Consensus and Disagreement in Collective Homing Problems: A Mean Field Games Formulation

Rabih Salhab, Roland P. Malhame and Jerome Le Ny

October 8, 2015

Introduction

Goal:
To model biological collective decision mechanisms. Examples:

- Honey bees searching for a new colony
- Collective navigation in fish schools

The two important properties that characterize such systems are

- Aggregation of agents
- Decentralized control

Table of Contents

Table of Contents

The model: the dynamics

N agents live in $\mathbb{R}^n$ over a time interval $[0, T]$.

They follow identical and independent linear dynamics, for $i = 1, \dots N$:

$$\begin{cases} \dot{x}_i = Ax_i + Bu_i \\ x_i\big|_{t=0} = x_i^0 \end{cases} \tag{1}$$

where $x_i^0 \in \mathbb{R}^n$, $u_i \in \mathbb{R}^m$.

The model: the costs

The agents want to finish at one of two destinations $p_a, p_b \in \mathbb{R}^n$ while not using too much effort. Agent i wants to minimize the cost functional

$$J_i(u_i; \bar{x}, x_i^0) = \int_0^T \frac{q}{2}\|x_i - \bar{x}\|^2 + \frac{r}{2}\|u_i\|^2 \, \mathrm{d}t + \frac{M}{2} \min\left(\|x_i|_{t=T} - p_a\|^2, \|x_i|_{t=T} - p_b\|^2\right) \quad (2)$$

where $q, r, M > 0$ and M is large compared to the other stuff.

Note:
The agents are cost coupled.

The model: strategy towards solution

How do we find/approximate a Nash equilibrium to this problem?

Strategy proposed in the paper:

- (Decentralization) Introduce a path x^* that all agents respond to instead of the mean $\bar{x}$.
- (Consistency) Choose x^* so that it will be replicated by $\bar{x}$ when players are optimally tracking x^*.
- (Optimization) Find the optimal control u_i^* while tracking x^*
- Apply u^* to the original problem and get an $\epsilon-$Nash equilibrium.

Table of Contents

Tracking problem: decentralized problem

Assume that x^* is a continuous path and that agent i solves the control problem

$$\begin{cases} \inf_{u_i} J_i(u_i; x^*, x_i^0) = \min\left(\inf_{u_i} J_i^a(u_i, x^*, x_i^0), \inf_{u_i} J_i^b(u_i, x^*, x_i^0)\right) \\ dx_i = Ax_i + Bu_i \\ x_i\big|_{t=0} = x_i^0 \end{cases} \tag{P}$$

where, for $e \in \{a, b\}$,

$$J_i^e(u_i; x^*, x_i^0) = \int_0^T \frac{q}{2}\|x_i - x^*\|^2 + \frac{r}{2}\|u_i\|^2 \; dt + \frac{M}{2}\|x_i(T) - p_e\|$$

Tracking problem: the optimal control 1

The optimal control is

$$u_i^* = \begin{cases} u_i^a & \text{if } J_i^a(u_i^a, x^*, x_i^0) \leq J_i^b(u_i^b, x^*, x_i^0) \\ u_i^b & \text{if } J_i^a(u_i^a, x^*, x_i^0) > J_i^b(u_i^b, x^*, x_i^0) \end{cases}$$

Note:
The choice of either u_i^a or u_i^b is made at time 0 and kept throughout the game. The choice depends only on the initial position x_i^0 and the parameter values.

Tracking problem: the optimal control 2

Solving the optimal control yields

$$u_i^e(t) = -\frac{1}{r} B^\mathsf{T} \left(\alpha(t) x_i + \beta^e(t)\right), \quad e \in \{a, b\}$$

with corresponding optimal cost

$$J_i^e(u_i^e, x^*, x_i^0) = \frac{1}{2}(x_i^0)^\mathsf{T} \alpha(0) x_i^0 + \beta^e(0)^\mathsf{T} x_i^0 + \delta^e(0)$$

where α, β, δ are solutions to three coupled ODEs.

