15th ICCRTS
"The Evolution of C2"

Course of Action Ontology for Counterinsurgency Operations

Primary Topic: Topic 8 – C2 Assessment Metrics and Tools

Secondary Topic: Topic 3 – Information Sharing and Collaboration Processes and Behaviors, Topic 5 – Experimentation and Analysis

Timothy Darr, Perakath Benjamin, Richard Mayer

Knowledge Based Systems, Inc.
1408 University Drive East
College Station, TX 77840
(979)574-3189

Point of Contact: Timothy Darr (tdarr@kbsi.com)

Course of Action Ontology for Counterinsurgency Operations

Timothy P. Darr[1], Ph.D., Perakath Benjamin, Ph.D., and Richard Mayer, Ph.D.

Knowledge-Based Systems, Inc.

College Station, TX

Abstract

This paper presents the design and implementation of a course of action (COA) ontology for counterinsurgency (COIN) operations. This ontology supports the commander's C2 decision making processes for COA design, analysis and selection and represents the key elements of the COA planning process; specifically measures-of-performance (MOP) and measures-of-effectiveness (MOE) for describing COA plan states and objectives. These metrics are used in the generation of the plan as well as for assessment of the plan as it unfolds, which provides the decision maker with the ability to change course, if needed, during the execution of the plan. Using these metrics, it is possible to normalize the effects of a COA activity or task, allowing disparate activities and plans to be rationally compared. A utility-theoretic preference model represents the decision making trade-offs that a specific socio-cultural group makes over conflicting plan objectives. This preference model enables the twofold assessment of the plan: from the perspective of the blue force plan creator, and from the perspective of interested specific socio-cultural group, such as an Afghanistan tribal group.

I. Introduction

This paper presents a course-of-action (COA) ontology to support Deep Maroon, a decision-support tool for course of action (COA) design, analysis and selection[2] in the counterinsurgency (COIN) domain. Deep Maroon provides *middleware* capabilities to assist COA planners in *gap analysis* using interleaved forward (from COAs to meet commander's objective) and backward (from commander's objective to possible COAs) reasoning methods. The Deep Maroon COA ontology represents COAs, phases, activities, states, outcomes, measures-of-performance (MOP) and measures-of-effectiveness (MOE). The motivation behind this objective is the recognition of the complexity of the COIN domain and the attendant need to build decision aids that embody the commander and his planning staff as active collaborators. In a nut shell Deep Maroon will take as input the Commanders Intent, current situation assessment, COIN options, and knowledge of historical and situation defined preferences and reason to the "critical" gaps where the planners creativity needs to be focused or additional field experience needs to be acquired. Extensions to the COA ontology can support COIN doctrine, stability operations doctrine, and information operations (IO) doctrine. The use of MOP and MOE to describe outcomes and states, respectively, provide a normalized representation for comparing and analyzing disparate COA elements (states, outcomes, etc.).

Using the Deep Maroon preference model that represents the trade-offs that a decision-maker employs to resolve conflicting objectives, gap analysis will help the decision-maker to:

- rank and assess COA plan elements (activities, tasks, states, COA phases and outcomes),
- identify blind alleys and black holes[3],
- validate or challenge assumptions that are implicit in the COA or preference model,

[1] Contact author, tdarr@kbsi.com.

[2] Many of the examples in this white paper are for illustration purposes only. Specifically, activities, COA phases, MOEs and MOPs are based on discussions with SMEs and military publications ([FM 3-13], [FM 3-24]). More realistic examples, utilizing official documentation ([UJTL2009]) will be developed in future research.

[3] In Deep Maroon, a black hole is an undesirable state from the perspective of the U.S. forces; the state may meet military objectives, but has very grave non-military consequences. An example is a course-of-action for pacifying a village that involves aggressive destruction of homes and public buildings in an effort to root out insurgents; it may find the insurgents, but will likely turn the local population against the U.S. in the process. A blind alley is a state with a highly uncertain outcome from a socio-cultural standpoint. An example is an action that leads to a state that violates "honor" of an elder in an honor-based society.

- validate or modify Human Social Cultural Behavior (HSCB) models of the adversary, local population, the "unaligned middle", or other group of interest, and
- assess information operations (IO) MOP and MOE.

Deep Maroon employs a utility-theoretic preference model to compare and analyze COA elements. In this model, a linear utility function is defined over states or outcomes, where the utility function terms are the MOP or MOE that describe the states or outcomes, respectively. Utility theory, as applied to the Deep Maroon COIN domain, enables the following capabilities:

- resolve multiple, conflicting objectives
- decide among a set of complex alternatives
- evaluate trade-offs among multiple objectives
- game and project COA outcomes
- incorporate the perspective of multiple interest groups

The emerging Deep Maroon capability will provide support in environments where there are short decision cycles and incomplete knowledge of both the adversary and the interpretation of the indigenous population of COIN actions. Deep Maroon will be architected to handle situations where the "complete" plan is not known in advance, supporting re-planning as the situation unfolds (and the adversaries get their votes). Deep Maroon will serve as an aide and a critic raising to the forefront the consequences of COA options available right to the commander at the decision time, in decision real-time. In particular, Deep Maroon is envisioned to adapt well even in dynamic IO environments and should dramatically improve missions related to IO planning (both kinetic and non-kinetic), PSYOP, precision influence targeting, and even strategic communication.

Deep Maroon provides the following tangible and practical benefits to the COA planner:

- well-defined MOP and MOE for describing COA plan states and outcomes
- MOP and MOE-based metrics for evaluating the progress of the plan as it unfolds
- normalization of the effects of an activity, described as changes in state or outcomes, allowing disparate activities and plans to be compared
- a catalog of COIN activities defined by the states in which an activity applies and the expected states that result after application of the activity
- a utility-theoretic preference model that represents the trade-offs that a group of interest (blue forces, insurgents, unaligned middle, etc.) makes over conflicting objectives
- the ability to assess COA plan states, activities and outcomes from the perspective a specific interest group

Finally, Deep Maroon will provide institutional preference knowledge in the form of "leave behind" knowledge products [Crider2009]. It is often the case that a commander will arrive in an area of responsibility with the objective to "increase connection between the local population and the local government." The proposed COA to achieve this objective will be to (a) build roads so that the local population can more easily reach out to the local government, and (b) distribute humanitarian aid through the local government. Unfortunately, the commander finds out that the COA had no effect because the local population believes that the local government is not trustworthy and does not have the resources to deliver the promised aid. Because there are no leave behind knowledge products of the lessons learned, a new commander arrives in the same (or similar) area of responsibility and attempts the exact same COA, which yields identical results. Deep Maroon will mitigate this lack of communication and knowledge sharing by delivering preference knowledge products, contextualized to a given culture or geographic area, that a new commander can reference when formulating a COA plan.

A. Paper Organization

The remainder of this paper is organized as follows. Section II describes the planning context within which Deep Maroon operates. Section III describes the practical value that Deep Maroon brings to the war fighter. Section IV presents the COA ontology that guides the Deep Maroon capability. Section V illustrates the Deep Maroon technical approach. Section VI discusses performance and scalability considerations of the proposed solution. Section VII outlines the data requirements needed for Deep Maroon. Section VIII presents a summary, conclusions and future work for Deep Maroon.

II. PLANNING PROCESS CONTEXT

Figure 1 shows the COA planning challenge. Given a commander's objective in which essential services are to be restored in an urban or semi-urban environment, the decision-maker is faced with the challenge of achieving interim objectives to achieve the ultimate goal. Specifically, the decision-maker must establish some sort of local civil control in order for essential services to be restored. In order for some sort of local civil control to be established, however, it is necessary to establish security. The final COA plan to achieve the commander's objective consists of three phases: establish security, establish civil control and restore essential services.

A COA plan is designed to achieve the commander's final and intermediate objectives. One possible COA for achieving the commander's objective is shown in Figure 2. Each phase is terminated by an outcome that serves as a milestone for measuring progress of the plan. Each phase contains a sequence of activities that are performed to achieve the end-phase outcomes. The activities can be sequential, as shown in the establish security and restore essential services phases; or branch-and-sequence as shown in the establish civil control phase[4].

