Proceedings of the 2007 IEEE Symposium on Computational Intelligence in Multicriteria Decision Making (MCDM 2007)

Exploring Robustness of Plans for Simulation-Based Course of Action Planning: A Framework and an Example

B. Chandrasekaran
Department of Computer Science and Engineering
The Ohio State University, Columbus, OH 43210 USA

Mark Goldman
Aetion Technologies LLC
Business Technology Center
1275 Kinnear Road
Columbus, OH 43212 USA

Abstract- **Planning requires evaluating candidate plans multi-criterially, which in turn requires some kind of a causal model of the operational environment, whether the model is to be used as part of evaluation by humans or simulation by computers. However, there is always a gap - consisting of missing or erroneous information - between any model and the reality. One of the important sources of gaps in models is built-in assumptions about the world, e.g., enemy capabilities or intent in military planning. Some of the gaps can be handled by standard approaches to uncertainty, such as optimizing expected values of the criteria of interest based on assumed probability distributions. However, there are many problems, such as military planning, where it is not appropriate to choose the best plan based on such expected values, or where meaningful probability distributions are not available. Such uncertainties, often called "deep uncertainties," require an approach to planning where the task is not choosing the optimal plan as much as a robust plan, one that would do well enough even in the presence of such uncertainties. Decision support systems should help the planner explore the robustness of candidate plans. In this paper, we illustrate this functionality, robustness exploration, in the domain of network disruption planning, an example of effect-based operations.**

I. Simulation Models and Reality

Planning generally entails generation and evaluation of several alternate plans and selecting one of the plans based on multiple, potentially conflicting criteria. Evaluating plans along multiple criteria requires, except in trivial situations, simulation of the plans. Simulation, whether done in a computer or in one's thought process, requires a causal model of the operational environment, including, in adversarial planning, models of the adversaries.

However, any planning based on simulations has to face an *intrinsic* problem of simulation models, namely, that there is an inevitable gap between the models and reality. Comparing and selecting plans on the basis of simulation results can be problematic in the face of simulation model gaps, uncertainties and errors.

The concept of model errors is easy to understand: the model has incorrect information about reality. Gaps are missing aspects of reality. These are harder to protect against, since generally we don't know what we don't know. Uncertainties are normally handled in a probabilistic fashion. However, as Bankes [Bankes] points out, difficult decisions about the future cannot be made on the basis of expected value for at least two reasons. The first is that if the outcome measures are correlated, then individual expected values will give a false picture of the future, but this can be taken care of by a more sophisticated stance towards computing the joint expected values. More seriously, however, expected values-based assessments fail to indicate both dangers and opportunities that may lie nearby, and possibilities for driving the future states to avoid dangers and exploit opportunities. The decision maker needs to drive robustness exploration in the larger context of actions available. For example, for a proposed military course of action (COA), suppose there is great uncertainty in the mind of the planner about the likelihood of rain. Perhaps with some expenditure of effort better information about the likelihood of rain can be obtained, but if COA is robust with respect to rain, the resources that would be used in obtaining better information about rain could be put to better use. Even more importantly, the decision maker might be able to determine that while the probability of achieving the objective varies considerably depending on whether it rains, if subgoal S is achieved, then rain doesn't affect the final outcome very much, and that increasing a certain COA parameter above a particular value can ensure that S can be achieved, rain or shine.

Another example is a potential large uncertainty in enemy position. The decision maker might explore a how to make a COA robust with respect to this uncertainty. It would be good for him to discover that it is possible to design a COA such that the initial stages of the COA are largely oriented towards getting more information about the location, and such that based on the information, later stages of the COA can be assigned appropriate parameters. Again, the goal is less to simply obtain estimated values for the outcomes than to explore the space around a COA to see whether it is possible to make it more robust.

The gap between model and reality is not an issue for computer simulations alone – everyday human planning uses models as well. However, humans have evolved several commonsense heuristics to cope with the various problems with the models we just identified. Humans often know in their guts which aspects of their model of reality they are especially unsure of. For important aspects of reality they are unsure of, they may not simply play probabilities, but instead selectively explore the world for additional information. Or design plans that in their earlier steps combine achieving goals of the plan with information gathering, and have branch points in the plan based on the information. Or work back from especially worrisome outcomes back to identifying aspects of reality that may be relevant even though their immediate simulation model contained no information about them.

