

AFRL-IF-RS-TR-2007-84
Final Technical Report
March 2007



RISK ADVISED COURSE OF ACTION (COA) ANALYSIS

Synergia, LLC

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

STINFO COPY

**AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK**

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report was cleared for public release by the Air Force Research Laboratory Rome Research Site Public Affairs Office and is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (<http://www.dtic.mil>).

AFRL-IF-RS-TR-2007-84 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE DIRECTOR:

/s/

DAWN A. TREVISANI
Work Unit Manager

/s/

JAMES W. CUSACK, Chief
Information Systems Division
Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE				<i>Form Approved</i> OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Service, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.					
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) MAR 2007		2. REPORT TYPE Final		3. DATES COVERED (From - To) Aug 06 – Feb 07	
4. TITLE AND SUBTITLE RISK ADVISED COURSE OF ACTION (COA) ANALYSIS				5a. CONTRACT NUMBER FA8750-06-C-0196	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER 62702F	
6. AUTHOR(S) Gregg Courand and Wanderley Liu				5d. PROJECT NUMBER 459S	
				5e. TASK NUMBER N6	
				5f. WORK UNIT NUMBER 04	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Synergia, LLC 702 Marshall St Redwood City CA 94063				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFRL/IFSB 525 Brooks Rd Rome NY 13441-4505				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER AFRL-IF-RS-TR-2007-84	
12. DISTRIBUTION AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED. PA# 07-124					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The purpose of this project is to demonstrate the feasibility and value of Synergia's experimental "Risk-Advised Planning" (RAP) technology. This technology embodies an analytical tool called "Value of Informed Control over Events" (VOICE). VOICE manages simulation-based analysis for COA development and assessment. RAP integrates human systems modeling tools to develop causal conditions and trajectories for projected COA execution, with VOICE decision-theoretic tools to assess and comparatively evaluate candidate COAs. The results of VOICE guide COA design/redesign, and circumscribe whatever information gathering and human systems modeling may be required. The value of the experimental VOICE technology was successfully demonstrated as a formal mechanism for managing the potentially unbounded combinatorial complexity of simulation-based planning. VOICE clearly identifies the most important causal variables in a COA. It sharply circumscribes the simulation studies that are required to establish the kinds and frequency of outcomes pertinent to COA evaluation and redesign.					
15. SUBJECT TERMS Risk-Advised Planning (RAP), Risk-Advised Course of Action (COA) Analysis (RACA), COA generation analysis, Value of Informed Control over Events (VOICE), normative hierarchical planning					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UL	18. NUMBER OF PAGES 46	19a. NAME OF RESPONSIBLE PERSON Dawn Trevisani
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code)

Table of Contents

LIST OF FIGURES.....	ii
INTRODUCTION TO RISK-ADVISED COURSE-OF-ACTION ANALYSIS	1
EXECUTIVE SUMMARY	1
THE PLANNING PROBLEM	2
THE JAKARTA SCENARIO.....	2
<i>Phase I – Air Battle</i>	<i>4</i>
<i>Phase II – Clearing Paths after Troops Landing.....</i>	<i>5</i>
<i>Phase III – Securing Jakarta International Airport</i>	<i>7</i>
<i>Phase IV – Securing Presidential Palace</i>	<i>8</i>
DISCUSSION	9
<i>Planning Complexities</i>	<i>9</i>
<i>Multi-Stage Planning</i>	<i>10</i>
<i>Simulation Studies</i>	<i>10</i>
TECHNOLOGY.....	11
COA DEVELOPMENT	12
THE VOICE TECHNOLOGY	14
THE HIERARCHICAL RISK ADVISED PLANNING.....	17
THE HRAP METHODOLOGY.....	18
ILLUSTRATION	20
SUMMARY OF DEVELOPMENTS	20
INITIAL COA MODELING	20
<i>Master-COA Model.....</i>	<i>21</i>
<i>COA Model Refinements.....</i>	<i>22</i>
<i>COA Modeling and JSAF Simulations</i>	<i>22</i>
<i>Synthesizing Sub-COA Models</i>	<i>24</i>
<i>Note on Value Assessment</i>	<i>25</i>
VOICE ANALYSIS	25
COA IMPROVEMENT	28
<i>Action Redesign.....</i>	<i>28</i>
<i>COA Model Revision.....</i>	<i>29</i>
<i>Next Steps</i>	<i>30</i>
<i>Remarks about Complexities</i>	<i>31</i>
CONCLUSIONS	32
REFERENCES	33
APPENDIX A: OVERVIEW OF THE RISK-ADVISED PLANNING METHODOLOGY.....	34
CO-FORMATION OF THE ANALYSIS COA AND THE GUIDING CHOICE.....	34
EVALUATION BASED ON THE GUIDING CHOICE	37
SUMMARIZING RISK-ADVISED PLANNING	40

List of Figures

<i>Figure 1: The Jakarta Scenario</i>	3
<i>Figure 2: Jakarta Operations Phase I – Air Battle</i>	5
<i>Figure 3: Jakarta Operations Phase II – Clearing Paths after Troops Landing</i>	6
<i>Figure 4: Jakarta Operations Phase III – Securing Jakarta International Airport</i>	7
<i>Figure 5: Jakarta Operations Phase IV – Securing the Presidential Palace</i>	8
<i>Figure 6: Sample COA Model (Phase II)</i>	13
<i>Figure 7: Sample Actor Model (Indonesian Air Force)</i>	14
<i>Figure 8: Sample V of C Graph</i>	16
<i>Figure 9: The Hierarchical Risk-Advised Planning</i>	17
<i>Figure 10: H-RAP Methodology</i>	19
<i>Figure 11: Summary of Developments</i>	20
<i>Figure 12: Master-COA Model</i>	21
<i>Figure 13: Sub-COA Modeling (Master-COA Refinement)</i>	22
<i>Figure 14: COA Modeling and JSAF Simulation</i>	23
<i>Figure 15: Synthesizing Sub-COA Models</i>	24
<i>Figure 16: VOICE Analysis Results (Master-COA)</i>	26
<i>Figure 17: VOICE Graph for Phase II</i>	27
<i>Figure 18: VOICE Analysis Results (Phase II)</i>	28
<i>Figure 19: Revised Choice Model for Phase II</i>	29
<i>Figure 20: COA Evaluation Results</i>	30

Introduction to Risk-Advised Course-of-Action Analysis

Executive Summary

The purpose of the RACA project is to demonstrate the feasibility and value of Synergia's experimental "Risk-Advised Planning" technology. This technology embodies an analytical tool we call VOICE, for "Value Of Informed Control over Events." We use VOICE in this project to manage simulation-based analysis for COA development and assessment. Risk-Advised Planning (RAP) integrates human systems modeling tools, to develop causal conditions and trajectories for projected COA execution, with VOICE decision-theoretic tools, to assess and comparatively evaluate candidate COAs. VOICE results guide COA design/redesign, and circumscribe whatever information gathering and human-systems modeling may be required.

VOICE enables planners to identify the most critical elements for planning and re-planning—i.e. those variables in the choice model (evaluating a COA) that most repay improved control. VOICE quantifies risks and opportunities in the neighborhood of current best estimates for the effects of COA execution. Analysis of VOICE results enables users to identify:

1. the most critical information to collect because it is most diagnostic of the best COA or COA-element to choose, among alternatives;
2. the actions most worth redesigning because small increments in control imply large gains or because small losses of control (e.g. underestimate adversary capability) imply large value losses;
3. possible gaps in value models, by identifying significant outcomes and trajectories for which no value feature (dimension of cost or benefit) has yet been defined.

