Hierarchical Probabilistic Models for Operational-Level Course of Action Development

Lucia Falzon, Lin Zhang and Mike Davies
Information Technology Division
Defence Science and Technology Organisation
PO Box 1500, Salisbury, SA 5108, Australia
Email: lucia.falzon@dsto.defence.gov.au

Abstract

A thorough understanding of the relationships between a centre of gravity (COG) and its underlying critical capabilities and requirements is crucial to the development of a sound military plan. In this paper we describe the concepts underlying the COG Network Effects Tool (COGNET), which uses Bayesian (or causal probabilistic) networks to investigate this relationship structure. COGNET provides a visual representation of the causal structure and provides an effects-based analysis capability, which facilitates the determination of the critical vulnerabilities that have to be degraded or negated to influence the COG. Furthermore it provides a framework, which can serve as a knowledge base representing generic causal relationships to aid knowledge reusability and knowledge transfer.

1. Introduction

The initial stage of any operational-level planning process 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 ([1], Chapter 4). Mission analysis relies heavily on input from the intelligence preparation of the battlespace, in particular intelligence on the enemy centre of gravity (COG) and the likely enemy courses of action (COA). 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* [1].

Having determined the enemy COG, planners must now generate suitable COAs. Suitability refers to whether it meets the objectives as detailed in the mission analysis step. Since directly targeting the enemy COG is not usually feasible a critical capability analysis is conducted at this stage of the planning process. A critical capability 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* [1]. Each critical capability might have a number of associated critical requirements, which are essential for it to be fully functional. These requirements may be further decomposed into critical vulnerabilities: elements that are potentially vulnerable ([1], Chapter 8). The idea behind critical capability analysis is to identify which aspects of the threat critical requirements can be targeted in order to influence the enemy critical capabilities and hence the COG.

A good understanding of the key concepts of operational art is as essential for military operational planners as it is for developers of planning support tools. The qualitative relationship between these planning concepts has already been explored [3,10]. In this paper we describe the concepts underlying the COG Network Effects Tool (COGNET), which uses Bayesian (or causal probabilistic) networks that reflect the relationships among the critical capabilities and requirements. Using this model it is possible to investigate the effect that a set of actions has on the enemy centre of gravity. The model facilitates the drafting of a COA and determination of the critical vulnerabilities that have to be degraded or negated to influence the COG. It can therefore provide the principal components to an effects-based modelling capability.

2. Bayesian Networks

Bayesian networks (BN) are graphical representations of causal relations in a particular domain. They are typically used to model a domain that has inherent uncertainty due to a combination of incomplete knowledge of the domain and randomness in the environment [4]. The network may be represented by a directed acyclic graph whose nodes correspond to random variables, which can take on two or more values, and which are linked by causal dependencies. The causal direction is represented by the direction of the arcs in the graph. Nodes that have arcs directed towards them are called destination nodes while nodes with arcs directed away from them are known as origin nodes. Internal nodes are both origin and destination nodes, whereas the nodes at the edge of the network are either root nodes (they only have arcs directed away from them) or terminal nodes (purely destination nodes).

The strength of the causal relationship is expressed as a conditional probability. Each node has associated with it a set of two or more potential values or states. The probability of being in each particular state is conditioned on the state of each of its neighbouring origin nodes. The probability distribution of a Bayesian net is specified by assigning to each root node an initial probability of being in each state and all other nodes are assigned conditional probabilities, given all possible combinations of the states of all neighbouring origin nodes. As Pearl [8] points out, the advantage of this graphical representation is that it allows a specification of direct dependencies representing the fundamental qualitative relationships. The network structure then augments these relationships with a consistent set of induced indirect dependencies, which are stable and independent of the numerical probability estimates. It is then possible to calculate the probabilities of the states of the terminating nodes each time the probability values of the root nodes change.

The numbers required for a Bayesian network are normally 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, which allows us to determine whether the model is complete and consistent. Another advantage of using Bayesian nets is that determining 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.