That robustness of decisions is important is not novel, but the in-principle nature of the gap between simulation models and reality, and how it forces a shift away from optimality to robustness is not widely appreciated, specifically in computer-based decision support. More and more complex simulation models are being built as the enabling technology for decision making, with nary a warning or awareness of the problems that we have referred to. The problem is compounded by the fact that simulation models in many complex planning problems are composed from bits and pieces written by several different people, and the planner may not even be aware of possible gaps. This leads to a tendency to place too much reliance on the results of the simulation, and hence on the plans based on them.

What the planner needs, and what decision support systems must help with, is a sense of what model assumptions have what kind of effect on what dimensions of outcome, and correspondingly, how to make the desired outcomes be less sensitive to the gaps and uncertainties. That is, the planner needs to develop a sense of how robust the plan is with respect to possibly problematic aspects of the model and to modify the plan to be more robust if it is not robust enough.

Handling the gaps as well the uncertainties requires a shift in point of view from optimality to robustness. In order to realize the full potential of the vastly increased search spaces made possible by computer simulation, it is essential that the decision support system empower the planner to explore the plans for robustness of the selected plans. The research challenge is to identify, and incorporate as part of decision support systems, a variety of techniques by which the selected plan can be tested for sensitivity to various model assumptions and uncertainties. This is a tall order.

This paper is an initial effort to illustrate the idea of robustness exploration in the context of a multi-criterial decision making framework.

II. The Seeker-Viewer-Filter Architecture for Multi-Criterial Decision Making

This research is being conducted in the context of the development an integrated planning support system in which the goal is to support the entire planning cycle, as illustrated in Fig. 1. A core of this is the *Seeker-Filter-Viewer* architecture [Josephson], where the Seeker generates decision alternatives and evaluates them on multiple criteria, the Filter produces the Pareto-optimal subset with respect to the criteria, and the Viewer provides a variety of interactive visual means by which to view the alternatives on one or two criteria, examine the tradeoffs, select and narrow the choices. The Viewer is also a means of understanding the decision space: e.g., the user can see at what points in the decision space, how a change in one of the variables is related to change in selected other variables. In fact, this feature of the Viewer will be used in the robustness investigation in the example later in the paper.

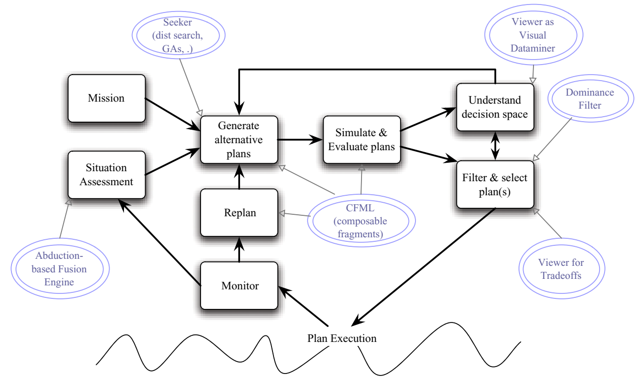

Fig. 1. An integrated planning framework

III. Network Disruption Planning: An Illustrative Example

Effects-Based Operations (EBO's) [Davis] are military operations that focus "on planning, executing, and assessing military activities for the effects produced rather than merely attacking targets or simply dealing with objectives," and complement traditional target-based courses of action (COA's). Examples include so-called "psywar" (psychological warfare) operations and disrupting aspects of infrastructure, such as communication and transportation networks, that support enemy military operations. In the example we consider in this paper, the robustness of a network disruption plan is explored with respect to various assumptions built into the model of the enemy network used in simulating the plan.

The network we used in our experiments is sketched in Fig. 2. The network includes power, communication and control nodes, and the links may be power or communication or both.