We have successfully demonstrated the value of the experimental VOICE technology, as a formal mechanism for managing the potentially unbounded combinatorial complexity of simulation-based planning. VOICE clearly identifies the most important causal variables in a COA. In this way it sharply circumscribes the simulation studies that are required to establish the kinds and frequency of outcomes pertinent to COA evaluation and redesign.

Another major accomplishment of the project is the invention of a ***normative hierarchical planning*** approach. We refined the method underlying the RAP technology to enable a

hierarchical approach that we call H-RAP (Hierarchical Risk-Advised Planning.) Hence, under the guidance of our technology and method, planners are able to tackle a complex planning problem, partition it in chunks, and refine these chunks as needed.

We apply our hierarchical planning approach on the domain of AFRL’s Joint Synthetic Battlespace and demonstrate its value to manage the scope and complexity of planning-driven simulation studies. We illustrate its operation on a simulation scenario—the Jakarta scenario—that has been developed under the Joint Semi-Automated Forces (JSAF) environment.

We organize this document as follows. First, we describe our revised version of the Jakarta scenario and use it as the background to illustrate the complexities associated with the development of plans and the design of simulation studies. Then, we summarize the technology and methods employed in this demonstration, including the VOICE technology and the H-RAP method. Next we present the demonstration itself; we describe each step of our hierarchical planning approach as we apply it on the Jakarta scenario. Finally, we present a detailed discussion of Risk-Advised Planning in Appendix A.

The Planning Problem

The Jakarta Scenario

The Jakarta scenario consists of a series of events following the assassination of the Indonesian president and key government leaders that ultimately led to a military intervention of the Australian forces (blue forces) on Indonesia (red forces.) In the aftermath of the assassinations, an Indonesian general—himself aligned with a known terrorist separatist group that wishes to create an Islamic Indonesian state—seizes control of the Indonesian government. Because of its geographic proximity and its history of political tensions with Indonesia, Australia decides to send troops to Indonesia to restore order in the region. The planning problem is, thus, how to develop the proper COA for the Australian forces so that they succeed in their intervention.

Figure 1 provides an overview of the planning problem for the Jakarta Operations. The figure shows a map of Indonesia with the position of the red troops as they occupy three major strategic points for the operations: the landing area for the Australian forces, the Jakarta (Bandar Idara) International Airport, and the presidential palace.




Figure 1: The Jakarta Scenario

The figure suggests some critical issues that planners may need to address while elaborating the COA. Planners need to plan the proper allocation of resources (equipment, type and number of units) for the whole operation and the distribution of these resources among the various strategic points. Planners also need to elaborate the plans considering what they hypothesize about the adversary—its strategy, (post-coup) status, and skills. Here are some examples of how the (insufficient) knowledge about the adversary may affect the outcome of the operation. The causal import of each and the quality of our knowledge of each combine within formal COA models to reveal control requirements (leading to simulation and replanning):

- Planners are unsure about the strategy of the adversary. If the rebel government expects an invading force, it may take some measures (e.g. plant minefields) to protect potential target areas, such as the presidential palace.
- Planners are unsure about the actual amount of resistance that the blue forces may encounter. In particular, planners are unsure about the commitment of the Indonesian troops to the new regime and whether the new regime has induced the departure or the adhesion of members to the armed forces.
- Similarly, planners are unsure about the effects of the coup on the competence level of the troops. The departure or the joining of new members may shift the average

competence level of whole units; however, planners are unsure whether this shift actually occurs and, if it did, to which direction.

Planners may deal with all these uncertainties throughout the whole military campaign, although they may occur with more or less prominence depending on the particular battle front. For example, one may expect that the units assigned to protect the presidential palace are likely to be composed of the most skilled and committed troops. However, the level of resistance outside the most strategic areas may be more subject to the effects aforementioned. In order to capture these local variations of the issues aforementioned, and also to make the COA development more tractable, we divide the planning of the Jakarta operation into phases. We describe these phases as follows.

Phase I – Air Battle

The first phase of the Jakarta Operation (Figure 2) is the air combat where blue and red forces battle for the control of the Indonesian air space. Most military interventions begin with air strikes because whoever controls the air space has a substantial advantage over the adversary in the ground operations.

The Australian military superiority manifests clearly in Phase I. Australia has, in general, better equipment and better aircraft—Australia’s inventory contains several squads of F-18 Hornets—than Indonesia. Australia has also better trained and more competent pilots who have accumulated substantial experience from their missions in Afghanistan. On the other hand, although with a less impressive inventory that includes dozen of F-16 fighters, and with less competent pilots on average, Indonesia has the strength of numbers in its favor. Indonesia has several air bases spread throughout the country that enables it to send timely reinforcements as needed, an advantage that the Australian forces don’t share because their flights need to originate from Australia mainland—Australia does not currently have an aircraft carrier. Thus, for the first phase of the operation, planners need to develop the COA accounting for the uncertainty about the actual sizes of both the blue and red forces—they need to draft their plans considering various hypotheses about both unit sizes. In addition, for the reasons discussed before, planners also need to consider the impact of the levels of competence for both blue and red forces on the success of the COA.




Figure 2: Jakarta Operations Phase I – Air Battle

Phase II – Clearing Paths after Troops Landing

The second phase of the Jakarta Operation (Figure 3) initiates after the air combat and after the landing of the Australian forces on an inhabited area along the Jakarta seashore¹. The objective of Phase II is to clear the pathways for the ingress of the ground forces into the Indonesian territory so that they can successfully reach their engagement targets. In Phase II, the landing forces are split into two targets: the Jakarta International Airport (Westbound) and the presidential palace (Eastbound.) Planners expect the red forces to engage the blue forces at any point along the pathways to their (the blue forces') final destinations.

Like its air force, Australia also exhibits a substantial military superiority of its ground forces over Indonesia: Australia has better equipment and weapons—including several M1A1 tanks—and also more competent personnel. Indonesia still has the strength of numbers in its favor and the ability to send timely reinforcements, if needed. Hence, most of the issues raised in Phase I—uncertainties about unit size and levels of competence—are also valid for Phase II.

¹ Presumably the blue forces won the naval battle and successfully secured the Jakarta port. We didn't simulate this scene because the ship behaviors in JSAF are very little automated.

An important issue of the planning for Phase II is to understand the causal relationships of Phase II with the other phases. On the one hand, planners need to consider the impact of preceding phases on Phase II outcomes. In the Jakarta scenario, the outcomes of Phase I (i.e. whoever controls the Indonesian air space) define pre-conditions that causally affect the outcomes of Phase II. Hence, Phase II is less likely to yield a favorable outcome if the Indonesian forces maintain their control of the air space (and vice-versa.)

On the other hand, planners need also to attend the impact of Phase II outcomes to subsequent phases. Phase II is the initial point for the ground operations; thus, choices made at this phase may potentially promote or preclude the success of the subsequent phases. For example, an excessive allocation of ground units to the West front battle (the airport-bound pathway) may increase the chances of having more troops reaching the airport but may preclude a critical contingency of troops to reach the presidential palace for its take-over—and may ultimately risk a mission failure.




Figure 3: Jakarta Operations Phase II – Clearing Paths after Troops Landing

Phase III – Securing Jakarta International Airport

In Phase III, the objective is to take over and secure the Jakarta International Airport so it can serve as the operation base for the Australian forces. Having the airport secured has a significant strategic value for several reasons. First, it enables a continuous flow of supplies and resources for the ground troops, including the arrival of reinforcements. Second, it also enables the continuity of the air operations because fighters have now runways for take off and landing, and a base for inland refueling. The latter point is of significant strategic value because Australia currently does not possess sea carriers and, thus, its fighters need to take off—and perhaps refuel—on Australia mainland.




