A Constructive Approach to Pursuit-Evasion Games

Emiliano Cristiani[1]

Maurizio Falcone [2]

[1]CEMSAC - Salerno & IAC-CNR Roma
cristiani@mat.uniroma1.it

[2]SAPIENZA - Università di Roma
falcone@mat.uniroma1.it

EPSRC Workshop "Game Theory for Finance, Social and Biological Sciences"
Warwick, April, 14–17, 2010

Outline

Introduction

Our approach is based on the Dynamic Programming Principle which allows to derive a nonlinear first order partial differential equation describing the value function of the game (or the upper/lower value function).

The theory of viscosity solutions gives the correct framework to characterize the value function (or its upper/lower version) as the unique weak solution of the Isaacs equation.

This characterization has been used also to construct numerical schemes for the value function and to synthetize optimal feedbacks.

Introduction

Our approach is based on the Dynamic Programming Principle which allows to derive a nonlinear first order partial differential equation describing the value function of the game (or the upper/lower value function).

The theory of viscosity solutions gives the correct framework to characterize the value function (or its upper/lower version) as the unique weak solution of the Isaacs equation.

This characterization has been used also to construct numerical schemes for the value function and to synthetize optimal feedbacks.

Introduction

Our approach is based on the Dynamic Programming Principle which allows to derive a nonlinear first order partial differential equation describing the value function of the game (or the upper/lower value function).

The theory of viscosity solutions gives the correct framework to characterize the value function (or its upper/lower version) as the unique weak solution of the Isaacs equation.

This characterization has been used also to construct numerical schemes for the value function and to synthetize optimal feedbacks.

Pursuit-Evasion Games without State Constraints

Hamilton-Jacobi-Isaacs equation

$$v(x) + \min_{b \in B} \max_{a \in A} \{-f(x,a,b) \cdot \nabla v(x)\} - 1 = 0 \quad x \in \overline{\mathbb{R}^n} \backslash \mathcal{T}$$

where A, B are the admissible controls for the first and for the second player, f is the vectorfield and $\mathcal{T}$ is the (given) target set.

Typically the HJI equation is complemented with Dirichlet boundary conditions

$$v(x) = g(x) \quad \text{for } x \in \partial\mathcal{T}$$

Pursuit-Evasion Games without State Constraints

Uniqueness of viscosity solutions holds under rather general assumptions (Evans-Souganidis, Bardi). Another approach have been proposed by A. Subbotin (minmax solutions). This approach has been used successfully for problems without state constraints and has produced accurate results for 2-dimensional games.

However its extension to games with state constraints is non trivial and few results are available, mainly via the viability approach proposed by Aubin (viability kernel solutions).

Pursuit-Evasion games with state constraints in $\mathbb{R}^{2N}$

$$\begin{cases} \dot{y}(t) = f(y(t), a(t), b(t)), & t > 0 \\ y(0) = x \end{cases}$$

where $x = (x_P, x_E)$, $\quad y = (y_P, y_E)$, $\quad a \in A, b \in B$

$$f(x, a, b) = f(x_P, x_E, a, b) = \begin{pmatrix} f_P(x_P, a) \\ f_E(x_E, b) \end{pmatrix}, \qquad f_P, f_E \in \mathbb{R}^N$$

State Constraints

$y_P(t) \in \overline{\Omega}_1, \quad y_E(t) \in \overline{\Omega}_2$

Target

$\mathcal{T} = \{(x_P, x_E) \in \mathbb{R}^{2N} : |x_P - x_E| \leq \varepsilon\}, \quad \varepsilon \geq 0$

Pursuit-Evasion games with state constraints in $\mathbb{R}^{2N}$

$$\begin{cases} \dot{y}(t) = f(y(t), a(t), b(t)), & t > 0 \\ y(0) = x \end{cases}$$

where $x = (x_P, x_E)$, $\quad y = (y_P, y_E)$, $\quad a \in A, b \in B$

$$f(x, a, b) = f(x_P, x_E, a, b) = \begin{pmatrix} f_P(x_P, a) \\ f_E(x_E, b) \end{pmatrix}, \qquad f_P, f_E \in \mathbb{R}^N$$

State Constraints

$y_P(t) \in \overline{\Omega}_1, \quad y_E(t) \in \overline{\Omega}_2$

Target

$\mathcal{T} = \{(x_P, x_E) \in \mathbb{R}^{2N} : |x_P - x_E| \leq \varepsilon\}, \quad \varepsilon \geq 0$