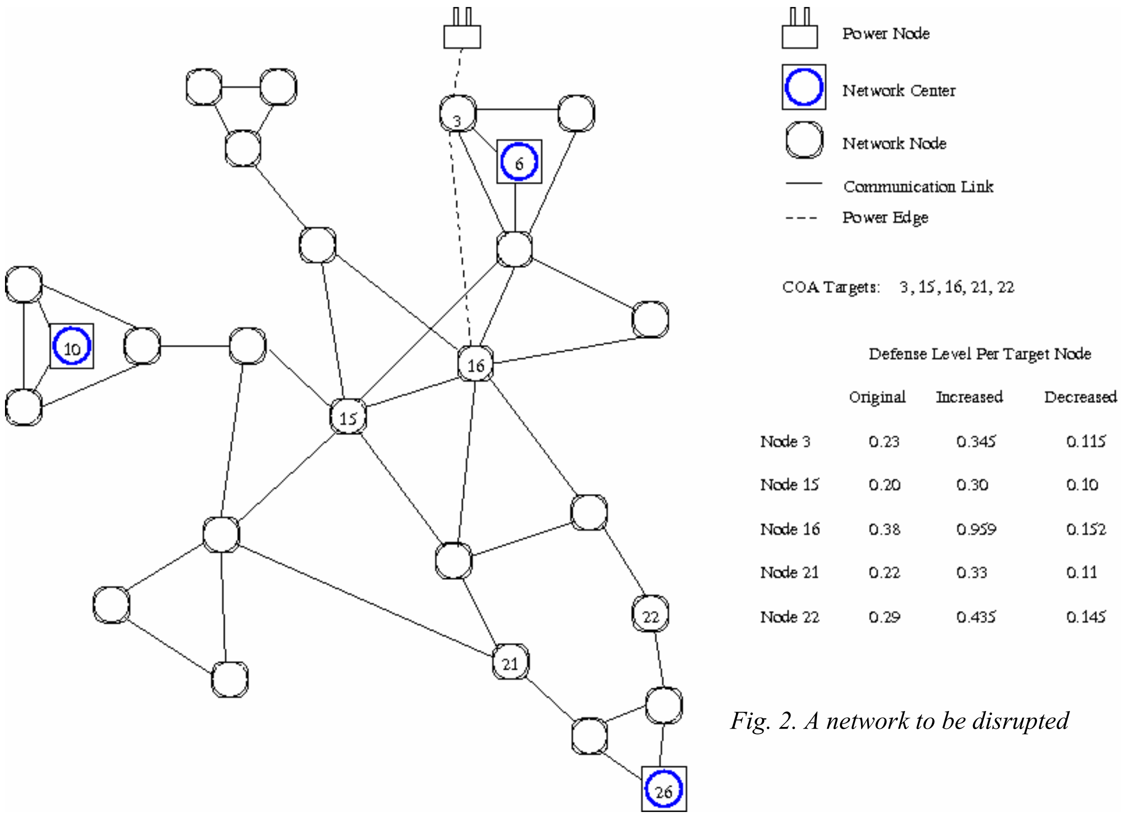

Defense Level Per Target Node

	Original	Increased	Decreased
Node 3	0.23	0.345	0.115
Node 15	0.20	0.30	0.10
Node 16	0.38	0.959	0.152
Node 21	0.22	0.33	0.11
Node 22	0.29	0.435	0.145

Fig. 2. A network to be disrupted

Some of the links are directed (one-way), others bi-directional. Nodes 6, 10 and 26 are command and control nodes. A network disruption plan (or plan from now on) is abstractly a list of nodes and links to be attacked. The goal of the disruption plan is to reduce the bandwidths for communication between these nodes. These three bandwidths provide three criteria, and expected friendly and enemy casualties provide two further criteria in our experiments. In the experiments, we use the number of fragments into which the network is split as a useful proxy for bandwidths. Even though we don't use it in this illustrative experiment, collateral damage could be an especially important criterion in practice in such EBO operations.

In order to assess the potential performance of a plan, it needs to be simulated, and this requires a fully specified model of the relevant aspects of reality: the structure of the network, the number and type of nodes, their connectivity, the intentions and capabilities of the defenders and the attackers, and the repair capabilities are just some examples. Here we draw attention to one set of parameters in the model that will be particular interest in our experiments: "defense strength" parameters (denoted by d_i for node i) associated with nodes. This parameter represents the probability that the node will survive an attack. It also determines the casualties suffered by the attacker of the node. Because the model is stochastic, each run might produce different values, so several runs are performed and expected values for the criteria of interest are generated.

