Some Aspects of Nonzero-Sum Differential Games

John J. Lawser and Richard A. Volz
Department of Electrical Engineering
The University of Michigan
Ann Arbor, Michigan

Dick, it has been fun thinking about the good old days and I thank you for all of your help. What I learned from you has had a major impact on my work over the years and I thank you for your dedication and example.

John

First International Conference on Differential Games, September 1969, University of Massachusetts, Amherst

SOME ASPECTS OF NONZERO-SUM DIFFERENTIAL GAMES*

John J. Lawser and Richard A. Volz
Department of Electrical Engineering
The University of Michigan
Ann Arbor, Michigan

ABSTRACT

In this paper some definitions from classical game theory and optimal control theory are used to define various aspects of differential games. Several examples are given to clarify the definitions and to illustrate some interesting properties of nonzero-sum differential games. One example was programmed on a hybrid computer and some experimental results are included.

INTRODUCTION

The study of differential games was initiated by Isaacs in the early 1950's. Except for a series of Rand Corporation Memos [3] issued in 1954 and 1955, however, his work was not published until his book [4] appeared in 1965. Isaacs was concerned with continuous games of pursuit, i.e. games in which one player, called the pursuer, tries to minimize a functional and the other player, called the evader, tries to maximize the same functional. He made no attempt to develop a rigorous theory; instead he gave formal solutions for specific problems which illustrated questions the theory should answer. He rationalized that future investigators would add rigor to his work. Indeed, this has been the case; most of the papers written on differential game theory have dealt with questions raised by or related to the problems studied by Isaacs. Even so, there is still no satisfying theory for this class of problems. Moreover, pursuit-evasion games represent only a small fraction of what should be called differential game theory.

Case [1] was one of the first authors to consider differential games with more than two players. His work is based on the equilibrium point solution. Starr and Ho [6] have pointed out that there are many other interesting features of nonzero-sum differential games. Contrary to the situation in pursuit-evasion games in which a saddle point is generally accepted as a solution to the game, the nonzero-sum game admits to no simple definition of a solution. Indeed, the equilibrium point is but one of many points one might consider as a solution.

This question of what is a solution to the nonzero-sum game has also reared its head in discrete games and has been discussed at length in this context [5], [8]. Starr and Ho[6] have made one of the first efforts at examining this question for differential games. In this paper we discuss the nature of a "solution" to the nonzero-sum differential game, extend many of the ideas for solution for discrete games to differential games, and introduce several concepts involving dynamic bargaining, threats, cooperation and external competition among players. In addition, an example game has been programmed on a hybrid computer and some preliminary tests run.

DEFINITION AND SOLUTIONS OF A DIFFERENTIAL GAME

Let

$$\dot{x} = f(x,t,u_1,\cdots u_N) \quad x(t_o) = x_o \tag{1}$$

be a vector differential equation that describes the dynamics of a system controlled by N players (identified $1,2,\ldots,N$) who each choose a control $u_i(t)$ from restraint sets Ω_i. The game starts at $t = t_o$ and ends when $t = t_f$, a fixed terminal time, or when $x(t) \in C(t)$, a terminal constraint set. Each player receives a payoff of the form

$$J_i = \int_{t_o}^{t_f} g_i(x,t,u_1,\ldots,u_N)dt + K_i(x(t_f)) \tag{2}$$

The object of each player is to maximize his payoff. Let J be a vector with entries J_i.

Definition 1. The triple (eq(1),N,J) is said to be a differential game in extensive form.

The term extensive form is used because the differential equation defines how the game evolves in the same sense that a game tree [5, Chp. 3,] defines how a finite stage game evolves. If we look at the evolution of the game at some time t, we will refer to eq.(2) with t_f replaced by t as the current value of the payoff.

If N = 1 then Def. 1 is interpreted as an optimal control problem. There is then no question about what constitutes a solution for this problem (though there may be questions of existence and uniqueness); an optimal control for the player is simply one that maximizes his payoff J.

