Long time average of mean field games

Alessio Porretta
Universita' di Roma Tor Vergata

Equadiff '13, Prague 29/08/2013

Outlines of the talk

- Brief description of the Mean Field Games model system.
  Coupling viscous Hamilton-Jacobi & Fokker-Planck.

- "Long time behavior" of Mean Field Games: natural questions and setting. The ergodic problem and expected behavior.

- Main results obtained:

  (i) long time average convergence: a matter of energy estimates

  (ii) exponential rate of convergence

  (joint works with P. Cardaliaguet, J-M. Lasry, P-L. Lions)

- Links with optimal control problems in the long horizon: a general *turnpike behavior*.

  (joint work E. Zuazua)

On the Mean Field Games theory

The Mean Field Games model was introduced by J.-M. Lasry and P.-L. Lions [CRAS '06, Cours Collège de France since 2006] and independently by M. Huang-P. Caines-R. Malhamé.

Main goal: *describe games with large numbers (a continuum) of agents whose strategies depend on the distribution of the agents.*

Typical features of the model:

- players act according to the same principles (they are indistinguishable and have the same optimization criteria).

- players have individually a minor (infinitesimal) influence, but their strategy takes into account the mass of co-players.

Roughly: players are particles but have strategies

Goal: introduce a macroscopic description through a mean field approach as the number of players $N \to \infty$.

The simplest form of the continuum limit is a coupled system of PDEs

$$\begin{cases} (1) \qquad -u_t - \Delta u + H(x, Du) = F(x, m) & \text{in } (0,T)\times\Omega \\ (2) \qquad m_t - \Delta m - \mathrm{div}(m\, H_p(x, Du)) = 0 & \text{in } (0,T)\times\Omega\,, \end{cases}$$

- (1) is the Bellman equation for the agents' value function u.
- (2) is the Kolmogorov-Fokker-Planck equation for the state of the agents. $m(t)$ is the probability density of the state of players at time t.

Roughly, each agent (infinitesimal) controls the dynamics

$$dX_t = \alpha_t\, dt + \sqrt{2}\, dB_t$$

where B_t is a d-dimensional Brownian motion, in order to minimize, among controls α_t, some cost

$$J(\alpha) := \mathbb{E}\left[\int_0^T [L(X_s, \alpha_s) + F(X_s, m(s, X_s))\; ds + u_T(X_T)\right]$$

where L is the Legendre transform of H and u_T a final pay-off.

If u solves the Bellman equation it gives the best value:

- $\inf_\alpha J(\alpha) = \int u(x,0)dm_0(x)$,

where m_0 is the probability distribution of X_0.

- the optimal control is given by the feedback law:

$\alpha^*_t = -H_p(X_t, Du(t,X_t))\,, \qquad H_p := \frac{\partial H(x,p)}{\partial p}$.

In turn, if

$$dX_t = \alpha(X_t)dt + \sqrt{2}dB_t$$

the probability measure $m(t)$ (distribution law of X_t) satisfies

$$m_t - \Delta m + \text{ div } (\alpha\, m) = 0$$

Hence, the evolution of the state of the agents is governed by their optimal decisions α^*_t, and m satisfies

$$m_t - \Delta m - \text{ div } (m\, H_p(x, Du)) = 0$$

This is the Mean Field Games system (with horizon T):

$$\begin{cases} (1) \qquad -u_t - \Delta u + H(x, Du) = F(x, m) & \text{in } (0,T) \times \Omega \\ (2) \qquad m_t - \Delta m - \operatorname{div}(m\, H_p(x, Du)) = 0 & \text{in } (0,T) \times \Omega\,, \end{cases}$$

usually complemented with initial-terminal conditions:

-$m(0) = m_0$ (initial distribution of the agents)

-$u(T) = u_T$ (final pay-off)

$+$ boundary conditions (here for simplicity assume periodic b.c.)

Main novelties are:

- the backward-forward structure.
- the interaction in the strategy process: the coupling $F(x, m)$
  Two coupling regimes are usually considered:

  (i) Nonlocal coupling with smoothing effect (ex. convolution): $F : \mathbb{R}^N \times \mathcal{P}_1 \to \mathbb{R}$ is smoothing on the space of probability measures. Ex: $F(x, m) = \Phi(x, k \star m)$

  (ii) Local coupling: $F = F(x, m(t, x))$.
  (regularity of sol.'s is a big issue)

Long time behavior

Pb: What is the behavior of the MFG system when the horizon $T \to \infty$?

$$\begin{cases} -u^T_t - \Delta u^T + H(x, Du^T) = F(x, m^T)\,, & \text{in } (0,T) \\ m^T_t - \Delta m^T - \operatorname{div}\,(m^T\, H_p(x, Du^T)) = 0\,, & \text{in } (0,T) \\ m^T(x,0) = m_0(x)\,, \quad u^T(x,T) = u_T\,. \end{cases}$$

- To fix the ideas, we work in the periodic setting.