Pursuit-Evasion games with state constraints in $\mathbb{R}^{2N}$

$$\begin{cases} \dot{y}(t) = f(y(t), a(t), b(t)), & t > 0 \\ y(0) = x \end{cases}$$

where $x = (x_P, x_E)$, $\quad y = (y_P, y_E)$, $\quad a \in A, b \in B$

$$f(x, a, b) = f(x_P, x_E, a, b) = \begin{pmatrix} f_P(x_P, a) \\ f_E(x_E, b) \end{pmatrix}, \qquad f_P, f_E \in \mathbb{R}^N$$

State Constraints

$y_P(t) \in \overline{\Omega}_1, \quad y_E(t) \in \overline{\Omega}_2$

Target

$\mathcal{T} = \{(x_P, x_E) \in \mathbb{R}^{2N} : |x_P - x_E| \leq \varepsilon\}, \quad \varepsilon \geq 0$

The Tag-Chase game with state constraints

Two boys P and E which run one after the other in the same 2-dimensional domain (courtyard), so that the game is set in $\overline{\Omega} = \overline{\Omega}_1^2 \subset \mathbb{R}^4$.
P and E can run in every direction with velocity V_P and V_E respectively.

$$\begin{cases} \dot{x}_P = V_P \, a & a \in B_2(0,1) \\ \dot{x}_E = V_E \, b & b \in B_2(0,1) \end{cases}$$

A simple example

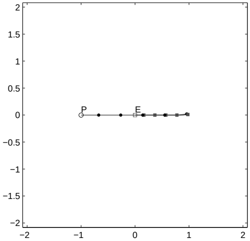

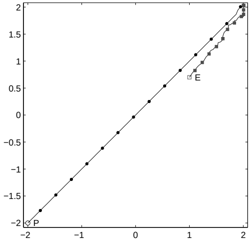

Optimal exploitation of natural resources (Jorgensen)

Let us consider a common property fishery.
$m(t)$ = biomass of a particular stock of fish at time t
Assume that the stock is harvested by two fishermen A and B.
Fish has a natural growth, which follows the logistic function.
State dynamics

$$\dot{m}(t) = m(1-m) - (\alpha(t) + \beta(t))$$

α and β are the controls of the two players, which correspond to their harvest rate.
We impose two two natural constraints:

$$0 \leq \alpha(t), \beta(t) \leq C_{max} \text{ and } m(t) \geq 0 \text{ for any } t$$

Under these hypotheses, we have $m \in [0,1]$ for any t.

Optimal allocation of natural resources

Each fisherman wants to maximize his payoff

$$J_A(\alpha) = \int_0^\infty \left(-\frac{1}{\alpha(s)}\right) e^{-\lambda s} ds\,, \quad J_B(\beta) = \int_0^\infty \left(-\frac{1}{\beta(s)}\right) e^{-\lambda s} ds$$

where $\lambda \in (0,1)$.

Every fisherman would like to choose its harvest rate as large as possible, but he must be careful not to drive the stock to zero because then he has to stop fishing.

WARNING: for $m = 0$, the constraint $m \geq 0$ imposes $\alpha = 0$ and $\beta = 0$.

Isaacs equation

$T(x) :=$ capture time under optimal (non-anticipating) strategies of both players.

THEOREM (Koike, 1995)

$v(x) = 1 - e^{-T(x)}$ is the unique viscosity solution of

$$\begin{cases} v(x) + \min\limits_{b \in B(x)} \max\limits_{a \in A(x)} \{-f(x,a,b) \cdot \nabla v(x)\} - 1 = 0 & x \in \overline{\Omega} \backslash \mathcal{T} \\ v(x) = 0 & \text{on } \partial\mathcal{T} \end{cases}$$

For $A(x) = A$ and $B(x) = B$

$v(x) = 1 - e^{-T(x)}$ is the unique viscosity solution of

Hamilton-Jacobi-Isaacs equation

$$\begin{cases} v(x) + \min_{b \in B} \max_{a \in A}\{-f(x,a,b) \cdot \nabla v(x)\} - 1 = 0, \quad x \in \overline{\Omega} \backslash \mathcal{T} \\ v(x) = 0 \text{ on } \partial\mathcal{T} \end{cases}$$

Fully discrete scheme (without constraints)