If N = 2 and $J_2 = -J_1$ for all u_1, u_2 and the goal of player I(II) is to maximize his payoff J_1 (loss J_1), then Def. 1 defines the two person zero-sum game studied by Isaacs. In this case there is also an obvious definition of a solution.

Definition 2. If controls u_1^* and u_2^* exist such that

$$J_1(u_1,u_2^*) \leq J_1(u_1^*,u_2^*) \leq J_1(u_1^*,u_2)$$

for all admissible u_1,u_2, then the pair of controls (u_1^*,u_2^*) is said to be a saddle point for the game.

This definition of a solution agrees well with our intuition; if a saddle point (u_1^*,u_2^*) exists, then a player can only hurt himself by departing from his saddle point control if the other player uses his saddle point control.

If $N \geq 2$ (and $J_1 \neq -J_2$ for N = 2) then it is no longer clear what we should call a solution. The most pessimistic viewpoint a player can take is that the other players are out to hurt him, that they will disregard their own payoffs and will seek to minimize his payoff. Thus, for each player we de-

*Supported in part by the Army Research Office, Durham, N. C. under contract DA-ARO-31-124-G767.

fine a maxmin or security level payoff.

Definition 3. The i^{th} player's security payoff, or maxmin payoff, denoted by J_i^s is given by

$$J_i^s = \max_{u_i \in \Omega_i} \left\{ \min_{u_1 \in \Omega_1} \cdots \min_{\substack{u_j \in \Omega_j \\ j \neq i}} \cdots \min_{u_N \in \Omega_N} \right\} J_i \qquad (3)$$

Another possible definition of a solution [1] can be obtained by generalizing the idea of the saddle point mentioned earlier.

Definition 4. If controls u_i^e exists for each player such that

$$J_i(u_1^e,\cdots,u_i^e,\cdots,u_N^e) \geq J_i(u_1^e,\cdots,u_{i-1}^e,u_i,u_{i+1}^e,\cdots u_N^e)$$

for $i = 1,\cdots,N$ and all admissible u_i, then $(u_1^e,\dots,u_N^e)$ is called the equilibrium solution and $J_i^e = J_i(u_1^e,\cdots u_N^e)$ is called the equilibrium payoff.

One attribute of the equilibrium solution is that if $N = 2$ and the game is zero-sum the equilibrium solution is identical to the saddle point solution. However, in the nonzero-sum game direct competition between the players need not occur. Hence the existence of an equilibrium solution $(u_1^e,\cdots,u_N^e)$ does not preclude the existence of controls $(u_1',\cdots,u_N')$ such that

$$J_i(u_1',\dots,u_i',\dots,u_N') > J_i(u_1^e,\dots,u_N^e) \quad \text{for } i=1,\dots,N \qquad (4)$$

Thus by cooperating all the players may benefit. Whether this will indeed occur is a subject for further research.

We now consider two examples to illustrate these definitions and point out some of the difficulties of considering the equilibrium point as a solution. The first example is actually a discrete game. However it serves to illustrate some of the difficulties and is used in the discussion of dynamic bargaining in the next section.

Example 1. The payoff matrix below is an example of the well-known prisoner's dilemma game.

	C_2	D_2
C_1	5,5	-10,10
D_1	10,-10	-5,-5

(Player 1 chooses a row strategy, player 2 a column strategy; the entries in the matrix give the payoff to player 1 and player 2, respectively). Since each player's D strategy dominates his C strategy (i.e., gives a higher payoff for either of the other player's strategies), the equilibrium strategy for each player is his D strategy and both the equilibrium and security payoff are -5 each. However, if both players depart from their equilibrium strategies and play C the payoff is +5 each and both benefit. Thus there is difficulty in accepting the equilibrium strategies as the "solution" to the game.

