

AD-A062 104

INSTITUTE FOR DEFENSE ANALYSES ARLINGTON VA PROGRAM --ETC F/G 15/7
DECISION MODELING IN LARGE SCALE CONFLICT SIMULATIONS.(U)
OCT 78 L B ANDERSON


UNCLASSIFIED

P-1355


IDA/HQ-78-20580

NL

| OF |
AD
A062104



END
DATE
FILMED
3-79
DDC



MICROCOPY RESOLUTION TEST CHART
 NATIONAL BUREAU OF STANDARDS-1963-A

(12) ^{NEW} LEVEL II

IDA PAPER P-1355

DECISION MODELING IN LARGE SCALE
CONFLICT SIMULATIONS

Lowell Bruce Anderson

DDC FILE COPY

ADA062104

DDC
RECEIVED
DEC 14 1978
B

October 1978

DISTRIBUTION STATEMENT A
Approved for public release
Distribution Unlimited

78 11 27 075



INSTITUTE FOR DEFENSE ANALYSES
PROGRAM ANALYSIS DIVISION

The work reported in this document was conducted under IDA's Central Research Program. Its publication does not imply endorsement by the Department of Defense or any other government agency, nor should the contents be construed as reflecting the official position of any government agency.

This document is unclassified and suitable for public release.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER IDA PAPER P-1355	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) DECISION MODELING IN LARGE SCALE CONFLICT SIMULATIONS	5. TYPE OF REPORT & PERIOD COVERED FINAL	
	6. PERFORMING ORG REPORT NUMBER P-1355	
7. AUTHOR(s) Lowell Bruce Anderson	8. CONTRACT OR GRANT NUMBER(s) IDA Central Research Program	
9. PERFORMING ORGANIZATION NAME AND ADDRESS INSTITUTE FOR DEFENSE ANALYSES 400 Army-Navy Drive Arlington, Virginia 22202	10. PROGRAM ELEMENT PROJECT TASK AREA & WORK UNIT NUMBERS N/A	
11. CONTROLLING OFFICE NAME AND ADDRESS	12. REPORT DATE October 1978	
	13. NUMBER OF PAGES 80	
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)	15. SECURITY CLASS (of this report) UNCLASSIFIED	
	15a. DECLASSIFICATION DOWNGRADING SCHEDULE N/A	
16. DISTRIBUTION STATEMENT (of this Report) This document is unclassified and suitable for public release.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Decision Modeling, Combat Simulation, Models of Conflict, Game Theory		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This paper presents various options for modeling decisions in conflict simulations, and discusses the advantages and limitations of using these options in large scale simulations. It is argued that there are many decisions that must be modeled and that how these decisions are modeled can significantly affect the results of a simulation. It is important to be aware of the characteristics of the various options that are available for making decisions in combat situations.		

DDC
RECEIVED
DEC 14 1978
RECEIVED

B

DD FORM 1473 JAN 73 EDITION OF 1 NOV 68 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

14

IDA PAPER P-1355

6

DECISION MODELING IN LARGE SCALE
CONFLICT SIMULATIONS,

9

Final repl.,

10

Lowell Bruce/Anderson

11

October 1978

12

487.

18

IDA/HQ,
SBIE

19

78-20580,
AD-E500 032



INSTITUTE FOR DEFENSE ANALYSES
PROGRAM ANALYSIS DIVISION
400 Army-Navy Drive, Arlington, Virginia 22202

IDA Central Research Program

403 219

mt

ACCESSION for	
NTIS	Write Section <input checked="" type="checkbox"/>
DDC	Brief Section <input type="checkbox"/>
UNANNOUNCED	<input type="checkbox"/>
JUSTIFICATION	
BY	
DISTRIBUTION/AVAILABILITY CODES	
Dist.	AVAIL. and/or SPECIAL
A	

CONTENTS

FOREWORD v

ACKNOWLEDGMENT vii

A. INTRODUCTION 1

B. DECISIONS 1

 1. A Context for Decision Modeling 1

 2. Examples of Decisions That Must Be Modeled 4

 a. Small Scale Simulations 5

 b. Intermediate Level Simulations 5

 c. Large Scale Simulations 7

 3. Impact of Decision Modeling 8

 4. Decisions and Technical Details 11

C. ALTERNATIVE APPROACHES TO DECISION MODELING 12

 1. Optimizing "Can Do" Methods 14

 2. Non-Optimizing "Can Do" Methods 16

 3. "Will Do" Methods 17

 4. Simple Methods 18

 5. Modeling the Human Mind 20

 6. Using Humans as Models 20

D. "CAN DO" METHODS VERSUS "WILL DO" METHODS 21

E. CONCLUSION 27

APPENDIX

Extract From: METHODOLOGY FOR USE IN MEASURING THE EFFECTIVENESS OF GENERAL PURPOSE FORCES



FIGURES

1	Two Simple Models of Attrition	3
2	Potential Missions for Combat Aircraft	9
3	Alternative Methods for Making Decisions in Combat Simulations	13

FOREWORD

This paper presents various options for modeling decisions in conflict simulations, and discusses the advantages and limitations of using these options in large scale simulations. It is argued that there are many decisions that must be modeled and that how these decisions are modeled can significantly affect the results of a simulation; and so it is important to be aware of the characteristics of the various options that are available for making decisions in combat situations.

This paper was presented as part of a seminar series on Operations Analysis and Force Planning at the Hochschule der Bundeswehr München in the Federal Republic of Germany.

ACKNOWLEDGMENT

Discussions with Dr. Jerome Bracken, Mr. Edward Kerlin, Dr. Roy Shanker (now at Resource Planning Associates, Inc.), and many other members of IDA contributed significantly to the formulation of the ideas presented in this paper. The opinions and judgments presented here are those of the author.

PRECEDING PAGE BLANK-NOT FILMED

A. INTRODUCTION

The purpose of this paper is not to explain in detail a few "right" or "preferred" ways to model decision making in large scale conflict simulations. Instead, the purpose is (1) to argue that there are important decisions that must somehow be modeled, and (2) to present some options for modeling decision making in large scale conflict simulations. The goal here is to discuss various options for modeling decisions so that: (a) in building a model or performing a quantitative analysis, selection among these options will be facilitated, and (b) in evaluating an analysis performed by others, the assumptions and options they used can be better understood.

B. DECISIONS

This paper should begin by addressing two questions: what are large scale conflict situations? And what are the decisions to be modeled in these simulations? Let me start with the "what are the decisions" question, because there are some basic points in answering that question alone. The first point is as follows:

In virtually all simulations of conflict that are complex enough to be useful, there are decisions that must be modeled.

1. A Context for Decision Modeling

To make what I am saying a little more precise, let me first give some examples of models that are simple enough that there are no decisions that must be modeled, then I will give examples of models that require decisions. In giving these examples, I will concentrate on conventional ground and air

combat, especially air combat, but one could easily give examples of decision modeling in naval conflict or nuclear wars.

First, for a model that requires no decisions, consider the "extreme example" of a "model" that says that with probability P_k a particular weapon will be killed during an engagement. There may be many decisions that have to be modeled (or made somehow) to determine what the value of P_k should be; but given this P_k , no more decisions need to be made in this very simple model.

