Pursuit-Evasion Problems with Robots

Rafael Murrieta Cid

Centro de Investigación en Matemáticas (CIMAT)

April 2016

Outline

Previous work

Target capturing in an environment without obstacles

- R. Isaacs. Differential Games: A Mathematical Theory with Applications to Warfare and Pursuit, Control and Optimization. John Wiley and Sons. Inc., 1965.
- Y.C. Ho et. al., Differential Games and Optimal Pursuit-Evasion Strategies, IEEE Transactions on Automatic Control, 1965.

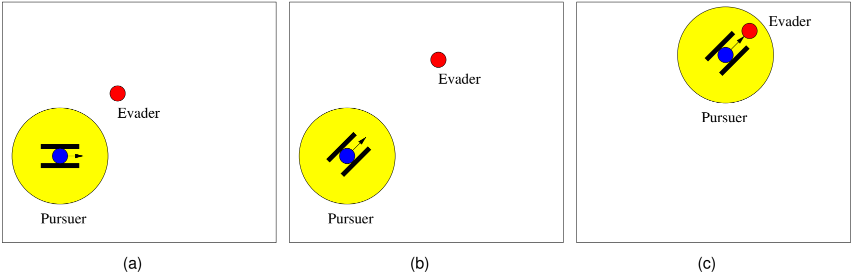

Figure: Target capturing.

Previous work

Target tracking in an environment with obstacles

- S. M. LaValle et. al., ”Motion strategies for maintaining visibility of a moving target”, in Proc. IEEE Int. Conf. on Robotics and Automation, 1997.
- H.H. González-Baños et. al., Motion Strategies for Maintaining Visibility of a Moving Target In Proc IEEE Int. Conf. on Robotics and Automation, 2002.
- R. Murrieta-Cid et. al., Surveillance Strategies for a Pursuer with Finite Sensor Range, International Journal on Robotics Research, Vol. 26, No 3, pages 233-253, March 2007.
- S. Bhattacharya and S. Hutchinson , On the Existence of Nash Equilibrium for a Two Player Pursuit-Evasion Game with Visibility Constraints, The International Journal of Robotics Research, December, 2009.

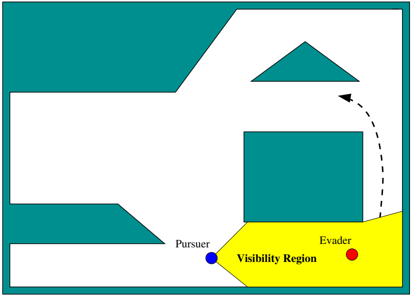

Figure: Target tracking.

Previous work

Target finding in an environment with obstacles

- V. Isler et. al., “Randomized pursuit-evasion in a polygonal enviroment”, IEEE Transactions on Robotics, vol. 5, no. 21, pp. 864-875, 2005.
- R. Vidal et. al., “Probabilistic pursuit-evasion games: Theory, implementation, and experimental evaluation”, IEEE Transactions on Robotics and Automation, vol. 18, no. 5, pp. 662-669, 2002.

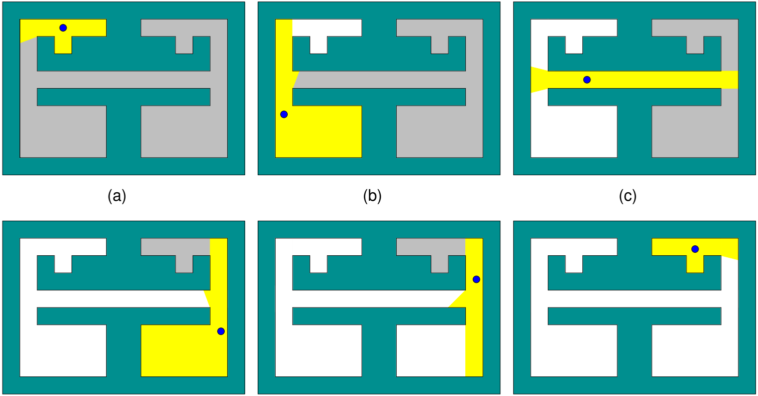

(a) (b) (c) (d) (e) (f)

Figure: Target finding.

Target finding in an environment with obstacles

- L. Guibas, J.-C. Latombe, S. LaValle, D. Lin and R. Motwani, "Visibility-based pursuit-evasion in a polygonal environment", International Journal of Computational Geometry and Applications, vol. 9, no. 4/5, pp. 471-494, 1999.