Example 2. Now consider a continuous version of the above game in which the matrix entries are payoff rates. Let the players choose constant controls 0 or 1 and the payoffs be

$$V_i(u,v) = \int_0^1 F_i(u,v)dt \quad i=1,2 \qquad (5)$$

where $V_1(V_2)$ is the payoff to player 1(2), and

$$\begin{aligned} F_1(u,v) &= \mathrm{sgn}(1-u)\{5\ \mathrm{sgn}(1-v) - 10\ \mathrm{sgn}(v)\} \\ &\quad + \mathrm{sgn}(u)\{10\ \mathrm{sgn}(1-v) - 5\ \mathrm{sgn}(v)\} \\ F_2(u,v) &= \mathrm{sgn}(1-u)\{5\ \mathrm{sgn}(1-v) + 10\ \mathrm{sgn}(v)\} \\ &\quad + \mathrm{sgn}(u)\{-10\ \mathrm{sgn}(1-v) - 5\ \mathrm{sgn}(v)\} \end{aligned} \qquad (6)$$

Reasoning as in example 1 it can be argued that the equilibrium controls are u=1 and v=1. Similarly, if both players depart from their equilibrium, they both gain. The same dilemma would seem to exist here as in the discrete game. However, as will be discussed later, here there are opportunities available which were not available with the single play of the discrete game.

It is clear from the foregoing discussion that the notion of a maxmin payoff or equilibrium solution is not entirely satisfactory; there generally exist solutions in which all players can improve their payoff over these points. It seems more appropriate to look at sets of solutions which in some sense are "reasonable." In the following paragraphs we explore this further and introduce more explicit definitions of reasonableness. We begin by stating:

Definition 5. The set of attainable payoffs, $\wp$ for a differential game is the set of all payoffs $(J_1,\dots,J_N)$ which result from admissible controls.

To illustrate this, the set of attainable payoffs for Example 2 is shown in Fig. 1. In this case the security payoff and the equilibrium solution are identical. Similarly, during the play of a game it is important to know what payoffs are possible given the past history of the game. Let $\hat{u}_i(t)$, $i=1,\dots,N$ be a given set of functions which are restrictions of admissible controls to the interval $[t_o,\tau]$.

Definition 6. The conditional attainable payoff set $\wp(\tau/\hat{u})$ is the set of all payoffs $(J_1,\dots,J_N)$ which can result from admissible controls $u_i(\tau) \in \Omega_i$ subject to the conditions $u_i(t) = \hat{u}_i(t) \quad t \in [t_o,\tau]$, $i=1,\dots,N$.

One set defined in classical game theory which might be considered reasonable is the subset of attainable payoffs $P \subset \wp$ such that no player can receive a higher payoff except at the expense of other players.

Definition 7. The Pareto optimal set P is the set

$$P = \{(J_1,\dots,J_N):(J_1,\dots,J_N) \in \wp \text{ and } \nexists (J_1',\dots,J_N') \in \wp \ni (J_1',\cdots,J_N') \neq (J_1,\cdots,J_N) \text{ and } J_i' \geq J_i \text{ for } i=1,\cdots,N\}$$

This is also shown in Fig. 1. One can argue that no player should settle for less than his security payoff. This leads one to consider that portion of the Pareto optimal set which dominates the security payoff. This was introduced by von Neumann and Morgenstern[8] for discrete two player games. It seems reasonable to consider it here as well.

Definition 8. The negotiation set η is the set

$$\eta = \{(J_1,\dots,J_N): (J_1,\dots,J_N) \in P \text{ and } J_i \geq J_i^s \quad i=1,\dots,N\}$$

Similarly, we feel one should consider that portion of the Pareto optimal set which dominates the equilibrium payoff as reasonable potential solutions.

Definition 9. The dominant negotiation set η' is

$$\eta' = \{(J_1,\dots,J_N): (J_1,\dots,J_N) \in \eta \text{ and } J_i \geq J_i^e,\ i=1,\dots,N\}$$

Solving for these various points or sets is a difficult task. In general, even existence theorems are not known. Only in some special cases, are results available. Starr and Ho [6] have determined the maxmin and equilibrium controls and controls that lead to the Pareto optimal set for linear games with quadratic costs. From these the Pareto optimal, negotiation and dominant negotiation sets can be determined. Da Cunha [2] has obtained necessary conditions for vector valued performance measures which may be applied, and Case [1] has studied the problem of finding equilibrium points in detail. Rather than proceed along these lines, we shall examine various philosophies which the players may take toward choosing their strategies.

