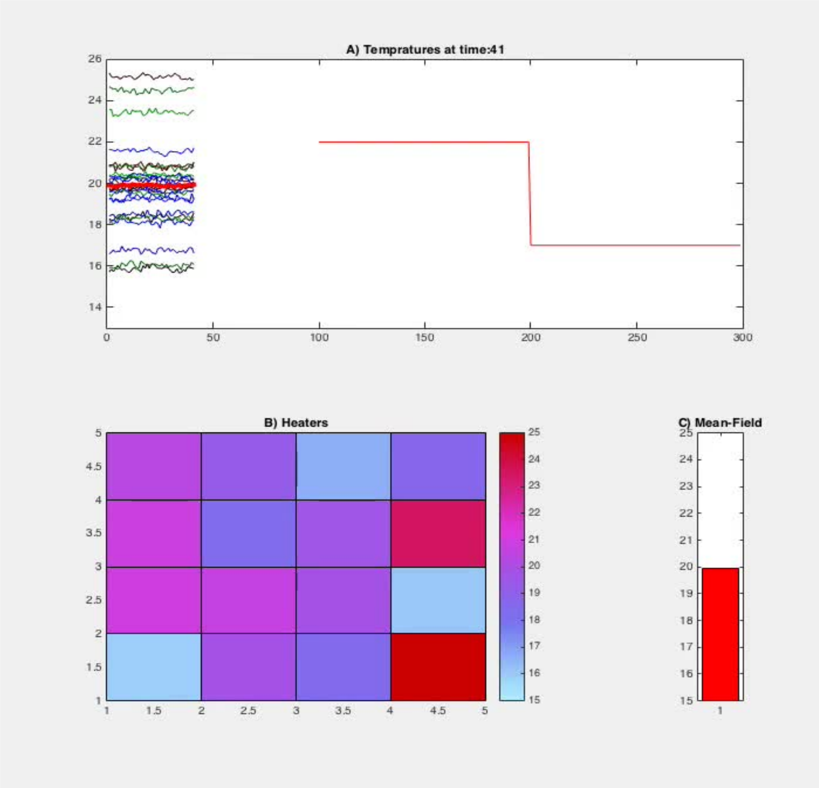

Everyone follows the optimal strategy

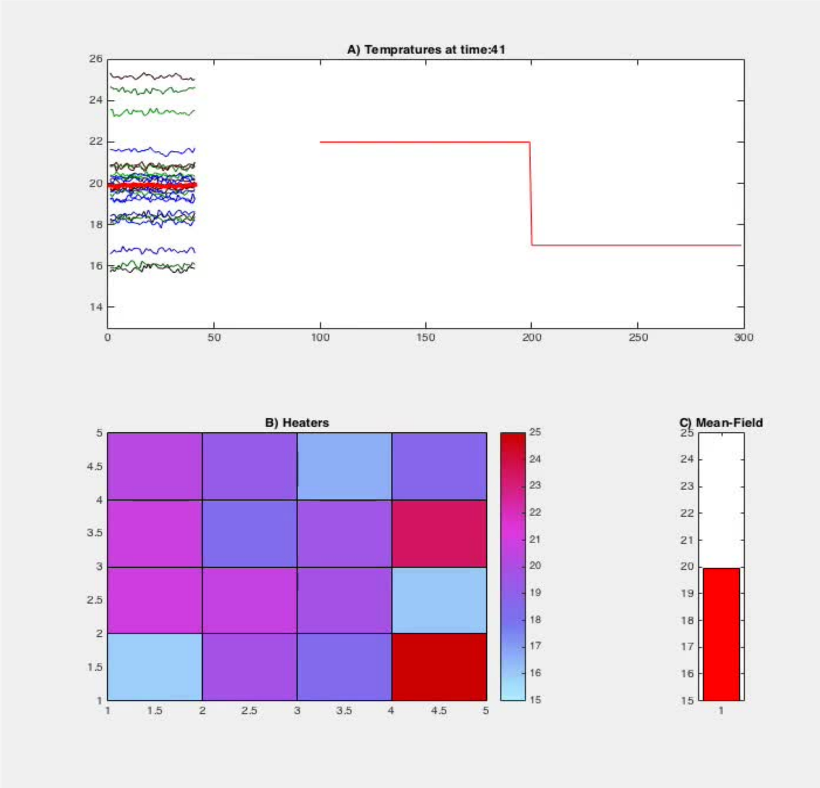

Summary

System with exchangeable agents

Dynamics $\mathbf{x}_{t+1} = f_t(\mathbf{x}_t, \mathbf{u}_t, \mathbf{w}_t)$ with per-step cost $c_t(\mathbf{x}_t, \mathbf{u}_t)$.