Tracking problem: the optimal control 3

Lemma:

The tracking problem (P) has a unique optimal control

$$u_i^*(t) = \begin{cases} -\frac{1}{r} B^{\mathsf{T}} \left(\alpha(t) x_i + \beta^a(t) \right) & \text{if } x_i^0 \in D_a(x^*) \\ -\frac{1}{r} B^{\mathsf{T}} \left(\alpha(t) x_i + \beta^b(t) \right) & \text{if } x_i^0 \notin D_a(x^*) \end{cases}$$

where $\alpha, \beta^e, \delta^e$ are the unique solutions to the coupled ODEs for $e = a, b$ and

$$D_a(x^*) = \left\{ x \in \mathbb{R}^n \,\middle|\, \left(\beta^a(0) - \beta^b(0) \right)^{\mathsf{T}} x \leq \delta^b(0) - \delta^a(0) \right\}$$

Note:

Given any x^*, there exists a basin of attraction $D_a(x^*)$ such that all players initially present in this region will go to p_a.

Table of Contents

Replicated by the mean: the mean's dynamics

For consistency we require $x^* = \bar{x}$. The dynamics of $\bar{x}$ when all players are optimaly tracking x^* is

$$\dot{\bar{x}}(t) = -K^{\mathsf{T}}\bar{x} - \frac{q}{r}BB^{\mathsf{T}} \int_T^t \phi_K(t,\sigma)x^*(\sigma)\, d\sigma + \frac{M}{r}BB^{\mathsf{T}}\phi_K(t,T)p_\lambda \tag{3}$$

<u>Note:</u>
All data in (3) is given except the tracked path x^* and λ, the number of players initially in $D_a(x^*)$.

For any $\lambda \in \{0, \ldots, N\}$ let $T_\lambda : C([0,T], \mathbb{R}^n) \to C([0,T], \mathbb{R}^n)$ where $T_\lambda(x^*)$ is the unique solution of (3), that is

$$\bar{x} = T_\lambda(x^*), \ \lambda \text{ players initially in } D_a(x^*)$$

Replicated the mean: fixed point of T_λ

Lemma: Let $\lambda \in \{0, \dots, N\}$. The map T_λ has a unique fixed point equal to

$$R_1(t)\bar{x}_0 + R_2(t)p_\lambda.$$

Let the players be ordered according to their initial positions,

$$\beta_0^\mathsf{T} x_1^0 \leq \cdots \leq \beta_0^\mathsf{T} x_N^0$$

Theorem (replication): A path x^* that is replicated by the mean when all players are optimally tracking it exists if and only if there exists $\lambda \in \{0, \dots, N\}$ such that

$$\beta_0^\mathsf{T} x_\lambda^0 - \delta_0 - \theta_1 \leq \lambda\theta_2 < \beta_0^\mathsf{T} x_{\lambda+1}^0 - \delta_0 - \theta_1. \tag{Ineq}$$

In this case x^* is the unique fixed point of T_λ.

Table of Contents

Nash equilibria: u^* gives ϵ-Nash

Assume that $\|x_i^0\|$ is uniformly bounded from above.

Theorem (existence):
Assume that $\lambda \in \{0, \ldots, N\}$ that satisfies (Ineq). Let Σ be the set of decentralized controls that generates a fixed point of T_λ. Then Σ is an ϵ-Nash equilibrium with respect to the costs $J_i(u_i, \frac{1}{N}\sum_{j=1}^N x_j(u_j), x_i^0)$ where $\epsilon = o(1/N)$.

Theorem (uniqueness):
Assume that exists N_0 such that if $N \geq N_0$ then

$$\max_\lambda \|x_{\lambda+1}^0 - x_\lambda^0\| \leq k\frac{1}{N},$$

where k is independent of N. Then for all $N \geq N_0$ there exists at most one ϵ-Nash equilibrium.

Summary of problem solution

Given data: $(x_1^0, \ldots, x_N^0)$, A, B, q, r, M, p_a, p_b

1. Find λ that solves (Ineq)
2. By replication theorem, λ gives x^* that is a fixed point for T_λ
3. For $i = 1, \ldots, N$, find u_i^* that solves (P) given x^* and x_i^0
4. By replication theorem, these u_i^* gives a mean agent trajectory that replicates x^*.
5. By existence theorem, these u_i^* constitute an ϵ-Nash equilibrium to the original problem.

Table of Contents

Simulations: setup

Simulations from the paper will be presented.

