Parabolic PDEs and Deterministic Games

Robert V. Kohn
Courant Institute, NYU

Joint work with Sylvia Serfaty

ICIAM07, Zurich, July 2007

Preamble

I'll be exploring connections between motion by curvature and optimal control. Brief reminder about their importance:

MOTION BY CURVATURE arises in

- materials science, as model of surface-energy-driven coarsening; and
- image processing, as scheme for denoising images without blurring edges.

DETERMINISTIC AND STOCHASTIC OPTIMAL CONTROL are connected with

- Hamilton-Jacobi equations, for example through the Hopf-Lax solution formula; and
- finance, for example through Merton's analysis of portfolio optimization.

Goals and perspective

Part 1: Motion by curvature has a deterministic control interpretation.

Part 2: So do $u_t = u_{xx}$ and $u_t = f(Du, D^2u)$!

	deterministic control	stochastic control	steepest descent
1st order nonlinear	Hamilton-Jacobi		
2nd order linear	today, part 2	random walk	$\int \|\nabla u\|^2$
2nd order nonlinear	today, part 1	controlled diffusion	perimeter

Deterministic control and first-order PDE

Focus on basic example:
motion with velocity 1

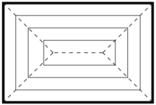

Arrival time: The function $u(x) =$ *time of arrival to x* solves $|\nabla u| = 1$ in Ω with $u = 0$ at $\partial\Omega$

Level set method: If $v(x,t)$ is such that *level sets move with velocity* 1 then $v_t/|\nabla v| = -1$

Equivalence: $v(x,t) = u(x) - t$

First-order HJ eqn: characteristics and shocks

Kinetic: models decay (or growth)

Optimal control: $u(x) = \min_{\text{velocity} \leq 1, \text{ starting at } x} \{\text{time to } \partial\Omega\}$

Motion by curvature

Analogous basic example:
motion in R^2 with velocity κ

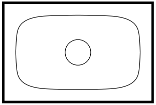

Arrival time: The function $u(x) =$ *time of arrival to x* solves
$\text{div}\,(\nabla u/|\nabla u|) = -1/|\nabla u|$

Level set method: If $v(x,t)$ is such that *level sets move with velocity* κ then $v_t/|\nabla v| = \text{div}\,(\nabla v/|\nabla v|)$

Equivalence: $v(x,t) = u(x) - t$

Second-order parabolic: no characteristics

Thermodynamic: steepest descent for perimeter

Goal, for next 20 minutes

These evolutions are actually very similar

- Surprising: parabolic pde are very different from HJ eqns.
- Or maybe not: level-set methods used in both settings.
- Argument provides something like characteristics.
- Anlaysis extends to other geometric motions.

Work with Sylvia Serfaty. (Related work on similar lines, focused on applications to image processing: Catté, Dibos, Koepfler; Guichard; Cao; Pasquignon.)

The game

Paul wants to exit; Carol wants to stop him; step size is $\varepsilon > 0$

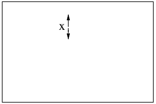

- Paul chooses dirn $|v| = 1$
- Carol may reverse it $b = \pm 1$
- Paul goes $x \to x + \sqrt{2}\varepsilon b v$

Can Paul exit? Yes!

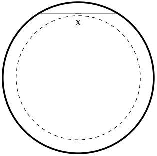

If domain is B_R, he can exit in one step from ball of radius approx $R - \varepsilon^2/R$

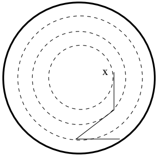

Sets from which Paul can exit in j steps shrink with velocity $1/R_j$ (after normalization)

Carol can't stop Paul, but she can certainly slow him down

Some properties

Geometric interpretation

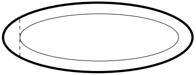

Set from which Paul can exit in one step is traced by midpoints of secants of length $2\sqrt{2}\varepsilon$

Nonconvex case

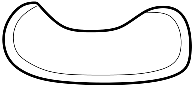

Same as above, but he cannot exit from concave part of $\partial\Omega$

He can exit in $O(\varepsilon^{-2})$ steps

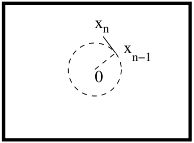

Trial strategy: use $v \perp x$. Then $|x_{k+1}|^2 = |x_k|^2 + 2\varepsilon^2$.

