

DECENTRALIZED CONTROL OF MULTI-AGENT SYSTEMS USING LOCAL DENSITY FEEDBACK


ACC 2021

Shiba Biswal

UCLA


Joint work with K. Elamvazhuthi (UCLA) and S. Berman (ASU)

A SIMPLE GAME



- Distribute agents amongst two blocks
- Desired distribution $\mu^d = [0.5, 0.5]$.

A SIMPLE GAME




- Distribute agents amongst two blocks
- Desired distribution $\mu^d = [0.5, 0.5]$.
- Rules:
 - Each agent is identity free
 - Decentralized law, i.e. no central computer
 - No inter-agent communication

SOLUTION

- Solution: Make it **Stochastic!**


- $$P = \begin{bmatrix} p_{11} = H & p_{12} = T \\ p_{21} = H & p_{22} = T \end{bmatrix}$$



SOLUTION

- Solution: Make it **Stochastic!**

- $$P = \begin{bmatrix} p_{11} = H & p_{12} = T \\ p_{21} = H & p_{22} = T \end{bmatrix}$$




- Con: Agents keep transitioning!

SOLUTION

- Solution: Make it **Stochastic!**

- $$P = \begin{bmatrix} p_{11} = H & p_{12} = T \\ p_{21} = H & p_{22} = T \end{bmatrix}$$



- Con: Agents keep transitioning!

GIST OF OUR WORK

Control agent distribution when agents follow some Markov process

MEAN-FIELD MODEL 101

- Graph based methods popular in multi-agent systems.

MEAN-FIELD MODEL 101

- Graph based methods popular in multi-agent systems.
- *Alternative approach*: All agents follow same dynamics, independent of agent identities $\xrightarrow{N \rightarrow \infty}$ fluid approximation = mean-field model/Macroscopic model.

MEAN-FIELD MODEL 101

- Graph based methods popular in multi-agent systems.
- *Alternative approach*: All agents follow same dynamics, independent of agent identities $N \rightarrow \infty$ **fluid approximation = mean-field model/Macroscopic model**.



- Each agent follows an identical Markov process \implies mean-field behavior determined by the *Kolmogorov forward equation*.
- Model is independent of agent population size, therefore, scalable.


ALTERNATIVE VIEWPOINT

- Classical ODE (Deterministic) can be seen as evolution of Dirac delta function

$$\dot{x}(t) = f(x(t), u(t)).$$

- Markov process is evolution of distributions (pdfs):

$$\dot{\mu}(t) = F(\mu(t), k(t)).$$



KOLMOGOROV FORWARD EQUATION

$$\dot{\mu}(t) = F(\mu(t), k(t)). \quad (1)$$

- States (μ) are probability distributions, transition rates/probabilities as controls k

KOLMOGOROV FORWARD EQUATION

$$\dot{\mu}(t) = F(\mu(t), k(t)). \quad (1)$$

- States (μ) are probability distributions, **transition rates/probabilities** as controls k
- **Stabilizability**: System is stabilizable if given μ_d , there exists a map $k(\cdot)$ such that μ_d is asymptotically stable for (1)


KOLMOGOROV FORWARD EQUATION

$$\dot{\mu}(t) = F(\mu(t), k(t)). \quad (1)$$

- States (μ) are probability distributions, **transition rates/probabilities** as controls k
- **Stabilizability**: System is stabilizable if given μ_d , there exists a map $k(\cdot)$ such that μ_d is asymptotically stable for (1)
- Challenge: $\mu \in \Sigma$ - infinite dimensional.

POTENTIAL APPLICATIONS

Redistributing large number of agents for environmental monitoring, surveillance, autonomous construction.



Harvard University/Michael Rubenstein

GOALS

- **Stabilizability** - Existence of control law(s) k to stabilize desired distributions.
- **Equilibrium of the Microscopic model** - Zero state transitions at equilibrium to minimize energy expended by agents at equilibrium.

GOALS

- **Stabilizability** - Existence of control law(s) k to stabilize desired distributions.
- **Equilibrium of the Microscopic model** - Zero state transitions at equilibrium to minimize energy expended by agents at equilibrium.

