31. Differential Games

Neil Roques

May 19, 2005

Contents

1 Introduction

Classical, or discrete, game theory has been an area of great interest in the mathematical community during the last century, showing remarkable depth as a mathematical entity as well as producing results with far reaching implications in finance, military strategy, and of course, recreational games. It is only natural to explore extensions to this theory, and one such extension is to allow our players to select the value of a continuous variable rather than picking one of a number of discrete options.

At the same time a class of problems arose, particularly from military applications, but also from economics, in which we want to find the optimum behaviour for competing with another intelligent entity in an evolving, continuous space. Initially theorists attempted to solve problems such as guided missile pursuit as optimisation problems, an approach which produced the subject of control theory, but it was soon realised by mathematicians such as R. Isaacs and D. L. Kelendzeridze that the strategy of pursuit must consider the strategy of evasion employed by the target.

Differential games are the meeting point of these two: the formulation of a game between opposing intelligent players taking on the tools of calculus to tackle problems operating on continuous variables, their conception and solution will be the subject of this essay.

We will begin by reviewing the relevant concepts of game theory (§2) before taking Isaacs'[1] simple pursuit game as an example to motivate the formal definition of a differential game (§3). We will then seek to develop the techniques that will enable us to solve differential games, first by reducing the problem of finding the optimal strategies to one of 'minimax' on the value of the game - the result of the game when both sides play correctly. From this we are able to construct a system of linked differential equations we can solve in order to solve the game.

Armed with these tools we take two problems of my own construction which illustrate new features of the field, and dispatch them.

2 Game Theory Review

Since our approach to differential games will extensively borrow not only terminology, but also concepts from discrete game theory it will be instructive to briefly review the central concepts of discrete game theory.

A game is played by two opposing players, which we shall call, for reasons which will become apparent, Player $+$ and Player $-$. The players are opposing in the sense that whatever one player gains the other loses (so called zero sum games) hence the payoff at the end of a round may be represented by a single number, positive for Player $+$'s advantage, negative for Player $-$'s. Hence $+$ adopts a strategy to maximise the payoff, $-$ to minimise.

For each round each player secretly decides on one of a finite number of moves, then both are revealed and the payoff associated with this play is awarded. Let us illustrate this with a simple example:

Each player has two possible moves available to him. The associated payoffs are given by the payoff matrix:

$$\begin{array}{ccc} & & - \\ & & \begin{array}{cc} \text{I} & \text{II} \end{array} \\ + & \begin{array}{c} \text{I} \\ \text{II} \end{array} & \begin{pmatrix} 1 & -2 \\ -3 & 4 \end{pmatrix} \end{array}$$

To solve this game we want to find the best strategy for each player and the resulting payoff, called the value of the game.
It is well worth taking time to ask what exactly what it means for a strategy to be the 'best'. The sense we mean is that strategy which *forces* the greatest concession from our opponent, the strategy which, assuming our opponent plays the best moves he can, still forces the best payoff we can achieve. Of course if our opponent plays poorly this optimal strategy might not be the one that takes most advantage of our opponent's error, but it is a boring game in which one strategy is the best response to everything. For example it is possible to show that in scissors-paper-stone the optimal strategy is to pick each option with equal probability, and with this strategy we expect the scores of the two players to be equal regardless of the strategy our opponent plays. If we knew our opponent was fixated by scissors and played them a disproportionately often (as most people do) we might be tempted to always play rock, but our opponent could at any moment develop an obsession with paper that would destroy us, so all rock cannot be called the 'best' strategy. Let us give some definitions:

Definition Let $e(\mathbf{p}, \mathbf{q})$ be the expected payoff when player $+$ plays strategy $\mathbf{p}$ and player $-$, $\mathbf{q}$.

Definition The optimal strategy for player $+$ is that which maximises over his strategies the minimum payoff over his opponents. Conversely the optimal strategy for player $-$ is that which minimises over his strategies the maximum payoff over his opponents. In algebra: $\hat{\mathbf{p}}$ is the optimal strategy for $+$ if and only if

$$\min_{\mathbf{q}} e(\hat{\mathbf{p}}, \mathbf{q}) = \max_{\mathbf{p}} \min_{\mathbf{q}} e(\mathbf{p}, \mathbf{q}) = v$$

Similarly $\hat{\mathbf{q}}$ is the optimal strategy for $-$ if and only if

$$\max_{\mathbf{p}} e(\mathbf{p}, \hat{\mathbf{q}}) = \min_{\mathbf{q}} \max_{\mathbf{p}} e(\mathbf{p}, \mathbf{q}) = u$$

v is the maximum payoff $+$ can force, u the minimum payoff $-$ can force. Clearly $v \not> u$ as otherwise when both players played optimally we would find ourselves, for example, achieving a payoff less than 1, but greater than 2 - an obvious contradiction. Hence $v \leq u$ (as you might suspect they are in fact equal).

Definition A strategy for each player forms an equilibrium pair if neither player can increase their payoff by deviating from their current strategy. Put formally: $\mathbf{p}^*$ and $\mathbf{q}^*$ form an equilibrium pair if and only if

$$e(\mathbf{p}, \mathbf{q}^*) \leq e(\mathbf{p}^*, \mathbf{q}^*) \leq e(\mathbf{p}^*, \mathbf{q}) \quad \forall \mathbf{p}, \mathbf{q}$$

Theorem A pair of strategies are optimal for each player if and only if they form an equilibrium pair.

Proof $\Leftarrow$:

$$\begin{aligned}
e(\mathbf{p},\mathbf{q}^*) \leq \quad & e(\mathbf{p}^*,\mathbf{q}^*) & \leq e(\mathbf{p}^*,\mathbf{q}) \quad \forall \mathbf{p},\mathbf{q} \\
\Rightarrow \max_{\mathbf{p}} e(\mathbf{p},\mathbf{q}^*) \leq \quad & e(\mathbf{p}^*,\mathbf{q}^*) & \leq \min_{\mathbf{q}} e(\mathbf{p}^*,\mathbf{q}) \\
u = \min_{\mathbf{q}} \max_{\mathbf{p}} e(\mathbf{p},\mathbf{q}) \leq \max_{\mathbf{p}} e(\mathbf{p},\mathbf{q}^*) \leq \quad & e(\mathbf{p}^*,\mathbf{q}^*) & \leq \min_{\mathbf{q}} e(\mathbf{p}^*,\mathbf{q}) \leq \max_{\mathbf{p}} \min_{\mathbf{q}} e(\mathbf{p},\mathbf{q}) = v
\end{aligned}$$

However, $v \leq u \Rightarrow$ all the inequality signs above are in fact equalities, hence $\mathbf{p}^*$ and $\mathbf{q}^*$ are optimal strategies.

$\Rightarrow$:

$$e(\mathbf{p},\hat{\mathbf{q}}) \leq \max_{\mathbf{p}} e(\mathbf{p},\hat{\mathbf{q}}) = \min_{\mathbf{q}} \max_{\mathbf{p}} e(\mathbf{p},\mathbf{q}) = e(\hat{\mathbf{p}},\hat{\mathbf{q}}) = \max_{\mathbf{p}} \min_{\mathbf{q}} e(\mathbf{p},\mathbf{q}) = \min_{\mathbf{q}} e(\hat{\mathbf{p}},\mathbf{q}) \leq e(\hat{\mathbf{p}},\mathbf{q})$$

Hence $(\hat{\mathbf{p}},\hat{\mathbf{q}})$ is an equilibrium pair.□
Let's apply this knowledge by solving the game we gave above.
Player + plays I with probability p and II with probability 1-p
Player − plays I with probability q and II with probability 1-q
Player + wishes to select p so as to get:

$$\begin{aligned}
V &= \max_{0\leq p\leq 1}\{\min_{0\leq q\leq 1}(pq - 2p(1-q) - 3(1-p)q + 4(1-p)(1-q))\} \\
&= \max_{0\leq p\leq 1}\{\min_{0\leq q\leq 1}(pq - 2p + 2pq - 3q + 3pq + 4 - 4p - 4q + 4pq)\} \\
&= \max_{0\leq p\leq 1}\{\min_{0\leq q\leq 1}(10pq - 6p - 7q + 4)\} \qquad (1) \\
&= \max_{0\leq p\leq 1}\{4 - 6p + \min_{0\leq q\leq 1}(q\,(10p-7))\}
\end{aligned}$$

$$\begin{aligned}
p \geq \frac{7}{10} \Rightarrow q = 0 \Rightarrow V = \max_{0\leq p\leq 1}\{4-6p\} \Rightarrow p = \frac{7}{10} \Rightarrow V = -\frac{1}{5} \\
p \leq \frac{7}{10} \Rightarrow q = 1 \Rightarrow V = \max_{0\leq p\leq 1}\{4p-3\} \Rightarrow p = \frac{7}{10} \Rightarrow V = -\frac{1}{5}
\end{aligned}$$