Joel Spencer introduced this "pusher-chooser" game in 1986, as heuristic for certain combinatorial problems.

Going beyond pictures

Value function:

$$u_\varepsilon(x) = \varepsilon^2 \text{ [min number steps to exit, starting at x]}$$

Simplest theorem: for convex plane domains,

$$lim_{\varepsilon \to 0} u_\varepsilon = \text{arrival time of motion by curvature}$$

Main tool is dynamic programming principle:

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

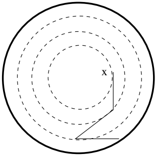

Level-set PDE is the HJB eqn assoc with this game

Formal derivation of Hamilton-Jacobi-Bellman equation:

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

(1) Ignore dependence on ε and use Taylor expansion:

$$u(x) \approx \min_{|v|=1} \max_{b=\pm 1} \left\{ u(x) + \sqrt{2}\varepsilon b v \cdot \nabla u + \varepsilon^2 \langle D^2 u\, v, v \rangle + \varepsilon^2 \right\}$$

(2) Order ε term dominates unless Paul chooses $v \cdot \nabla u = 0$. Thus:

$$u(x) \approx \min_{v \perp \nabla u} \left\{ u(x) + \varepsilon^2 \langle D^2 u\, v, v \rangle + \varepsilon^2 \right\}$$

(3) Simplify:

$$\langle D^2 u \cdot \frac{\nabla u^\perp}{|\nabla u|}, \frac{\nabla u^\perp}{|\nabla u|} \rangle + 1 = 0$$

Equivalent in R^2 to arrival-time formulation of motion by curvature:

$$|\nabla u| \operatorname{div}\left(\nabla u / |\nabla u|\right) + 1 = 0$$

PDE is 2nd order since 1st order expansion was insufficient

Level-set PDE is the HJB eqn assoc with this game

Formal derivation of Hamilton-Jacobi-Bellman equation:

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

(1) Ignore dependence on ε and use Taylor expansion:

$$u(x) \approx \min_{|v|=1} \max_{b=\pm 1} \left\{ u(x) + \sqrt{2}\varepsilon b v \cdot \nabla u + \varepsilon^2 \langle D^2 u\, v, v \rangle + \varepsilon^2 \right\}$$

(2) Order ε term dominates unless Paul chooses $v \cdot \nabla u = 0$. Thus:

$$u(x) \approx \min_{v \perp \nabla u} \left\{ u(x) + \varepsilon^2 \langle D^2 u\, v, v \rangle + \varepsilon^2 \right\}$$

(3) Simplify:

$$\langle D^2 u \cdot \frac{\nabla u^\perp}{|\nabla u|}, \frac{\nabla u^\perp}{|\nabla u|} \rangle + 1 = 0$$

Equivalent in R^2 to arrival-time formulation of motion by curvature:

$$|\nabla u| \operatorname{div}(\nabla u / |\nabla u|) + 1 = 0$$

PDE is 2nd order since 1st order expansion was insufficient

Level-set PDE is the HJB eqn assoc with this game

Formal derivation of Hamilton-Jacobi-Bellman equation:

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

(1) Ignore dependence on ε and use Taylor expansion:

$$u(x) \approx \min_{|v|=1} \max_{b=\pm 1} \left\{ u(x) + \sqrt{2}\varepsilon b v \cdot \nabla u + \varepsilon^2 \langle D^2 u\, v, v \rangle + \varepsilon^2 \right\}$$

(2) Order ε term dominates unless Paul chooses $v \cdot \nabla u = 0$. Thus:

$$u(x) \approx \min_{v \perp \nabla u} \left\{ u(x) + \varepsilon^2 \langle D^2 u\, v, v \rangle + \varepsilon^2 \right\}$$

