Proceedings of the 1988 Winter Simulation Conference
M. Abrams, P. Haigh, and J. Comfort (eds.)

Modeling qualitiative issues in military simulations with the RAND-ABEL™ language

P. D. Allen
The RAND Corporation
1700 Main Street
Santa Monica, CA 90406

and

B. A. Wilson
The RAND Corporation
1700 Main Street
Santa Monica, CA 90406

ABSTRACT

This paper describes a methodology for developing combat simulations that may be readily tailored to specific study issues, and structured so as to treat both qualitative and quantitative variables using the natural language of military planners. The methodology was developed for the RAND Strategy Assessment System (RSAS) and exploits a new programming language called RAND-ABEL™. RAND-ABEL™ can be understood and used by analysts who are not expert programmers, and allows key relationships to be expressed in a form that is more easily reviewable by non-analyst military officers.

The RAND-ABEL™ language has been used to build several combat simulations, which we refer to generically as Referee models. In Referee models, both the assessment and decision processes are rule-based, thereby making the models both transparent and flexible. Through the use of a selective RAND-ABEL™ interpreter, the models may be modified even while it is running. This new approach to modeling has been successfully applied to several projects.

INTRODUCTION AND OVERVIEW

This new approach to modeling has appeared in different forms and been given several names. As part of the RSAS, it has been called the Secondary Land theater model (S-Land), and later renamed as the CAMPAIGN-ALT (for alternative theater) model. It has also been used within the RAND Analytic Modeling Platform (RAMP), which is an unclassified descendent of the RSAS. For purposes of this paper, we will refer to the class of these models as Referee models, since they act as a gaming referee in combat assessment. The Referee approach to modeling and analysis was developed independently of any of the qualitative process theory work that has appeared in the literature over the last few years.

This paper is arranged in the following sequence: First, the needs and problems faced by the military simulation community are presented. Next, the features of the new RAND-ABEL™ language are described, along with a few examples of the code. It was decided that a completely new approach to combat modeling was necessary to satisfy both the needs of the military simulation community and to fully exploit the features of the RAND-ABEL™ language. Therefore, the third section describes the implementation techniques that have been successfully used in building military simulations in RAND-ABEL™. The fourth section presents new analysis techniques that have been developed to fully utilize the power of both the language and the model design. Some of these new techniques include interactive sandtabling (as opposed to batch runs), new measures of effectiveness for qualitative variables, and aspects of sensitivity analysis, verification and validation.

The fifth section briefly describes several applications that have successfully employed the Referee models. The last section relates some of the observations made by users of the models.

SIMULATION NEED AND APPROACH

The simulation and modeling community has been troubled in its efforts to find the elusive balance between model realism and usability. The usual answer in the area of air and ground combat modeling has been to go to higher and higher resolution at the expense of transparency, flexibility, and overall utility. In addition, the demands placed by these models on extensive quantitative data bases have made the care-and-feeding expenses of most large simulations prohibitive. Even when most of the data is obtainable, many of the qualitative questions, such as the effects of surprise or the value of intelligence, are still not measured by these models. As the size of the models increase, their ability to respond to short-deadline projects becomes severely impaired.

Most combat simulations require a large quantity of hard data to perform properly. Many people beyond the military modeling community would be shocked by how little hard data actually exist. Furthermore, much of the existing historical combat data depend upon qualitative factors, such as combined arms effects, surprise, and operational art. Most analysts in this community would agree that warfare is still an art, and that the simulation of warfare is an even more difficult art. Consultants from academia are often startled by the fact that most of the combat processes simulated are not only subject to debate, but are frequently modified in a given model to address different study requirements.

A research institution (such as RAND) will inevitably face a wide range of study requirements. An issue one month may be whether or not a given weapon system can have a desired effect in a theater of operations. Next month, one may want to study in more detail the supportability of such a system, a task that requires a model with a significantly different emphasis. Unfortunately, most places tend to suffer from the "if I have a hammer, everything looks like a nail" syndrome. Models are often applied because they exist, regardless of how applicable they are to the problem. The two most important questions to remember are: (1) what specific study questions are being asked? and (2) how much time is available? The answers to these two questions determine the tools that should be applied to the problem.

Rather than trying to design the one model to solve everyone's problems, we chose to create a modeling methodology that could be rapidly prototyped and quickly adapted for a wide variety of study issues.

