

UNCLASSIFIED

AD NUMBER

AD090492

LIMITATION CHANGES

TO:

Approved for public release; distribution is unlimited.

FROM:

Distribution authorized to U.S. Gov't. agencies and their contractors;
Administrative/Operational Use; 30 NOV 1954.
Other requests shall be referred to Deputy Chief of Staff Research and Development, Department of the Air Force, Washington, DC 20330.

AUTHORITY

RAND ltr 28 Jun 1967

THIS PAGE IS UNCLASSIFIED

UNCLASSIFIED

190492

Armed Services Technical Information Agency

Reproduced by

DOCUMENT SERVICE CENTER

KNOTT BUILDING, DAYTON, 2, OHIO

This document is the property of the United States Government. It is furnished for the duration of the contract and shall be returned when no longer required, or upon recall by ASTIA at the following address: Armed Services Technical Information Agency, Document Service Center, Knott Building, Dayton 2, Ohio.

NOTICE: WHEN GOVERNMENT OR OTHER DRAWINGS, SPECIFICATIONS OR OTHER DATA ARE USED FOR ANY PURPOSE OTHER THAN IN CONNECTION WITH A DEFINITELY RELATED GOVERNMENT PROCUREMENT OPERATION, THE U. S. GOVERNMENT THEREBY INCURS NO LIABILITY, NOR ANY OBLIGATION WHATSOEVER; AND THE FACT THAT THE GOVERNMENT MAY HAVE FORMULATED, FURNISHED, OR IN ANY WAY SUPPLIED THE SAID DRAWINGS, SPECIFICATIONS, OR OTHER DATA IS NOT TO BE REGARDED BY IMPLICATION OR OTHERWISE AS IN ANY MANNER LICENSING THE HOLDER OR ANY OTHER ENTITY, CORPORATION, OR CONVEYING ANY RIGHTS OR PERMISSION TO MANUFACTURE, USE OR SELL ANY PATENTED INVENTION THAT MAY IN ANY WAY BE RELATED THERETO.

UNCLASSIFIED

90492

FILE COPY

90492

U.S. AIR FORCE
Project
RESEARCH

RAND
MEMORANDUM

FC

This is a working paper. It may be expanded, modified, or withdrawn at any time. The views, conclusions, and recommendations expressed herein do not necessarily reflect the official views or policies of the U. S. Air Force.

The RAND Corporation
SANTA MONICA • CALIFORNIA

U S AIR FORCE
PROJECT RAND
RESEARCH MEMORANDUM

DIFFERENTIAL GAMES II:
THE DEFINITION AND FORMULATION

by

Rufus Isaacs

RM-1399

30 November 1954

Assigned to _____

This is a working paper. It may be expanded, modified, or withdrawn at any time. The views, conclusions, and recommendations expressed herein do not necessarily reflect the official views or policies of the United States Air Force.

Copy to RAND 14 April 55

The RAND Corporation
1700 MAIN ST • SANTA MONICA • CALIFORNIA

Copyright, 1955
The RAND Corporation

SUMMARY

This report presents the mathematical definition and concept of a differential game.

It is an intermediate chapter on its subject. It is to be preceded by a chapter discussing the general scope and applications of differential games, and followed by chapters giving methods of solution, examples, etc.

CONTENTS

	Page
1. The Kinematic Situation	1
2. Termination of the Game	4
3. The Payoff	7
4. Games of Kind and Games of Degree	10
5. Strategies	12
6. Transformations	18
7. Two Basic Assumptions	23
8. More About Strategies; Examples	26

DIFFERENTIAL GAMES: II:
THE DEFINITION AND FORMULATION

Rufus Isaacs

To build a theory of differential games, we need an accurate mathematical definition. Because the field is young and we know not all the applications that may still come to light, we wish to keep this definition flexible. Below we set down a canonical form of a differential game, but we intercalate discussion and motivation, so that the reader will find his way should modification be necessary.

Throughout we deal with games of perfect information. That is at all times each player knows the values of the descriptive variables (defined below).

1. The Kinematic Situation

Our theatre of operations is \mathcal{E} , a region in Euclidean n -space and its boundary. This boundary is to consist of pieces of certain surfaces (we mean by surface an $(n-1)$ -dimensional manifold*). We think of a particular point $x = (x_1, \dots, x_n)$ to be in motion in \mathcal{E} , its path being governed by what we shall call the kinematic equations:

$$(1) \quad \dot{x}_j = f_j(x_1, \dots, x_n, \phi_1, \dots, \phi_p, \psi_1, \dots, \psi_\lambda) \quad (j=1, \dots, n)$$

or more briefly

$$(1) \quad \dot{x} = f(x, \phi, \psi) \quad .$$

* Supposed piecewise smooth.

The functions f are given; we suppose them to be of a simple character and in the sequel we shall not hesitate to speak of any partial derivative (of any order) of the f_j that we have occasion to need. We term ϕ and ψ the navigation variables. They are at all times one each under the control of a player. Thus the motion of x is to be thought of as influenced by the wills of two individuals. If they seek conflicting objectives -- and only such cases are of interest -- the situation assumes something of the nature of a game. As suggested by game theory, we will speak of a particular x in \mathcal{E} as a position; we call the x_1, \dots, x_n the descriptive variables in that they describe this position; the two individuals are the players.