MODEL

Discrete-Time, Continuous State-Space Markov Process (\mathbb{R}^n or some manifold)

Ideal for modeling robots as robots don't evolve on graphs

PROBLEM FORMULATION

AGENT MODEL

$$x_{n+1} = F(x_n, u_n), \quad x_0 \in \Omega$$


- $x_n \in \Omega$, $u_n \in U$, $F : \Omega \times U \rightarrow \mathbb{R}^d$
- $(u_n)_{n=1}^{\infty} \in U$ such that $F(x_n, u_n) \in \Omega$

PROBLEM FORMULATION

AGENT MODEL

$$x_{n+1} = F(x_n, u_n), \quad x_0 \in \Omega$$

- $x_n \in \Omega$, $u_n \in U$, $F : \Omega \times U \rightarrow \mathbb{R}^d$
- $(u_n)_{n=1}^{\infty} \in U$ such that $F(x_n, u_n) \in \Omega$



ASSUMPTIONS

- Domain $\Omega \subset \mathbb{R}^d$ is closed, bounded, connected, $\delta\Omega$ is 'regular'
- Controls $U \subset \mathbb{R}^d$ is closed, bounded (compact)
- F is continuous, C^1 , non-singular
- **Locally controllable condition**: there exists $r > 0$ such that, for every $x \in \Omega$, $B_r(x) \cap \Omega \subseteq F(x, U)$.

PROBLEM FORMULATION

- N - number of '*identity free*' agents
- The empirical distribution of the N agents is $\frac{1}{N} \sum_{k=1}^N \delta_{x_n^k}$

PROBLEM FORMULATION

- N - number of '*identity free*' agents
- The empirical distribution of the N agents is $\frac{1}{N} \sum_{k=1}^N \delta_{x_n^k}$
- Goal:

$$\frac{1}{N} \sum_{k=1}^N \delta_{x_0^k} \xrightarrow{n \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \delta_{x^{k,d}} \approx f_d \in L^\infty(\Omega)$$

PROBLEM FORMULATION

- N - number of 'identity free' agents
- The empirical distribution of the N agents is $\frac{1}{N} \sum_{k=1}^N \delta_{x_n^k}$
- Goal:

$$\frac{1}{N} \sum_{k=1}^N \delta_{x_0^k} \xrightarrow{n \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \delta_{x^{k,d}} \approx f_d \in L^\infty(\Omega)$$

- $\frac{1}{N} \sum_{k=1}^N \delta_{x_n^k}$ not a state variable of the agent dynamics; in order to make it a state variable:

PROBLEM FORMULATION

- N - number of 'identity free' agents
- The empirical distribution of the N agents is $\frac{1}{N} \sum_{k=1}^N \delta_{x_n^k}$
- Goal:

$$\frac{1}{N} \sum_{k=1}^N \delta_{x_0^k} \xrightarrow{n \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \delta_{x^{k,d}} \approx f_d \in L^\infty(\Omega)$$

- $\frac{1}{N} \sum_{k=1}^N \delta_{x_n^k}$ not a state variable of the agent dynamics; in order to make it a state variable:
- Take $N \rightarrow \infty$

PROBLEM FORMULATION I

FORWARD EQUATION = MEAN FIELD MODEL:

$$\mu_{n+1} = P\mu_n, \quad \mu_0 \in \mathcal{P}(\Omega) = \text{Probability measures on } \Omega$$

PROBLEM FORMULATION I

FORWARD EQUATION = MEAN FIELD MODEL:

$$\mu_{n+1} = P\mu_n, \quad \mu_0 \in \mathcal{P}(\Omega) = \text{Probability measures on } \Omega$$

PROBLEM

Given $\mu_d \in \mathcal{P}(\Omega)$ and F can we construct an operator P such that

$$\lim_{n \rightarrow \infty} P^n \mu_0 = \mu_d?$$

such that μ_d is exponentially stable equilibrium.

PROBLEM FORMULATION I

FORWARD EQUATION = MEAN FIELD MODEL:

$$\mu_{n+1} = P\mu_n, \quad \mu_0 \in \mathcal{P}(\Omega) = \text{Probability measures on } \Omega$$

PROBLEM

Given $\mu_d \in \mathcal{P}(\Omega)$ and F can we construct an operator P such that

$$\lim_{n \rightarrow \infty} P^n \mu_0 = \mu_d?$$

such that μ_d is **exponentially stable** equilibrium.