In addition, the modeling system must be flexible, so that it can be changed by the analyst even while the model is running. The RAND-ABEL™ interpreter allows interactive operations in which the analyst changes rules and algorithms, not just data. This is in contrast to old fashioned batch operations with their voluminous output and the batch requirement to restart at the beginning of the run. The interactive process (called "sandtabling" in the RSAS) has greatly reduced the time necessary to develop, test and use the model. For example, one can gain insight by running an "incomplete" simulation, and then selectively "flesh out" the more promising area of interest as time and resources permit.

This selective resolution handles the question of "granularity" referred to in the literature. One can start with a low level of resolution (for example, very qualitative factors) and then selectively add resolution (such as more quantitative factors) as the model is refined. Since one is never really finished when improving a model, the process of refinement typically ends when there is no more time available to address the study issues. One can continue to build off the most refined model for the next project, or one may choose to start from a simpler base model for ease of use.

Restraint is emphasized while selecting which factors will be included in the model to prevent the model from becoming too complex. Model transparency and usability is emphasized over high resolution "realism." One should be selective in what is included, and be guided by the study issues. One should avoid the desire to include everything in the model because it is part of "reality." One does not get "truth" from models, only insights. The quote "All models are false; some models are useful" applies at all times.

Not every issue being analyzed need be modeled explicitly. Offline analysis should be used extensively with the model to account for important details not worth the effort to model thoroughly. For example, actually simulating the flight path of every aircraft is not cost effective if one is simply trying to determine whether or not a system will have an effect at all. This combination of simulation and offline analysis is particularly useful in option screening analysis.

The Referee model is designed to start with concepts that are familiar to the sponsors and the analysts. For example, one starts with one or more concepts of operations in a theater, and tailors the details of the model around those concepts. Terms are defined such as air control, surprise, and phases of battle.

Finally, due to the lack of "hard" quantitative data in most aggregate military simulations, the Referee models have a distinct advantage. Quantitative data is used whenever available, but in its absence, qualitative human judgement is used instead. The RAND-ABEL™ language is emminently suited for this task, as shown in the next section.

THE RAND-ABEL™ LANGUAGE

The RAND-ABEL™ language was developed at RAND as part of the RSAS development program. RAND identified the need for a language that could handle flexible and transparent decision processes quickly without requiring a special machine (like a LISP machine). The RAND-ABEL™ language "compiles" into the "C" language and runs on the UNIX operating system. This makes the RAND-ABEL™ processing time relatively quick--no more than three times slower than normal "C" code. In addition, there is a selective "interpreter" process that interprets only those files modified by the analyst. Although interpreted files take longer to run, the ability to change the model while it is running more than offsets the increased runtime.

Probably the strongest feature of the RAND-ABEL™ language is the table structure. Originally designed to describe an automated player's decision process in a simple, tabular form, it has since been applied to the model's assessment processes as well. The following sample RAND-ABEL™ tables defining qualitative factors such as surprise and air control and the quantitative factor of a loss rate were taken from the airborne operation assessment process.

Table 1 shows actual RAND-ABEL™ code. The input variables are the two variables to the left of the slashes. The output variable is to the right of the slashes. There may be a number of input and output variables, but they are usually limited to what can be seen on a screen or on a normal sheet of paper.

```
TABLE 1: Local Degree of Surprise.

Decision Table

 strategic-    air-control- / local-degree-
 air-surprise  over-target  / of-surprise
 ============  ============ / =============
 High          Attacker       High
 Medium        Attacker       High
 Low           Attacker       Medium
 High          Defender       Medium
 Medium        Defender       Low
 Low           Defender       Low
 High          --             High
 Medium        --             Medium
 Low           --             Low
 --            --             Low

[End Table].
```

The table is read as follows: If the value of the first variable (strategic air surprise) has the value "High" and the variable "air control over target" has the value "Attacker," then set the output variable "local degree of surprise" to "High." Otherwise, then examine the next row until one of the rows is true. Note that there is an implied "and" statement between input columns, and an implied "else if" statement between rows.

The symbol "--" means that we do not care about the value of this variable when examining this row. Therefore, in rows seven through ten, the air control over the target does not matter when determining the local degree of surprise. This is because the other two values for the variable "air control over target" are "Contested" and "Neither." Therefore, only the variable "strategic air surprise" determines the value of the output variable. The last row of the table is the default row, which sets the value of the output variables in case the earlier part of the table did not form a complete partition. This last row can also be used to indicate an error condition if so desired.