(3) Simplify:

$$\langle D^2 u \cdot \frac{\nabla u^\perp}{|\nabla u|}, \frac{\nabla u^\perp}{|\nabla u|} \rangle + 1 = 0$$

Equivalent in R^2 to arrival-time formulation of motion by curvature:

$$|\nabla u| \mathrm{div}\left(\nabla u / |\nabla u| \right) + 1 = 0$$

PDE is 2nd order since 1st order expansion was insufficient

Paul's paths are like characteristics

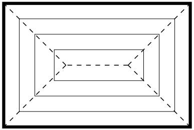

$|\nabla u| = 1, \; u = 0 \text{ at } \partial\Omega$

chars are shortest paths to bdry
PDE becomes ODE along path

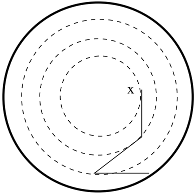

$u_\varepsilon = \varepsilon^2$[# steps to exit]

value increases by ε^2 each time
Paul takes a step

Selected extensions

Same game in R^3?

Boundary moves with velocity = largest principal curvature

Can we get v = mean curvature?

Yes, with a modified game. In R^3,

- Paul chooses two orthog dirns $|v| = |w| = 1$, $v \perp w$
- Carol may reverse either (or both), $b = \pm 1$, $\beta = \pm 1$
- Paul goes $x \to x + \sqrt{2}\varepsilon b v + \sqrt{2}\varepsilon \beta w$

What if Ω is not convex?

$\lim_{\varepsilon \to 0} u_\varepsilon$ is arrival time of flow with velocity κ_+.

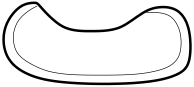

A related stochastic game

Consider the stochastic game in which

- Carol just flips coins
- Paul seeks to exit with prob one in min time

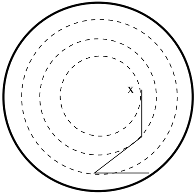

Then Paul's optimal strategy is the same.

Actually, stochastic version was studied first (in cont's time): Buckdawn, Cardaliaguet & Quincampoix; Soner & Touzi

The analogous final-time game

To get a time-dependent PDE, consider our deterministic game with a fixed final-time T. Same rules but Paul has a different goal:

$$v_\varepsilon(x,t) = \min_{\text{starting from x at t}} \Phi(y_\varepsilon(T))$$

$$\begin{aligned} y_\varepsilon(s) &= \text{Paul's path} \\ \Phi &= \text{Paul's objective} \end{aligned}$$

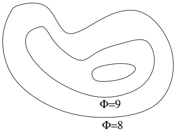

Associated HJB equation is

$$v_t + \langle D^2 v \cdot \frac{\nabla v^\perp}{|\nabla v|}, \frac{\nabla v^\perp}{|\nabla v|} \rangle = 0 \quad \text{for } t < T$$

with $v = \Phi$ at $t = T$.

Each level set of v moves by curvature (backward in time)

Two paths to rigorous analysis

Focus on original minimum-time problem. Two alternatives:

(1) Viscosity solutions

$$\lim_{\varepsilon \to 0} u_\varepsilon = \text{unique viscosity solution of 2nd order PDE}$$

- very general
- gives convergence but no rate
- uses uniqueness of viscosity solution

(2) Verification argument

$$u(x) - C\varepsilon \leq u_\varepsilon(x) \leq u(x) + C\varepsilon$$

- less general – requires u to be C^3
- stronger result (linear rate)
- elementary (each bound proved by considering one player's optimal strategy)

Transition to Part 2

Obvious questions:

- Is this idea limited to geometric problems?
- Or might every 2nd order PDE have a deterministic game interpretation?
- For example, what about the linear heat equation . . .

Thanks to Soner for a crucial hint.

What about the linear heat equation?