Similarly for Player − working from (1)

$$\begin{aligned}
V &= \min_{0\leq q\leq 1}\{4 - 7q + \max_{0\leq p\leq 1}(2p\,(5q-3))\} \\
q \leq \frac{3}{5} \Rightarrow p = 0 \Rightarrow V &= \min_{0\leq q\leq 1}\{4-7q\} \Rightarrow q = \frac{3}{5} \Rightarrow V = -\frac{1}{5} \\
q \geq \frac{3}{5} \Rightarrow p = 1 \Rightarrow V &= \min_{0\leq q\leq 1}\{3q-2\} \Rightarrow q = \frac{3}{5} \Rightarrow V = -\frac{1}{5}
\end{aligned}$$

In all this V is the value of the game, the expected payoff from optimal strategies - we see that in this game Player − has the advantage. We now have the optimal strategy for each player together with the value for the game:
Player +: $(7/10, 3/10)$ Player −: $(3/5, 2/5)$ Value $V = -1/5$

For all their richness discrete games of this form have a number of limitations, in particular no state persists between rounds, each one begins the game anew, and we can only effectively analyse a finite number of choices, with the result that we cannot address situations where players are able to set the value of a continuous variable. However, in reality, situations with both of these features arise, and require solutions, the theory of differential games rises to meet that need.

3 What is a Differential Game?

3.1 Motivating Example - Simple Pursuit

When formulating and examining the formal definition of a differential game it will be helpful to have an example in mind and so to this end we will begin with one of the simplest examples of a differential game which retains the key features*.

The two players reside in $\mathbb{R}^2$. Each chooses at each moment the direction they move in, and their speed, between 0 and m for player $-$, p for player $+$. The game ends when player $-$ comes within a distance l of player $+$. Player $-$ wants this to happen as soon as possible, player $+$ wants to put off his capture for as long as possible.

In the form of differential equations this becomes:

$$\begin{array}{lcl} \dot{x}_1 &=& a\cos\psi \\ \dot{y}_1 &=& a\sin\psi \\ \dot{x}_2 &=& b\cos\phi \\ \dot{y}_2 &=& b\sin\phi \end{array} \qquad \begin{array}{ccc} 0 \leq & \psi & \leq 2\pi \\ 0 \leq & \phi & \leq 2\pi \\ 0 \leq & a & \leq m \\ 0 \leq & b & \leq p \end{array}$$

$$\text{Payoff} \quad P = \int_0^T dt = T \quad \text{Where } T \text{ is the time when the game ends}$$

Here (x_1, y_1) is the position of player $-$, a his speed and ψ his direction. Similarly (x_2, y_2) is the position of player $+$, b his speed and ϕ his direction. $\dot{}$ is the time derivative. The game ends when

$$(x_1 - x_2)^2 + (y_1 - y_2)^2 = l^2$$

We assume that $m < p$ since otherwise it is obviously the case that player $+$ can always evade player $-$ indefinitely.

This form suggests the formal definition which follows, after which we will develop the techniques for solving differential games, and apply them to this example.

3.2 Definition

A differential game consists of the following objects:

*This example, along with the language of the definition, is taken from Isaacs[1], thought he offers only a geometric solution rather than employing the techniques of differential game theory.

Space The space in which the action takes place, ε, a connected subset of an n-dimensional Euclidean space. The state of the game is given by $\mathbf{x} = (x_1, x_2, \ldots, x_n)^T$, where the $x_1, x_2, \ldots, x_n$ are called the state variables, is a position in ε. In the example above x, y, u, v are the state variables.

Terminal Surface ζ, the subset of ε on which the game terminates. We require that ζ form a boundary of ε, which consists of piecewise smooth surfaces, i.e. $n-1$ dimensional manifolds in the underlying n-dimensional euclidean space. ζ does not need to bound ε in the sense of containing it, we will often find ε to be unbounded, as in the example above

$$\zeta = \left\{(x_1, y_1, x_2, y_2) : (x_1 - x_2)^2 + (y_1 - y_2)^2 = l^2\right\}$$

We require ζ to be an $n-1$ dimensional surface as it will form the boundary conditions for our analysis, but exposition of this will have to wait for §5.

Control variables $\mathbf{\Phi} = (\boldsymbol{\phi_1}, \boldsymbol{\phi_2}, \ldots, \boldsymbol{\phi_m})^{\mathbf{T}}, \mathbf{\Psi} = (\boldsymbol{\psi_1}, \boldsymbol{\psi_2}, \ldots, \boldsymbol{\psi_q})^{\mathbf{T}}$ where $\mathbf{\Phi}$ is controlled by player $+$, $\mathbf{\Psi}$ by player $-$, are called the control variables. This is the sense in which our construction is a game - the two players strive against one another to accomplish their desired outcome through the choice of their control variables. While the control variables may be bounded, typically in the form $a_i \leq \phi_i \leq b_i{}^\dagger$, within those bounds the players must be able to chose any value at each instant. In the example above $\Phi = (b, \phi)^T, \Psi = (a, \psi)^T$

Kinematic Equations During the course of the game $\mathbf{x}$ will change according to the decisions made by each player, this change is specified by the kinematic equations which take the form

$$\dot{x_i} = f_i(\mathbf{x}, \mathbf{\Phi}, \mathbf{\Psi})$$

We require that each f_i be smooth with respect to the state variables, and continuous with respect to the control variables.

Payoff As in discrete game theory this is the result of the game, the quantity which player $+$ seeks to maximise and player $-$ to minimise. We will consider payoffs coming in two parts, an amount that accrues over the course of the game (perhaps fuel used by a jet intercepting another), and a terminal payoff based on the precise termination achieved (to continue the example, how far into our territory the enemy bomber was able to get before being intercepted).

$$\text{Payoff} \quad P = \int_0^T G(\mathbf{x}, \mathbf{\Phi}, \mathbf{\Psi}) dt + H(\mathbf{s})$$

We insist on at least one of G and H being non-zero, the zero situation, being equivalent to both players dropping from a high building whilst allowed to choose only what tune they hum on the way down, has little resemblance to a game. If $H = 0$ we say the

†In principle we might encounter constraints of the form $a_i(\mathbf{x}) \leq \phi_i \leq b_i(\mathbf{x})$, but in such circumstances we would usually just perform a change of variables using some appropriate map to obtain new control variables with constant bounds.

game has integral payoff. If $G = 0$ we say the game has terminal payoff. The example above has integral payoff with $G = 1, H = 0$.

3.3 Solution

As in the discrete games discussed in §2, solving a game consists of obtaining the optimal strategy for each player - that is, how he should set his control variables (again equilibrium pairs), and finding the value of the game. The principal complication is that when we have a game as above, we do not really have one game, but a whole family of games, each characterised by their initial conditions. As a case in point, in our example, where the players begin with respect to one another will have a profound impact on how the game develops, and on the eventual payoff. We address this complication by making the parts of our solution functions of the state variables and insisting that they be defined on all of ε, i.e.

$$\mathbf{\Phi} = \mathbf{\Phi}(\mathbf{x}), \mathbf{\Psi} = \mathbf{\Psi}(\mathbf{x}), V = V(\mathbf{x})$$

Observe that this parameterisation is precisely what we need to have strategies which change over the course of the game as the state variables change during play. In the course of seeking optimal strategies we will find ourselves dealing with non-linear partial differential equations. While it would be exciting to resolve issues at the level of the Clay Millennium Prize Problems this does not seem a realistic goal for the essay, so we will not address issues of uniqueness of optimal strategies, instead being satisfied to find any optimal strategy.
We have stated what we have (a differential game) and what we wish to achieve (its solution). In the next two chapters we will develop the method for finding that solution.

