Australian Government
Department of Defence
Defence Science and
Technology Organisation

The Centre of Gravity Network Effects Tool: Probabilistic Modelling for Operational Planning

Lucia Falzon and Jayson Priest

Command and Control Division
Information Sciences Laboratory

DSTO-TR-1604

ABSTRACT

The centre of gravity (COG) Network Effects Tool (COGNET) uses Bayesian networks to represent the COG causal structure. Its impact analysis capability facilitates the determination of the critical vulnerabilities that have to be degraded or negated to influence the COG. COGNET provides a modelling framework and a generic model database to aid knowledge reuse and knowledge transfer. Its graphical user interface is tailored to the military user and provides a user-friendly capability for populating and interacting with the models. In this report we discuss the methodology, development and implementation of the COGNET suite. The importance of this work is that it uses existing planning process concepts to facilitate the construction of comprehensive models in which uncertainties and subjective judgements are clearly represented, thus enabling future re-use and traceability.

Published by

DSTO Information Sciences Laboratory
PO Box 1500
Edinburgh South Australia 5111 Australia

Telephone: (08) 8259 5555
Fax: (08) 8259 6567

AR-013-155
July 2004

APPROVED FOR PUBLIC RELEASE

The Centre of Gravity Network Effects Tool: Probabilistic Modelling for Operational Planning

Executive Summary

A number of research projects are currently under way in DSTO's Command & Control Division motivated by a need for systematic and rigorous support for operational-level planning. The research is focussed on developing decision-making tools with an underlying conceptual framework that complements the ADF planning process but with a theoretical underpinning derived from decision analysis. All these tools feature interfaces that take into account the military user's level of expertise and do not require a background in operations research. They are being integrated using well-founded effects-based concepts to form an Integrated Modelling Environment (In-MODE), featuring a knowledge framework, which enables storage for future analysis and re-use.

This report addresses one of the In-MODE research projects: the Centre of Gravity Network Effects Tool (COGNET), which uses causal probabilistic networks to represent the relationships among the critical capabilities and requirements for a Centre Of Gravity (COG) construct. These networks provide a visual representation of the COG causal structure to clarify thinking and provide a useful way to record and impart this thinking; and Bayesian techniques to determine which actions are most likely to achieve a desirable end-state. COG analysis is an integral and cognitively demanding aspect of military operational planning. It involves identifying the enemy and friendly COG and subsequently determining the critical vulnerabilities that have to be degraded or negated to influence the COG of each side. A thorough understanding of the relationships between a COG and its underlying capabilities and requirements is crucial to the development of a sound military plan. The relationship structure is often complex and not always easy to determine. COGNET goes a long way to facilitate this task and it provides an effects-based analysis capability.

COGNET embraces the COG analysis concepts in military planning thereby encouraging structured problem solving. It provides a graphical representation of complex relationships between capabilities and requirements that facilitates a shared understanding and has a user-friendly capability for populating, evaluating and interacting with the models. The suite includes an impact analysis tool that provides a measure of effectiveness in terms of the impact that the critical requirements have on a COG as well as an ability to determine the amount of influence that needs to be achieved over relevant critical requirements in order to obtain a predetermined effect. It provides a generic models database to facilitate the construction of comprehensive models and enable future re-use and traceability, and has provision for compiling large complex networks. We describe a modelling framework based on the causal relationships among the critical capabilities and requirements for an operation. We

then use this framework as a basis for the construction, population and analysis of Bayesian networks to support a rigorous and systematic approach to COG analysis. The importance of this work is that it uses existing planning process concepts to facilitate the construction of comprehensive models in which uncertainties and subjective judgements are clearly represented.

Authors

Lucia Falzon
Command and Control Division

Lucia Falzon is a Senior Research Scientist in DSTO's Systems Simulation and Assessment Group. Her current research interests include research into probabilistic modelling techniques to support military planning, in particular network models for supporting Command and Control decision-making.

Jayson Priest
Command and Control Division

Jayson Priest is a Research Scientist in the Command and Control Division. He holds a PhD in Physics from the University of New England. His previous research activities are in the areas of plasma surface modification and modelling and simulation of gaseous electronic phenomena. His current research interests include modelling, simulation and analysis of command and control processes for development of decision support systems in a Joint operational environment.

Contents

Abbreviations

AAR	Air-to-Air Refuelling
ADF	Australian Defence Force
ADFWC	Australian Defence Force Warfare Centre
AHP	Analytic Hierarchy Process
ANP	Analytic Network Process
BAI	Battlefield Air Interdiction
BDA	Battle Damage Assessment
CAP	Combat Air Patrol
CAST	Causal Strengths
CC	Critical Capability
COA	Course of Action
COA-Sim	COA Simulation
COAST	COA Scheduling Tool
COG	Centre of Gravity
COGNET	COG Network Effects Tool
CPT	Conditional Probability Table
CR	Critical Requirement
CV	Critical Vulnerability
GUI	Graphical User Interface
HQ	Headquarters
In-MODE	Integrated Modelling Environment
IO	Information Operations
JMAP	Joint Military Appreciation Process
JOPC	Joint Operations Planning Course
POL	Petroleum, Oil and Lubricants
SIAM	Situational Influence Assessment Module

1. Introduction

A number of research projects are currently under way in DSTO's Command & Control Division motivated by a need for systematic and rigorous support for operational-level planning. The research, which has been influenced by the Joint Operational Planning Course (JOPC) instructors at the ADF Warfare Centre (ADFWC) as well as our observations of operational planning teams, is focussed on developing decision-making tools with an underlying conceptual framework that complements the Australian Defence Force (ADF) planning process but with a theoretical underpinning derived from decision analysis. Among these is the Course Of Action Scheduling Tool (COAST) [Zhang *et al*. 2002], which provides a mathematical representation of a course of action (COA) to enable quantitative analysis for sequencing and scheduling of tasks in an optimised COA; and COA Simulation (COA-Sim) [Matthews & Davies 2003], a software agent-based wargaming tool to explore the feasibility, effectiveness and risk of an operational-level COA. All these tools feature interfaces that take into account the military user's level of expertise and do not require a background in operations research. They are being integrated using well-founded effects-based concepts to form an Integrated Modelling Environment (In-MODE), featuring a knowledge framework, which enables structured data storage for future analysis and re-use.

This report addresses one of the In-MODE research projects: the Centre of Gravity Network Effects Tool (COGNET), which uses causal probabilistic networks to represent the relationships among the critical capabilities and requirements for a Centre Of Gravity (COG) construct [Falzon *et al*. 2000 and 2001; Priest *et al.* 2002, Falzon 2004]. COGNET models typically represent the functional decomposition of the centre of gravity to identify its influencing elements and to categorise them into a hierarchy: COG, critical capabilities and lower-level capabilities and requirements. This type of decomposition ensures that the direction of influence travels up the functional hierarchy. In other words, targeting a critical requirement at the bottom of the hierarchy produces an effect on all related elements higher up. Using this model it is possible to investigate the effect that a set of actions has on the COG.

In the following sections we give a comprehensive description of the concepts underlying COGNET. The structure of the report follows the historical structure of the development of COGNET, starting with a look at some of the well-tried techniques of decision analysis to lead into the motivation for the adoption of Bayesian modelling for our problem of interest. A good understanding of the key concepts of operational art is as essential for military operational planners as it is for developers of operational planning support tools. In Section Three we describe the concepts underlying the relevant aspects of military planning, as an introduction to a description of the COGNET suite and how it is envisaged that COGNET will be used in operational planning.

2. Probabilistic Models for Decision Support

2.1 Systematic decision making

Decision analysis typically draws on Operations Research techniques such as probabilistic modelling, optimisation and game theory. These form a basis for analytical methods to evaluate and structure incomplete knowledge in order to reason about how an action taken in the current decision would lead to a result. Decision analysis provides a systematic method for structuring a complex decision problem. It is motivated by a need to understand the current state of knowledge, its limitations and implications. The analysis is normally performed with some decision-making criterion in mind, such as the valuation of outcomes in terms of benefits and costs or reduction of risks, independent of costs and benefits, and the criteria adopted are explicitly declared. Uncertainties about quantities or about the appropriate model structure should be clearly defined and a systematic sensitivity and uncertainty analysis should be conducted in order to determine the effect that changes in input values and model topology have on the concluding analysis [Morgan and Henrion 1990].

The kind of complex decision problems under consideration here, i.e. military planning for joint operations, have multiple dimensions of cost and impact and are typified by an environment of uncertainty. The systems analysis approach of disaggregating the problem systematically and evaluating the expected utilities associated with multiple attributes (including inherent risk) provides a way around this problem. In this way individual domain experts can make a judgement in their own area of expertise rather than an overall judgement in a complex domain. Judgements about values are by nature subjective and vary among individuals. In order to solve a complex decision problem we need a structure that models the components within the system and the causal influences and effects among the components so that it enables elicitation of judgements and a way to represent them quantitatively. In the approach advocated by Saaty, judgement is based on paired comparisons relative to a common criterion or goal [Saaty 1996]. Similarly, probabilistic modelling techniques, such as Bayesian belief networks and influence diagrams, rely on the ability of probability theory to process context-sensitive beliefs [Pearl, 1988]. These models are populated with conditional probabilities (the probability that A is true given the context, C), which are easier to estimate than absolute probabilities.

Barclay *et al.* present a methodology [Barclay *et al.* 1977] which is based on four elements: a set of initial courses of action, a set of possible consequences for each initial act, the value of each act in terms of money, utility or some other unit and the likelihood that a particular act will result in a particular consequence. The first two elements are an integral part of Course of Action (COA) development in military operational level planning. Ideally, the last two elements should form the basis of a systematic COA analysis.

In this section we briefly describe some of the techniques used in quantitative decision analysis for modelling decisions under uncertainty. We begin with Bayesian networks, which is the modelling technique of choice for the class of problems discussed in this report. We include the others for completeness.

2.1.1 Bayesian networks and influence diagrams

A Bayesian network is used to model a domain that has inherent uncertainty due to a combination of incomplete knowledge of the domain and randomness in the environment. The network may be represented by a directed acyclic graph whose nodes correspond to random variables linked by causal dependencies. The causal direction is represented by the direction of the arcs in the graph. Each node has associated with it a set of potential states. A node in a Bayesian network is called a parent of another node if there is an arrow from the former pointing to the latter. Each node has associated with it a conditional probability of the node being in a specified state given the states of its parent nodes. If a node has no parents it is assigned an initial probability of being in each of its potential states. A Bayesian network can be constructed from the qualitative causal knowledge, estimates of probabilities and subjective beliefs held by a domain expert.

Influence diagrams are graphical models for structuring predefined sequences of actions and observations [Jensen 2001]. They were originally developed as a compact representation of decision trees but can be thought of as an extension of Bayesian networks. Making a decision is modelled as choosing a set of decision variables in a Bayesian network and fixing their values unambiguously [1988]. This would alter the probability distribution of the consequences of the decisions and determine the expected utility associated with the decision chosen. Any Bayesian network can be converted into an influence diagram by adding decision variables, representing the choices available to the decision-maker, and the expected utilities associated with each decision. Arcs pointing to utility and chance nodes represent functional dependence whereas arcs pointing to decision nodes show which variables will be known to the decision-maker before the decision is made, implying time precedence. The value of each decision variable is imposed from the outside to meet some objective. Influence diagrams have two main limitations: the graph must contain a directed path encompassing all decision variables; and the next set of observation and decisions in the sequence must be independent of the current set.

2.1.2 The Analytic Hierarchy Process [Saaty 1994]

The Analytic Hierarchy Process (AHP) derives ratio scales of relative magnitudes of a set of elements by making paired comparisons with respect to importance, preference or likelihood of a property they have in common. Decision making with the AHP is based on ranking activities in terms of relative ratio scales.