A deterministic game for the 1D linear heat equation:

$$u_t + u_{xx} = 0 \text{ for } t < T, \quad u = \Phi \quad \text{at } t = T.$$

Paul's initial position is x at time t.

- Paul chooses $\alpha \in R$, then Carol chooses $b = \pm 1$
- Paul moves $x \to x + \sqrt{2}\varepsilon b$ and pays penalty $\sqrt{2}\varepsilon\alpha b$
- clock steps forward ε^2
- at final time T Paul gets bonus $\Phi(x(T))$

Paul's value function: $u_\varepsilon(x, t) = \max\{\text{bonus - accumulated penalty}\}$

Dyn prog: $u_\varepsilon(x, t) = \max_{\alpha \in R} \min_{b=\pm 1} u_\varepsilon(x + \sqrt{2}\varepsilon b, t + \varepsilon^2) - \sqrt{2}\varepsilon\alpha b$

Formal HJB eqn: $0 = \max_{\alpha \in R} \min_{b=\pm 1} \sqrt{2}\varepsilon b(u_x - \alpha) + \varepsilon^2(u_t + u_{xx})$

So optimal $\alpha = u_x$ and limiting value function solves $u_t + u_{xx} = 0$.

What about the linear heat equation?

A deterministic game for the 1D linear heat equation:

$$u_t + u_{xx} = 0 \text{ for } t < T, \quad u = \Phi \quad \text{at } t = T.$$

Paul's initial position is x at time t.

- Paul chooses $\alpha \in R$, then Carol chooses $b = \pm 1$
- Paul moves $x \to x + \sqrt{2}\varepsilon b$ and pays penalty $\sqrt{2}\varepsilon\alpha b$
- clock steps forward ε^2
- at final time T Paul gets bonus $\Phi(x(T))$

Paul's value function: $u_\varepsilon(x,t) = \max\{\text{bonus - accumulated penalty}\}$

Dyn prog: $u_\varepsilon(x,t) = \max_{\alpha\in R} \min_{b=\pm 1} u_\varepsilon(x + \sqrt{2}\varepsilon b, t + \varepsilon^2) - \sqrt{2}\varepsilon\alpha b$

Formal HJB eqn: $0 = \max_{\alpha\in R} \min_{b=\pm 1} \sqrt{2}\varepsilon b(u_x - \alpha) + \varepsilon^2(u_t + u_{xx})$

So optimal $\alpha = u_x$ and limiting value function solves $u_t + u_{xx} = 0$.

What about the linear heat equation?

A deterministic game for the 1D linear heat equation:

$$u_t + u_{xx} = 0 \text{ for } t < T, \quad u = \Phi \quad \text{at } t = T.$$

Paul's initial position is x at time t.

- Paul chooses $\alpha \in R$, then Carol chooses $b = \pm 1$
- Paul moves $x \to x + \sqrt{2}\varepsilon b$ and pays penalty $\sqrt{2}\varepsilon\alpha b$
- clock steps forward ε^2
- at final time T Paul gets bonus $\Phi(x(T))$

Paul's value function: $u_\varepsilon(x,t) = \max\{\text{bonus - accumulated penalty}\}$

Dyn prog: $u_\varepsilon(x,t) = \max_{\alpha\in R} \min_{b=\pm 1} u_\varepsilon(x + \sqrt{2}\varepsilon b, t + \varepsilon^2) - \sqrt{2}\varepsilon\alpha b$

Formal HJB eqn: $0 = \max_{\alpha\in R} \min_{b=\pm 1} \sqrt{2}\varepsilon b(u_x - \alpha) + \varepsilon^2(u_t + u_{xx})$

So optimal $\alpha = u_x$ and limiting value function solves $u_t + u_{xx} = 0$.

What about the linear heat equation?

A deterministic game for the 1D linear heat equation:

$$u_t + u_{xx} = 0 \text{ for } t < T, \quad u = \Phi \quad \text{at } t = T.$$

Paul's initial position is x at time t.