Figure 4: Jakarta Operations Phase III – Securing Jakarta International Airport

Planners need to develop the Phase III COA mostly considering the outcomes of Phase II (i.e. the size of the blue forces reaching the airport) and the adversary strategy and status (i.e. the size and competence level of the red units.)

Phase IV – Securing Presidential Palace

The objective of Phase IV is to secure the Presidential palace, remove the rebel government from power, and restore the control of the Indonesian government to the rightful authorities. Phase IV is the last phase and its goal is the ultimate goal of the whole Jakarta operation.

Like the preceding phases, Phase IV planners need also to consider the impact of both relevant outcomes from preceding phases (i.e. the number of blue forces reaching the palace) and other factors more local to the phase (e.g. red unit size and competence levels.) One may assume, however, that the Presidential palace is well-secured and guarded with the most skilled and committed combatants. Thus, planners may hypothesize that the red unit size is less likely to be reduced and red competence level is likely to remain high. In addition, planners may also hypothesize that the red units may introduce some safe-guards against potential attempts to take over the palace, such as the placement of minefields surrounding the palace area. The introduction of minefields adds another layer of complexity to problem, as the outcome of Phase IV now also depends on the outcome of Phase III because one solution is to have aircraft clear a path for the ground units through the minefields through bombing.




Figure 5: Jakarta Operations Phase IV – Securing the Presidential Palace

Discussion

Planning Complexities

The preceding description of the Jakarta scenario illustrates several issues of complexity involved in the elaboration of successful plans for the operation. Complexities arise, in general, from the planners' need to understand the planning problem thoroughly, and to deal with it effectively. Building a sufficient understanding of the problem is a non trivial task because of various uncertainties that may arise during the planning process. We pointed out in the preceding section that planners need to account for uncertainties about the adversary, including its defensive strategies and the status of its military—e.g. the size, the competence, and the cognitive state of its members.

Likewise, dealing with the planning situation, once we understand it, is also non-trivial because of the (potentially) unbounded number of possibilities for action and the choice situations that these may entail. In addition, each new action is likely to bring new factors and uncertainties that require further analytical efforts for the planners. For example, in the Jakarta scenario, planners face several hard decisions with respect to resource allocation (e.g. how many, which types, and where to position troops), for both air and ground operations and for each phase of the operation. Planners may also attempt to develop alternatives to deal with the size of the red forces (e.g. engage in an information operation to promote the desertion of the less committed members of the military) and understand their implications (e.g. one may hypothesize that the information operation is likely to have less impact on the forces protecting the presidential palace than on the forces guarding the airport) and side-effects (e.g. the information operations may increase the average competence level of the troops if it induces the departure of the least skillful members of the red forces.)

The issue aforementioned becomes more salient if we place the planning problem into the context of a more realistic resource and time constrained environment. Complexity ultimately arises from the need to develop and improve the COAs so they can actually succeed *given the amount of resources and time available*. In an ideal world, planners could simply consider all the possible issues for analysis and develop all action possibilities so they can produce the best possible COA. However, in the real world, resources are bounded and planners are often

required to devise a successful plan under severe time constraints. Under these conditions, the planning activity becomes more sensitive to the problem of unbounded numbers and the need to address this complexity problem becomes substantially more critical.

Multi-Stage Planning

The issues of complexity discussed above become even more critical in the case of the Jakarta scenario because we are dealing with an instance of a multi-stage (multi-phase) planning. The complexities involved in multi-stage planning are substantially larger for several reasons. First, formulating a multi-stage planning problem is harder because it inherently entails a larger plan (e.g. each phase of the Jakarta scenario is a full planning problem.) Second, planners need to understand the compounding complexity that arises from the causal interdependencies among phases. In multi-stage operations, *the outcomes of a phase often define the starting conditions for—thus, affecting the outcome of—subsequent phases*. For example, in the Jakarta scenario, the outcomes of Phase I (i.e. who wins the air combat) defines the pre-conditions for Phase II (i.e. the beginning of the ground operations.) Hence, planners need to develop the COA for Phase II considering two possible scenarios: (a) the blue forces win the air battle, and (b) the red forces win the air battle.

Finally, the complexities arising from the unbounded numbers for action possibilities and choice situations are substantially magnified in multi-stage planning. To illustrate this, suppose that each phase of the operation contains one choice with two alternatives, and three uncertainties, each with two possibilities. Hence, we have four variables, and because the operation has four phases, planners need to consider $65,536 (=2^{16})$ possible combinations for analysis. Adding another factor (with two possibilities) doubles the number combinations to 131,072. In other words, the number of possible combinations grows exponentially as planners add more issues to their COAs. Given the context of bounded time and resources, the problem of unbounded numbers in multi-stage planning may impair substantially the planners' ability to produce thorough analysis and high quality plans.

Simulation Studies

Simulations, in the context of the RACA project, are planning-driven; that is, simulations are instruments for planners to produce projections that they need for the COA development. In

general, planners run simulation experiments to qualify and, if pertinent, generate quantitative assessments for the COAs under given conditions. In the Jakarta scenario, planners may design JSAF simulations to assess the success of a COA—e.g. Phase II COA—under the various possible scenarios for that COA (e.g. under the various settings for red unit size, red unit competence, and air space control.) Simulations are, thus, accessories to COA development and the question of what one should simulate is ultimately a question of what COA one should develop.

The issues of complexities discussed in the preceding section have important implications for the management of simulation studies. First, the number of simulations needed for a given COA is a function of the complexity of that COA. COAs that contain more factors—actions and uncertainties—that affect their outcomes are likely to require more simulation studies. Second, given the exponential growth effect discussed before, the number of required simulations grows exponentially as the complexity of the COA increases. Adding one additional factor with N possibilities to the developing COA (e.g. one additional layer of final outcomes) increases the number of required simulations by N -fold. Adding R factors, each with S possibilities, increases the number of required simulations by R^S . Third, actor models contain random variables, as does the simulation control discipline, so that each simulation scenario needs to be repeated a certain amount of times for the purpose of producing statistically significant results. The higher the number of simulation runs, the more significant the results are.

The implications listed above illustrate the problem of unbounded numbers in the design of planning-drive simulation studies. Given these implications and considering the issues of complexity on both multi-stage and resource-and-time-bounded planning, the “ideal-world” strategy of developing all possible COAs and select the best of them is prohibitively expensive and, most of time, unfeasible.

Technology

Synergia’s method for assisting planning is Risk-Advised Planning (RAP). RAP assists planners in various facets of COA development. Our technology to assist planners carrying out RAP includes tools that facilitate the management and interpretation of primary data, the development of models of humans and collectivities, and the formation and evaluation of plans. (See Appendix A for a detailed review of the risk-advised planning methodology.)

During the development of the RAP technology, we invented a normative method to guide COA development in which planners iteratively revise and improve COAs so that they ultimately succeed. This normative method has evolved into what we now call the VOICE (Value Of Informed Control over Events) technology. VOICE addresses some of the planning complexity issues raised in the preceding sections as it identifies the most critical opportunities and risks worth the planners' attention so planners don't need to consider all possible factors and develop all possible alternatives.

In the course of the RACA project, we faced a new complexity challenge of having to deal with a multi-stage planning problem—i.e. the Jakarta scenario. In effect, every action and outcome in a multi-stage problem for all actions and events in all later stages—vastly increasing the complexity of the analysis and the scope of simulation that might, in principle, be required. In response to this challenge, we extended RAP and created a normative hierarchical planning approach that we are calling H-RAP (i.e. Hierarchical Risk Advised Planning.) We revised our VOICE-based, iterative refinement method for COA improvement so that it now supports hierarchical planning. In this section, we provide a brief description of the RAP technology with emphasis on the tools to assist COA development and on the VOICE technology. Then we describe the H-RAP approach and the H-RAP methodology for COA improvement.