Let the constraint be given by $\Omega = \Omega_1 \cap \Omega_2$.
We define

$$\begin{aligned}
\beta &= e^{-h} \\
I_{in} &= \{i : x_i \in \Omega \setminus \mathcal{T}\} \\
I_{\mathcal{T}} &= \{i : x_i \in \mathcal{T} \cap \Omega\} \\
I_{out_1} &= \{i : x_i \notin \Omega_2\} \\
I_{out_2} &= \{i : x_i \notin \Omega_2 \setminus \Omega\}
\end{aligned}$$

Fully discrete scheme (without constraints)

The discretization in time and space leads to a fully discrete scheme

$$
\begin{aligned}
w(x_i) =& \max_{b \in B} \min_{a \in A} [\beta w(x_i + hf(x_i, a, b))] + 1 - \beta \quad \text{for } i \in I_{in} \\
w(x_i) =& 1 \qquad \text{for } i \in I_{out_2} \\
w(x_i) =& 0 \qquad \text{for } i \in I_{\mathcal{T}} \cup I_{out_1}
\end{aligned}
$$

Fully discrete scheme (without constraints)

Theorem

The operator $S : [0,1]^L \to [0,1]^L$.
S satisfies the following properties:

$$U \leq V \Rightarrow S(U) \leq S(V)$$

S is a contraction map.

Let U^* be the unique fixed point, we define

$$w(x_i) = U_i^* \qquad \forall i$$
$$w(x) = \sum_j \lambda_{ij}(a,b) w(x_j) \text{(linear interpolation)}$$

Convergence (without constraints)

Naturally w depends on the discretization steps, h and k.

Theorem

Let $\mathcal{T}$ be the closure of an open set with Lipschitz boundary, "$\mathrm{diam}\ \Omega \to +\infty$" and v continuous. Then

$$w^{h,k} \to v \quad \text{on compact sets of } \mathbb{R}^N$$

for $h \to 0^+$ and $\frac{k}{h} \to 0^+$.

Convergence: discontinuous value

Let w_n^ε be the sequence generated by the numerical scheme with target $\mathcal{T}_\varepsilon = \{x : d(x, \mathcal{T}) \leq \varepsilon\}$

Theorem

For all x there exists the limit

$$\overline{w}(x) = \lim_{\substack{\varepsilon \to 0^+ \\ n \to +\infty \\ n \geq n(\varepsilon)}} w_n^\varepsilon(x)$$

and it coincides with the lower value V of the game with target $\mathcal{T}$, i.e.

$$\overline{w} = V$$

The convergence is uniform on every compact set where V is continuous.

Time-discrete scheme for P-E games with SC

Definition: Admissible controls

$$A(y) = \{a \in A : \exists\, r > 0 \text{ such that } y_P(t; y'_P, a) \in \overline{\Omega}_1 \\ \text{for } t \in [0, r] \text{ and } y'_P \in B(y_P, r) \cap \overline{\Omega}_1\},$$

$$B(y) = \{b \in B : \exists\, r > 0 \text{ such that } y_E(t; y'_E, b) \in \overline{\Omega}_2 \\ \text{for } t \in [0, r] \text{ and } y'_E \in B(y_E, r) \cap \overline{\Omega}_2\}.$$

$A(x)$ and $B(x)$ are the admissible controls sets at x w.r. t. the constraints.

Time-discrete scheme for P-E games with SC

Definition: Admissible controls (discrete time version)

$$A_h(x) := \left\{a \in A : x_P + hf_P(x_P, a) \in \overline{\Omega}_1\right\}, \qquad x \in \overline{\Omega}$$

$$B_h(x) := \left\{b \in B : x_E + hf_E(x_E, b) \in \overline{\Omega}_2\right\}, \qquad x \in \overline{\Omega}.$$

$$\begin{cases} v_h(x) = \max\limits_{b \in B_h(x)} \min\limits_{a \in A_h(x)} \{\beta v_h(x + hf(x,a,b))\} + 1 - \beta & x \in \overline{\Omega} \setminus \mathcal{T} \\ v_h(x) = 0 & x \in \mathcal{T} \end{cases}$$

where $\beta = e^{-h}$.