In Table 1, it is assumed that the values of the input variables were set before reaching this decision table. The strategic air surprise, for example, can be set either by fiat (one may want to assume strategic air surprise for study purposes) or by assessing an earlier table that defines strategic air surprise as a function of aircraft alert and dispersal rates. Similarly, local air control is a function of whether or not each side's air control range extends over this specific target. The air control range for each side is a function of the most forward operating base, the sustained operating range of fighter aircraft, the quantity of aircraft flying, and whether or not friendly AWACs are present.

Note that one can add additional columns as input or output variables even while the model is running by using the RAND-ABEL™ interpreter. This allows the analyst to selectively increase the resolution of this assessment process interactively.

Tables may be sequenced so that the outputs of one table become the inputs to subsequent tables. To continue our example, the value of "local degree of surprise" is used as an input variable in Table 2.

The first row of Table 2 is read as follows: If the number of defensive counterair (DCA) sorties flying in the theater is less than or equal to 50, and the ratio of escort to DCA sorties is at least one-to-four, and the local degree of surprise is High, then the lift loss rate is only two percent of the lift aircraft in this operation. The twelfth row shows that if there are more than 50 DCA sorties and there are not enough escorts, and the local degree of surprise is Low, then the loss rate is 25 percent of the lift aircraft. If one does not agree with the numbers in the table, one can insert better numbers, thereby refining the model by including better data over time. As more quantitative data becomes available, the definition and impact of qualitative parameters can be selectively improved.

TABLE 2: Determine Airborne Lift Loss Rate

Decision Table

DCA-sorties	escort-sorties	local-degree-of surprise	/ lift-loss-rate
<=50	>=(.25*DCA)	High	0.02
<=50	< (.25*DCA)	High	0.03
> 50	>=(.25*DCA)	High	0.07
> 50	< (.25*DCA)	High	0.10
<=50	>=(.25*DCA)	Medium	0.05
<=50	< (.25*DCA)	Medium	0.09
> 50	>=(.25*DCA)	Medium	0.12
> 50	< (.25*DCA)	Medium	0.15
<=50	>=(.25*DCA)	Low	0.10
<=50	< (.25*DCA)	Low	0.15
> 50	>=(.25*DCA)	Low	0.20
> 50	< (.25*DCA)	Low	0.25
--	>=(.25*DCA)	--	0.20
--	< (.25*DCA)	--	0.25

[End Table].

Again, these two tables are actual RAND-ABEL™ code with only minor modifications to allow them to fit into the format of this article. For example, the actual code has another output column for the fraction lost on ingress. That means all of the important model assessment variables, logic, and assumptions are kept within the context of the code. Some key intensity variables, such as combat intensity, sortie generation rates, and movement rates are available in the Data Editor for easy access, but the analyst no longer needs a single enormously long list of variables for all model modification. *All* of the variables in the model can be changed by the analyst. *There are no hardwired numbers in the model.* In addition, keeping most of the assessment variables in the code allows the analyst to see the effects of changing the value of a parameter in the context of the code.

Note also that the tables may include enumerated variables (which have values such as High and Low), numerical values, equations (such as in Table 2), and even calls to other functions. This allows the analyst to call upon a more detailed assessment in selected situations of particular interest.

Due to these features of RAND-ABEL™ tables, and to the model design that placed all significant variables and logic in a table format, the model acts as its own audit trail and corporate memory. For example, a new analyst may be asked "Did the model run last year account for air defense assets present at the drop zone?" All that is needed is to check the table (in this case Table 2) for an input column accounting for air defense assets. Although the answer in this case is no, one could either add another column to Table 2 to account for air defense assets, or add a third table to increase the lift loss rate as a function of air defense assets.

IMPLEMENTATION TECHNIQUES

Due to the advantages of RAND-ABEL™, it was perceived that new modeling techniques were required to fully exploit the features of this new language. Furthermore, one could design the Referee model around existing military concepts and planning procedures.