COA Development

RAP technology contains a suite of tools that assists planners in various facets of COA modeling and evaluation. Our approach to COA modeling and evaluation is decision-theoretical based and consists of embedding COAs within a formal choice model that we call the *guiding choice* of the COA. This representation of COAs enables a proper, quantifiable characterization of the uncertainties, as well as precise models of costs and benefits. The choice model consists of a tree-like structure where we expand all possible ramifications of actions and events—i.e. alternatives and uncertainties, respectively. Actions are restricted to contemplated steps in blue COAs; adversary actions are unknowns depicted as uncertainties in the choice. Each path of the model represents a prospect of the COA—a possible combination of actions and events organized in a cause-effect flow of execution. Typically, a prospect describes an execution flow trajectory, from initiating conditions, through actions, to feasible outcomes of these actions. We

annotate events with their likelihoods (i.e. probabilities.) We also score each prospect and annotate it with its respective value.




Figure 6: Sample COA Model (Phase II)

Figure 6 illustrates a sample choice model for the Phase II of the Jakarta scenario. The COA itself is represented by a blue start. The parallelogram represents the action² and the ovals represent the uncertainties. We attach a diamond-shaped node to the end of each prospect to annotate its value.

The RAP technology also provides tools to support the construction of the choice model structures. For example, Figure 7 shows a screenshot of the Actor Model tool portraying the behavioral profile of the Indonesian Air Force. Creating a human systems model, especially models of the adversary, is critical for evaluating military plans because some of the greatest sources of the uncertainties in military operations are uncertainties about the enemy. In other words, the choice models used to evaluate COAs are summaries of cause-effect relations determined within actor models. Therefore, we use VOICE to identify key variables to control, and thus *derive* elements of actor models we need to develop, so that we understand the

² This choice model is a degenerative case where there is only one alternative for the guiding choice

causality. Simulation then works through the potential inter-dependencies created by the possibilities for practice interactions of the subject (blue-red) human system.

Note that we may start with subjective estimates for actor interactions, then use VOICE to determine which ones need to be refined through actor mapping and modeling. *This means that VOICE is actually helping not only to determine action redesigns, but also, by identifying phenomena that need to be controlled, identifying modeling requirements. In these cases, control is derivative of better understanding first, then redesign when that understanding is in place.*




Figure 7: Sample Actor Model (Indonesian Air Force)

The VOICE Technology

The core technology of our approach—the one that we employ to address some issues of planning complexity—is the VOICE technology. The VOICE technology assists planners to identify the most critical elements—the elements that induce the greatest gains of loss of value—for COA improvement. These elements include:

- The most critical information worth gathering
- The actions most worth of redesign and those events most worth of newly-introduced actions
- most prominent risks and opportunities in evaluating and improving COAs

- Causal loci for which clarification of value (features and quantities) may be most useful

VOICE technology addresses the complexity of unbounded numbers under bounded time and resources because it relieves the planners from having to consider all possible factors and develop all possible alternatives. It ensures that planners direct their attention and resources are geared only towards the identified critical elements as those provide the best ratio of value versus planning efforts.

The essential idea in VOICE technology is to recognize that probabilities in a choice model are judgments. They reflect the quality of our understanding of a domain—both the cause/effect relations, and the likelihood of any particular configuration of that event relations. In general, any given system of causes may be placed under more “control”—we may invest resources and action to better determine the actual state of affairs at the given moment. Naturally, here we aim to control to more favorable outcomes, and to place more probability on those outcomes—i.e. get more of what we want, with more confidence. Alternatively, our understanding may be shallow (e.g. the adversary may be better equipped, or more competent, than we imagine.) It is possible that the world will eventuate, on average, in a way that is worse than that expressed in the model. In this case we would wish to control events so to minimize the worse outcomes.

VOICE supports this analysis formally through calculations that successively manipulate the probability density functions for each variable in a choice model, and compute the resultant expected value of the optimal choice³. Then we produce a simple graphic that shows the successive improvement in value (control probability toward the best prospect emanating from the event space being controlled) and the successive loss in value (control probability to the worst prospect emanating from the event space being controlled). Figure 8 shows our graphical visualization of the VOICE results for the size of the red unit in Phase II. The right side of the graph shows that the best opportunities in terms of red unit size are found in actions that aim to reduce the size of the red forces. Conversely, the left side shows the greatest risks occur when the adversary takes action to improve their chances of having their troops reinforced.

³ A useful property of VOICE is that it doesn't require a fully-specified choice in a COA model. It is possible to perform the analysis entirely with a collection of prospective trajectories of events, provided they are annotated with a value model.



Figure 8: Sample V of C Graph

The interpretation for planning is that each point in the VOICE curves accords with a set of possible worlds in which the biased probability density has been brought about—e.g., by redesign of the friendly COA. VOICE offers several diagnostics that enable the analyst to interpret these results and identify redesign requirements—e.g., a table shows which dimensions of value most contribute to the best prospects, and which most contribute to the worst outcomes, so that COAs can be redesigned to better carry favorable value features, and better ward off conditions that are laden with the costly value features. In addition, when there are several COA alternatives, the VOICE results can reveal landmarks—inflection points where value is substantially gained, or lost, if probability happens to be shifted to those landmarks.

VOICE is flexible about control, offering three main alternatives. It is possible to control a variable outright, to control it contingent on other variables having eventuated in a given way, and to control multiple variables.

Finally, each point on a VOICE result graph indicates a degree of control. For any such point, provided there are at least 2 action alternatives, the VOICE calculations additionally produce the quantity that information on the variable of interest would be worth. A formal (and

common sense) result is that information has zero value at the points of perfect control, or loss of control.

The Hierarchical Risk Advised Planning

A major accomplishment of the project is the development of a normative hierarchical approach that we call the Hierarchical Risk-Advised Planning (H-RAP.) We refined the methods and concepts underlying the RAP technology so that they enable the partition of a complex planning project into more manageable chunks and apply our tools and method to refine the chunks. The H-RAP directly addresses the complexities of multi-stage planning. The hierarchical method controls the exponential escalation of complexity because it enables planners to work on smaller, less complex chunks at a time instead of working on the whole planning project.




Figure 9: The Hierarchical Risk-Advised Planning

Figure 9 illustrates the H-RAP approach. The H-RAP defines two levels of scoping for COA modeling: a high level master COA model and lower level sub-COA (phase) models. The master COA model, in general, provides an overview of the overall plan of a given operation by emphasizing its major components (i.e. the sub-COAs) and the causal relationships among these components. The sub-COA model refines and expands the master COA model for a given phase and introduces critical details that are (and should be) omitted from the master COA model.

We change scopes in H-RAP through two operations:

- Top-down model refinement

- Bottom-up model summarizing (or synthesis)

The conversion from master COA to sub-COA is a top-down refinement. In this operation, the master COA “passes” only information that is relevant to the sub-COA, such as

- Pre-conditions or scenarios. E.g. the master COA specifies to Phase II COA that it occurs under two scenarios: blue forces control air space and red forces control air space
- Specification for outcomes relevant to other phases. E.g. the master COA specifies to Phase II COA that its needs to include two outcomes: the size of troops moving towards the airport and towards the palace
- Relevant quantities (e.g. probabilities for pre-conditions and values for outcomes)

The conversion from sub-COAs to master COA is a process of model synthesis. The sub-COA introduces model details that are irrelevant at the master COA level. For example, the sub-COA contains figures for outcome probabilities and values (e.g. cost from casualties) assessed under various phase conditions, such as red unit size and competence levels. Because these details are irrelevant at the master COA level, these figures need to be summarized before they are passed to a higher level of abstraction.