Next, consider a model of an engagement between a number of identical shooters on one side, and a number of identical targets on the other side. (This is a "one-sided" engagement in that only one side can shoot and only the other side can be killed.) Suppose that each shooter selects one of the targets to fire at, independently of all of the other shooters, and according to a uniform distribution over the targets. Suppose further that each shooter then fires one shot at the target he has selected and that the probability of killing that target is P_k (and that only targets fired at can be killed). Then it can be shown that the expected number of targets killed is given by the top equation (Model A) in Figure 1.

Note that assuming that each shooter selects a target independently of the other shooters means that it will sometimes happen that two or more shooters are shooting at the same target, while no shooters are shooting at another target. Now instead of this independent selection of targets, suppose that shooters are assigned targets by a central controller and, for simplicity, assume that there are at least as many targets as shooters. If the central controller assigns each shooter to a different target, then the expected number of targets killed is maximized and Model B with Decision 1 results.

The point I want to make here is that if the shooters have no way to coordinate, then Model A is a direct extension of the

<p>Notation: S = number of shooters (one shot per shooter) T = number of targets P_k = single shot probability of kill E[T_k] = expected number of targets killed</p>	<p>Model A: Shooters Select Targets Independently</p> $E[T_k] = T(1 - (1 - \frac{P_k}{T})^S)$	<p>Model B: Shooters Assigned to Targets by Central Controller</p> <p>(Suppose $T \geq S$)</p> <p>Decision 1 Assign Shooters to Maximize E[T_k]</p> $E[T_k] = P_k S$ <p>Decision 2 Assign Shooters Independently Same Result as Model A</p> <p>Decision 3 Assign Shooters to Minimize E[T_k]</p> $E[T_k] = 1 - (1 - P_k)^S \leq 1$
--	---	--

Figure 1. TWO SIMPLE MODELS OF ATTRITION

one-on-one P_k "mini-model" of the first example, and there are no decisions involved. But if there is a central controller that has the capability to assign any shooter to any target, then there is a decision required. Namely, if there are S identical shooters and T identical targets (with $T \geq S$) and if any of the S shooters can shoot at any of the T targets, then how should the shooters be assigned to the targets? Assigning each shooter to a separate target is one way, and this way maximizes the expected number of targets killed.

Two other decisions that could be made are to assign shooters to targets independently (Model B, Decision 2), which results in the same equation as Model A; or to assign all the shooters to the same target (Model B, Decision 3), which minimizes the number of targets killed (given that all S shooters shoot at one of the targets).

Of course, for Model B, the choice is simple. If the shooters have a central controller that has the capability to make any assignment then, all other things being equal, Decision 1 is clearly the most reasonable decision. But again, the the point of this figure is to point out the difference between assuming that there is no capability to do something different, in which case there is no decision to be modeled, versus the case where there is a capability to do something in several different ways, in which case there is a decision to be made.

2. Examples of Decisions That Must Be Modeled

I have just talked about some very simple models, two of which did not require any decision to be made. Let me now give some examples of more complex (and more realistic) cases where decisions are required. By complex, I don't mean only large scale. One aircraft versus one aircraft certainly is small scale--but a one aircraft versus one aircraft model can be very complex if all the aerodynamic factors and the radars and the

weapon characteristics, and so on, of the two aircraft are considered. On the other hand, Lanchester's differential equations can be used as models of one side's force in a theater versus an opponent's force in the theater. (Homogeneous Lanchester-square differential equations, for example, are $\dot{B}(t) = -k_1 R(t)$ and $\dot{R}(t) = -k_2 B(t)$.) These equations are very simple, perhaps taken by themselves, too simple to be realistic, yet they can be considered as being a large scale simulation in that they attempt to model one force versus another force in a theater. Of course, one could also have simple small scale models, and complex large scale ones--all combinations are possible.

a. Small Scale Simulations

Concerning small scale simulations, consider some of the decisions to be made in modeling one aircraft versus another aircraft: When and which way to turn, when and how much to accelerate and to decelerate, when and how much to change altitude, which weapon should be fired (if there is a choice, say, between infrared and radar missiles) and when to fire, and when to "bug out", i.e., when to try to escape the engagement unharmed, but without killing the opponent.

And all these decisions and many more need to be made in modeling two aircraft versus one enemy. The single aircraft has to decide which opponent to go after first, and when to change and attack the other opponent instead. The two aircraft have many combinations of tactics they could try to employ. Similar decisions must be made for modeling, say, two tanks versus one tank--or m weapons of any type on one side versus n weapons on the other side.

b. Intermediate Level Simulations

The decision level builds when one attempts to model a battle of, say, a division versus an enemy division, where each division has several subunits with different types of weapons and

supporting fires. Or, consider an air raid by bombers and escorts and SAM-suppressors against an enemy force of interceptors, long range SAMs, and short range SAMs. At this level of simulation, one might simplify away the decision of each pilot as to which way to turn during an engagement--and just model the outcome of an engagement by a simple P_k (as discussed earlier). But then new decisions have to be made. For the defenders, where should the long range SAMs be located (clumped up in a belt, or spread out over an area, or some other way), and where should the short range SAMs be located (near the long range SAMs, spread out around all potential targets of the attacker, or concentrated around the highest value targets, for example)? And should SAMs fire all their missiles at the first aircraft they see, or should they hold fire occasionally to wait for possible second waves of aircraft? And should the interceptors be on air patrol or on ground alert, and should interceptors be allowed to operate in the same airspace as defended by SAMs, at the risk of the SAMs shooting down their own interceptors as well as enemy aircraft?

For the attacker, what flight paths and at what altitudes should the escorts and bombers fly? How should the attacker position his escorts relative to his bombers? Should all SAM suppressors attack the first SAMs they locate, or should some stay with the bombers as they penetrate farther to the rear? As the bombers approach the targets, when should they drop their bombs, realizing that the closer in they come, the better their target acquisition and accuracy will be, but the more likely it becomes that they will be shot down by enemy SAMs and AAA? And what about feints and false attacks? Should the attackers send a few bombers one way with the majority of the force going another way? And how should the defender react to a potential feint (one which might be a feint or might be a full attack)? Finally, what targets should the bombers attack? Should they strike deep at airfields or shallow at units in

combat? These are all decisions that have to be made in a simulation of an air raid, which, like a division versus a division, is an intermediate level simulation. But the last decision I mentioned--which targets should the bombers attack--leads to large scale simulations, with correspondingly more sets of decisions to be made.

c. Large Scale Simulations

A large scale simulation of *ground* forces in a conventional war must address such decisions as: which forces should be in reserve and which in combat (as a function of the status of the combat) throughout the war; and of the forces in combat, where should they be, what posture should they be in, when should they change positions, and how should they maneuver. Also, how should each side use supporting forces, especially their close air support and interdiction sorties.

As another example, a large scale simulation of *air* combat might attempt to aggregate many of the decisions (that I just mentioned were required for a one-raid model) into parametric values and equations. But the aforementioned last decision, which targets should be attacked, becomes one part of a major decision in a large scale simulation of air combat. This major decision is:

Which missions should aircraft fly, and which aircraft should fly these missions on which days of the war being simulated?

The answer may seem simple at first. Let the bombers do bombing missions, the escorts do escort missions, and the interceptors do interception missions. Well, there are some aircraft specifically built to be bombers (or light attackers), but what should these aircraft bomb? Should they attack enemy units in combat (that is, fly close air support missions), or bomb units in reserve and logistical areas (interdiction missions), or bomb airbases (airbase attack missions), or bomb the

SAMs that are defending one of these three types of targets (that is, fly SAM-suppression missions in support of close air support, or interdiction, or airbase attack missions)? And virtually every modern aircraft that can fly either escort or defensive missions can also fly the other mission, and can also be loaded with bombs to become a fighter/bomber. And aircraft can switch back and forth between missions at any time. This is summarized in Figure 2.