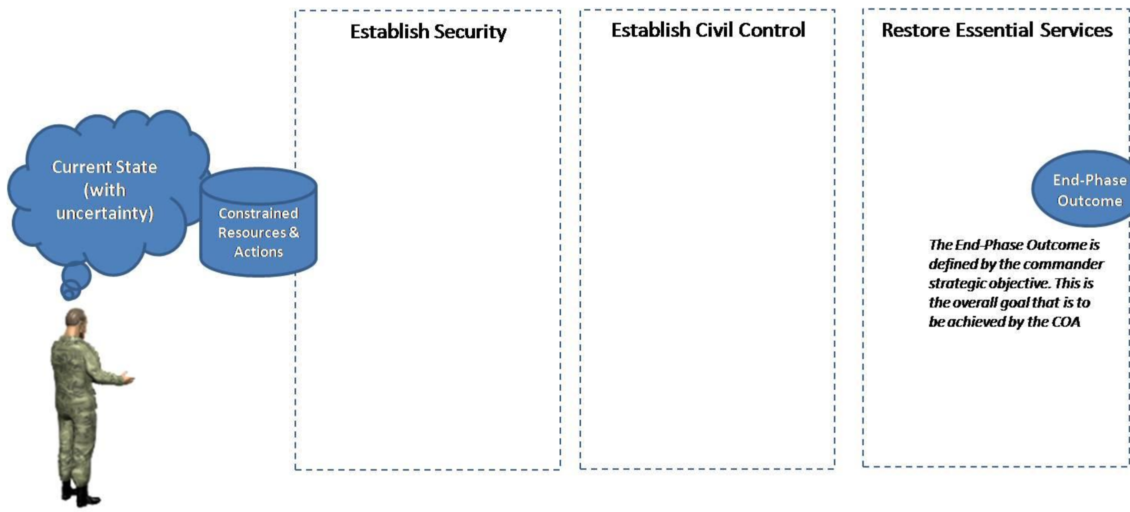

Figure 1 – COA Planning Challenge

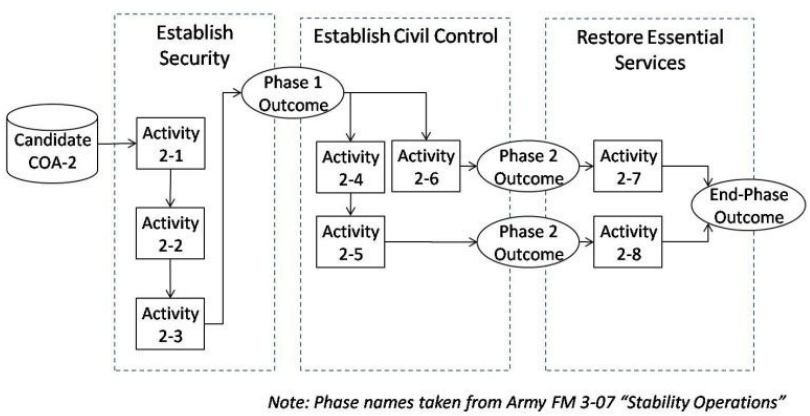

Figure 2 – COA Example

III. PRACTICAL VALUE TO THE WAR FIGHTER

Deep Maroon is principally aimed at supporting the counterinsurgency (COIN) mission. It does this by aiding the commander in determining the insurgent networks present within his area of responsibility (AOR) as well as possible connections outside his AOR. The following is the vision of how Deep Maroon will be used to support the commander in the development of an appropriate COA and in modifying the execution of an operation as new

[4] It is also possible to have concurrent activities, though not shown in the figure.

information is developed. The framework for this discussion is acquired from FM 3-24/MCWP 3-33.5 COUNTERINSURGENCY manual [FM 3-24].

In this scenario we will make the following assumptions and limitations:

1. A Marine Expeditionary unit of at least Regimental Combat Team size is to be deployed in a premeditated manner into an AOR in which a normal pre-deployment training and work up profile for the AOR is completed.
2. The deployment will be the first into the area.
3. The work up has allowed the deploying unit to conduct a detailed Intelligence Preparation of the Battlefield (IPB) but not at the detail level (i.e. local individual's names, stations in society, names of insurgent leaders, unique issues and facets of the community, etc.).
4. Interagency support is available.
5. The host nation government, while not proactive, is willing to participate even though its capabilities are poor.
6. The scenario will begin after kinetic operations have reached a level in which COIN activities can be initiated.
7. The scenario will focus on a narrow selection of COIN activities at the initial phase for demonstrative purposes as opposed to a full spectrum COIN operations plan from start to completion.

As kinetic activities decline, it is now believed that COIN operations can commence. In the initial phase of the plan the following short term goals are desired by the commander:

1. Establish security thru active presence in the AOR.
2. Establish movement control mechanisms.
3. Initiate active contact with the population to determine societal and political leadership structures and population needs.
4. Reestablish the host nation government involvement in the area of operations.
5. Initiate development of intelligence networks within the area to identify insurgents and activities.

In developing the plan for these goals, a COA which employs a logical line of operations (LLO) which are mutually supportable across the span of activities must be developed. To that end, the support staff must work from the desired plan end state as stated by the commander back towards the current state and from the current state and capabilities towards the desired plan end state simultaneously. Once the LLO's have been decided upon, Deep Maroon would assist in the implementation by helping to highlight the hidden nets, friendly, enemy, and alternative, that exist within the population and insurgency.

IV. COA Ontology

The Deep Maroon COA ontology supports COA planning. This section describes the COA ontology, which extends previously published material [Darr2009]. This ontology applies to the COA planning processes defined for the United States Army and Marine Corps for multiple domains, to include stability operations planning [FM 3-07], counterinsurgency planning [FM 3-24] and information operations planning [FM 3-13]. The core ontology includes definitions of the common concepts and properties for defining COA plans, including: COAs, COA activities, COA phases, measures-of-performance (MOP) and measures-of-effectiveness (MOE).

The COA Ontology consists of multiple sub-ontologies, each of which contains a small number of concepts and properties and can easily be integrated into other ontologies. For example, we have defined a measure-of-effectiveness (MOE) ontology that can be imported as a stand-alone ontology into other ontologies for the purpose of providing a common definition for representing MOEs. The urban Counterinsurgency (COIN) ontology is an example of how to extend the core ontology to define activities, MOP, MOE and phases that are tailored to a specific domain.

Qualitative MOPs and MOEs are used to describe plan states and outcomes. It is often the case that objectives are defined not in terms of fixed numeric quantities or ranges, but rather in qualitative terms such as: "reduce the number of attacks against coalition forces", or "increase the level of activity in the central market during daylight hours." MOPs and MOEs provide a way to normalize activities, states and outcomes so that disparate activities (provide electricity to a neighborhood, perform a detailed census) can be meaningfully compared.

A. *Decision Theory*

The Deep Maroon COA ontology is based on multi-attribute decision theory [Keeney1976]. This section defines the decision theory concepts, which are extended in the following sections to apply to the COA COIN planning domain. In this decision theory formulation, decision makers are presented with a set of alternatives $A = \{a_i\}$, where each a_i is described by a set of attributes. An evaluator $X_k(a_i)$ returns a numeric attribute level for an attribute k for the alternative a_i. Typically, these attributes are objectives that the decision maker wants to achieve and are either minimized or maximized. In the COIN domain, decision theory attributes are measures-of-effectiveness, as described below. The problem for the decision maker is to select an alternative that trades off the multiple, conflicting objectives as represented by the decision problem attributes. To compare alternatives that are described by disparate objectives, it is necessary to normalize the attribute levels. This is done using a value function $V_k(a_i)$ that returns a value for attribute k for the alternative a_i. Typically, the value returned by the value function ranges from 0.0 to 1.0.

Figure 3(a) and (b) shows the representation of decision theory attributes and alternatives in the Deep Maroon COA ontology. In this ontology, there are two types of attributes: objective attributes, which are described with a raw numeric value; and subjective attributes, which are described by a raw non-numeric value (for example; "high", "medium", or "low"). The raw objective and subjective attributes are translated to an attribute level using an evaluator. The attribute levels are translated to a value using a value function.

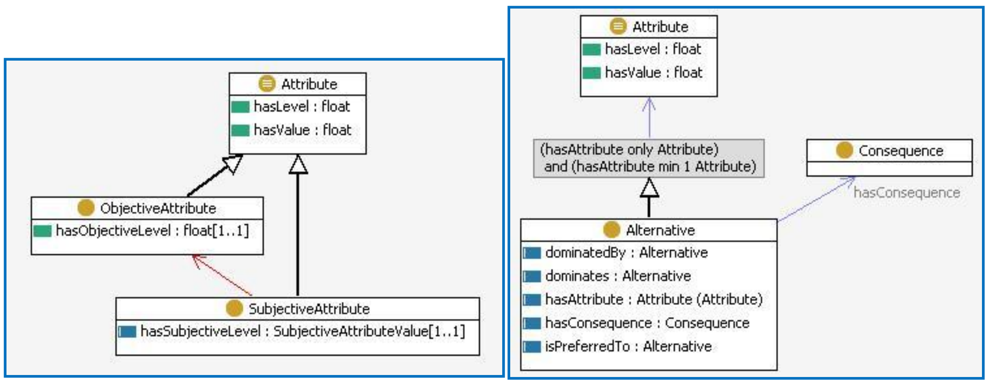

(a) Decision Theory Attributes (b) Decision Theory Alternatives

Figure 3 – Decision Theory Attribute and Alternative Concepts

Table 1 shows an example decision problem that consists of three alternatives, each described by four attributes. The table gives the name of the alternatives and the attribute levels for each attribute. The X1, X2, and X3 attributes are "more is better attributes", while the X4 attribute is a "less is better attribute."

Table 1 – Example Alternatives (from [Keeney1976], pg. 118)

Alternative	X1	X2	X3	X4
a1	7.5	344	0.47	12.15
a2	3.7	268	0.79	12.20
a75	6.7	250	0.24	12.92

Figure 4 shows the decision problem of Table 1 represented in the decision theory ontology.