4 The Main Equation

4.1 Derivation

The first stage in determining the solution is determining the equation obeyed by the solution where it is smooth*, that is where $\mathbf{\Phi}, \mathbf{\Psi}$ are continuous and $V(\mathbf{x}) \in C_1$, by which we mean:

$$V(\mathbf{x} + \mathbf{u}) = V(\mathbf{x}) + \mathbf{u}.(\nabla V) + O(|\mathbf{u}|^2)^\dagger$$

Theorem Where the solution to the game is smooth it obeys the main equation:

$$\min_{\mathbf{\Psi}} \max_{\mathbf{\Phi}} \{\mathbf{f}(\mathbf{x}, \mathbf{\Phi}, \mathbf{\Psi}).[\nabla V(\mathbf{x})] + G(\mathbf{x}, \mathbf{\Phi}, \mathbf{\Psi})\} = 0$$

where the choice of $\mathbf{\Phi}$ and $\mathbf{\Psi}$ which extremise this give the optimal strategies.

Proof Suppose we begin the game in the state $\mathbf{x}$, we play for the short time δt by which point we have reached the nearby point $\mathbf{y}$. From $\mathbf{y}$ both players play optimally until

*Both Theorem and Proof are derived from Isaacs[1]

†Here and elsewhere we use the standard notation of vector calculus, including, where appropriate, summation convention.

the game terminates. The payoff at the end of the game will be integral payoff accrued between the times 0 and δt, and the payoff of the path followed from $\mathbf{y}$. However, since the play from $\mathbf{y}$ is optimal the payoff from $\mathbf{y}$ for the remainder of the game will amount to the value of the game from $\mathbf{y}$. Hence we get:

$$P(\mathbf{x}) = \int_0^{\delta t} G(\mathbf{x}(t), \mathbf{\Phi}, \mathbf{\Psi}) dt + V(\mathbf{y})$$

Since δt is small and we have insisted on various smoothness conditions we can apply the central method of applied maths: Taylor expand everything in sight.

$$\begin{aligned}
\mathbf{y} &= \mathbf{x} + \delta t \mathbf{f}(\mathbf{x}, \mathbf{\Phi}, \mathbf{\Psi}) + O(\delta t^2) \\
\Rightarrow V(\mathbf{y}) &= V(\mathbf{x}) + \delta t \mathbf{f}(\mathbf{x}, \mathbf{\Phi}, \mathbf{\Psi}).[\nabla V(\mathbf{x})] + O(\delta t^2) \\
\int_0^{\delta t} G(\mathbf{x}(t), \mathbf{\Phi}, \mathbf{\Psi}) dt &= \delta t G(\mathbf{x}, \mathbf{\Phi}, \mathbf{\Psi}) + O(\delta t^2) \\
\Rightarrow P(\mathbf{x}) &= V(\mathbf{x}) + \delta t \left\{ G(\mathbf{x}, \mathbf{\Phi}, \mathbf{\Psi}) + \mathbf{f}(\mathbf{x}, \mathbf{\Phi}, \mathbf{\Psi}).[\nabla V(\mathbf{x})] \right\} + O(\delta t^2)
\end{aligned}$$

Now we seek the optimal strategies for both sides by minimising $P(\mathbf{x})$ over $\mathbf{\Psi}$ and maximising $P(\mathbf{x})$ over $\mathbf{\Phi}$. But

$$\min_{\mathbf{\Psi}} \max_{\mathbf{\Phi}} P(\mathbf{x}) = V(\mathbf{x}) \quad \text{by definition. Then:}$$

$$\begin{aligned}
V(\mathbf{x}) &= V(\mathbf{x}) + \delta t \min_{\mathbf{\Phi}} \max_{\mathbf{\Psi}} \left\{ G(\mathbf{x}, \mathbf{\Phi}, \mathbf{\Psi}) + \mathbf{f}(\mathbf{x}, \mathbf{\Phi}, \mathbf{\Psi}).[\nabla V(\mathbf{x})] \right\} + O(\delta t^2) \\
0 &= \min_{\mathbf{\Psi}} \max_{\mathbf{\Phi}} \left\{ G(\mathbf{x}, \mathbf{\Phi}, \mathbf{\Psi}) + \mathbf{f}(\mathbf{x}, \mathbf{\Phi}, \mathbf{\Psi}).[\nabla V(\mathbf{x})] \right\} + O(\delta t)
\end{aligned}$$

In the limit $\delta t \to 0$ we obtain the main equation.$\square$

4.2 Sums of Sines and Cosines

We will illustrate the use of the main equation by finding the main equation of our example of simple pursuit from §3.1. In order to do so we will need a result on the sums of sines and cosines which finds frequent application in physically inspired differential games and so is worth stating and proving as a lemma.

Lemma Let $\rho = \sqrt{a^2 + b^2}$ then:

$$\max_{\phi}(a \cos\phi + b \sin\phi) = \rho \quad \text{when} \quad \cos\phi = \frac{a}{\rho}, \sin\phi = \frac{b}{\rho}$$

$$\min_{\psi}(a \cos\psi + b \sin\psi) = -\rho \quad \text{when} \quad \cos\psi = -\frac{a}{\rho}, \sin\psi = -\frac{b}{\rho}$$

Proof Pick θ s.t. $\cos\theta = a/\rho,\ \sin\theta = b/\rho$ then

$$\begin{aligned} a\cos\phi + b\sin\phi &= \rho(\cos\theta\cos\phi + \sin\theta\sin\phi) \\ &= \rho\cos(\theta - \phi) \end{aligned}$$

Obviously this has a max of ρ at $\phi = \theta$ and a min of $-\rho$ at $\phi = \theta - \pi$. Since

$$\begin{aligned} \cos(\theta - \pi) &= \cos\theta\cos\pi + \sin\theta\sin\pi = -\cos\theta \\ \sin(\theta - \pi) &= \sin\theta\cos\pi - \sin\pi\cos\theta = -\sin\theta \quad \square \end{aligned}$$

4.3 Simple Pursuit

We are now in a position to find the optimal strategies for the simple pursuit game outlined in §3.1 in terms of derivatives of V using the main equation. We will first restate the kinematic equations for ease of reference, then determine the main equation.

$$\begin{array}{lllll} \dot{x}_1 = a\cos\psi & 0 \le \psi \le 2\pi & & \\ \dot{y}_1 = a\sin\psi & 0 \le \phi \le 2\pi & G = 1 \\ \dot{x}_2 = b\cos\phi & 0 \le a \le m & H = 0 \\ \dot{y}_2 = b\sin\phi & 0 \le b \le p & \end{array}$$

$$\min_{\psi,a}\max_{\phi,b}\{a\cos\psi V_{x_1} + a\sin\psi V_{y_1} + b\cos\phi V_{x_2} + b\sin\phi V_{y_2} + 1\} = 0$$

$$\min_{\psi,a}\{a(V_{x_1}\cos\psi + V_{y_1}\sin\psi)\} + \max_{\phi,b}\{b(V_{x_2}\cos\phi + V_{y_2}\sin\phi)\} - 1 = 0$$

min over a, ψ attained by:

$$\begin{aligned} a &= m \\ \cos\psi &= -\frac{V_{x_1}}{\rho_1} \\ \sin\psi &= -\frac{V_{y_1}}{\rho_1} \\ \rho_1 &= \sqrt{V_{x_1}^2 + V_{y_1}^2} \end{aligned}$$

max over b, ϕ attained by:

$$\begin{aligned} b &= p \\ \cos\phi &= \frac{V_{x_2}}{\rho_2} \\ \sin\phi &= \frac{V_{y_2}}{\rho_2} \\ \rho_2 &= \sqrt{V_{x_2}^2 + V_{y_2}^2} \end{aligned}$$

The optimal behaviour is to always travel at the maximum velocity, as the astute reader had no doubt, already guessed. This is always the case in such problems, and while worth proving once, we shall in future assume the players are doing so, rather than considering the option of traveling at a slower velocity. The main equation thus reduces to:

$$-m\rho_1 + p\rho_2 - 1 = 0$$

5 The Path Equations

In the previous chapter we derived the main equation which enabled us to determine optimal strategies in terms of the value of the game and its derivatives. However, this does not solve the game since all we know about V is the value it takes on ζ, where it equals H, and a partial differential equation which is typically non-linear for even the simplest of problems (witness the example of simple pursuit from the previous chapter). Further our solution will entail following a path in ε, something the main equation alone is incapable of doing, hence we need the path equations*.