We developed the conversion algorithms in the course of the RACA project. These algorithms are mathematically founded on decision theory and Bayesian Statistics and uses basic operations from these disciplines (e.g. decision tree roll back and Bayesian inference.)

The HRAP Methodology

The RAP Methodology (Figure 10) consists of three major steps:

4. Produce the Initial Choice Models for the COAs
5. Perform VOICE Analyses on the Models
6. Improve the COAs and Revise the Models, based on the results of the VOICE Analyses
7. Re-start the cycle (go back to 1.)




Figure 10: H-RAP Methodology

The method begins with the development of the choice models for the COAs at both levels of scope: the master COA model and the sub-COA models. The RAP technology provides several tools to assist the development of each individual COA. Here, simulations have the important role of supporting the quantitative assessments needed for COA evaluation: the COAs defines the simulation experiments (e.g. JSAP “scenarios”) by specifying the set of required projections. The simulations, in turn, feed the projected quantities (e.g. probability and value assessments) back to the COAs.

The second step is to perform VOICE analyses on the model to determine the critical factors—the most promising opportunities and the greatest risks—that are worth our attention. We perform VOICE analysis on the master-COA to determine which of the sub-COAs is the most worth improving. Then we perform VOICE analysis on the model for the critical sub-COA and identify its critical factors for its improvement.

Based on the results of the VOICE analyses—i.e. the critical sub-COA and the critical factors for that sub-COA—we initiate the improvement of the COA. We improve the COA by redesigning the actions—revising the current actions or inventing new ones—based on the identified critical factors of the critical sub-COA. Finally, we revise the choice models and

evaluate them, so we can have a proper assessment of the expected value gains (if any) from the modifications.

In the next section, we provide an illustration of the method based on the Jakarta scenario.

Illustration

Summary of Developments

We provide an illustration of the H-RAP methodology. Figure 11 shows a summary of our developments as we apply one pass of the methodology on the Jakarta Scenario.




Figure 11: Summary of Developments

Initial COA Modeling

We begin with the development of proper choice models for the Jakarta scenario, including the master-COA model and the four refined sub-COA models.

Master-COA Model

At the master-COA level, we limit the COA representation so that only the most important actions and events—those most relevant to understand the causal relationships among phases—are included in the model.




Figure 12: Master-COA Model

Figure 12 shows an excerpt of the master-COA model for the Jakarta scenario representing the two initial phases of the operation. Phase I and Phase II are expressed only in terms of their respective summarized actions and the consequences of these actions that are relevant for the subsequent phases. This excerpt of the model informs us that planners need to develop and analyze Phase II model under these two scenarios: (a) the blue forces cleared the Indonesian air space in Phase I and (b) they didn't.

Because this is the first pass of the method, we lack some quantitative figures—probabilities and values—to complete the model. These figures are better assessed at the sub-COA level where the more detailed treatment ensures a more accurate assessment for these figures.

COA Model Refinements

We proceed by refining the master-COA model into more detailed sub-COA models, one for each phase of the operation. At the sub-COA level, we expand the phase model so as to include a more detailed chain of causal influences on the outcomes of the phase. These causal include conditions set by preceding phases and conditions that are more local to the phase (e.g. size and competence level of the red troops.) Figure 13 shows excerpts of the refined sub-COA models for the Jakarta scenario.




Figure 13: Sub-COA Modeling (Master-COA Refinement)

Once the structure of the choice model for the COA is complete, we shift our attention to the production of the proper quantitative assessments. These assessments may come from several sources, including the collection of statistical data, (the fusion of) experts' opinions, and projections through simulations. We elaborate the latter point in the next section.

COA Modeling and JSAF Simulations

As mentioned before, simulations have the role of providing projections needed for the quantitative assessments of event likelihoods in a choice model. Figure 14 shows a more detailed illustration of the relationship between the COA model and simulation studies. The COA model

defines a set of required quantitative projections that we can translate directly into a set of requirements for simulation experiments. Once we set the requirements, we design and run the simulation experiments, collect the results, and translate the results into projected assessments—probability distributions and/or value figure—for the phase outcomes. We then feed these projections into the COA model for later evaluation.

In the figure, we design simulation experiments for Phase I so that the results inform us the probabilities of winning or losing the air combat and the number of blue casualties for cost computation. In general, for each phase of the operation, we define one simulation experiment (i.e. a JSAF scenario) per COA prospect. Thus, for Phase I, we designed 54 simulation experiments considering 54 prospects: 3 sizes for red troops, 3 levels for red competence, 3 levels for blue competence, and 2 sizes for blue troops.⁴




Figure 14: COA Modeling and JSAF Simulation

In general, we should perform several simulation runs per simulation experiment for statistical significance. For the Jakarta scenario, we typically performed 5 to 10 runs per experiment. We were able to assess, for each simulation experiment, the mean and the standard

⁴ Planners may not know in advance how many units and which pilots will actually be available for the operation; hence, the information about blue unit size and competence may also be uncertainties

deviation for various simulation parameters (e.g. the number and the types of the remaining units for red and blue forces.) We used these parameters to build our needed projections (i.e. the probability distribution for the outcomes and the expected number of blue casualties.)

Synthesizing Sub-COA Models

The assessed quantities—either obtained through simulations or assessed by other means—are fed into each of the sub-COA models so now we can properly evaluate them. The next step is to update the master-COA model so that it reflects the quantities (or even structural changes, if pertinent) expressed in the sub-COA models.

Figure 15 illustrates this operation. In Phase I, we synthesize the 54 assessments for the win-lose probability distributions (one for each prospect) into one win-lose distribution that summarizes those 54 assessments. We also synthesize 54 assessments for casualty costs of a phase into one unique value for cost that embodies all the 54 assessments.




Figure 15: Synthesizing Sub-COA Models

Having both the structure and the quantitative assessments completed for all COA model, we complete the development of the initial COA and initiate the process of improving it.

Note on Value Assessment

For the value modeling in the Jakarta scenario, we estimated numerical values based on the costs of units and potential casualties for the blue forces. The inventory of the Australian forces consists of F-18 fighters, M1A1 tanks, and ground forces. We used the following figures to guide the assessments:

- The cost of an F-18 fighter (blue forces) is approximately U\$ 24 million
- The cost of a M1A1 tank is approximately U\$ 4.3 million
- The cost of a dead ground troop is approximately U\$ 0.25 million.

We estimate the total worth of the mission as being approximately U\$ 2 billion (i.e. about the price of 100 air fighters.) Regarding the value of a life of 250K, we have seen many reports that insurance underwriters and designers (building, vehicles etc.) use this assessment. It is used to assess whether the total cost of bringing an innovation to market (from concept and design all the way through to production) is less than the value of lives saved by the innovation⁵.

VOICE Analysis

We proceed by applying our core technology, VOICE analysis, to identify the critical elements of the COA—the critical risks to prevent and the highest-valued opportunities to capture—that are most worth our attention. Our hierarchical approach provides us the opportunity to apply the VOICE technology at various hierarchical levels. We can apply the VOICE technology at a higher level to determine which of the sub-COAs is the most critical. We can also apply the VOICE technology on a sub-COA to identify the critical factors for that sub-COA.

First, we perform the VOICE analysis on the master COA. Figure 16 shows the results of the analysis in tabular form.