To simplify the presentation, I will consider a reference case:
$H(x, Du) = \frac{1}{2}|Du|^2$, initial data m_0, u_T smooth, $m_0 > 0$.

Long time behavior is a very natural question in the viewpoint of SDE.
In the long horizon, agents are expected to behave in a way to minimize the average (ergodic) cost, regardless of the initial distribution.

Recall the case of a single equation (with no coupling):
(see e.g. [Bensoussan-Frehse], [Namah-Roquejoffre], [Barles-Souganidis])

$$\begin{cases} -u_t^T - \Delta u^T + \frac{1}{2}|Du^T|^2 = F(x) & \text{in } (0,T) \\ u^T(x,T) = G(x) \end{cases}$$

(i) $\frac{u^T(x,0)}{T}$ converges uniformly to a constant $\bar{\lambda} \in \mathbb{R}$, which is the ergodic (minimal) cost

$$\bar{\lambda} = \inf_{\alpha} \lim_{T\to\infty} \frac{1}{T}\mathbb{E}\left\{\int_0^T \frac{1}{2}|\alpha(X_s)|^2 + F(X_s)]ds\right\}$$

(ii) $u^T(x,0) - \bar{\lambda}T \to \bar{u}$, periodic solution of the "ergodic problem"

$$\bar{\lambda} - \Delta\bar{u} + \frac{1}{2}|D\bar{u}|^2 = F(x)\,.$$

• If $Du^T \to D\bar{u}$, then m^T converges to the unique invariant measure $\bar{m}$ associated to the process

$$dX_t = -D\bar{u}(X_t)dt + \sqrt{2}dB_t$$

What happens for Mean Field Games ?

Good news: if the coupling $F(x,\cdot)$ is monotone, then the ergodic problem is well posed ([Lasry-Lions '07]).

There exists a unique couple $(\bar{u},\bar{m})$ and a unique constant $\bar{\lambda}$ which solve

$$\begin{cases} \bar{\lambda} - \Delta \bar{u} + \frac{1}{2}|D\bar{u}|^2 = F(x,\bar{m})\,, & \int_\Omega \bar{u} = 0 \\ -\Delta \bar{m} - \operatorname{div}(\bar{m} D\bar{u}) = 0\,, & \int_\Omega \bar{m} = 1 \end{cases}$$

Moreover, $\bar{u}$, $\bar{m}$ are smooth, $\bar{m} > 0$

Expected long time behavior: $u^T/T \to \bar{\lambda}$ and $m^T \to \bar{m}$.

However:

- one can not use the arguments of the single equation: there are no standard/simple comparison arguments, gradient estimates, etc...

- forward-backward conditions: there is not just an evolution forward in time! Some boundary layer could appear at $t = 0$ or $t = T$.

$\to$ stability will appear in a large transient time $[\delta T, (1-\delta)T]$

The kind of results which we prove [Cardaliaguet-Lasry-Lions-P.] :

- (ergodic behavior) $\frac{u^T(x,0)}{T} \to \bar{\lambda}$
- (Du^T, m^T) is close *in average* to $(D\bar{u}, \bar{m})$:

$$\frac{1}{T}\int_0^T\int_\Omega |Du - D\bar{u}|^2 + (F(x,m) - F(x,\bar{m}))(m - \bar{m})\,dx \to 0$$

Eventually, under some stronger assumption we also get:

- (Du^T, m^T) are exponentially close to $(D\bar{u}, \bar{m})$ in the transient time:

$$\|Du^T(t) - D\bar{u}\| + \|m^T(t) - \bar{m}\| \leq C\left(e^{-\kappa(T-t)} + e^{-\kappa t}\right) \quad \forall t \in (a, T-a)\,,$$

where a, C may depend on initial-terminal conditions.

The norms of the above convergences may vary according to local/nonlocal coupling.

Rmk: This is a *turnpike result*, in the terminology of math. economics, since the work of Nobel Price P. Samuelson in 1949:
an efficient expanding economy should for most of the time be nearly an equilibrium path

Convergence in average

Main ingredient: energy equality [Lasry-Lions] $\rightarrow$ uniqueness, stability of the system when $F(x, \cdot)$ monotone.