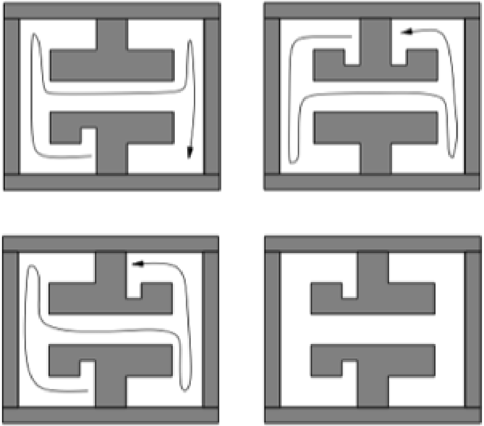

Time-Optimal Motion Strategies for Capturing an Omnidirectional Evader using a Differential Drive Robot

Time-Optimal Motion Strategies for Capturing an Omnidirectional Evader using a Differential Drive Robot

Problem formulation

- A Differential Drive Robot (DDR) and an omnidirectional evader move on a plane without obstacles.
- The game is over when the distance between the DDR and the evader is smaller than a critical value l.
- Both players have maximum bounded speeds $V_p^{\max}$ and $V_e^{\max}$, respectively. The DDR is faster than the evader, $V_p^{\max} > V_e^{\max}$.
- The DDR wants to minimize the capture time t_f while the evader wants to maximize it.
- We want to know the time-optimal motion strategies of the players that are in Nash Equilibrium.

The Homicidal Chauffeur Problem

- A driver wants to run over a pedestrian in a parking lot without obstacles.
- The pursuer is a vehicle with a minimal turning radius (car-like).
- The question to be solved is: under what circumstances, and with what strategy, can the driver of the car guarantee that he can always catch the pedestrian, or the pedestrian guarantee that he can indefinitely elude the car?

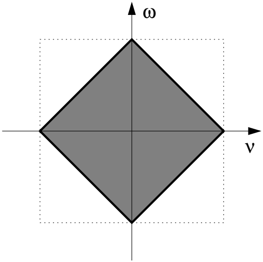

(a) DDR

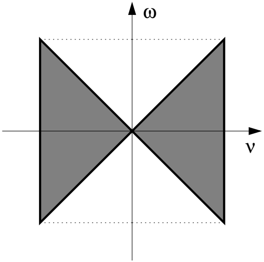

(b) Car-like

Figure: Control domains

Model

Reduced space

The problem can be stated in a coordinate system that is fixed to the body of the DDR. The state of the system is expressed as $\mathbf{x}(t) = (x(t), y(t)) \in \mathbb{R}^2$.

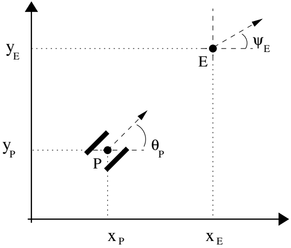

(a) Realistic space

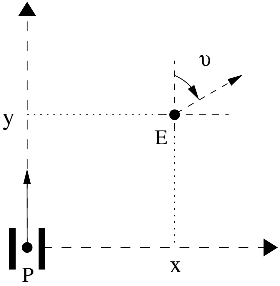

(b) Reduced space

Model

The evolution of the system in the DDR-fixed coordinate system is described by the following equations of motion

$$\begin{aligned}
\dot{x}(t) &= \left(\frac{u_2(t) - u_1(t)}{2b}\right) y(t) + v_1(t) \sin v_2(t) \\
\dot{y}(t) &= -\left(\frac{u_2(t) - u_1(t)}{2b}\right) x(t) - \left(\frac{u_1(t) + u_2(t)}{2}\right) + v_1(t) \cos v_2(t)
\end{aligned} \tag{1}$$

This set of equations can be expressed in the form $\dot{\mathbf{x}} = f(t, \mathbf{x}(t), u(t), v(t))$, where $u(t) = (u_1(t), u_2(t)) \in \widehat{U} = [-V_p^{\max}, V_p^{\max}] \times [-V_p^{\max}, V_p^{\max}]$ and $v(t) = (v_1(t), v_2(t)) \in \widehat{V} = [0, V_e^{\max}] \times [0, 2\pi)$.