The Referee modeling approach strongly encourages focusing first on the broad scope of issues, accounting for the first-order effects of the functional area processes. These functional areas in combat simulations include: ground, naval and air combat; ground, coastal and air control of, next to, or over targets; special operations; simple logistics functions; and basic automated local commander functions. Heavy emphasis is placed on the representation of operations that rely on factors other than just mass, such as amphibious landings, airborne operations, and unconventional warfare operations (e.g., special forces). Rather than simply determining the success or failure of an operation on a simple probability (say, a 50-50 chance of success), *the focus is on the factors that help contribute to the success or failure of the operation,* such as surprise and air control. Additional resolution may be added later in these areas as a function of the study requirements. Starting with the outputs in each of the functional areas is a good way to avoid the "kitchen sink" approach to modeling.

The kitchen sink approach assumes one must include every possible factor in a model for the model to be "realistic," regardless of how important each factor may be to the overall outcome of the combat assessment. By focusing on the first-order effects *and all of their interactions,* one can retain a top-level view of the simulation. (Freqently, representatives of a particular function, whether it be pilots, tank commanders, or logisticians, feel that unless their favorite function is explicitly modeled, then the model is "invalid." This view is prevalent regardless of the study requirements, the resolution of the model, or the fact that their favorite function is four levels "down in the weeds." Too many models have become so loaded down with detail that they are no longer useful for purposes of analysis.)

After defining the first-order effects in each functional area, one should ensure that the situations that are modeled in each functional area define a partition. (A partition is a division of possible outcomes that covers the space, and where no outcomes overlap. For example, the partition of the set of outcomes of a coin toss is "heads" and "tails.") That is, make sure that nothing falls through the cracks, even if one has to define a category called "other situations." Additional resolution to each functional area may be added as needed later, subject to time, resource, and data constraints. The literature sometimes refers to this as "granularity."

The addition of more detailed resolution to a basic Referee model may take many forms. For the functional areas, it may mean distinguishing more clearly between different types of situations that were previously lumped together. Or it may mean replacing a qualitative definition with a more quantitative one. For example, strategic air surprise may be changed from only the aircraft alert and dispersal rate to more precise warning times based upon strategic intelligence assets. In models that include qualitative factors, it is relatively easy to add increased quantitative resolution.

Another area where additional resolution may be added is in the model's geographic resolution. For example, one may wish to add features that were identified as being important after the model was run intially. In Table 3, the geographical resolution distinguishes between the northern European theater (NEUR) and all other theaters, and between a target called "Pass" at a point called "Wedge." This refers to a strategic fortified mountain pass known as the "Finnish Wedge." (We usually have more column space in the model, so we spell out the key point names for easy reference.)

TABLE 3: Determine Type Battle

Decision Table

theater-name	point-name	type-target	prep-days	/	type-battle
NEUR	Wedge	Pass	>=2		Fortified
NEUR	Wedge	Pass	<2		Prepared
NEUR	--	--	>=3		Prepared
NEUR	--	--	<3		Hasty
--	--	--	>=4		Prepared
--	--	--	<4		Hasty

[End Table].

This example also distinguishes between whether or not the defender has had at least two days to prepare for the upcoming battle. If so, the first row triggers, and the type of battle is "Fortified." Otherwise, the type of battle is "Prepared," which is less advantageous for the defender than "Fortified." For any other point in this theater, the type of battle is at best "Prepared" after three days of defender preparations, or "Hasty" otherwise. Note that this table also defines at least four days defender's preparation time to achieve a "Prepared" defense for any other theater except "NEUR."

To add geographic resolution, one can identify additional points of strategic interest. For example, the first two rows may be duplicated, and "Wedge" replaced by "Kiruna," representing a different strategic mountain pass in the NEUR theater. The conditions for the days of required preparations and the effect on the type of battle may be varied for this newly defined point. Similarly, one could define a whole new theater by adding rows below those beginning with NEUR and before the last two "default" rows.

Note that the Referee model requires data to be stored only on an exception basis. This reduces the data requirements, as well as the model compilation and run times. Furthermore, the model has default values that allow the analyst to run with only the "bare essentials" of terrain definition. Being able to run the model early in the process allows for the interactive sandtabling to begin as soon as possible. In that way, only those areas of interest to the analyst need to be refined, while the remaining areas require little "model overhead."

One way to keep RAND-ABEL™ tables manageable is to distinguish between those factors that must be assessed simultaneously and those that may be assessed sequentially. Since each table is supposed to define a partition, the net effect is that all of the input variables are being assessed simultaneously to produce the desired output. Tables that follow one another are assumed to be processed sequentially. The output of an earlier table may be used in a later table, but not vice versa.