Any couple of solutions (u_1, m_1) and (u_2, m_2) satisfy

$$-\frac{d}{dt}\int_\Omega (u_1 - u_2)(m_1 - m_2)dx =$$

$$\int_\Omega \frac{(m_1 + m_2)}{2}|Du_1 - Du_2|^2 + (F(x, m_1) - F(x, m_2))(m_1 - m_2)\ dx$$

Apply the energy equality to (u, m) and $(\bar{u}, \bar{m})$ between 0 and T:

$$\int_0^T \int_\Omega \frac{(m + \bar{m})}{2}|Du - D\bar{u}|^2 + (F(x, m) - F(x, \bar{m}))(m - \bar{m})\, dx$$
$$= -\left[\int_\Omega (u - \bar{u})(m - \bar{m})dx\right]_0^T \quad \overset{?}{\leq} \quad C$$

Bounds at $t = 0$ and $t = T \Rightarrow$ convergence in average.

Main point: obtain an estimate $\|Du^T(0)\|$ independent of T

Typically, $u^T(0) \sim C\,T$. However, if we set $\langle u \rangle := \int u\,dx$, we have

$$\int_\Omega u^T(0)(m_0 - \bar{m})dx = \int_\Omega (u^T(0) - \langle u^T(0) \rangle)(m_0 - \bar{m})dx \leq c\|Du^T(0)\|_{L^2}$$

$\Rightarrow$ it is enough to bound $\|Du^T(0)\|$ independently of T.

We get estimates differently according to local or nonlocal coupling.

(i) Smoothing coupling $F(x,m)$: we use a (uniform in time) semiconcavity estimate $\Rightarrow$ Lipschitz bound for u^T.

(ii) Local coupling $F(x,m)$: we use the property that the system has an Hamiltonian structure $\Rightarrow$ there exists an invariant (constant in time)

$$\mathcal{E}(u,m) = \int_\Omega [\frac{1}{2} m|Du|^2 + Du \cdot Dm - \mathcal{F}(x,m)]dx$$

Thanks to this fact, we obtain a bound on $\|Du^T(0)\|_{L^2}$.

Exponential rate of stability

The exponential rate of convergence may come from two possible ingredients:

- stronger coercivity of the coupling $F(x,\cdot)$
- stability of the linearized pb. (if sol.'s are smooth !)

1. Local coupling $F(x,m)$.

We strengthen the monotonicity condition

$$(F(x,m_1)-F(x,m_2))(m_1-m_2) \geq \gamma(m_1-m_2)^2 \tag{1}$$

Theorem

Under assumption (1), there is some $\kappa > 0$ (independent of T) such that (we denote $\tilde{u} = u - \langle u \rangle$)

$$\|\tilde{u}(t)-\bar{u}\|_{L^1} + \|m(t)-\bar{m}\|_{L^1} \leq C\left(e^{-\kappa(T-t)} + e^{-\kappa t}\right) \quad \forall t \in (1, T-1)\,,$$

2. Nonlocal smoothing coupling.

We strengthen the regularizing property of the coupling term

$$\|F(x,m_1) - F(x,m_2)\|_{\mathcal{C}^{1+\alpha}} \le \bar{C}\|m_1 - m_2\|_{H^{-1}} \qquad \forall m_1, m_2 \qquad (2)$$

for some $\alpha > 0$.

Theorem

Under assumption (2), there exists $\kappa > 0$ (independent of T) such that

$$\|\tilde{u}(t) - \bar{u}\|_{C^{3,\alpha}} \le C\left(e^{-\kappa(T-t)} + e^{-\kappa t}\right) \qquad \forall t \in (a, T-a),$$

$$\|m(t) - \bar{m}\|_{C^{2,\alpha}} \le C\left(e^{-\kappa(T-t)} + e^{-\kappa t}\right) \qquad \forall t \in (a, T-a)$$

(C, a depend on initial-terminal conditions).