Pair of exchangeable agents

Agents i and j are exchangeable if

- $\mathcal{X}^i = \mathcal{X}^j$, $\mathcal{U}^i = \mathcal{U}^j$, $\mathcal{W}^i = \mathcal{W}^j$.
- $f_t(\sigma_{ij}\mathbf{x}_t, \sigma_{ij}\mathbf{u}_t, \sigma_{ij}\mathbf{w}_t) = \sigma_{ij}\big(f_t(\mathbf{x}_t, \mathbf{u}_t, \mathbf{w}_t)\big)$
- $c_t(\sigma_{ij}\mathbf{x}_t, \sigma_{ij}\mathbf{u}_t) = c_t(\mathbf{x}_t, \mathbf{u}_t)$.

Set of exchangeable agents

A set of agents is exchangeable if every pair in that set is exchangeable

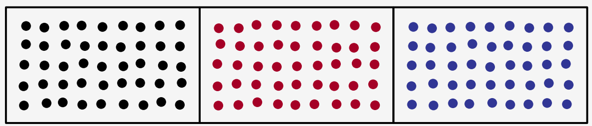

System with partially exchangeable agents

. . . is a multi-agent system where the set of agents can be partitioned into disjoint sets of exchangeable agents.

Summary

System with exchangeable agents

Linear quadratic system with partially exchangeable agents

Dynamics $\mathbf{x}_{t+1} = A_t \mathbf{x}_t + B_t \mathbf{u}_t + \mathbf{w}_t$

Cost $\displaystyle\sum_{t=1}^{T} \left[\mathbf{x}_t^\top Q_t \mathbf{x}_t + \mathbf{u}_t^\top R_t \mathbf{u}_t\right]$

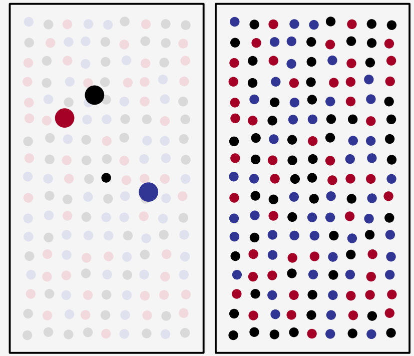

Irrespective of the information structure
such a system is equivalent to a mean-field coupled system

Agent dynamics in sub-population k $\quad x^i_{t+1} = A^k_t x^i_t + B^k_t u^i_t + D^k_t \bar{\mathbf{x}}_t + E^k_t \bar{\mathbf{u}}_t + w^i_t$

Cost $\displaystyle\sum_{t=1}^{T} \left[\sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{N}^k} \frac{1}{|\mathcal{N}^k|} \left[(x^i_t)^\top Q^k_t x^i_t + (u^i_t)^\top R^k_t u^i_t\right] + \bar{\mathbf{x}}_t^\top P^x_t \bar{\mathbf{x}}_t + \bar{\mathbf{u}}_t^\top P^u_t \bar{\mathbf{u}}_t \right]$

Summary

System with exchangeable agents

Linear quadratic system with partially exchangeable agents

Assume centralized information and use certainty equivalence

	Local States	Mean-field state
Dynamics	$\breve{x}^i_{t+1} = A^k_t \breve{x}^i_t + B^k_t \breve{u}^i_t + \breve{w}^i_t$	$\bar{\mathbf{x}}_{t+1} = A_t \bar{\mathbf{x}}_t + B_t \bar{\mathbf{u}}_t + \bar{\mathbf{w}}_t$
Cost	$(\breve{x}^i_t)^\top Q^k_t \breve{x}^i_t + (\breve{u}^i_t)^\top R^k_t \breve{u}^i_t$	$(\bar{\mathbf{x}}_t)^\top (P^x_t + Q_t) \bar{\mathbf{x}}_t + (\bar{\mathbf{u}}_t)^\top (P^u_t + R_t) \bar{\mathbf{u}}_t$
Control Law	$\breve{u}^i_t = \breve{L}^k_t \breve{x}^i_t$	$\bar{\mathbf{u}}_t = \bar{L}_t \bar{\mathbf{x}}_t$
Gains	$\breve{L}^k_t = -(\cdots)^{-1} (B^k_t)^\top \breve{M}^k_{t+1} A^k_t$	$\bar{L}_t = -(\cdots)^{-1} (\bar{B}_t)^\top \bar{M}_{t+1} \bar{A}_t$
Riccati Equation	$\breve{M}^k_{1:T} = \mathrm{DRE}(A^k_{1:T}, B^k_{1:T}, Q^k_{1:T}, R^k_{1:T})$	$\bar{M}_{1:T} = \mathrm{DRE}(\bar{A}_{1:T}, \bar{B}_{1:T}, \bar{Q}_{1:T} + P^x_{1:T}, \bar{R}_{1:T} + P^u_{1:T})$