DYNAMIC BARGAINING AND COMPETITION

In this discussion we assume that each player has complete knowledge of the system dynamics and the evolution of the payoff vector up to the present time. We assume further that he can compute instantaneously the equilibrium point control and controls to reach the Pareto optimal set, and is able to tell whether or not the other players are following any of these controls.

In a single play of the simple matrix game of example 1, each player must make an open loop decision for his control. There is no opportunity for a player to observe the progression of the game and modify his control accordingly, i.e. use a feedback law for the control. In the differential game with the above assumption, however, this opportunity is present, and may offer a way around the prisoner's dilemma effect noted in examples 1 and 2. Since there is no generally accepted unique solution to the nonzero-sum differential game, however, it appears that there is at present no way to specify a deterministic feedback law for the players. Rather, we will look at various philosophies the players may adopt in different situations. In particular, we shall be interested in situations in which one player uses his control to try to induce the other player to adopt a control more favorable to himself (or possibly to all players). We refer to this type of behavior as dynamic bargaining since it takes place as the game progresses. We shall also be interested in situations in which some players are interested in beating the others, that is, they place a value on competition beyond their interest in the game payoff.

A first type of dynamic bargaining occurs when, as in examples 1 and 2, it is possible for all players to improve their equilibrium or security payoffs by cooperating. In this situation, it is possible for one player to take the lead and hint at cooperation by adopting, for a short time, a policy more favorable to the other players. By so doing, he hopes to induce them to adopt controls more favorable to him.

To illustrate this consider the game of example 2. It was seen that the equilibrium controls were u=1 and v=1, and the corresponding payoffs were $J_1=J_2=-5$, as shown in Fig. 1. Yet if both players could be induced to use a control of 0, they could both benefit. If both players are initially using their equilibrium controls either has an opportunity to hint at cooperation by adopting the control 1 for a short period of time. This will give the other player a temporary gain at the expense of himself. If the other player takes the hint, both will then gain. If he doesn't, the first player may return to his equilibrium control. Since he can make his use of the control 0 very short in duration, little relative loss occurs.

A similar situation occurs in repeated plays of the discrete matrix game of example 1, called a supergame. What is a solution for the supergame? Offhand one would expect two rational players eventually to lock in on the C_1, C_2 strategies and be content to collect 5 units/play. However, arguments can be presented against this "solution" [5]; it is easy to recognize the opportunity for one player to "double cross" another. In [7] experimental results of the supergame are described. The subjects were paid a nominal amount which was modified by the net payoff of the supergame. The most frequent lock in strategy in many runs of the supergame was the C_1, C_2 point but occasionally two players settled on the D_1, D_2 point.

Similar experimental tests with a differential game are not yet available. It is clear, however, that in either the differential game or the supergame, the player who first hints at cooperation must be willing to make some initial sacrifice. If the number of plays of the discrete game is small, the initial sacrifice of the player may be a significant portion of the total payoff and one wonders whether or not a player might be reluctant to take the step away from the equilibrium. With the continuous game this reticence is less likely to occur.

Hinting at cooperation is not the only type of dynamic bargaining which can take place in example 2. A notion of threatening can also be introduced. Suppose in this example that both players are using the control zero. If one player tries to double cross the other to try to achieve a higher payoff rate, the second player can threaten him by temporarily returning to his equilibrium solution also. If the first player refuses to acknowledge the threat, the second has little choice but to remain at the equilibrium. There is no advantage for the first player to be stubborn, however, and if he is assumed to be rational, he should return to the zero control.

The extent to which hinting at cooperation or threatening is a reasonable philosophy for a player is largely determined by the nature of the game. In some cases, one player is completely at the mercy of the other. The following example illustrates this situation and the futility of hinting at cooperation.

Example 3. Consider the payoff matrix [5,p.139] below, with the entries being payoff rates.