The x_j are descriptive in the following sense. If a play of a differential game is halted before completion, the values of the x_1, \dots, x_n at the time of interruption supply all the data needed to reconstruct the finish of the play. We mean that if a new game is commenced starting with these x_j , it will be tantamount to the part of the original that would have occurred after the interruption.

In particular, the values of x_j at the outset supply starting data. Thus when we use the term game, we are not speaking of a single game but of a collection. There is a distinct game emanating from each point of \mathcal{E} .

In general the ϕ and ψ are (individually) subject to certain constraints which depend at most on x . The sets to which they are constrained are closed and usually convex. It will always be understood that we are interested only in ϕ and ψ which satisfy the

constraints and explicit mention of them in the future will be considered unnecessary.

With ψ fixed and x fixed in \mathcal{E} , the set of vectors $f_j(x, \phi, \psi)$ for all ϕ will be called a vectogram or a ϕ -vectogram (likewise for a ψ -vectogram). A full vectogram allows both ϕ and ψ to take all values.

Presupposing the numerical payoffs to be introduced in Section 3, we name the players:

P, controlling ϕ and minimizing.

E, controlling ψ and maximizing.

The names relate to the Pursuer and Evader of pursuit games, an important instance of the theory.

A player may have more or less control of his present and future, but no one can affect the past. Thus we interpret the left side of (1) as a forward time derivative. At each instant during the course of a play, the players are faced with a full vectogram. If we think of each as choosing a value of his navigation variable, the choice results in the selection of a constituent velocity vector along which x will travel in the immediate future. Thus what corresponds to the moves in a discrete game is here a continuous and unrelenting choice of ϕ or ψ . The reader may justly protest that we are demanding of the players feats beyond human ability. We will assuage him in Section 4. In the meanwhile let us be content with the intuitive picture of x moving in \mathcal{E} , the motion being under the partial control of the players.

Certain parts of the boundary of \mathcal{E} may have to be upheld by special constraints. Suppose, for example, in a war game, x_1 represents the supply of certain munition held by a player. The equation (1) with $j = 1$ tells us how this supply increases or diminishes in terms of some navigation variables, which may, for example, be the rate at which the player himself exhausts or replenishes the munition or the rate at which his opponent destroys it. Plainly x_1 cannot be negative, a fact which imposes part of the boundary of \mathcal{E} . But possibly the first and natural formulation of (1) will permit negative x_1 , and some readjustment is necessary.

In general, we may either make the constraints on the ϕ_1 and ψ_1 special at such boundary points of \mathcal{E} to prevent x from going outside, or we must reformulate the f_j so that when, say, an x_1 has reached its lowest allowed limit there can be no further decrease. But this matter is not likely to be a serious obstruction and we can leave its handling to individual cases.

2. Termination of the Game

There is a portion of a surface \mathcal{C} , called the terminal surface, which is part of the boundary of E . When x reaches \mathcal{C} , the game is over.

We take this form of termination as part of the canonical definition, but feel obliged to defend our motivation. Why a surface? Why part of the boundary of E ?

The termination of a pursuit game usually is capture, which at first glance, we might take to mean the coincidence of P , with

E^* . If the x_1 are the totality of variables descriptive of the positions of both P and E, the subset of \mathcal{E} corresponding to capture will normally be of dimension $< n-1$. We reject this definition of capture on two grounds.

First, it is unrealistic. In applications, the point P or E will be some fixed spot on a missile (aircraft, ship, torpedo, etc.) of tangible bulk, intended to serve as an index of the missile's location. On these grounds alone, P and E will never coincide; but in tactical situations often all that is required for captures is less even than physical contact -- just a certain proximity. Thus a more reasonable criterion of capture is, say, to specify a positive δ and capture occurs when the distance P to E is δ . The set of all capture positions is specified by one equation and thus constitutes a surface in \mathcal{E} .

The second ground applies generally. The technique we shall use entails differential equations. A surface provides just the number of dimensions needed for initial conditions so as to obtain unique solutions. A lesser dimensional manifold generally gives rise to singular points of the solutions. (We are acquainted with several examples of games which are pathological because of this

* Simple pursuit games involve the motion of two points, one pursuing the other, as explained in the ensuing paragraph. We shall use P and E for the names of these points as well as for the names of the controlling players.

dimensional deficiency.) Should we be presented with a game with a terminal set \mathcal{C}' of too small a dimension, we amend it by using the boundary of a δ -neighborhood of \mathcal{C}' for a terminal surface. If desirable, we can investigate the limiting situation as $\delta \rightarrow 0$.

Suppose that, in the formulation of a game fresh from the physical situation, the surface \mathcal{C} is not on the boundary of \mathcal{E} , but is interior to \mathcal{E} . Locally it separates \mathcal{E} or, relatedly, \mathcal{C} has two "sides". Often we will wish to count as termination only the cases where x reaches \mathcal{C} from a particular side. For example, let us return to the above pursuit game, and suppose we started from a position where the distance P to $E < \delta$. Clearly we would not want to consider a subsequent occurrence of [distance $PE = \delta$] as capture. What we do here is to exclude all positions with [distance $PE < \delta$] from \mathcal{E} . Then \mathcal{C} will be part of the boundary.