Of course, how well an aircraft can do on these various types of missions depends on the type of aircraft, the training of the pilot, and the type of equipment and munitions available. But training can be done and munitions and equipment can be bought. The important point is that if an analysis says that an aircraft is flying a particular mission, then this is a result of a decision that has somehow been made or modeled --it is not due to the inherent capabilities of the aircraft involved. I think that this is a very important point, because I believe that the decisions about which missions to fly, on which days of a simulation of combat, can have a major impact on the results of that simulation; and yet, these decisions are often made implicitly, and without justification or discussion of the impact of making alternative decisions.

3. Impact of Decision Modeling

- In virtually all simulations of conflict that are complex enough to be useful, there are decisions that must be modeled.
- How these decisions are modeled frequently has a major impact on the results of the simulation.

This finishes a comment made above (which had only the first point) by adding the second point.

Concerning these two points, let me make the following comments. First, these points apply to all simulations, large scale or not.

Virtually every modern NATO and Warsaw Pact aircraft can fly (or can easily be modified to fly) several of the following missions:

Offensive Missions

Close Air Support -- Direct Attack
Close Air Support -- Escort
Close Air Support -- SAM-Suppression
Interdiction -- Direct Attack
Interdiction -- Escort
Interdiction -- SAM-Suppression
Airbase Attack -- Direct Attack
Airbase Attack -- Escort
Airbase Attack -- SAM-Suppression

Defensive Missions

Battlefield Defense
Rear Area Defense

And many aircraft can fly all of these missions (that is, they can fly any one of these missions on one day and any other mission the next day).

Figure 2. POTENTIAL MISSIONS FOR COMBAT AIRCRAFT

Second, when I say that decisions must be "modeled" I mean this statement in a very general sense. And I mean to include, for example, making decisions by using an interactive simulation, where the simulation poses decisions to be made to humans (who are usually called "players"). These players then make the decisions and the model processes the impact of these decisions, which generally leads to posing more decisions to be made. In this case, the "model of decision making" is the human player. That is, the human players represent (and so are models of) the military commanders and their staffs (or, in a smaller scale simulation, they represent the pilots, the tankers, etc.).

A third comment about these points is that I admit to being vague when I use the words "frequently" and "major impact". I haven't done a statistical survey to determine "how frequently" or "how big of an impact". These statements are my judgments based on my experience with simulations (mostly, but not exclusively, large scale ones). But I should note that I don't mean just that one can invent an implausible set of decision rules that is so bad for one side that that side does much worse than it would have done with a good set of decisions. Nor do I mean that different plausible decision rules can produce different intermediate outcomes but not different final results. I mean that it is my judgment that different decision rules (or decision models) that are plausible, at least by *some* standards, frequently (but not always) have a major impact on the final results and conclusions drawn from using a simulation.

My last comment on these points is that one might say "so what--doesn't everybody believe that these two points are true and isn't the real problem what to do about them?" Well, I agree that what to do about them is a problem, but the first step is to recognize that the problem exists, and I think that it is not widely enough recognized that the problem exists.

I frequently hear about results from analyses of weapons systems where either the existence of decision making is ignored in evaluating these weapons, or it is claimed that the decision making has only a small impact on the results of the evaluation. But, as I've claimed, there always are decisions being made (sometimes implicitly) and quite often the sensitivity of the results to alternative decisions is not tested, so no one knows, not even the user of the model, whether different decisions would have had a major impact or not. My judgment is based on several cases where alternative decision rules have been tested, and frequently they made a difference.

By "frequently" I don't mean always. There are also many cases where, if the decisions are kept within plausible bounds, then the details of these decisions do not have a major impact on the overall results. Sometimes different decisions affect intermediate outcomes but not the final conclusions. But if one does not test different plausible decision rules, at least judgmentally (I'll say more about "judgmental testing" later), then how does one know whether the particular decisions being used are affecting the results or not?

4. Decisions and Technical Details

I gave one example where I strongly believed that different decision rules might likely have an impact on the conclusion of a study: namely, the decision of how to assign multi-purpose aircraft. Other decision rules might not have such an effect on a simulation. Rather than go through a list of examples, let me make a relevant point here. There is a strong tendency in building a simulation to put more and more detail into it, rather than to use average values. And there is a strong tendency in evaluating an analysis to rate it higher, and act on it more strongly, if it includes more detail.

Adding more detail to a simulation usually makes it more saleable in that it adds more "apparent realism". But usually more detail inside a simulation also means that more decisions must be made inside the simulation. And if the effects of these decisions are not tested, then the results of the simulation may, in general, be less realistic because the results could change significantly if different, but still plausible, decisions were made.

In other words, more detail usually means more decisions. If these decisions are tested or are easy to make, then an analysis is better than it would have been without the additional detail. If the results of the simulation change based on the decisions made with the additional detail, and if an analysis reports the impact of these changes, then the analysis is better for it. But suppose an analysis says "Because of the increased detail, we could only examine one (or two) different decisions, and there are other plausible decisions that we did not examine and we do not know if these different decisions would produce different results." Then what good is the increased detail over using average values as inputs?

Just as one has to question a relatively more aggregated model as to the impact of its aggregation of details, so one has to question a relatively more detailed model as to the impact of its internal decisions.

And because most people who build, use, and analyze simulations are trained in technology, not decision modeling, I think that this latter question is not asked often enough.

C. ALTERNATIVE APPROACHES TO DECISION MODELING

So far I've said that there are decisions to be made or modeled somehow in combat simulations, and that how these decisions are made can be very important. Some different ways that one can make decisions in simulations of combat are listed in Figure 3.

<u>Method</u>	<u>Type of Method</u>
1. Make game-theoretic decisions based on some global MOE a. Use optimization methods b. Try a few selected options and pick the best	"Can Do"
2. Make decisions using an algorithm with a local goal a. Pseudo-optimal methods b. Sub-optimal methods c. Methods based on reasonable rules	"Can Do"
3. Make decisions by modeling expected tactics a. Fixed input decisions b. Algorithms based on "will do" rules	"Will Do"
4. Make decisions using simple rules a. Monotonic rules based on some global MOE b. Non-monotonic but conceptually simple rules	Mixed
5. Model human decision processes	Mixed
6. Use humans in interactive roles	Mixed

Figure 3. ALTERNATIVE METHODS FOR MAKING DECISIONS IN COMBAT SIMULATIONS
(Different Methods Could be Used for Different Decisions in
the Same Simulation)

At the outset, let me note that this list is not meant to be mutually exclusive and collectively exhaustive. There will be some grey areas and overlap between the various methods on the list. What I am trying to point out by listing different methods, is that there are different basic ideas on how to model decision making. And different methods can be used to make different decisions in the same simulation.

1. Optimizing "Can Do" Methods

The first approach for making decisions in a combat simulation is to make optimal game-theoretic decisions based on some global measure-of-effectiveness (MOE). By a "global MOE", I mean, in a large scale simulation of a war, that global MOE would be an outcome of the war (such as territory captured, or cumulative exchange ratio, or cumulative NATO minus Warsaw Pact ordnance delivered), not an outcome of one battle or one air raid. And in a middle-level simulation of a battle or air raid, a global MOE would be a measure of the outcome of the whole battle or raid, not, for example, the outcome of an individual tank or aircraft engagement.