Preliminaries

Payoff

A standard representation [Isaacs65, Basar95] of the payoff is

$$J(\mathbf{x}(t_s), u, v) \quad = \quad \int_{t_s}^{t_f} \underbrace{L(\mathbf{x}(\bar{t}), u(\bar{t}), v(\bar{t}))}_{\text{running cost}} \, d\bar{t} \quad + \quad \underbrace{G(\mathbf{x}(t_f))}_{\text{terminal cost}}$$

For problems of *minimum time* [Isaacs65, Basar95], as in this game, $L(\mathbf{x}(t), u(t), v(t)) = 1$ and $G(t_f, \mathbf{x}(t_f)) = 0$. Therefore in our game, the payoff is represented as

$$J(\mathbf{x}(t_s), u, v) = \int_{t_s}^{t_f(\mathbf{x}(t_s),u,v)} d\bar{t} = t_f(\mathbf{x}(t_s), u, v) - t_s \tag{2}$$

Note that $t_f(\mathbf{x}(t_s), u, v)$ depends on the sequence of controls u and v applied to reach the point $\mathbf{x}(t_f)$ from the point $\mathbf{x}(t_s)$.

Preliminaries

Value of the game

For a given state of the system $\mathbf{x}(t_s)$, $V(\mathbf{x}(t_s))$ represents the outcome if the players implement their optimal strategies starting at the point $\mathbf{x}(t_s)$, and it is called the *value of the game* or the *value function* at $\mathbf{x}(t_s)$ [Isaacs65, Basar95]

$$V(\mathbf{x}(t_s)) = \min_{u(t)\in\widehat{U}} \max_{v(t)\in\widehat{V}} J(\mathbf{x}(t_s), u, v) \tag{3}$$

where $\widehat{U}$ and $\widehat{V}$ are the set of valid values for the controls at all time t. $V(\mathbf{x}(t))$ is defined over the entire state space.

Preliminaries

Open and closed-loop strategies

Let $\gamma_p(\mathbf{x}(t))$ and $\gamma_e(\mathbf{x}(t))$ denote the closed-loop strategies of the DDR and the evader, respectively, therefore $u(t) = \gamma_p(\mathbf{x}(t))$ and $v(t) = \gamma_e(\mathbf{x}(t))$.
A strategy pair $(\gamma_p^*(\mathbf{x}(t)), \gamma_e^*(\mathbf{x}(t)))$ is in closed-loop (saddle-point) equilibrium [Basar95] if

$$\begin{aligned} J(\gamma_p^*(\mathbf{x}(t)), \gamma_e(\mathbf{x}(t))) &\leq J(\gamma_p^*(\mathbf{x}(t)), \gamma_e^*(\mathbf{x}(t))) \\ &\leq J(\gamma_p(\mathbf{x}(t)), \gamma_e^*(\mathbf{x}(t))) \forall \gamma_p(\mathbf{x}(t)), \gamma_e(\mathbf{x}(t)) \end{aligned} \tag{4}$$

where J is the payoff of the game in terms of the strategies. An analogous relation exists for open-loop strategies.

Necessary Conditions for Saddle-Point Equilibrium Strategies

Theorem (Pontryagin's Maximum Principle - PMP)

Suppose that the pair $\{\gamma_p^*, \gamma_e^*\}$ *provides a saddle-point solution in closed-loop strategies, with* $\mathbf{x}^*(t)$ *denoting the corresponding state trajectory. Furthermore, let its open-loop representation* $\{u^*(t) = \gamma_p(\mathbf{x}^*(t)), v^*(t) = \gamma_e(\mathbf{x}^*(t))\}$ *also provide a saddle-point solution (in open-loop polices). Then there exists a costate function* $p(\cdot) : [0, t_f] \to R^n$ *such that the following relations are satisfied:*

$$\dot{\mathbf{x}}^*(t) = f(\mathbf{x}^*(t), u^*(t), v^*(t)),\ \mathbf{x}^*(0) = \mathbf{x}(t_s) \tag{5}$$

$$H(p(t), \mathbf{x}^*(t), u^*(t), v(t)) \leq H(p(t), \mathbf{x}^*(t), u^*(t), v^*(t)) \leq H(p(t), \mathbf{x}^*(t), u(t), v^*(t)) \tag{6}$$

$$p(t) = \nabla V(\mathbf{x}(t)) \tag{7}$$

$$\dot{p}^T(t) = -\frac{\partial}{\partial x} H(p(t), \mathbf{x}^*(t), u^*(t), v^*(t))\ \textbf{(Adjoint Equation)} \tag{8}$$

$$p^T(t_f) = \frac{\partial}{\partial x} G(t_f, \mathbf{x}^*(t_f))\ \textit{along}\ \zeta(\mathbf{x}^*(t)) = 0 \tag{9}$$

where

$$H(p(t), \mathbf{x}(t), u(t), v(t)) = p^T(t) \cdot f(\mathbf{x}(t), u(t), v(t)) + L(\mathbf{x}(t), u(t), v(t))\ \textbf{(Hamiltonian)} \tag{10}$$

and T *denotes the transpose operator.*

Time-Optimal Motion Primitives

Optimal controls

Lemma

The time-optimal controls for the DDR that satisfy the Isaacs' equation in the reduced space are given by

$$\begin{aligned} u_1^* &= -sgn\left(\frac{-yV_x}{b} + \frac{xV_y}{b} - V_y\right) V_p^{\max} \\ u_2^* &= -sgn\left(\frac{yV_x}{b} - \frac{xV_y}{b} - V_y\right) V_p^{\max} \end{aligned} \tag{11}$$