The robustness exploration project that we are reporting on is part of a larger project that includes EBO plan space exploration, including generation and simulation of multiple plans, Pareto-filtering on multiple criteria, and further down-selecting by the human, based on trade-offs. A companion paper submitted to this conference [Carroll] describes the use of an evolutionary algorithm for generating a plan. For purposes of this paper, however, we assume that a good candidate plan has been generated. Specifically, the plan is the list of nodes to be disrupted (attacked): {3, 15, 16, 21, and 22}.

IV. Robustness With Respect to Model Assumptions

Almost any aspect of the model used to simulate an attack plan on the network is an "assumption" about the world of interest, the impact of which on plan performance might be of interest. However, the planners generally have more confidence in some aspects of the model than others. To illustrate the methodology, we take the set of defense strength parameters as the assumptions that the planner is most unsure of[1]. This is not unrealistic – the resources available to the

[1] While in some situations the uncertainty about the parameters can itself be treated probabilistically, in others such an approach is not appropriate. If by "most unsure of," the planner means that the source of information about the defense strength is unreliable and may have questionable motives, the uncertainty of the planner is best treated as an example of deep uncertainty. Treating it by just putting a

defenders are likely to be among the most common sources of uncertainty. The model has some assumptions of these defense strength parameters built in. In Fig 2, the first column of the Table to the right lists these assumptions. Given the plan "Attack nodes {3, 15, 16, 21, and 22}," we would like the robustness exploration to help us answer questions such as the following.

(1) Which nodes' defense parameter assumptions are most critical to the outcomes?

(2) If we know that the outcomes are especially sensitive to (i.e., not robust with respect to) assumption about the parameter for node *i*, is there any other assumption whose correctness would reduce the plan's sensitivity to the assumption about node *i*?

The answer to the first question might suggest adding resources to strengthen the attack on node *i*. Or, the planner might focus some of the information gathering resources on checking the assumption about node *i*. Suppose it turns out that it is hard to get additional information about node *i,* say because the information about the defense allocation to the node is even better defended that the node itself. In this case, the answer to the second question may be useful. For example, suppose the robustness investigation reveals that if node *j*'s parameter is in a certain range, then the plan is less sensitive, i.e., more robust with respect, to the assumption about node *i*, then it may be worth redirecting the information resources to find out more about node *j*.

In passing we note that there is no implication from the example question (2) that only pair-wise interactions are important. There could be cases where the planner might have reason to hypothesize that a parameter's sensitivity might depend on two other factors. In general, however, humans find it is easier to make hypotheses regarding the interaction of variables when the numbers are small rather than large.

The Methodology

The general methodology is simple: Vary the model assumption in question over its range, for each variation simulate the plan and compute the criteria of interest. Examine how sensitive are the criteria of interest to the variations. This examination can be done in several ways. We first change only one model assumption and investigate how sensitive the criteria are to this change. The user might then investigate possible *interaction*s between assumptions. For example, while the plan may turn out to be not robust with respect to assumption *A1*, it might turn out to be acceptably robust with respect to *A1*, if another model assumption, say *A2*, is changed in a certain way.

Given a system characterized by independent variables $\{x_1, ..x_i…x_n\}$, dependent variables $\{c_1,…c_j…c_m\}$, and a set of samples $\{s_1,…s_k\}$ consisting of instances (candidates, runs..) of the system behavior for various values of *x*'s, a common way to estimate the sensitivity of a dependent variable to the independent variables is to compute the *correlation coefficient* based on the samples. For our application, for Question (1), this requires that we compute, from the data from simulation runs, the correlation coefficient between defense strength parameter d_i and the outcomes of interest. For interactions between assumptions -- to decide how sensitive a dependent variable c_j is to independent variable x_i, given that specified other independent variables have a certain values or occur in certain ranges, the answer for our Question (2) -- the correlation coefficient would be computed between x_i and c_j, using only the runs that satisfy the constraints on the other independent variables.