5.1 Derivation

Let $\bar{\mathbf{\Phi}}, \bar{\mathbf{\Psi}}$ be optimal strategies for players + and − respectively. From the main equation we know that these are function of $\mathbf{x}$ and ∇V only. Thus the kinematic equations become:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \bar{\mathbf{\Phi}}(\mathbf{x}, \nabla V), \bar{\mathbf{\Psi}}(\mathbf{x}, \nabla V))$$

And from the main equation becomes:

$$\mathbf{f}(\mathbf{x}, \bar{\mathbf{\Phi}}, \bar{\mathbf{\Psi}}).\nabla V(\mathbf{x}) + G(\mathbf{x}, \bar{\mathbf{\Phi}}, \bar{\mathbf{\Psi}}) = 0$$

To this equation we want to apply ∇, hence we consider

$$\begin{aligned} \frac{\partial}{\partial x_i}\left(\mathbf{f}(\mathbf{x}, \bar{\mathbf{\Phi}}, \bar{\mathbf{\Psi}}).\nabla V(\mathbf{x}) + G(\mathbf{x}, \bar{\mathbf{\Phi}}, \bar{\mathbf{\Psi}})\right) &= f_j \frac{\partial^2 V}{\partial x_i \partial x_j} + \frac{\partial f_j}{\partial x_i}\frac{\partial V}{\partial x_j} + \frac{\partial G}{\partial x_i} \\ &\quad + \frac{\partial \bar{\phi}_k}{\partial x_i}\frac{\partial}{\partial \bar{\phi}_k}(\mathbf{f}.\nabla V + G) + [\bar{\mathbf{\Phi}} \leftrightarrow \bar{\mathbf{\Psi}}] \end{aligned}$$

$$\begin{aligned} \text{But} \quad f_j &= \frac{dx_j}{dt} \\ \Rightarrow f_j \frac{\partial^2 V}{\partial x_i \partial x_j} &= \frac{\partial}{\partial x_j}\left(\frac{\partial V}{\partial x_i}\right)\frac{dx_j}{dt} \\ &= \frac{d}{dt}\left(\frac{\partial V}{\partial x_i}\right) \\ &= \frac{\partial \dot{V}}{\partial x_i} \end{aligned}$$

Since each ϕ_k is chosen so as to extremise $\mathbf{f}.\nabla V + G$, by the main equation, either:
a) An interior minimum is obtained in which case $\partial/\partial\bar{\phi}_k\,(\mathbf{f}.\nabla V + G) = 0$
b) $\bar{\phi}_k$ takes an end point value, but since this is constant† $\partial\bar{\phi}_k/\partial x_i = 0$

$$\text{Either way} \quad \frac{\partial \bar{\phi}_k}{\partial x_i}\frac{\partial}{\partial \bar{\phi}_k}(\mathbf{f}.\nabla V + G) = 0$$

The same of course applies to $\bar{\mathbf{\Psi}}$. Hence:

$$\frac{\partial}{\partial x_i}(\mathbf{f}.\nabla V + G) = \frac{\partial \dot{V}}{\partial x_i} + \frac{\partial f_j}{\partial x_i}\frac{\partial V}{\partial x_j} + \frac{\partial G}{\partial x_i}$$

*The derivation of the path equations is based on Isaacs[1], but he treats different examples, these I have solved myself.

†In §3 we insisted that the bounds on the control variables be constant.

$$\text{But} \quad \nabla\left[\mathbf{f}(\mathbf{x},\bar{\mathbf{\Phi}},\bar{\mathbf{\Psi}}).\nabla V(\mathbf{x}) + G(\mathbf{x},\bar{\mathbf{\Phi}},\bar{\mathbf{\Psi}})\right] = 0$$
$$\Rightarrow \nabla\dot{V}(\mathbf{x}) = -\nabla\mathbf{f}(\mathbf{x},\bar{\mathbf{\Phi}},\bar{\mathbf{\Psi}}).[\nabla V(\mathbf{x})] - \nabla G(\mathbf{x},\bar{\mathbf{\Phi}},\bar{\mathbf{\Psi}})$$

Together with the kinematic equations these give $2n$ equations for $2n$ unknowns. However in order to solve differential equations we need initial conditions and we only have final conditions: we know V on ζ, the terminal surface. Hence we change the progressive variable from t to $\tau = T - t$, the time till termination. In these terms we have a system of differential equations we can solve.

$$\text{Denoting} \quad \frac{\partial x}{\partial \tau} = \check{x} \quad \text{so that} \quad \check{x} = -\dot{x} \quad \text{we obtain:}$$

$$\begin{aligned}
\check{\mathbf{x}} &= -\mathbf{f}(\mathbf{x},\bar{\mathbf{\Phi}}(\mathbf{x},\nabla V),\bar{\mathbf{\Psi}}(\mathbf{x},\nabla V)) \\
\nabla\check{V}(\mathbf{x}) &= \nabla\mathbf{f}(\mathbf{x},\bar{\mathbf{\Phi}},\bar{\mathbf{\Psi}}).[\nabla V(\mathbf{x})] + \nabla G(\mathbf{x},\bar{\mathbf{\Phi}},\bar{\mathbf{\Psi}})
\end{aligned}$$

The Retrogressive Path Equations (RPEs).

5.2 Solving Simple Pursuit

We wish to apply the Retrogressive Path Equations to our simple pursuit example to solve it. However, this requires us to analyse V over ζ and in our present formulation ζ is a complicated three dimensional figure embedded in a four dimensional space. While we can solve the problem in this form, the detail makes the example uninstructive and so we instead first exploit some of the symmetries of the problem to reduce the dimension of ε. We do this by observing that the termination of the game and its value care only about the distance between the two players, not their absolute position. Hence we can fix the origin to one of the players by:

$$\begin{aligned}
x &= x_2 - x_1 \\
y &= y_2 - y_1
\end{aligned}$$

The equations then become:

$$\begin{aligned}
\dot{x} &= p\cos\phi - m\cos\psi & \zeta &= \left\{(x,y) : x^2 + y^2 = l^2\right\} & H &= 0 \\
\dot{y} &= p\sin\phi - m\sin\psi & \rho_1 &= \rho_2 = \sqrt{V_x^2 + V_y^2} = \rho & G &= 1
\end{aligned}$$

The main equation becomes:

$$\begin{aligned}
\min_{\psi}\max_{\phi}\{(p\cos\phi - m\cos\psi)V_x + (p\sin\phi - m\sin\psi)V_y + 1\} &= 0 \\
m\min_{\psi}\{-V_x\cos\psi - V_y\sin\psi\} + p\max_{\phi}\{V_x\cos\psi + V_y\sin\psi\} + 1 &= 0
\end{aligned}$$

$$\begin{aligned}
\Rightarrow \cos\phi &= \cos\psi = \frac{V_x}{\rho} \\
\sin\phi &= \sin\psi = \frac{V_y}{\rho}
\end{aligned}$$

So we see that the two players always travel in the same direction, and that the main equation is:

$$\begin{aligned}\rho &= \frac{1}{m-p}\\ \Rightarrow \cos\phi &= \cos\psi = (m-p)V_x\\ \sin\phi &= \sin\psi = (m-p)V_y\end{aligned}$$

The Retrogressive Path Equations are then:

$$\begin{aligned}\check{x} &= (m-p)^2 V_x \quad (2) & \nabla\check{V} &= \mathbf{0}\\ \check{y} &= (m-p)^2 V_y \quad (3) & V &= 0 \quad \text{on} \quad \zeta\end{aligned}$$

Now we need to establish our initial conditions - in particular we need V_x and V_y on ζ. ζ is a circle of radius l centered at at the origin, hence we can parameterise it by the angle θ. The game then terminates with:

$$\begin{aligned}x &= l\cos\theta & \Rightarrow && \frac{dx}{d\theta} &= -l\sin\theta\\ y &= l\sin\theta & && \frac{dy}{d\theta} &= l\cos\theta\end{aligned}$$

Since $V = 0$ uniformly on ζ it follows that $V_\theta = 0$, but

$$\begin{aligned}V_\theta &= \frac{\partial V}{\partial x}\frac{dx}{d\theta} + \frac{\partial V}{\partial y}\frac{dy}{d\theta}\\ 0 &= -\sin\theta V_x + \cos\theta V_y\end{aligned}$$

Taking the main equation and squaring it we get the equations

$$V_x^2 + V_y^2 = \frac{1}{(m-p)^2} \tag{4}$$

$$V_x\sin\theta = V_y\cos\theta \tag{5}$$

on ζ. Substituting (5) into (4) for V_y we obtain

$$V_x^2(1+\tan^2\theta) = \frac{1}{(m-p)^2} \quad \text{but} \quad 1+\tan^2\theta = \sec^2\theta \Rightarrow V_x^2 = \frac{\cos^2\theta}{(m-p)^2}$$