⁵ In addition, our choice modeling tool enables the value estimates to be developed graphically, so that the user can anchor some values and then scale others visually until satisfied

View Table		
Event Space	Event	Value Shift
Phase I Air Combat	Blue Air	20.7
Phase I Air Combat	Red Air	-94.32
Phase II	W, E Cleared	141.05
Phase II	W, E Partially Cleared	-409.6
Phase III	Airport Secured	88.95
Phase III	Airport Not Secured	-198.0
Phase IV	Palace Secured	120.93
Phase IV	Palace Not Secured	-1894.5

Phase II is Most Critical

- Most Potential Gains
- 2nd Most Potential Losses

Figure 16: VOICE Analysis Results (Master-COA)

The results show that Phase II provides the highest-valued opportunity for improvements and also the second highest risk for losing value, when compared to the other phases. In Phase II, we can accrue a value gain of approximately U\$ 141 million if we can guarantee that blue forces prevails in both the East (palace-bound) and West (airport-bound) front battles. However, we can also lose up to U\$ 410 million in Phase II, if we allow the red forces to prevail in these fronts.

Figure 17 provides a more detailed figure of the risks and opportunities in Phase II. For example, the amount of probability increase that we need to produce the maximum gain of U\$141 million is 0.75 which is relatively high. Conversely, the amount of probability change to achieve the maximum loss is 0.9, which is also high. We can also observe the slopes of the curve and conclude that, in Phase II, we may lose value faster than we may gain. These analyses informs us that we a lot of room for improvement (or worsening) and we should seek actions with substantial impact on the success odds for Phase II so we can capitalize the most from the potential gains or avoid big surprises leading to potential losses.




Figure 17: VOICE Graph for Phase II

We identified Phase II as the critical phase so now we apply the VOICE analysis on the Phase II COA model. Figure 18 shows the table summarizing the results. For Phase II, the best opportunity for value gains is to ensure that a large blue unit size reaches the airport (i.e. that the blue forces win the battle on the West front.) The greatest risk is to have a reduced blue unit size on the palace pathway (i.e. the blue forces lose the battle on the East front.) We use these results to guide the redesign of the actions for Phase II to improve the COA.

View Table		
Event Space	Event	Value Shift
Airport Pathway Troop Size	Airport Path: Large B Troops	120.5
Red Size	Red Reduced Size	70.06
Palace Pathway Troop Size	Palace Path: Large B Troops	53.8
Red Competence	Red Competence: Low	40.01
Air Space Cleared	Air Space Cleared	20.7
Red Competence	Red Competence: High	-26.57
Airport Pathway Troop Size	Airport Path: Reduced B Troops	-38.05
Air Space Cleared	Air Space Not Cleared	-94.32
Red Size	Red Base Size	-70.06
Palace Pathway Troop Size	Palace Path: Reduced B Troops	-360.02

For Phase II

- **Best Opportunity:** Have Large Blue Unit Size on Airport Pathway (West)
- **Riskiest Outcome:** Have Reduced Blue Unit Size on Palace Pathway (East)

Figure 18: VOICE Analysis Results (Phase II)

COA Improvement

Action Redesign

The next step is to improve the COA of the critical phase by revising its constituent actions and providing alternatives in light of the results of VOICE analysis. Hence, we initiate the search for feasible alternatives with two criteria in mind:

- Criterion 1: maximize the likelihood of having a large blue unit size on the West front (best opportunity)
- Criterion 2: minimize the likelihood of having a reduced blue unit size on the East front (greatest risk)

Planners need to evaluate the alternative actions according to their feasibility and their ability to achieve both (preferably) or either of the criteria above. For example, planners may consider an increase of the total units that arrives on the Indonesian seashore so they could increase troop size on both fronts and satisfy both criteria. However, planners may conclude that this alternative is unfeasible, given perhaps the timing required to initiate the operation. We proposed changing the distribution of troops along the fronts so more troops are allocated westbound. This strategy passes the feasibility test and satisfy the first criterion but is opposite to

the second criterion. Hence, we explored an incremental shift of troops and defined two alternative actions:

- *Alternative 1*: Shift 10% of the units from the East to the West front
- *Alternative 2*: Shift 20% of the units from the East to the West front

Next, we revised the Phase II COA model by adding these two alternatives so they can be properly evaluated against the base case where no units are reallocated.

COA Model Revision

Figure 19 shows an excerpt of the revised choice model for Phase II COA when we add the two new alternatives to the base case (the base case is the uppermost alternative in the model.)

Structurally, we developed each of the alternatives following the templates of the base case; hence their outcomes are subject to the same causal factors present in the base case (i.e. the size and competence level of the red units, and the outcomes of the air battle.) For the quantitative assessments, we designed and ran new sets JSAF simulations to assess probabilities and costs, also following the template of the base case. We recorded the simulation results, inserted the quantitative assessments in the model, and evaluated the alternatives.




Figure 19: Revised Choice Model for Phase II

Figure 20 provides a graphical representation of the COA evaluation results for the base case and for the two new alternatives. The results show that *Alternative 1* is optimal: shifting

10% of the troops to the West front yields a net gain of approximately U\$ 30 million in total value; however, further shifting the troops produces diminishing returns with total gains of approximately U\$ 20 million. *Alternative 1* capitalizes on the opportunity for gains expressed by the first criterion without accruing substantial losses from the risks expressed in the second criterion. On the other hand, *Alternative 2* also capitalizes on the same opportunity for value gains but has the risk factor more pronounced—reducing 20% of the blue unit size eastbound risks a least favorable outcome on the East front which offset part of the accrued gains.




Figure 20: COA Evaluation Results

We complete the methodology cycle by accepting a new Phase II COA where the allocation of troops is slightly biased to the West front (*Alternative 1*) and by updating the master COA accordingly⁶.

Next Steps

Alternative 1 produces a net value gain of approximately U\$ 30 million out of the possible U\$ 120 million that we assessed through the VOICE analysis—i.e. it yields 25% of the potential

⁶ Planners may update the master COA model considering only the optimal alternative. They may also update the master COA model so it now a choice of three alternatives for Phase II. We opted for the former because the *Alternative 1* statistically dominates the other alternatives; that is, each possible prospect in *Alternative 1* has a higher score than its respective counterparts in the other alternatives

gains—suggesting that there is room for further improvements in the phase. Hence, planners may continue their search for additional revisions of their COAs for Phase II.

However, after revising the COAs, it may be possible that Phase II is no longer the most critical. Indeed, performing VOICE analysis on the revised master COA model shows that Phase IV is now the most critical and all the analysis and developments should now be centered on this phase. We conclude that a good discipline when following our methodology is to always perform a VOICE analysis on the master COA before initiating the process of improving COAs.

Remarks about Complexities

We conclude our demonstration with a few remarks about how our method deals with complexity. First, we address the exponential growth of complexity of multi-stage planning with our hierarchical method. The H-RAP decreases substantially the number of possibilities for assessments and analyses so that, instead of the 65,536 possibilities in a four-phase plan with 4 variables per phase and 2 outcomes per variable, planners need to deal with only 64 combinations ($4^2 = 16$ possibilities per phase, 4 phases).

Second, rather than exploring all four phases, we were able to concentrate on Phase II. And for Phase II, we only needed to be concerned with the relative distribution of troops east and west (i.e., allocation to the mission of securing the airport, versus the palace). This meant that, in contrast to the many simulation studies we might have run for Phase II (or the 65,536 for all four phases!), we just needed about 200 per each alternative.

The combination of the two previous points is simply that, even if we loop through the entire risk-advised planning method several times (e.g., we next focus on optimizing Phase IV), the amount of total developments and simulation runs would be substantially lower than the “ideal analysis” approach of elaborating all alternatives before evaluating and selecting them. Yet, our method ensures that the value gains from this development are among the highest possible – and can be developed within the amount of time and resources available.