In practice, there are many problems in simply computing the correlation coefficients for our needs, problems that make the kind of visual interface that we will soon describe an attractive option. First, correlation coefficient formulae look for a linear relation between the variables. If there is a positive correlation over a certain range of x_i and a negative correlation for another range – quite often the case in real world situations -- such computations might return zero as the value for the correlation coefficient. Second, while it is easy enough to calculate correlation coefficients for each single independent variable and each dependent variable, the number of combinations to be considered for pairs or triples of such variables can grow very large. Interactive data analysis, where the analyst uses domain-specific expertise to generate hypotheses, which he then interactively tests, which then results in additional hypotheses is often an effective way to fight the combinatorics. These considerations lead to the need for a visual interactive interface for exploring the structure of simulation data to form and test hypotheses about robustness.

A. The Viewer

The Viewer is described in [Josephson] and the thesis by Iyer [Iyer]. The Viewer enables the analyst to have on one screen a combination of several types of displays of data. The first is a *spectrum* display. For a given variable, the corresponding spectrum displays all the candidates, each at the location of its value for the variable. Two dimensional scatter diagrams can be created for pairs of variables, and again the candidates are displayed at the location determined by the values that each candidate takes for the two variables. These two display types work well when the variables are continuous. When a number of such displays are created in one screen, the same data item will appear in different locations in different diagrams depending on the values that the item takes for the variables in the display.

When the variables take only a small number of discrete values, as in our application, histograms are used to represent the number of entities in the run that take a specific value (in the spectrum), and in the scatter diagram boxes whose areas are proportional to the number of entities that take the

---

probability on the reliability of the source is exactly what we propose is inappropriate.

specified discrete values are placed in the coordinate location, as in Figure 3.

A key functionality is cross-linking: a set of alternatives can be chosen in one of the diagrams – the selected ones change color from blue to red – not only in the selected diagram but in all the other diagrams.

B. Exploring the Robustness of the Plan

Simulation. The attack plan was simulated a total of 5000 times (producing 5000 runs), under various defense parameters for the nodes. The first group, the *base group*, consisted of runs with the defense parameters all at the assumed values. The second group consisted of a sequence of sets of runs: in each set, one of the parameters was increased to various higher values[2] from those in the base model, and the rest were unchanged; one of the parameters was decreased to various lower values from those in the base case and the other parameters were left unchanged. This was done for each of the nodes. The third group consisted of pair by pair changes: e.g., d_{15} and d_{16} were each increased, d_{15} was increased and d_{16} decreased, and d_{15} was decreased and d_{16} increased. This was done for all pairs of nodes in the plan. It is worth noting that the relation between the d_i parameters and the probability of the node being destroyed can be complex, and it itself part of the set of model assumptions.

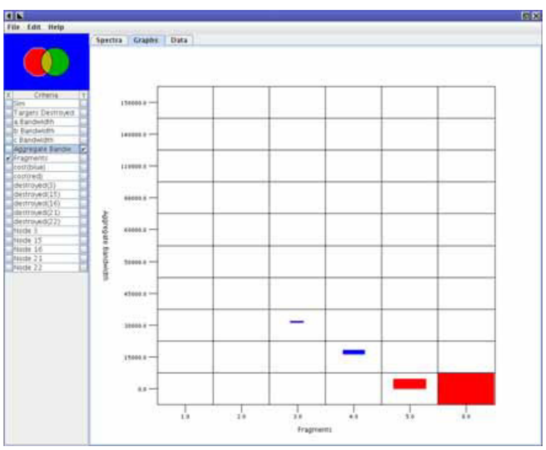

Fig. 3. The simulations are plotted as the number of network fragments vs the aggregate bandwidth of the network resulting from the execution of the COA.

C. Assessing Sensitivity.

Which assumption about nodes' defense parameters are most critical for what aspects of the plan's performance? To start with, the number of fragments and average bandwidths are both measures of how well the plan is performing: the higher the number of fragments, or the lower the bandwidth, the better is the plan. Let us look at Fig. 3, which plots the number of fragments against the bandwidth. It appears that there really is not much reason to distinguish between 5 and 6 fragments; both of them reduce the bandwidth to 0. In fact, using the additional consideration that a bandwidth value below 15,000 units is quite good, fragments 4, 5 and 6 are more or less equally good enough.

