A DETERMINISTIC GAME APPROACH TO CURVATURE FLOWS AND FULLY NON LINEAR PDEs

Note Title 07/07/2009

The aim of this minicourse is to review some classical links between optimal control and first order PDES and to present some new connections, via simple 2-player repeated games, with second order PDES.

I/ Deterministic control and first order PDES

1) First examples: the exit time from a domain

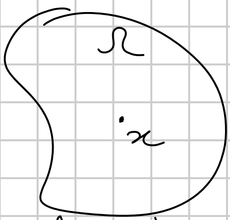

Ω domain in $\mathbb{R}^n$

x = starting point

What are the trajectories starting from x which minimize the time to exit Ω with constraints on the velocity?

$$\begin{cases} \dot{y}(s) = \alpha(s) \\ y(0) = x \end{cases}$$

with α = velocity $\in \mathcal{A}$

α: "control"; $\mathcal{A}$: admissible set

α can be chosen freely in $\mathcal{A}$.

For example $\mathcal{A} = B(0,1)$ $\rightsquigarrow$ constraint $|\alpha| \le 1$

The trajectories minimizing the exit time are then the straight segments originating at x and minimizing the distance to $\partial\Omega$.

So setting

$$u(x) = \min_{\alpha \in A} T(\alpha) \qquad T(\alpha) = \text{exit time for the choice } \alpha(s).$$

we have $u(x) = \operatorname{dist}(x, \partial\Omega)$

We can notice this u is solution to the PDE

$$\begin{cases} |Du| = 1 & \text{in } \Omega \\ u = 0 & \partial\Omega \end{cases}$$

but it is not the only solution. ex in dim 1

$\Omega = [0,1]$

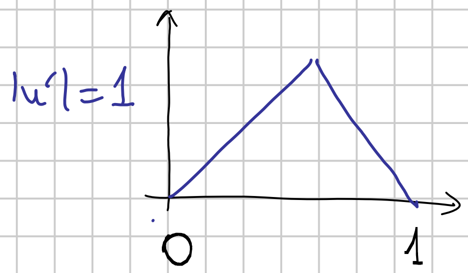

another possible solution
(there are infinitely many!)

Among all the solutions to this PDE, the control problem selects one. Note that they all have singularities (points of nonsmoothness = shocks)

The game of Paul and Carole (invented by Joel Spencer)

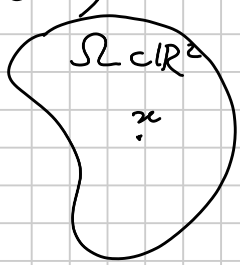

Paul starts from $x \in \Omega$ and wants to exit from Ω in the minimal possible time. Carole is his opponent and wants to obstruct him. A small number ε is given.

- Paul chooses a vector v of norm 1
- then Carole chooses a number $b = \pm 1$

- then Paul moves from x to $x + \sqrt{2}\,\varepsilon\, b v$
- the previous steps are repeated from that position, until exit.

What is the minimal exit time for Paul? Can he exit? What is his optimal exit strategy?

A variant with objective function can be given: u_0 continuous bounded function on $\mathbb{R}^2$ is given. A horizon time T is given. Starting from x at time t, Paul wants to minimize his final "score" $u_0(y(T))$

$y(T)$ ↳ position at final time T

with the same motion rules.

We will see that

$$u(x,t) = \min \{ u_0(y(T)) \,,\ \text{starting from } x \text{ at } t \}$$

is solution to a PDE.

2) the general setting for the classical theory

We are given a controlled dynamical system

$$\begin{cases} \dot{y}(s) = f(y(s), \alpha(s)) \\ y(0) = x \quad \text{or} \quad y(t) = x \end{cases}$$

x is a starting point, t a starting time, T a final time.

f is continuous and Lipschitz with respect to x

$\alpha(\cdot) \in \mathcal{A} = \{\alpha : [0, +\infty[\longrightarrow A\}$ A = control set = compact metric space (given).

The goal of the game is always to optimize = maximize or minimize a certain quantity (score, cost, minimal time ...)
These problems arise in engineering, economics, finance. Then the scores are called "utility" (economics), costs (engineering)...

Possible variants:

- control pb with finite horizon time
a final time T is given, functions r and g (regular enough) are given. The cost starting from x at time t is then

$$C_{x,t}(\alpha) = \int_t^T \underbrace{r(y(s), \alpha(s))}_{\text{running cost along trajectory}} ds + \underbrace{g(y(T))}_{\text{cost at final time}}$$

- Control pb with infinite horizon time:

$$C_x(\alpha) = \int_0^\infty r(y(s), \alpha(s)) \underbrace{e^{-s}}_{\text{ensures convergence of the } \int} ds$$

- minimal time problem
Given a closed target set $\mathcal{G}$
the minimal time to reach $\mathcal{G}$ starting

from x at time 0 is

$$t_x(\alpha) = \min \{ s \,/\, y(s) \in \mathcal{G} \}$$

- minimal time with discounting

$$E_x(\alpha) = \int_0^{t_x(\alpha)} e^{-s}\, ds$$

The discounting factors e^{-s} or e^{-rs} can represent the discounting of wealth under the effect of inflation: an amount x today will be worth $e^{-rs}x$ after time s. (Conversely money can be invested with an interest rate ρ)

Each time we introduce the <u>value function</u> associated to the problem:

- finite horizon

$$u(x,t) = \inf_{\alpha \in \mathcal{A}} C_{x,t}(\alpha)$$

minimal cost starting from x at t

- infinite horizon

$$u(x) = \inf_{\alpha \in \mathcal{A}} C_x(\alpha)$$

minimal cost starting from x

- minimal time

min time starting from x

$$u(x) = \inf_{\alpha} t_x(\alpha) \quad \text{or} \quad \inf_{\alpha} E_x(\alpha)$$

3) The dynamic programming principle

It's the key point for relating control problems and PDEs.