There are, in a sense, two ways to make these optimal decisions. One way (1a) is to use mathematical techniques to calculate optimal game-theoretic decisions. This method usually picks the best option from a large set of options and uses optimization techniques (like linear programming) to avoid explicitly enumerating all of the options. This method is usually computationally tied to the conflict simulation, so that one computer run consists of both the optimization technique and the simulation of conflict.

The advantage of this method is that it can give the best measure of what each side "can do". (Shortly, I will discuss the differences between measuring what each side can do versus what each side will do). If this method is used to make a

decision, then an analyst can say that the resulting decision is the best (from a game theory point of view) that the side could possibly make out of many different options. The major disadvantage of this method (assuming that one wants to measure what each side can do) is that using mathematical optimization techniques forces the simulation to be relatively simple--or the optimization will not work. If the simulation has to be oversimplified in order to allow the optimization to work, then the results can be meaningless. (What is "too oversimplified" and what is an "acceptable aggregation of details" depends on the issues being analyzed and on many other factors.)

An alternative to using formal optimization techniques in an aggregated simulation is to examine the effect of a few different decisions for each side, using a less aggregated simulation, and then to find the best (or game-theoretic optimal) decision from this small set. For example, one could pick four different options for NATO and four different options for the Warsaw Pact, make sixteen runs of the simulation (one for each possible pair of decisions), and select the game-theoretic optimal decisions that result from the corresponding four by four matrix game. The advantage of this method is that it allows the use of a more detailed combat simulation. The obvious disadvantage of this method is that there are usually more than four options for each decision, and many more than one decision per side to be made. But this disadvantage can be reduced if the analyst can use judgment to lessen the number of options for a few basic decisions, and then investigate in greater depth only the "winners" from that first look. This method is not mathematically the same as investigating all options for all decisions, and one cannot be sure that some options for some decisions that were not investigated would not be better than the best ones that were investigated (the role of the analyst's judgment is to reduce these possibilities). But one can know for sure that the decisions resulting

from this method are the optimal ones over the limited set of decisions considered. And, depending on the analyst's judgment, a limited set can be good enough to address the important factors that affect the results of the simulation. (One member of this limited set of decisions could be the decision that was found to be optimal for a simpler model by using Method 1a.)

2. Non-Optimizing "Can Do" Methods

The second approach is also a "can do" approach. Method 2a is similar to Method 1a in that the model attempts to do a mathematical optimization. By "attempts a mathematical optimization" I mean that the method deals with gradients, Lagrange multipliers, and so forth, not with military considerations like attempting to achieve a certain force ratio or to reduce losses. The distinction that I want to point out here is that a mathematical approach fits into Method 1a *if it has been proven* to produce optimal solutions, and it fits into Method 2a *if it has not been proven* to do so (and especially if it has been shown to produce non-optimal results in some cases). Method 2a is frequently used because it is hard to make the simulation realistic enough and yet have the mathematical structure simple enough to be able to prove optimally. And sometimes Method 2a is used with the argument that "we can't prove it works but we've never found it to be wrong" (implying that it is probably right). But I don't think that Method 2a should be used because when it goes wrong (i.e., produces non-optimal decisions), it can be very wrong (and, in general, one doesn't know if it is right or very wrong). And I think that the cases where people haven't found a pseudo-optimal method to be wrong occur because they haven't looked very hard for errors.

Let me skip Method 2b for a minute and talk about Method 2c. Method 2c might also be used when the mathematical structure is too complex to guarantee optimality. Method 2c would use

algorithms based on intermediate goals and input parameters that are consistent with military judgment. For example, such an algorithm might move ground units to attempt to achieve an input desired force ratio, or it might reassign aircraft from a mission with a high loss rate to one with a low loss rate. This approach also can produce non-optimal decisions, but I think can *sometimes* be useful because it can produce military reasonable, if not optimal, decisions, and reasonable decisions are often good enough.

I have included Method 2b mostly for completeness. Method 2b is a cross between 2a and 2c in that it optimizes an intermediate MOE. Sub-optimization does not, in general, find a global optimum, so it is not the same as 1a. Depending on the details of the simulation and the particular intermediate MOE used, Method 2b might be more like an incorrect optimization (i.e., like Method 2a) or more like a militarily reasonable but non-optimal approach (Method 2c).

I will discuss a major disadvantage of Approach 2 in conjunction with Method 4a. below.

3. "Will Do" Methods

Approach 3 is what some people have called a "will do" approach and is the opposite of a "can do" approach. Method 3a is to attempt to find out from intelligence organizations what they think the Warsaw Pact will do. Intelligence sources might give some general tactics expected to be used by the Warsaw Pact, and this method would try to make these tactics as specific as possible so that they can be used to make decisions--and these decisions would be fed into the simulation as inputs. Similarly, Method 3a would attempt to find out from NATO commanders and planners as much detail as possible on NATO's expected tactics--that is, what we think

we will do. And then these tactics would be converted into fixed input decisions to be fed into the simulation.

Method 3b is similar in concept to Method 3a, but is similar in mathematical structure to Method 2c. It is like Method 2c in that algorithms, parameters, and goals are used to make decisions. The difference is that in Method 2c these algorithms are used to try to find a reasonable measure of the best each side can do. In Method 3b these algorithms are intended to reflect the decision processes that we think each side will make. For example, a decision process that tries to move forces on each side to achieve a certain force ratio is an example of Method 2c, while a decision process that requires NATO to make individual replacements and the Warsaw Pact to make unit replacements, based on some algorithm, is an example of Method 3b--providing that each side could use the other replacement method if it wanted to. (But, for example, if the Warsaw Pact couldn't make individual replacements because it doesn't have the logistical structure required to do so, then the restriction to unit replacements is a measure of capability, not a tactical decision.)

I separated the first three methods to make decisions into "can do" or "will do" methods. The difference between these methods is discussed in Section D below. The remaining methods discussed here are called "mixed", because they are in-between "can do" and "will do" in that they can be very close to one or the other, depending on how they are used.

4. Simple Methods

The fourth approach contains two different methods. I put them both under the same category because they have one characteristic in common--a type of simplicity.

Method (4a) makes decisions using a global MOE (like that discussed in conjunction with Approach 1). The difference here is that instead of seeking the optimum game-theoretic decision, the goal is only to make monotonic decisions. By a monotonic decision, I mean the following:

There are certain resources and effectiveness parameters that have the property that the more NATO has (for resources) or the better NATO is (for effectiveness parameters), the better off NATO should be in terms of global MOEs (and the same holds for the Warsaw Pact). For example, if NATO has more tanks or better tanks, it ought to do better than it would in a case where it has fewer tanks or poorer tanks. This sounds trivial, but if any of the methods under Approach 2 (or Method 3b) is used to make any decisions, things can go wrong. And if they can go wrong, they eventually will go wrong, maybe at the worst possible time. For example, a better NATO tank may mean that an algorithmic decision is made to keep reserves out of combat for a day or two longer (with the better tank, they are not needed as quickly to achieve some force ratio). But it may turn out that, over the course of the war, it is much better to move reserves up early no matter what tank NATO has. In this situation, the case with the poorer NATO tank would do better for NATO than the case with the better NATO tank, because of the decision process. This type of "bad result" can be avoided only if all decisions can be guaranteed to be (what I've called) monotonic. If all decisions are the game-theoretic optimal decision, then monotonicity occurs. But finding optimal decisions is hard with a complex model. The idea here is to allow the use of a complex model by using decisions that are simple enough that they can be shown to be monotonic (but not optimal). For example, a decision that $1/3$ of all forces are always in reserve and $2/3$ are always in combat would (in many combat simulations) be monotonic. Of course, proving monotonicity might, in some

cases, be as hard as proving optimality. But I think that monotonicity is usually much easier to prove (because the decision processes can be made more and more simple until monotonicity is obvious). And even if it can't be proven, and even if the decision turns out to be non-monotonic, I think that the consequences are not as severe as an error in optimality, because I think that a non-monotonic "approximation" is usually very close to monotonic, while a non-optimal "approximation" can be far from optimal.