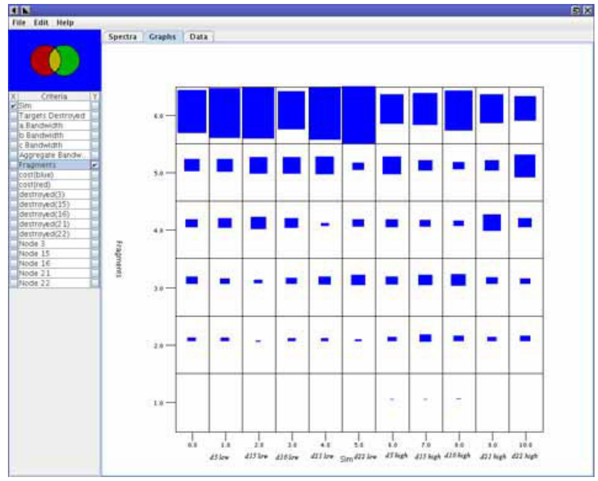

Fig. 4. Various simulation runs (X-axis) plotted against the number of fragments (Y-axis). X-axis: left-most (label 0) is the base case, the next 5 (labeled 1 through 5) correspond to the parameters for nodes 3, 15, 16, 21 and 22 being set low respectively, and the next 5 (labeled 6 through 10) to the parameters for the same nodes set high respectively.

In Fig. 4, the analyst can see that the runs with labels 1, 2, 4 and 5 seem to maximize the number of cases with 4 or more fragments. Also, simulations 1 and 2 have the fewest number of runs with 3 or fewer fragments. Conversely, though simulation 8 is fine with respect to having a large number of instances with 6 fragments, it also has a large number of cases with non-fragmented network. Simulation 6 has both a small number of 6 fragments and a large number of non-fragmented cases. Labels 1 and 6 correspond to low and high values for d_3, respectively. One hypothesis that the analyst can make is that the COA's performance changes significantly as the value of d_3 changes, flagging it as one of the parameters with respect to which the COA is possibly not robust. Different lines of analysis might identify different nodes, but, given that the analyst is concerned with the worst cases as well as the best cases, the identification of node 3 as a candidate for robustness concern is reasonable.

Another line of analysis (we don't show the figures because they consume too much space) leads the planner to wonder about node 16. Specifically, when the analyst looks at cases where the d parameters of pairs of nodes are increased, he notices that there are a couple of cases that stand out in regards to preventing the network from fragmenting: one where d_3 and d_{16} are increased and the other where d_{16} and d_{21}

[2] The details of how these values are chosen are not important for the illustrative point we aim to make.

are increased. This and other similar analysis fingers d_{16} as a parameter that the COA is somewhat sensitive to, especially with respect certain cases of bad worst-case behavior.

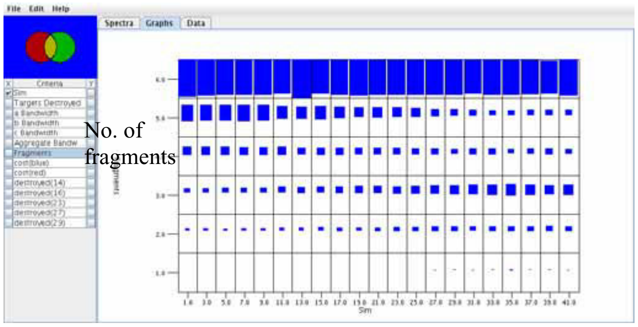

X-axis shows simulation run labels: The values of d_{16} increase along this axis, with the base value case in the middle.

Fig. 5. Value of d_{16} plotted against the number of fragments into which the network splits.

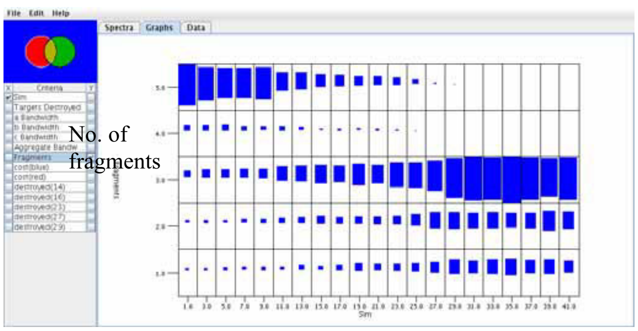

X-axis shows simulation run labels: The values of d_{16} increase along this axis, with the base value case in the middle.