We have that both controls are always saturated. The controls of the evader in the reduced space are given by

$$v_1^* = V_e^{\max}, \ \sin v_2^* = \frac{V_x}{\rho}, \ \cos v_2^* = \frac{V_y}{\rho} \tag{12}$$

where $\rho = \sqrt{V_x^2 + V_y^2}$. *The evader will move at maximal speed.*

Decision problem
Usable part and its boundary

The portion of the terminal surface where the DDR can guarantee termination regardless of the choice of controls of the evader is called the *usable part* (UP) [Isaacs65]. From [Isaacs65], we have that the UP is given by

$$\mathrm{UP} = \left\{ \mathbf{x}(t) \in \zeta : \min_{u(t)\in\widehat{U}} \max_{v(t)\in\widehat{V}} \mathbf{n} \cdot f(\mathbf{x}(t), u(t), v(t)) < 0 \right\} \tag{13}$$

where $\widehat{U}$ and $\widehat{V}$ are the sets of valid values for the controls, and $\mathbf{n}$ is the normal vector to ζ from point $\mathbf{x}(t)$ on ζ and extending into the playing space.

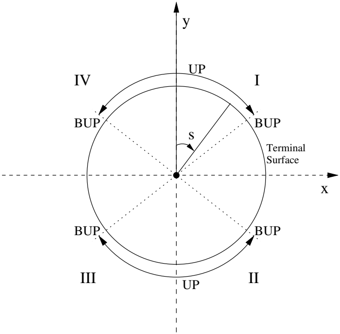

Decision problem

Theorem

If $V_e^{\max}/V_p^{\max} < l|\tan S|/b$ the DDR can capture the evader from any initial configuration in the reduced space. Otherwise the barrier separates the reduced space into two regions:

1. *One between the UP and the barrier.*
2. *Another above the barrier.*

The DDR can only force the capture in the configurations between the UP and the barrier, in which case, the DDR follows a straight line in the realistic space when it captures the evader.

Partition of the reduced space

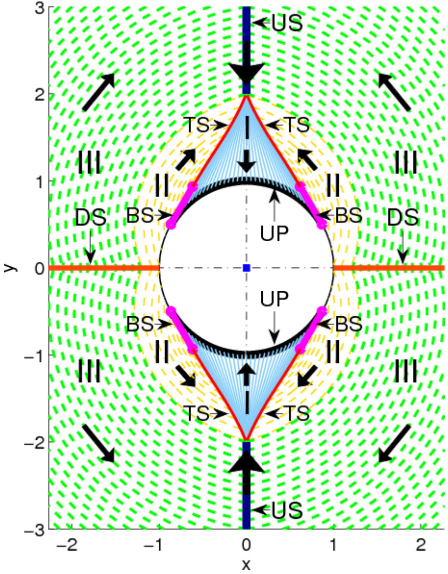

Partition of the first quadrant

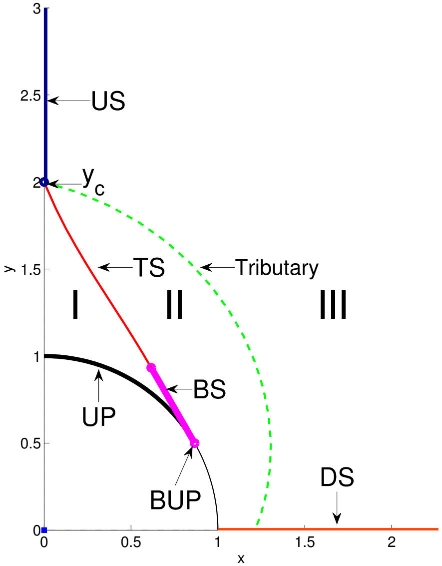

Global optimality

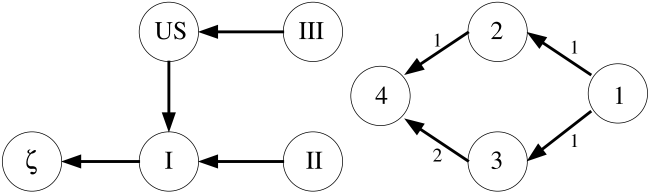

(c) (d)

Figure: Graphs

Simulations - Optimal Strategies

The parameters were $V_p^{max} = 1$, $V_e^{max} = 0.5$, $b = 1$ and $l = 1$. Capture time $t_c = 1.2s$.