One could also just use a simple, easily explained decision rule with the disadvantages of not being optimal, not being monotonic, and not being based on intelligence sources, but with the advantage that the rule is easy to explain and to understand. This is what I mean by the Method 4b. I think that the major disadvantage of Method 4b is that, in a large scale simulation, the *impact* of even a simple rule is hard to predict or understand.

5. Modeling the Human Mind

In small scale simulations one can attempt to model the human mind. For example, one could model how a pilot would make decisions during a dogfight by testing various pilots. I think one should be aware of this method, but I don't think that it has reasonable applications to large scale simulations.

6. Using Humans as Models

Finally, as I mentioned earlier, one can use humans in interactive simulations to make decisions. This approach has several advantages--primarily in the complexity of the decisions that can be given to humans to make. (It is very hard to automate complex decisions, especially in a two-sided game context.) The major disadvantages of this approach are: (1) it takes a lot of time to train people to play and to make

decisions in the simulation, (2) it takes a lot of people playing these games to be sure to obtain a result that depends on the issues being analyzed, and not on the game-playing skill of the particular players (if just a few runs are made with just a few players, then one frequently ends up evaluating only the players' ability to play computerized games), and (3) it takes a long time with trained players just to play through one game--let alone to repeat the game many times with different players--and so examining many alternative cases and doing parametric analyses is impossible. (Also, if the game is a Monte-Carlo simulation, then it would take very long just to play enough trials to ensure reasonable statistical validity for even one or two cases.) For these reasons, Method 6 is frequently difficult or impossible to use for force analysis. (Of course, Method 6 can be very useful for training.)

D. "CAN DO" METHODS VERSUS "WILL DO" METHODS

The distinction between "can do" and "will do" methods is as follows: A "can do" method uses intelligence sources and our own data to determine the effectiveness of the weapons and the capabilities of the forces on each side, but not to determine the decisions that will be made by the commanders on each side. Instead, a "can do" method attempts to make decisions based on the capabilities of weapons and forces, and uses the principles of game theory. A "will do" method uses intelligence sources and our own plans both to determine capabilities and to model decisions for the two sides, and does not particularly worry about whether better decisions could be made.

This distinction would not be important if the results of a simulation were about the same for a "can do" or a "will do" approach. But in large scale analyses, these approaches may frequently give very different results; and so the basic

decision as to whether a "can do" or a "will do" method is to be used is a very important and fundamental decision in doing any large scale analysis.

Let me briefly point out some of the arguments for and against each approach.

Some arguments in favor of the "can do" approach are: (1) we really don't know what either side would do if a war broke out next month, let alone if a war occurred in some future planning time so a "will do" approach is really just guessing; (2) if we do analyses based on the best each side can do, and then it turns out that the Warsaw Pact does something else, then NATO will be better off than planned. But if we plan on an estimated Warsaw Pact strategy, and it turns out that the Pact could do much better using a different strategy, then NATO could be much worse off than expected. That is, one should plan as if one is facing an intelligent enemy, not a stupid one; (3) if one is using a combat simulation, for example, to compare two different NATO force structures, then one should evaluate each force structure by how well it does if it is used as best it can be used against the enemy (who in turn, is trying to do his best to beat it)--not by how well it does if it is used poorly. Indeed, any force structure can be made to look good if good decisions are made for it, while bad decisions are made for the enemy; and any force structure can be made to look bad if bad decisions are made for it, while good decisions are made for the enemy.

More elaboration on these points is given in a report written by the Assistant Chief of Staff for Studies and Analysis of the United States Air Force in 1971. This report presents a good argument for the "can do" approach over the "will do" approach. (Unfortunately, it goes on to present a somewhat simple simulation and a pseudo-optimal method--type 2A, and this report has been criticized for these points, with the

argument for "can do" versus "will do" being frequently overlooked.) The part of this report that discusses "can do" versus "will do" is reproduced in the appendix below.

Some arguments in favor of the "will do" approach are as follows: First, it just sounds very natural for an analyst to say "you tell me what weapons each side will have and how they will use them, and I'll do an analysis to tell you how good those weapons are at that use." The problem with this argument is that, while it sounds reasonable, it frequently isn't reasonable because it evaluates the weapons and the expected use of the weapons together, and what is usually needed is not a combined evaluation, but a comprehensive evaluation of how good the weapons are without limitations on their use. For example, a large scale analysis might compare a force of 100 F-15s and 400 F-16s to a force of 200 aircraft of each type. Now there may be a current plan on how to use these aircraft and how the Warsaw Pact will use their aircraft in opposition. But if an analysis is to be made of these force structures and not of the plans (plans can be changed after the aircraft are bought, but what you buy, hopefully, is what you get), then the interesting question is which force structure is better to have if each force is used as best as it can be used against an intelligent enemy.

A second argument (which I also don't accept) is as follows: "We can't use game-theoretic methods because we don't really know what goals the Warsaw Pact have. And since we don't know what they want to do, we should use intelligence sources to estimate their plans, and then optimize our defenses against their plans." The reason I don't accept this argument is that its proponents never seem to want to examine several different but reasonable goals for the Warsaw Pact to see if the conflict in question really forms a non-zero sum game, or to see if there are several different goals and both sides want

to get the best they can for all the goals. I think that the non-zero sum game case is unlikely, but if it occurs there are game-theoretic methods to handle it. The multiple goals case is likely and should be examined more carefully, not arbitrarily dismissed. So I reject this argument, not because we know exactly what the Pact's goals are, but because we can estimate several reasonable goals for both NATO and the Pact, and then proceed on a "can do" approach from there.

A third, but hypothetical argument in favor of the "will do" approach is as follows: Proponents could argue that they are absolutely sure that the Pact will use certain non-optimal tactics, and they want to exploit this knowledge rather than settle for results based on a "can do" analysis. (In general, if one knows for sure what an opponent will do, then one can achieve a more favorable outcome than the game-theoretic solution.) But this argument is hypothetical; I have never heard anyone use it in practice, perhaps because no responsible person would say that NATO can ever know in advance that the Pact will make a particular erroneous (non-optimal) decision (which we should therefore exploit). Instead of this "I know what errors the Pact will make" argument, the second argument just discussed, "We don't know what goals the Pact have", is usually used to support a particular position.