Fig. 6. Same as Fig. 5, except that in all these runs node 3 is not destroyed

The behavior of node 16 gives rise to an example of Question (2): Are there any conditions on other nodes that might reduce the plan's sensitivity to those values of d_{16} that correspond to the worst case behavior? Note that this is indeed an issue of robustness with respect to model assumptions: if we could be sure that the value of d_{16} is not one that corresponds to bad performance, we don't need to worry. But given the uncertainty, the planner might look for ways to moderate the effects. One way is to identify another node that he do something about in the COA to ensure good performance, in spite of the uncertainty about d_{16}. If he feels more confident about the assumptions regarding this node, or he has reason to think that it would be easier to acquire better information about this node, then identifying such a node would be useful.

Fig. 5 is a display that shows on the x-axis simulations ordered by the value of d_{16}: in the middle is the simulation for the base value and the value of d_{16} increases from left to right. The y-axis is the number of fragmentations, with the bin sizes corresponding to the numbers of runs that produced the corresponding number of fragments. Fig. 6 is a similar display, except that only the runs where node 3 was not destroyed are included. It is clear that when node 3 is not destroyed, the attack plan overrides any effects of high values of d_{16}, including those corresponding to poor worst-case performance. This is a reassurance to the planner that as long as he can take steps either to acquire more reliable knowledge about the value of d_3 to make sure it is not high, or to modify the COA to ensure that the chances of destroying node 3 are high, he can worry less about not knowing enough about d_{16}.

V. Discussion

This paper has two parts. In the first, we make a general case for a shift in perspective in multi-criterial decision-making from optimal decisions to robust decisions. This is because almost all decision-making involves modeling the world of interest and such models are bound have gaps and errors, and it is more important to ensure that the decisions we make are not fatally compromised by such model errors than to optimize based on extreme confidence in our world models. In the second part, we illustrated the approach with an example drawn from a planning domain where we gave examples of the use of a visual interface to generate hypotheses about how robust a proposed plan is with respect to various model assumptions.

It is useful to list a number of caveats regarding the example and the approach used. First, the example was used to *illustrate* some of the ideas in the first part and not as a realistic instance in the network disruption planning domain. Second, using the visual interface is a good way to start making interesting hypotheses, but statistical work would need to be done to verify or reject the hypotheses. The advantage of the visual interface is that it can help identify correlations over subranges of values that are unlikely to show up in correlation calculations over the entire data set. In addition to making available histograms as in Figures 3-6, the Viewer, the interface described in the paper, provides other ways, such as "spectra," and scatter plots, to display and analyze data to explore various robustness hypotheses. Third, robustness analysis is computation-intensive so the planner has to be reserve it for assumptions that he has reasons to worry about. In fact, it is impossible to test a plan with respect all aspects of a model: if one is in such a state of uncertainty about the domain, planning with any hope of achieving success is a fool's errand.

Having an idea about the robustness of a decision can help improve the decision in many ways. Information gathering resources can be selectively focused on acquiring better information regarding those model assumptions that the plan is most sensitive to. The plan may be modified to overcome any negative effects of lack of knowledge. For example, in our example, the plan was quite sensitive to assumptions about node 3. One way to respond to this knowledge might be to add plan resources so it succeeds even if node 3 happens to much better defended than assumed. In the absence of this specific information about node 3 being crucial, information-gathering or plan enhancement resources would need to be allocated to confirm all possible assumptions, reducing the benefits. Having the information enables the planner to focus the resources specifically to node3.

In our illustrative example, we started with a plan already generated, presumably selected multi-criterially from a set of plan candidates, and described evaluating the robustness of the selected plan as a last stage. In fact, however, generation and modification of plans and robustness evaluation can interact with other in various ways. The planner might be unhappy with what the robustness analysis revealed, and might change the plan, or create a new plan, and follow this up with another robustness analysis. Thus, generation and robustness analysis may form a cycle until satisfaction is obtained.