In order to take the square root we must determine the sign. In order for the game to terminate x must decrease in magnitude, so in the RPEs we expect it to increase in magnitude. If the game terminates with positive x (i.e. $\cos\theta \geq 0$) then it must have arrived from the positive x direction, so we expect (2) to show x increasing, but (2) gives $\check{x}$ the sign of V_x, so $\cos\theta$ positive $\Rightarrow V_x$ positive. We can repeat the argument for $\cos\theta$ negative $\Rightarrow V_x$ negative. Thus we have

$$V_x = \frac{\cos\theta}{m-p},\ V_y = \frac{\sin\theta}{m-p}$$

Then (2) and (3) give:

$$\begin{aligned} \check{x} &= (m-p)\cos\theta & x &= [l+(m-p)\tau]\cos\theta \\ \check{y} &= (m-p)\sin\theta & y &= [l+(m-p)\tau]\sin\theta \end{aligned}$$

So if we begin the game at (x, y) then we find

$$\begin{aligned} \tan\theta &= \frac{y}{x} & V_x &= \frac{x}{(m-p)\sqrt{x^2+y^2}} \\ \cos\theta &= \frac{x}{\sqrt{x^2+y^2}} & V_y &= \frac{y}{(m-p)\sqrt{x^2+y^2}} \\ \sin\theta &= \frac{x}{\sqrt{x^2+y^2}} & V &= \frac{\sqrt{x^2+y^2}}{(m-p)} + C \end{aligned}$$

$$\text{But} \quad V(x,y) = 0 \quad \forall (x,y) \in \zeta \quad \Rightarrow V = \frac{\sqrt{x^2+y^2}-l}{(m-p)}$$

The optimal strategies are given by $\phi = \psi = \arctan(y/x)$

5.3 Observations

It is appropriate to observe a number of points about the example we have just solved.

Technique Of course we could have solved the example above immediately on learning that the players traveled in the same direction, but our approach would be have been limited to that case, we have chosen to crack this particular nut with a sledgehammer in order to be sure of our sledgehammers.

Reduced Space The technique used above, of reducing the dimension of ε by exploiting symmetries in the payoff, is of frequent use in problems of differential games. Because all the game really cares about is the value, whenever a symmetry appears in both G and H, we may reduce the dimension of the problem by ignoring changes in state up to that symmetry. In fact in this example we have a remaining rotational symmetry and could reduce the dimension of the problem to one (the distance between the players), we will do this in §7. The space of minimum dimension for the problem is called the reduced space, the problem formulated in that space, the reduced problem. The name 'realistic problem' is applied to formulations in more dimensions. Usually it is best to solve the reduced problem, but sometimes the reduction makes the problem opaque. Simple pursuit is easier to solve in the reduced space, but to do so is less instructive, and the meaning of the result less clear.

6 Transition Surfaces and Industrial Espionage

Thus far we have discussed how to solve only games where the control variables are smooth, however, many games, and much of the interest and variety in the subject, arises from those singular surfaces where this is not the case. A good introduction to these are transition

surfaces - surfaces in ε where the control variables are not continuous. This is best elucidated by an example, through which we can examine a couple of other points*.
Two companies, + and − are competing for a government contract to make paper clips. Both companies are given one month to produce as many paper clips as they can, and at the end of that month the company who has produced the most paper clips wins. Each company has two assets: an automated production line, and a crack team of rogue engineers. A team can divide itself between improving the existing production line, thereby increasing the rate at which paper clips are produced by that company, or breaking into their rival's warehouses and stealing boxes of paper clips.

Higher Derivatives The theory we have outlined has assumed that the kinematic equations are of first order. However, a team devoted to improving the process of production affects the second derivative of the number of paper-clips, how are we to model this? We can proceed without changing our formulation by introducing another state variable for the first derivative, so that if x is the mass of paper-clips

$$\begin{aligned} \dot{x} &= u \\ \dot{u} &= f(\mathbf{x}, \mathbf{\Phi}, \mathbf{\Psi}) \end{aligned}$$

Payoff The game as constructed is a so called 'game of kind', that is the players care only to win the contract, how much they win it by is not really of concern to them. However, usually when we find such a game we are best to convert it into a game of degree (where the payoff is continuous, as in our simple pursuit example) as in most cases games of kind are rather silly when treated as such. For example, if we treat the simple pursuit game as a matter merely of whether capture occurs or not, then every strategy up to an arbitrarily high time is optimal - whatever happens − can still capture, so he may as well write his name with his path for three days before actually going in for the kill. Such behaviour makes the problem highly singular, it is also somewhat absurd - we would hardly describe as optimal the play of a man who, in a game of chess, having a forced checkmate, decided to putter around with an irrelevant knight for twenty moves first. Hence where we are dealing with a deterministic game† it is often wise to generalise the game of kind to one of degree, for example replacing the question of capture with the time to capture. In this case we replace the matter of who has most by a terminal payoff of the difference between the two, positive for company + winning the bid, negative for company −‡.

Fixed Termination How shall we model a game which ends at the fixed time $t = T$? We make T one of the state variables, with the kinematic equation $\dot{T} = -1$, then

*The approach to transition surfaces is inspired by Isaacs[1], but the example, and so of course its solution, is my own invention.

†Games of kind with a probabilistic element, for example, advancing gun duels, don't suffer these issues since we can maximise the probability of a successful outcome, and this gives a continuous payoff.

‡If this seems too abstract a justification imagine that each company suspects the other of planning to bring in a secret pre-produced stock of paper-clips of unknown size, and is acting to maximise their chances of winning in spite of this.

$\zeta = \{\mathbf{x} : T = 0\}$. We will ignore T so far as the main equation goes, because in this setup expressions involving V_T will always turn out to be vacuous.

Let's construct the equations for this example. Let x_1 be the mass of paper-clips possessed by player $+$, e_1 the effectiveness of his team as engineers, r_1 their effectiveness as robbers; x_2, e_2 and r_2 those for player $-$, all of course positive, and let the control variables indicate the attention that company is giving to stealing, then:

$$\begin{aligned} \dot{u}_1 &= e_1(1-\phi) \\ \dot{x}_1 &= u_1 + \frac{r_1}{2}\phi - \frac{r_2}{2}\psi \\ \dot{u}_2 &= e_2(1-\psi) \\ \dot{x}_2 &= u_2 + \frac{r_2}{2}\psi - \frac{r_1}{2}\phi \\ \dot{T} &= -1 \end{aligned} \qquad \begin{aligned} 0 \leq \ &\psi \ \leq 1 \\ 0 \leq \ &\phi \ \leq 1 \\ G &= 0 \\ H &= x_1 - x_2 \end{aligned}$$

We see immediately that there is a symmetry - we don't actually care how many paper-clips each side has, only the difference between them. We can therefore reduce the dimension of our space by setting

$$\begin{aligned} x &= x_1 - x_2 \\ u &= u_1 - u_2 \end{aligned}$$

our equations then become:

$$\begin{aligned} \dot{u} &= e_1(1-\phi) - e_2(1-\psi) \\ \dot{x} &= u + r_1\phi - r_2\psi \\ \dot{T} &= -1 \end{aligned} \qquad \begin{aligned} 0 \leq \ &\psi, \phi \ \leq 1 \\ G &= 0 \\ H &= x \end{aligned}$$

Since the symmetric case turns out to be less interesting let us suppose that $+$'s S.W.A.T. team are better engineers than they are thieves, at least compared to $-$'s, i.e.

$$\frac{r_1}{e_1} \leq \frac{r_2}{e_2}$$

Let's construct the main equation:

$$\begin{aligned} \min_{0\leq\psi\leq 1} \max_{0\leq\phi\leq 1} \{(e_1 - e_2 + e_2\psi - e_1\phi) V_u + (u + r_1\phi - r_2\psi) V_x\} &= 0 \\ (e_1 - e_2)V_u + uV_x + \min_{0\leq\psi\leq 1} \{(e_2 V_u - r_2 V_x)\psi\} + \max_{0\leq\phi\leq 1} \{(r_1 V_x - e_1 V_u)\phi\} &= 0 \end{aligned}$$