It is a relation on the value function u that expresses a sort of induction, saying that the value u is optimal if

- one lets the system evolve by a (small) time h choosing a control α
- one continues with the optimal strategy starting from the new state
- one sums the two scores/costs and optimizes over the initial choice of α.

Prop°: the value functions defined above satisfy $\forall h>0$

- finite horizon time T $\quad h \leq T-t$

$$u(x,t) = \inf_{\alpha\in\mathcal{A}} \left[\int_t^{t+h} r(y,\alpha)\, ds + u(y(t+h), t+h) \right]$$

- infinite horizon time

$$u(x) = \inf_{\alpha\in\mathcal{A}} \left[\int_0^h r(y,\alpha) e^{-s}\, ds + u(y(h)) e^{-h} \right]$$

- minimal time

$$u(x) = \inf_{\alpha\in\mathcal{A}} \left[h + u(y(h)) \right] \quad \text{or} \quad \inf \left[\int_0^h e^{-s} ds + e^{-h} u(y(h)) \right]$$

// this proposition is easily proved by gluing together trajectories on $[0,h]$ and $[h,T]$ and using the property of time translation invariance. //

In addition, we have natural final time/boundary conditions:

- finite horizon time $u(x,T) = g(x)$
- minimal time $u(x) = 0$ if $x \in \mathcal{T}$

4) Formal derivation of the PDE

Assuming u is regular enough to use a Taylor expansion (Note: it is not in g^{al}, cf $dist(\cdot, \partial\Omega)$) we let formally $h \to 0$ in the relations above.

For the case of finite horizon time

$$u(x,t) = \inf_{\alpha \in \mathcal{A}} \left[\int_t^{t+h} r(y,\alpha)\, ds + u(y(t+h), t+h) \right]$$

$$\cancel{u(x,t)} = \inf_{\alpha \in \mathcal{A}} \left[h\, r(x,\alpha) + \cancel{u(x,t)} + \nabla u \cdot \frac{dy}{dt} h + h \partial_t u + o(h) \right]$$

letting $h \to 0$

$$0 = \inf_{\alpha \in A} \left[r(x,\alpha) + \nabla u \cdot f(x,\alpha) + \partial_t u \right]$$

$$\to \quad \partial_t u + H(x, \nabla u) = 0$$

where $H(x,p) = \inf_{\alpha \in A} \left[r(x,\alpha) + p \cdot f(x,\alpha) \right]$

with the terminal time condition $u = g$

$$(HJB) \begin{cases} \partial_t u + H(x, \nabla u) = 0 & t < T \\ u = g & t = T \end{cases}$$

is called the Hamilton-Jacobi-Bellman eq. associated to the control pbl. H is called the Hamiltonian. It is a concave function in p (a Legendre transform)

(HJB) is a <u>1st order</u> PDE.

For the other pbls, we obtain similarly
- infinite horizon $\quad u(x) + H(x, Du) = 0$
- minimal time $\quad \begin{cases} H(x, Du) = 1 & x \notin \mathcal{G} \\ u = 0 & x \in \mathcal{G} \end{cases} \qquad \begin{cases} H(x, Du) + u = 0 \\ u = 0 \quad \mathcal{G} \end{cases}$

One can in fact show rigorously (see below) that u is a solution to this equation. Moreover all 1st order PDEs with concave hamiltonians can be obtained this way. $\rightarrow$ the control pbl provids the existence of a solution and a representation formula for it

$$u(x,t) = \inf_{\alpha} C_{x,t}(\alpha) = \inf_{\alpha} [\ \cdots\ ] \qquad \underline{\text{representation formula}}$$

Allows to prove properties on solutions, calculate them numerically ...

Conversely, the knowledge of the PDE allows to solve the control problem, since the u = minimal cost is identified, and from it one may reconstruct the optimal α (the one for which there is equality in the definition of H).

So this is a sort of 1-to-1 correspondance.

Note: to obtain all possible 1st order eq with non concave hamiltonians, the optimisation pbl needs to be replaced by a "two player differential game" i.e by a min-max

There are then 2 players and 2 controls $\alpha(s)$, $\beta(s)$ for each of them.

$$\begin{cases} \frac{dy}{ds} = f(y, \alpha, \beta) \\ y(t) = x \end{cases}$$

again a cost function $C_{x,t}(\alpha, \beta)$

the player 1 chooses α to minimize $C_{x,t}$

——— 2 ——— β — maximize $C_{x,t}$

This is called a <u>differential game</u>

To write the DPP one needs to be more careful and define the notion of <u>strategy</u>: which player makes their choice first? What information do they have at each time? Do they know only past choices of the other player? All this modifies the DPP.

Assuming that each player only knows the choices of the past ("nonanticipating strategies") we can define two value functions for the game:

- lower value function

$$v(x) = \inf_{\alpha \in A} \sup_{\beta \in B} C_{x,t}(\alpha, \beta)$$

corresponds to the first player announcing his choice first

- upper value function

$$u(x) = \sup_{\beta \in B} \inf_{\alpha \in A} C_{x,t}(\alpha, \beta)$$

corresponds to the second player announcing his choice first. We can check that $v(x) \leq u(x)$.