Fully-discrete scheme for P-E games with SC

We build a regular triangulation of $\overline{\Omega}$ denoting by:
X the set of its nodes x_i, $i = 1, \dots, L$
$k := \max_j \{diam(S_j), S_j simplexofthetriangulation\}$

$$\begin{cases} v_h^k(x_i) = \max_{b \in B_h(x_i)} \min_{a \in A_h(x_i)} \{\beta v_h^k(x_i + hf(x_i, a, b))\} + 1 - \beta & x_i \in (X \backslash \mathcal{T}) \\ v_h^k(x_i) = 0 & x_i \in \mathcal{T} \cap X \\ \\ v_h^k(x) = \sum_j \lambda_j(x) v_h^k(x_j)\,, \quad 0 \le \lambda_j(x) \le 1\,, \quad \sum_j \lambda_j(x) = 1 & x \in \overline{\Omega} \end{cases}$$

Discrete Reachable Sets

Definition

$\mathcal{R}_0 := \mathcal{T}$

$$\mathcal{R}_n := \left\{ x \in \overline{\Omega} \setminus \bigcup_{j=0}^{n-1} \mathcal{R}_j : \text{ for all } b \in B_h(x) \text{ there exists} \right.$$

$$\left. \hat{a}_x(b) \in A_h(x) \text{ such that } x + hf(x, \hat{a}_x(b), b) \in \mathcal{R}_{n-1} \right\}, \quad n \geq 1.$$

Main result

Theorem

Let Ω be an open bounded set, f be continuous and Lipschitz continuous w.r. t. x. Assume P always reaches E and let $\min_{x,a,b} |f(x,a,b)| \geq f_0 > 0$ and $0 < k \leq f_0 h$. Then, we have:

a) $v_h(x) \leq v_h(y)$, for any $x \in \bigcup_{j=0}^{n} \mathcal{R}_j$, for any $y \in \overline{\Omega} \setminus \bigcup_{j=0}^{n} \mathcal{R}_j$;

b) $v_h(x) = 1 - e^{-nh}$, for any $x \in \mathcal{R}_n$;

c) $v_h^k(x) = 1 - e^{-nh} + O(k) \sum_{j=0}^{n} e^{-jh}$ for any $x \in \mathcal{R}_n$;

d)

$$|v_h(x) - v_h^k(x)| \leq \frac{Ck}{1 - e^{-h}}, \quad \text{for any } x \in \overline{\Omega}.$$

for some positive constant $C > 0$.

Convergence for P-E games with SC

Coupling the two results we can prove that our approximation scheme converges to the value function.

Cristiani, F. (2006)

Under the assumptions of our main result, $v_h^k \longrightarrow v_h$

$$+$$

Bardi, Koike, Soravia (2000)

Under regularity assumptions on the sets of constraints, $v_h \longrightarrow v$

$$\Downarrow$$

Under the above assumptions, $v_h^k \longrightarrow v$ uniformly for $k = O(h^{1+\alpha})$ and $h \rightarrow 0$.

Feedback controls for games

The algorithm computes an approximate optimal control couple (a^*, b^*) at each point of the grid. By w we can also compute an approximate optimal feedback at every point of Q.

$$(a^*(x), b^*(x)) \equiv \text{argminmax}\{e^{-h} w(x + hf(x, a, b))\} + 1 - e^{-h}$$

In case of multiple solutions we can select a unique couple, e.g. minimizing two convex functionals.

We can also introduce an inertia criterium to stabilize the trajectories, i.e. if a at step $n + 1$ the set of optimal couples contains (a_n^*, b_n^*) we keep it.