3. The COG Network Effects Tool (COGNET)

The causal networks we consider represent the centre of gravity (own or threat) and all the elements that influence it. Functional decomposition [3] of the centre of gravity is used to identify the influencing elements and to categorise them into a hierarchy: COG, critical capabilities and critical requirements. Evaluation of the Bayesian net enables a systematic analysis that helps to identify possible critical vulnerabilities among the critical requirements. Thus, in our model, the only terminal node represents the COG while possible critical vulnerabilities appear as root nodes at the edge of the network. Such a decomposition ensures that the direction of influence travels up the hierarchy. In other words, targeting a critical vulnerability in order to change its state produces an effect on all related elements higher up in the hierarchy and hence the COG. All the networks presented in this paper were produced in HUGIN [4], a software tool for building Bayesian networks, which forms a Bayesian engine for the COG Network Effects Tool (COGNET). The probability-updating algorithms used in HUGIN yield exact probabilities.

In addition to HUGIN, COGNET will include:

- An underlying database for the flexible management of COG element categorisation;
- A higher-level user interface that aids military end-user model construction, maintenance and interaction;
- The ability to define desired effects in terms of influence on COG;
- Software utilities for recursive sensitivity analysis and metrics;
- Alternative, high-level visualisation of COG dependencies, including three-dimensional views;
- Mechanisms for integration of COGNET within a broader spectrum of models for COA development and analysis that includes modelling and simulation tools for COA scheduling and resource allocation [11], target systems and element analysis, and synthetic environments for war gaming and rehearsal [7].

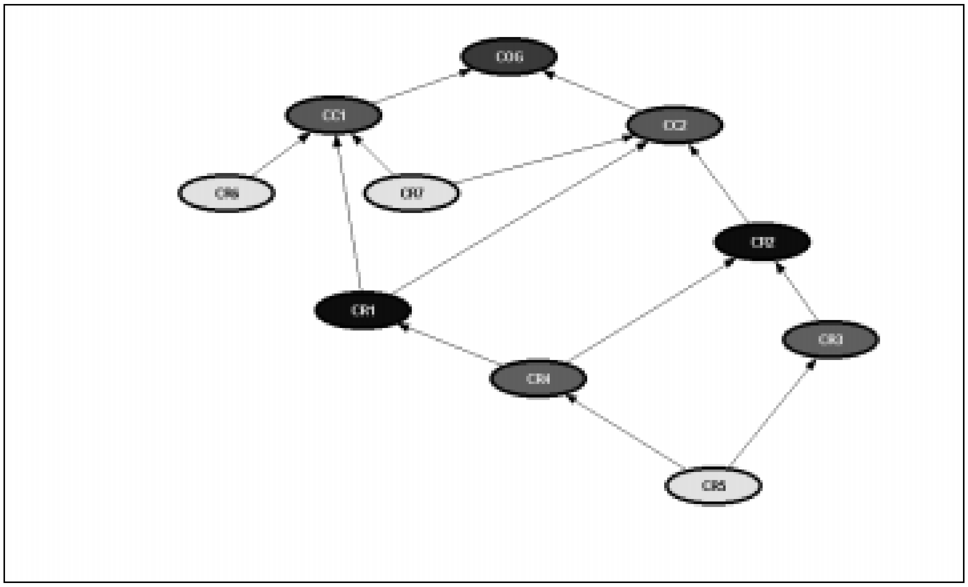

Figure 1: A simple representative COG network

The net depicted in Figure 1 represents a simple case of a COG with two critical capabilities, each of which has a number of critical requirements depending on other requirements, in turn. For logical convenience, we further decompose the critical requirements into a hierarchy ensuring that root nodes represent possible critical vulnerabilities. The node colours, as shown in Figure 2, represent the levels of the hierarchy. Thus, the critical requirements represented by the yellow nodes (i.e. CR5, CR6 and CR7) are possible critical vulnerabilities. In this case we have four levels of critical requirements as well as the root-node level, but the number of levels is arbitrary. Once the model has been populated it is possible to calculate a relative measure of the effect the state of each of the possible critical vulnerabilities has on the centre of gravity. Software utilities for sensitivity analysis will enable desired influence (effect) on COG to drive determination of a set of candidate target critical vulnerabilities.