We obtain Hamilton-Jacobi-Isaacs PDE's for u and v that involve hamiltonians H and $\tilde{H}$

$$H(x,p) = \min_{\beta \in \mathcal{B}} \max_{\alpha \in \mathcal{A}} \left[r(x,\alpha,\beta) + f(x,\alpha,\beta) \cdot p \right]$$

$$\tilde{H} = \max \min \underline{\qquad}$$

In some cases one can show that $\tilde{H} = H$ (Isaacs condition) and then $u = v$.

This allows in practice to obtain all 1st order PDES.

Now note that on the other hand, the Paul-Carol game will lead to a $\underline{2^d}$ order equation; more precisely

$$\partial_t u = |Du| \operatorname{div}\left(\frac{Du}{|Du|}\right)$$ contains second derivatives

Later on we will introduce another game which allows to represent solutions to general second order PDES

$$-\partial_t u + f(x,t,u,Du,D^2u) = 0$$

where f is a general type of nonlinearity (conditions to be specified later)

5) A word on stochastic representations

We have seen there exists a correspondance between viscosity solutions of first order Hamilton-Jacobi equations and deterministic control problems.

So deterministic control / games seem to always lead to first order equations.

To interpret second order equations, a natural way is to use stochastic control interpretation, i.e. via probabilistic processes such as Brownian motion. Instead of optimizing over trajectories, one computes expectations along random trajectories.

example: the heat eq and Brownian motion: Brownian motion is a stochastic process (i.e. a map from $[0,\infty) \times \Omega \to \mathbb{R}^n$) such that $B_t - B_s$ is independent of the past and follows a normal law with variance $(t-s)I$.

Denoting by $p(t, x, s, A)$ the probability to be in A at time s after being at x at time t, then

$$p(t, x, s, A) = p(0, 0, s-t, A-x)$$

$$= \frac{1}{\sqrt{2\pi(t-s)}^n} \int_A \exp\left(-\frac{|y-x|^2}{2(t-s)} dy\right)$$

Given a function $h(x)$ and a terminal time T, set $u(x,t) = \mathbb{E}_{x,t}\left[h(y(T))\right]$ expectation starting from x at time t

$$= \frac{1}{2\pi(T-t)^{n/2}} \int_{\mathbb{R}^n} h(y)\, e^{-\frac{|y-x|^2}{2(T-t)}}\, dy$$

✓ where $y(t)$ is the evolution of the Brownian motion starting from x at t.

u solves the PDE

$$\begin{cases} \partial_t u + \frac{1}{2}\Delta u = 0 \\ u = h \quad \text{at } t = T \end{cases}$$

this can be found by a direct calculation from the formula above, (one may recognize the heat kernel)

→ this provides a representation formula for the heat eq $\quad u(x,t) = \mathbb{E}\left(h(y(T))\right) \quad$ (Feynman-Kac)

II/ The Paul-Carol game and mean curvature flow

1) The game with exit time

Paul is sure to exit in one step if he starts outside of the dotted line consisting of the set of points which are midpoints of segments of length $2\sqrt{2}\varepsilon$ whose endpoints are both on $\partial\Omega$

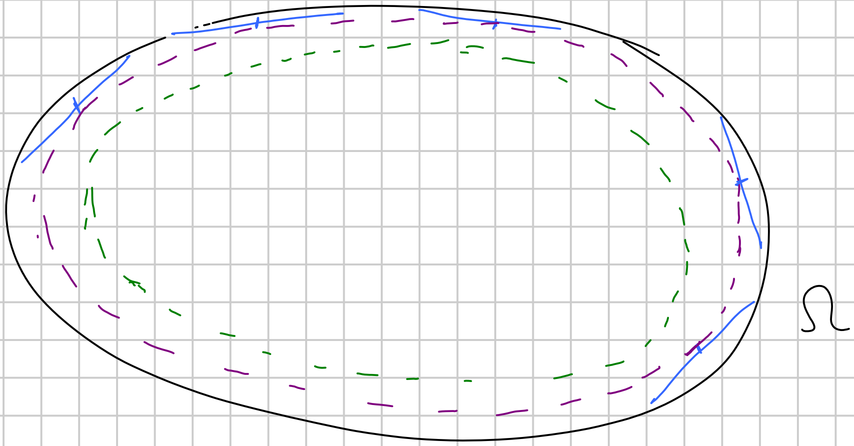

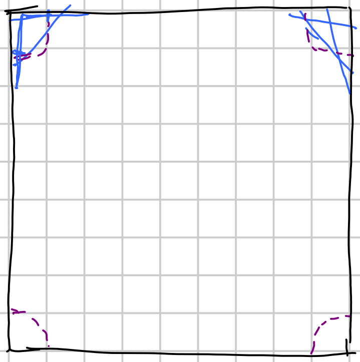

Then one can iterate the reasoning : the set of points from which he is sure to escape in 2 steps is the one from which he can get in one step outside of the purple curve → green curve.

We work by analogy with the control problems studied in the first section and set a value function

$V_\varepsilon(x)$ = min exit time for Paul starting from x (under Carol's best opposition)

where ε is the step-size (parameter of the game)

We have $U_\varepsilon(x) = 0$ if $x \in \mathbb{R}^2 \setminus \Omega$
How do we know that $U_\varepsilon(x) < \infty$ or in other words that Paul can be sure to exit?