The Deep Maroon COA ontology implements evaluators and value functions as SPARQL[5] queries over attributes. Figure 5(a) shows a default evaluator for objective attributes and Figure 5(b) shows an evaluator for subjective attributes whose values are "high", "medium" and "low." The objective attribute evaluator translates the raw objective level to the attribute level, if the raw objective level exists; if the raw objective level does not exist, the attribute level is set to 0.0. The subjective attribute evaluator translates the "high" raw level value to 500.0, the "medium" raw level to 100.0 and the "low" raw level to 50.0. These assignments for the subjective attribute are domain or context dependent and typically are specified by a domain expert.

[5] SPARQL is a query language for querying and constructing RDF graphs (http://www.w3.org/TR/rdf-sparql-query/).

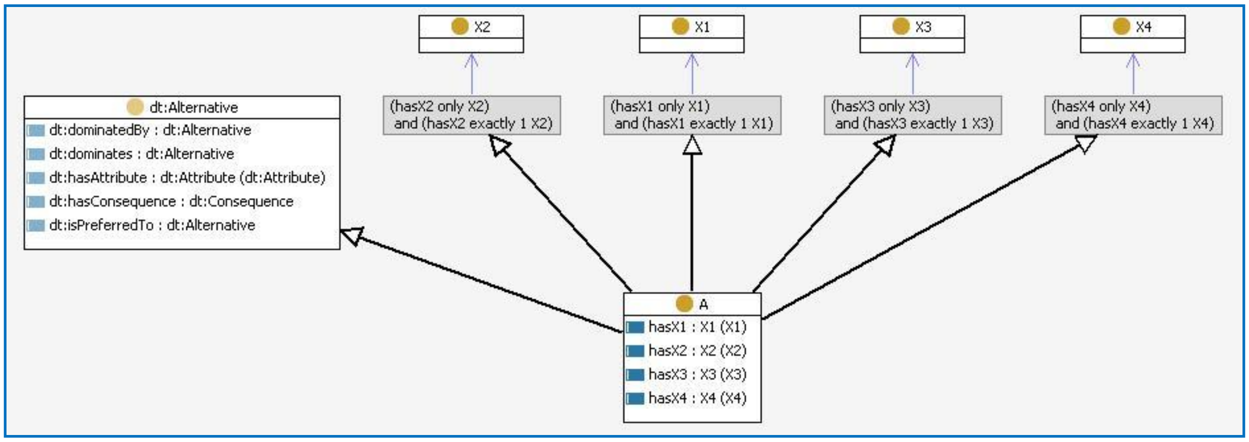

Figure 4 – Example Decision Problem

(a) Objective Attribute Level Evaluator (b) Subjective Attribute Level Evaluator

Figure 5 – Decision Theory Evaluators

Figure 6 shows an example value function for the X1 attribute in the example presented in Table 1. This is a linear, piecewise value function defined as:

$$V_{X_1}(a_i) = \begin{cases} 0.0 \; if \; V_{X_1}(a_i) \leq 5.0 \\ 0.5 \; if \; 5.0 < V_{X_1}(a_i) < 7.0 \\ 1.0 \; if \; V_{X_1}(a_i) \geq 7.0 \end{cases}$$

The CONSTRUCT clause in the SPARQL queries assign the value of the attribute and the WHERE clause filters on the level according to the piecewise linear value function.

B. Measures of Performance

The measures-of-performance ontology defines the measures of performance (Figure 7(a)) and state concepts (Figure 7(b)). According to COIN doctrine [FM 3-24], a measure of performance is defined as "a criterion to assess friendly actions that is tied to measuring task accomplishment." MOPs in the COA ontology describe states and have a time stamp, value, and value direction (increasing, decreasing, stable). States in the COA ontology are described by one or more MOPs. Example MOPs in the COIN domain include: response time to calls for help, number of hours of patrol coverage within a neighborhood, and amount of money distributed to the local population for infrastructure projects.

C. *Measures of Effectiveness*

The measures-of-effectiveness ontology defines measures of effectiveness (Figure 8(a)) and outcomes (Figure 8(b)). According to COIN doctrine [FM 3-24], a measure of effectiveness is defined as "a criterion used to assess changes in system behavior, capability, or operational environment that is tied to measuring the attainment of an end state, achievement of an objective, or creation of an effect." MOEs in the COA ontology are decision theory objective attributes and describe outcomes, including the commander's objective (end-of-plan outcome), objectives to be achieved at the end of each COA phase, and objectives to be achieved after each COA activity is applied at a given state. The MOEs are influenced by a set of MOPs[6]. Outcomes in the COA ontology are decision theory alternatives and are described by a set of MOEs. Example MOEs in the COIN domain include: number of local human intelligence (HUMINT) reports from the local population, level of interaction with the local population (children, elders, etc.), freedom of movement by the local population, and number of attacks against the coalition forces or local population.

Figure 6 – Decision Theory Value Function

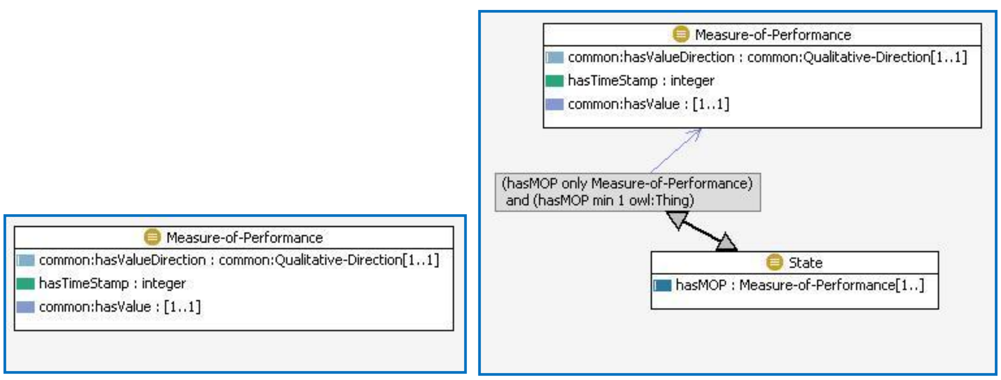

(a) Measure-of-Performance (b) State

Figure 7 – Measure-of-Performance Concepts

Figure 9 shows an example MOE (type of movement through checkpoints) and the MOPs that influence the MOE. The MOE represents movement such as legal commerce going through the checkpoints (indicating increased level of security), and people returning to their homes (also indicating an increased level of security). This MOE is positively influenced by the degree of movement control (checkpoints, barriers, recording who comes and goes) and the level of completion of the census for the area behind the checkpoint (the U.S. forces know who should be coming and going through the checkpoint and for what reason). Future research will develop models for how the MOPs influence the MOE.

[6] How MOEs are influenced by MOPs is an area of future research. The idea is that the value of a MOE can be modeled as some function of MOPs, using systems dynamics modeling, or some other modeling method.

D. Activities

The activities ontology defines COA activities (Figure 10). Activities are the tasks or steps within a COA. A given activity only applies in certain states, described by MOPs. The result of applying an activity results in a new state, described by MOPs. Within a COA, there are activities that precede and succeed a given activity. An activity has a target, which is the thing (person, group, community, etc.) that is the object of the activity[7]. Example activities in the COIN domain include: providing 24-hour patrol coverage and performing a census.

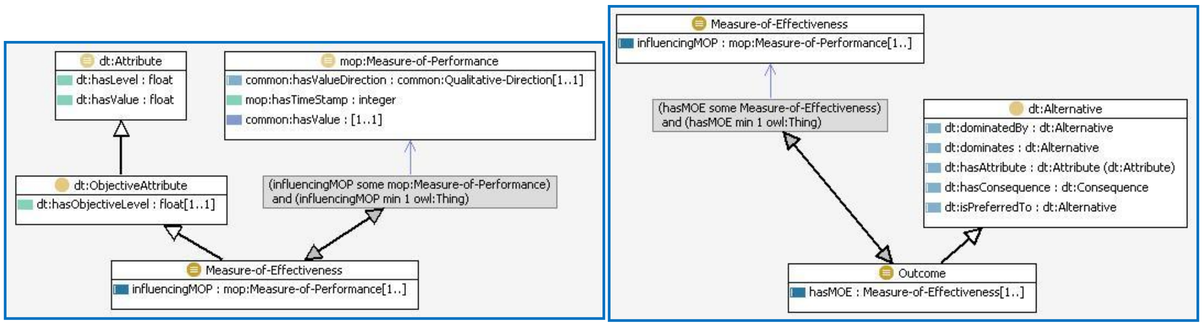

(a) Measure-of-Effectiveness (b) Outcome

Figure 8 – Measure-of-Effectiveness Concepts

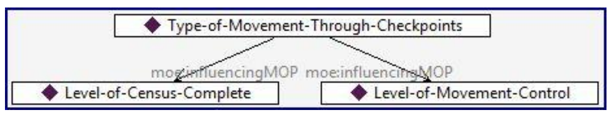

Figure 9 – Example Influencing MOP

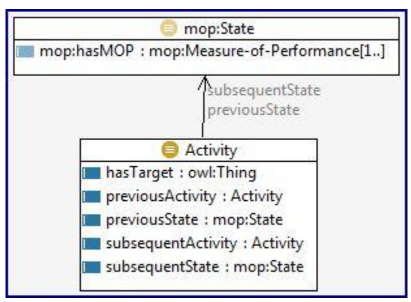

Figure 10 – Activity Concept

E. Courses of Action

The course-of-action ontology defines COA phases (Figure 11(a)) and COAs (Figure 11(b)). A COA phase is defined by a sequence of activities within the phase, an end-phase outcome, described by MOEs, and preceding and succeeding COA phases. A COA is defined by a sequence of COA phases, which may or may not have clearly-defined boundaries. Example COA phases within the COIN domain include: establishing security, establishing civil control and restoring essential services.