The various elements of a decision problem are organised into a multiple-level hierarchy. Each level has multiple nodes with respect to which the alternatives on the next level are compared. The first step in the analysis is to compare the elements in each level in pairs according to their contribution to the parent node in the level above. A pairwise scale of relative priorities is derived from the pairwise comparisons for the group. This is repeated for all groups on all levels. A weighting process uses these priorities to rank the overall importance of the criteria.

Judgements for comparing the criteria of a particular level can be represented by a reciprocal (i.e. $[a_{ij}] = [1/a_{ji}]$) square matrix whose elements reflect the relative importance of the criterion represented by the row compared to the criterion represented by the column. The components of the eigenvector of the matrix, termed the priorities vector, represent the conversion of the pairwise comparison of the criteria into a ratio scale (the sum of these numbers must be one). The principle eigenvalue gives a measure of the consistency of the judgements. The procedure is repeated for every criterion. The sub criteria under each criterion are compared with respect to that criterion to obtain their local priorities. The importance of each sub criterion is weighted by the priority of the parent criterion to obtain its global priority. At the lowest level the global priorities are summed to obtain the overall priorities.

The AHP can be extended to an Analytic Network Process (ANP) to incorporate dependencies and feedback [Saaty 1996]. While hierarchies are concerned with the extent of a quality among the elements being compared, a network is concerned with the extent of influence on some element with respect to a given quality. A network is well suited to modelling dependence relations among components. It makes it possible to represent and analyse interactions and to synthesise their mutual effects by a single logical procedure.

2.1.3 Markov decision processes

Markov decision processes use the fundamental properties of Markov processes to find optimal strategies for decision making using dynamic programming techniques. Each decision taken incurs a cost or reward and affects the decision maker's state, so affecting future choices. A Markov decision process is a controlled stochastic process, in which costs are assigned to state transitions. There are four main components in these decision processes: a set of states, a set of possible actions, the immediate reward of the actions and the effects of the actions. In essence the decision-maker observes the current state and must choose among a finite set of possible actions incurring a possible cost (or reward) for each action chosen. The costs and state transition probabilities are functions only of the last state and subsequent action (i.e. they satisfy the Markov property). The set of rules by which the decision-maker chooses alternatives at each stage of the process is called a policy. It determines the transitions that optimise outcomes according to some performance criteria, for example, maximising the expected aggregate reward.

In Bayesian decision theory, the underlying Markov process has uncertain transition probabilities and rewards. It uses the Bayesian theory of probability to characterise degrees of uncertainty in terms of subjective probabilities. The decision-maker can update information about the current state by performing experiments.

2.2 Bayesian modelling

Bayesian methodology is based on conditional probabilities: if A and B are not independent then the belief in A given that B is known is denoted by $P(A|B)$ and described as "the probability of A given B". Probability theory defines this conditional probability as being equal to the probability of A and B, divided by the probability of B, $P(A|B) = P(A,B)/P(B)$ - it represents the degree of belief in the state of A when the state of B is known. Similarly the probability of B given A can be calculated in the same way thus yielding *Bayes' Law*,

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)} .$$

This rule is the very basis of Bayesian analysis. It allows information updating in response to new evidence. Three steps are involved in Bayesian modelling: developing a probability model that incorporates prior knowledge about the probabilities required; updating knowledge by conditioning probabilities on observed data; evaluating the model with respect to the data and the sensitivity of the conclusions to the assumptions.

Bayesian networks [see Pearl 1988 and Jensen 2001] are directed acyclic graphs representing the causal relations in a particular domain. The topology of the directed graph defines the conditional independence relationships among the variables in the network represented by the nodes. Each variable has associated with it a set of two or more potential values or states. The probability of being in each particular state of a node is conditioned on the states of each of its parent nodes, that is, the strength of the causal relationships among the nodes is expressed as a conditional probability. Let x_i denote some value for variable X_i and pa_i denote some set of values for X_i's parents, then $P(x_i|pa_i)$ denotes the conditional distribution.

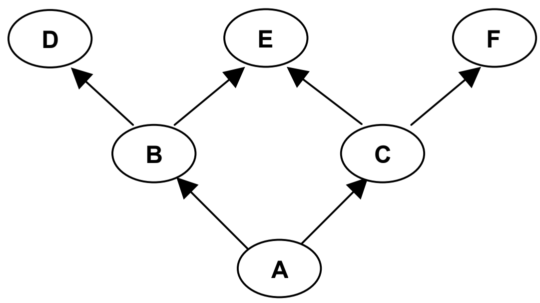

Figure 1: An example Bayesian network

In Figure 1, node *C* is a parent of nodes *E* and *F* and a child of node *A*, representing the fact that the state of node *F* is conditioned on the state of node *C*, which is in turn conditioned on the state of node *A*. For each node a conditional probability distribution $P(x_i|pa_i)$ must be specified. If the node has no parents then its unconditional probability $P(x_i)$ must be specified instead. For the network in Figure 1 we require the probabilities *P(A), P(C|A), P(B|A), P(F|C), P(E|B,C), P((D|B)* in order to compute the joint distribution,

$$P(A,B,C,D,E,F) = P(A)P(B \mid A)P(C \mid A)P(E \mid B,C)P(D \mid B)P(F \mid C) \ .$$

This equation is valid because it is assumed in these networks that each variable is conditionally independent of its non-descendants in the network, given its parents. Variables X_1 and X_2 are said to be conditionally independent given X_3 if $P(X_1 \mid X_2, X_3) = P(X_1 \mid X_3)$.

As Pearl [Pearl, 1988] points out, the advantage of this graphical representation is that it allows a specification of direct dependencies representing the fundamental qualitative relationships. In fact the network structure and link direction defines the conditional independencies among the variables in the network according to a criterion called d-separation, which is loosely defined in terms of causal dependencies with reference to Figure 1 as follows. For paths traversing diverging arrows ($D \leftarrow B \rightarrow E$) or serial arrows ($A \rightarrow B \rightarrow E$) the connection between the variables at each end of the path is considered blocked (i.e. they are d-separated) if *B* is known. However, the path traversing $B \rightarrow E \leftarrow C$ (converging arrows) should not be interpreted as transmitting information between *B* and *C* until *E* is instantiated. *B* and *C* are considered marginally independent; they become mutually dependent once evidence on *E* is received.

The conditional probabilities required for a Bayesian network can be elicited from a domain expert. They may be completely subjective estimates of the likelihood of an event. However, in Bayesian formalism the measures must obey the fundamental axioms of probability theory. The network is a graphical representation of a decision maker's subjective and uncertain knowledge of a domain.

An added advantage of using these models for reasoning is due to the conditional independence property described above. In order to determine the conditional probabilities $P(x_i|pa_i)$, we can ignore all the relationships in the network except for the ones between X_i and its parents. Determining such context-dependent probabilities is much more compatible with human reasoning than estimating absolute probabilities. In the statement "the probability of *A* given *B*", *B* serves as a context of the belief attributed to *A* and is much easier to determine than "the probability of *A* and *B*". Probabilities provide the means for drawing inferences from causal connections and the relative strengths of those connections.

2.2.1 Reasoning about uncertainty

Building a Bayesian network involves three tasks [Drudzel &Van der Gaag 2000]. First one must identify the variables that need to be included in the model and possible states or values have to be specified. The second task is to identify the relationships among the variables and construct a directed acyclic graph to represent the direction of causal dependence, adding more variables as needed to construct a comprehensive model. Bayesian nets are by convention Markovian: each variable is independent of its non-descendents, conditional on its parents. This property helps us to structure the network when it comes to deciding which variables should be explicitly modelled among the parents of a particular variable. The final task is to assign conditional probabilities to each fragment (a child node and its parents) and unconditional probabilities to the parentless nodes, which we will interchangeably call leaf or initial nodes. These probabilities might come from various sources: statistical data, literature or experiential knowledge from domain experts. The latter source is useful not only for providing the probabilities required but also for fine-tuning and verifying numbers obtained from other sources. However, eliciting probabilities from experts is time-consuming and is susceptible to biases and inconsistencies. In most cases of interest to us there is not enough historical data available and initial probabilities as well as conditional probabilities must be elicited from subjective area experts with little or no knowledge of the statistical aspects of data, thus making consistency an important requirement. By consistency we mean that elicited probabilities do not contradict each other, for instance the subset of an event having a higher probability than the event itself.

The properties inherent in Bayesian modelling: the graphical structure reflecting direct dependencies; the ability to clearly reflect uncertainty using probabilities; and the fact that these probabilities are context-dependent and therefore easier to determine, make Bayesian networks very suitable for COG representation.

3. COG Representation Using Bayesian Nets

3.1 Operational planning concepts

The Joint Military Appreciation Process (JMAP) has been adopted as the basis for operational level planning in the ADF. The initial stage of any operational-level planning process, such as the JMAP, typically includes some form of mission analysis. This involves identifying and analysing the superior commander's intent in order to ensure that commanders and staff can determine which tasks are essential to achieve the operational objective. Correct assessment of the objective is deemed to be crucial to success at the operational level. In current ADF thinking the objective can be achieved by targeting the enemy's centre of gravity (COG) through their vulnerabilities while protecting one's own, so that the operational objective and the COG are inexorably linked. The COG, a key concept of operational art, is defined as *that characteristic,* *capability or locality from which a military force, nation or alliance derives its freedom of action, strength or will to fight at that level of conflict. The centre of gravity at each level of conflict might consist of a number of key elements* [ADDP 5.0.1 2002].

Once the enemy COG has been determined, the planners must generate suitable COA. Suitability refers to whether it meets the objectives as detailed in the mission analysis step. In the operational context a COG is typically a high-level capability or characteristic, such as “Operational Sustainability” or “Force Projection Capability” (indeed, it has also been described as “a *focal point* that gives a force purpose and direction” [Echevarria 2003]). Therefore directly targeting the enemy COG is not usually feasible, and a critical capability analysis is conducted at this stage of the planning process. A critical capability (CC) is defined to be *a characteristic or key element of a force that if destroyed, captured or neutralised will significantly undermine the fighting capability of the force and its centre of gravity* [ADDP 5.0.1 2002]. Each CC might have a number of associated critical requirements (CR), which are essential for it to be fully functional. These requirements may be further decomposed into critical vulnerabilities (CV): elements that are potentially vulnerable [ADDP 5.0.1 2002, Chapter 2].

Related to these two concepts is the notion of a decisive point: *a major event that is a precondition to the successful disruption or negation of a COG of either combatant.* Each COA developed in this step must target the enemy COG by exploiting the enemy’s critical vulnerabilities in a sequence of actions known as a *line of operation*. The planners must also identify critical vulnerabilities and decisive points from the enemy’s perspective, that is, related to the friendly COG. In summary, this stage of the JMAP should identify the enemy’s COG, a number of approaches to undermine and neutralise it, the decisive points and lines of operation for each of these approaches and the critical vulnerabilities contributing to each decisive point. In addition the planners should determine their own force’s COG and related decisive points and critical vulnerabilities.

It is also stipulated that “throughout COA development the staff must consider the ‘cost-benefit’ that results in apportioning capabilities and rates of effort to achieve objectives and tasks”, and they should also “identify and analyse the consequences of potential risks and how they may impact on own and higher missions” [ADDP 5.0.1 2002, Section 2.41]. All four elements of Barclay’s methodology defined in Section 2.1 are thus represented. It could be argued that what is missing is a systematic comparative assessment based on the costs, benefits and risks of each COA, broken down into its component actions. A comparative assessment is done in Step 4 of JMAP, the Decision and Execution step. The decision criteria used in this step are derived from the war game analysis in Step 3 (COA analysis). They are measured over the whole COA rather than over each action in the lines of operations. This type of decision-making is categorised as a holistic decision rule in the literature. Sage [Sage 1981] describes it as reasoning by intuitive affect and typically absorbing information by looking at the situation as a whole rather than disaggregating it into its component parts. The two approaches are complementary and compatible and, ideally, both types of assessment should be carried out.