	β_1	β_2
α_1	1,4	-1,-4
α_2	-4,-1	4,1

One might expect that due to the symmetry of the payoff matrix the players would agree to coordinate the game so each could collect 5/2 units per unit time. But the column player has an advantage; if he threatens to always play β_1 what can the row player do but play α_1? One can envision the row player trying to "teach" the column player to alternate his control so he (the row player) could accumulate at the rate of 4 units/time unit for at least part of the game. But would this happen? We think not, for the column player should always play β_1. If the row player adopts α_2 or switches to α_2 during play of the game he is hurt more than the column player, soon gives up and returns to α_1.

Another aspect of playing the game is competition. It has been suggested that if a player expects the remaining players to choose controls to hurt him he should choose his maxmin or security control. This thought is not really in keeping with the object of the game; each player is to maximize his payoff, not try to hurt the other players. Yet it does point out an interesting attitude which may prevail, and was, in fact, observed in some of the experimental tests. Even though the stated object of the game is for each player to maximize his payoff, players may play the game not only to maximize their own payoff, but to beat the other player as well, in spite of the fact that there may be a certain competitive feature built into the original game. In effect they are modifying their payoff to include a term reflecting the amount by which they "beat" their opponents. As a first approximation to modeling this situation, one can define a modified game.

Definition 10. Let

$$J' = MJ$$

where $M = [m_{ij}]$ and

$$m_{ij} = \begin{cases} -\alpha_{ij} & i \neq j \\ 1 + \sum_{\substack{k=1 \\ k\neq i}}^{N} \alpha_{ik} & i = j \end{cases} \tag{8}$$

and J is as in Def. 1. The triple (eq.(1),N,J') is said to be a modified game.

The α_{ij} are called coefficients of competitiveness which reflect the desire of the ith player to beat the jth player. This, of course, leads to new security and equilibrium payoffs and to new dominant negotiation sets. In order to see the effect that playing the modified game has on the payoffs to the original game one would like to reflect the controls leading to the modified equilibrium point and dominant negotiation set back to the original payoff and compare the set of payoff vectors thus obtained with the equilibrium and dominant negotiation sets of the original game. Let η'_r be the reflection of the dominant negotiation set of the modified game. From the definition of η' it is clear that η' dominates η'_r. Furthermore, in general, equality cannot hold. Thus, at many points the dominance will be strict. This means that the players, by introducing competition external to that inherent in the game may, in many situations, actually reduce their payoffs to the original game.

This line of thought leads to another interesting question. Suppose that the competetiveness coefficients are known and a game A is given. How does one design a pseudo-game B so that when the players are presented with the modified game B they actually play game A? For the simple mechanism of introducing the external competition given here, one would choose the payoffs for the psuedogame B as

$$J' = M^{-1}J \tag{9}$$

When the players play the game as modified in eq. (9) they then are actually playing the given game (assuming M is nonsingular).

SIMULATION OF A TWO PERSON NONZERO-SUM GAME

In order to investigate some of the ideas discussed above a two player nonzero-sum differential game was simulated on a hybrid computer. The game considered is given by

$$\dot{x} = u_1 + u_2 \quad x(0) = 0 \quad 0 \leq u_i \leq 1 \quad i=1,2 \tag{10}$$

$$J_i = \frac{1}{T}\left\{\int_0^T -u_i^2\, dt + x(T)\right\} \quad i=1,2 \tag{11}$$

Each player benefits from having the terminal payoff x(T) as large as possible and each is penalized for his control effort in enlarging x. It can be shown that the security controls for the game are $u_1^s(t) = 1/2$ and $u_2^s(t) = 1/2$ and the security payoff to each player is $J_1^s = J_2^s = 1/4$. The equilibrium controls of the players are the same as the security controls, but the equilibrium level payoff is 3/4 each, (i.e., $u_1^e(t) = u_2^e(t) = 1/2$ and $J_1^e = J_2^e = 3/4$.) However, if both players cooperate, deviate from their equilibrium controls and use $u_1(t) = u_2(t) = 1$ the payoff is 1 each.