Simulations - Evader avoids capture

The parameters were $V_p^{max} = 1$, $V_e^{max} = 0.787$, $b = 1$ and $l = 1$.

State Estimation using the 1D Trifocal Tensor

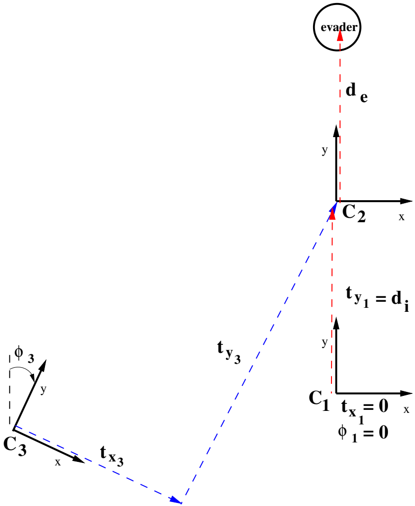

A bound for the angle delimiting the field of view of the pursuer

Theorem

If the evader is in position (r_0, ϕ_0) in the reduced space at the beginning of the game with $\phi_0 < \phi_v$ and $S < \phi_v$ then, if the pursuer applies its time-optimal feedback policy the evader's position (r, ϕ) will satisfy $\phi < \phi_v$ at all times until the capture is achieved regardless of the evader's motion strategy.

Feedback-based motion strategies for the DDR

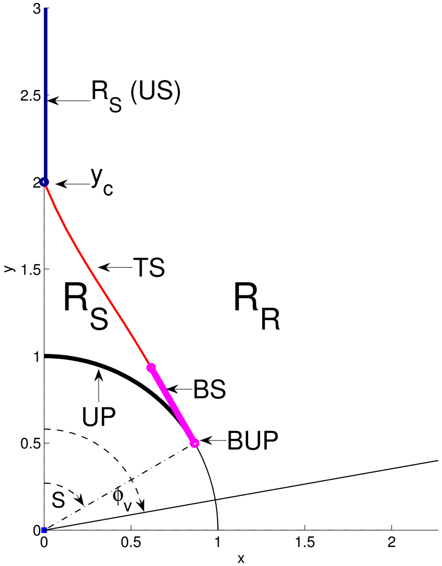

State Estimation

3D View

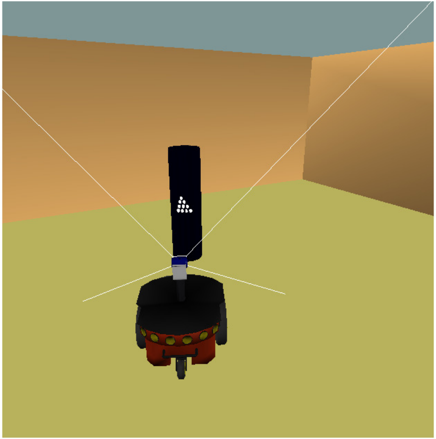

Simulations

Maintaining Visibility of an Evader in an Environment with Obstacles

The Objective

We address our target tracking problem as a game of kind consisting in the next decision problem: is the pursuer able to maintain surveillance of an evader at all time?

Classical Visibility

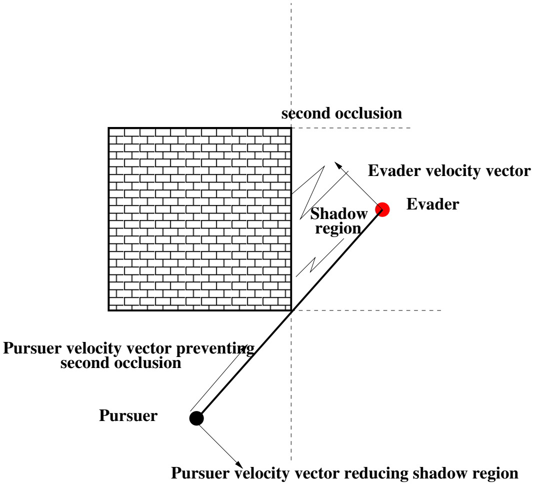

Strong Mutual Visibility

Definition

Two regions R and R' are said to be *strongly mutually visible* (SMV) if visibility holds for all points x and x' such that $x \in R$ and $x' \in R'$.