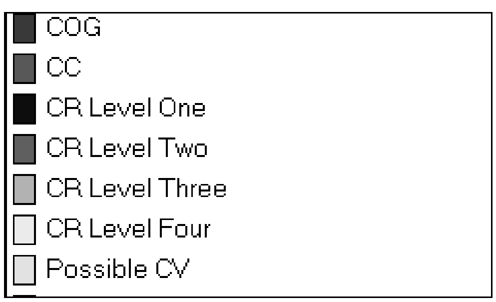

Figure 2: Node-colour/hierarchy-level mapping

Acquiring the probabilities for a Bayesian network is sometimes problematic. Take the simple example shown in Figure 1: CC2 is dependent on three critical requirements. If each of the requirements can be in two possible states, then eight conditional probabilities would have to be specified for CC2, one for each configuration of the states of CR1, CR2 and CR7. Some of these configurations may be too specific for any expert, who might be able to specify the probability of CC2 being in a certain state given the state of CR1 but might have difficulty estimating the probability of CC2 having to consider the states of all three nodes. However, in practice, conceptualisation of causal relationships can be simplified by creating hierarchies of small clusters of variables whose interactions can be categorised into prototypical structures for combining different combinations. One such structure (Pearl's Noisy OR [8]) arises when any one of the dependent requirements is likely to produce the same effect. Higher-level interfaces that facilitate elicitation of conditional probabilities are also being developed.

Let us demonstrate a Bayesian Network representation of a typical critical capability analysis. The node colours represent the same hierarchy levels as shown in Figure 2. The network shown in Figure 3 represents a COG analysis exercise observed by the authors. It is based on a fictitious scenario used for training purposes. The perceived threat was an imminent invasion of an island belonging to an ally. The assigned mission was to deter or dislodge the enemy and return sovereignty of the island to its government. The threat COG was assessed to be the enemy's ability to project force, and this was further assessed to be dependent on the threat sustainment, amphibious landing and airborne force capabilities. Associated critical requirements appear as neighbouring origin nodes. The root nodes represent the critical elements that are potential targets and hence vulnerable. Each node can take on two values, normally called strong and weak, which describe the current state of the capability or requirement corresponding to that node. Note that this model is not very extensive. There are possibly three reasons for this: a) the exercise participants were trainees – not experienced in Joint Planning; b) there were tight time constraints; c) it is a rather tedious exercise to derive a comprehensive model. However, the model can be shown to be consistent as necessitated by the formalism inherent in the modelling.