However, there may be cases where \mathcal{C} is desirably in the interior of \mathcal{E} . We will then distinguish between approaches of x to the two sides of \mathcal{C} . We can think of \mathcal{E} as "slit" along \mathcal{C} and \mathcal{C} itself as two-sheeted. Thus, in a sense, we have restored to \mathcal{C} its role of boundary.

What shall we do if x never reaches \mathcal{C} ? A reasonable and quite practical thing is to supply a stop rule. That is we select some large value T of time and decree that the game is over should T elapse. We can bring this situation into the canonical picture by introducing time as a new descriptive variable x_{n+1} . We enlarge

(1) with

$$(2) \quad \dot{x}_{n+1} = 1 \quad ,$$

and take for the new \mathcal{E} the direct product of the old by $[0, T]$. The new \mathcal{C} is the direct product of the old by $[0, T]$ as well as the part of $x_{n+1} = T$ bounding the new \mathcal{E} . We only consider plays of the enlarged game which start from an x with $x_{n+1} = 0$.

3. The Payoff

We will be concerned with two types.

a. Integral Payoff: There is a given function $G(x, \phi, \psi)$ defined for the same range of variables as the f_j and of about the same degree of complexity. The payoff is

$$(3) \quad \int G(x, \phi, \psi) dt$$

the range of integration being over the path traversed by x in a play, extending from the starting point to the point where the path meets \mathcal{C} .

b. Terminal Payoff: There is given a function H of x defined on \mathcal{C} . The payoff is the value of H at the point where x meets \mathcal{C} .

Of course we can also consider payoffs which are combinations of the above two types, but experience has shown that most cases arising in practice fall under (a) or (b).

It would not be difficult to amalgamate the above two concepts. Take G independent of ϕ and ψ and zero in all of E except for a thin layer neighboring \mathcal{C} . Then what we have is nearly a terminal payoff. This approximation could be made exact through the use of a Stieltjes integral in (3).

If we have not imposed a stop rule, it will be necessary to assign a numerical value to the payoff to apply when the play does not terminate. We do not exclude $\pm \infty$ as possibilities.

It is possible to derive from any game with integral payoff an equivalent one with terminal payoff. We adjoin a new positional variable x_{n+1} . The new \mathcal{E} and the new \mathcal{C} consist of the direct product of the old by the entire new axis. We take H to be the value of x_{n+1} on the new \mathcal{C} . We enrich (1) by adjoining

$$(4) \quad \dot{x}_{n+1} = G(x, \phi, \psi) .$$

After analysing the new game, we discard all starting positions save those with $x_{n+1} = 0$. It is easy to perceive the equivalence to our original game. Any path (that is, play) in the derived game becomes one for the old by suppressing its x_{n+1} component, and the payoff is the integral (3) over the latter path*.

There are other types of payoffs subsumable in the above categories.

Suppose t (time) effectively appears among the arguments of the f_j , of G , or even of H . In the latter case, the payoff is a function of the time of termination as well as the place. Then

* It is thus possible to uniformize further by taking the terminal payoff case as canonical. However integral payoffs arise naturally in many applications and often they are more pleasant to visualize and handle.

we adjoin $x_{n+1} = 1$ to (1), take the new \mathcal{E} and \mathcal{C} as the direct products of the old by the full interval $((-\infty, +\infty))$ of x_{n+1} , and replace the argument t in f_j , G or H by x_{n+1} . When the revised game has been analysed we discard all starting points except those with $x_{n+1} = 0$.

There are applications with payoffs

$$(5) \quad \int_0^T G(x, \phi, \psi) dt$$

where T is some prescribed positive value of the time ($t = 0$ means the starting time). Such a T is essentially a descriptive variable. We adjoin to (1)

$$(6) \quad x_{n+1} = -1$$

and take the new \mathcal{E} as the direct product of the old by $[0, \infty)$. For \mathcal{C} we take that part of the boundary of the new \mathcal{E} where $x_{n+1} = 0$. We play the revised game with an integral payoff with integrand G . We utilize only starting points with $x_{n+1} = T$.

Suppose we are given a function $K(x)$ defined in \mathcal{E} and the payoff is to be the value of $K(x)$ at the end of a prescribed time T . We treat this case similarly to the preceding, but use a terminal payoff.

Another type payoff which can be reduced to the standard form not always, but at least in simple cases, is as follows. Let $K(x)$ be given in \mathcal{E} . The payoff will be the minimum of $K(x)$ which occurs during the play. Example: In a pursuit game, how close can the pursuer get to the evader?

Let \mathcal{E}_1 be that subset of \mathcal{C} in which E can cause $K(x)$ to increase whatever P may do. That is \mathcal{E}_1 is the set of x for which

$$(7) \quad \text{Max}_{\psi} \text{Min}_{\phi} \sum_j K_{x_j} f_j(x, \phi, \psi) > 0 \quad .^*$$

Let \mathcal{C} be the boundary of \mathcal{E}_1 . It is clear that if a minimum of K occurs at all, against optimal opposition from E , it will occur on \mathcal{C} .** Thus we may reduce matters to a terminal payoff with H being the value of K on \mathcal{C} . The reader can easily construct examples in which P can achieve low minima only by causing x to enter \mathcal{E}_1 and leave it again; in such cases the above idea will present its difficulties.