Let $S_i = r_i V_x - e_i V_u$ then

$$\bar{\phi} = \begin{cases} 0 & \text{if} \quad S_1 < 0 \\ 1 & \text{if} \quad S_1 > 0 \end{cases} \qquad \bar{\psi} = \begin{cases} 0 & \text{if} \quad S_2 < 0 \\ 1 & \text{if} \quad S_2 > 0 \end{cases}$$

$$(e_1 - e_2)V_u + uV_x + S_1\bar{\phi} - S_2\bar{\psi} = 0$$

Now we construct the RPEs:

$$\begin{aligned} \check{u} &= e_2(1-\bar{\psi}) - e_1(1-\bar{\phi}) \\ \check{x} &= -u + r_2\bar{\psi} - r_1\bar{\phi} \\ \check{T} &= 1 \end{aligned} \qquad \begin{aligned} \check{V}_u &= V_x \\ \check{V}_x &= 0 \end{aligned}$$

On ζ $V = H = x \Rightarrow V_x = 1,\ V_u = 0 \Rightarrow S_i = r_i > 0 \Rightarrow \bar{\phi} = \bar{\psi} = 1$ so both players finish the contest devoting all their attention to stealing. Parameterise ζ by $x = y,\ u = v$ The equations in this region are then:

$$\begin{aligned} \check{u} &= 0 & \Rightarrow \quad u &= v & \check{V}_u &= 1 \quad \Rightarrow & V_u &= \tau \\ \check{x} &= -u + r_2 - r_1 & x &= y + \tau(r_2 - r_1 - v) & \check{V}_x &= 0 & V_x &= 1 \end{aligned}$$

Due to the independence of the RPEs for V_x and V_u from x, u, ϕ and ψ the expressions for V_x and V_u are valid everywhere, and so we shall not calculate them again in this example. The rest of the behaviour continues only until one of the S_is change sign. This occurs at $\tau = r_i/e_i$, which by $r_1/e_1 \leq r_2/e_2$ occurs first* at $\tau = r_1/e_1 = \tau_1$. We proceed by establishing boundary conditions on the surface $\tau = \tau_1$ as we did on ζ and then propagating back from there.

From above we have on $\tau = \tau_1$ $u = v,\ x = z = y + \tau_1(r_2 - r_1 - v)$. The region we are entering has $S_1 < 0,\ S_2 > 0 \Rightarrow \bar{\phi} = 0,\ \bar{\psi} = 1$. Let $\eta = \tau - \tau_1$, of course $\partial/\partial\eta = \check{}$. The RPEs are then:

$$\begin{aligned} \check{u} &= -e_1 & \Rightarrow \quad u &= v - e_1\eta & \Rightarrow \quad & x = z + (r_2 - v)\eta + \tfrac{e_1}{2}\eta^2 \\ \check{x} &= -u + r_2 & \check{x} &= e_1\eta - v + r_2 \end{aligned}$$

This continues until $\tau = \tau_2 = r_2/e_2$ when S_2 changes sign, so we now have $\bar{\phi} = \bar{\psi} = 0$, i.e. if the competition lasts long enough both companies begin devoting their resources entirely to the production process. On this surface

$$u = w = v - e_1(\tau_2 - \tau_1), \quad x = s = z + (r_2 - v)(\tau_2 - \tau_1) + \frac{e_1}{2}(\tau_2 - \tau_1)^2$$

Let $\mu = \tau - \tau_2$, noting again that $\partial/\partial\mu = \check{}$ we see that the RPEs are

$$\begin{aligned} \check{u} &= e_2 - e_1 & \Rightarrow \quad u &= w + (e_2 - e_1)\mu & \Rightarrow \quad & x = s - w\mu - \tfrac{(e_2 - e_1)}{2}\mu^2 \\ \check{x} &= -u & \check{x} &= -w - (e_2 - e_1)\mu \end{aligned}$$

All that remains is the tedious business of substitution to determine the point of impact on ζ for a given (x, u). While we could find V by integrating its derivatives we would need to determine a constant equivalent to the backwards substitution, which will render us the value anyway. This task is done below for the sake of completeness, but the reader will probably find his life happier if he elects to skip the details.

$$w = v - e_1\left(\frac{r_2}{e_2} - \frac{r_1}{e_1}\right) \tag{6}$$

*In the case of equality both S_1 and S_2 change sign simultaneously and the middle phase is of length zero, nonetheless the equations take all this in their stride and our solution is still valid.

$$
\begin{aligned}
u &= w + (e_2 - e_1)\left(T - \frac{r_2}{e_2}\right) \\
\text{Hence } w &= u - (e_2 - e_1)\left(T - \frac{r_2}{e_2}\right) && (7) \\
(6)\&(7)\ v &= u - (e_2 - e_1)\left(T - \frac{r_2}{e_2}\right) + e_1\left(\frac{r_2}{e_2} - \frac{r_1}{e_1}\right) \\
&= u + r_2 - r_1 - (e_2 - e_1)T && (8) \\
x &= s - w\left(T - \frac{r_2}{e_2}\right) - \frac{e_2 - e_1}{2}\left(T - \frac{r_2}{e_2}\right)^2 && (9) \\
(9)\&(7) \quad &= s - \left(u - (e_2 - e_1)\left(T - \frac{r_2}{e_2}\right)\right)\left(T - \frac{r_2}{e_2}\right) - \frac{e_2 - e_1}{2}\left(T - \frac{r_2}{e_2}\right)^2 \\
\Rightarrow s &= x + u\left(T - \frac{r_2}{e_2}\right) - \frac{e_2 - e_1}{2}\left(T - \frac{r_2}{e_2}\right)^2 && (10) \\
\text{By Def } s &= z + (r_2 - v)\left(\frac{r_2}{e_2} - \frac{r_1}{e_1}\right) + \frac{e_1}{2}\left(\frac{r_2}{e_2} - \frac{r_1}{e_1}\right)^2 && (11) \\
(8)(10)(11)\ z &= x + u\left(T - \frac{r_2}{e_2}\right) - \frac{e_2 - e_1}{2}\left(T - \frac{r_2}{e_2}\right)^2 \\
&\quad -(r_1 - u + (e_2 - e_1)T)\left(\frac{r_2}{e_2} - \frac{r_1}{e_1}\right) - \frac{e_1}{2}\left(\frac{r_2}{e_2} - \frac{r_1}{e_1}\right)^2 \\
&= x + u\left(T - \frac{r_1}{e_1}\right) - (e_2 - e_1)T\left(\frac{T}{2} - \frac{r_1}{e_1}\right) - \frac{1}{2}\left(\frac{r_2^2}{e_2} - \frac{r_1^2}{e_1}\right) && (12) \\
\text{By Def } z &= y + \frac{r_1}{e_1}(r_2 - r_1 - v) && (13) \\
(13)\&(8) \quad &= y - r_1 T + \frac{r_1}{e_1}(e_2 T - u) && (14) \\
(12)\&(14)\ y &= x + u\left(T - \frac{r_1}{e_1}\right) - (e_2 - e_1)T\left(\frac{T}{2} - \frac{r_1}{e_1}\right) - \frac{1}{2}\left(\frac{r_2^2}{e_2} - \frac{r_1^2}{e_1}\right) \\
&\quad + r_1 T - \frac{r_1}{e_1}(e_2 T - u) \\
&= x + uT - \frac{e_2 - e_1}{2}T^2 - \frac{1}{2}\left(\frac{r_2^2}{e_2} - \frac{r_1^2}{e_1}\right) && (15)
\end{aligned}
$$

Since $V(x,u) = y$ (15) gives us the value of the game. It is worth pausing to check the result. First we see that the form is that which we expected from the value (i.e. $V_x = 1,\ V_u = T$), and that in the symmetric situation $x = u = 0; \quad r_1 = r_2,\ e_1 = e_2 \quad V = 0$, i.e. a company will always draw with itself, just as we would expect! Note that we cannot examine the result in the limit $T \to 0$ where we would expect $V = x$ as we assumed in our formulation of the solution that $T \geq r_2/e_2$. Solutions for small times are easy enough to find after this analysis, we just include only those phases the game passes into.

7 Construction of the Reduced Problem

7.1 Uniqueness

The formulations of simple pursuit in §3.1 and §4.3 show that the realistic space is not really a well defined concept even in dimension, the example of the last chapter raises similar questions for the reduced problem. Consider the problem:

$$\begin{aligned} \dot{u} &= e_1(1-\phi) - e_2(1-\psi) \\ \dot{T} &= -1 \end{aligned} \qquad \begin{aligned} 0 \leq \ \psi,\phi \ &\leq 1 \\ G &= u + r_1\phi - r_2\psi \\ H &= 0 \end{aligned}$$

This problem is equivalent to the previous for the initial condition $x = 0$ and can be adapted to the condition $x = s$ simply by adding s to the payoff.