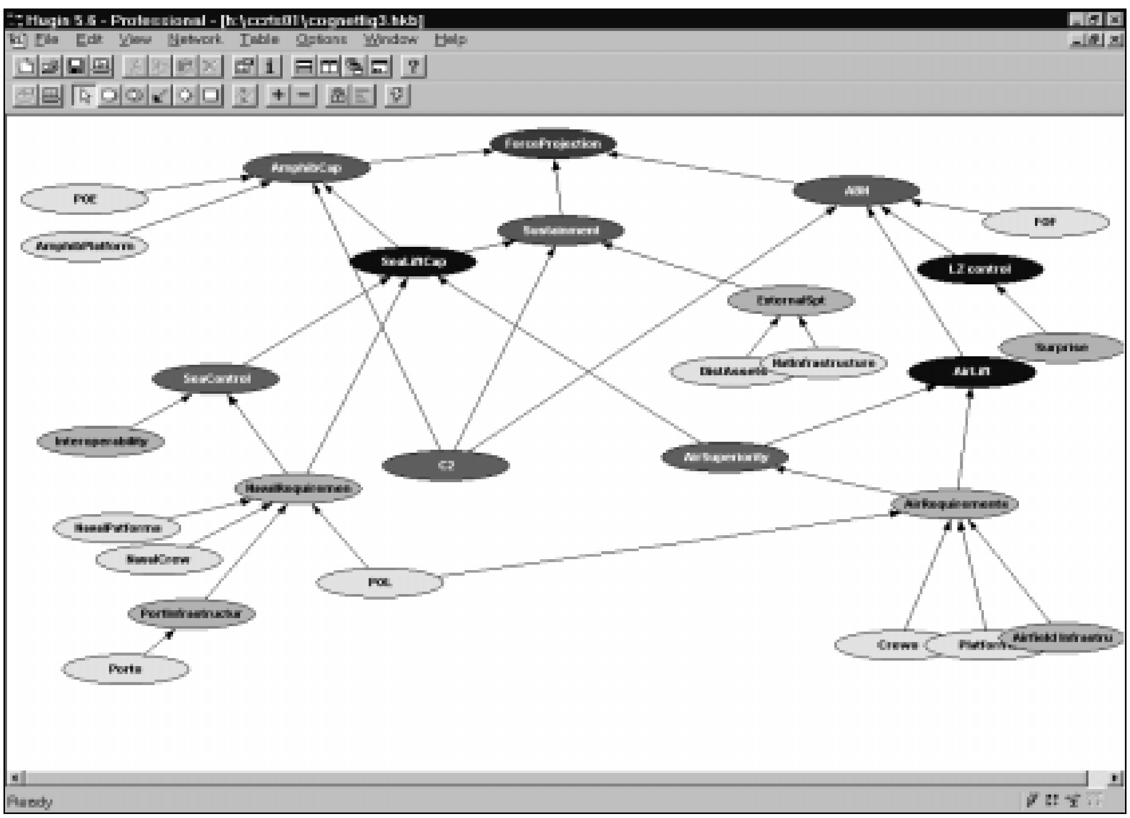

Figure 3: An example threat COG network

Figure 4 shows the initial probability distribution assigned to the network. Each root node is assigned a probability of being in each state, independent of the status of all the other nodes. All the remaining nodes have a conditional probability table defining the probability of being in each state conditioned on the states of its neighbouring origin nodes. For example, the probability that the enemy Air Requirements capability is in a strong state is conditioned on whether Petroleum Oil and Lubricants (POL) stocks, Air Crews, Air Platforms and Airfield Infrastructure are in a strong or a weak state. The conditional probability table is shown in Figure 5.