4. Games of Kind and Games of Degree

When we speak of a game of degree we will mean one with a continuum of possible payoffs. A game of kind has finitely many, usually only two, the outcome of the game depending on whether or not one of the players can achieve some objective. For example, in a pursuit game the objective might be capture; in a battle game, complete extermination of the opponent.

If a stop rule is imposed a game of kind becomes one with terminal payoff for which H assumes only finitely many values. The

* Unspecified limits of summations will always be 1 and n .

** We are supposing all given functions continuous, differentiable, etc. In general, \mathcal{E}_1 will be open and \mathcal{C} a surface. In fact \mathcal{C} will be defined by (7) with $>$ replaced by $=$.

game falls within our compass and no special treatment is required. Nevertheless it is often possible and desirable to imbed the game of kind within one of degree and deal with the latter.

The solution to a game of kind may be tremendously indeterminate. It results in the division of \mathcal{E} into two (sometimes more) aliquot subsets (some possibly null), one favorable to each player. If the starting point lies in a player's set, then he can attain his objective. Then usually any strategy is optimal for him as long as it lets him remain in his set, while any strategy at all is optimal for his opponent.

Let us take two species of games of kind:

A). A pursuit game with capture of the objective.

B). The same game with the objective capture before a stipulated time T . Such would be the case, say, if P were an interceptor missile with a limited fuel supply.

In both cases we would lose nothing and might gain much if the time of capture is made the payoff, which is taken as ∞ if capture does not occur. We can then expect definite optimal strategies instead of a sprawling class delineated only by inequalities. In case A, the strategy will not only instruct P as to how to capture, but will show him how to do it as quickly as possible. Similarly, it tells E how to delay it. If we have Case B, we need only look at the value and see whether or not it exceeds T . We will have solved Case B for all possible values of T at once.

However we do not advocate complete abandonment of the game of kind. There are cases where the direct solution is much simpler than the embedding procedure just suggested and supplementary information of little value. Sometimes, too, a game of kind appears as a phase of a game of degree. For example, a player may not be able to terminate favorably unless he first surmounts some obstacle. The question of whether this can or cannot be done may constitute a game of kind whose solution is a preliminary to that of the whole game.

5. Strategies

In the theory of discrete games, a strategy is defined as a set of decisions for a player, one for each position that may arise. If each player chooses a strategy, the ensuing play, and particularly the payoff, is uniquely determined.

We recognize a somewhat analogous circumstance existing here. The election of a decision for each possible position amounts to a player's choosing his navigation variables as functions of the descriptive variables. If the players each so select $\phi(x)$ and $\psi(x)$ and these values are inserted in the kinematic equations (1), the latter become differential equations. Recalling that the data of a game must include a starting value of x , we see that this value plays the role of an initial condition. Thus, under reasonable circumstances, we may expect the paths — and hence the payoff — to be uniquely determined. This concept requires some refinement which will be supplied below.

In the definition of a discrete game the decisions that each player may make are expressed in terms of moves. The concept of a

strategy does not appear until we are done with the definition and are ready for the analysis. With differential games things appear otherwise. In Section 1, we saw the difficulty entailed in defining for differential games the analogue of a move. Hence we must promote strategies from the ranks of tools for analysis to a place in the rules of the game.

An attempt to define strategies in the form $(\phi(x), \psi(x))$ leads immediately to difficulties. First we require assurance that the differential equations to which (1) are converted are integrable. We recall that the left sides are to be construed as forward derivatives. Now we have never seen in the literature an existence criterion for forward differential equations. Certainly, it could be much broader than for the usual variety. The difficulty of delineating the criterion suitably for our purpose will perhaps be discovered in the two examples below.

$$\dot{x}_1 \text{ (forward derivative)} = f_1(x_1, x_2) \quad (i = 1, 2)$$

where (f_1, f_2)

$$= (1, 1) \text{ when } x_1 < 0$$

$$= (0, 2) \text{ when } x_1 = 0$$

$$= (-3, 0) \text{ when } x_1 > 0 \quad . . .$$

The reader will see that this system has exactly one solution starting from each point in plane. Later we will find that functions of this genre are no strangers to the solutions of differential games.

Now let us take for f :

(1, 1) when $x_1 \leq 0$

(-3, 0) when $x_1 > 0$.

We are frustrated should we start from or arrive at a point with $x_1 = 0$.

We will say that $\phi(x)$ is playable if the forward differential equations

$$\dot{x} = f(x, \phi(x), c)$$

have a solution starting from each point of \mathcal{E} for each constant c^* ; correspondingly for $\psi(x)$.

In the sequel we shall develop methods of solution. The results will include values of ϕ , ψ , say $\bar{\phi}(x)$ and $\bar{\psi}(x)$ which in general will be playable. But there remain a second difficulty. An assertion that, say, $\bar{\phi}$ is optimal entails knowledge of its performance when opposed by a certain class of ψ . What class? It must be such that the pair $\bar{\phi}$, ψ will always lead to an integrable (1) and that all ψ representative of an actual player's tactics are included.