Third, we offer this critique of what we have referred to as the “ideal analysis,” the analysis that *from the beginning* depicts all possibilities for all four Phases. The presumption is that the analyst can realistically envision all reasonable prospects and inter-dependencies. But our experience using VOICE reveals very strongly that it is the ability to focus on selected

elements of the causality that makes it feasible to invent/represent the new alternatives. Hence as a practical matter, we find that the critique, exploration and design activities promoted by VOICE are not just an approximation of an imagined totalizing analysis. Instead, they create the possibility of working with as much complexity as is *actually possible*, on the most important causal loci. Hence if we are mindful of human boundedness, VOICE analysis is actually a better approximation of the ideal than a totalizing form initially considered to be “ideal analysis.”

Conclusions

We have demonstrated the use of Synergia's VOICE technology to create a new form of hierarchical planning. Normative assessment guides iterative refinement of constituent elements in a multi-stage COA in which the outcomes of earlier stages determine the entry conditions for later stages.

We have demonstrated that we can manage the scope and complexity of simulation-based analysis, based on the results of the VOICE calculations. These calculations reveal critical risks to prevent, and high-value opportunities to capture. Hence these assessments determine those elements of a plan that most deserve redesign, as well as those elements of its execution conditions and consequences that most repay better understanding. These two factors drive data collection, actor modeling (self, adversaries), and the design of simulation studies. Simulation studies are used to project future conditions as a function of actor models and associated COAs, as well as exogenously-arising conditions. Simulation results update the COA assessment framework (a formal choice model), and a next round of VOICE analysis determines the elements of understanding and action most worth improving, leading to the next simulation studies worth developing. This methodology is entirely general; for the current work we are managing the use of the JSAF simulation environment.

References

The following sources were consulted and are referenced for this work.

1. Courand, Fehling, Grace, Liu: *Risk-Advised Planning: Final Project Report*. March 2, 2006.
2. Kussard: *JSAF Scenario Design, Jakarta Indonesia*. June 10, 2005.
3. Kussard: *Jakarta Scenario Summary*. October 30, 2005.
4. Kussard: *JSAF Scenario Design (AFRL Outbrief)*. October 30, 2005.
5. Synergia LLC: *ACCORD User Documentation*. August, 2006.
6. USJFCOM: *JSAF User's Guide, Volume 1 and 2*. November 12, 2002.


Appendix A: Overview of the Risk-Advised Planning Methodology

We summarize primary elements of the Risk-Advised Planning methodology here. This methodology places candidate COAs within a formal choice structure as a basis for producing quantitative estimates of the quality of the COA (or, comparative estimates among COAs).

Co-Formation of the Analysis COA and the Guiding Choice

Our approach is based on developing and elaborating a formal choice structure, called the “guiding choice.” The **guiding choice** is the basis for COA evaluation; it synthesizes the following ingredients: an analysis-COA; event spaces that determine the admissible causes and effects associated with the analysis-COA; probability densities over each event space; dimensions of worth as a function of types of activities and events; and quantification of worth measures. An **analysis-COA** is representation of a COA, produced by planning technology or processes, and suitable for choice modeling and analysis. In general, the analysis-COA depicts all key choices and points of uncertainty in the COA, all loci of value creation or loss, and points that may be sensitive to adversary initiatives and responses. Other elements of the COA are summarized as events that provide context for the analysis-COA.⁷

The guiding choice is not an assemblage of pieces developed by other means. Instead, the analysis-COA is an anchor around which the guiding choice (i.e., all other ingredient structure) is built. Generally, we will have the COA as a starting point for the analysis-COA, and



develop the guiding choice according to the analyses in the graphic shown here. A typical path through this process is as follows. The COA is summarized as an analysis-COA. For example, it

⁷ Note that formal choice models are directed *acyclic* graphs. This creates a modeling requirement, because some activities are conducted for as long as a condition is true, and others until a condition becomes true – a logic that would be captured by the underlying practices. As such, summarization of behavior circumscribes around loops/cycles generated by the underlying practices, or flattens the loops by injecting milestone events and marking them in time. Often, this requirement has already been met by the computational planning operators or human planning processes that created the COA itself.

might be just re-stated directly, with any elements not thought relevant to risk analysis left out, and with any merging thought acceptable put in place. Then the analyst has the option of developing the guiding choice, or of modeling the underlying human system. Often, the user will begin with the former⁸.

To develop the guiding choice at this point is to begin with the guiding values (the actual concerns and objectives that determined the COA), and convert them into a preliminary worth model. This will determine kinds of events that need to be included in the guiding choice. The user installs all event structure if it is previously known (some is given by the effects intended to be created by the COA).

As needed, the analyst gathers data and conducts human-systems modeling, and simulation as necessary, to establish or to validate the event spaces that trigger and influence action, and its products. Usually, human-systems modeling is necessary whenever the effects (intended to be) created depend sensitively on the state or activities of the adversary. This may advise the guiding choice directly, or, may lead to a reformulation of the analysis-COA, which in turn may advise the guiding choice.

We show a sample guiding choice below, to illustrate the role of human-systems modeling. (Please zoom on the soft copy to see the details.) We imagine that a U.S. and Afghan coalition seeks to capture a narco-terrorist actor; a primary concern is whether he will seek to cross the border into safe haven, making his capture much more difficult. In this case, the coalition has identified two COAs. One is the standard operating procedure, which announces the intent to capture this individual, publicly offers monetary incentives, and uses Afghan police forces to locate and apprehend the subject, with U.S. backup. The second is a covert strategy that makes no announcements, and only very judicious and private use of monetary incentives. In this case,


⁸ Modeling is never an end in itself – information is in service to action, and for this work, choice as a way of evaluating candidate action designs. Hence we favor forming the guiding choice as early as possible, since normative calculations can then be used to constrain the modeling and analysis activities required to develop and validate the event spaces necessary for evaluation. Of course, as modeling occurs, the choice changes, and the cycle may continue. This process continues so long as a) worth as per the improvement of the choice and ultimately the COA is greater than the cost of modeling and plan implementation, and b) time prior to execution of the COA permits the modeling and analysis activities. Informally, clarity is increasing as one cycles through this process. Formally, risks are being identified, quantified, and reduced, and opportunities are being identified, quantified, and exploited.

the analysis-COAs are simply the names of the two COAs, with probability judgments reflecting their relevant properties.

The probability of capture depends on the behavioral state of the subject – i.e., whether he is conducting business as usual, is hiding in the country, or is exiting to the border. In turn, the behavioral state of the subject depends on his cognitive state – his beliefs about whether he is being sought, and his estimate of threat to his freedom.

Friendly activities are causal for his cognitive state, which is causal for his behavioral state, which then is causal for whether friendly forces detect and capture⁹ him. Hence the causality that intervenes between the friendly COA and outcomes of the friendly COA traces through the mind and behavior of the subject.

Therefore, the analyst needs to model this subject's likely information (sensor), interpretation (perception of threat), and decision-making (behavior selection) practices, to establish the cause/effect relations in the choice, to comparatively assess the two COAs.



Note that the risk-advised methodology here has also identified the causal relations that are most important to understand, and so has circumscribed the actor modeling that is necessary.