Figure 12 shows an example COA with three phases. The initial phase is establish security, which is followed by establish civil control, followed by restore essential services. Each phase has an outcome that is defined by one or more MOEs. The establish security phase outcome is described by the MOEs: attacks against US forces, and type of movement through checkpoints. The idealized objective of the establish security phase is to reduce the number of attacks against the US forces and increase the type of positive movement through checkpoints (commerce, people returning to their homes, etc.).

Figure 13 shows an example COA phase, establish security, with the activities: establish 24-hour patrols, perform census, and control movement. The example further shows the previous (State-X) and subsequent (State-Y) states for the control movement activity. The previous state is described by the MOPs low level of census and no

[7] Activity targets have not been detailed in the current version of the ontology.

movement control, both with values LOW and STABLE. This means that prior to the control movement activity, there is no census and no movement control and there is no discernable improvement on these MOPs. The (expected) subsequent state is described by the MOPs acceptable movement control and acceptable census level, both with values MEDIUM and INCREASING. This means that after the control movement activity, there is moderate levels of census and movement control and these MOPs are increasing.

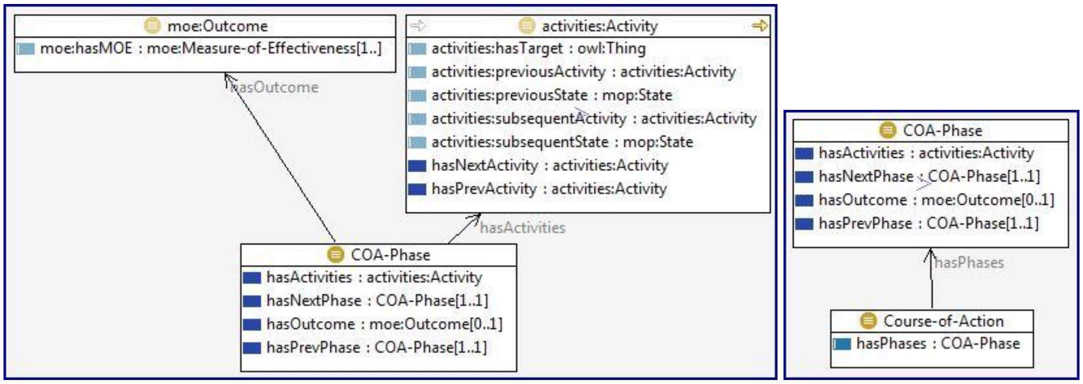

(a) COA-Phase (b) Course-of-Action

Figure 11 – COA Concepts

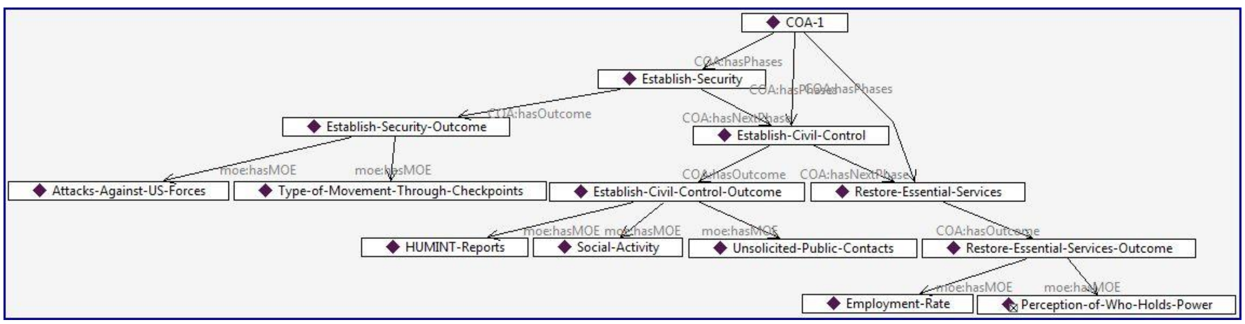

Figure 12 – Example COA Phases With Outcomes

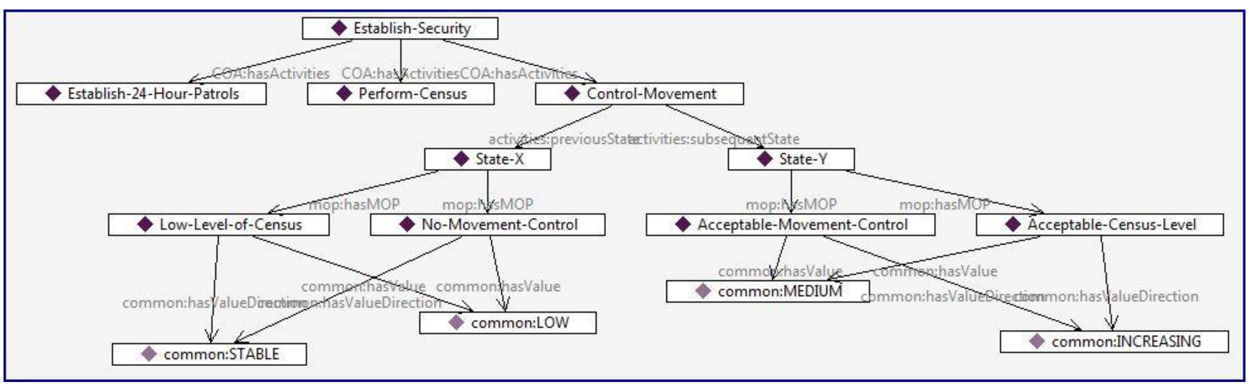

Figure 13 – Example COA Phases With Activities

F. Utility

Utility theory was originally developed in economics to measure the desirability of a good or alternative from the perspective of an agent [Keeney1976]. In the Deep Maroon domain, a state or outcome replaces a "good or alternative" in the economics application, and an interest group replaces an "agent" in the economics application.

A utility function is given by:

- $u: O \rightarrow \mathbb{R}$, where $O$is a Deep Maroon state or outcome and $\mathbb{R}$ is a real-valued number

A common form of utility function is a weighted sum of attribute values:

- $u(o_i) = \sum_k w_k * a_k[o_i]$, where w_k is an attribute weight, $a_k[o_i]$ is an attribute-value score that assigns a real-valued number to the attribute value a_k for outcome o_i. In the Deep Maroon domain, the $a_k[o_i]$ is the value of an MOP or MOE that describes a state or outcome, respectively.

The use of utility theory addresses how to compare disparate COAs, activities, or sequences of activities using normalized states and outcomes described by MOPs and MOEs. Consider two possible COAs, each defined by two activities, for establishing security:

- COA-1
  - Activity-1: perform census
  - Activity-2: US control movement to neighborhoods via checkpoint
- COA-2
  - Activity-1: U.S. performs 24 hour patrols
  - Activity-2: Local forces control movement to neighborhoods via checkpoint

Using utility theory, COAs are directly comparable by comparing the outcomes that result from the COA activities, which are normalized using MOEs. Table 2 shows possible MOE values for each of the COAs.

For COA-1, we assume that performing a census and having the U.S. control movement to neighborhoods will result in a high, but decreasing, number of attacks against U.S. forces and a moderate, but stable level of positive movement (commerce, people returning to their homes, etc.) through checkpoints. The rationale for the attacks against U.S. forces MOE is that checkpoints will be easier targets for suicide attacks and similar methods; the rational for the type of movement through the checkpoint MOE is that locals will be less trusting of U.S. forces at the checkpoints.

For COA-2, we assume that having 24-hour patrols and having local forces control movement to neighborhoods will result in a moderate but stable number of attacks against U.S. forces and a high, but stable level of positive movement through checkpoints. The rationale for the attacks against the U.S. forces MOE is that the U.S. forces will be more visible in possibly hostile areas; the rationale for the type of movement through the checkpoint MOE is that locals will be more trusting of local forces at the checkpoints.

Table 2 Comparison of COAs Using MOEs

Outcome	**MOE** **Attacks-Against-US-Forces**	**MOE** **Type-of-Movement-Through-Checkpoints**
COA-1 Outcome	HIGH, DECREASING	MEDIUM, STABLE
COA-2 Outcome	MEDIUM, STABLE	HIGH, STABLE

G. Preferences

The preferences ontology defines preferences over COA outcomes or states (Figure 14). A preference in this context is a relation between two outcomes or states in which one of the outcomes or states is preferred to the other outcome, given the perspective of a commander, decision maker, social / cultural group or other entity. These preferences are typically asserted by an SME while role playing a specific HSCB perspective or community group. Example preferences within the COIN domain include: in an agricultural community in which there is little or no electricity, a COA whose outcome involves restoration of economic self-sufficiency via the activity of building or restoring a canal system for crop irrigation, will be *preferred to* a COA in which the same outcome is achieved via the activity of providing electrical power to the local market. Figure 15 shows an example preference.