It is clear what's happening here: in the main equation the same term appears, previously multiplied by V_x, but since $\mathbf{f}$ is independent of x we know that $\dot{V}_x = 0$ and by $H = x$ we always have $V_x = 1$ on ζ, and hence everywhere, so of course we find the same optimal strategies. Further

$$P_{old} = H(x) = \int f_1 dt = \int G_{new} dt = P_{new}$$

so the value of the game is the same.

7.2 Towards an Algorithmic Construction

However, we can make progress with the matter of defining the reduced space in an algorithmic manner. Suppose we have a differential game proposed in the usual form, that is:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{\Phi}, \mathbf{\Psi}),\ H = H(\mathbf{x}),\ G = G(\mathbf{x}, \mathbf{\Phi}, \mathbf{\Psi}),\ \zeta = \{\xi(\mathbf{x}) = 0\}$$

We can describe ζ as such because it is an $n-1$-dimensional surface, so can be described by one constraint.

We now construct the reduced problem:

$$\begin{aligned} y_1 &= H(\mathbf{x}) \\ y_2 &= \xi(\mathbf{x}) \end{aligned}$$

Assuming h and ξ are non-constant, and one not the function of the other (otherwise we may neglect one or both).

Next we seek to replace the state variables in G with a combination of our existing y_is and the minimum number of additional state variables. If G can be decomposed with respect to $\mathbf{\Phi}$ and $\mathbf{\Psi}$, that is if G is of the form $G = \sum g_i(\mathbf{x})\varpi_i(\mathbf{\Phi}, \mathbf{\Psi})$, then we may do this by, for each i, either expressing $g_i(\mathbf{x}) = \gamma_i(y_1, y_2, \ldots, y_p)$ (initially $p = 2$) or if this is not possible defining $y_{p+1} = g_i(\mathbf{x})$. A peculiarly complex dependence on the control variables can upset this process, but this is not typically the case with games of interest.

We have now selected the state variables that the problem cares about, we need only add those necessary for determining the evolution of these variables. This is the most complex part of the process, while we may begin by writing $\dot{y}_i = \dot{x}_j(\partial y_i/\partial x_j)$ we will sometimes find that clever transformations borne of good geometrical intuition can reduce the problem far further than we might initially expect. Shortly we will illustrate this by constructing the reduced problem for simple pursuit. In any case we construct $\mathbf{l}$ such that $\dot{\mathbf{y}} = \mathbf{l}(\mathbf{y}, \mathbf{\Phi}, \mathbf{\Psi})$.
If the remaining problem is such that $\mathbf{l}$ is independent of y_1 (and hence y_i for $i > 1$ independent of y_1) then we may reduce the problem by the transformation:

$$\begin{aligned} H &\rightarrow c \quad \text{a constant, the initial value of } y_1 \\ G &\rightarrow G + l_1 \\ \mathbf{l} = (l_1, l_2, \ldots)^T &\rightarrow (l_2, l_3, \ldots)^T \end{aligned}$$

This is then the reduced problem, with the set of possible $\mathbf{y}$ forming the reduced space.

7.3 Reducing Simple Pursuit

If we use the above algorithm on the simple pursuit problem of §5.2 then the only state variable we acquire is $r = \sqrt{x^2 + y^2}$ from $\zeta = \{r = l\}$.

$$\begin{aligned} \dot{r} &= \frac{\partial r}{\partial x}\dot{x} + \frac{\partial r}{\partial y}\dot{y} \\ &= \frac{x}{r}(p\cos\phi - m\cos\psi) + \frac{y}{r}(p\sin\phi - m\sin\psi) \end{aligned}$$

If at all possible we want $\dot{r}$ to be independent of x and y. To this end introduce

$$\begin{aligned} \theta &= \arctan\frac{y}{x} \\ \dot{r} &= \cos\theta(p\cos\phi - m\cos\psi) + \sin\theta(p\sin\phi - m\sin\psi) \\ &= p\cos(\phi - \theta) - m\cos(\psi - \theta) \end{aligned}$$

But whatever θ the setting of ϕ and ψ has full control over the value of $\phi - \theta$ and $\psi - \theta$, so we might as well make these the new control variables, with the result that we have reduced the problem to the one dimensional:

$$\dot{r} = p\cos\phi - m\cos\psi \qquad \zeta = \{r = l\} \qquad H = 0,\ G = 1$$

The solution is now trivial from the main equation: $\phi = \psi = 0$, that is $\arctan(y/x)$ in the old co-ordinates.

$$\left.\begin{array}{c} (p - m)V_r + 1 = 0 \\ V(l) = 0 \end{array}\right\} \Rightarrow V(r) = \frac{r - l}{m - p}$$

Corresponding to the results we obtained in §5.2.

8 Manoeuvrability in Pursuit

We will conclude the essay by taking the sledgehammers we have thus far beaten on nuts, and applying them to a problem worthy of their full force*: A drunk cyclist has taken it upon himself to play tag with an unwilling pedestrian. The pedestrian moves with simple motion at speed p, the cyclist travels in the direction he is pointing with speed $m > p$, but may alter this direction at up to a rate r. The pedestrian wants to put off being tagged as long as possible, the cyclist of course wants to capture in as short a time as possible, not least because he's already feeling a bit queasy from the alcohol†. The cyclist has a reach of l in every direction. In the realistic space the equations are then:

$$\begin{aligned} \dot{x}_1 &= p\cos\phi & 0 \leq &\ \phi \leq 2\pi \\ \dot{y}_2 &= p\sin\phi & -1 \leq &\ \psi \leq 1 \\ \dot{x}_2 &= m\cos\theta & \zeta &= \{(x_1 - x_2)^2 + (y_1 - y_2)^2 = l^2\} \\ \dot{y}_2 &= m\sin\theta & G &= 1 \\ \dot{\theta} &= -r\psi & H &= 0 \end{aligned}$$

We can reduce this problem to a space of two dimensions by placing the cyclist at the origin and making the direction of the cyclist the x-axis. We then measure ϕ with respect to the orientation of the cyclist. The impact of ψ on the kinematic equations is then one of rotation. Using infinitesimal notion for convenience we observe that if we rotate the position of the pedestrian $(x, y)^T$ by an angle $d\theta$ in the anti-clockwise direction it becomes:

$$\begin{pmatrix} \cos d\theta & -\sin d\theta \\ \sin d\theta & \cos d\theta \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} x\cos d\theta - y\sin d\theta \\ x\sin d\theta + y\cos d\theta \end{pmatrix}$$

$$\text{Hence} \quad \frac{\partial \begin{pmatrix} x \\ y \end{pmatrix}}{\partial\theta} = \frac{\begin{pmatrix} x\cos d\theta - y\sin d\theta \\ x\sin d\theta + y\cos d\theta \end{pmatrix} - \begin{pmatrix} x \\ y \end{pmatrix}}{d\theta} = \begin{pmatrix} -y \\ x \end{pmatrix}$$

The θ dependence of e.g. $\dot{x}$ is given by:

$$\frac{\partial x}{\partial\theta}\frac{\partial\theta}{\partial t} = yr\psi$$

We are now in a position to give the kinematic equations in the reduced space:

$$\begin{aligned} \dot{x} &= p\cos\phi + yr\psi - m \qquad \zeta = \{x^2 + y^2 = l^2\} \\ \dot{y} &= p\sin\phi - xr\psi \end{aligned}$$

From which we construct the main equation:

*While inspired by Isaacs 'Homicidal Chauffeur' game [1], this example is my own invention.