3.2 Modelling key concepts

Our aim is to develop a relational model between the key concepts currently used by operational planners in COA development: end-state, centre of gravity, critical vulnerabilities, decisive points and lines of operation. It is motivated by a need for a rigorous methodology for COA analysis based on these concepts. In [Zhang *et al.* 2000] we draw out these relationships and apply them to develop a modelling framework. In this section we use this framework as a basis for an influence model to support COA development. The models are structured to reflect the relationships among the planning concepts detailed above.

At the COA development stage, the planners develop sequences of actions to connect the initial and military end-state through time and space using available means [Zhang *et al.* 2002]. To do this they must first determine ways to influence the operational centre of gravity, directly or indirectly. It is important at this stage to keep the COG in mind. As explained in [Giles and Galvin 1996], "exploiting weaknesses and vulnerabilities are clearly important considerations; however, doing so will not cause the deteriorating effect desired unless it influences the centre of gravity".

The relational model we present here aims to support this activity. It represents the COG and all the elements that influence it. Functional decomposition of the centre of gravity is used to identify the influencing elements and to categorise them into a hierarchy: COG, critical capabilities (typically abstract functions) and lower-level capabilities and requirements, such as general functions, processes and physical systems. As we will see below, such decomposition ensures that targeting a physical system produces an effect on all related elements higher up in the hierarchy. COGNET seeks to exploit the benefits of systematic modelling and, as will be shown in the next section, Bayesian nets in particular. In addition, COGNET uses existing operational planning process concepts as an underlying framework, making it easier to embed its usage by military planners.

3.3 Characteristics of COGNET models

COGNET models are causal probabilistic networks that represent the functional decomposition of the centre of gravity to identify its influencing elements and to facilitate critical capability analysis. A typical COGNET model would be structured as shown in Figure 2, a network produced in HUGIN [www.hugin.com], a software tool for building Bayesian networks, which forms the Bayesian engine for COGNET.

The COG node at the highest level is dependent on its parent nodes representing high-level critical capabilities, which, in turn, are dependent on other critical capabilities or requirements, as expressed in the definitions of these concepts given in Section 3.1. This representation is useful for several reasons: a) the combined states of the low-level requirements, which are usually observable elements, allow us to estimate the state of unobservable (possibly abstract) higher-level capabilities; b) the relationship among such requirements and capabilities is often uncertain and is better represented by a probability rather than a unique deterministic quantity; c) the conditional independence property in Bayesian networks allows us to reason about these probabilistic strengths in a context-dependent way.

In critical capability analysis we are trying to determine the effect that lower-level nodes have on the higher-level nodes rather than attempting to infer the converse. Indeed, in general, it is easier to make a judgement about the state of low-level nodes, which typically represent physical resources that contribute to a capability or abstract function.

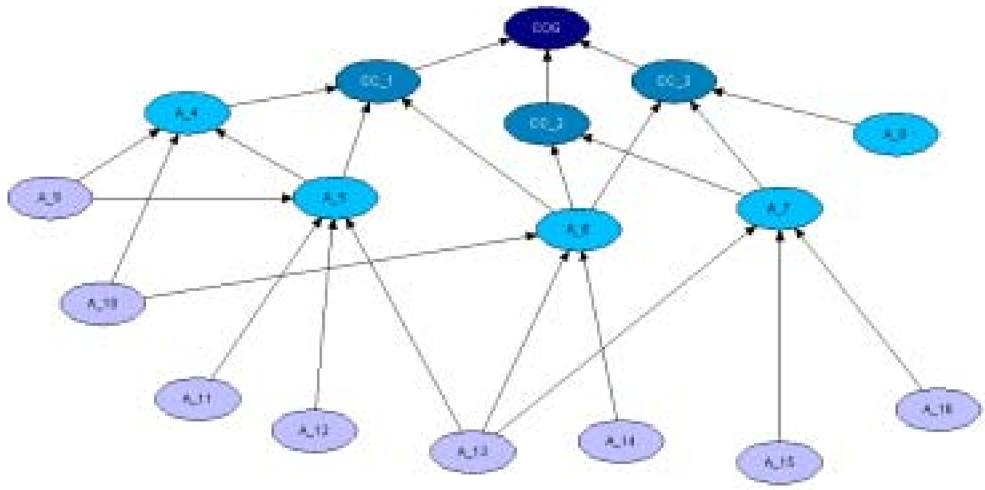

Figure 2: A simple Bayesian network representing a typical COGNET structure

Figure 3 shows an example Bayesian network representation of a typical COG/CC/CR/CV analysis. The network represents the results of a COG analysis exercise conducted by students at a Joint Operations Planning course. It is based upon a fictitious scenario used for training purposes in which the perceived threat was an imminent invasion of an island belonging to an ally. The threat COG was assessed as the ability of the enemy to project force, which was subsequently broken down into its associated critical capabilities and requirements. The leaf nodes represent critical elements that are potential targets and hence potentially vulnerable. In this particular network each node can be in one of three states, named “strong”, “degraded” and “unavailable”, although it is possible to assign any number (at least two) of states to a node and give each state any name as required.

It can be shown that, in such networks, leaf nodes that have a significant impact on the COG are typically ones that have many paths leading to its node. This network represents a simple example of a COG network, yet even in this case it can be seen that it is possible to trace a path from certain high-impact nodes, such as "petroleum, oil and lubricants" (POL), in several different ways. Conversely, low-impact nodes are typically ones that have a single path to the COG. Network structure conveys important conceptual information. It allows users to distinguish between direct and indirect dependencies.

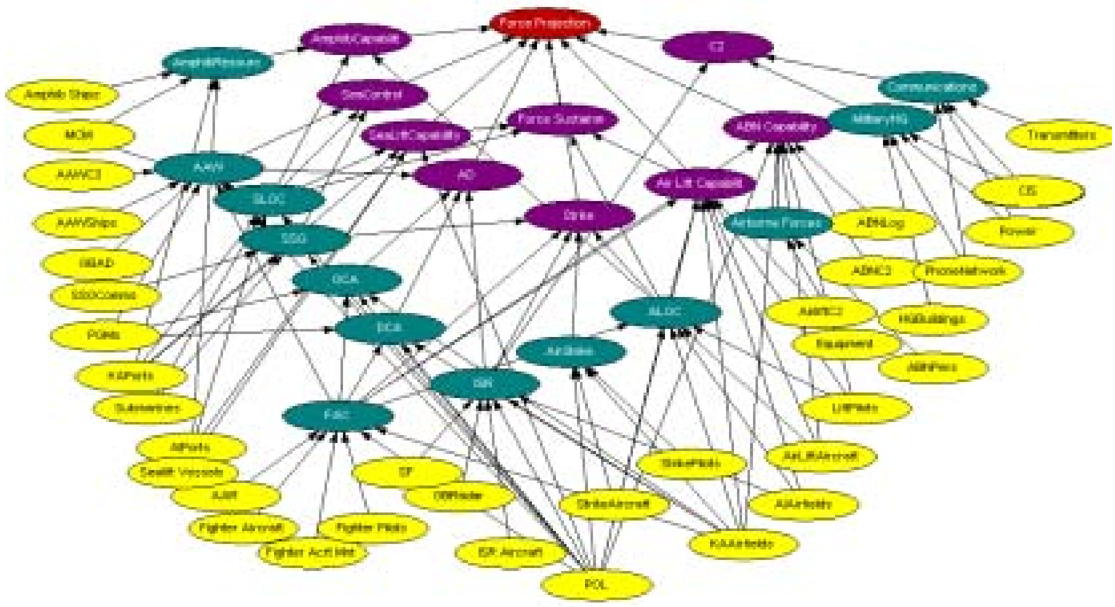

Figure 3: A typical COG network

3.4 Other planning tools based on probabilistic models

Decision theory techniques, based on systems analysis and probabilistic modelling, have been used in different ways for such military problems as planning counter-measures for terrorist threats [Paté-Cornell & Guikema 2002]; planning naval operations [Kidd 2002]; and collaborative strategic planning [Rosen & Smith 2000]. Evans *et. al*. use *Dynamic Bayesian Nets* to represent the causal relationship between lower-level friendly tasks and higher-level effects on adversary systems in order to guide plan generation and to analyse the observed impact of planned military actions during plan execution [Evans *et. al*. 2003]. Another probabilistic modelling tool, the Situational Influence Assessment Module (SIAM) [Rosen & Smith 1996], uses graphical models known as *Influence Nets*, which were specifically developed for analysing the causal relations for complex situations. SIAM nets represent events that influence other events, whose relationships are modelled as promoting and inhibiting influences, or causal strengths, rather than strict conditional probabilities. SIAM is typically used in an unstructured planning session to brainstorm different ideas. Its proprietary population and analysis mechanisms mimic a special case of Bayesian networks (influence diagrams, see Section 2.1.1) and use an approximating algorithm to convert user-defined "causal strengths" into conditional probabilities [de Poel *et al.* 2002]. COGNET, on the other hand, has been designed (in collaboration with ADFWC planning instructors) to support the ADF COG analysis as defined in doctrine. It is based on classical Bayesian network modelling techniques, which use conditional probabilities and model functional and probabilistic dependencies. The goals that motivated the development of all these tools are very similar. Graphical probabilistic models such as SIAM's influence nets, Bayesian nets and influence diagrams, provide a visual representation that facilitates reasoning and enhances shared understanding of complex situations. Moreover, this reasoning can be better imparted and recorded for future reference. The analysis functions are well developed and add significant value to the modelling

The advantages of these systematic techniques are many: a structured approach clarifies the complex interrelationships among the variables of interest and facilitates "what if" analysis; the problem can be handled in a piecemeal approach without losing sight of the whole structure; the representation of the problem is an ideal way to clarify and store concepts arising from planning sessions, enabling future re-use and traceability; uncertainties and subjective judgements are clearly represented. Like the other probabilistic planning tools COGNET seeks to exploit the benefits of systematic modelling. In addition, COGNET uses existing ADF planning process concepts as an underlying framework. As a result, it fits directly into ADF planning doctrine, making it easier to embed its usage by military planners into standard operating procedures. A significant proportion of the project has been dedicated to developing an interface and a suite of tools that match users' requirements.

4. The COGNET Suite

COGNET provides a visual representation of the centre of gravity causal structure and an impact analysis capability, which facilitates the determination of the critical vulnerabilities that have to be degraded or negated to influence the COG. Its graphical user interface is tailored to the military user and provides a user-friendly capability for populating and interacting with the models. COGNET provides a framework and database structure, which can serve as a knowledge base representing generic causal relationships to aid knowledge reuse and knowledge transfer. In this section we describe each aspect of the COGNET suite and discuss how it is envisaged that military planners will use them.

4.1 Generic models

During the course of the development of COGNET the need for a modular knowledge representation became evident [Falzon *et al.* 2001]. Building a comprehensive COG network systematically takes more time than is normally available during a crisis situation or even during an exercise. However, there is scope to build such models as part of the ongoing deliberate planning for the ADF. Deliberate (long-term) planning at the operational level aims to develop plans that can be adapted, when a conflict or situation arises, to meet the objectives set out by strategic guidance. While the COG for a particular force may change according to circumstance, a relatively fixed network can reflect the current force structure and capabilities over a fixed set of critical capabilities depending on a fixed set of requirements. The network structure is invariant for a range of problems but the conditional probabilities may vary with respect to the specific COG being considered.