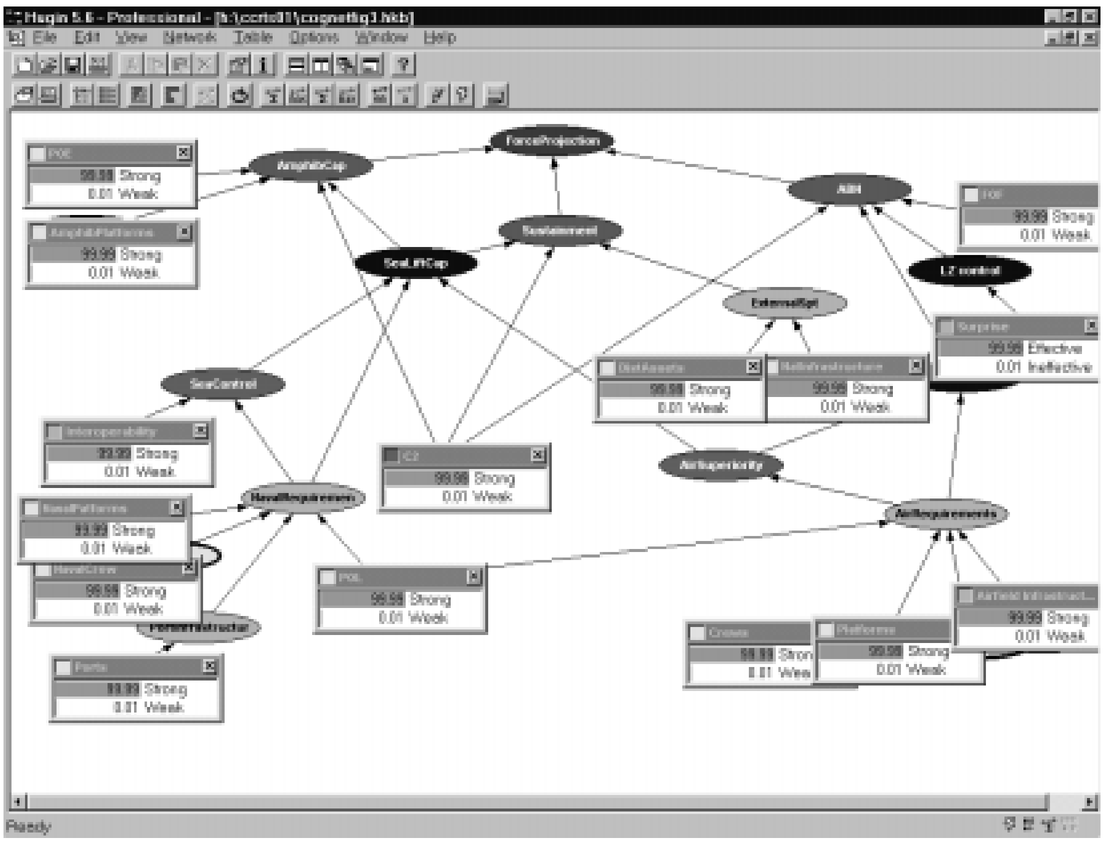

Figure 4: Initial probability distribution for the root nodes

POL	PortInfrast...	NavalCrew	NavalPatf...	Strong	Weak
Strong	Strong	Strong	Strong	0.9999	1.0E-4
Strong	Strong	Strong	Weak	0.2	0.8
Strong	Strong	Weak	Strong	0.2	0.8
Strong	Strong	Weak	Weak	0.2	0.8
Strong	Weak	Strong	Strong	0.2	0.8
Strong	Weak	Strong	Weak	0.1	0.9
Strong	Weak	Weak	Strong	0.1	0.9
Strong	Weak	Weak	Weak	0.05	0.95
Weak	Strong	Strong	Strong	0.1	0.9
Weak	Strong	Strong	Weak	0.1	0.9
Weak	Strong	Weak	Strong	0.1	0.9
Weak	Strong	Weak	Weak	0.1	0.9
Weak	Weak	Strong	Strong	0.05	0.95
Weak	Weak	Strong	Weak	0.05	0.95
Weak	Weak	Weak	Strong	0.05	0.95
Weak	Weak	Weak	Weak	1.0E-4	0.9999

Figure 5: Conditional probabilities for the Air Requirements capability

The numbers were estimated by the modellers and are only indicative. However, they still demonstrate the potential power of the tool and its usefulness. In Figure 6 we see the initial probability distribution for all the random variables. Note that all the root nodes are assigned a 99.99% probability of being in a strong state. This is as one might expect at the start of a conflict. As the probabilities propagate up to the terminal node the probability that the COG is in a strong state is calculated to be 99.89%. This is therefore the base case for our problem.