A straightforward test for the SMV relation using a convex hull computation is given in the following expression:

Regions R and R' are strongly mutually visible if and only if

$$\mathrm{int}[\text{convex-hull}[(R \cup R')]] \subset W$$

where W is the polygon representing the workspace.

Environment Partition and Graphs

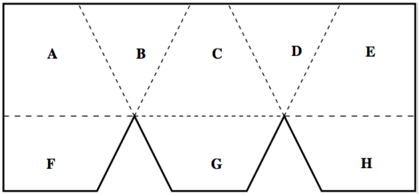

(a) Environment partition

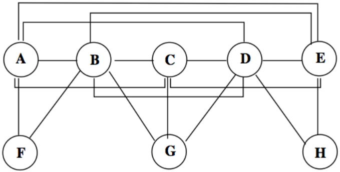

(b) Mutual visibility graph

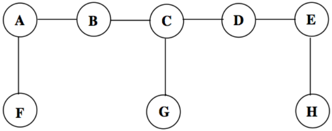

(c) Accessibility graph

Figure: Strong Mutual Visibility

Workspace Partition

$SMV(R_i)$ is the set of regions that are SMV with region R_i.
The total area of these regions is then given by

$$\mathcal{V}_{\mathrm{sm}}(R_i) = \sum_{R_k \in SMV(R_i)} \mu(R_k)$$

$\mu(R_k)$ denotes the area of region R_k.
$\sum \mathcal{V}_{\mathrm{sm}}(R_i)$ is the summation of each $\mathcal{V}_{\mathrm{sm}}(R_i)$ done over all region R_i, this is a global measure intrinsic to a given partition.

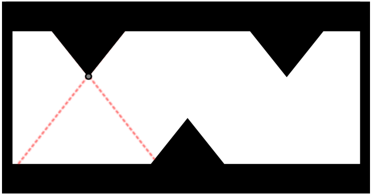

(a)

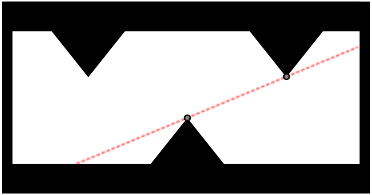

(b)

Figure: (a) Reflex rays and (b) Extended bitangent segment

Workspace Partition

Theorem

For a given reflex vertex v, for any other procedure ς for partitioning R_1 and R_2 into two new regions each, apart from drawing pivot segments, there exists a pivot segment $\mathcal{S}$ that partitions both R_1 and R_2 that yields a bigger value of $\sum \mathcal{V}_{\text{sm}}(R_i)$ than the one related to the partition obtained by applying ς.

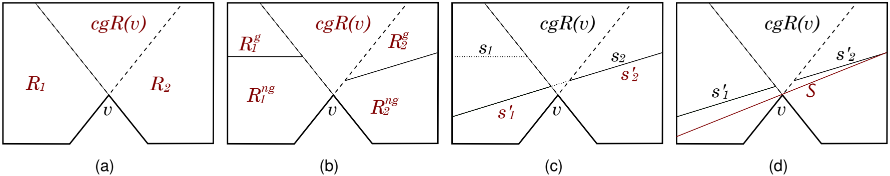

Figure: Different partitioning procedures

Definition

A *guard polygon* for a given point q is the set of all regions in which each of them is mutually visible to all the regions that own the point q. Let $Q(q) = \{R : q \in R\}$, a guard polygon $gP(q)$ for a given point q is defined by:

$$gP(q) = \{R : (R, R_k) \in MVG, \forall R_k \in Q(q)\} \tag{14}$$

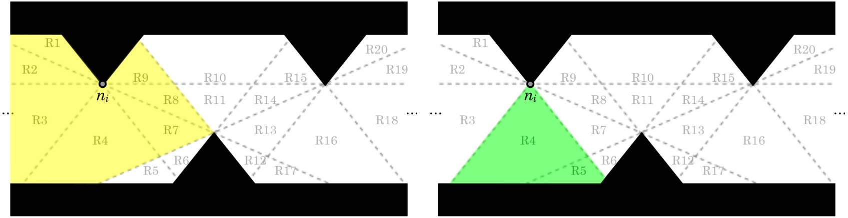

(a) If the evader stands on n_i, it is simultaneously over regions $\{R1, R2, R3, R4, R7, R8, R9\}$

(b) The guard polygon for point n_i is $gP(n_i) = \{R4, R5\}$

Figure: Guard polygon

$$\mathbf{R}_{i,i+1} = \{(v, w) : t_p(v, w) \leq t_e(q_i, q_{i+1}) \text{ where } v \in gP(q_i) \text{ and } w \in gP(q_{i+1})\} \tag{15}$$

Reduced Visibility Graph (RVG)

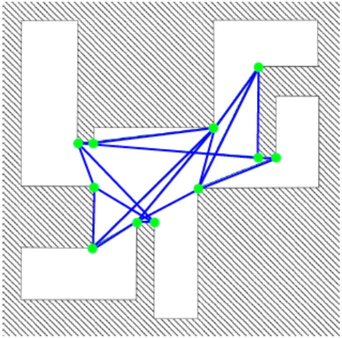

- The vertices in the RVG are the reflex vertices of the environment.
- An edge between two vertices in the RVG is generated if the two vertices are endpoints of the same edge of an obstacle.
- Or if a bitangent line can be drawn between such vertices.