Sam Karlin has advanced an idea which obviates these troubles. A strategy for P is defined as both the choice of a function $\phi(x)$, now subject to no restrictions (save the constraints) and a sequence $\sigma_t = \{t_0=0, t_1, t_2, \dots\}$ of increasing values of time with $\lim_j t_j = \infty$.

* Of course, constants outside the range of any ψ subject to the constraints can be ignored.

Such will be called a K-strategy. In playing it, suppose P at time t_k finds x to be at $x^{(k)}$ ($x^{(0)}$ is the starting position). Then in the interlude $[t_k, t_{k+1})$ P holds ϕ to the constant value $\phi(x^{(k)})$.

Let a K-strategy, $\psi(x)$ and $\sigma_t = \{t_0 = 0, t_1', \dots\}$ be also defined for E. We subdivide time by both the t_j and the t_j' . In each subinterval the equations (1) are obviously integrable. We build the path, using as the initial values for each later interval the final x from the previous.

Thus for each starting point and pair of K-strategies, the path of x and consequently the payoff is uniquely determined. We define the value in usual way as the $\sup \inf$ (= $\inf \sup$) of the payoff.

It is manifest that, in general, the K-strategies will not yield optimal strategies but only ϵ -optimal strategies; that is, strategies will attain within ϵ of the value (this being done by increasing the fineness of the temporal subdivision). However we can expect to be able to obtain from a sequence of ϵ -strategies, with the $\epsilon \rightarrow 0$, say, for P a single $\bar{\phi}(x)$ which is a common part of all of them. We will say that $\bar{\phi}(x)$ is an optimal tactic. In most cases, it will turn out to be a playable function and can be viewed in the same light as a few pages back when we discussed strategies more primitively.

When employing K-strategies it is imperative that the game be in a form in which the constraints on the navigation variables are independent of x . See (10) and (11) below.

It is clear that if we make changes of variable on the ϕ or ψ , the K-strategies will be effectively altered. Thus a definite labelling of the navigation variables must be fixed at the outset. We have at our disposal some choice. Often it is possible to exercise it so that an optimal tactic bears a constant numerical value. Then the K-strategies will actually be optimal. See Example 4 in Section 8.

We have now a well defined, self-contained mathematical situation and it is quite possible to stand pat on it. But the reader may raise the following justifiable question, stemming from the fact that we have been forced to incorporate strategies into the rules of a game. Suppose a player, say E, follows the dictates of a policy S_E^* , not a strategy as we have defined it. For example, S_E^* may entail a ψ which is a function of x_j , other higher derivatives of the x_j (E somehow provides for cases where they don't exist), past values of the x_j , integrals over such past values, etc. If E pits S_E^* against an optimal strategy S_P for P, how do we know he will not emerge with better than the value?

We will endeavor to reply in two ways. The first is heuristic. It is based on the fact that the descriptive variables are truly descriptive of the position in the sense discussed in Section 1.

To illustrate let us consider in part a pursuit game in which P is a point moving in a plane. Let x_1, x_2 be his coordinates. First we will consider a very simple kind of motion: P is constrained only to move with speed v . Then $d\phi = 1$, and those of the kinematic equations relating to P are

$$\begin{aligned}\dot{x}_1 &= v \cos \phi_1 \\ \dot{x}_2 &= v \sin \phi_1 .\end{aligned}$$

We claim that the only rational policy for E is to base his actions on x_1 and x_2 alone. He might have them depend on, say, \dot{x}_1 , \dot{x}_2 , past values of the x_j etc. as well, such as would be the case if he endeavored to extrapolate P's future positions. But P's velocities - according to the way we have framed the problem - are at all times subject to abrupt change without notice. It is impossible for E to rely on any prediction or indeed to derive any constructive knowledge from anything other than x_1 and x_2 .

Let us now make P's motion a bit more complicated (and more realistic). We suppose that now he regulates the acceleration ϕ_1 , ϕ_2 (subject to some bounds that we won't mention). The kinematic equations in part are now

$$\begin{aligned}\dot{x}_1 &= x_3 \\ \dot{x}_2 &= x_4 \\ \dot{x}_3 &= \phi_1 \\ \dot{x}_4 &= \phi_2 .\end{aligned}$$

Now P can no longer abruptly switch velocities and there are sound grounds for E's basing his policy upon them. But the velocities are x_3 , x_4 and appear among the descriptive variables. However, the same argument as before shows that E could be misled if he based his decisions on, say, P's acceleration.

We can proceed thus, creating a chain of more and more complex types of motion for P and, in fact, can introduce many variants and

offshoots along the way. In each case we single out those data of the motion on which it appears. E can rationally rely when making his decisions, and in each case they appear among the descriptive variables.

The second reply is mathematical. Suppose P plays S_P and E plays S_E^* from some starting point. On the resulting path we will have the ψ arising from S_E^* defined as a function of t . We take the liberty of supposing that this function is piecewise continuous. Then there will be a strategy S_E for E which will agree with S_E^* whenever a play results in this same path. Thus E will reap the same yield if he plays S_E or S_E^* as long as P adheres to S_P . As S_P is optimal, E cannot do better than the value.

From the standpoint of mathematical rigor, the liberty we took above is admittedly a weakness. But from the standpoint of applications, it is hard to see how a human being could devise a practicable scheme of playing which would result in a navigation variable being a more complicated function of time than piecewise continuous.