Sometimes, one can break down very large tables into smaller, more manageable tables without violating the model design. For example, all of the factors that contribute to ground combat theoretically must be considered all at once to obtain a "realistic" outcome. However, some subset of the input factors (e.g., the size of the battle, the climate, and the aggressiveness of each side) may be considered as a single input variable called "combat intensity." Therefore, one can define combat intensity in an earlier table, and replace the three or four columns in a large table with a single column, thereby reducing its size.

As the simulation is run, the analyst observes the sequence and outcomes of each situation to ensure credibility and to identify new issues for analysis. The emphasis is on defining the scope of situations to be allowed in the course of the study, and on selecting or defining the best representation. For example, the assessment process for a "breakthrough" type of battle may be much different than for a "prepared defense" or a "flank counterattack." The assessment process selected in the tables may differ for each type of battle. Too many combat simulations have only a single combat assessment algorithm that is applied to all combat situations, no matter how ridiculous that algorithm may be when applied to a particular type of engagement.

ANALYSIS TECHNIQUES

Just as new modeling techniques were invented to take advantage of the RAND-ABEL™ features, new analysis techniques were created to take advantage of the Referee model features. Many different types of postprocessing displays can be created from qualitative outputs. All of the standard quantitative displays commonly used in theater-level combat simulations are available, such as FLOT traces, force ratios, and assets over time. Several useful qualitative displays may be used as well. For example, air control range over time is a useful display to show how each side's efforts to attain air superiority progress over time. Or, one could plot the air control over a key target over the course of the conflict. Most of these standard quantitative and qualitative displays can be created through either the RSAS "Graphics Tool" or "Maptool" displays at the end of the run, or at any time during the run.

One of the more useful classes of displays we recently created includes "statistical" displays. Since a large number of different types of battles are described by the model, we decided to analyze this data with some standard statistical methods. The types of battles currently represented in the model are listed in Table 4. (NOTE: This table is truncated. Mirror-image the engagements with Blue as attacker to obtain the lower half of the table.) These types of engagements may be described as being Red favorable, Blue favorable, or neutral (where Red and Blue denote opposing sides in the simulation).

TABLE 4: Ranked Types of Battles

Decision Table

TYPE ENGAGEMENT	RANKING
Red at final objective	Most Red favorable
Red unopposed advance	Very Red favorable
Red pursuit of Blue	Very Red favorable
Red breakthrough	Red favorable
Blue delay of Red	Neutral
Blue hasty defense	Neutral
Blue prepared defense	Blue favorable
Blue fortified defense	Blue favorable
Red cannot attack Blue	Blue favorable

[End Table].

Histograms showing the frequency of each type of engagement over the course of a 30-day conflict are shown in Figure 1. Two different scenarios are represented in Figure 1, allowing the analyst to compare the difference in the number and degree of Blue favorable (BF) or Red favorable (RF) engagements. For example, the first case, (A), may be a surprised Blue scenario, while the second case, (B), may be a prepared Blue scenario.

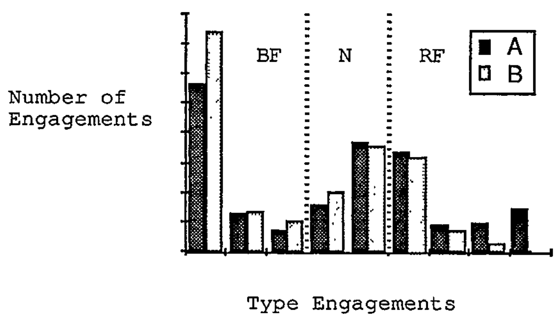

FIGURE 1: Distribution of Engagements

These histograms can be useful in several ways. As one example, if a specific asset is very useful in a certain type of engagement, then its potential usefulness can be estimated from the frequency that these engagements occur. The inclusion of this hypothetical asset in the simulation may also change the distribution of engagements. The difference in the distributions provides a way to measure the benefit of the new asset in both qualitative and quantitative terms (the number and type of engagements). Most traditional measures of combat value have been limited to attrition and movement rates, which has made it consistently difficult to measure the benefits of non-lethal assets, such as intelligence, communications, command and control, and logistics support systems.