The Homicidal Chauffeur

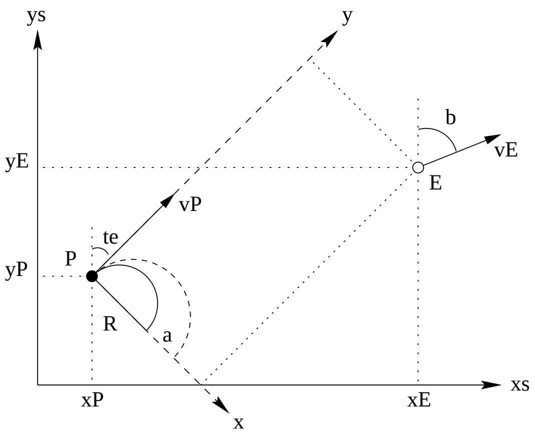

Trajectories 1

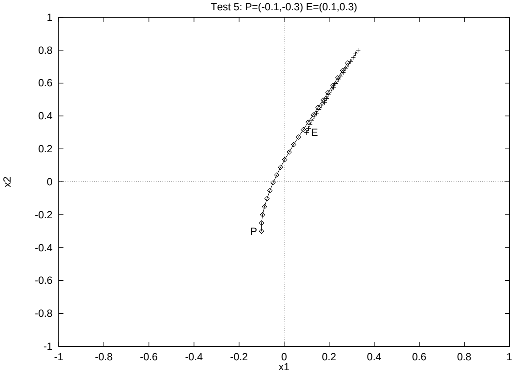

Trajectories 2

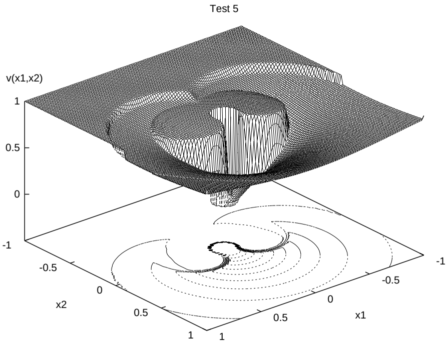

Optimal Trajectories (Merz Thesis)

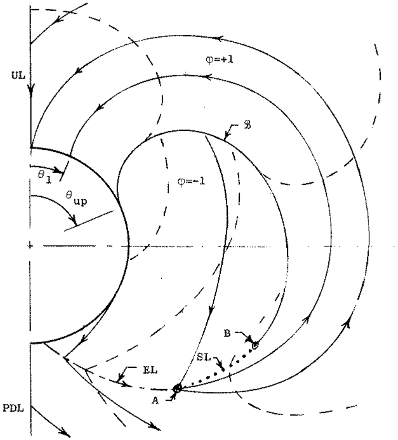

Optimal Trajectories (computed)

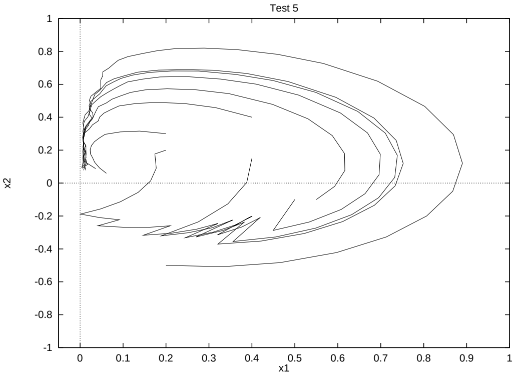

Solution of the fishery game (symmetric)

Considering the symmetry of the problem, one possibility is that they make an agreement and we expect the same optimal strategy for the two players.
In this case we can consider the two players as a unique entity and deal with a standard infinite-horizon optimal control problem.

Solution of the fishery game (symmetric)

We solve numerically the HJB equation

$$\lambda v(x) + \max_{a \in (0, C_{max}]} \{-f(x,a) \cdot \nabla v(x) - \ell(x,a)\} = 0 \,, \quad x \in [0,1]$$

where

$f(x,a) = x(1-x) - a$ and $\ell = 1/a$
the new cost $\widehat{J} = \int_0^\infty \ell(m(s), \alpha(s)) e^{-\lambda s} ds$ must be minimized.

Once the value function v is computed, the optimal control in feedback form and the optimal trajectory are reconstructed

Solution of the fishery game (symmetric)

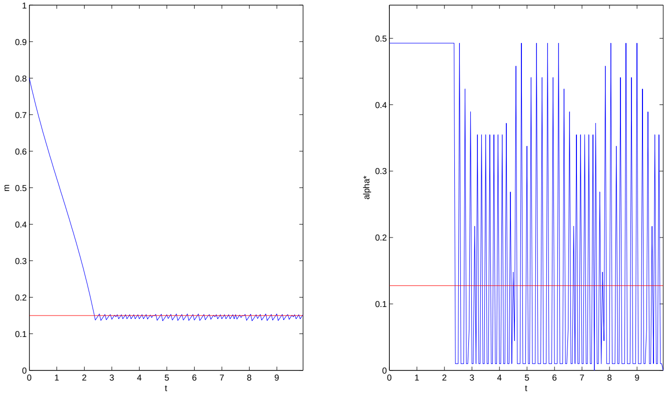

Figure: Optimal trajectory (left) and optimal feedback control (right)

Solution of the fishery game (asymmetric)

In order to avoid the state constraint $m \geq 0$, and we introduce an asymmetry in the problem. we slightly modify the state dynamics

$$\dot{m}(t) = m(1-m) - (\alpha(t) + \beta(t))m(t).$$

Here the two controls represent the "fishing power" used by the two players. We define the cost functional in such a way the problem can be seen as a two-player zero-sum differential game,

$$J(\alpha, \beta) = \int_0^{\infty} e^{-\lambda s}\Big((\beta m - c_B \beta) - (\alpha m - c_A \alpha)\Big) ds$$

with $0 \leq c_A,\ c_B \leq 1$.