†Our cyclist is presumably a mathematician of considerable talent to be able to adopt the correct game theoretic strategy even in his inebriated state.

$$\begin{aligned}
\min_{\psi}\max_{\phi}\{(p\cos\phi + yr\psi - m)V_x + (p\sin\phi - xr\psi)V_y + 1\} &= 0 \\
p\max_{\phi}\{V_x\cos\phi + V_y\sin\phi\} + 1 - mV_x + r\min_{\psi}\{(yV_x - xV_y)\psi\} &= 0
\end{aligned}$$

Using the lemma of §4.2 we discover:

$$\begin{aligned}
\rho &= \sqrt{V_x^2 + V_y^2} \\
\cos\phi &= \frac{V_x}{\rho} \\
\sin\phi &= \frac{V_y}{\rho} \\
\psi &= -\mathrm{sgn}(yV_x - xV_y)
\end{aligned}$$

$$p\rho + 1 - mV_x - r|yV_x - xV_y| = 0$$

Observe that ψ takes one of two extreme values, this leads us to expect transition surfaces. We now construct the RPEs:

$$\begin{aligned}
\check{x} &= -p\frac{V_x}{\rho} - yr\psi + m \qquad & \check{V_x} &= -r\psi V_y \\
\check{y} &= -p\frac{V_y}{\rho} + xr\psi & \check{V_y} &= r\psi V_x
\end{aligned}$$

Now we examine V_x and V_y on ζ where $x = l\cos\theta,\ y = l\sin\theta$. From our solution of simple pursuit in §5 we know that on ζ

$$V_x\sin\theta = V_y\sin\theta \Rightarrow yV_x - xV_y = 0$$

While this prevents us from immediately fixing ψ, it does reduce the main equation to:

$$\begin{aligned}
p\rho + 1 - mV_x &= 0 \quad \text{For this to be satisfied we need } V_x \geq 0\,\forall\theta \text{ where impact is possible} \\
p^2(V_x^2 + V_y^2) &= m^2V_x^2 - 2mV_x + 1 \\
p^2V_x^2 &= \cos^2\theta(m^2V_x^2 - 2mV_x + 1) \\
0 &= V_x^2(m^2\cos^2\theta - p^2) - 2mV_x\cos^2\theta + \cos^2\theta \\
V_x &= \frac{m\cos^2\theta \pm \sqrt{m^2\cos^4\theta - \cos^2\theta(M^2\cos^2\theta - p^2)}}{m^2\cos^2\theta - p^2} \\
&= \frac{m\cos^2\theta \pm \sqrt{p^2\cos^2\theta}}{(m\cos\theta + p)(m\cos\theta - p)} \\
&= \frac{\cos\theta}{m\cos\theta \mp p}
\end{aligned}$$

Impact can only occurs on the front of the cyclist - you cannot run over somewhere behind you while traveling forwards! Hence $\cos\theta \geq 0$, but for $V_x \geq 0$ for all θ in this range we need to take the negative root. Hence

$$
\begin{aligned}
V_x &= \frac{\cos\theta}{m\cos\theta + p} \\
V_y &= \frac{\sin\theta}{m\cos\theta + p} \\
\rho &= \frac{1}{m\cos\theta + p} \\
\phi &= \theta
\end{aligned}
$$

We are now in a position to integrate the RPEs for ∇V.

$$
\begin{aligned}
\check{V}_x &= -r\psi V_y \\
\check{V}_y &= r\psi V_x \\
\check{\check{V}}_x &= -r^2\psi^2 V_x \\
V_x &= A\cos r\psi\tau + B\sin r\psi\tau \\
\check{V}_x &= -Ar\psi\sin r\psi\tau + Br\psi\cos r\psi\tau \\
\text{At} \quad \tau = 0 \quad V_x &= A = \frac{\cos\theta}{m\cos\theta + p} \\
\check{V}_x &= Br\psi = -r\psi\frac{\sin\theta}{m\cos\theta + p} \\
\Rightarrow \quad V_x &= \frac{\cos\theta\cos r\psi\tau - \sin\theta\sin r\psi\tau}{m\cos\theta + p} \\
&= \frac{cos(\theta + r\psi\tau)}{m\cos\theta + p} \\
\text{Similarly} \quad V_y &= \frac{sin(\theta + r\psi\tau)}{m\cos\theta + p} \\
\Rightarrow \rho &= \frac{1}{m\cos\theta + p} \quad \text{i.e. constant} \\
\phi &= \theta - r\psi\tau
\end{aligned}
$$

Before we can proceed we must discover what value ϕ initially takes as we move away from ζ. crudely, the significant change just before impact is the growth of x, so $\phi = -\mathrm{sgn}(-V_y) = \mathrm{sgn}(\sin\theta)$, i.e. the cyclist turns towards his target, as we might expect. We will treat the case $0 < \theta < \pi/2$, i.e. $\psi = 1$, the remainder can be easily deduced from symmetry‡.
Our equations are then:

$$
\begin{aligned}
\check{x} &= -p\cos(\theta + r\tau) - yr + m \\
\check{y} &= -p\sin(\theta + r\tau) + xr \\
\check{\check{x}} + r^2 x &= 2pr\sin(\theta + r\tau)
\end{aligned}
$$

‡The case of $\theta = 0$ is likely to involve a new type of surface, the so called universal surface, and is hence beyond the scope of this essay.

$$
\begin{aligned}
x &= A\sin r\tau + B\cos r\tau - p\tau\cos(\theta + r\tau) \\
x(0) &= B = l\cos\theta \\
\check{x}(0) &= rA - p\cos\theta = -p\cos\theta - rl\sin\theta + m \\
x &= \left(\frac{m}{r} - l\sin\theta\right)\sin r\tau + l\cos\theta\cos r\tau - p\tau\cos(\theta + r\tau) \\
&= \frac{m}{r}\sin r\tau + (l - p\tau)\cos(\theta + r\tau) \\
\check{\check{y}} + r^2 y &= rm - 2pr\cos(\theta + r\tau) \\
y &= C\sin r\tau + D\cos r\tau + \frac{m}{r} - p\tau\sin(\theta + r\tau) \\
y(0) &= \frac{m}{r} + D = l\sin\theta \\
\check{y}(0) &= rC - p\sin\theta = -p\sin\theta + rl\cos\theta \\
y &= l\cos\theta\sin r\tau + \left(l\sin\theta - \frac{m}{r}\right)\cos r\tau + \frac{m}{r} - p\tau\sin(\theta + r\tau) \\
&= \frac{m}{r}(1 - \cos r\tau) + (l - p\tau)\sin(\theta + r\tau)
\end{aligned}
$$

We now have expressions for everything inside the region between the upper active region of ζ and our first impact on a transition surface, we should therefore ask where the first transition surface is, we do this by examining

$$
\begin{aligned}
\phi &= \operatorname{sgn}(xV_y - yV_x) \\
xV_y - yV_x &= \frac{1}{\rho}\left\{\left[\frac{m}{r}\sin r\tau + (l - p\tau)\cos(\theta + r\tau)\right]\sin(\theta + r\tau)\right. \\
&\qquad \left. - \left[\frac{m}{r}(1 - \cos r\tau) + (l - p\tau)\sin(\theta + r\tau)\right]\cos(\theta + r\tau)\right\} \\
&= \frac{m}{r\rho}(\cos\theta - \cos(\theta + r\tau))
\end{aligned}
$$

As claimed this is 0 for $\tau = 0$, and positive for $0 < \tau \ll 1$. It is unclear whether the solution for the whole space will need universal surfaces, or an infinite number of transition surfaces, in either case the equations across a transition surface do not vary significantly from those solved above, so further elucidation of them is unnecessary. As for the local solution, this we have done above, aside from the tedious business of backwards substitution which was carefully explored in §6.

9 Conclusion

In the course of the essay we have explored the construction of differential games, and the primary techniques for their solution, which we applied with considerable success to diverse examples. However, in the latter we saw that apparently innocuous problems hinge upon a menagerie of singular surfaces, a subject we were barely able to scratch the surface of, though Isaacs[1] devotes hundreds of pages to them. Nonetheless the structure of a differential game

has proved a powerful modelling tool, and the techniques derived, while somewhat arduous, have proved equally powerful, there is however much more to be explored in the subject.
Important directions in discrete game theory invite similar advances in her differential equivalent, particularly the generalisations to non-zero sum and multiplayer games, this is touched on by Battiston and Weisbuch[2]. It may be that problems such as air traffic control can be successfully formulated and resolved in this manner.
Finally there remains outstanding the issue of the reduced problem which we explored in §7. Although a powerful technique for making difficult problems tractable there is no reliable method for obtaining the reduced problem. While we explored an algorithm to the problem (§7.2) the matter begs an approach which does not depend on significant geometrical insight*. It may be that the resolution of this issue requires a robust method for identifying equivalent games (such a notion might also provide valuable insight to the differing species of singular surface).

10 Bibliography

[1] Rufus Isaacs - Differential Games, 1965
[2] Stegano Battiston, Gérard Weisbuch - Co-evolutionary Differential Games on Networks, 2003

*It is noteworthy that game theorists have shied away from formulating even the most natural problems in more than two dimensions, and many studies of dogfighting rather arbitrarily limit the aircraft involved to a fixed height.