In addition to histograms, one can also condense this information into a "percent of time" chart. One calculates the percent of time or fraction of engagements that were Red favorable, Blue favorable, or neutral, as displayed in Figure 2. The scenarios may represent different levels of combat intelligence capabilities. This format allows one to readily compare how well each side performed over a wide variety of scenarios.

One can create several other types of useful displays for qualitative factors. Some examples include air control range for each side over time, the air control over selected targets over time, or the degree or frequency of strategic or tactical surprise over time. The main point is that it is just as easy, if not easier, to display qualitative factors in a useful manner as it is to display quantitative factors. Furthermore, one can define the qualitative factors around those issues that are important to the sponser, rather than simply those values that the analyst could traditionally measure.

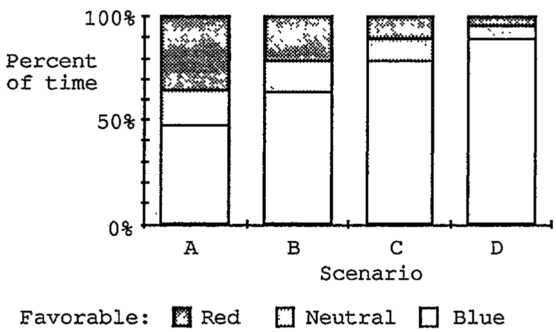

FIGURE 2: Percent of Time Favorable to Each Side Over Four Scenarios

Sensitivity analysis is a key part of this modeling process. Since an analyst using this methodology is usually dealing with a low resolution model, one should frequently examine the conditions under which certain events will or will not take place. This is especially true if one is using the model as part of a screening process to reduce the set of feasible options. Sensitivity analysis is often done in "batch" mode, rather than with interactive sandtabling. The model is designed to facilitate multiple runs in batch mode through the global variable "run number." This variable may be added as an input column to any existing table. If one is performing, say, ten runs with different parameter setting for each run, the run number determines which parameter setting will be used for each run, and acts as an audit trail for the parameter values used in each run.

Model verification and validation are no more difficult with qualitative modeling, since there are usually fewer variables to examine. One can vary the qualitative factors as easily as quantitative factors. It may even be easier to perform model verification with this model than with purely quantitative factors, since it may take several different quantitative factors to adequately define a qualitative factor. For example, one may vary the single value of "tactical surprise" in the Referee model, whereas many quantitative variables may need to be varied in other models to achieve the same model variation.

Another important need in the military modeling community has been for model assessment by independent evaluators. To assist evaluators, sponsors, and analysts in assessing a given model, the model parameters, logic, and assumptions must be understandable. As shown above, the RAND-ABEL™ table structure is very helpful in this process.

RECENT APPLICATIONS

The RAND-ABEL™ based Referee models have been successfully applied to several long term and short term analysis projects at RAND, as described below.

(1) The Referee modeling approach was developed for the RAND Strategy Assessment System (RSAS). The RSAS needed a generic, quick, and flexible theater-level model to actually become a global military assessment model. The first fully operational theater-level Referee simulation was completed in 12 months. This model was called "S-Land" for Secondary Land theater model, since it was applied to all theaters in the RSAS except Central Europe and Korea (which is handled by the CAMPAIGN-MT model). Since then, the capabilities of S-Land expanded enough to participate in the RSAS effort to become a variable resolution model. Since the original S-Land theater may now be applied to all theaters of operation including Central Europe and Korea, its new RSAS name is CAMPAIGN-ALT. (The CAMPAIGN-MT model is a more detailed representation of Central European and Korean combat than is CAMPAIGN-ALT. The CAMPAIGN-ALT theater is suggested when one needs to tailor a model to satisfy specific study needs.)

(2) The S-Land model was recently applied to two quick-response studies at RAND. The first dealt with measuring the effectiveness of an airborne conventional weapons platform performing nontraditional

missions in various theaters of operation. The study was performed in six weeks by RAND personnel who were not the authors of the model. The results were successfully briefed to the Commander-in-Chief (CINC) of a major air command. The second quick response study involved measuring the operational impact of special operations in different theaters of operation. These study results were also well-received. In both of these cases, the S-Land model was used in the interactive sandtabling mode, with extensive offline analysis performed between model runs.

(3) Longer studies have also been supported by S-Land analysis. Southwest Asia is a perennial favorite for analysis. Another study required the inclusion of integrated warfare (i.e., chemical and nuclear) assessment models.