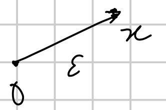

Let 0 be Paul's starting point and y_n the position after n turns. Let's always choose inductively the next vector v to be a unit vector perpendicular to y_n. Then no matter what Carol's choice is, by Pythagore's Th

$$\|y_{n+1}\|^2 = \|y_n + \sqrt{2}\varepsilon v\|^2 = \|y_n\|^2 + 2\varepsilon^2$$

So $\|y_n\|^2 \geq 2n\varepsilon^2 \rightarrow$ in enough steps $\left(\frac{1}{\varepsilon^2}\right)$ Paul can exit from any ball so from any domain. This is an exit strategy for him, but <u>not</u> the best exit strategy.

We can write a principle of dynamic programming:

$$\boxed{U_\varepsilon(x) = \min_{|v|=1} \max_{b=\pm 1} \left[U_\varepsilon(x + \sqrt{2}\varepsilon b v) + \varepsilon^2 \right]}$$

where we decide 1 step of time $= \varepsilon^2$
this "induction relation" characterizes U_ε.

By a formal Taylor expansion we can formally derive a PDE:

$$\cancel{U_\varepsilon(x)} = \min_{|v|=1} \max_{b=\pm 1} \Big[\cancel{U_\varepsilon(x)} + \sqrt{2}\,\varepsilon\, b v \cdot DU_\varepsilon(x) + \varepsilon^2 \langle D^2 U_\varepsilon v, v \rangle + \varepsilon^2 + o(\varepsilon^2) \Big]$$

$$0 = \min_{|v|=1} \Big[\sqrt{2}\,\varepsilon\, |v \cdot DU_\varepsilon(x)| + \varepsilon^2 \langle D^2 U_\varepsilon(x) v, v \rangle + \varepsilon^2 + o(\varepsilon^2) \Big]$$

The first order as $\varepsilon \to 0$ is the most important term so to minimize it we must take $v = \pm \frac{D^\perp U_\varepsilon}{|DU_\varepsilon|}$

$$0 = \varepsilon^2 \Big[\Big\langle D^2 U_\varepsilon \frac{D^\perp U_\varepsilon}{|DU_\varepsilon|}, \frac{D^\perp U_\varepsilon}{|DU_\varepsilon|} \Big\rangle + 1 + o(1) \Big]$$

letting $\varepsilon \to 0$, and assuming $U_\varepsilon \to U$ we find

$$\begin{cases} 0 = \Big\langle D^2 U \frac{D^\perp U}{|DU|}, \frac{D^\perp U}{|DU|} \Big\rangle + 1 & \Omega \\ U = 0 & \partial\Omega \end{cases}$$

This is a second order degenerate parabolic PDE. It is equivalent (in 2D) to

$$0 = 1 + |DU| \operatorname{div}\Big(\frac{DU}{|DU|}\Big)$$

The reason why we got a second rather than first order PDE is that the first order term in the calculation degenerated, because of the min-max representation. This seems to require both having two players and having a discrete game and an asymptotic representation. (the representation we obtain only comes as an $\varepsilon \to 0$ limit)

But this representation is deterministic, in contrast with usual stochastic representations of second order PDEs

The PDE above is related to mean curvature flow, a classic geometric motion. To see it it is more convenient to introduce the finite horizon time version of the game:

2) The game with horizon time

A function $u_0(x)$ is given in $\mathbb{R}^2$ and Paul's objective is to reach the minimal value of u_0 at $y(T)$ where T is the final time and $y(T)$ his position at that time. Otherwise the rules of motion are the same.

Choosing $u_0 > 0$ in Ω, $u_0 = 0$ in Ω and $u_0 < 0$ in $\mathbb{R}^2 \setminus \Omega$ pushes Paul to try to exit, as in the other version of the game.

We then set again a value function

$$u_\varepsilon(x,t) = \min \left[u_0(x(T)), \text{ starting from } x \text{ at time } t \right]$$

A dynamic programming principle can be written:

$$u_\varepsilon(x,t) = \min_{|v|=1} \max_{b=\pm 1} u_\varepsilon(x + \varepsilon b v, t + \varepsilon^2)$$

The same reasoning of formal Taylor expansion gives that if $u_\varepsilon \to u$ then u solves

$$\begin{cases} \partial_t u + \left\langle D^2 u \dfrac{\nabla^\perp u}{|\nabla u|}, \dfrac{\nabla^\perp u}{|\nabla u|} \right\rangle = 0 & \mathbb{R}^2 \\ u(x,T) = u_0(x) \end{cases}$$

equivalent to $\begin{cases} \partial_t u + |\nabla u| \operatorname{div}\left(\dfrac{\nabla u}{|\nabla u|}\right) = 0 \\ u(x,T) = u_0(x) \end{cases}$

this is the PDE of mean curvature motion, solved backwards in time. This PDE was studied by Evans-Spruck, Chen-Giga-Goto, in the framework of viscosity solutions.