K equations, one for each sub-population

1 equation for all mean-fields

Summary

System with exchangeable agents

Linear quadratic system with partially exchangeable agents

Assume centralized information and use certainty equivalence

Partial mean-field sharing information structure

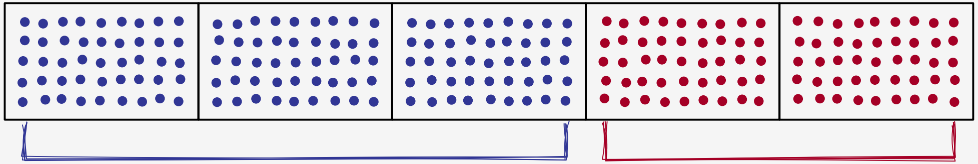

Set $\mathcal{S}$: MF observed   Set $\mathcal{S}^c$: MF not observed

Estimated mean-field $\quad z_t = (z_t^1, \dots, z_t^K) = \mathbb{E}[\bar{\mathbf{s}}_t \mid \{\bar{s}_t^k\}_{k \in \mathcal{N}}],$

$$\text{where } z_{t+1}^k = \begin{cases} \bar{s}_{t+1}^k, & k \in \mathcal{S} \\ A_t^k z_t^k + (B_t^k \bar{L}_t^k + D_t^k + E_t^k \bar{L}_t) z_t, & k \notin \mathcal{S} \end{cases}$$

Summary

System with exchangeable agents

Linear quadratic system with partially exchangeable agents

Assume centralized information and use certainty equivalence

Partial mean field sharing information structure

Certainty equivalence controller and its performance

Exact Performance

$$\hat{J} - J^* = \mathrm{Tr}(\tilde{X}_1 \tilde{M}_1) + \sum_{t=1}^{T-1} \mathrm{Tr}(\tilde{W}_t \tilde{M}_{t+1}) \text{ where } \tilde{M}_{1:T} = \mathrm{DLE}(\tilde{A}_{1:T}, \tilde{Q}_{1:T})$$

Performance bound

Let $n = \min_{k \notin S}\{|\mathcal{N}^k|\}$. Suppose all noises are independent. Then, there exists a matrix C such that $\tilde{X}_1 \leqslant C/n$ and $\tilde{W}_t \leqslant C/n$. Thus,

$$\hat{J} - J^* \in \mathcal{O}\left(\frac{T}{n}\right),$$

Infinite horizon

Results extend to infinite horizon setup under standard assumptions. For both discounted and average cost setup:

$$\hat{J} - J^* \in \mathcal{O}\left(\frac{1}{n}\right), \qquad \text{c.f. } \mathcal{O}\left(\frac{1}{\sqrt{n}}\right) \text{ in MFG}$$

Summary

System with exchangeable agents

Linear quadratic system with partially exchangeable agents

Assume centralized information and use certainty equivalence

Partial mean-field sharing information structure

Certainty equivalence controller and its performance

Solution generalizes to . . .

Major-minor setup: One major agent and a population of minor agents.

Tracking cost function:

$$\sum_{k\in\mathcal{K}}\sum_{i\in\mathcal{N}^k}\frac{1}{|\mathcal{N}^k|}\Big[(x^i_t-\mathring{x}^i_t)^\top Q^k_t(x^i_t-\mathring{x}^i_t)+(u^i_t)^\top R^k_t u^i_t\Big] + (\bar{\mathbf{x}}_t-r_t)^\top P^x_t(\bar{\mathbf{x}}_t-r_t)+\bar{\mathbf{u}}^\top_t P^u_t\bar{\mathbf{u}}_t$$

Systems coupled through weighted mean-field:

$$\bar{x}^k_t=\frac{1}{|\mathcal{N}^k|}\sum_{i\in\mathcal{N}^k}\lambda^i x^i_t,\qquad \bar{u}^k_t=\frac{1}{|\mathcal{N}^k|}\sum_{i\in\mathcal{N}^k}\lambda^i u^i_t.$$

Decentralized control with exchangeable agents–(Arabneydi and Mahajan)

Conclusion

Salient Features

- The solution complexity depends only on the number of sub-populations; not on the number of agents.
- Agents don't need to be aware of the number of agents.
- Same performance as centralized information.