6. Transformations

By performing changes of variables on the navigation variables, the descriptive variables, or the time, we can alter the format of a differential game without changing its nature. There are essentially two desiderata for such transformations:

- 1). Bringing the game to a canonical form for the purpose of obtaining general theorems.

2). Facilitating the solution of particular problems. The applied mathematician does not need to be told of the dividends in simplicity that at times result from an adroit choice of a coordinate system.

The two objectives need not coincide. Judgement is required by the investigator in Case 2) and tolerance in Case 1), as the youth of our science renders premature the adoption of too rigid a standard form.

We consider the navigation variables. For definiteness we speak only of ϕ . First we may suppose there are no redundancies; we mean that to each member of a ϕ -vectogram there corresponds exactly one value of ϕ . If the original problem is not so framed, we may, after a possible variable change on the ϕ_1 , purge the superfluous components. Then, for games with integral payoff we are assured that

$$(8) \quad \mu \leq n, \quad \lambda \leq n .$$

For terminal payoffs, we see that only the directions and not the magnitudes of vectors in a vectogram are significant. Thus in this case we will have

$$(9) \quad \mu \leq n-1, \quad \lambda \leq n-1 .$$

We recall that ϕ suffers constraints which depend continuously on x . We always suppose that the set to which ϕ is confined is closed. We may also take it to be simply connected^{*}. Under

* See the convexity assumption in the next section.

reasonable circumstances, it will also be bounded. Then we can find a continuous, biunique mapping

$$(9a) \quad \phi_j = \Phi_j(x, \phi'_j)$$

which takes the constraining set for each x into the hypercube

$$(10) \quad -1 \leq \phi_j \leq 1 \quad . \quad (j = 1, \dots, d\phi)$$

Interpreting (9a) as a change of variable, the constraints assume the form (10), and we so take them for our canonical model.

If the constraining set is not bounded, then there will be at least a way, for some j , to attain in place of (10)

$$(11) \quad \phi_j \text{ unbounded}$$

or

$$(11') \quad \phi_j \geq 0 \quad .$$

This case is not likely to occur in practice. Unbounded ϕ_j which cannot be mitigated by a variable change imply vectors of unbounded magnitude in some vectograms. It may be that good strategies never require P to utilize the very large vectors; then nothing is lost if we curtail the recalcitrant ϕ_j at some large, synthetic bound. In the contrary case, we can expect to be confronted by a game with no solution. Probably the physical circumstances to which it relates are pathological or our mathematical model falsely describes them.

Redundancies may also appear among the descriptive variables when the formulation is fresh from the physical application. Consider, for example, a simple pursuit game in which P and E are two points on

the same straight line with x_1, x_2 as their single coordinates. The kinematics are independent of position and, if the payoff is integral, so is G . Then we need only one descriptive variable, namely $x_1 - x_2$. However if the payoff is terminal, we need two, x_1 and x_2 .

Another type may arise through the kinematic equations if $d\phi + \lambda < n-1$. From a fixed starting point, all paths may find themselves confined to a subset \mathcal{E}' of \mathcal{E} with dimension $< n$. Then we may use \mathcal{E}' in place of \mathcal{E} . We illustrate by

Example 1. Let $n = 3, d\phi = 1, \lambda = 0$. The kinematic equations are

$$(12) \quad x = \alpha(x)\phi_1 + \beta(x) \quad (-1 \leq \phi_1 \leq 1)$$

with $\gamma(x) = \alpha \times \beta \neq 0$. Then, if

$$(13) \quad \gamma \cdot \text{curl } \gamma = 0$$

a known result of classical analysis tells us that \mathcal{E} is covered by a family of surfaces such that everywhere the vectors α and β lie in the tangent plane. Then x must always remain in the same one of these surfaces from which it started. We can use this surface for \mathcal{E} .

When the number of x_1 has been made satisfactory, we can if desired make transformations to attain such convenient ends as having \mathcal{C} lie in the surface: $x_1 = 0$, etc.

If $G = 1$, in a game with integral payoff, the payoff is the time until termination. We can always bring about something like

this situation by a change of variable on t . We can cause (1) to be replaced by

$$(14) \quad \dot{x}_j = \frac{f_1(x, \phi, \psi)}{|G(x, \phi, \psi)|} *$$

with the payoff becoming

$$(15) \quad \int \operatorname{sgn} G(x, \phi, \psi) dt .$$

When such is done, we will speak of a time payoff. If always $G > 0$, then the time payoff is time to termination.

For games with terminal payoff, we have already remarked that only the directions of the members of a vectogram count. A change of variable on t , then, does not affect the game in any essential way. We can bring about such conveniences as making the vectors in a ϕ vectogram all have the same length or making the f_1 linear in ϕ .

* The vanishing of G is not drastic. If $G = 0$ everywhere, the original game was trivial. If G vanishes only on certain surfaces (for all ϕ and ψ) these surfaces may well be singular, as will appear in some later examples. If for each x , G vanishes for a discrete set of ϕ or ψ and these are used for an interlude in the play, the resulting part of the path may still exist as a curve (namely, the same curve as the path would lie on before we made the transformation (14)). The "infinite velocities" on these arcs simply mean that for them (15) is 0.