A knowledge representation framework expressing the invariant functional and causal relationships is being constructed for each specific operational capability. This serves as a knowledge base expressing generic relationships whose probabilistic strengths may be adjusted to tailor each model to a particular situation. The framework is organised in hierarchies of sub networks that reflect generic military categories such as Command & Control, Protection, Deployment etc and their underlying requirements. They can be combined as required to construct a comprehensive Bayesian net for a specific scenario from a library of modular subnets that reflect the hierarchical structure and capture the stable patterns of probabilistic relationships. Relevant entity data for these models are stored in COGNET as a relational database system.

4.2 The COGNET system and user interface

The COGNET database provides a framework for representing generic causal relationships to aid knowledge reuse and knowledge transfer.
Figure 4 depicts the COGNET system. Upon entry into COGNET a user can select a data-source from those available, each of which may have been created and managed by different headquarters (HQ) or organisations. Each data-source consists of three main sets of tables defining entities, types and scenarios. The entity tables contain the list of entities and their relationships to one another according to the capability models developed. The generic model database is structured according to operational-level capability categories, which range from standard warfare capabilities such as Battlefield Air Interdiction (BAI) to Information Operations (IO). Development of the database is ongoing.

The relationships between entities needed to form capability Bayesian networks are stored in the database in the form of parent and children associations. In addition, country associations can be added to entities or capabilities in the database. For example, a country can have BAI (and its appropriate entities) associated with it if it has that capability. The ability to tag individual entities with a country association also allows the removal (or addition) of entities from a capability model for that particular country. As an example let us consider two countries A and B that are in a situation that requires BAI. Let us also consider in this situation that country A has Air-to-Air Refuelling (AAR) available and country B does not (AAR allows BAI to be used in situations where its operating base is a significant distance from the Area of Operations). The remaining entities required for BAI are the same for both countries. In COGNET this is implemented by tagging the BAI entity and each of its sub-elements, with the exception of AAR, with countries A and B. AAR is only tagged with country A. The ability to associate countries with entities in this manner allows several country capability models to be generated from a single generic 'catch all' model.

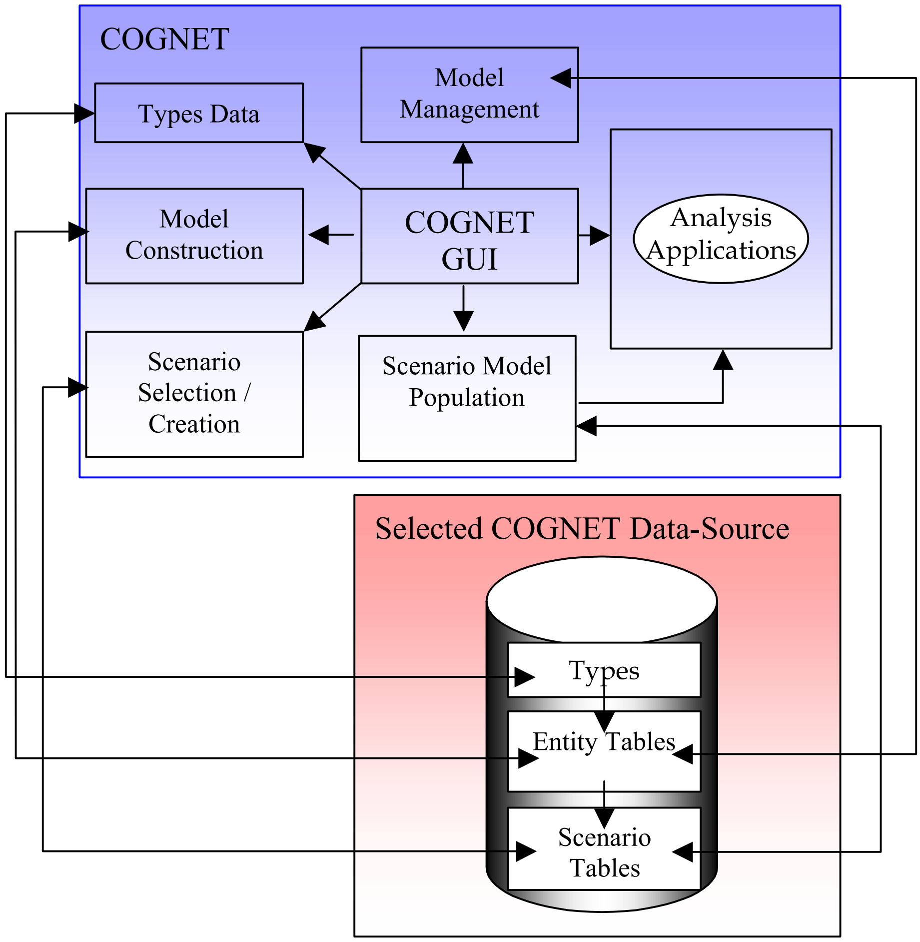

Figure 4: COGNET System block diagram

COGNET includes a graphical user interface (GUI), which allows the user to interact with the database easily and efficiently. Figure 5 shows a typical screen shot from the COGNET GUI when a user is interacting with a chosen data-source. The user can choose to look at the entire data-source without filtering; or choose to filter by country (using the country dialogue box); or by scenario (using the scenario dialogue box). The Bayesian network shown in Figure 5 represents an entire data-source without filtering (hence the scenario field is <Unspecified> and the country field is <All>). If the user has appropriate access they may also add entities and relationships to the generic database by creating nodes and links. This is done by right clicking on the screen to invoke a pop-up menu and selecting “add-entity” or “add-links”. Each time a user selects “add-entity” they are presented with the dialogue box shown in Figure 6 where required entity information is entered. Selecting “add-links” invokes the linking tool, which allows users to add links from node to node. All rules of Bayesian Networks are still enforced e.g. no cycles.

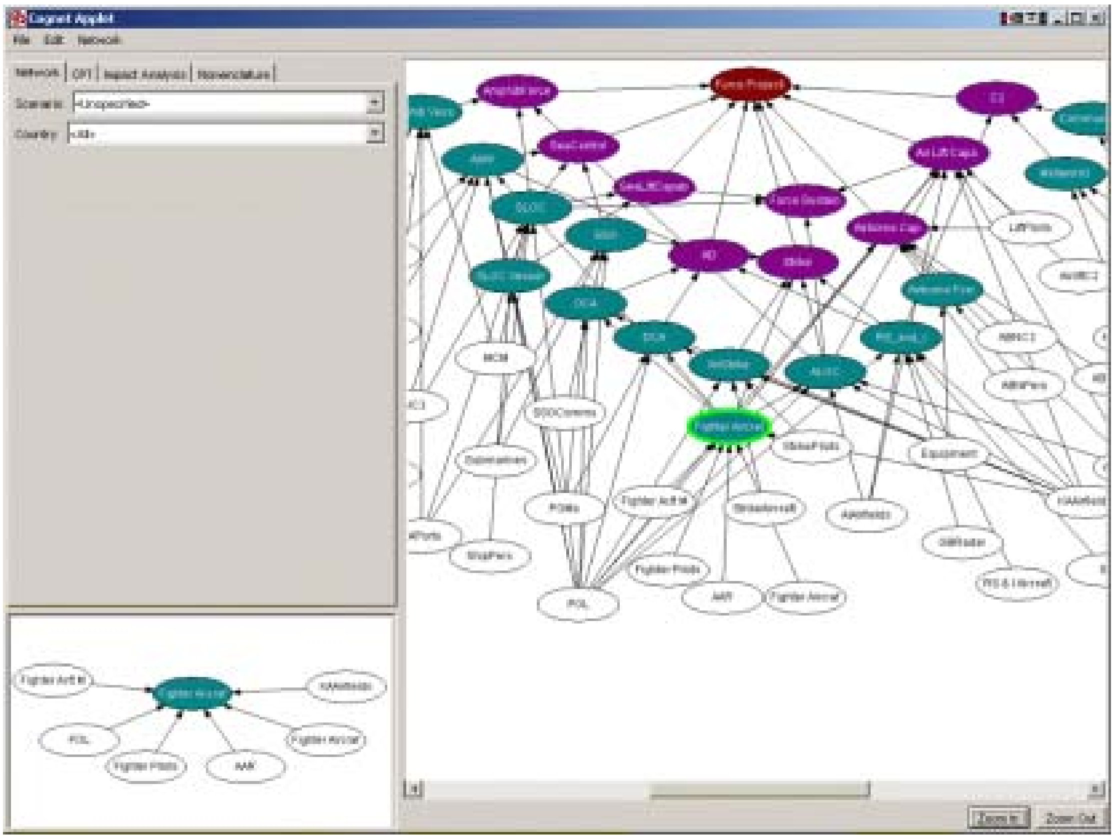

Figure 5: The COGNET GUI seen once a user has selected a data-source

Figure 6: : Window seen in COGNET when adding or editing a node.

4.2.1 Individual and group user views

The amount and type of model information required by various individuals and groups of the planning team differ in both content and detail. The commander may only be interested in the enemy and own perceived centres of gravity and the critical capabilities for those COG; while, at the other end of the spectrum, the targeting planning group may require more detailed information on a critical requirement to determine possible critical vulnerabilities. For example, consider airfields as a critical requirement. A breakdown of airfields by the targeting group, such as that shown in Figure 7, may show that the true critical vulnerability is the airfields fuel resupply system. This, however, is a low-level tactical target and would not normally be of direct interest to the Commander. To facilitate the ability for each planning group to focus on their area of interest COGNET allows expansion and contraction of parts of the network as required. Figure 8 shows the model from Figure 5 but it has been expanded in some parts and contracted in others due to the interest of the user. The highlighted green node indicates the selected node that has been expanded, in this case fighter aircraft capability. Nodes may be expanded (to show one or more levels of ancestors) or collapsed (to hide all ancestors) by double-clicking on the appropriate child-node. Right clicking on a node and selecting "show-all" expands the node and all its ancestors to show its complete structure.

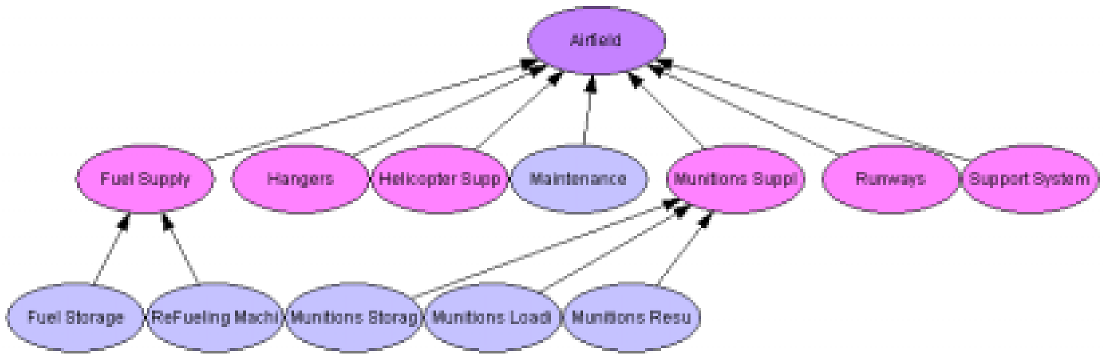

Figure 7: Example view of a breakdown of Airfields seen by a targeting planner.