Preference reasoning provides a way to rank-order outcomes or states from the perspective of a given interest group (counterinsurgents, insurgent group, religious or ethnic group, etc.). An inference algorithm can use these preferences to reason about assessment of how a given outcome or state will be perceived and can assist a planner in the identification of black holes or blind alleys.

H. Inference Support

This section presents some examples of the types of inference supported by the Deep Maroon ontology. Figure 16 shows SPARQL-based inference rules for linking an activity with the previous and subsequent states. These rules assume that the previous state already exists and that the subsequent state is constructed dynamically. Figure 16(a) is the inference rule that applies to a Control-Movement activity and links the activity to any previous state in which the Movement-Control MOP is LOW and STABLE and the Census-Level MOP is LOW and STABLE. Figure 16(b) is the inference rule that constructs a new subsequent state in which the Movement-Control MOP is MEDIUM and INCREASING and the Census-Level MOP is MEDIUM and INCREASING. Note that the subsequent state inference rule dynamically creates the new state and the MOPs that describe it. The inference rules illustrates in Figure 16 apply to a specific activity. Each activity in the ontology would have similar rules for linking the activity with the previous and subsequent states.

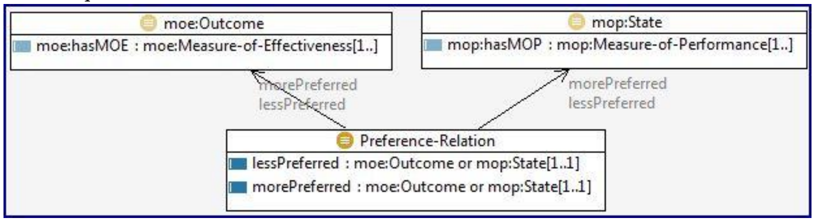

Figure 14 – Preference Concept

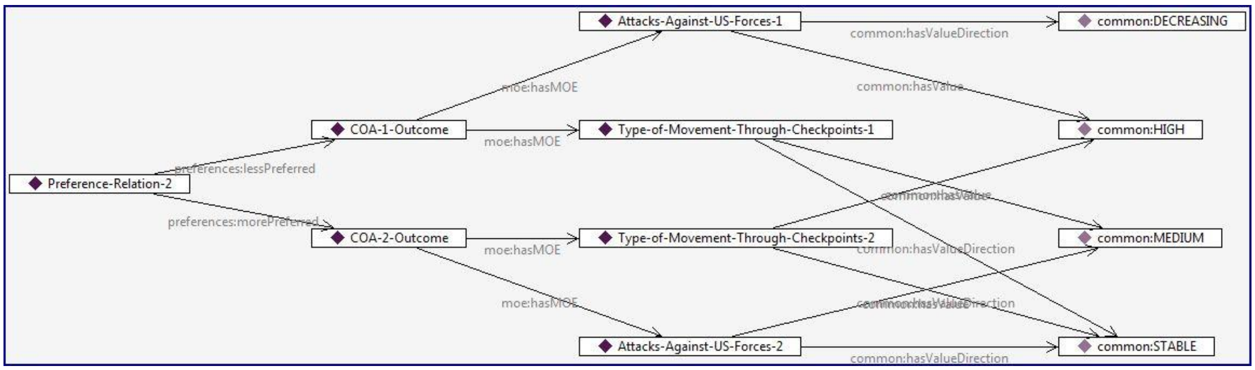

Figure 15 – Example Preference

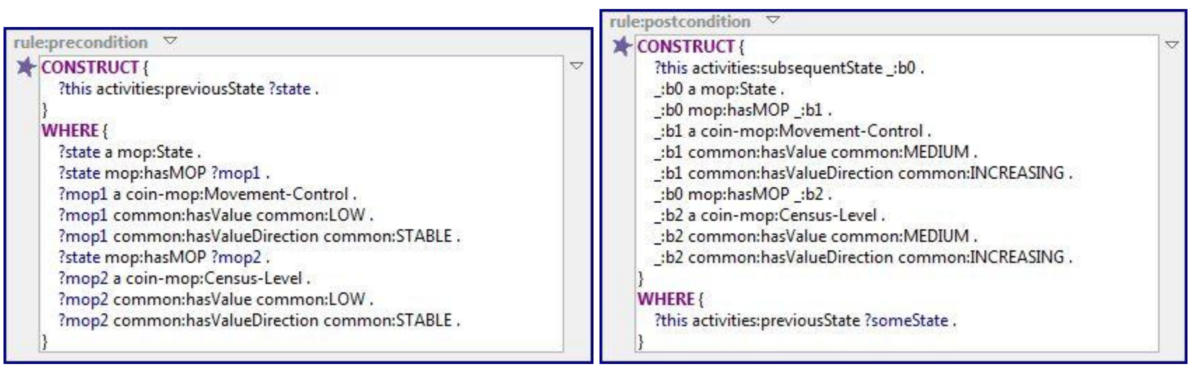

(a) Activity Previous State (b) Activity Subsequent State

Figure 16 – Deep Maroon Inference

The concept of preferential dominance is important to reasoning about preferences as it supports identification of black holes and blind alleys and allows outcomes or states to be pruned very efficiently, thereby reducing the computational complexity, at very little cost, of searching through potentially very large state and outcome spaces. Intuitively, one outcome (or state) dominates another outcome (or state) if it is "better" than the other outcome (or state) along all the values of the MOEs (or MOPs) that describe the outcome (or state). Figure 17 shows an example SPARQL rule for implementing dominance identification for the example given in Table 1. Recall that the X1, X2 and X3 attributes are "more is better" and the X4 attribute is "less is better." The CONSTRUCT clause establishes the dominates relationship for any pair of alternatives that satisfy the WHERE clause. The WHERE clause binds the values of the X1 through X4 attributes for a pair of alternatives. The FILTER clause checks the relationships between each of the attributes. The execution of this rule results in the inference of the relationship that a1 dominates a75, which is the expected result.

One possible inference method, currently under development, to determine the preference relationship between two alternatives is based on Imprecisely Specified Multi-Attribute Utility Theory (ISMAUT) [White1984]. ISMAUT is based on the observation that it is unrealistic, because of the inherent uncertainty in most decision problems, to specify with precision the trade-off weights that are required by the use of utility theory. In Deep

Maroon, uncertainty is present when trying to describe the current state of the environment, when trying to predict the expected effects of a given action, and when trying to predict how a target cultural group will perceive an action. The ISMAUT method deals with uncertainty by using the preferences between alternatives asserted by an SME, and any additional inequality relationships among the weights, to constrain the possible set of weights for the decision problem. Given two alternatives for which a hypothesized preference has been stated, it is possible to confirm the hypothesized preference by solving a system of linear inequalities that contains the SME-specified constraints. In this way, the SME is implicitly taking into account uncertainties in the specification of the preferences. Details of the ISMAUT method and how it can help deal with uncertainty can be found in [White1984].

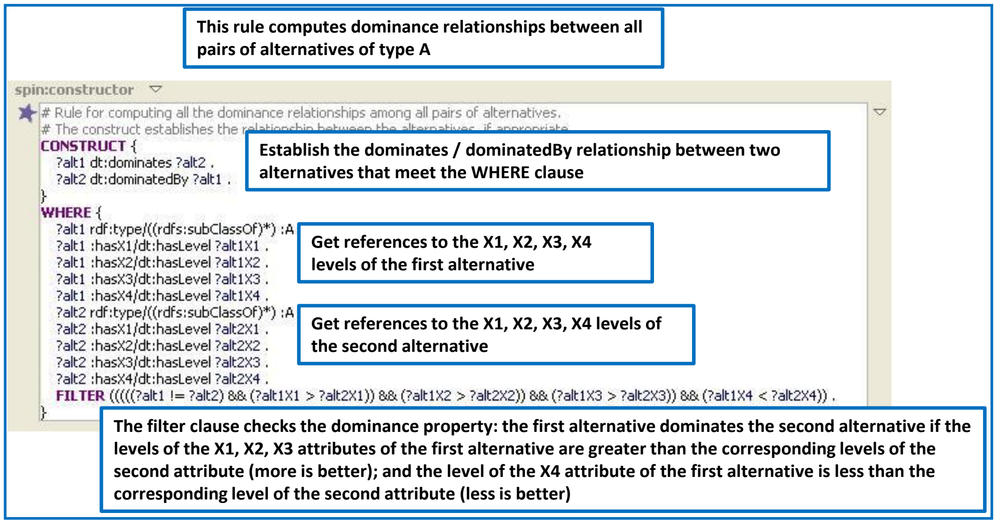

Figure 17 – Example Dominance Rule

V. TECHNICAL APPROACH

A variety of techniques are used to assess a set of candidate COAs. Typically, a small number of COAs are developed by the commander's staff based on a mission statement, the commander's intent and the commander's planning guidance. A subset of the developed COAs are designated by the commander for war gaming. During war gaming, the commander's staff determines the advantages and disadvantages of each designated COA, based on the enemy response (most likely, most dangerous to the blue forces, most advantageous to the blue forces) and battle space. The results of the war game are documented and provided to the commander for final decision. [MCWP 5-1]

Figure 18 shows a common strategy for assessing and war-gaming a COA by iteratively thinking forward from the current state to the desired end state outcome and thinking backward from the desired end state outcome to the current state. In thinking forward from the current state, the possible activities that are possible in a given state are determined. In thinking backward from the desired end-state, the possible states that lead to a given outcome, and the possible activities that can achieve those states, are determined. Thus, reasoning over a COA in this fashion is a state-space search in which activities apply in a given state and the application of an activity results in a new state. The complexity of this state-space search is compounded by the fact that the domain is fraught with uncertainty; not only is it uncertain what the current state is; but also how the adversary will respond to a given activity and the dependent change in state. In the non-kinetic world of COIN where the point of uncertainty in the forward reasoning (means-end) does not match with a state that connects to the backward reasoning (end-means) the "gap" materializes.