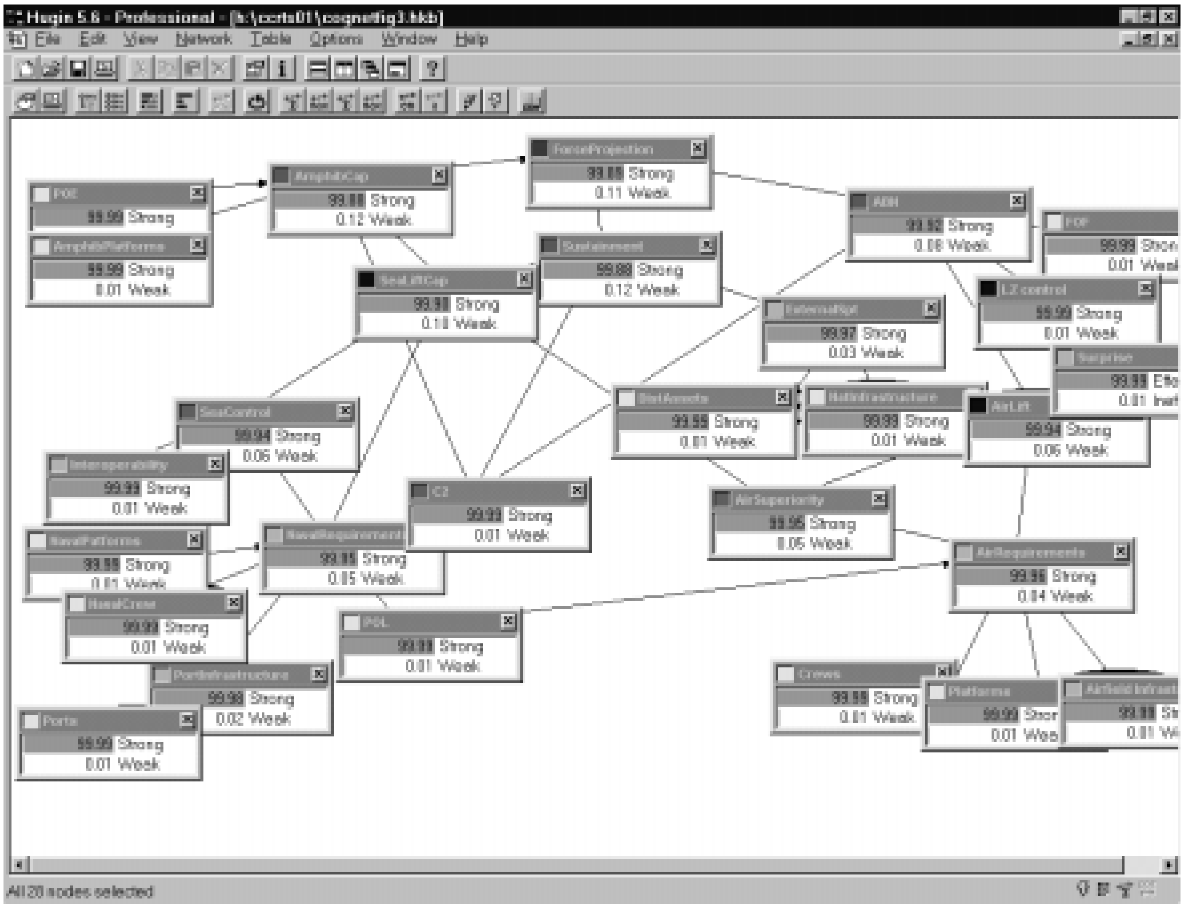

Figure 6: Initial probability distributions – the base case

We are now in a position where we can answer some “What if” questions. Based on the scenario and our own force’s strengths and weaknesses we can choose which of the enemy critical requirements can be targeted. This information can be fed into the COGNET model to investigate the effect of such actions on the COG. In Figure 7 we show the effects of targeting the enemy’s POL stocks, Command & Control, Air Platforms and Amphibious Platforms. The probability that the COG is in a strong state has been updated to 34.58%, given that the probability distributions of the targeted requirements are as shown. Software utilities that will enable sensitivity analysis to determine target critical vulnerabilities based on the desired effect on the COG are under development. Once again we note that the network is overly simplified. In the next section we outline a description of a comprehensive generic model. Of course, when the network is large and complex a more extensive sensitivity analysis would be necessary to calculate these effects for as many combinations of targets as is plausible.