We offer a few additional remarks on modeling the human system that is the context for plan implementation. This system is composed of the interdependent practices of: friendly forces that enact the COA, coalition partners, neutral parties, and adversaries. In general, an actor, and so that actor's practices, is relevant if and only if its behavior affects plan realization. It may do so by initiating activities (creating conditions), responding to friendly activities, or participating in joint activities. The particular practices that we model are those that generate or respond to


⁹ These latter two event spaces are not shown because they have been formally subsumed in the value nodes.

events *critical to the criteria of the subject actor*¹⁰. In short, we bootstrap from the events and behaviors implied by the analysis-COA, to all practices that are required to realize the analysis-COA, to the events they depend on and create. Some events will be deemed critical, in terms of the strength or prominence of the criteria in the practices. Then we determine which other actors' practices create, prevent, or respond to the critical conditions, and map these practices. This modeling process continues until the analyst establishes a system of practices that is properly coupled and in which the outermost initiating events are directly available (e.g., observable, known by the analyst, installed in a database, susceptible to subjective estimation, expressible as a distribution over some event space, etc.) Event and activity descriptions are elaborated until the analyst is satisfied that probability estimates can be formed for all event spaces. Then the elements of the analysis-COA and its surrounding event spaces are abstractions that summarize behaviors and events that may be generated by this human system.

Evaluation Based on the Guiding Choice

Risk-advised planning is a critical method, in that it provides for well-founded, validated critique and so cyclic development of all constituents in the analysis. Once the guiding choice is fully specified – the analyst is satisfied with the analysis-COA, event conditioning, probability assessments, and worth model – the desired evaluations are produced. These are the ones of direct interest to planners – e.g., the comparative expected worth of two alternative COAs – *as well as* normative assessments Synergia has invented to quantify the value of action redesign (minimize risks and maximize opportunities).

We show the risk-advised planning methodology here, completed to show the role of critical analysis of



¹⁰ In this way we ensure that we evaluate the COA in light of the most powerful intentions and capabilities of the adversary. By carefully elucidating the practices of our adversary, we guard against the tendency to test our plans with invalid models of the adversary. Analysts need to guard against models that are too optimistic (or pessimistic!) about the self/adversary opposition, as well as models that implicitly treat the adversary as if they behave like we do.

the guiding choice. The evaluations derived from the choice articulate critiques of the analysis-COA. These critiques stimulate refinement of event structures as determined by the human-systems models, the analysis-COA, worth modeling, and the guiding choice – i.e., they lead to revised evaluations and additional modeling and analysis. Finally, it is possible that a critical evaluation will lead to a recommendation to refine the COA itself.

Here is a summary of evaluations that are pertinent to COAs. All of them are produced in Synergia’s technology.

- The probability of success of the COA, which may be developed as the sum of the probability of all prospects that produce positive outcomes.
- The probability of highly negative outcomes, up to and including failure, for the COA.
- The distribution of worth that may be produced, considering whatever conditions the activities in the COA might face (at the time of execution), and might create. This is derived from the products of prospect-worth with prospect-probability.
- Sensitivity results per each COA or its associated policies – e.g., the degree to which the optimal COA is sensitive to judgments about the probability of one or another variable (depicted as an event space). This sensitivity is defined for a given event space (i.e., its general influence on the choice), and also for the event space *contingent on* some selected trajectory of events (i.e., the influence of the event space, but conditional on some preceding context).

The final evaluation is the basis for Synergia’s Value of Control metric. Value of Control quantifies the expected worth of improving control – increasing the probability of favored outcomes for some activity. This calculation also exhibits the threat posed by having less control than one supposes – e.g., due to bad luck (adversary luck), adversary skill, or their combination – which increase the probability of harmful outcomes.


Value of Control calculations are delivered through Synergia’s VOICE (“Value Of Informed Control over Events) technology. VOICE computes the worth¹¹ associated with

¹¹ The word “value” has many meanings. We use “worth” to signify “conditions of benefit or harm/cost” to a decision-maker. A worth model is a relation on: a set of dimensions or qualities of benefit or harm, a set of event

favorable or unfavorable perturbations of events (“Value of Control”). It also computes, for each perturbation, the value of optimizing information state (“Value of Information, for a given perturbation in Control”). We refer to the overall metric as “informed control” because we are computing the value of information, conditional on a given amount of control.

Value of Control is formalized as the rate of gain/loss in worth (for the guiding choice as a whole) as a function of incrementally improving or losing control over an event space. Each incremental gain or loss in control is modeled by a transformation of the probability density over the event space, such that some probability mass δ is moved to best or worst outcomes, and all other relative likelihoods are preserved (maximum entropy assumption). Equivalently, each change in control implies an intervention that moves our best understanding of the future to those possible worlds

that admit the revised probability density. Each incremental probability shift is then evaluated for its contribution to overall worth, creating graphs like the ones shown here.



Variables may offer large or small gains or losses, as a function of probability shifts. In the graphs above, the leftmost variable happens to result in the greatest overall gains and the greatest overall losses. Also, the slope for each of the gain and loss is maximal for that graph – so that the investment of effort to move the world (i.e., intervene to effect changes that accomplish the probability shift) has the greatest effect on this variable – the ‘most bang for each increment of probability buck.’ Therefore, it is the variable most worth controlling.

spaces that ‘carry’ or ‘produce’ those kinds of benefits or harm, and a quantification of the relative merit of events in the event space, as per each dimension of worth (e.g., utility function on an event space, reflecting the dimensions).

The “value of information” calculation specifies the increase in expected worth accrued, given information about an event space (if at least two action options are being contemplated). It shows how much can be gained by optimizing one’s *selection* among candidate actions, if a forecast is created for events mentioned in the choice. VOICE computes the value of information, for the current view, and for each increment of control in the control graph. The standard “Value of Information” calculation in decision analysis is, for VOICE, the degenerate case of zero additional control effort to shape events.

We provide for perfect and imperfect sources of information. The value of information may be greatest at a point reflecting non-zero control.¹² A formal property of VOICE is that the value of information is minimized for points at which control over events is maximized (or minimized) – i.e., at the extrema of the control curve. This accords with common sense – e.g., if control over a variable is perfect, so that an actor (self, or other) can absolutely guarantee a state of affairs, then there is no value to be added by learning more about how it might eventuate.

In summary, the positive side of a control curve reflects the rate at which we accrue worth as a result of good fortune, or, by investing in control of the given variable. The negative side reflects the impact for us, of bad luck, or of flawed estimates of adversary capabilities and commitment to invest itself in control of the subject event space (towards its opposed ends).

Summarizing Risk-Advised Planning

In general, there are many paths through the collection of tasks we are calling “risk-advised planning.” This is true in the case of one analyst, and vivid in the (more realistic) organizational case that there are several analysts, collaborating over each of the aspects in this method. In this view, each element of the method is being concurrently upgraded, based on all available information – that portion that is developed in that task, and any information/results communicated from the other tasks. Periodically, each element advises the others (i.e., as results stabilize and their significance and confidence warrant communication to others).

An important property of this critical method is that it inherently adapts to changing conditions. For example, as results of partial execution occur, the event spaces and human-

¹² Value of Control is the fundamental metric: a) more control implies less reliance on information, b) many information-development actions in the real world affect the world (e.g., the adversary notices!), and so have a control aspect, and c) less needs to be specified to quantify the value of control, versus the value of information.

systems models can be revised. This updates the guiding choice to then-current conditions, and then critical evaluations motivate next steps. This means that our standard on this risk-advised planning process is not convergence, but rather, one of adaptive rationality. We don't seek to finish it so much as to improve our understanding of the COA in the most effective way, and as a function of all real-time constraints, for the life of the COA.

In summary, Risk-Advised Planning is

1. Expressing COAs as a program of action – an analysis-COA – within a guiding choice.
2. Developing the necessary human-systems models to capture the feasible prospects, with emphasis on all influential events that participant actors create that provide for, contextualize, or result from activities in the COA.
3. Eliciting guiding values and formalizing them as a worth model.
4. Producing critical guidance in the form of normatively-grounded metrics derived from the guiding choice, and using these critiques to guide next steps.