> Thus, centralized performance can be achieved by simply sharing the mean-field (empirical mean) of the states!

Generalizations

- Noisy observation of mean-field
- Delay in the observation of mean-field
- Controlled Markov processes

arXiv:1609.00056

Optimal decentralized control: Applications and Theory

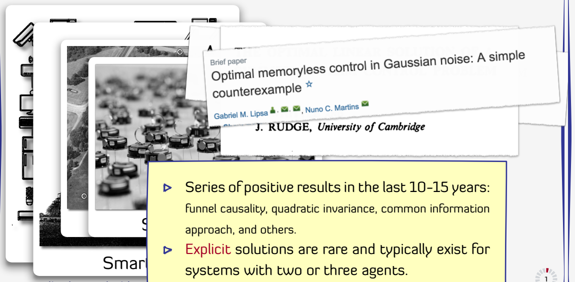

- Series of positive results in the last 10-15 years: funnel causality, quadratic invariance, common information approach, and others.
- Explicit solutions are rare and typically exist for systems with two or three agents.

System with exchangeable agents

Dynamics $x_{t+1} = f_t(x_t, u_t, w_t)$ with per-step cost $c_t(x_t, u_t)$.

Pair of exchangeable agents

Agents i and j are exchangeable if

- $\mathcal{X}^i = \mathcal{X}^j$, $\mathcal{U}^i = \mathcal{U}^j$, $\mathcal{W}^i = \mathcal{W}^j$.
- $f_t(\sigma_{ij} x_t, \sigma_{ij} u_t, \sigma_{ij} w_t) = \sigma_{ij}(f_t(x_t, u_t, w_t))$
- $c_t(\sigma_{ij} x_t, \sigma_{ij} u_t) = c_t(x_t, u_t)$.

Set of exchangeable agents

A set of agents is exchangeable if every pair in that set is exchangeable

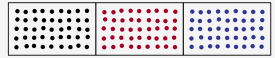

System with partially exchangeable agents

... is a multi-agent system where the set of agents can be partitioned into disjoint sets of exchangeable agents.

Linear quadratic system with partially exchangeable agents

Dynamics $x_{t+1} = A_t x_t + B_t u_t + w_t$

Cost $\sum_{t=1}^{T} \left[x_t^\top Q_t x_t + u_t^\top R_t u_t \right]$

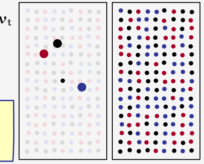

Irrespective of the information structure such a system is equivalent to a mean-field coupled system

Agent dynamics in sub-population k: $x_{t+1}^i = A_t^k x_t^i + B_t^k u_t^i + D_t^k \bar{x}_t + E_t^k \bar{u}_t + w_t^i$

Cost $\sum_{t=1}^{T} \left[\sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{N}^k} \frac{1}{|\mathcal{N}^k|} \left[(x_t^i)^\top Q_t^k x_t^i + (u_t^i)^\top R_t^k u_t^i \right] + \bar{x}_t^\top P_t^x \bar{x}_t + \bar{u}_t^\top P_t^u \bar{u}_t \right]$

A surprisingly simple solution ...

Parallel axis Theorem
$$\frac{1}{|\mathcal{N}^k|} \sum_{i \in \mathcal{N}^k} (x_t^i)^\top Q_t^k x_t^i = \frac{1}{|\mathcal{N}^k|} \sum_{i \in \mathcal{N}^k} (\breve{x}_t^i)^\top Q_t^k \breve{x}_t^i + (\bar{x}_t^k)^\top Q_t^k \bar{x}_t^k,$$
where $\breve{x}_t^i = x_t^i - \bar{x}_t^k$.

Decoupled Per-step cost
$$\sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{N}^k} \frac{1}{|\mathcal{N}^k|} [(\breve{x}_t^i)^\top Q_t^k \breve{x}_t^i] + \bar{x}_t^\top (\bar{Q}_t + P_t^x) \bar{x}_t$$
+ similar u-terms

Noise coupled Dynamics
$$\breve{x}_{t+1}^i = A_t^k \breve{x}_t^i + B_t^k \breve{u}_t^i + \breve{w}_t^i, \quad \bar{x}_{t+1} = A_t^k \bar{x}_t + B_t^k \bar{u}_t + \bar{w}_t$$