4.2.2 Constructing a scenario

The COG for a particular force may change as the scenario changes but the force dependency structure will typically remain the same. This can be represented by a relatively fixed causal network, in which only the causal strengths vary according to the scenario. Using COGNET it is possible to tailor models to a particular situation using the generic framework. For example, a user may create a new scenario by importing entities from the generic database with the relevant countries' associations (a scenario may have many countries associated with it). The network then consists of several subnets of capabilities and requirements from the database for the country of interest. For example, if a scenario involves Australia's Force Projection capability it can be selected and imported to the workspace from the country selection dialogue box. COGNET will then add the Force Projection node and its networked ancestors to the required scenario. In the process any duplicate nodes are deleted[1] and links are rearranged as necessary to retain the causal structure of both the imported network and the network under construction. The user can subsequently scan each of the capabilities and/or entities and delete the nodes that are irrelevant to that scenario as they wish. In addition any other nodes that are only influenced by these deleted nodes are also automatically deleted. Alternatively a user may choose to construct a new network by selecting separate parts of one or more generic models and adding them to the new scenario, which can be saved as a new network for population and analysis. Data on dependence parameters of any subnets imported from the database may be re-examined, in light of the specific problem at hand, in the conditional probability table (CPT) generation tool, which is invoked through the CPT tab (Figure 8).

[1] At present duplicate nodes are assumed to have identical names. Future research will focus on naming conventions to simplify this process and on sophisticated database management techniques to detect multiple entities by other attributes.

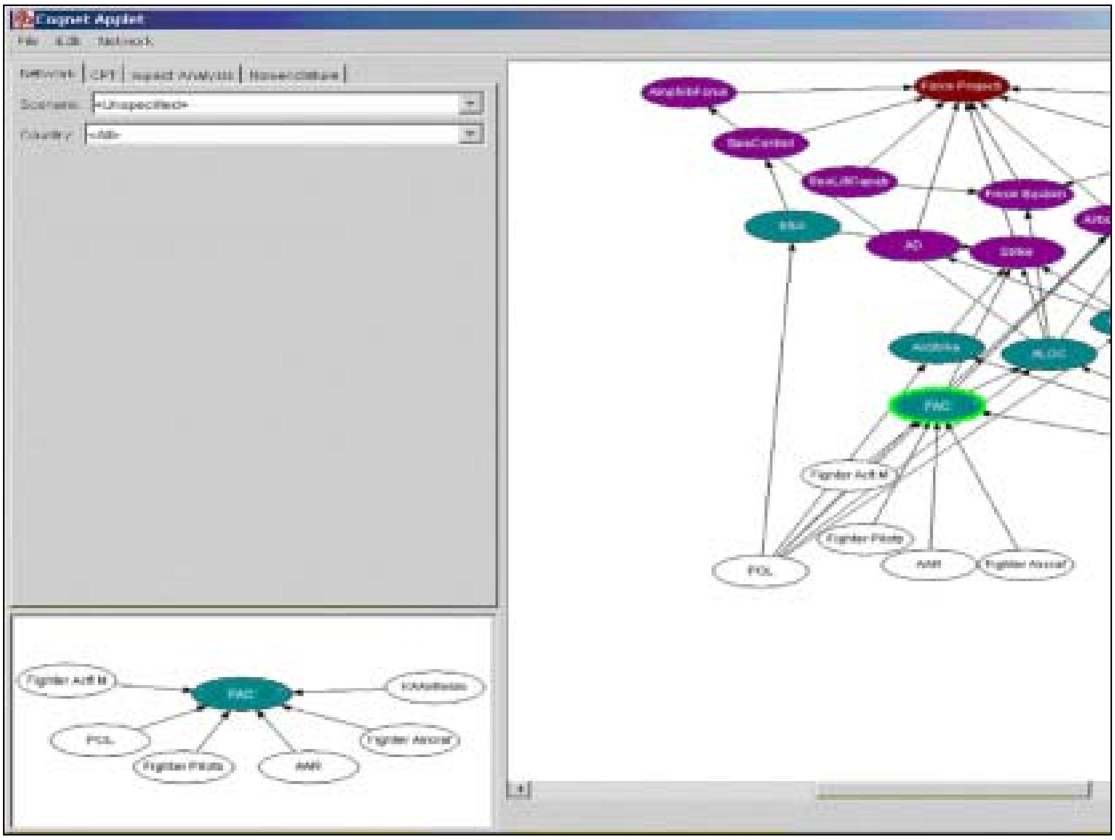

Figure 8: Demonstrating the ability to expand and contract in COGNET

4.3 Generating the conditional probabilities

Of course, analysing the COG Bayesian networks is not possible until the model has been populated, which involves specifying conditional probability tables (CPT). These probabilities can, in principle, be extracted from historical data or elicited from a domain expert. So far, we have solely relied on the judgement of domain experts to populate COGNET models since there are no appropriate or valid historical data available. However, this entails a large cognitive workload resulting in potentially inconsistent models. If we consider a two-state node with six parents, each having two states, then sixty-four conditional probabilities must be specified for just this network fragment. For nodes with more parents, and possibly more states per node, the number of distributions to be subjectively assigned by the user increases exponentially. Research into this problem has resulted in the development of a tool with a high-level user interface and built-in algorithms that generate a CPT with much less cognitive effort, particularly for those cases in which time is critical.

In the course of our investigations we considered the following algorithms for approximating conditional probabilities. The first algorithm is an integral part of SIAM, which also performs automatic CPT generation and for much the same reasons. SIAM's CPT generator is based on a purpose-built Causal Strengths (CAST) Logic [Rosen & Smith 1996]. In SIAM, domain experts create an influence net of a complex situation or scenario and provide information about the cause-effect relationships in the model. Using the CAST algorithm, this causal strength information is then converted into quantitative approximations of the conditional probability values required for analysis of the model. The CAST algorithm significantly reduces the number of inputs required to populate these nets, as does the algorithm developed for Noisy-OR or *disjunctive interaction* models [Pearl 1988; Jensen 2001]. The latter class of models is used to represent events and their causes as binary variables. They need to satisfy two main assumptions: the first requires that if all causes of a certain event are false then so is the event itself; the second requires that all inhibitors of a certain event must be mutually independent. The assumptions required for Noisy-OR models and for the influence nets in SIAM were considered too restrictive for our purposes.

The algorithm we have adopted so far implements a weighted sum technique and is fully described in a separate paper [Das 2004]. Briefly, this algorithm reduces the cognitive load required to generate the CPT by allowing the user to consider the effects on the child node for each parent node in-turn, which reduces the input to be specified by the user. The first step in the weighted sum technique is to define a probability distribution matrix for each parent-child combination. Next, the user is asked to consider the importance of each parent node on the child node in terms of relative weights. Thus, if a critical requirement A is considered to be twice as important to capability C as critical requirement B it is assigned twice the weight. The weighted sum algorithm generates conditional probabilities from the relative weights and the probability distributions. For example, if a node has six parent nodes the algorithm reduces the number of probabilities that need to be specified from 64 (2^6) to 12 (2×6) by automatically generating 64 probabilities from the six weights and the twelve probability distributions assigned. The CPT generator in COGNET also requires the user to specify whether a set of parent nodes (possibly a singleton) is critical to the child node under consideration: a subset of parent nodes is defined to be critical to a child if the latter is totally dependent on it. In this case all conditional probability entries for which P[*child = weak|all nodes in the critical parent [set] = weak*] are automatically set to 1, regardless of the given state of the rest of the parent nodes.

This algorithm does not restrict us to three-state nodes; it can handle any number of states and the number of distributions to be specified grows linearly, rather than exponentially, with the number of states. The algorithm also makes it easier to ensure that the conditional probabilities are consistent, in the sense that two probabilities do not contradict each other: for instance, the subset of an event having a higher probability than the event itself. Once the CPT has been generated the user may then test it with the COGNET model verification tool.

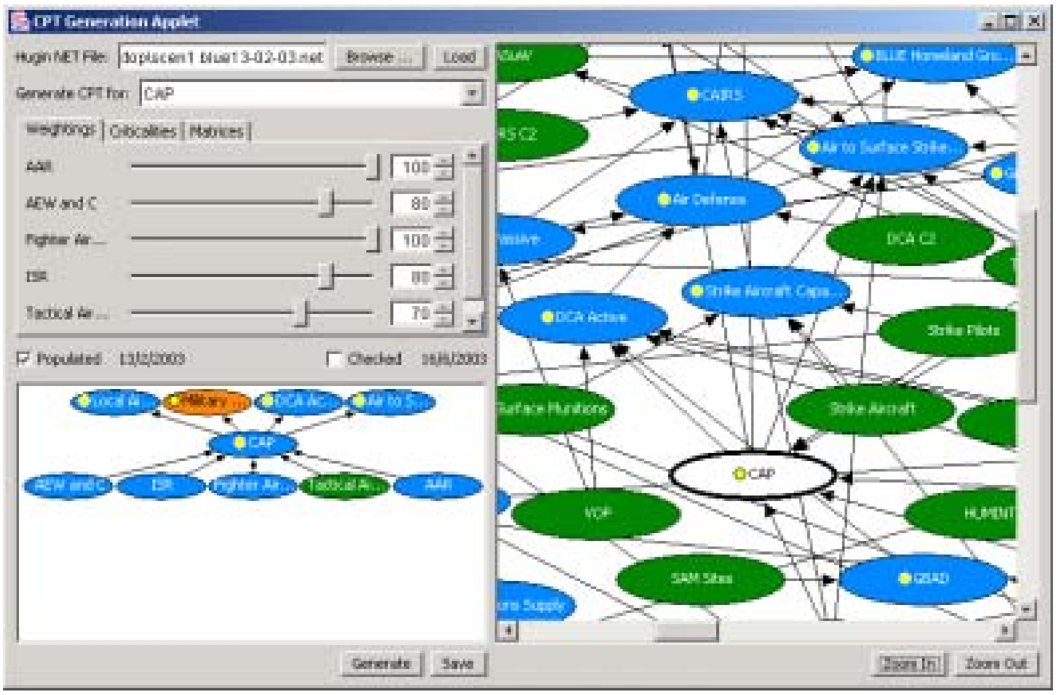

Figure 9: Example of COGNET CPT utility

The CPT generator interface used in COGNET is shown in Figure 9. When a user selects a node (such as CAP shown in Figure 9) while in the CPT generator utility a list of its parents appears in the upper left corner of the utility. Next to each parent name is a slider bar, which is used for assigning weights that are needed for populating the CPT using the algorithm described above. In addition to the weights the user must also assign the probability distributions and criticality information, also required for population, under each of the tabs next to the weights tab (see Figure 9). Once a user has populated the fragment they can then tick the populated check box, which will change the indicator dot on the node from red to yellow (as can be seen in Figure 9). This gives the user a quick check as to which fragments have been populated and those that have not. Once all fragments are populated model evaluation can be undertaken.

4.4 Model evaluation

Model checking is an essential part of Bayesian analysis. It is the general procedure for comparing final conclusions with our knowledge of reality. There are two aspects of a Bayesian net that need to be assessed. First we need to be confident that the topology is a true representation of the system being modelled. Are all significant variables present and have all relevant relationships been captured? Once we are confident of the model structure we need to determine the sensitivity of the marginal probabilities to changes in the initial distributions. It is important to investigate the level of accuracy of the numbers required for a probabilistic model and the extent to which their inaccuracy affects its output. We would also like to know whether the probabilities assigned to these nodes reflect reality. The latter might have to be a subjective assessment based on the judgement of a domain expert.

Initial probabilities describe our current belief of the likelihood of each possible state of the initial or leaf nodes. In order to understand how these numbers contribute to our probability of interest we would like to be able to rank these nodes in terms of the impact they might have on the probability of interest. This can be done in COGNET by invoking the relative impact analysis function. Obvious errors of judgement or inconsistencies should be discovered if the ranking order defies understanding, for example, if the relative impact analysis assesses ground-based radar to have a much larger effect on the red COG than POL in our scenario. A better understanding of sensitivity to inaccuracies in a probabilistic network helps determine the sets of probabilities that are most sensitive. The relative sensitivity of these numbers can help the modeller decide how much effort is worth investing into probability assessment [Henrion *et. al.* 1996].