Input data:

$$n = m = 2,\ N = 20,\ A = B = I_2,$$
$$p_a = -p_b = (-10, 0),\ \ T = 1$$
$$(x_1^0, \ldots, x_N^0) \text{ given for each case}$$

Cases 1&2: $q = r = 1,\ M = 10000$

Case 3: $q = 10,\ r = 1,\ M = 1000$

Simulations: Case 1

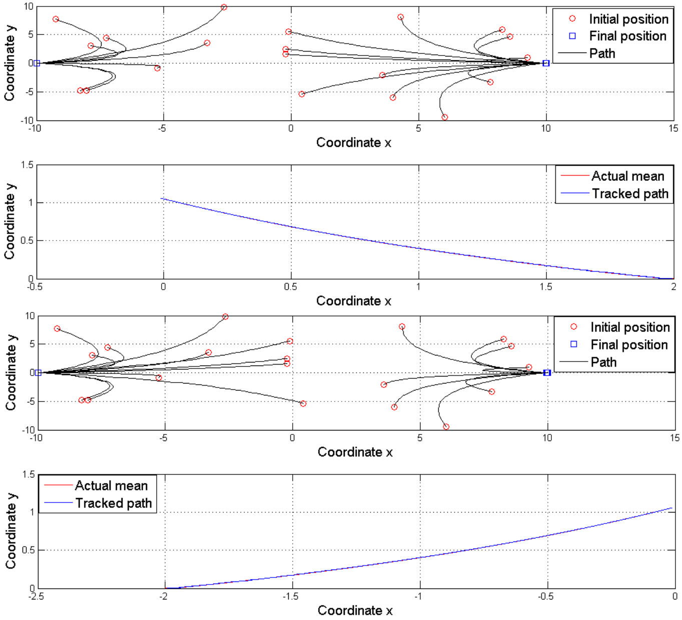

Simulations: Case 2

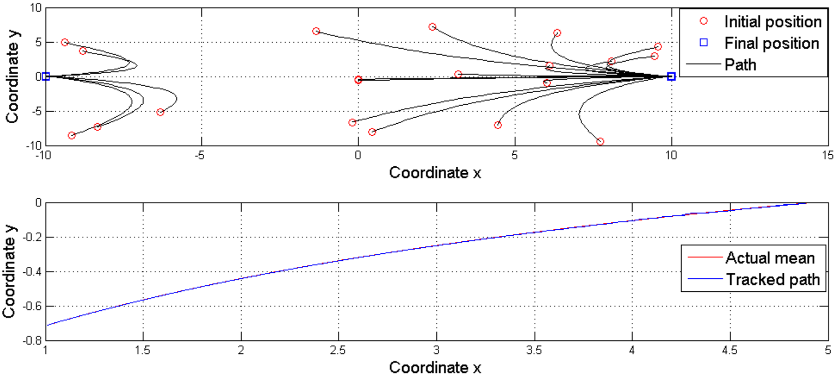

Simulations: Case 3

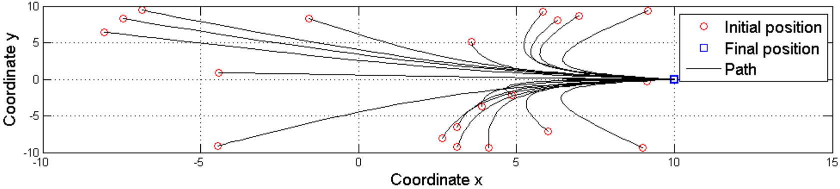

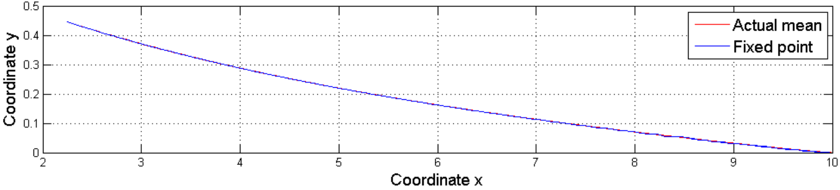

What then?

What will happen if we introduce

- only statistical knowledge of the initial positions?
- noise in the dynamics/costs?

What then?

What will happen if we introduce

- only statistical knowledge of the initial positions?
- noise in the dynamics/costs?

Other ideas:

- $N \to \infty$?
- Moving destinations?
- When can we have decentralization?

Thank you!