The following argument for the "will do" approach over the "can do" approach is, I think, a good argument. This argument is that it is much easier to use the "will do" approach (given that a certain standard of technical detail must be met), and that the sponsors of research either want or will be satisfied with a "will do" approach. The reasons that the sponsors will be satisfied with a "will do" approach are first, they also are usually trained in technology rather than decision analysis, and second, the results of a "can do" approach are likely to be more controversial and be more directly useful to a high level

decision maker (which is a problem if the sponsor is a representative of a middle or low level decision maker, and does not want to take on high level decisions). The reason that the "will do" approach is usually easier (for a constant level of detail) is that the "can do" approach usually requires expanding the scope of the study, and more interactions need to be simulated, and so either more data and a bigger model are required, or the level of detail must be reduced (which would also reduce the credibility of the study from a technical viewpoint).

These arguments in favor of "will do" approaches are rarely explicitly stated in studies. More commonly, no justification at all is given for using a "will do" approach over a "can do" approach. It's just done, with no justification, and frequently with no criticism from sponsors or reviewers.

Let me give two examples to illustrate the difficulty of choosing between a "can do" and "will do" approach. First, suppose NATO is deciding between buying one of two types of SAMs to protect an area, and suppose one of the SAMs costs twice as much but kills three times the number of Warsaw Pact bombers than the other SAM. It would seem that the more expensive SAM is better--but what if the Warsaw Pact could mitigate the advantage of the expensive SAM by spending some more money itself--perhaps by buying more bombers? Here the problem is that NATO has to take money from something else to pay for the more expensive SAM, and the Warsaw Pact could also take money from something else to defeat the more expensive SAM, so the simple answer of "three times the kills for twice the money" alone doesn't say which SAM NATO should buy. It may depend, in this example, on comparing the cost to NATO of buying the more expensive SAM versus the cost to the Pact of defeating this SAM. And even if we assume that the Pact cannot directly mitigate the effect of the more expensive SAM by

buying more bombers (or other increased spending), the Pact could decide to use their bombers to attack somewhere else, if the more expensive SAMs cannot defend all possible attack areas. My point here is not that the scope of any study can be expanded so much that the study can't be done. My point is that an analysis that recommends spending more and getting more (or spending less and getting less) must worry both about its scope and about the decisions it has explicitly and implicitly made for each side, and addressing these decisions can be very hard if a "can do" approach is used.

My second example is based on the consideration of multiple missions for aircraft. Suppose one is doing an air defense study and is comparing SAMs with aircraft. Then it is natural to look at a Warsaw Pact raid against NATO and see which system (the SAM or the aircraft) does better. But the Warsaw Pact might not fly a raid into NATO territory-- they might use their aircraft to fly defensive missions over their own territory. In this case, our SAMs would be relatively useless, but our aircraft could fly offensive missions against the Pact. The following is a quote taken from an FRG publication:

Thus it is the aim of equipment planning to identify an optimum mix of equipment that will accomplish the entire spectrum of tasks of the *Bundeswehr*. To this end, the relative importance of the various groups of tasks must be assessed against the overall mission of the *Bundeswehr*. ...

This poses a difficult problem for the planner since almost all tasks require several mutually complementary weapon systems and, conversely, a number of weapon systems are employed for several tasks.

Example: A specific air defence missile system is employed only for the air defence task, whereas

a combat aircraft may be employed for reconnaissance, close air support, and air defence.¹

Yet most air defense studies pay little attention to the offensive capabilities of NATO aircraft--after all, they would say, we're studying air defense not attack. And expanding the scope of a study to consider both the defensive and the offensive capabilities of both side's aircraft, expands very much the data requirements, the size of the simulation, and the time, work, and risk involved. It is not an easy choice.

E. CONCLUSION

In summary, Figure 3 gives a list of methods that can be used for modeling decisions, and I have just discussed some advantages and disadvantages of each of these methods, especially as they relate to large scale conflict simulations. My only firm recommendation is that one should be aware that there are decisions that must be made inside conflict simulations, that how these decisions are made can potentially be very important, and that there are different methods available for making these decisions. Which options are best under which circumstances depend on the circumstances and the people involved. However, let me conclude with my general preferences for large scale analyses.

If an analyst is investigating an issue and has judgmentally determined that the answers to that issue are likely to be the same over a wide set of possible decisions, then these decisions should, in general, be made using monotonic decision rules (Method 4b).² The advantage of monotonic rules

¹White Paper 1975/76, The Security of the Federal Republic of Germany and the Development of the Federal Armed Forces. Published by the Federal Minister of Defence on Behalf of the Federal Government.

²Fixed decisions are usually monotonic, so Method 3a (fixed decisions based on expected tactics) is usually a special case of a monotonic rule. Thus, Method 3a can also be used here.

over optimal methods (Methods 1a or 1b) is that they allow a complex model to be used without requiring iteration through all important options. And their advantage over algorithmic methods (Methods 2 or 3b) is that the details of their decision logic won't trap the analysis into inadvertently making better decisions for one case than for another. The point of this recommendation is that the analyst will have judgmentally tested the impact of other decisions, and will have judgmentally determined that, while other decisions might affect intermediate outcomes or other issues, other decisions would probably not affect the results for the issue at hand.

However, if the analyst's judgmental testing of other decisions indicates that different decisions will likely produce different results for the issues at hand, then I recommend using Method 1b. (Method 1a can be used provided that an appropriate simulation is simple enough to accommodate an optimization process.) This recommendation is easy to make but it is frequently hard to follow because, for the reasons just given, using Method 1 is generally harder than using Methods 2, 3, or 4; and so, for a constant level of effort, some detail must be sacrificed to use Method 1. But when the results depend on the decisions, then Method 1 is the best way, and sometimes the only reasonable way, to model decision making in conflict simulations.

APPENDIX

This Appendix is an extract from the Summary
and Introduction to

METHODOLOGY FOR USE IN MEASURING THE
EFFECTIVENESS OF GENERAL PURPOSE FORCES
(An Algorithm for Approximating the Game Theoretic
Value of N-Staged Games)
SABER GRAND (ALPHA)

UNITED STATES AIR FORCE
ASSISTANT CHIEF OF STAFF, STUDIES AND ANALYSIS
MARCH 1971

This study produced a model called TAC CONTENDER.

SUMMARY

The methodology developed in this paper is a portion of an overall research program conducted by the Office of the Assistant Chief of Staff for Studies and Analysis, HQ USAF, to develop improved methods to measure and assess the relative value of alternative general purpose force structures.

The employment of general purpose forces, as distinct from the employment of strategic forces, is characterized by the opportunity for successive decisions by the theater commander as the overall campaign progresses. These decisions are based at each stage in the campaign upon the residual forces available to both sides, the expectation of reinforcements at later times, the status of the campaign in terms of objectives sought, and the time period available to achieve those objectives. At each decision point, which may be hourly or daily or even longer, the commander has to allocate his available aircraft to various tasks. He must judge the utility of any allocation in terms of its contribution in the immediate time period to the objectives sought and its effect on conserving his own forces in order to use them in subsequent time periods toward achieving his objectives. Similarly, he must judge the utility of his allocation in terms of preventing enemy forces from achieving their objectives in the immediate time period versus destroying enemy forces so that they cannot be used in future time periods. Finally, the commander must judge any allocation in terms of enhancing his own objectives versus diminishing the enemy's ability to achieve his objectives.