The empirical approach to sensitivity analysis of a probabilistic network entails varying the probabilities systematically and recording the affects on a particular probability of interest [Coup, et al. 2000]. We may also take an analytical approach and investigate sensitivity by expressing the probability of interest as a function of the parameter being varied [Kipersztok and Wang, 2001]. Such an approach is illustrated in [Laskey, 1995], in which *sensitivity values* are defined as the partial derivatives of the probability of interest with respect to the conditional probabilities being varied. These values measure the sensitivity of the model outputs to small changes in model inputs, thus enabling automatic computation of model sensitivity.

Another analytic method is illustrated in [Castillo *et. al.*, 1997]. They represent initial probabilities as well as conditional probabilities as symbolic parameters instead of actual numbers and present an algorithm for expressing the probabilities of interest as algebraic functions of these parameters. They present an earlier result that the prior marginal probability of any node in the network is a polynomial function of the parameters. The degree of the polynomial is less than or equal to the minimal number of parameters or nodes but it is a first-degree polynomial with respect to each parameter. Once new evidence is injected into the network, the posterior marginal probability of any node is a ratio of two polynomial functions of the parameters such that the denominator is the same for each node and does not need to be explicitly computed each time. They exploit the dependency structure of the network to eliminate those variables that do not contribute to the polynomial function under investigation, thus making the computation more efficient and more feasible. They are also able to calculate lower and upper bounds for all the marginal probabilities.

These analytical methods enable an efficient analysis of the whole network compared to the computationally intensive empirical approach. However, they are not appropriate for military users, whose expertise is in the system being modelled rather than the underlying mathematics. Our approach aims to build a model-checking tool that exploits the structure of COGNET models and can be easily used by a decision maker who is not necessarily an expert in Bayesian modelling techniques. It performs empirical rather than analytical sensitivity analysis and is better described as an informal process of interactive verification according to the modeller's intuition rather than a systematic evaluation of the model's statistical properties [Gill, 2002]. This is justified by the fact that these probabilities are determined from expert judgement rather than statistical data. Although it would be helpful to have a better understanding of the relative sensitivity of the conditional probabilities our motivation for model evaluation is to increase the users' confidence in their model and to make the consequences of their subjective assumptions clear. We aim to encourage military users to verify that analysis results match their knowledge of the environment being modelled.

The model checker is invoked through the COGNET user interface. The user steps through the model, in a top-down or bottom-up approach, populating each fragment (a child node and its parents). Once the CPT of a fragment is complete the model checker performs relative importance analysis on the fragment with the node of interest set as the child node. If the analysis results do not match the user's judgement, they may pursue iterative refining of the weights, dependency matrices, critical node information or individual conditional probability table entries until the model verification output is satisfactory. The software requires the user to manually register the fact that each fragment has been populated and checked before impact analysis of the whole network is allowed.

4.5 Impact analysis

Although constructing a graphical representation of the COG and its underlying capabilities and requirements helps to clarify thinking and provides a useful way to record this thinking, the real power of using Bayesian nets for COG analysis lies in the resulting capability to perform impact analysis, that is, to determine the actions that would achieve a desirable end-state. In this case a desirable end-state is to degrade the enemy's COG while protecting our own.

Impact analysis allows the user to investigate the potential impact that the modelled capabilities have on the enemy (or own) COG. COGNET models show the causal link between the result of an action, such as a degraded military capability, and the resulting effect on the state of the COG. This provides a measure of effectiveness of planned actions as well as an assessment of possibly undesirable side effects. We define such a measure of effectiveness as the change in the marginal probability that the COG is in its strongest state, as a result of the change in state of any one of the leaf nodes. For example, in our example network of Figure 3, we might compare the probability, *P(COG = "strong" | all leaf nodes are "strong")* to *P(COG = "strong" | POL = "degraded" and all other leaf nodes are "strong")* and consider this as the impact of POL on the COG. Similarly, an assessment of the weakening of the friendly COG as a result of the same action by the friendly force can be made in the same way. For example, as a result of striking enemy fuel stocks the friendly force might lose one or more strike aircraft. In order to measure the effect this would have on the friendly COG we might compare *P(Friendly COG = "strong"|all leaf nodes are "strong")* to P*(Friendly COG = "strong"|Strike Acft Capability = "degraded" and all other leaf nodes are "strong").*

The current version of COGNET, which has been designed on the basis of our observations of operational-level planning, assumes that the node states are ordered from strongest to weakest (or vice-versa) and the analysis is based on this assumption. Future requirements might necessitate implementation of more flexible analysis to accommodate other users (e.g. intelligence analysts or target systems analysts). The COGNET impact analysis tool currently provides the following types of analyses.

Base case analysis: Propagate initial distributions and/or available evidence and observe the base case probability values for the node of interest, which may be the COG or a capability node of particular interest. Alternatively, in order to plan for a worst-case scenario, assume all leaf nodes in the adversary COG model are in strongest state and let this be the case for analysis. The user may stipulate what the base case should be.

"What if" analysis: Conduct exploratory testing through "what if" analysis by instantiating selected leaf nodes to the weakest state (or any other state as required) and observing how the probabilities of the higher-level capabilities and the COG itself change, as shown in Figure 10. This is currently the most flexible type of analysis but also the most tedious, particularly for large networks.

Evidence-based analysis: Instantiate leaf nodes with new evidence and re-calculate distributions of all other nodes. If evidence on an intermediate node is available this is also instantiated but all links from this node to its ancestors are deleted so that it becomes a leaf node.

Relative impact analysis: Instantiate each leaf node to its weakest state in turn, observe effect on node of interest when compared to the selected base case results.

This analysis has been automated in COGNET so that a user can generate a list of leaf nodes ordered by the potential effect on the node of interest. Figure 11 shows an example of relative impact analysis results generated by COGNET. In this example, COG is the node of interest and the base case assumes that the leaf nodes are initially in their strongest state. The bar chart to the right of Figure 11 shows the initial nodes that have the greatest effect on the COG. The operator, after analysis, has the option of saving the results to file for presentation.

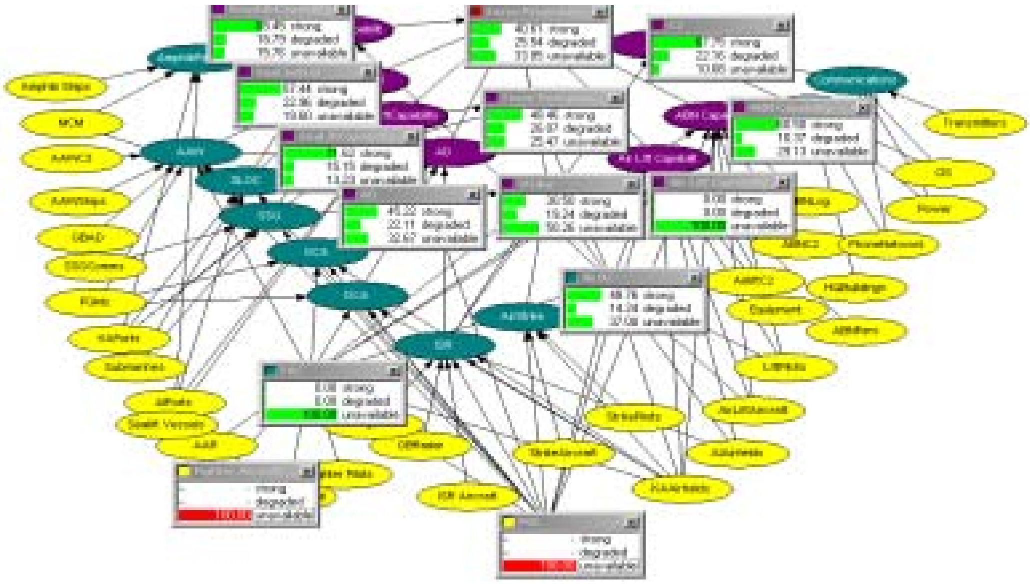

Figure 10: The effects of targeting certain critical requirements

The following sections describe the analysis algorithm currently implemented in COGNET. In the techniques described below, the user must first select the node they wish to impact, which may be the COG itself. The user must then select the leaf nodes that they want to investigate. The selected nodes being the ones they wish to analyse in terms of the impact they are likely to have on a particular node, which we will call the impacted node.

4.5.1 The relative impact analysis algorithm.

The analysis shown in Figure 11 was obtained using a relative impact analysis technique, which is undertaken by the following method:

1) All leaf nodes are instantiated to their strongest state.
2) The marginal probabilities for the impacted node are re-calculated to find the initial base case probability $\mathbf{B_r}$. This gives the probability of the impacted node being in its strongest state.
3) All selected nodes are then, in turn, instantiated to their weakest state and all other nodes left in the previous settings of step 1.
4) The probabilities are then propagated through to find the revised probability, $\mathbf{N_i}$, of the impacted node due to the change in the selected node i. The impact is calculated from the % difference from the base case i.e.

$$\text{Relative impact resulting from degrading node } i = \frac{B_r \quad N_i}{B} \times 100 .$$

5) The algorithm continues through all selected nodes in this way until all have been analysed and then the results are sorted according to the relative impact.

The relative impact analysis algorithm gives a best-case scenario for effect on possible targeted nodes. That is it gives the largest possible impact that can be achieved against a state of greatest strength.

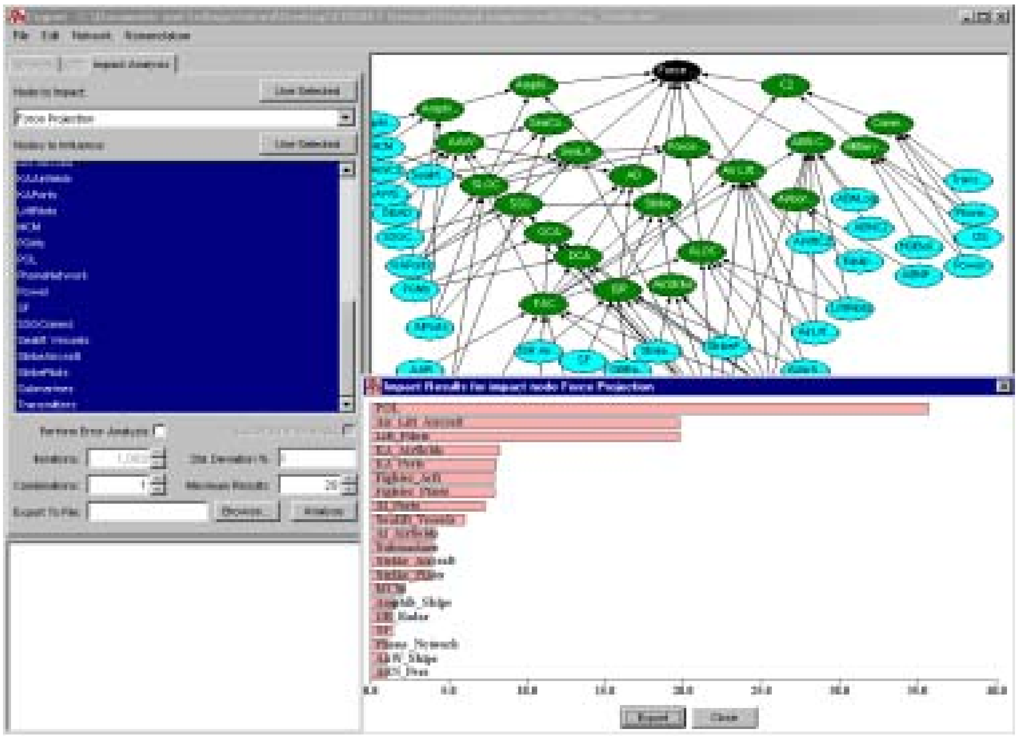

Figure 11: The relative effect of targeting certain critical requirements

4.5.2 Dealing with unknown entities and validated data

The relative impact analysis algorithm is useful at the initial planning stages - it is, in fact, a type of automated "what-if" analysis, in which one assumes that all adversary capabilities are at their strongest at the beginning of a conflict. Once evidence is available, whether this takes the form of uncertain evidence based on an analyst's subjective beliefs or hard evidence from reliable intelligence feeds or Battle Damage Assessment (BDA), they can be injected into the leaf nodes in order to determine the marginal probabilities of all the other nodes in the network. In cases where evidence or estimates are not available or are at best sketchy, analysis with unknown and validated states may be of use. However, if the evidence for an initial node is unknown then the probability mass is apportioned among its states. For example in a two-state system if a user selects a node and tags (by right clicking on the node) it as unknown the probability would be set to 50% strong/50% weak. In future we could consider adding an "unknown" state to each node in order to represent explicitly a lack of intelligence on the state of a node [Das and Davies 2002]. If evidence for an initial node is validated the user can also tag the node with validated evidence e.g. "weak" or "unavailable". This will then set that node to that state for the complete analysis.