- Look at the linearized system around $(\bar{m}, \bar{u})$:

$$\begin{cases} -v_t \overbrace{-\Delta v + D\bar{u}Dv}^{Av} = F'(\bar{m})\mu \\ \mu_t \underbrace{-\Delta\mu - \operatorname{div}(\mu\, D\bar{u})}_{A^*\mu} = \underbrace{\operatorname{div}(\bar{m}\, Dv)}_{-Kv} \end{cases}$$

We show that $w := K^{-\frac{1}{2}}\mu$ satisfies

$$\frac{d^2}{dt^2}\|w(t)\|_2^2 \geq \omega_0^2\ \|w(t)\|_2^2$$

for some $\omega_0 > 0$.

$$\|w(t)\|_2^2 \leq \max\{\|w(0)\|_2^2, \|w(T)\|_2^2\}\left(e^{-\omega_0 t} + e^{-\omega_0(T-t)}\right).$$

$$\Rightarrow \|\mu(t)\|_{H^{-1}}^2 \lesssim \|w(t)\|_2^2 \leq \max\{\|m_0\|_2^2, \|m(T)\|_2^2\}\left(e^{-\omega_0 t} + e^{-\omega_0(T-t)}\right)$$

- Through a fixed point argument, we can preserve such property for the nonlinear problem:

$$\|m(t) - \bar{m}\|_{H^{-1}}^2 \lesssim C\left(e^{-\omega_0 t} + e^{-\omega_0 (T-t)}\right)$$

- Using

$$\|F(x, m(t)) - F(x, \bar{m})\|_{\mathcal{C}^{1+\alpha}} \leq \bar{C}\|m(t) - \bar{m}\|_{H^{-1}}$$

we bootstrap the estimates between the two equations, using the exponential decay of the operators.

Links with optimal control problems

MFG as optimality system (for a bilinear control problem).

Ex: Optimize in terms of the field α

$$\inf_\alpha \int_0^T [\int_\Omega \tfrac{1}{2}\, m\, |\alpha|^2 + \mathcal{F}(m(s))] ds$$

$$\text{state eq.} \qquad m_t - \Delta m - \text{ div } (\alpha\, m) = 0\,, \qquad m(0) = m_0$$

Optimality gives:

$$\begin{aligned} &\alpha_{opt} = Du(t,x) \\ &-u_t - \Delta u + \tfrac{1}{2}\alpha_{opt} \cdot Du = F(m) \end{aligned} \quad \Leftrightarrow \quad -u_t - \Delta u + \frac{1}{2}\,|Du|^2 = F(m)$$

We proved: Controls $[Du^T]$ and trajectories $[m^T]$ which are optimal in $[0, T]$ are close to the corresponding steady-state ones.

The convergence holds in average and exponentially in the transient time.

Is this a general issue of optimality systems?

It turns out that similar exponential estimates hold for a wide class of optimal control problems in the long horizon
(joint work with E. Zuazua).

Ex: linear case $\iff$ minimize a quadratic cost

$$J(u) = \frac{1}{2}\int_0^T \left[\|Cx - z\|^2 + \|u\|^2\right] dt$$

over the dynamics

$$\begin{cases} x_t + Ax = Bu \\ x(0) = x_0. \end{cases}$$

where $A, B, C \in \mathcal{M}_N$, $z \in \mathbb{R}^N$ is some target observation.

The optimality system reads as

$$\begin{cases} x_t + Ax = -BB^* p \\ -p_t - A^* p = C^* Cx - C^* z \end{cases}$$

Theorem

If (A, B) is controllable and (A, C) is observable, then there exist $\kappa > 0$ and K:

$$|u^T(t) - \bar{u}| + |x^T(t) - \bar{x}| \leq K(e^{-\kappa t} + e^{-\kappa(T-t)}) \qquad \forall t \in [0, T],$$

where (u^T, x^T) and $(\bar{u}, \bar{x})$ are the evolution and the stationary optimal control and state.

- Actually, κ is characterized as the exponential rate of the dynamics stabilized through the solution of algebraic Riccati equation.
- We extend the same approach to infinite dimensional setting (at least for a large class of examples, ex. heat and wave equations). Stabilization and observability estimates play a crucial role.

  The linearized MFG system (around the ergodic solution) is an example of this kind.
- Properties of the linearized systems + fixed point arguments are a possible approach for nonlinear systems, as in MFG.

Conclusions

- Under mild monotonicity conditions, we have shown, as $T \to \infty$

  (i) the convergence of $\frac{u(t)}{T}$ to $\bar{\lambda}(T-t)$

  (ii) the convergence of $u(t) - \int_\Omega u(t,y)dy$ to $\bar{u}$

  (iii) the convergence of $m(t)$ to $\bar{m}$

  expressed in different norms or scales.

- Under either stronger monotonicity in the local case or stronger continuity in the nonlocal case we have shown that the convergence has exponential rate in the transient time.

- The results obtained are consistent with a general behavior of optimality systems in the long horizon. The structure of the linearized system explains the exponential stability and suggests more general viewpoints

Thanks for the attention !