THEOREM

*There exists a control law that stabilizes a **class** of μ_d !*

Any μ_d with density $f_d, f_d^{-1} \in L^\infty(\Omega) = \{g : g \text{ is bounded a.e.}\}$.

SOLUTION

- **Transition kernel** \sim Transition Probability in discrete state space
 $K(x, W)$ - From x , the probability of choosing a set of controls W

SOLUTION

- **Transition kernel** \sim Transition Probability in discrete state space
 $K(x, W)$ - From x , the probability of choosing a set of controls W
- Proved existence of K such that $f_d \in L^\infty(\Omega)$ can be stabilized

SOLUTION

- **Transition kernel** \sim Transition Probability in discrete state space
 $K(x, W)$ - From x , the probability of choosing a set of controls W
- Proved existence of K such that $f_d \in L^\infty(\Omega)$ can be stabilized
- For the discretized model, constructed optimal K to drive $\|f_0 - f_d\|_2 \rightarrow 0$ exponentially
- Open-Loop Problem for the mean-filed model, state feedback for the agent system

SOLUTION

- **Transition kernel** \sim Transition Probability in discrete state space
 $K(x, W)$ - From x , the probability of choosing a set of controls W
- Proved existence of K such that $f_d \in L^\infty(\Omega)$ can be stabilized
- For the discretized model, constructed optimal K to drive $\|f_0 - f_d\|_2 \rightarrow 0$ exponentially
- Open-Loop Problem for the mean-field model, state feedback for the agent system

PROBLEM FORMULATION II

PROBLEM II

Given $\mu_d \in \mathcal{P}(\Omega)$ and F can we construct an operator P such that

$$\lim_{n \rightarrow \infty} P^n \mu_0 = \mu_d, \quad \mu_0 \in \mathcal{P}(\Omega)$$

$P = I$ at $\mu_d \implies$ agents stop switching at μ_d

FORWARD EQUATION = MEAN FIELD MODEL:

$$\mu_{n+1} = P(\mu_n) \mu_n, \quad \mu_0 \in \mathcal{P}(\Omega) = \text{Probability measures on } \Omega$$

P is now nonlinear, function of the current distribution μ_n

THEOREM

There exists a *time-dependent* control law that satisfies above conditions.

Any μ_d with density $f_d \in L^\infty(\Omega) = \{g : g(x) < \infty \text{ a.e.}\}$

STOCHASTIC FEEDBACK LAW

Let $k : \Omega \times U \rightarrow [0, 1]$ be in $L^\infty(\Omega \times U)$ and

$$k(x, u) \begin{cases} > 0 \text{ for } m\text{-a.e. } x \in \Omega, u \in U \text{ st. } F(x, u) \in \Omega; \\ = 0 \text{ otherwise;} \end{cases}$$

$$\int_U k(x, u) du = 1 \text{ for } m\text{-a.e. } x \in \Omega.$$

STOCHASTIC FEEDBACK LAW

Let $k : \Omega \times U \rightarrow [0, 1]$ be in $L^\infty(\Omega \times U)$ and

$$k(x, u) \begin{cases} > 0 \text{ for } m\text{-a.e. } x \in \Omega, u \in U \text{ st. } F(x, u) \in \Omega; \\ = 0 \text{ otherwise;} \end{cases}$$

$$\int_U k(x, u) du = 1 \text{ for } m\text{-a.e. } x \in \Omega.$$

Assuming $F(x, 0) = x$ (can be generalized to $F(x, V(x)) = x$).

$$K_\mu(x, W) = a_{f_\mu}(x) \int_W k(x, u) du + (1 - a_{f_\mu}(x)) \delta_0(W).$$

$$a_f(x) = \begin{cases} \frac{f(x) - f_d(x)}{f(x)} & \text{for } m\text{-a.e. } x \text{ if } f(x) - f_d(x) > 0; \\ 0 & \text{otherwise.} \end{cases}$$