More formally, the "think forward" strategy is forward chaining reasoning using deductive methods and the "think backward" strategy is backward chaining reasoning using abductive reasoning, as shown in Figure 19. Goal-directed forward and backward chaining reasoning provides a way to reason about the desired trajectory of the plan over time (forward chaining), and given an end state, determine a set of starting states that would result in that end state (backward chaining).

In forward chaining, the sequence of activities that are available at each plan state can be determined by matching activity preconditions with the current state and asserting the new state that results from the application of the activity. Deductive methods allow the derivation of the possible tasks or activities that can be applied to a given state, and the resulting next state once a given task or activity has been applied. In a sense, the designated COAs represent the deduction path of going from the current state to the end-phase outcome, without the states explicitly identified. The Deep Maroon deductive methods will use the COA ontology to deduce the possible next states for each task or activity, and provide an assessment of each state against the preference model for a given interest group.

In backward chaining, the possible states that can achieve a given outcome are determined, followed by the activities that can achieve that state. Interleaving forward and backward chaining, with preference-based filtering, can help to mitigate the complexity of developing and analyzing realistic plans[8]. Abductive methods allow the inference of what must be true for an MOE to be achieved, or a task or activity to be applied. Abduction is often referred to as "reasoning by best explanation". That is, if state S is known and action A is known to result in a state S then we assume A occurred to produce S. Abductive reasoning is non-monotonic in that as additional information is acquired the assumption may need to be retracted. In addition, there may be (and often are), multiple actions that could result in the state S. The selection of which one to entertain becomes a matter of economy [Pierce 1955]; in Deep Maroon we have access to the "preference" knowledge for establishing a basis for that choice. Often, abductive methods are used in the initial development of a COA plan. The planners begin with the commander's intent and ask themselves "what must be the state of the world for the commander's intent to be realized?" In the current example, the abductive reasoning chain at the COA level would be:

- "In order to restore essential services, we must first establish civil control"
- "In order to establish civil control, we must first establish security"

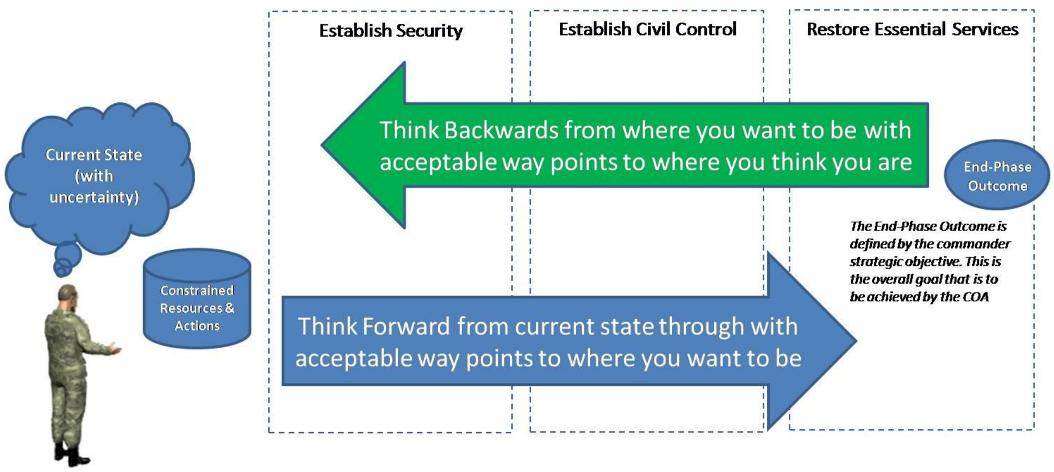

Figure 18 – COA Planning Strategies

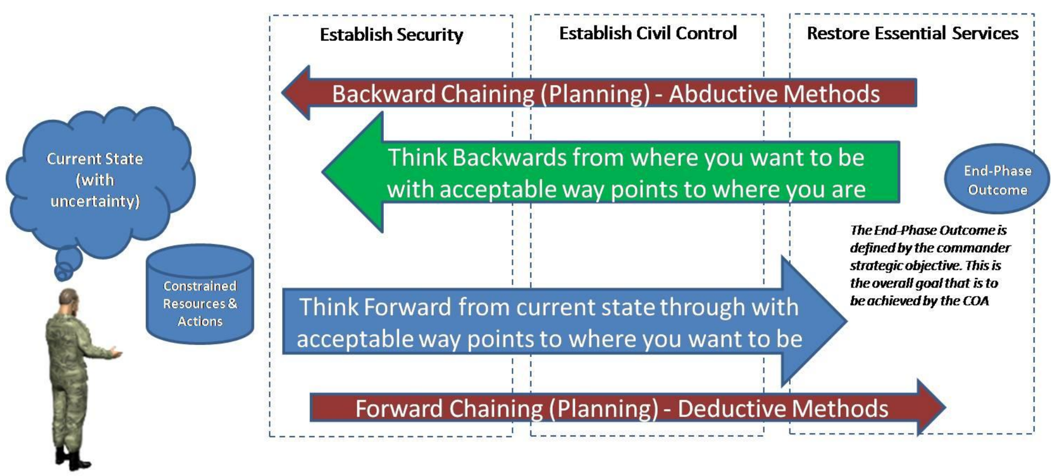

Figure 19 – COA Planning Strategies

[8] How the complexity is mitigated in this way is beyond the scope of this paper.

Once the COA phases are outlined, a finer level analysis uses abductive methods to develop the tasks / activities within a COA phase. For example, the MOEs define success at each end-of-phase outcome, the planner uses these MOEs to determine the states that must be true for the MOEs to be achieved, and the tasks / activities that can bring about the desired states. This may require a sequence of activities or tasks, or a branch-and-sequence pattern to model and reason about uncertainty.

Within the Deep Maroon forward and backward chaining reasoning methods, preference knowledge is used to identify black holes and blind alleys associated with each activity, state or end-phase outcome. In the context of a COA plan, a black hole is a state that once you get into it, you can never get out of it. An example of a black hole is an activity that leads to an inflammatory situation such as civil war or increased intra-militia violence. A blind alley is a state that is unproductive in that there is no feasible next state or no path to a goal state. Unfortunately, in the COIN context blind alleys often turn into black holes as you may not be able to retrace to an earlier state.

Example black holes include:

- An establish security activity (A1) in the COA plan that leads to a state in which the Attacks-Against-US-Forces are MEDIUM, STABLE is a black hole (all else being equal) if there exists another establish security activity (A2) that leads to a state in which the Attacks-Against-US-Forces are MEDIUM, DECREASING. In this case, A1 would never be chosen, since A2 results in a better outcome (decreasing attacks).

Example blind alleys include:

- An establish security activity in the COA plan that results in a state that does not match any of the activity preconditions of known activities cannot support additional forward chaining.
- An activity which results in a state that has already been pruned by the preferential dominance would be a blind alley.
- An activity that results in a state that cannot be assessed given the preference knowledge would be considered a blind alley. In this case, it would be wise to revisit the preference knowledge to gather more preferences to be able to make an assessment about the activity.
- In the backward reasoning process an activity for which the precondition state(s) have all been pruned by the preferential dominance would be a blind alley.

Figure 20 and Figure 21 illustrate a forward reasoning chain and the capabilities that Deep Maroon provides to the COA planner in more detail[9]. In Figure 20, three complete COA plans are shown[10]. Each COA consists of end-phase objectives ("Establish Security End-Game", "Establish Civil Control End-Game", etc.), activities within a phase that lead to the end-phase objectives, and the overall COA plan objective[11]. The end-phase objectives are described in terms of MOEs. This example also shows a branch and sequence after the establish security phase to go to activity 2-4 or 2-6 to handle situations in which things do not go as planned.

Figure 21 shows the use of the Deep Maroon preference knowledge to identify black holes and blind alleys in the forward reasoning process, using the Deep Maroon preference reasoning. For the first COA under consideration, the expected outcome that is the result of the application of activity 1-1 and 1-2 is uncertain (blind alley) when taking into account the preferences of the U.S. forces, local population, or some other interest group. For the second COA under consideration, activities 2-5 and 2-7 are infeasible. For the third COA under consideration, the expected objective at the end of the establish civil control phase does not meet that objective.