The set of attainable payoffs $\mathcal{P}$ for the game is shown in Figure 3. The upper boundary of $\mathcal{P}$ is given parametrically by

$$J_1 = \sigma \qquad 0 \leq \sigma \leq 1$$

$$J_2 = 1 + \sigma - \sigma^2$$

Since the game is symmetric, reversing the payoff indices describes the right hand boundary of P. The equilibrium payoff, security payoff, Pareto optimal set and dominate negotiation set are also indicated in Figure 2.

Given that the game has been played to time t, $0 \leq t \leq T$, the payoffs at the end of the game can be written

$$J_i(T/t) = \frac{1}{T}\left\{\int_0^t -u_i^2(\tau)d\tau + x(t)\right\} + \frac{1}{T}\left\{\int_t^T -u_i^2(\tau)d\tau + x(T)-x(t)\right\} \quad i=1,2 \tag{12}$$

The first term depends only on $u_i(\tau)$ for $\tau \leq t$ and in a sense represents the payoff accumulated at time t. It is easy to see that the set of possible values of the second terms in eq. (12) is the attainable payoff set for the game reduced in size by the factor $\frac{T-t}{T}$. The conditional payoff set for the game is just this reduced set translated by the first term in eq. (12). Thus, as the game is played the set of attainable payoffs translates according to how the payoff builds up and shrinks at a constant rate.

The above game was programmed on an Applied Dynamics AD/4 hybrid computer. The players each viewed an oscilloscope screen on which was displayed the current value of the payoff, the projected payoff assuming they maintained a control constant at the present level, and the attainable set conditional upon proceeding from the present point. From the projected payoff each player could see the effects of any changes in control by either player. All controls which result in payoffs on the curved portion of the boundary of the attainable set are reached by the use of constant controls; thus, the means of projecting the final payoff is reasonable. The players selected their controls by varying a potentiometer setting.

The center trace in the photograph in Figure 3 is a multiple exposure of the oscilloscope trace (taken at t = 0, .2T, .4T,...) with both controls set on the equilibrium control 1/2. The left trace is a similar multiple exposure with both controls set to 1. The trace on the right was obtained by setting one player's control to 1/2 (equilibrium) and making exposures with the other player's control set to 0, .2,..., 1. (restarting the game for each new setting). Then this process was repeated with the other player's control set to 1/2. The sequence of dots in the photograph clearly illustrates the meaning of an equilibrium solution.

Limitations on time and resources precluded playing games with a sufficient number of players to achieve any statistically significant result regarding the types of strategies used. However, a small group of people, all of whom understood the dynamics of the game, did play the game. All plays of the game were begun with the equilibrium control. Since funds were not available to provide cash payoffs to all of the players, two small prizes of equal value were awarded to the players with the highest total score on two runs. Interestingly enough most of the philosophies of play indicated earlier did appear in the various plays.

Figure 4 shown the controls used by two players as a function of time to go (i.e., t=0 is on the right). Note how player B tried to encourage player A to move off the equilibrium control at points (in time) 1 and 2. At point 5 it looks like A was ready to cooperate; note how quickly B started to follow A's lead. At 4 B made a small threat. Soon after the threat, at 5 they began to cooperate; from 5 to 6 both increased their controls steadily. Note, however, how close their controls are for any point in time; this means the projected value of the game would be on the 45° line of equal payoff to each player. At 7 A began to back off and B, when he noticed A's action also began to back off. In a second play of the game player B achieved much the same results as in the above run.

It is interesting that in both the above runs B "lost" the run by a small amount .9073 to .9000 and .9870 to .9584. However, his total for the two runs was far above that of the other participants.

The run shown in Figure 5 illustrates the effect of a competitive player vs one who tries to hint at cooperation. Both players began to cooperate at the start of the play. At 1 D was close to using his full control. E tried to move further up on the payoff set to beat D by reducing his control. Eventually E realized that both players were driving themselves, toward a lower final payoff and began to cooperate again. As soon as a reasonable level of cooperation was reached, player E again tried to achieve an advantage by backing off on his control. This is particularly evident during the time between points 3 and 4. The result was a cyclic effect in the controls used. Throughout the play it was player D who made the opening moves toward cooperation and player E who tried to gain competitive advantage over D by reducing his control. At 7 E reduced his control to zero, reducing his rate of accumulating payoff somewhat but hurting D much more. E's payoff for the game was .8087, and D's was .6759. E obtained his competitive objective of beating D. But notice that both payoffs were significantly below that obtained by players who cooperated more willingly. It seems likely that if player E had played someone as compeitive as himself they would have driven themselves to a much lower payoff.