For a tactical air campaign carried out in conjunction with a land campaign, the tasks to which aircraft can be allocated can be usefully characterized as follows:

- a. Offensive ground support of friendly (Blue) forces by direct air strike on enemy (Red) ground forces.
- b. Offensive air strikes against enemy (Red) airfields to destroy enemy aircraft.
- c. Defense of the battlefield area to reduce the effectiveness of enemy (Red) attacks on Blue ground forces.
- d. Defense of friendly (Blue) airfields against attack by Red aircraft.

The overall outcome of an air campaign constructed of these four tasks can be stated in terms of the cumulative amount of ordnance that each air force is able to deliver on opposing ground forces. It is easy to show that, with various arbitrary (but apparently reasonable) employment strategies among the four tasks, the outcome can vary over a wide range of values. In fact, the amount of Red ordnance delivered on Blue troops can range anywhere from zero to some "Red-maximum"; and, similarly, the amount of Blue ordnance delivered on Red troops can vary from zero to some "Blue-maximum." There are two difficulties associated with this wide range of outcomes. First, the outcome is altogether ambiguous except for the two bounds, Red-maximum and Blue-maximum. Second, the values of these bounds do not embody all aspects of the combat capabilities of the forces. In particular, the air-to-air combat capabilities of the opposing forces have no effect upon the values of Red-maximum and Blue-maximum.

In investigating the nature of the employment strategies which lead to this wide range of outcomes, it is clear that many of the outcomes require "cooperation" between the opposing sides. For example, if Red uses all of its forces to attack Blue ground troops, then Red can deliver its Red-maximum amount of ordnance on Blue troops only if Blue obliges by employing

all of its air forces in attack on Red ground troops (in which case Blue delivers its Blue-maximum ordnance), or in defense of Blue airfields (in which case Blue delivers zero ordnance). Put another way, Red cannot "enforce" the Red-maximum outcome unless Blue cooperates. Such cooperation is generally not in Blue's best interest. This leads us to the concept of a "mutually enforceable outcome" in which each side chooses an employment strategy such that each side has an "enforceable" minimum outcome independent of the other side's employment strategy. That such solutions exist is a well-known result of a branch of mathematics called game theory which has been developed over the last 20 years.

This paper, therefore, formulates the four-task air campaign outlined above as a "two-person, zero-sum, N-staged game." It is called "two-person" because there are two players, Red and Blue. It is "zero-sum" because we have chosen the payoff to Blue as the difference (Blue ordnance on Red troops minus Red ordnance on Blue troops) and the payoff to Red as being the negative of this difference. That is, Blue's gain is Red's loss, and vice versa. It is called an "N-staged game" because we allow each player (theater commander) to make N successive choices as to how to allocate his residual forces among the four tasks. The period of time between successive allocations is typically a day, and the number N was typically between 10 and 90. As is well known from game theory, there are three possible types of allocations on any one day: one, allocation of the entire force to a single task; two, allocation of a specific portion to each task (a split strategy); and three, allocation among the tasks on a probabilistic basis where the frequency distribution is known but the actual choice on any given play of the game is chosen at random.

In making each day's choice, the theater commander must, as mentioned before, balance the contribution to the overall payoff function achieved today (through allocations to ground

support and battlefield defense) and the relative force strengths that the forces will have to conduct the remainder of the campaign after today. The key to making this decision depends clearly upon a knowledge of the potential value of both friendly and enemy aircraft that will be available to fight the remainder of the campaign, but that in turn depends upon the employment strategy that will be used by both sides for the remainder of the campaign. The classical solution to the problem of evaluating the potential value of residual aircraft is to solve the problem backwards, starting with the last day. Then, knowing the employment strategies for the last day, solve for the allocations on the second-to-last day. Then, knowing the employment strategies for the last two days, solve for the third-to-last day, and so on.

There are two practical problems associated with this classical solution that led to the present research. The first problem is that analytic solutions can be found only for very simplistic interactions between the forces, whereas more realistic interactions did not appear to have analytic solutions. Second, numerical solutions could be obtained only for given final values of Red and Blue forces, but what was needed was a numerical solution which started with given initial values of the forces....

INTRODUCTION

The methods of numerical analysis presented in this paper were developed in the course of an investigation of improved methods for evaluating the effectiveness of general purpose forces, and of tactical air forces in particular. These evaluations can be considered as answers to three types of questions:

- What is the value of adding a new weapon or more weapons to a force? Or
- Which of several alternative forces is the best force? Or
- What is the minimum force required to meet an opposing threat?

The method to be described here focuses on the combat interaction of opposing forces over a protracted, but finite time period--that is, a campaign. This method is distinct from previous analyses that seek to evaluate general purpose forces in terms of targets killed, sorties flown, firepower possessed, or other such measures of merit that ignore the sequential, interacting nature of a tactical campaign.

The method also focuses on the problem of force employment, which is often neglected in comparison to matters like the level and composition of opposing forces and the geographical and political scope of the scenario. By force employment we mean the commander's allocation of his forces among various geographical sub-regions and among various military missions or tasks. The entire set of such allocation decisions, made sequentially over time during the conduct of the campaign, may be called the strategy.

As most operational commanders of general purpose forces know, the ways such forces are employed usually have a greater impact on the outcome of a conventional tactical campaign than do the characteristics of the two opposing forces. While it is true that larger and better equipped forces win over inferior forces more often than they lose, history is replete with examples where superior forces were defeated as a result of the unwise employment of the superior forces, and the wise employment of the inferior forces. Using combat simulations, it is also very simple to demonstrate the defeat of a superior force through a combination of "smart" and "dumb" force employment by the victor and the loser, respectively. Even when a "black and white" win or lose is not the issue, it is clear, both from historical and analytical perspectives, that a wide range of outcomes in conventional tactical campaigns can be expected, depending upon how the forces are used. Unfortunately, as was stated previously, the force employment decisions are usually given less than careful consideration.

Most simulations or force comparisons begin by stating the enemy strategy (e.g., 10% of the aircraft will be allocated to air defense, 40% to airfield attack, 5% to interdiction, and 45% to ground support). At the worst, the allocations are "thought up in order to have something to make the program run." At the best, the allocations are a considered "best estimate" of the enemy's behavior. Typically, a similar procedure is used to develop the strategy to be followed by the friendly forces. These strategies are then "fixed" or held constant throughout the evaluation of several general purpose forces.¹

¹In most cases, as in the example cited, the allocations are held fixed throughout the campaign. However, this is in general incorrect. In a conventional campaign each day's combat will entail casualties and movement of units. As we will see in the body of this paper, both change the situation in ways that alter the preferred allocations for both sides.

Even when the strategies chosen are given "careful thought" and analysts choose "good strategies," a procedure that uses the same strategy in comparing alternative forces makes invalid comparisons. To see why this is so, consider a Blue force that is to be compared against a Red threat. "Good" strategies are selected for both sides, and a campaign is fought. Let us call the outcome, in terms of some measure of merit, V_1 . Suppose now that a new weapon system is added to the Blue force and the campaign is fought again, with the same strategies as before, and a more favorable outcome, V_2 , is achieved. The statement usually made at this point is, "The value of the weapon system is the incremental difference in the outcomes, $V_2 - V_1$." However, the quantity, $V_2 - V_1$, could well be an overestimation of the effects of the weapon system. The reason is simple. In a conventional tactical campaign, there generally exist alternative strategies for Red that, when employed, could at least partially counter the new weapon system. Conversely, $V_2 - V_1$ could well be an underestimation since there probably exist alternative Blue strategies that would make better use of the new weapon system. Thus, to say that in a conventional tactical campaign (where both forces have a wide range of strategy options) that the value of a weapon system is $V_2 - V_1$, is to make an incorrect and misleading statement. Of course, the problem is the same when alternative Blue forces are compared. The best Red strategy against Blue alternative force A is probably different from the best strategy against Blue alternative force B. Similarly, Blue's best strategies will be different.