Figure 22 and Figure 23 illustrate a backward reasoning chain and the capabilities that Deep Maroon provides to the COA planner in more detail. In Figure 22, reasoning backward from the commander's intent objective, through each end-phase outcome is shown. The arrows from each outcome to an activity represent the possible "previous states" and the activities that achieve that state, which can possibly lead to the outcome. Multiple activity threads can lead to the same outcome; for example; activity 1-4 and 2-8 both lead back to the same establish civil control outcome. The same outcome can lead to multiple activities; for example, the establish civil control outcome in the lower portion of the figure can be preceded (enabled) by both activity 2-6 and 3-3. This differs from the deductive reasoning chain previously described in that there are fewer establish civil control outcomes.

[9] Not shown are the probabilities associated with a transition from activity to activity. These are left out to simplify the discussion of the basic concepts.

[10] The COA plans are distinguished by activity number: Activity n-m indicates the nth COA and the mth activity within that COA.

[11] For simplicity, the states within a phase are not shown. See section IV for a detailed discussion of the relationships between states and activities.

Figure 23 shows the use of the Deep Maroon preference knowledge to identify black holes and blind alleys in the backward reasoning process, using the Deep Maroon preference reasoning. We see that activity 2-7 does not follow from the establish civil control outcome as in the deductive reasoning process, but follows from the establish civil control outcome as defined in the third COA. Similarly, using backward reasoning, the projected establish civil control outcome shown in black does not achieve the objective as originally thought. Also, the Deep Maroon preferential knowledge provides a way to identify the non-kinetic effects of an activity. This is shown in COA 2 for the establish security phase in which activities 2-1, 2-2, and 2-3 have non-kinetic effects that are not obvious from a purely kinetic analysis. As highlighted in the upper left of the figure, using backward reasoning, it is possible to identify gaps in the plan; missing activities such that if the activity were present, a complete plan would be possible.

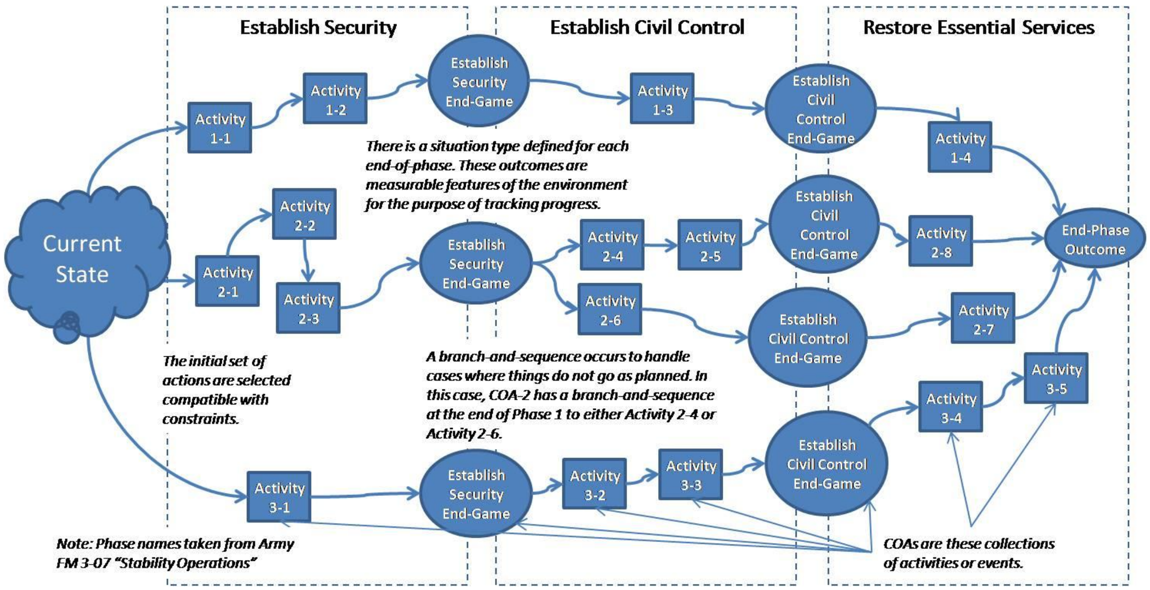

Figure 20 – Deep Maroon Forward Reasoning

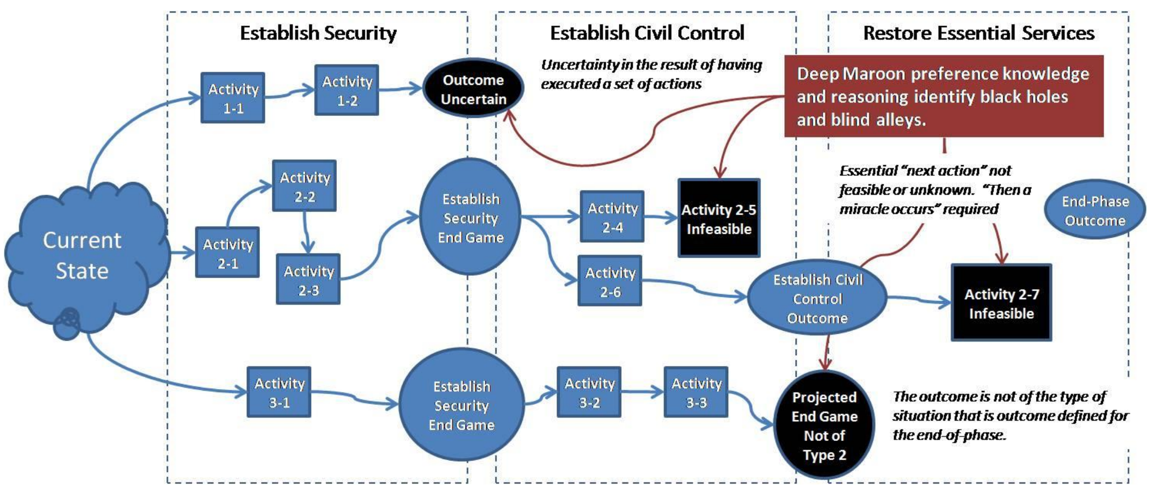

Figure 21 – Forward Reasoning Results

The forward and backward reasoning chains illustrated in this section are created using the COA ontology previously described in section IV. Specifically, the sequence of activities and states are constructed using SPARQL inference rules as illustrated in section IV.H. The construction of these chains are interleaved with preference knowledge to identify black holes and blind alleys, using utility-theoretic properties such as dominance as described in section IV.H.

The combination of forward and backward reasoning as described in this section provides recommendations to the planner, using the Deep Maroon preferential reasoning. Each reasoning method will reveal black holes and blind alleys that must be resolved by the planner. More importantly, the application of Deep Maroon forward and backward reasoning results in a *gap analysis* of the COA plan(s). Ideally, forward and backward reasoning will yield the same states, activities, and end-phase outcomes that were present in the original plan. More likely, there will be gaps in the plan, or disagreements in the results of the forward and backward reasoning. If there are disagreements, then this indicates that there are assumptions in the creation of the preference models that must be challenged, the models need to be modified in some way, the MOPs and MOEs that describe states and outcomes must be revisited, or the activities need to be analyzed against the assumed previous and subsequent states.