SUMMARY

A differential game was defined as a game whose extensive form description is given by a set of integro-differential equations. This emphasized that differential games are a part of what might be called classical game theory. However, differential games can also be considered as generalized optimal control problems and much of the formulation reflected this. Various points and sets which may be considered as "solutions" to the game were defined.

The concepts of dynamic bargaining and competition were discussed. It was shown that in nonzero-sum differential games it is frequently possible for a player to hint at a cooperative solution which would benefit several players. It was also shown that it is sometimes possible to make threatening moves to try to induce other players to adopt a mutually beneficial control. The notion of players artificially introducing an external weighting on competitiveness was introduced, and it was pointed out that in many cases this may actually hurt the final payoff of the players involved. These ideas were illustrated by simulating a two player game on a hybrid computer and having a limited number of people play the game.

ACKNOWLEDGEMENTS

The authors are grateful to R.M. Thrall for his suggestions and advice and to the Simulation Center of the University of Michigan for the use of their hybrid computer.

REFERENCES

1. Case, J. H., Equilibrium points of N-person differential games, Doctoral Thesis, The University of Michigan, 1967.
2. Da Cunha, Nelson O., Constrained Minimization Under Vector-Valued Criteria, Doctoral Thesis, University of California, Berkeley, 1966.

3. Isaacs, R., Differential Games I, II, II, IV, V, the Rand Corporation Research Memoranda RM-1391, 1399, 1411, and 1486, 1954, 1955.

4. Isaacs, R., Differential Games, John Wiley, New York, 1965.

5. Luce, R. D., and Raiffa, H., Games and Decisions John Wiley, New York, 1957.

6. Starr, A. W., and Ho, Y. C., "Nonzero-sum differential games," Technical Report No. 564, Division of Engineering and Applied Physics, Harvard University, Cambridge, Massachusetts, May 1968.

7. Rapoport, A. and Chammah, A. M., Prisoner's Dilemma, University of Michigan Press, Ann Arbor, Michigan, 1965.

8. Von Neumann, J. and Morgenstern, O., Theory of Games and Economic Behavior, Princeton University Press, Princeton, New Jersey, 1946

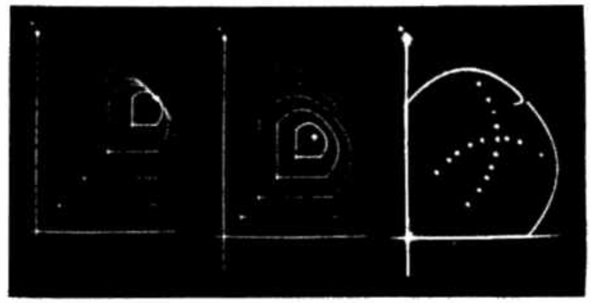

Figure 3. Oscilloscope traces of simulated game.

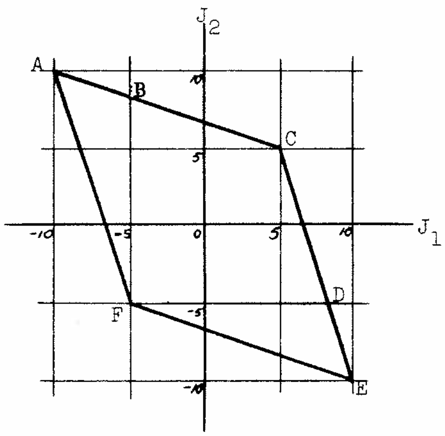

Security payoff F
Pareto optimal set ABCDE
Negotiation & dominant negotiation sets BCD

Figure 1. Set of attainable Payoffs 𝒫 for example 2.