- Paul chooses $\alpha \in R$, then Carol chooses $b = \pm 1$
- Paul moves $x \to x + \sqrt{2}\varepsilon b$ and pays penalty $\sqrt{2}\varepsilon\alpha b$
- clock steps forward ε^2
- at final time T Paul gets bonus $\Phi(x(T))$

Paul's value function: $u_\varepsilon(x, t) = \max\{\text{bonus - accumulated penalty}\}$

Dyn prog: $u_\varepsilon(x, t) = \max_{\alpha \in R} \min_{b=\pm 1} u_\varepsilon(x + \sqrt{2}\varepsilon b, t + \varepsilon^2) - \sqrt{2}\varepsilon\alpha b$

Formal HJB eqn: $0 = \max_{\alpha \in R} \min_{b=\pm 1} \sqrt{2}\varepsilon b(u_x - \alpha) + \varepsilon^2(u_t + u_{xx})$

So optimal $\alpha = u_x$ and limiting value function solves $u_t + u_{xx} = 0$.

What about R^n? Fully nonlinear PDE?

A slightly different approach extends to fully nonlinear PDE in R^n:

$$u_t + f(Du, D^2u) = 0 \text{ for } t < T, \quad u = \Phi \quad \text{at } t = T$$

assuming parabolicity: $f(p, \Gamma) \leq f(p, \Gamma')$ if $\Gamma \leq \Gamma'$.

Paul's initial position is again x at time t.

- Paul chooses $p \in R^n$ and $\Gamma \in R^{n \times n}_{\rm sym}$, then Carol chooses $w \in R^n$
- Paul moves $x \to x + \varepsilon w$ and pays $\varepsilon p \cdot w + \frac{\varepsilon^2}{2}\langle \Gamma w, w\rangle - \varepsilon^2 f(p, \Gamma)$
- clock steps forward ε^2; Paul gets bonus $\Phi(x(T))$ at final time T

Paul's value function: $u_\varepsilon(x, t) = \max\{\text{bonus - accumulated penalty}\}$

Formal HJB eqn: usual argument gives

$$0 = \max_{p,\Gamma} \min_{w} \varepsilon w \cdot (\nabla u - p) + \varepsilon^2 \left(\tfrac{1}{2}\langle (D^2u - \Gamma)w, w\rangle + f(p, \Gamma) + u_t\right)$$

Paul prefers $p = Du$ and $\Gamma \leq D^2u$ to neutralize w-terms. He prefers $\Gamma = D^2u$ by parabolicity. So value function solves $u_t + f(Du, D^2u) = 0$

What about R^n? Fully nonlinear PDE?

A slightly different approach extends to fully nonlinear PDE in R^n:

$$u_t + f(Du, D^2u) = 0 \text{ for } t < T, \quad u = \Phi \quad \text{at } t = T$$

assuming parabolicity: $f(p, \Gamma) \leq f(p, \Gamma')$ if $\Gamma \leq \Gamma'$.

Paul's initial position is again x at time t.

- Paul chooses $p \in R^n$ and $\Gamma \in R^{n\times n}_{\text{sym}}$, then Carol chooses $w \in R^n$
- Paul moves $x \to x + \varepsilon w$ and pays $\varepsilon p \cdot w + \frac{\varepsilon^2}{2}\langle \Gamma w, w\rangle - \varepsilon^2 f(p, \Gamma)$
- clock steps forward ε^2; Paul gets bonus $\Phi(x(T))$ at final time T

Paul's value function: $u_\varepsilon(x, t) = \max\{\text{bonus - accumulated penalty}\}$

Formal HJB eqn: usual argument gives

$$0 = \max_{p,\Gamma} \min_{w} \varepsilon w \cdot (\nabla u - p) + \varepsilon^2 \left(\tfrac{1}{2}\langle (D^2u - \Gamma)w, w\rangle + f(p, \Gamma) + u_t\right)$$

Paul prefers $p = Du$ and $\Gamma \leq D^2u$ to neutralize w-terms. He prefers $\Gamma = D^2u$ by parabolicity. So value function solves $u_t + f(Du, D^2u) = 0$

What about R^n? Fully nonlinear PDE?