Mean curvature motion corresponds to the evolution of curves in the plane (or hypersurfaces in $\mathbb{R}^n$) under

$$V_n = \kappa = \text{mean curvature}$$

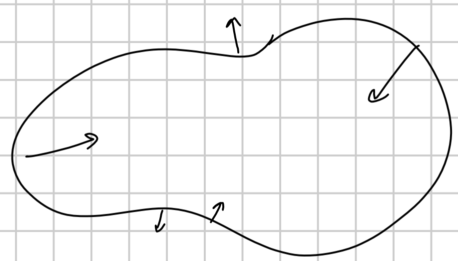

Gage-Hamilton, Grayson: curves become convex, rounder and rounder and shrink to a point

A convenient way to analyse this is the level-set formulation (Osher-Sethian): view the curve Γ_t evolving in time as the 0-level set of a function $u(x,t)$: $\Gamma_t = \{x \,/\, u(x,t) = 0\}$

Then normal velocity of the curve is $V_n = \dfrac{\partial_t u}{|\nabla u|}$

and curvature $\kappa = \operatorname{div} n = \operatorname{div} \frac{\nabla u}{|\nabla u|}$

$$\rightsquigarrow \quad \frac{\partial_t u}{|\nabla u|} = \operatorname{div}\left(\frac{\nabla u}{|\nabla u|}\right)$$

this is the PDE we have found (except solved backwards in time)

What is U? There is a formal correspondance between u and U via $u(x, U(x)) = 0$ or in other words $U(x)$ is the time it takes for the "front" of MCF to pass through x (but only for convex domains)

The optimal strategy for Paul was given by $v = \frac{\nabla^\perp u}{|\nabla u|}$ (in the formal reasoning) so Paul should always move along the tangent to the level set of u or U at the current point.

For nonconvex domains the situation is a bit different: Paul cannot exit from the concave parts of $\partial\Omega$

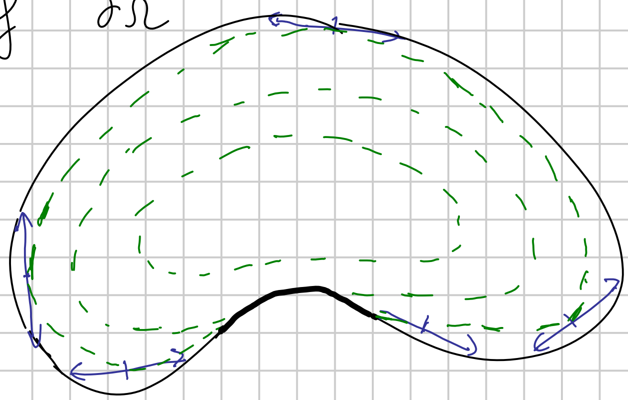

The concave part of the bdry is like a wall

The game with exit time and the associated PDE $1 + |Du| \operatorname{div}\left(\frac{Du}{|Du|}\right) = 0$ really correspond to motion by the positive part of the curvature, or $V_n = \kappa_+$. It's different from the horizon time game because in the latter the game doesn't stop if Paul exits from Ω, he can be led to exiting then re-enter. It does however match the game with horizon time if the rule $|v|=1$ is replaced with $|v| \leq 1$ which allows Paul to stop temporarily if he likes.

The solution to the ε-game (midpoints of segments of size $2\sqrt{2}\varepsilon$ joining points on the bdry) provide a numerical scheme (already introduced by Catté-Dibos-Koepfler) for mean curvature (or κ_+) motion (discrete in time).

The limit of u_ε would have been the same if the action of Carole was replaced by flipping a coin to determine b.

This can be extended to higher dimensions, and more general curvature flows. The game has to be modified appropriately: for example for MCF in dim n, Paul picks an o.n. frame $v_1, \dots, v_{n-1}$, Carol picks $n-1$ numbers ± 1, and Paul moves by $\varepsilon \sum_{i=1}^{n-1} b_i v_i$

3) The Mark-Helen game for general second order PDEs

The possibility of approximating solutions to 2^d order PDEs by discrete deterministic repeated games is <u>not</u> limited to geometric flows as above.

We can treat

$$\begin{cases} -\partial_t u + f(t,x,u,Du,D^2u) = 0 & \mathbb{R}^n \\ u(x,T) = u_0 \end{cases}$$

or

$$\begin{cases} f(x,u,Du,D^2u) = 0 & \Omega \\ u = g & \partial\Omega \end{cases}$$

with f "elliptic" i.e. decreasing in the matrix D^2u.

These are called elliptic / parabolic fully nonlinear equations. They include Δu but also more general dependence on second order derivatives.

We remove for simplicity the u-dependence here. The game with horizon time, introduced by Kohn-S is the following:

There are two opposing players, Mark, Helen

- we start from x at time t, T is given, h a function of space is given. z represents Helen's wealth, initially $z=0$.
- Helen chooses $p \in \mathbb{R}^n$ and Γ a symmetric $n\times n$ matrix

- Mark chooses $w \in \mathbb{R}^n$ and the game moves to $x + \varepsilon w$
- Helen's score is changed to
$$z - \left[\varepsilon p.w + \frac{\varepsilon^2}{2} \langle \Gamma w, w \rangle + \varepsilon^2 f(t, x, p, \Gamma) \right]$$
- time is reset to $t + \varepsilon^2$
- the previous steps are repeated until time T.
- At time T Helen collects a "bonus" $h(x(T))$

The value function for the game is

$u_\varepsilon(x,t)$ = Helen's best possible final wealth, starting from x at time t, under Mark's best opposition