Solution of the fishery game (asymmetric)

Player B wants to maximize J, whereas player A wants to minimize it.

βm (resp., αm) is proportional to the gain of player B (resp., player A).

$c_B\beta$ (resp., $c_A\alpha$) represents the cost paid by player B (resp., player A) to go fishing, we assume that this cost is linear w.r.t. the control variable.

Note that when m becomes too small, fishermen have no interest to fish, since the cost overcomes the gain.

Solution of the fishery game (asymmetric)

We solve numerically the HJI equation

$$\lambda v(x) + \min_{b \in [0, C_{max}]} \max_{a \in [0, C_{max}]} \{-f(x,a,b) \cdot \nabla v(x) - \ell(x,a,b)\} = 0$$

for $x \in [0,1]$ where

$$f(x,a) = x(1-x) - (a+b)x$$

$$\ell(x,a,b) = (bx - c_B b) - (ax - c_A a)$$

If $c_A = c_B$, the game is symmetric and the optimal value for J is 0.

If it is not symmetric, we a more interesting situation.

Choosing $c_A = 0.05$, $c_B = 0.025$, $C_{max} = 1$, $\lambda = 0.7$, and $m(0) = 0.6$, the second player (the most efficient one) is able to eliminate player A from the competition, leading the state below the threshold of the positive gain for its competitor.

Solution of the fishery game (asymmetric)

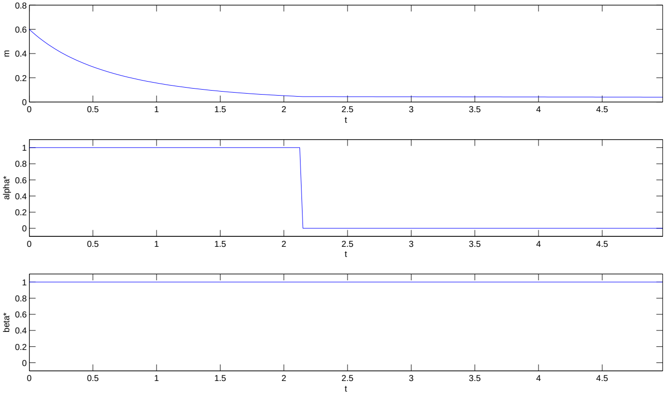

Figure: Optimal trajectory (above) and optimal controls (center and below)

The Tag-Chase game with state constraints

We consider two boys P and E which run one after the other in the same 2-dimensional domain, so that the game is set in $\overline{\Omega} = \overline{\Omega}_1^2 \subset \mathbb{R}^4$. P and E can run in every direction with velocity V_P and V_E respectively.

$$\begin{cases} \dot{x}_P = V_P \, a & \quad a \in B_2(0,1) \\ \dot{x}_E = V_E \, b & \quad b \in B_2(0,1) \end{cases}$$

The case $V_P > V_E$ was studied in [Alziary de Rocquefort, 1991].

Capturability in Tag-Chase game

THEOREM

Let the target be

$$\mathcal{T} = \{(x_P, x_E) \in \mathbb{R}^4 : |x_P - x_E| \leq \varepsilon\}, \qquad \varepsilon \geq 0.$$

and Ω_1 an open bounded set. Then,

- if $V_P > V_E$ then the capture time T is finite and bounded by

$$T(x_P, x_E) \leq \frac{|x_P - x_E|}{V_P - V_E}.$$

- If $V_P = V_E$, $\varepsilon > 0$ and Ω_1 is convex then the capture time T is finite.

The computation of the value function has been done in $\mathbb{R}^4$ since reduced coordinates can not be applied in the constrained problem.

We have exploited the symmetries in the Tag-Chase game and used a fast projection algorithm to interpolate in $\mathbb{R}^4$ (standard linear interpolation is too expensive).