4.5.3 Impact analysis with combinations

Impact analysis using singular nodes for analysis enables a user to answer "what if" questions regarding the degrading of singular nodes. The user can then use the results that are output to formulate a possible line of operations. However, while this approach allows the user the confidence that they will indeed (under their reasoning) degrade the COG by this possible line of operation it does not inform the user of the outcome of targeting combinations of nodes. Let us consider a simple example. If the impact analysis results indicated that the two entities that had the most effect on the COG were fighter aircraft and fighter pilots as they both contribute to fighter aircraft capability then it would be redundant to target both these entities in a line of operation, as it would have the same effect as targeting them singularly. Therefore, the COGNET impact analysis utility also incorporates the ability to analyse entities in combination. This system is identical to that previously described above but simply calculates impact for all combinations for the number of entities in combination specified by the user. The results are then sorted according to the total impact. The combinations toggle box can be seen in the left hand corner of Figure 11.

4.5.4 Impact analysis with intermediate states (or degraded states)

The relative impact analysis techniques described thus far also allow for investigation with intermediate states. To allow more flexibility in analysis with intermediate states COGNET allows the user to define, once they have selected one of the above algorithms, the state that they wish to use for analysis. For example if a model uses the three state system of "strong", "degraded" and "unavailable" the user may select to run the impact analysis using the degraded state instead of the unavailable state. Of course, this assumes that all nodes in the network have a uniform number of identically named states.

5. Compiling Large Complex Networks

5.1 Complex COGNET Networks

The Bayesian engine component of most software packages for Bayesian networks relies on efficient inference algorithms for probability updating. There are two classes of inference: exact and approximate. We are interested in the former, which (in principle) enable us to calculate marginal probabilities for each variable in the network. HUGIN uses an exact algorithm based on the construction of a junction tree of the triangulated graph derived from the directed acyclic graph representing the Bayesian network. These algorithms are complex and computationally intensive but are currently the most efficient method for exact probability updating in Bayesian networks. Other exact algorithms exploit algebraic schemes for variable elimination without using graph theoretic concepts. However, these algorithms can only compute marginal probability values for a given subset of variables [Cozman 2000].

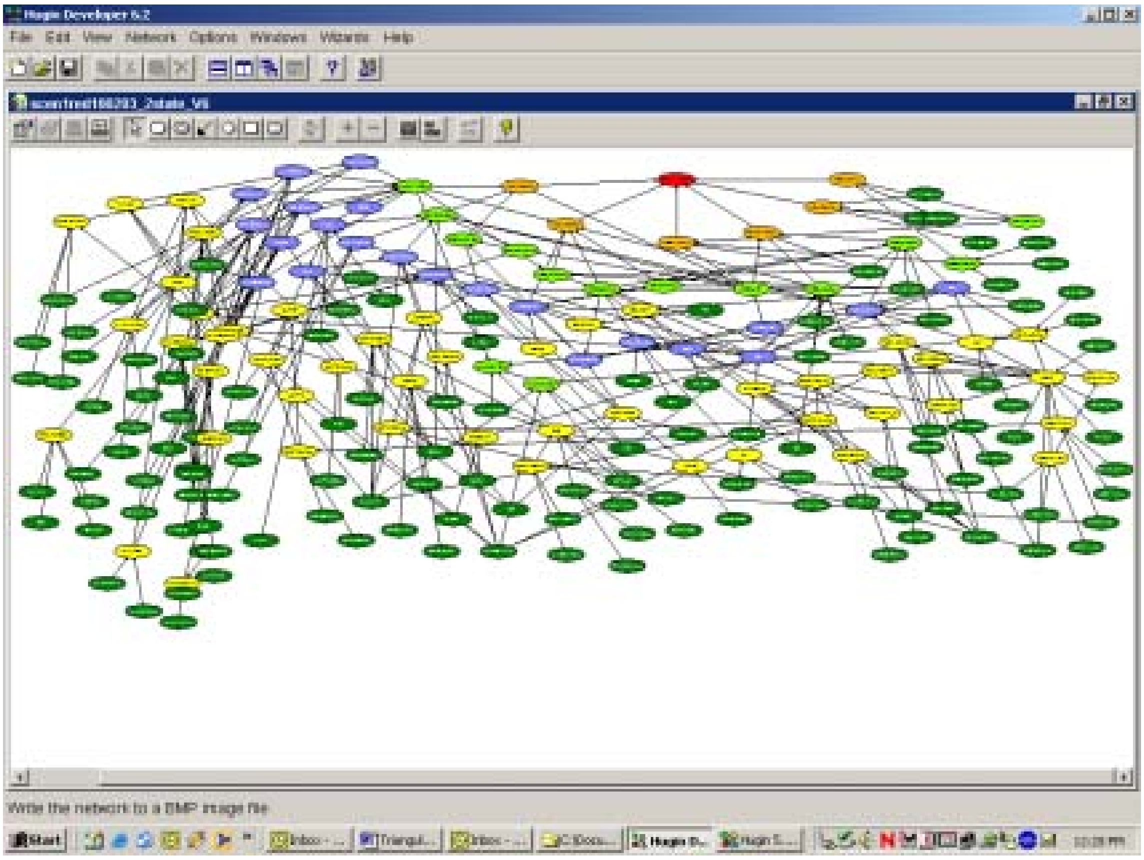

Figure 12: A more comprehensive COGNET model

The integration of generic subnets into a COGNET model makes them comprehensive but very large and complex, significantly larger than the networks we have shown so far. Although there is no problem with compiling large two-state networks we experienced some compilation problems when we started converting our models to three-state nodes in order to enable more detailed analysis. The exponential increase in memory requirement has meant that some of our more comprehensive networks cannot be compiled. Such a three-state network is shown in Figure 12. Looking at the structure of this COGNET model it is not difficult to understand why: many of the nodes have 5 or 6 parents and quite a few of these parent nodes have several children. These types of networks result in large junction trees, whose nodes are composed of large cliques (maximal subgraphs, all of whose nodes are pairwise linked). The HUGIN heuristic determines that the junction tree for this network will have 123 cliques, the smallest of which has 3 members and the largest has 18 members. Memory storage space for each clique grows exponentially with the number of members and the number of states per member. As a result the total clique cost is 621,060,111, requiring approximately 2.5 Gbytes of memory. See [Falzon 2003] and references therein for a description of the theory behind junction tree algorithms and details of the HUGIN inference propagation algorithm.

Fortunately there are ways to exploit the structure of COGNET models to make compilation of large three-state networks possible, provided we are only interested in forward propagation. Using the definition of d-separation [Pearl 1988; Jensen 2001], we can see that it is quite easy to separate a typical network into independent sub-networks that can be compiled separately. The network in Figure 2, for example, can be easily separated into 4 subnetworks as follows. One subnetwork, call it the high-level subnet in the rectangular frame in Figure 13, would consist of nodes COG, CC_1, CC_2 and CC_3; the other three would each consist of one of the CC nodes and all its ancestors (one of these networks is highlighted in Figure 13). The latter three subnets are d-separated if the leaf nodes are the only ones instantiated. Once these nodes have been instantiated, the three subnets can be compiled separately giving the marginal probabilities of the CC nodes, which are the bottom-level nodes of the high-level subnet; and forward propagation determines the marginal probabilities of the rest of the network. Larger, more complex networks can be handled in the same way provided the rules for d-separation are upheld.

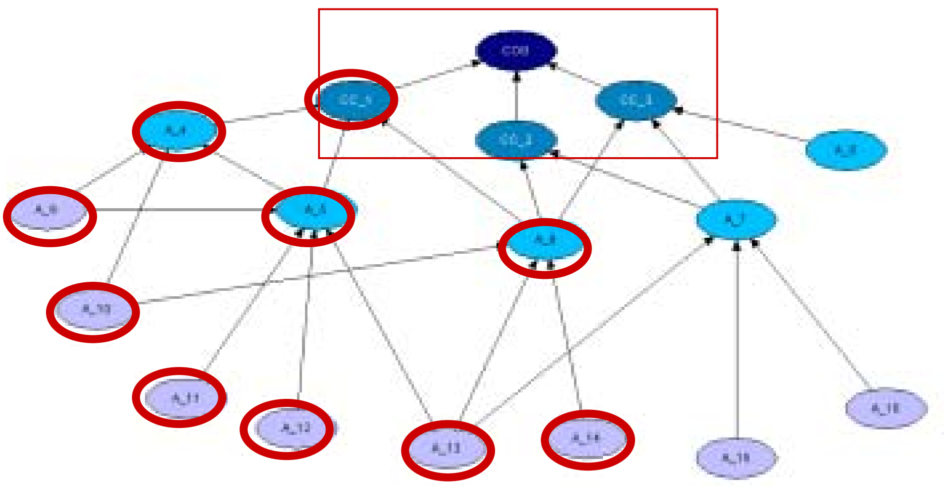

Figure 13: Compiling separate networks

Another aspect of COGNET models we can exploit is the fact that these are essentially causal networks, in which "interventions" can be represented by deleting links from the node whose state is altered by the intervention to all its parents. From the definition of Causal Bayesian Networks [Pearl 2001, Section 1.3.1], since the underlying relationships are causal (e.g. if a critical capability is "destroyed, captured or neutralised it will significantly undermine" the COG [ADDP 5.0.1 2002]), we can represent external or spontaneous changes as follows. Suppose we wish to investigate how neutralising the critical requirement represented by A_7 impacts on the COG. We represent this in the network by instantiating A_7 with this evidence and effectively deleting the links from A_7's parents to it; thus turning A_7 into a leaf node. By doing so we are reflecting the fact that since we know the state of A_7 we are no longer interested in the effect other nodes might have on it. We propagate this evidence through the rest of the network to see the effect on the nodes of interest, typically the COG.

As explained earlier, critical capability analysis attempts to determine the potential impact that the lower-level nodes have on the higher-level nodes rather than attempting to infer the converse. In other words, leaf node probabilities or evidence is propagated through the network to the nodes of interest higher up in the hierarchy. This type of analysis does not require inference to be made about the states of the leaf nodes from evidence gathered about high-level capabilities. For this reason the compilation methods described above are adequate. There might be other sorts of analyses, perhaps during the execution of a plan, in which such inference is required, mandating backward propagation of evidence. In this case the full capability of the HUGIN inference engine can still be used, provided the network size is suitably decreased, either by deleting nodes and links or by decreasing the number of states per node. Both these adjustments can be made easily in COGNET.