Safe Areas and RVG with Tree Topology

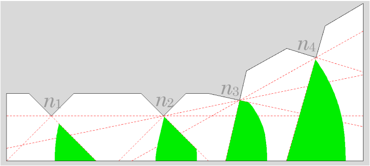

(a)

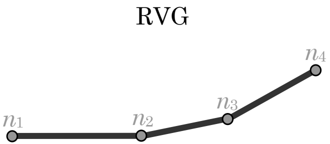

(b)

Figure: Example 1 with a tree topology *RVG* and its calculated safe areas, $\frac{V_p}{V_e} = 0.9$

Safe Areas and RVG with Tree Topology

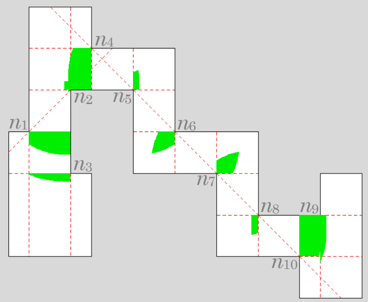

(a)

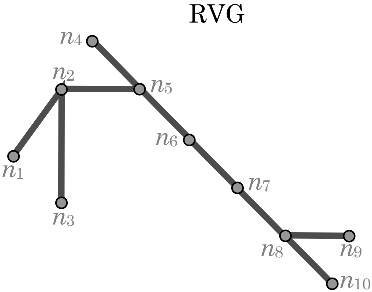

(b)

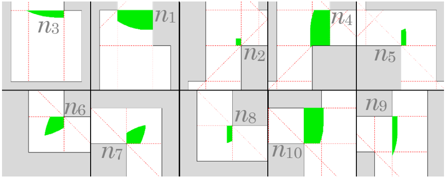

(c)

Figure: Example 2 with a tree topology *RVG* and its calculated safe areas, $\frac{V_p}{V_e} = 1.1$

RVG with Cycles

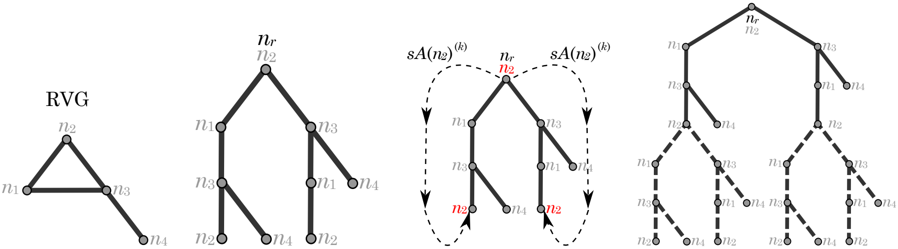

(a) Original RVG (b) Unfolded RVG (c) Feedback procedure (d) Tree that considers 2 laps

Figure: Cycles algorithm

Safe Areas and *RVG* with Cycles

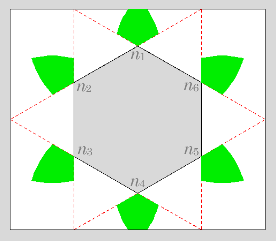

(a) $\frac{V_p}{V_e} = 1.2$

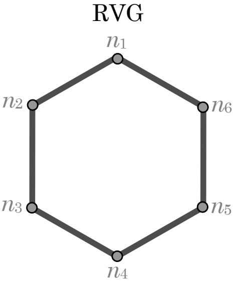

(b)