The synthesis of optimal controls in $\overline{\Omega} \backslash \mathcal{T}$ has been computed by the value function as

$$(a^*, b^*) = Arg \max_{b \in B_h(x)} \min_{a \in A_h(x)} \{\beta v_h(x + hf(x, a, b))\} + 1 - \beta$$

Test 1, $V_P > V_E$

$\varepsilon = 10^{-3}$, $V_P = 2$, $V_E = 1$, $n = 50$, $n_c = 48 + 1$. Convergence was reached in 85 iterations. The CPU time (IBM - 8 procs) was 17h 36m 16s, the wallclock time was 2h 47m 37s.

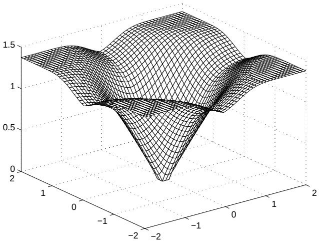

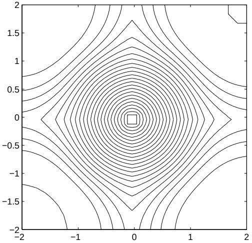

Value function $T(0, 0, x_E, y_E)$.

Test 1, $V_P > V_E$

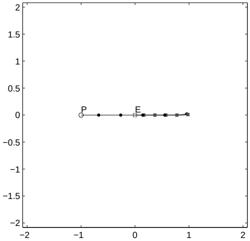

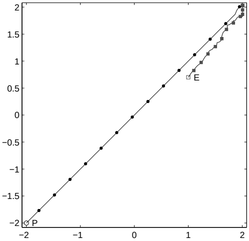

Test 1, $V_P > V_E$

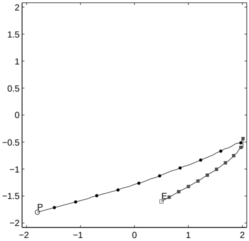

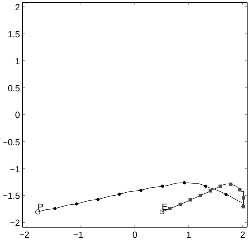

Test 2, $V_P > V_E$

architecture	wallclock time	speed-up	efficiency
IBM serial	26m 47s	-	-
IBM 2 procs	14m 19s	1.87	0.93
IBM 4 procs	8m 09s	3.29	0.82
IBM 8 procs	4m 09s	6.45	0.81
PC dual core, ser	1h 08m 44s	-	-
PC dual core, par	34m 51s	1.97	0.99

$$\text{speed-up} := \frac{T_{ser}}{T_{par}} \qquad \text{efficiency} := \frac{\text{speed-up}}{n_p}.$$

Test 3, $V_P > V_E$

In this test the domain has a square hole in the center. The side of the square is 1.06. $\varepsilon = 10^{-4}$, $V_P = 2$, $V_E = 1$, $n = 50$, $n_c = 48 + 1$. Convergence: 109 iterations. CPU time: 1d 00h 34m 18s, wallclock time: 3h 54m 30s.

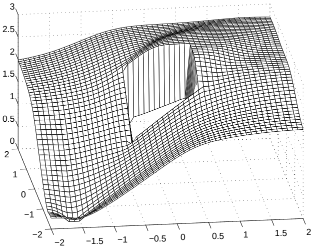

Value function $T(-1.5, -1.5, x_E, y_E)$.

Test 3, $V_P > V_E$

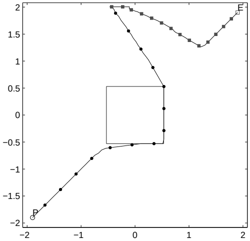

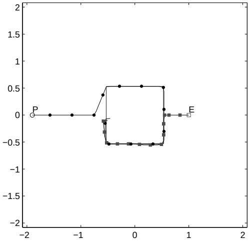

Test 4, $V_P > V_E$

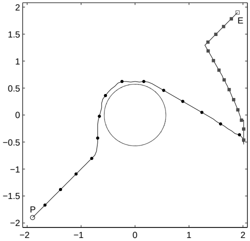

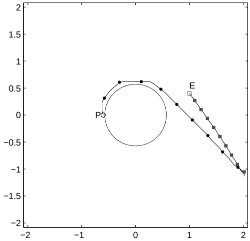

Test 5, $V_P = V_E$

v is discontinuous on $\partial\mathcal{T}$. No convergence results.
$\varepsilon = 10^{-3}$, $V_P = 1$, $V_E = 1$, $n = 50$, $n_c = 36$. Convergence was reached in 66 iterations.