We still have a non-classical information structure

Assume centralized information and use certainty equivalence

	Local States	Mean-field state
Dynamics	$\breve{x}_{t+1}^i = A_t^k \breve{x}_t^i + B_t^k \breve{u}_t^i + \breve{w}_t^i$	$\bar{x}_{t+1} = A_t \bar{x}_t + B_t \bar{u}_t + \bar{w}_t$
Cost	$(\breve{x}_t^i)^\top Q_t^k \breve{x}_t^i + (\breve{u}_t^i)^\top R_t^k \breve{u}_t^i$	$(\bar{x}_t)^\top (P_t^x + Q_t) \bar{x}_t + (\bar{u}_t)^\top (P_t^u + R_t) \bar{u}_t$
Control Law	$\breve{u}_t^i = \breve{L}_t^k \breve{x}_t^i$	$\bar{u}_t = \bar{L}_t \bar{x}_t$
Gains	$\breve{L}_t^k = -(\cdots)^{-1} (B_t^k)^\top \breve{M}_{t+1}^k A_t^k$	$\bar{L}_t = -(\cdots)^{-1} (\bar{B}_t)^\top \bar{M}_{t+1} \bar{A}_t$
Riccati Equation	$\breve{M}_{1:T}^k = \mathrm{DRE}(A_{1:T}^k, B_{1:T}^k, Q_{1:T}^k, R_{1:T}^k)$	$\bar{M}_{1:T} = \mathrm{DRE}(\bar{A}_{1:T}, \bar{B}_{1:T}, \bar{Q}_{1:T} + P_{1:T}^x, \bar{R}_{1:T} + P_{1:T}^u)$

K equations, one for each sub-population

1 equation for all mean-fields

Solution generalizes to ...

Major-minor setup: One major agent and a population of minor agents.

Tracking cost function
$$\sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{N}^k} \frac{1}{|\mathcal{N}^k|} \left[(x_t^i - \mathring{x}_t^i)^\top Q_t^k (x_t^i - \mathring{x}_t^i) + (u_t^i)^\top R_t^k u_t^i \right] + (\bar{x}_t - r_t)^\top P_t^x (\bar{x}_t - r_t) + \bar{u}_t^\top P_t^u \bar{u}_t$$

Systems coupled through weighted mean-field
$$\bar{x}_t^k = \frac{1}{|\mathcal{N}^k|} \sum_{i \in \mathcal{N}^k} \lambda^i x_t^i, \qquad \bar{u}_t^k = \frac{1}{|\mathcal{N}^k|} \sum_{i \in \mathcal{N}^k} \lambda^i u_t^i.$$

Partial mean-field sharing information structure

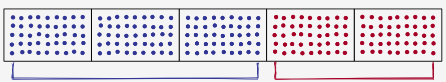

Set $\mathcal{S}$: MF observed

Set $\mathcal{S}^c$: MF not observed

Estimated mean-field $z_t = (z_t^1, \dots, z_t^K) = \mathbb{E}[\bar{s}_t \mid \{\bar{s}_t^k\}_{k \in \mathcal{N}}]$,

$$\text{where } z_{t+1}^k = \begin{cases} \bar{s}_{t+1}^k, & k \in \mathcal{S} \\ A_t^k z_t^k + (B_t^k \bar{L}_t^k + D_t^k + E_t^k \bar{L}_t) z_t, & k \notin \mathcal{S} \end{cases}$$

Certainty equivalence controller and its performance

Exact Performance
$$\hat{J} - J^* = \mathrm{Tr}(\tilde{X}_1 \tilde{M}_1) + \sum_{t=1}^{T-1} \mathrm{Tr}(\tilde{W}_t \tilde{M}_{t+1}) \text{ where } \tilde{M}_{1:T} = \mathrm{DLE}(\tilde{A}_{1:T}, \tilde{Q}_{1:T})$$

Performance bound

Let $n = \min_{k \notin \mathcal{S}} \{|\mathcal{N}^k|\}$. Suppose all noises are independent. Then, there exists a matrix C such that $\tilde{X}_1 \leqslant C/n$ and $\tilde{W}_t \leqslant C/n$. Thus,
$$\hat{J} - J^* \in \mathcal{O}\left(\frac{T}{n}\right),$$

Infinite horizon

Results extend to infinite horizon setup under standard assumptions. For both discounted and average cost setup:
$$\hat{J} - J^* \in \mathcal{O}\left(\frac{1}{n}\right), \qquad \text{c.f. } \mathcal{O}\left(\frac{1}{\sqrt{n}}\right) \text{ in MFG}$$

Everyone follows the optimal strategy