The dynamic programming principle is

$$u_\varepsilon(x,t) = \max_{p,\Gamma} \min_{w} \left[u_\varepsilon(x + \varepsilon w, t + \varepsilon^2) - \left[\varepsilon p.w + \frac{\varepsilon^2}{2} \langle \Gamma w, w \rangle + \varepsilon^2 f(t, x, p, \Gamma) \right] \right]$$

with $u_\varepsilon(x, T) = h(x)$

A formal derivation can be performed as before with Taylor expansion

$$u_\varepsilon(x,t) = \max_{p,\Gamma} \min_{w} \Big[u_\varepsilon(x,t) + \varepsilon \nabla u_\varepsilon . w + \frac{\varepsilon^2}{2} \langle D^2 u_\varepsilon w, w \rangle + \varepsilon^2 \partial_t u_\varepsilon - \varepsilon p.w - \frac{\varepsilon^2}{2} \langle \Gamma w, w \rangle - \varepsilon^2 f(t, x, p, \Gamma) \Big] + o(\varepsilon^2)$$

$$0 = \max_{p,\Gamma} \min_{w} \Big[\varepsilon w.(\nabla u_\varepsilon - p) + \varepsilon^2 \Big[\langle (D^2 u_\varepsilon - \Gamma) w, w \rangle + \partial_t u_\varepsilon - f(t,x,p,\Gamma) \Big] + o(\varepsilon^2) \Big]$$

looking at first order as the most important term the min$_w$ is $-\infty$ unless $p = \nabla u_\varepsilon(x)$ which forces the choice of p. Then there remains

$$0 = \max_{\Gamma} \min_{w} \Big[\langle (D^2 u_\varepsilon - \Gamma) w, w \rangle + \partial_t u_\varepsilon - f(t,x,\nabla u_\varepsilon, \Gamma) + o(1) \Big]$$

The min$_w$ can be made $-\infty$ unless $D^2 u_\varepsilon - \Gamma \geq 0$ as a symmetric matrix. Then the min is 0 and there remains

$$0 = \max_{\Gamma \leq D^2 u_\varepsilon} \Big[\partial_t u - f(t,x,\nabla u_\varepsilon, \Gamma) + o(1) \Big]$$

but by ellipticity of f, $f(t,x,\nabla u_\varepsilon \Gamma) \geq f(t,x,\nabla u_\varepsilon, D^2 u_\varepsilon)$ when $\Gamma \leq D^2 u_\varepsilon$ so the max is achieved for $\Gamma = D^2 u_\varepsilon$ and we get at the limit $\varepsilon \to 0$

$$0 = \partial_t u - f(t,x,\nabla u, D^2 u)$$

which is the desired equation.

In fact bounds on w, p, Γ and growth of f have to be imposed for this to work.

Almost optimal choices for Helen are $p = \nabla u$, $\Gamma = D^2 u$ given by the solution to the PDE.

- this is a deterministic control interpretation for all 2^d order eq - can serve as a semi-discrete numerical scheme.

- it has a "financial" interpretation: Mark stands for the market and Helen for the hedger $p = p_1 \dots p_n$ represents units of stocks $1, \dots n$ with prices $x_1 \dots x_n$.
Helen hedges her position by buying/selling to have $-p_i$ units of stock i. Mark sees her choices and sets the new prices to be $x_i + \varepsilon w_i$.
Helen gains or loses $-\varepsilon p.w$ due to the price change. She has to make choices which make her indifferent to the market fluctuations and there is only one such choice.
This is a deterministic analogue of the Black-Schols model for option pricing where the action of the market would be stochastic, but the best strategy for Helen would be the same.
The Γ and f ... represent here second order corrections which do not have a financial interpretation.

III/ Elements for rigorous proofs

The keyword here is that of viscosity solutions: first order control theory results are best proved via viscosity solutions. MCF eq has a viscosity solution framework, same for 2d order fully nonlinear equations. But until our games, there was no deterministic control correspondence with these viscosity solutions.

1) Deriving the PDEs via verification argument.

Historically, the first method for deriving rigorously the PDE was that of the "verification argument".
Its problem is that it assumes that one already knows a solution to the PDE which is differentiable.

example for the exit time problem:
Let $u(x)$ be the value function and U a (regular enough) solution to $\begin{cases} H(x, U) + 1 = 0 \\ u = 0 \quad \mathcal{G} \end{cases}$

We recall $H(x, DU) = \inf_{\alpha \in A} [f(x,\alpha) . DU] = -1$

so $\forall \alpha \in A \quad f(x,\alpha) . DU \geq -1$

Let y be any admissible trajectory starting from x
Let's compute

$$\frac{d}{ds}[U(y(s))] = DU . \frac{dy}{ds} = DU . f(y,\alpha) \geq -1$$

So $\underbrace{U(y(t_{fin}))}_{\in \mathcal{G}} - U(x) \geq - t_{fin}$

$\rightarrow t_{fin} \geq U(x)$

Taking the inf over all final times we find

$$u(x) \geq U(x)$$

On the other hand choosing α achieving the inf in the def of H we can find a trajectory such that equality holds and $t_{fin} = U(x)$. It follows that $u(x) \leq U(x)$. So $u = U$.

As we mentioned, solutions to HJB equations generally have shocks so this doesn't apply well.

2) Viscosity solutions

This is the appropriate rigorous framework. It was introduced by Crandall and Lions.

Motivation for viscosity solutions:

We want to solve $-\partial_t u + H(x, Du) = 0$ by adding a small "viscous" perturbation $\varepsilon \Delta u$

$$\begin{cases} -u_t^\varepsilon + H(x, Du^\varepsilon) - \varepsilon \Delta u^\varepsilon = 0 & t < T \\ u^\varepsilon = g & t = T \end{cases}$$

The idea is that as $\varepsilon \to 0$ u^ε will tend to a sol to the HJ eq, but not <u>any</u> solution, a "nice" one called viscosity solution, and unique.