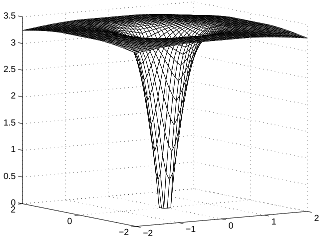

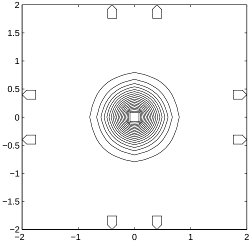

Value function $T(0, 0, x_E, y_E)$.

Test 5, $V_P = V_E$

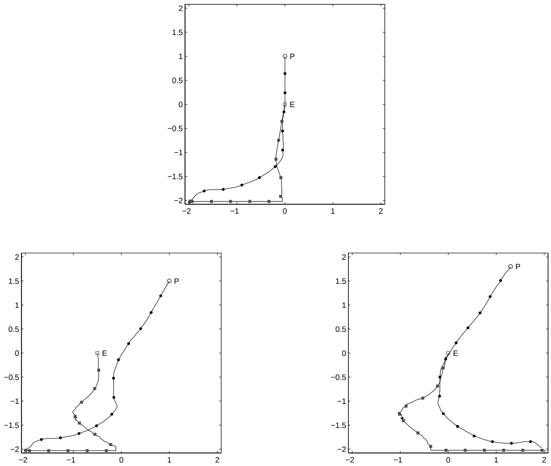

Test 6, $V_P = V_E$

The domain has a circular hole in the center. Non-convex domain, then no guarantee capture occurs.
Since v is equal to 1 in a large part of the domain this produces a strange behavior of some optimal trajectories.
$\varepsilon = 10^{-4}$, $V_P = 1$, $V_E = 1$, $n = 50$, $n_c = 48 + 1$. Convergence: 94 iterations. CPU time: 1d 12h 05m 22s.

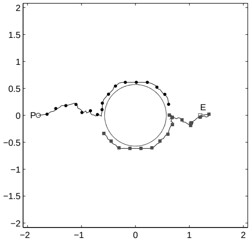

Fast Marching Methods for games

The main idea of Fast Marching method is based on the front propagation point of view.

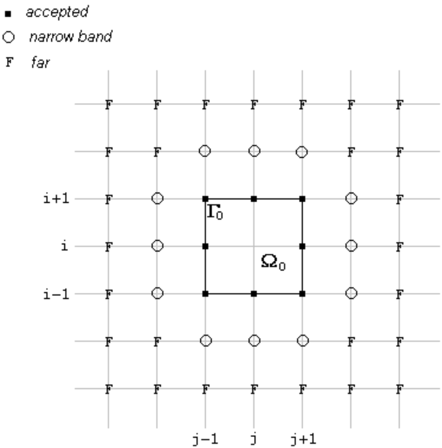

The evolution of the front at every time is given by the level sets of the function $T(x)$ solution of the

Eikonal equation

$$\begin{cases} c(x)|\nabla T(x)| = 1 & x \in \mathbb{R}^n \backslash \Omega \\ T(x) = 0 & x \in \partial\Omega \end{cases} \tag{1}$$

$T(x)$ is the arrival time of the front at x.

FM Algorithm. Test: differential games with state constraints

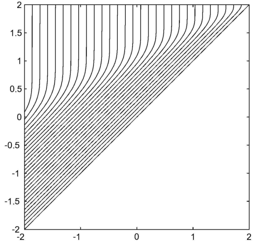

SL iter. for diff. games, grid= 101x101, 130 iter.

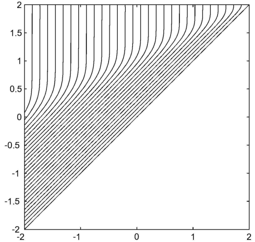

BFM for diff. games, grid= 101x101

SL iterative (left) and BFM (right)

Buffered FM Algorithm for differential games with state constraints

The L^1 error is computed with respect to the solution of the iterative algorithm rather than the exact solution.
2x2 controls, $tol = 10^{-6}$.

method	Δx	L^1 error	CPU time (sec)
SL (70 it)	0.08	-	21.5
BFM	0.08	0.002	9.16
FM-SL	0.08	0.23	1.59
SL (130 it)	0.04	-	157
BFM	0.04	0.002	60
FM-SL	0.04	0.42	6.17