In principle, during the multicriterial evaluation, comparison and selection stage, the plans's robustness properties might form additional criteria. There are two points to keep in mind, however. One is that assessing the robustness is an exceptionally computation-intensive process – recall that in our illustrative example, the selected plan was simulated thousands of times. Thus, while comparing plans with respect to robustness would be useful, the plan space would need to be significantly pruned, so that only a small number of plans that survive selection in the first selection round are assessed for robustness and then compared. More importantly, however, robustness conceptually is not the same sort of criterion as the ones that the plan that relate to the plan's purposes. Given two plans with approximately equal performance figures with respect to say the probability of mission success, it makes sense to choose the one with better robustness properties, but if one plan has an expected value of 50% for mission success probability with a variance (a robustness measure) of 5%, and another plan has success probability at 80% with a variance of 7%, we would of course choose the second plan, even though it is less robust than the first. Robustness is an important property, but it is conceptually not a criterion in then same sense as those that define the mission.

In the example in the paper, at least we knew what the parameters were about which we were making assumptions. How about missing knowledge? For example, what if there might be ways in which the enemy might foil our intentions but ways that we can't hypothesize at the moment? How to respond to missing knowledge in the models will be discussed in future papers.

Finally, we are not the first researchers to propose or discuss the issue of robustness of plans – [Bankes] is just an example of modern references that discuss approaches to this aspect of decision-making. Sensitivity analysis has long been an important topic in decision analysis. Economists in particular have discussed robust policies, since uncertainty about model knowledge is pervasive in that discipline. Our paper is a response to three observations. First, that the gap between models and reality is fundamental and is not restricted to cases clearly known to be full of uncertainty is not widely appreciated, shown by the fact that an overwhelming fraction of the literature in simulation and decision support still focuses on optimality, typically based on assumed probability distributions and expected values. There is usually little appreciation of the issue of missing knowledge or what has been called deep uncertainty. (Admittedly, our illustrative example doesn't deal with this issue.) Our goal is to add this view of robustness to a general framework on decision support as indicated in Fig. 1, so that *all* the issues that should be part of such a framework are identified and supported by underlying technologies that work together. For example, the fact that we can use the visual interface of the Viewer to perform some of this robustness analysis is important, since the Viewer is also part of the multi-criterial technology for selecting decision alternatives by applying trade-off judgments.

Second, we wish to raise awareness that robustness investigation is not simply a technical data analysis, but is motivated by the larger context of available actions, including various information-gathering and plan modification heuristics. We don't simply ask what assumptions the plan is most or least robust with respect to; the robustness analysis is partly driven by considerations of what we can do something about. Thus, if it turns out that the plan is most sensitive to assumption $A1$, but we have no means of verifying it, we then turn to exploring if there are any other assumptions we can do something about and which might moderate our uncertainty about $A1$. Another strategy that we mention in the earlier part of the paper is that of modifying plans so that the initial part of the plan is partly devoted to obtaining additional information about the assumptions that are uncertain and the plan is sensitive to. Proposing that we view robust decision-making not as a technical decision analysis issue but as embedded in a larger framework of action is another contribution of this paper.

ACKNOWLEDGMENT

This paper was prepared through participation in the Advanced Decision Architectures Collaborative Technology Alliance sponsored by the U.S. Army Research Laboratory under Cooperative Agreement DAAD19-01-2-0009, and by federal flow-through by the Department of Defense under contract FA8652-03-3-0005 (as a subcontract from Wright State University and Wright Brothers Institute). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U. S. Government. We thank John Josephson for useful discussions.

REFERENCES

[Bankes] Bankes, S.C., Tools and techniques for developing policies for complex and uncertain systems. *PNAS*, 2002. 99 (suppl. 3): p. 7263-7266.

[Carroll] Mark Carroll, John R. Josephson and James Russell, "Tradeoffs on the Efficient Frontier of Network Disruption Attacks," to be presented at this conference.

[Davis] Davis, P.K., Effects-Based Operations (EBO): A Grand Challenge for the Analytical Community. 2001, Santa Monica, CA: RAND Corporation

[Iyer] Iyer, N.S., A Family of Dominance Filters for Multiple Criteria Decision Making: Choosing the Right Filter for the Decision Situation, Ph. D thesis in Computer & Information Science. 2001, The Ohio State University: Columbus, OH. p. 169 + iv.

[Josephson] John R. Josephson, B. Chandrasekaran, Mark Carroll, Naresh Iyer, Bryon Wasacz, Giorgio Rizzoni, Qingyuan Li, David A. Erb, "An Architecture for Exploring Large Design Spaces," Proc. of the National Conf on AI (AAAI-98), AAAI Press/The MIT Press, pp. 143-150.