u may not be smooth but now u^ε is. Take ϕ a smooth test-function s.t. $u-\phi$ has a point of local maximum at (x_0, t_0). We may find $(x_\varepsilon, t_\varepsilon) \to (x_0, t_0)$ s.t. $u^\varepsilon - \phi$ has a local max at $(x_\varepsilon, t_\varepsilon)$. Then

$$\nabla(u^\varepsilon - \phi)(x_\varepsilon, t_\varepsilon) = 0 \qquad \partial_t (u^\varepsilon - \phi)(x_\varepsilon, t_\varepsilon) = 0$$

and $D^2(u^\varepsilon - \phi)(x_\varepsilon, t_\varepsilon) \leq 0 \to \Delta u^\varepsilon \leq \Delta \phi$ at $(x_\varepsilon, t_\varepsilon)$

It follows that $(\phi_t + H(x, D\phi))(x_\varepsilon, t_\varepsilon)$

$$= u_t^\varepsilon + H(x, Du^\varepsilon)(x_\varepsilon, t_\varepsilon) = \varepsilon \Delta u^\varepsilon (x_\varepsilon, t_\varepsilon)$$
$$\leq \varepsilon \Delta \phi (x_\varepsilon, t_\varepsilon)$$

Letting $\varepsilon \to 0$, since ϕ is smooth, we find

$$(\phi_t + H(x, D\phi))(x_0, t_0) \leq 0$$

II now we consider second order equations

(2) $\quad -\partial_t u + F(x, u, Du, D^2 u) = 0$

the same reasoning can be made provided F is increasing in $D^2 u$. Thus way $D^2 u^\varepsilon \leq D^2 \phi$ can be inserted into the eq as above.

This nondecreasing condition is called an ellipticity condition. Equations (2) such that this holds are called parabolic (elliptic if there is no time dependence)

This condition of ellipticity / parabolicity " is the appropriate one for 2^d order eqns in order to give a notion of viscosity solution (for 1st order no condition is needed).

<u>Def</u>: A lsc function $u: \Omega \to \mathbb{R}$ is a viscosity subsolution of
$$F(x,u,Du,D^2u)=0$$
$$\text{resp. } -\partial_t u + F(x,u,Du,D^2u)=0$$
if $\forall \phi \in C^2(\Omega)$, if x / (x,t) is a point of local maximum of $u-\phi$, then
$$F(x,u,D\phi(x),D^2\phi(x)) \leq 0$$
$$\text{resp. } -\partial_t\phi + F(x,u,D\phi,D^2\phi) \leq 0$$

An usc function $u: \Omega \to \mathbb{R}$ is a viscosity supersolution if $\forall \phi \in C^2(\Omega)$, if x / (x,t) is a point of local minimum of $u-\phi$ then

______________ ≥ 0

≥ 0

A function is a viscosity solution if it is both a viscosity sub and supersolution.

- Classical solutions are viscosity solutions
- boundary conditions or terminal time conditions can be added. With them and for a wide class of F's we have a <u>comparison th</u>:

<u>Th</u>: any subsolution is smaller than any supersolution

This allows to deduce that a viscosity solution is <u>unique</u> $\to$ this provides a mechanism for selecting a good solution of the PDE.

The viscosity solution coincides with the one obtained by the "vanishing viscosity method" above, i.e $\lim u^\varepsilon$.

- There exists a notion of boundary condition "in the sense of viscosity" which is different from the classical sense.

ex:
$$\begin{cases} |u'| = 1 & \text{on } [0,1] \\ u(0) = u(1) = 0 \end{cases}$$

viscosity solution can be found by
$$\begin{cases} |u'| - \varepsilon u'' = 0 \\ u(0) = u(1) = 0 \end{cases}$$

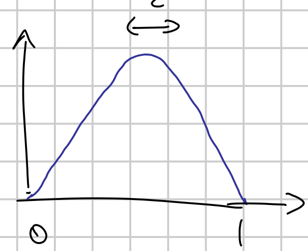

converges to

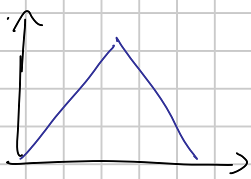

$u = \mathrm{dist}(x, \partial\Omega)$.

<u>the</u> viscosity solution of $|u'| = 1$.

ex
$$\begin{cases} u' = 1 & \text{on } [0,1] \\ u(0) = 0 \quad u(1) = 0 \end{cases}$$

first solve
$$\begin{cases} u' - \varepsilon u'' = 1 \\ u(0) = 0 \quad u(1) = 0 \end{cases}$$

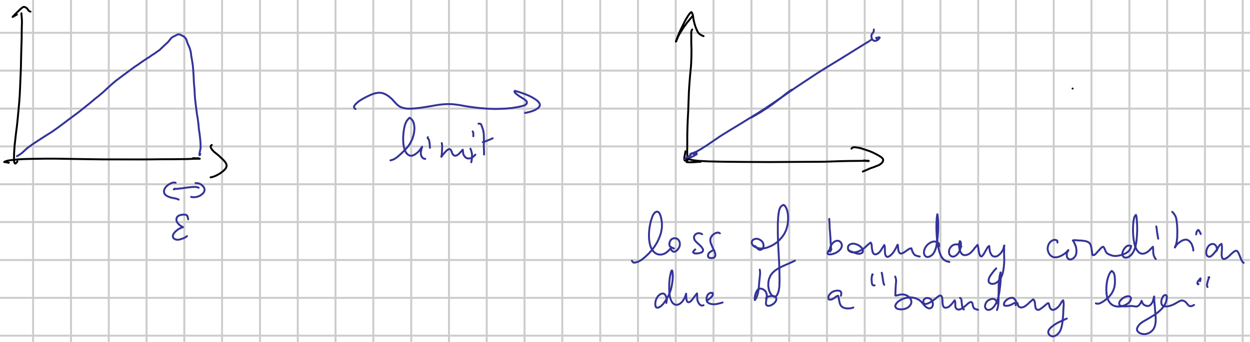

3) Main theorems

Th (Lions) the value functions $u(x)$ or $u(x,t)$ for the above optimal control problems is. the viscosity solution to the associated HJB eq.