6. Conclusion

A thorough understanding of the relationships between a COG and its underlying critical capabilities and requirements is crucial to the development of a sound military plan. The relationship structure is often complex and not always easy to determine. The COG Network Effects Tool goes a long way to facilitate this task and it provides an effects-based analysis capability. COGNET uses existing planning process concepts to create a knowledge framework, which fits directly into ADF planning doctrine. The network representation of COG analysis facilitates reasoning and enhances shared understanding of complex situations. In addition, probabilistic models ensure that uncertainties and subjective judgements are clearly represented. It provides a generic models database to facilitate the construction of comprehensive models and enable future re-use and traceability and has provision for compiling large complex networks. Its graphical user interface and the suite of tools it provides (for CPT generation, impact analysis and model checking) have been specially designed for military users, whose expertise is in the system being modelled rather than the underlying mathematics.

An important aspect of this work is integrating COGNET into the ADF planning process [Priest *et al.* 2002]. Operational planning instructors from the ADFWC, as well as practising operational planners, have been involved, since the inception of the project, in the conceptual development in order to ensure that the philosophy underlying the models is compatible with ADF doctrine and the planning process. The instructors have formed a major part of our domain expert pool, and have contributed to populating the generic models. We have also involved operational planning students during week-long planning exercises. Apart from the obvious benefit of having future planning officers become familiar with COGNET, trials of the tool also provide us with instant feedback of the performance and utility of the tool in realistic conditions and contribute to its evolution. During these trials the students are able to validate their analysis and review their lines of operation obtained through the JMAP. Observations demonstrate that use of the tool reinforces the requirement to maintain the link between critical capability analysis and course of action development. Introduction of COGNET to real-world planning groups is programmed once the robustness of the tool is validated over time. COGNET has been designed with operational-level concepts in mind, and therefore assumes that nodes and node states represent high-level capabilities. Future trials with Target Systems specialists and Information Operation planners will no doubt lead to further modifications as required. However, the conceptual foundation of the tool as described in this report has thus far withstood significant intellectual scrutiny.

Future research will include: linking results from the critical capability analysis conducted during mission analysis to the next step in the planning process in which COAST is used to derive a feasible line of operations from a sequence of tasks; using results from impact analysis in COGNET to measure the potential effectiveness of a sequence of tasks for plan refinement in COA-Sim; and deriving probabilities to populate COGNET models from results obtained in simulation runs of COA-Sim. These research proposals are based on the well-founded premise that the conceptual framework underlying COGNET and the other planning tools in In-MODE all serve to clarify the logical link between the essential tasks and the defined objective, defined as negating the threat COG while maintaining our own.

7. Acknowledgements

The authors wish to acknowledge the support and assistance provided by instructors of the Doctrine and Training Wing, ADF Warfare Centre. Members of DSTO's SSA Group particularly Balaram Das, Mike Davies, Melissa de Poel and Lin Zhang provided invaluable intellectual support and feedback.

8. References

Australian Defence Doctrine Publications, Plans Series (2002) *Joint Planning*, ADDP 5.0.1, Australian Defence Headquarters, Canberra, September 2002.

Barclay, S., Brown, R.V., Kelly, C. III, Peterson, C.R., Philips, L.D., Selvidge, J. (1977) *Handbook for Decision Analysis*, AD-A049-221 N00014-76-C-0074 TR-77-6-30, Advanced Decision Technology Program Office of Naval Research.

Castillo, E., Gutiérrez, J.M. and Hadi, A.S. (1997) "Sensitivity analysis in discrete Bayesian Networks", *IEEE Transactions on Systems, Man, and Cybernetics*, **27**(4) 412-423.

Coupe, V.H.M, van der Gaag, L.C. and Habbema, J.D.F. (2000) "Sensitivity Analysis: an Aid for Belief-network Quantification", *The Knowledge Review*, **15**(3).

Cozman, F.G. (2000) Generalizing variable elimination in Bayesian networks, *Workshop on Probabilistic Reasoning in Artificial Intelligence,* Atibaia, Brazil, November 2000.

Das, B. 2004. Generating Conditional Probability Tables for Bayesian Networks: Easing the Knowledge Acquisition Problem, to be published.

Das, B. and Davies, M. (2002) "Causes-Effects-Consequences Framework for the Analysis of Foreign Command and Control Systems", DSTO research report, DSTO-RR-0229.

de Poel, M., Priest, J. and Falzon, L. (2002) "The Applicability of the Situational Influence Assessment Module (SIAM) to the Centre of Gravity Network Effects Tool", DSTO Technical Note, DSTO-TN-0466.

Druzdzel, M. and van der Gaag, L.C. (2000) "Building Probabilistic Networks: Where do the numbers come from?", *IEEE Transactions in Knowledge and Data Engineering*, **12**(4) 481-486.

Echevarria, A.J., II (2003) "Clausewitz's Center of Gravity: It's Not What We Thought", *Naval War College Review,* Winter 2003, 108-123.

Evans, A., Graham, S., Jones, E.K., Pioch, N., Prendergast, M. and White, C.M. (2003) "Strategy Development for Effects-Based Planning", Military Operations Research Society official website, http://www.mors.org/meetings/ebo/ebo_read.htm, last accessed June 2003.

Falzon, L., Zhang, L. and Davies, M. (2000) "A Policy Analysis Approach to Operational Level Course of Action Analysis", *Proceedings of the 5th Command and*

Control Research and Technology Symposium, Australia War Memorial, Canberra ACT, Australia, October 2000.

Falzon, L., Zhang, L. and Davies, M. (2001) "Hierarchical Probabilistic Models for Operational-Level Course of Action Development", *Proceedings of the 6th Command and Control Research and Technology Symposium*, Annapolis, USA, June 2001.

Falzon, L. (2003) "The HUGIN algorithm for inference updating in Bayesian networks", DSTO discussion paper, DSTO-DP-0848.

Falzon, L. (2004) "Using Bayesian Network Analysis to Support Centre of Gravity Analysis in Military Planning", accepted for publication in the *European Journal of Operational Research.*

Giles, P.K., and Galvin, T.P. (1996) " Center of Gravity – determination analysis and application", Center for Strategic Leadership, U.S. Army College, Carlisle Barracks, PA.

Gill, J. (2002) *Bayesian Methods: A Social and Behavioural Science Approach*, Chapman & Hall/CRC, Boca Raton.

Henrion, M., Pradhan, M., Del Favero, B., Huang K., Provan, G. and O'Rorke, P. (1996) "Why is diagnosis using belief networks insensitive to imprecision in probabilities?" *Proceedings of the Twelfth Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann Publishers, San Francisco, CA.

HUGIN Expert Limited, URL - http://www.hugin.com.

Jensen, F. V. (2001) *Bayesian Networks and Decision Graphs*, Springer-Verlag, New York.

Kidd, M. (2002) "Applying Bayesian belief networks as a tool for structuring and evaluating the planning of naval operations", " *Military Operations Research,* **7**(4).

Kipersztok, O. and Wang, H. (2001) "Another look at sensitivity of Bayesian networks to imprecise probabilities", *Proceedings of the Eighth International Workshop on Artificial Intelligence and Statistics,* Morgan Kaufmann Publishers, San Francisco, CA.

Laskey, K.B., 1995. Sensitivity analysis for probability assessments in Bayesian networks, IEEE Transactions on Systems, Man, and Cybernetics, 25 (6), 901-909.

Matthews, K. and Davies, M. (2003) "Simulations for Joint Military Operations Planning", *Proc SimTecT03*, Adelaide Australia, 26-29 May 2003.

Morgan, M.G., and Henrion, M. (1990) *Uncertainty: a Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*, Cambridge University Press, Cambridge.

Paté-Cornell, E. and Guikema, S. (2002) "Probabilistic Modelling of Terrorist Threats: A Systems Analysis Approach to Setting Priorities Among Countermeasures" *Military Operations Research,* **7**(4).

Pearl, J. (1988) *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann Publishers, San Francisco, California.

Pearl, J. (2001) *Causality: Models, Reasoning and Inference*, Cambridge University Press, Cambridge.

Priest, J., Smallwood, R., Falzon, L., Zhang, L. and Lumsden, S. (2002) "A Centre of Gravity Analysis Tool to Support Operational Planning", *Proceedings of the 7th International Command and Control Research and Technology Symposium*, Canada.

Rosen, Julie A. and Smith, Wayne L. (1996) "Influence Net Modeling with Causal Strengths: An Evolutionary Approach", *Proceedings of Command and Control Research and Technology,* National Defense University, Monterey, California, June 1996.

Rosen, Julie A. and Smith Wayne L. (2000) "Influence Net Modeling for Strategic Planning: A Structured Approach to Information Operations", Phalanx, Vol. 33(4).

Saaty, T. L. (1994) *Fundamentals of Decision Making and Priority Theory with The Analytic Hierarchy Process*, RWS Publications, Pittsburgh, PA.

Saaty, T. L. (1996) *Decision Making with Dependence and Feedback: The Analytic Network Process*, RWS Publications, Pittsburgh, PA.

Sage, A. P. (1981) "Behavioral and Organizational Considerations in the Design of Information Systems and Processes for Planning and Decision Support", *IEEE Transactions on Systems, Man and Cybernetics*, **SMC-11** (9).

Zhang, L., Falzon, L. Davies, M. and Fuss, I. (2000) "On Relationships between Key Concepts of Operational Level Planning", *Proceedings of the 5th Command and Control Research and Technology Symposium*, Australia War Memorial, Canberra ACT, Australia, October 2000.

Zhang, L., Kristensen, L.M., Janczura, C., Gallasch, G. and Billington, J. (2002) "A Coloured Petri Net based Tool for Course of Action Development and Analysis", *Workshop on Formal Methods Applied to Defence System*s, Conferences in Research and Practice in Information Technology, **12,** Adelaide, Australia, June 2002.

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION
DOCUMENT CONTROL DATA

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA	1. PRIVACY MARKING/CAVEAT (OF DOCUMENT)
2. TITLE The Centre of Gravity Network Effects Tool: Probabilistic Modelling for Operational Planning	3. SECURITY CLASSIFICATION (FOR UNCLASSIFIED REPORTS THAT ARE LIMITED RELEASE USE (L) NEXT TO DOCUMENT CLASSIFICATION) Document (U) Title (U) Abstract (U)
4. AUTHOR(S) Lucia Falzon and Jayson Priest	5. CORPORATE AUTHOR Information Sciences Laboratory PO Box 1500 Edinburgh South Australia 5111 Australia

6a. DSTO NUMBER	6b. AR NUMBER	6c. TYPE OF REPORT	7. DOCUMENT DATE
DSTO-TR-1604	AR-013-155	Technical Report	July 2004

8. FILE NUMBER	9. TASK NUMBER	10. TASK SPONSOR	11. NO. OF PAGES	12. NO. OF REFERENCES
2004/1046515	JTW 02/196	Commandant ADFWC	45	34

13. URL on the World Wide Web	14. RELEASE AUTHORITY
http://www.dsto.defence.gov.au/corporate/reports/DSTO-TR-1604.pdf	Chief, Command and Control Division

15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT

Approved for public release

OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE, PO BOX 1500, EDINBURGH, SA 5111

16. DELIBERATE ANNOUNCEMENT

No Limitations

17. CITATION IN OTHER DOCUMENTS Yes

18. DEFTEST DESCRIPTORS

Operations research
Decision making
Military planning
Knowledge transfer

19. ABSTRACT

The centre of gravity (COG) Network Effects Tool (COGNET) uses Bayesian networks to represent the COG causal structure. Its impact analysis capability facilitates the determination of the critical vulnerabilities that have to be degraded or negated to influence the COG. COGNET provides a modelling framework and a generic model database to aid knowledge reuse and knowledge transfer. Its graphical user interface is tailored to the military user and provides a user-friendly capability for populating and interacting with the models. In this report we discuss the methodology, development and implementation of the COGNET suite. The importance of this work is that it uses existing planning process concepts to facilitate the construction of comprehensive models in which uncertainties and subjective judgements are clearly represented, thus enabling future re-use and traceability.