7. Two Basic Assumptions

For a vector $u = (u_1, \dots, u_n)$, write for short

$$Q = \sum_1 u_1 f_1(x, \phi, \psi) .$$

Vital for our work is the

Minimax Assumption. For all u and all x in \mathcal{E} , if the payoff is terminal

$$(16) \quad \min_{\phi} \max_{\psi} Q = \max_{\psi} \min_{\phi} Q$$

and if the payoff is integral

$$(17) \quad \min_{\phi} \max_{\psi} (Q+G) = \max_{\psi} \min_{\phi} (Q+G) .$$

In all applications encountered up to the time of writing, each f_1 and G have been separable; that is, of the form: function independent of ψ + function independent of ϕ . In this case (16) and (17) hold.

We will say a vectogram is convex if, whenever v_1, \dots, v_k belong to it, so does $\sum_{i=1}^k c_i v_i$ where $c_i \geq 0$, $\sum_{i=1}^k c_i = 1$.

Convexity Assumption (Integral Payoff). When the game has been reduced to time payoff*, all ϕ - and ψ -vectograms are convex.

Convexity Assumption (Terminal Payoff). All ϕ - and ψ -vectograms are convex.

* The reduction is unnecessary if G is independent of ϕ and ψ .

Should this assumption be violated there may be no solution.* We do not reject the game, but replace it by another in which the ϕ - and ψ -vectograms are the convex hulls of those of the old. If the new game can be solved, its solution will supply the essential information about the old.

An example will clarify matters. The reader will see how to apply its idea generally.

Example 2. In this game ψ does not appear, so that it really is a minimizing problem rather than a game. Here \mathcal{C} is the part of plane above the curve \mathcal{C} of Figure 1. The vectograms are the same for all x ; one is sketched. Let M be a high point of \mathcal{C} ; P is to start from x^0 , directly above M , and reach \mathcal{C} in the least time. Clearly a solution will entail a zigzag path arising from an alternate use of P 's two extreme velocities. There will be many solutions.




Figure 1.

* The game may have a value, but one or both optimal tactics may fail to exist.

Example 2'. Let L be a vertical line through M . We alter matters by letting the vectograms preserve their form but letting them decrease in magnitude with the distance from L . Then clearly P does the better with finer zigzags which stay closer to L . There is now no solution.*

Perform the alteration described above, replacing the vectograms by their convex hulls. The game now has a solution: P traverses L to M . We see in what sense this solution is approximated by those of the unaltered game.

In practice this assumption does not betoken much of an encumbrance. Generally, a bit of reflection will show that it is not sound policy for a player to employ a vector which is properly interior to his vectogram. For example, seldom in a pursuit or maneuvering game will it profit a player to use less than the top speed allowed him by the rules, and any exceptions will have evident causes. Thus in the examples to come later we shall usually ignore the convexity assumption in the literal formulation of the mathematical model, but we should always keep in mind that we are doing so and be prepared to make modifications when required.

A possibly noteworthy exception occurs when we deal with missiles (aircraft, torpedoes, etc.) with snap action controls. For example, the guidance mechanism may operate a two-position rudder. In the mathematical model it is best to allow intermediate positions, the situation being roughly similar to that of Example 2.

* If we think in terms of K -strategies, the practical difficulties become slight.

8. More About Strategies; Examples

In most quotidian applications optimal tactics can be used directly as optimal strategies. This may mean that they furnish practical instructions as how to navigate optimally or that the equations (1) are integrable when ϕ, ψ are optimal tactics. We will then speak of direct optimal play and call the arising paths direct optimal paths. However it happens so often that we can use direct play instead of K-strategies, that we shall seldom make a distinction. In fact, optimal strategies and optimal tactics are generally interchangeable terms. But the following example shows that this need not always be the case.

Example 3. We work in the plane: \mathcal{C} is the x_2 -axis, \mathcal{E} the half-plane $x_1 \leq 0$. The payoff is time of termination. This is to be a one-player game (or E is inactive), and the payoff is to be the time to termination. Two typical ϕ -vectograms