(4) The S-Land model has been used for several operational staff training exercises at the National Defense University. Course support was provided for the examination of operations in the AFSOUTH and AFNORTH regions of NATO. Results were briefed to both the current and previous CINCs of AFSOUTH.

(5) A study on the combat value of intelligence has recently been initiated at RAND. Due to the qualitative nature of intelligence, the Referee approach was selected as the analysis tool. This project will include a low-resolution model of the command and control process on each side, and the effects of different levels of intelligence support on the operational effectiveness of a force. A rudimentary Referee model of deception operations is also planned. This project will use the RAMP, described below.

(6) The global options project is an effort to look 25 to 30 years into the future and "forecast" likely issues of strategic importance. The Referee modeling approach will be used extensively to define and analyze the interactions of many nations involved in hypothetical global events.

(7) The RAND Strategy Assessment System and the qualitative modeling approach has been so successful at RAND that an unclassified version of the system has been developed. The RAND Analytic Modeling Platform (RAMP) has all of the system features of the RSAS, including the RAND-ABEL™ language, graphics tools, and data editor. In addition, prototype code similar in capabilities to the S-Land or CAMPAIGN-ALT model have been created for unclassified analysis purposes. The intent is to eventually perform technology transfer and make the RAMP available to nongovernment agencies. (At the moment, only government agencies may obtain copies of the RSAS.)

OBSERVATIONS

Analysts who have used the Referee model appreciate the small number of important parameters, as opposed to the massive number of parameters used in many larger models. The model parameters and logic are readily accessible to changes so that they may be tailored to the study needs. When changes are made to the data or model logic, they may be readily understood and explained to the client because there are few "hidden" assumptions in the model.

The sandtabling technique was developed as part of the RSAS and is becoming popular at RAND. One can understand what is occuring in the model because the analyst can watch events unfold. The analyst can stop the action, examine details, change simulated conditions, and then proceed with the run. This shortens the time it takes to learn the model, and allows for rapid prototyping of new features. An analyst can quickly produce preliminary runs with low resolution to gain quick insights, screen options, and identify issues of interest.

Due to the popularity and success of these features, it is expected that future analytic tools will more frequently have the features associated with the Referee model methodologies.

REFERENCES

Allen, P.D., and Wilson, B.A., *The Secondary Land Theater Model*, N-2625-NA, The RAND Corporation, July 1987

Bennett, B.M., Jones, C., Bullock, A.M., Davis, P.K., *Main-Theater Warfare Modeling in the RAND Strategy Assessment System*, N-2743-NA, The RAND Corporation, August 1988.

Davis, P.K., *Applying Artificial Intelligence Techniques to Strategic-Level Gaming and Simulation*, N-2752-RC, The RAND Corporation, 1988.

Shapiro, N.Z., Hall, H.E., Anderson, R.H., LaCasse, M., *The RAND-ABEL™ Programming Language: History, Rationale and Design*, The RAND Corporation, 1985.

Shapiro, N.Z., Hall, H.E., Anderson, R.H., LaCasse, M., Gillogly, M, Weissler, R., *The RAND-ABEL™ Programming Language: Reference Manual (first revision),* N-2367-1-NA, The RAND Corporation, 1988.

ABOUT THE AUTHORS

PATRICK ALLEN received his Ph.D. in mineral economics from the Colorado School of Mines (1980). He is an operations research analyst at RAND and project leader for the NATO Modernization Study, Air-Land Warfare Methodology Project, and involved in numerous other military simulation activities. He is also a Captain in the U.S. Army Reserves and formerly served at the U.S. Army Concepts Analysis Agency. Allen has consulted for several firms on a variety of topics, and authored or coauthored a number of articles on military simulation and related issues.

The RAND Corporation
1700 Main Street
Santa Monica, CA 90406
(213) 393-0411 x7340

BARRY A. WILSON received a B.S. in Geology from the California Institute of Technology (1981) and his M.S. in Computer Sciences from the Califronia State University at Fullerton (1984). He is a senior programmer analyst working for the RAND Corporation on the RAND Strategy Assessment System; in charge of software development and support for analytic warplans and the RAND-ABEL™ Referee combat model.

The RAND Corporation
1700 Main Street
Santa Monica, CA 90406
(213) 393-0411 x6472