This raises a philosophical question. Should one measure the effectiveness of a general purpose force (in terms of the measure of merit) against an opposing force following some assumed strategy or should one measure the effectiveness against opposing forces used to their very best advantage? The equally important converse question is: should our proposed

alternative forces be constrained to follow some assumed strategy¹ or should they be allowed to follow an "optimal" allocation policy which maximizes the force's effectiveness?

If the above two questions are answered in favor of the "optimal" alternative in both cases, then it is necessary to develop a technique to determine these "optimal" allocations for both the enemy force and the proposed alternative friendly forces. A well defined comparison of alternative forces can then be made.

Game theory has long held a theoretical answer to determining the optimal allocations by the two players whose interests are diametrical. When one defines the MINIMAX criteria² and uses game theory, one can derive the "optimal" strategies and the corresponding outcome, called the game value. Of real importance are the following characteristics of the game theory solution:

- If a player follows his optimal strategy, then the game value is the worst outcome he can expect (any strategy by the opponent other than the opponent's optimal strategy will produce a more favorable outcome for the player).
- Given that his opponent follows his optimal strategy, a player can do no better than the game value (any other strategy choice by the player will produce a less favorable outcome for him).

A further characteristic of a game theory solution is that any real change in force size or quality by Red or Blue will produce a well defined change in the outcome. As noted before,

¹The question of valid constraints on allowable strategies will be addressed in the following section.

²Put briefly, Blue first rates each of his strategy choices by the unfavorable outcome that could result from that choice. Then Blue chooses the strategy that has the most favorable of these "most unfavorable" ratings. Red does the same from his own point of view. The word MINIMAX is derived from the fact that a player is attempting to minimize his maximum losses.

the use of an arbitrary pair of strategies for Red and Blue may serve to overemphasize or underemphasize the value of an incremental change in one force relative to the value of the original force, or alternative changes.

Another way of interpreting a game value is to say that it is a mutually enforceable outcome. Any other outcome is not mutually enforceable in that one of the two players could always find a different strategy that would produce a more favorable result for him. Such outcomes are really not solutions, for, given the opportunity, players would change strategies and produce a new outcome, given that they are smart enough to figure out how to do so.

It should be noted at this point that we are dealing with a particular form of game theory--that for an "N-staged game." The name reflects the fact that there are a series of N decision points, or stages, in the conduct of a tactical campaign. The proper allocations of forces at each stage depends upon the cumulative outcome of the preceding stages and the length of time left in the campaign.

This paper presents a generalized algorithm which allows one to arrive at a numerical approximation to the mutually enforceable solution and associated strategies. Although the vehicle chosen to illustrate the technique is the problem of the proper employment of tactical fighter forces, we currently believe that the general technique can also find application in the very broad range of problems which are characterized by two opponents having conflicting objectives and making a series of strategy choices.

"CAN" VERSUS "WILL"

Inherent in the use of game theory and a "mutually enforceable" solution is the concept of measuring what effectiveness a force can achieve, as opposed to the effectiveness

a force will achieve. It is our strong belief that this is the correct measurement to make: what the effectiveness can be and not what someone thinks it will be. The reason is quite simple. In order to measure what will be the effectiveness of two tactical forces, one must postulate what the strategies for the two forces will be. Because the results are so sensitive to the strategies, the calculated effectiveness is driven by the postulated strategies that someone has decided will be followed by the two forces. Different "someones" inevitably decide on different strategies, and the scatter of outcomes falls to all extremes on any goodness scale.

In the past, analysts touting the superiority of our forces have postulated a Red strategy and a Blue strategy that produce very favorable results for Blue and thus show the Blue force to be superior. Other analysts, bent on showing the weakness of our forces, have countered by postulating different strategies that "will be followed" and calculating outcomes that are disastrous to Blue. Hence, the question of measuring the effectiveness often dissolved to a question of determining what the strategies of the two forces would be. Here, it was almost impossible to get an agreement.

Even for the "honest analyst" who advocates only the "truth," determining what the strategies would be is almost impossible. Assuming that intelligence estimates of the enemy's strategies are correct, and that our war plan estimates of our strategies are correct, an analyst can narrow the region of expected outcomes. By ruling out all strategies considered unlikely by the intelligence and war plan estimates, the analyst can eliminate all possible outcomes considered unlikely. He then has a narrower region from which to select what the outcome will be. Unfortunately, however, because of the sensitivity of the outcome to the strategies, even a slight

permutation of the strategies within the narrowed region of likely outcomes can produce widely differing outcomes. In short, the region is not as narrow as it might have seemed. Even when one outcome is chosen, the credibility will always be challenged by those postulating different strategies for the two forces. If, however, one measures the effectiveness based on what the forces can do, an analyst can only be questioned on "whether or not the forces could follow those strategies." That question is much easier to answer.

When one determines what strategies could be employed, one must, as stated above, consider the constraints upon the forces. Certain hardware characteristics (e.g., no bomb racks on some aircraft), basing policies (e.g., all aircraft based out of range of enemy airfields), or employment restrictions (e.g., air defense aircraft can not leave the homeland) constrain the strategies two forces might follow. These narrow the range of allowable strategies. Even then, however, one should attempt to measure the "mutually enforceable" solution over the range of allowable strategies and determine what effectiveness could be achieved.

If, for example, the campaign were to be fought tomorrow, then the constraints upon the forces might be fairly well defined. Intelligence might also provide additional information that could further eliminate some possible enemy strategies. The effectiveness of the forces should then be measured by gaming over the "allowable" strategy space for both forces. If the campaign were to be fought 10 years from now, then the outcome should be measured by gaming over the strategies available at that time. Careful thought must be given to the constraints one would want to impose at such a distant time.

Many of the constraints upon a force today could be removed by some time in the future. Interceptor aircraft could be modified to carry bombs, aircraft could be rebased or extra

fuel stores could be added, and policy decisions could be changed to permit aircraft currently designated as air defense aircraft to leave the homeland. Following the same reasoning which was used to measure what the effectiveness could be (as opposed to what it would be), we feel that the only constraints applicable for a future campaign are those that could not be removed by that time. Attempting to determine what constraints would still exist has all the pitfalls of determining what strategies will be followed.

In summary, since we are attempting to measure the effectiveness of a general purpose force, it is our belief that we ought to attempt to measure the maximum capability of a force. This implies that we ought to measure what the force can do in place of what someone thinks it will do. We will, therefore, remove all "restrictions" to strategy options that could be changed by the time period of interest. (If, with personnel or policy restrictions removed, the outcome is more favorable to the side removing the restrictions, then we have determined the cost of those restrictions, and policy makers might do well to consider the implications.) To measure what two opposing forces can do, a method must be found to determine the optimal interacting strategies for the two opposing forces.

Game theory is that method, for it determines the outcome of a campaign in which both forces are acting and reacting in an optimal manner. Unfortunately, however, game theory has not found wide application in general purpose force evaluation because of the difficulty of solving the resulting large and complex games....