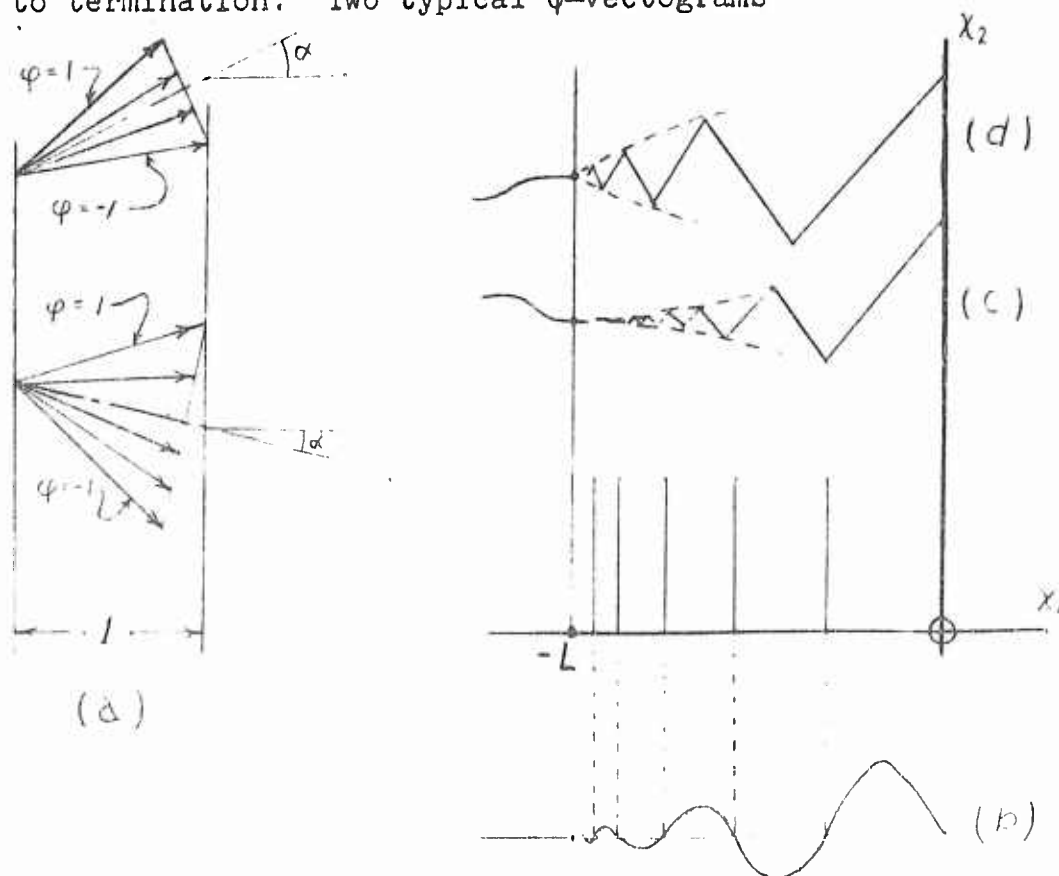


Figure 2

are sketched in Figure 2(a). They are always to be triangular in shape with lengths of the vectors so chosen that the length of the greatest horizontal projection is unity. The inclination α of the centerline is to be reasonably small and is a function of x_1 . We take $\phi^* = -1$ for the lowermost vector and, $+1$ for the uppermost. (The angle between these extreme vectors is not critical for our purposes.) It is clear that the optimal ϕ is -1 when $\alpha > 0$ and $+1$ when $\alpha < 0$.

Let us take for $\alpha(x_1)$ a function such as is graphed in Figure 2(b). For $x_1 \leq -L$, α is to be zero; to the right of this line α is to oscillate infinitely often. By diminishing the amplitude of α properly we can make f as smooth as we like at $x_1 = -L$. The set $[-L \leq x_1 \leq 0]$ will be divided into infinitely many strips in which the optimal ϕ will be alternately ± 1 . When $x_1 \leq -L$, the optimal ϕ is arbitrary. Two possibilities for "optimal paths" are shown at (c) and (d), either being attainable by a suitable spacing of the zeros of α .

We see that when x reaches $[x_1 = -L]$ that we have a reasonably playable situation in neither case. In case (c), a forward time derivative exists and the path is a solution of (1). In case (d), even this is not so.

However the use of K-strategies eliminates these troubles and, indeed, appears mandatory.

* In an example, when we encounter a vector with a single component, we shall omit the subscript.

It is clear that the value is $-x_1$, the x_1 referring to the starting point. The reader will see that a K-strategy based on such an optimal ϕ as we have described and with the t_k spaced with sufficient closeness will yield a payoff arbitrarily close to the value.

Remark that, despite the pathology of the strategies, the value is a perfectly smooth function of the starting position.

The next example illustrates how a K-strategy can be optimal (that is, we can directly play the optimal tactic), but only after a change of navigation variable has been made.

Example 4.

$$\begin{aligned} \text{K. E.:}^* \quad \dot{x}_1 &= \cos(\phi - x_3) \\ \dot{x}_2 &= \sin(\phi - x_3) \\ \dot{x}_3 &= -1 \end{aligned}$$

$$\text{C: } x_3 = 0, \quad \text{E: } x_3 > 0, \quad H = -x_1.$$

It is clear that P will strive to make \dot{x}_1 as large as possible; thus the optimal $\phi = x_3$. But this value will be constantly changing on any path and play according to a K-strategy is not direct. But let us make the replacement

$$\phi' = \phi - x_3.$$

Then the optimal ϕ' is always 0, and a K-strategy now achieves the value of the game.

The next example introduces a new phenomenon.

* We shall use such obvious abbreviations in the future.

Example 5. E: $-1 \leq x \leq 1$, C: $x = 1$ and -1 , Payoff = time of termination, K.E.: $\dot{x} = \phi + v \psi$ ($|\phi| \leq 1$, $|\psi| \leq 1$, and v is constant, $0 < v < 1$).

If the starting $x = x^0$ is positive, P will push x to 1, and E will endeavor to retard him. Thus the optimal $\phi = 1$, $\psi = -1$. The signs reverse if $x^0 < 0$.

What of the puzzling situation of $x^0 = 0$? Each player will wish his navigation variable to be of sign opposite that of his opponent's. Neither player can obtain the requisite information until a small but positive time has elapsed after starting.

The course is clear. The players will use mixed strategies. No havoc is wreaked on our concepts as long as this mixing - as is here the case - applies to one instant of the play only.