STOCHASTIC FEEDBACK LAW

Let $k : \Omega \times U \rightarrow [0, 1]$ be in $L^\infty(\Omega \times U)$ and

$$k(x, u) \begin{cases} > 0 \text{ for } m\text{-a.e. } x \in \Omega, u \in U \text{ st. } F(x, u) \in \Omega; \\ = 0 \text{ otherwise;} \end{cases}$$

$$\int_U k(x, u) du = 1 \text{ for } m\text{-a.e. } x \in \Omega.$$

Assuming $F(x, 0) = x$ (can be generalized to $F(x, V(x)) = x$).

$$K_\mu(x, W) = a_{f_\mu}(x) \int_W k(x, u) du + (1 - a_{f_\mu}(x)) \delta_0(W).$$

$$a_f(x) = \begin{cases} \frac{f(x) - f_d(x)}{f(x)} & \text{for } m\text{-a.e. } x \text{ if } f(x) - f_d(x) > 0; \\ 0 & \text{otherwise.} \end{cases}$$

Closed-Loop control for the Mean-Field model.

OPERATOR INDUCED BY K_μ

Define P via $K : \Omega \times \mathcal{B}(U) \rightarrow [0, 1]$.

DEFINITION (FORWARD OPERATOR (MEASURES))

$P : \mathcal{P}(\Omega) \rightarrow \mathcal{P}(\Omega)$

$$(P_\mu\mu)(E) = \int_\Omega \int_U \chi_E(F(x, u)) K_\mu(x, du) d\mu(x)$$

Note: $\chi_E(z) \begin{cases} = 1 & \text{if } z \in E \\ = 0 & \text{otherwise} \end{cases}$

SOLUTION


THEOREM

If P preserves distributions with L^2 densities then f_d is *globally asymptotically stable* in the $L^1(\Omega, m)$ norm, i.e.

$$\|f_n - f_d\|_1 \rightarrow 0 \text{ as } n \rightarrow \infty.$$

N-AGENT SYSTEM

- N agents on Ω evolving according to K .
- Empirical measure $m^N(x) = \frac{1}{N} \sum_{k=1}^N \delta_{x^k} \implies$ Does not have a density
- Smoothen the measure by convolving with a Mollifier.



MOLLIFIER

Standard bump function:


$$\phi(x) = \begin{cases} e^{-\left(\frac{1}{1-\|x\|^2}\right)}, & x \in (-1, 1), \\ 0, & \text{otherwise.} \end{cases}$$

Change the support, for some $h > 0$

$$\phi_h(x) = h^{-2} \phi\left(\frac{x}{h}\right).$$

$\int \phi_h = 1$ for any h .

(Gif)




MOLLIFIER

CONVOLUTION

Convolve the Dirac measure with a smooth (C^∞) function ϕ :

$$\phi * m^N = \int_{\Omega} \phi(x) dm^N = \frac{1}{N} \sum_{i=1}^N \phi(x - x^k).$$



SIMULATIONS: MEAN FIELD MODEL & $N = 100$

SIMULATIONS: MEAN FIELD MODEL & $N = 500$

SIMULATIONS: MEAN FIELD MODEL & $N = 1500$

AGENT TRAJECTORIES




FIGURE: $N = 100$




FIGURE: $N = 500$




FIGURE: $N = 1000$

THANK YOU

QUESTIONS?

Shiba Biswal, Karthik Elamvazhuthi, and Spring Berman. “Target Distribution Stabilization Using Local Density Feedback for Multi-Agent Systems.”, 2022. Accepted to *IEEE Transactions on Automatic Control*.

Shiba Biswal, Karthik Elamvazhuthi, and Spring Berman. “Stabilization of nonlinear discrete-time systems to target measures using stochastic feedback laws”, 2021. *IEEE Transactions on Automatic Control*.

Shiba Biswal, Karthik Elamvazhuthi, and Spring Berman. “Stabilization of Multi-Agent Systems to Target Distributions using Local Interactions.” Submitted to the 2020 *International Symposium on Mathematical Theory of Networks and Systems (MTNS)*, Cambridge, UK.

Shiba Biswal, Karthik Elamvazhuthi, and Spring Berman. “Fastest Mixing Markov Chain on a Compact Manifold.” *IEEE Conference on Decision and Control (CDC)*, Nice, France, 2019.