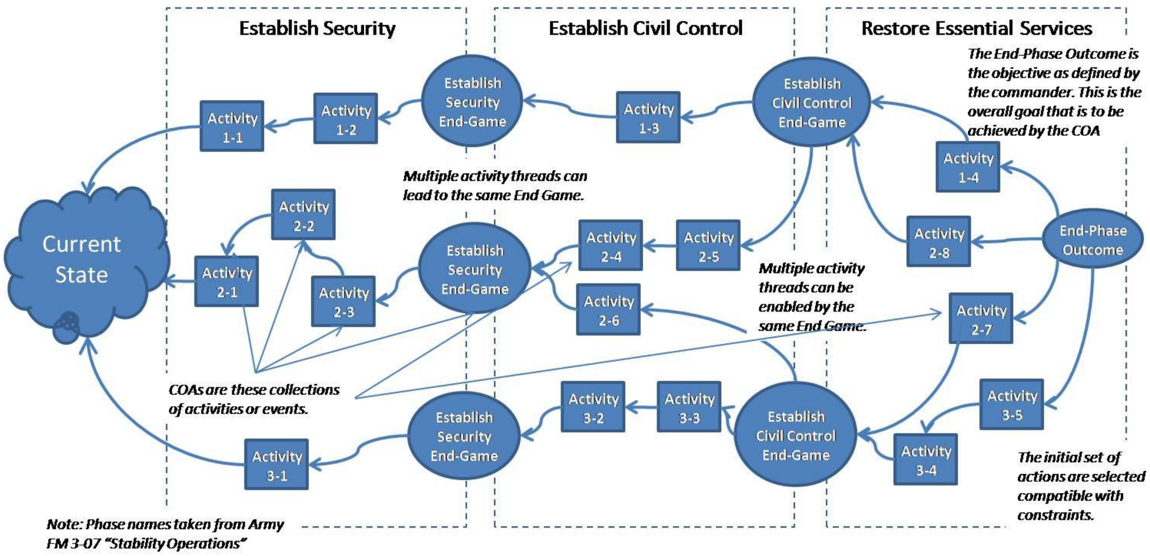

Figure 22 – Deep Maroon Backward Reasoning

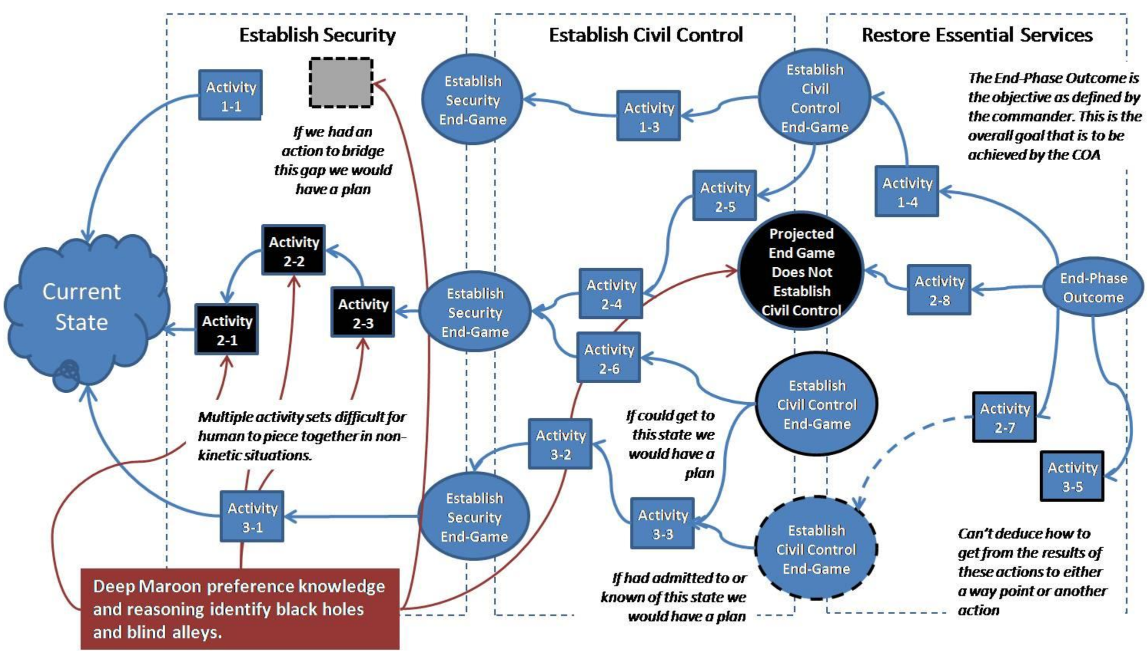

Figure 23 – Backward Reasoning Results

VI. Performance and Scalability Considerations

The computational complexity for reasoning about COA plans can easily become intractable as the number of activities and states grows. To mitigate this complexity, we will employ a factored envisioning approach [DeKleer2009]. Envisioning generates all possible qualitatively distinct states for how a plan might turn out. This maps naturally to the Deep Maroon activity / state / outcome model. Factored envisioning manages the potential computational explosion of this state space by (i) identifying non-overlapping, localized interactions and analyzing these interactions separately and (ii) utilizing *packing* to represent qualitatively similar interactions only once. For a kinetic COA application, factored envisioning has demonstrated a speedup as much as 10^6 over methods that do not use factored envisioning. While the Deep Maroon COIN domain is not the same as a kinetic COA domain, we expect similar performance improvements using these methods.

We have also performed some initial experiments on the expected gains using utility-theoretic methods, such as dominance, to prune paths in the state space envisionment graph. To determine the usefulness and feasibility of developing an algorithm to prune sub-optimal paths from a collection of possible paths, a computer simulation was created to test a candidate algorithm. The candidate algorithm serves to demonstrate our ability to find and eliminate paths that will not be of interest because each of them is completely dominated by at least one other path. A path consists of all the activities necessary to move from a given beginning state to a desired end state, but it does not include the beginning state or end state themselves.

The candidate algorithm works by comparing each step of each path. If a given path is dominated by another path then that path is eliminated (where dominance is defined such that one path dominates another if and only if it is at least as good as the other path on all steps, and, on at least one step, it is strictly better than the other path).

Using generated test data, we found we could prune away 60% - 70% of paths in a typical situation, which we felt was an excellent result[12]. Only in the situation where there were relatively large numbers of measures of effectiveness did the candidate algorithm perform poorly. The significance of this is that we can reduce the time and complexity of utility calculations by removing most suboptimal paths through the application of a simple, efficient algorithm.

VII. DATA REQUIREMENTS

Deep Maroon requires external data in the form of activities employed in the course of past COIN operations, MOPs and MOEs that describe states and outcomes, and any socio-cultural knowledge that can be used as input to the specification of preferences. Specifically, the following GIRH[13] / IPB[14] data would be useful:

- How misinformation and propaganda are distributed
- How the enemy fights in an urban area
- How the enemy employs weaponry (mines, IEDs, indirect fire, etc.)
- Circumstances under which an insurgent group would act and how
- How the enemy employs PSYOP
- How effects on the adversary are enhanced
- How receptive the population is to influences (authority, foreign messages, international entities, etc.)
- Who is the adversary
- What are the adversary's strategies and goals
- Cultural and social norms, factors, etc

Note that the list above contains mostly activities and context (cultural and social factors) that an insurgent group might employ. Deep Maroon would incorporate this data into its ontology and would contribute data back to the GIRH / IPB in the form of MOPs and MOEs for a given geographical and cultural area.

VIII. SUMMARY, CONCLUSIONS AND FUTURE WORK

This paper has described the preliminary technical approach for a course of action design, analysis and selection tool. We have developed an initial version of a COA ontology, including inference rules for building a sequence of states and activities, and performed some experiments to determine the tractability of this approach.

Future work includes the following:

- Continue to develop and refine the COA ontology, including MOPs and MOEs, inference rules, etc.
- Model the influence of MOPs on MOEs using systems dynamics models, qualitative simulation, or some other simulation method

[12] Complete test data is available upon request.
[13] Generic Intelligence Requirements Handbook
[14] Intelligence Preparation of the Battlespace

- Develop preference models for specific socio-cultural groups
- Complete the creation of inference rules for describing activity applicability - match the pre-state of an activity with a current state
- Develop capabilities to reason about preferences for the identification of black holes and blind alleys using ISMAUT [White1984]. Preferences are used to model HSCB perspectives for the purpose of supporting the decision maker(s) and COA planner(s) in COA development, war gaming, comparison and decision making
- Explore options for verification and validation of the Deep Maroon results. An initial approach is to rely on SME feedback; a more robust approach would be to mine historical data on past COA plans. The Deep Maroon approach would be used to generate a plan that satisfies the commander's objectives from the historical data and then compare the Deep Maroon results with the historical results. Assistance from SMEs would be needed in this approach as well to compare the Deep Maroon results with the historical results and assess how well the Deep Maroon plan would have performed with respect to the historical plan

ACKNOWLEDGEMENTS

This work was supported by the Office of Naval Research under Contract N00014-09-C-0334.

REFERENCES

[Crider2009] Crider, LTC Jim, "Inside the Surge", Center for a New American Security, Working Paper, June 2009.

[Darr2009] Darr, T. P., Benjamin, P. and Mayer, R., "Course of Action Planning Ontology", Ontology for the Intelligence Community Conference (OIC 2009), George Mason University, October 21-23, 2009.

[DeKleer2009] De Kleer, J., Forbus, K. D., Hinrichs, T., Yoon, S., Jones, E. K. "Factored Envisioning," International Workshop on Qualitative Reasoning (QR-2009), 2009.

[FM 3-07] "Stability Operations", Headquarters of the Department of the Army, Field Manual No. 3-07 (FM 3-07), October 2008.

[FM 3-13] "Information Operations: Doctrine, Tactics, Techniques and Procedures", Headquarters Department of the Army, Field Manual No. 3-13 (FM 3-13), November 2003.

[FM 3-24] "Counterinsurgency", Headquarters Department of the Army, Field Manual No. 3-24 (FM 3-07), Headquarters Marine Corps Combat Development Command Department of the Navy, Marine Corps Warfighting Publication No. 3-33.5 (MCWP 3-33.5), December 2006.

[Keeney1976] Keeney, R.L. and Raiffa, H. "Decisions with Multiple Objectives: Preferences and Value Tradeoffs", Wiley and Sons, New York, 1976.

[MCWM 5-1] "Marine Corps Planning Process", Marine Corps Warfighting Publication (MCWP) 5-1, September 24, 2001.

[Pierce 1955] Pierce, Charles Sanders, 1955. "Abduction and Induction", in Justus Buchler, editor, Philosophical Writings of Pierce, pp. 150-156, Dover Books, New York.

[UJTL2009] "Universal Joint Task List Database" (PDF version), Joint Staff, Washington, D. C., January 25, 2009.

[White1984] White, C. C., *et al*, "A model of multi- attribute decision making and trade-off weight determination under uncertainty", IEEE Transactions on Systems, Man and Cybernetics, Vol. SMC-14, No. 2, pp. 223-229, 1984.