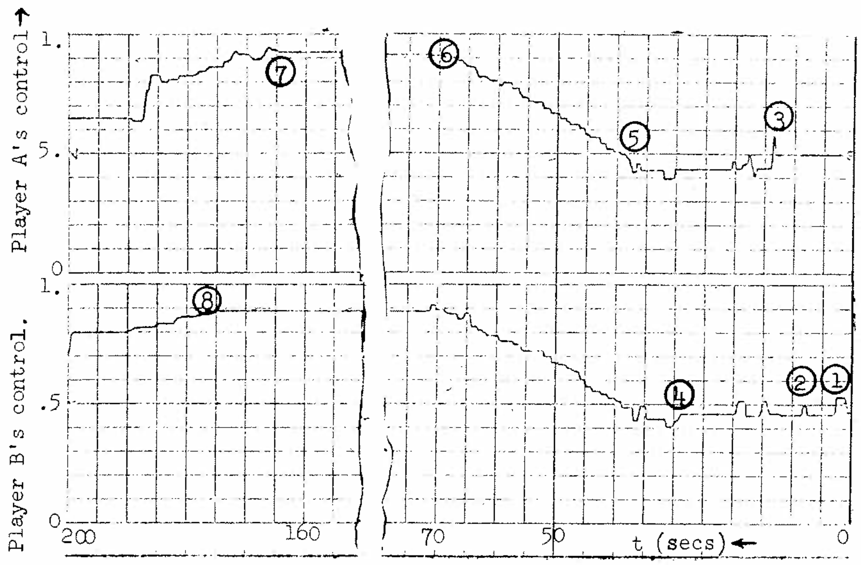

Figure 4. Controls used used by players A and B

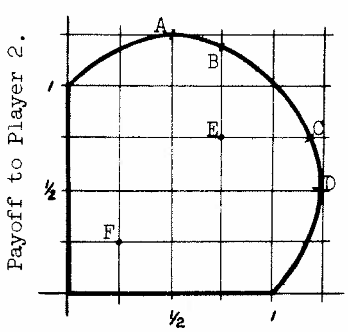

Payoff to Player 1.

Pareto optimal set ABCD
Negotiation set ABCD
Dominant negotiation set BC
Equilibrium payoff E
Security payoff F

Figure 2. Set of attainable payoffs 𝒫 for simulated differential game.

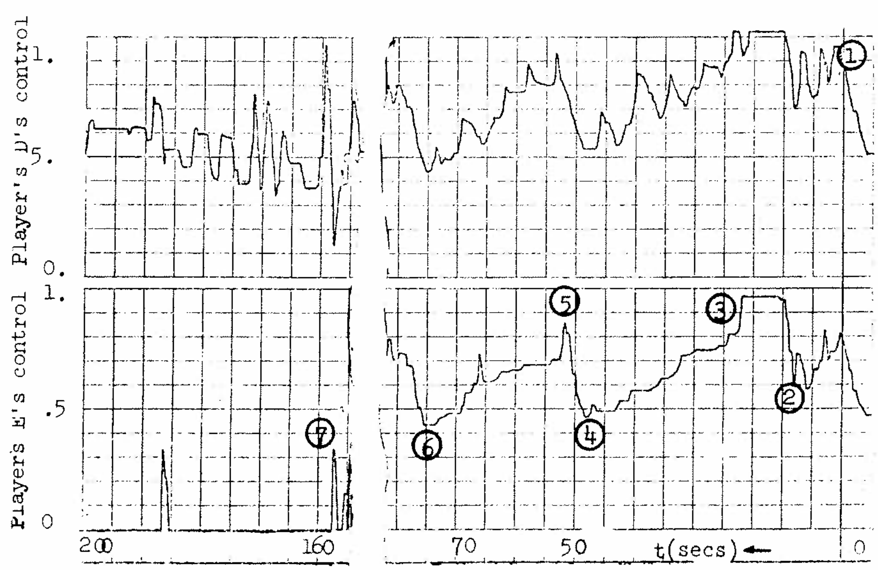

Figure 5. Controls used by players D and E