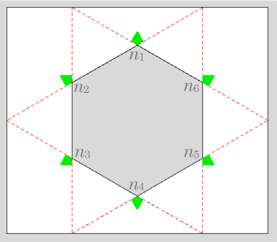

(c) $\frac{V_p}{V_e} = 1.05$

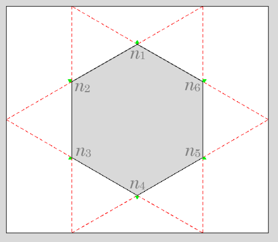

(d) $\frac{V_p}{V_e} = 1.015$

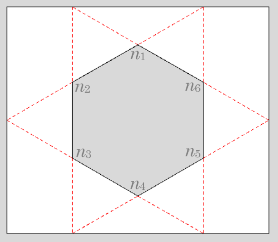

(e) $\frac{V_p}{V_e} = 0.999$

The S Set

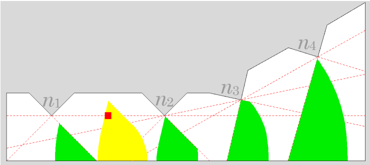

(a) $\frac{V_p}{V_e} = 0.9$

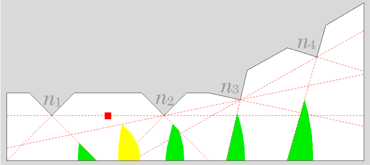

(b) $\frac{V_p}{V_e} = 0.78$

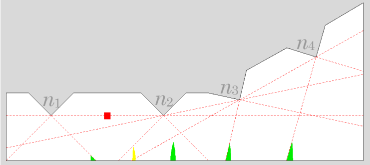

(c) $\frac{V_p}{V_e} = 0.71$

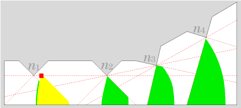

(d)

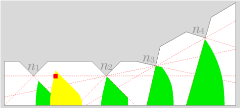

(e)

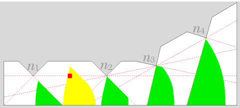

(f)

Figure: Example 4 with its calculated safe areas and sample S sets

Decidability and Complexity

Theorem

The proposed algorithm always converges in a finite number of iterations, hence, the problem of deciding whether or not a pursuer is able to maintain SMV of an evader that travels over the RVG, both players moving at bounded speed, is decidable.

Theorem

The problem of deciding whether or not the pursuer is able to maintain SMV of an evader that travels over the RVG, both players moving at bounded speed, is NP-complete.

Conclusions

In this work, we made the following contributions:

1. Pursuit/Evasion: DDR vs Omnidirectional Agent
- We presented time-optimal motion strategies and the conditions defining the winner for the game of capturing an omnidirectional evader with a differential drive robot.
2. Surveillance with Obstacles
- We proved decidability of this problem for any arbitrary polygonal environment.
- We provided a complexity measure to our evader surveillance game.

Future Work

1. Pursuit/Evasion: DDR vs Omnidirectional Agent
- Capturing an omnidirectional agent using two o more differential drive robots when one is not able to do it.
- To include acceleration bounds in the solution of the problem.
- Feedback motion policy based on the image space.
2. Surveillance with Obstacles
- A moving evader that is free to travel any path within the workspace.

Collaborators and Funding

- Israel Becerra (UIUC)
- David Jacobo (Google)
- Ubaldo Ruiz (CICESE)

- Hector Becerra (CIMAT)
- Seth Hutchinson (UIUC)
- Jean-Paul Laumond (LAAS/CNRS)
- Jose Luis Marroquin (CIMAT)
- Raul Monroy (ITESM CEM)

Funding: National Council of Science and Technology (CONACYT)
México

Publications

- Israel Becerra, Rafael Murrieta-Cid, Raul Monroy, Seth Hutchinson and Jean-Paul Laumond, Maintaining Strong Mutual Visibility of an Evader Moving over the Reduced Visibility Graph, *Autonomous Robots 40(2):395-423*, 2016.
- David Jacobo, Ubaldo Ruiz, Rafael Murrieta-Cid, Hector Becerra and Jose Luis Marroquin, A Visual Feedback-based Time-Optimal Motion Policy for Capturing an Unpredictable Evader, *International Journal of Control 88(4):663-681*, 2015.
- Ubaldo Ruiz, Rafael Murrieta-Cid and Jose Luis Marroquin, Time-Optimal Motion Strategies for Capturing an Omnidirectional Evader using a Differential Drive Robot, *IEEE Transaction on Robotics 29(5):1180-1196*, 2013.

Thanks... Questions?

murrieta@cimat.mx

- T. Başar and G. Olsder, *Dynamic Noncooperative Game Theory, 2nd Ed*. SIAM Series in Classics in Applied Mathematics, Philadelphia, 1995.
- R. Isaacs. *Differential Games: A Mathematical Theory with Applications to Warfare and Pursuit, Control and Optimization*. Wiley, New York, 1965.
- L. S. Pontryagin, V. G. Boltyanskii, R. V. Gamkrelidze, and E. F. Mishchenko. *The Mathematical Theory of Optimal Processes.* JohnWiley, 1962.