A slightly different approach extends to fully nonlinear PDE in R^n:

$$u_t + f(Du, D^2 u) = 0 \text{ for } t < T, \quad u = \Phi \quad \text{at } t = T$$

assuming parabolicity: $f(p, \Gamma) \le f(p, \Gamma')$ if $\Gamma \le \Gamma'$.

Paul's initial position is again x at time t.

- Paul chooses $p \in R^n$ and $\Gamma \in R^{n\times n}_{\mathrm{sym}}$, then Carol chooses $w \in R^n$
- Paul moves $x \to x + \varepsilon w$ and pays $\varepsilon p \cdot w + \frac{\varepsilon^2}{2}\langle \Gamma w, w\rangle - \varepsilon^2 f(p, \Gamma)$
- clock steps forward ε^2; Paul gets bonus $\Phi(x(T))$ at final time T

Paul's value function: $u_\varepsilon(x, t) = \max\{\text{bonus - accumulated penalty}\}$

Formal HJB eqn: usual argument gives

$$0 = \max_{p,\Gamma} \min_{w} \varepsilon w \cdot (\nabla u - p) + \varepsilon^2 \left(\tfrac{1}{2}\langle (D^2 u - \Gamma) w, w\rangle + f(p, \Gamma) + u_t\right)$$

Paul prefers $p = Du$ and $\Gamma \le D^2 u$ to neutralize w-terms. He prefers $\Gamma = D^2 u$ by parabolicity. So value function solves $u_t + f(Du, D^2 u) = 0$

What about R^n? Fully nonlinear PDE?

A slightly different approach extends to fully nonlinear PDE in R^n:

$$u_t + f(Du, D^2u) = 0 \text{ for } t < T, \quad u = \Phi \quad \text{at } t = T$$

assuming parabolicity: $f(p, \Gamma) \leq f(p, \Gamma')$ if $\Gamma \leq \Gamma'$.

Paul's initial position is again x at time t.

- Paul chooses $p \in R^n$ and $\Gamma \in R^{n \times n}_{\text{sym}}$, then Carol chooses $w \in R^n$
- Paul moves $x \to x + \varepsilon w$ and pays $\varepsilon p \cdot w + \frac{\varepsilon^2}{2}\langle \Gamma w, w\rangle - \varepsilon^2 f(p, \Gamma)$
- clock steps forward ε^2; Paul gets bonus $\Phi(x(T))$ at final time T

Paul's value function: $u_\varepsilon(x, t) = \max\{\text{bonus - accumulated penalty}\}$

Formal HJB eqn: usual argument gives

$$0 = \max_{p,\Gamma} \min_{w} \varepsilon w \cdot (\nabla u - p) + \varepsilon^2 \left(\tfrac{1}{2}\langle (D^2 u - \Gamma)w, w\rangle + f(p, \Gamma) + u_t\right)$$

Paul prefers $p = Du$ and $\Gamma \leq D^2 u$ to neutralize w-terms. He prefers $\Gamma = D^2 u$ by parabolicity. So value function solves $u_t + f(Du, D^2 u) = 0$

What's going on?

- Our games are semi-discrete numerical schemes, whose time-step problem is a max-min.
- Scheme for $u_t + f(Du, D^2u) = 0$ is like explicit Euler. Min over w picks out p and Γ as proxies for Du and D^2u, even if u is not differentiable.
- Not recommended for linear heat equation. But perhaps useful for nonlinear PDE whose solutions are not smooth.
- Paul behaves optimally $\Rightarrow$ he's indifferent to Carol's choices. She might as well choose randomly. Related to representation formulas via backward stochastic differential equations (Cheredito, Soner, Touzi, Victoir).

Conclusion

- New viewpoint on motion by curvature and related PDE's (Kohn-Serfaty, CPAM 2006)
- Semidiscrete numerical scheme for fully nonlinear PDE's (Kohn-Serfaty, in progress)