Th (Kohn-S) The value function $u_\varepsilon(x,t)$ for the Paul-Carole game with finite time horizon converges locally uniformly in space-time to u, the viscosity solution to
$$\begin{cases} \partial_t u + |\nabla u| \operatorname{div}\left(\frac{\nabla u}{|\nabla u|}\right) = 0 \\ u(x,T) = u_0(x) \end{cases}$$

The analogue holds for the value function $u_\varepsilon(x)$ for the Paul-Carole exit time game provided a comparison principle holds for the limiting PDE
$$\begin{cases} 1 + |\nabla u| \operatorname{div}\left(\frac{\nabla u}{|\nabla u|}\right) = 0 & \text{in } \Omega \\ u = 0 & \partial\Omega \end{cases}$$
and this is known only in the case Ω is starshaped (Barles-Da Lio) open otherwise.

Th (Kohn-S) the value function u_ε for the Mark-

Helen games converge locally uniformly as $\varepsilon \to 0$ to the viscosity solution of

$$\begin{cases} -\partial_t u + F(t, x, u, Du, D^2 u) = 0 \\ u(x,T) = h(x) \end{cases}$$

provided a comparison th holds for this equation.

4) Sketch of the proof

The proofs of the above theorems rely on the same ingredient: express the dynamic programming principle as a semigroup property, and check "monotonicity and consistency" of that semigroup.

More precisely, we may write in the continuous case (classical deterministic control)

(1) $\underline{u(\cdot,t) = S_h\left[u(\cdot, t+h)\right]} \quad \forall h \leq T-t$

where the operator S_h is defined for any function ϕ by

$$S_h[\phi](x) = \inf_{\alpha \in \mathcal{A}} \left[\int_0^h r(y(s), \alpha(s))\, ds + \phi(y(h))\right]$$

where as before $\begin{cases} \frac{dy}{dt} = f(y,\alpha) \\ y(0) = x \end{cases}$

We can then easily check that

- $S_0[\phi] = \phi \quad \forall \phi$
- $S_h[\phi + k] = S_h[\phi] + k \quad \forall$ cst k.
- S_h is monotone i.e if $\phi_1 \leq \phi_2$ $S_h[\phi_1] \leq S_h[\phi_2]$

- If ϕ is C^2 then
$$\lim_{h\to 0} \frac{S_h[\phi]-\phi}{h} = H(x, \nabla\phi(x)) \qquad (\text{"consistency"})$$

(this is essentially the Taylor calculation that we did before now legal since ϕ is C^2).

For the case of the Paul-Carole game or Mark-Helen games, we replace $S(h)$ by S_ε defined by
$$S_\varepsilon[\phi](x) = \min_{|v|=1} \max_{b=\pm 1} \phi(x+\varepsilon b v)$$

or
$$S_\varepsilon[\phi](x) = \max_{p,\Gamma} \min_{w} \left[\phi(x+\varepsilon w) - \left(\varepsilon p\cdot w + \frac{\varepsilon^2}{2}\langle \Gamma w, w\rangle + \varepsilon^2 f(x,p,\Gamma)\right)\right]$$

then the principle of dynamic programming are the same as writing

$$(2) \quad \underline{u_\varepsilon(\cdot, t) = S_\varepsilon\left[u_\varepsilon(\cdot, t+\varepsilon^2)\right]}$$

One may check as above

- $S_\varepsilon[\phi+k] = S_\varepsilon[\phi]+k$
- if $u\le v$ $\quad S_\varepsilon[u] \le S_\varepsilon[v]$
- for any $\phi\in C^2$

$$\lim_{\varepsilon\to 0} \frac{S_\varepsilon[\phi]-\phi}{\varepsilon} = \begin{cases} |\nabla\phi| \operatorname{div} \frac{\nabla\phi}{|\nabla\phi|} \\ -f(x, D\phi, D^2\phi) \end{cases} \quad \text{according to the case}$$

(again this is proved by the direct Taylor calculation)

Then the theorems follow from (1) or (2) and the above properties. Indeed: let $\phi \in C^2$ and let (x_0, t_0) be a point of local maximum of $u - \phi$, then we may always modify ϕ far from (x_0, t_0) to make it a global maximum (without changing $\phi(x_0, t_0)$, $D\phi(x_0, t_0)$ ...) and thus we may write $\forall x, \forall h$

$$u(x, t_0 + h) - \phi(x, t_0 + h) \leq u(x_0, t_0) - \phi(x_0, t_0)$$

Apply the operator S_h to this relation, by monotonicity we have

$$\underbrace{S_h[u(\cdot, t_0 + h)]}_{u(\cdot, t_0) \text{ by principle of dynamic programming}} - S(h)[\phi(\cdot, t_0 + h)] \leq u(x_0, t_0) - \phi(x_0, t_0)$$

Evaluating at x_0 we find

$$0 \leq S_h[\phi(\cdot, t_0 + h)](x_0) - \phi(x_0, t_0)$$

dividing by h and letting $h \to 0$ we find (with the "consistency" property and the fact that $\phi \in C^1$)

$$0 \leq H(x_0, D\phi(x_0, t_0)) + \partial_t \phi(x_0, t_0)$$

hence the desired property i.e u is a subsolution. The proof of supersolution works the same way, so does the proof for the two discrete games.