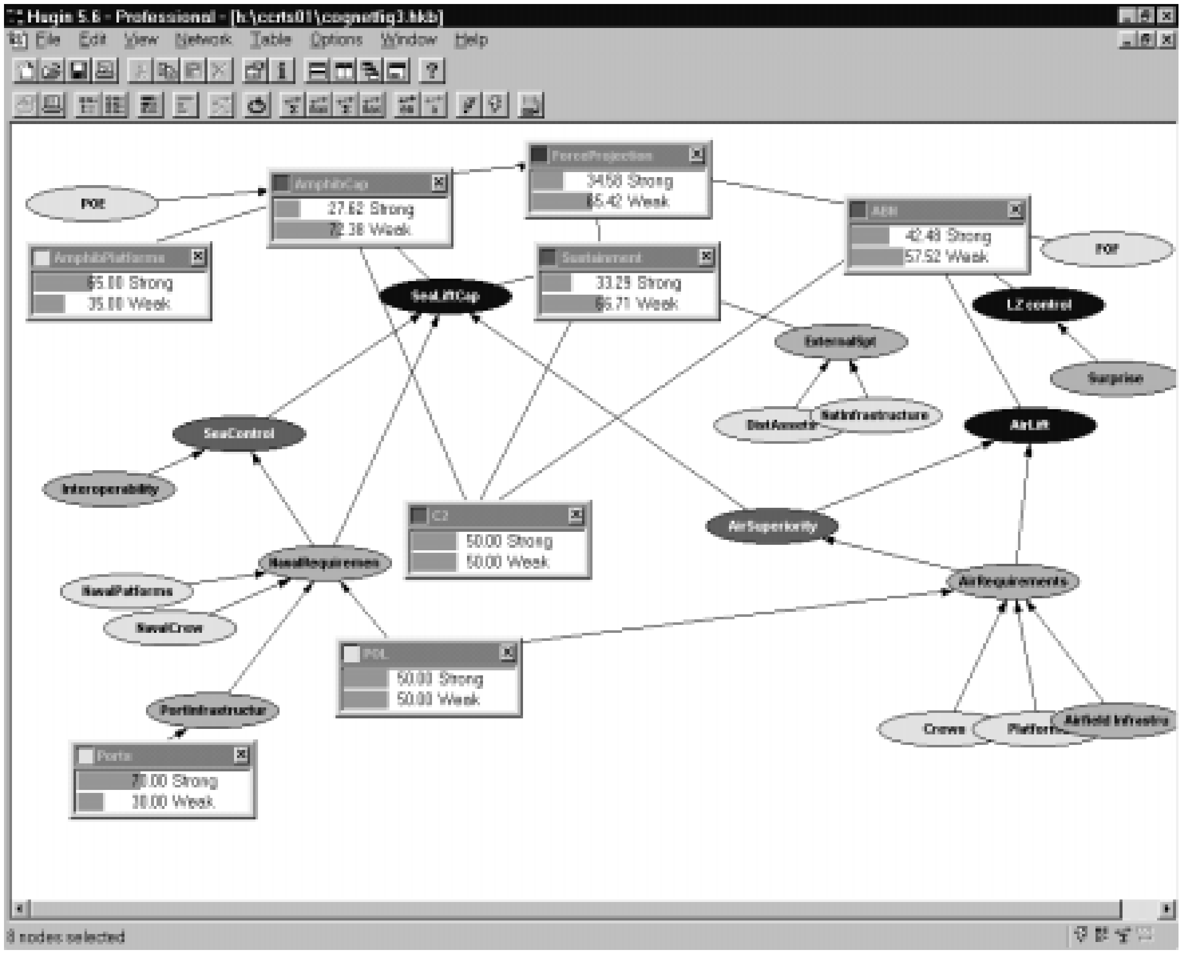

Figure 7: Examining the effects of targeting certain critical requirements

4. A Modular Knowledge Representation Framework

Deliberate (long-term) planning at the operational level aims at developing generic plans that meet the objectives set out by strategic guidance. It is conceivable that for a particular enemy the centre of gravity might change according to circumstance or type of conflict. However, the current force structure and capabilities can be reflected by a relatively fixed causal network over a fixed set of critical capabilities depending on a fixed set of requirements. The network structure would be invariant for a range of problem instances but the causal strengths may vary with respect to the specific COG being considered.

We propose a knowledge representation framework, expressing the invariant causal relationships, which can be constructed for each specific domain. This would serve as a knowledge base expressing generic causal relationships with probabilistic strengths integrated into the model to tailor it to a particular situation. A generic framework can be built on the basis of a categorisation of operational-level capabilities. A “natural” categorisation could be based on military functional areas such as Command & Control, Protection, Deployment etc and their underlying requirements, organised in hierarchies of subnetworks, which can be combined as required for each specific instance. In order to ensure that the generic model is sufficiently extensive and consistent, a hierarchically organised reference system such as a Joint Task List (JTL) is used as a basis. Task areas at the top level of the hierarchy are mapped to military capabilities, while subtasks at the lower levels are mapped to critical requirements wherever possible. The structure of a JTL is such that any task can be traced through the hierarchy to determine its contribution to higher-level tasks. In the same way our generic model can help determine which of the requirements are critical for a friendly or threat capability. Such a comprehensive Bayesian net will necessarily be large and complex but may be built from a library of modular subnets reflecting the hierarchical structure and capturing the stable patterns of probabilistic relationships. Each specific COG network would consist of a subnet of capabilities and requirements from the generic model – nodes that are irrelevant for a specific COG are simply deleted or ‘desensitised’ along with any nodes that only influence deleted nodes. Similarly the probabilistic interrelationships among relevant nodes are re-examined in light of the specific problem at hand. HUGIN facilitates establishing such a capability, however, COGNET includes additional utilities for the management and flexible configuration of COG network modules.

Figure 8 shows an example of a generic model. The model is based on a subset of the Universal Joint Task List [2]. It shows critical capabilities and Level One critical requirements. One of the Level One CRs, Offensive Ops, is expanded, as is one of its component requirements, Amphib Assault, which is also expanded all the way down to possible critical vulnerabilities. It is clear from the diagram that expanding a few more nodes will result in a sizeable complex network. A modular network structure is necessary for such a model to be feasible. The modules could be built for each specific domain and stored in a knowledge base for future use. Maintenance of random variables and conditional probability updates can be handled by using object-oriented Bayesian Networks [5], in which network fragments describe the probabilistic relations between the attributes of an object. An object here can be a random variable (a node in a BN) or may represent the relationships between attributes, which might themselves be objects. The object-oriented approach allows inheritance hierarchies as well as the ability to enclose objects within other objects.

This ability to encapsulate objects allows the user to ignore some of the detail at first. Domain experts often consider a related set of variables together. Representation of conceptually meaningful aggregates of variables and their interrelationships facilitates both knowledge elicitation and knowledge base maintenance. Similar modelling techniques for constructing and maintaining complex Bayesian networks have been previously applied to problems in military situation assessment [6] and fault diagnosis in engineering systems [9].

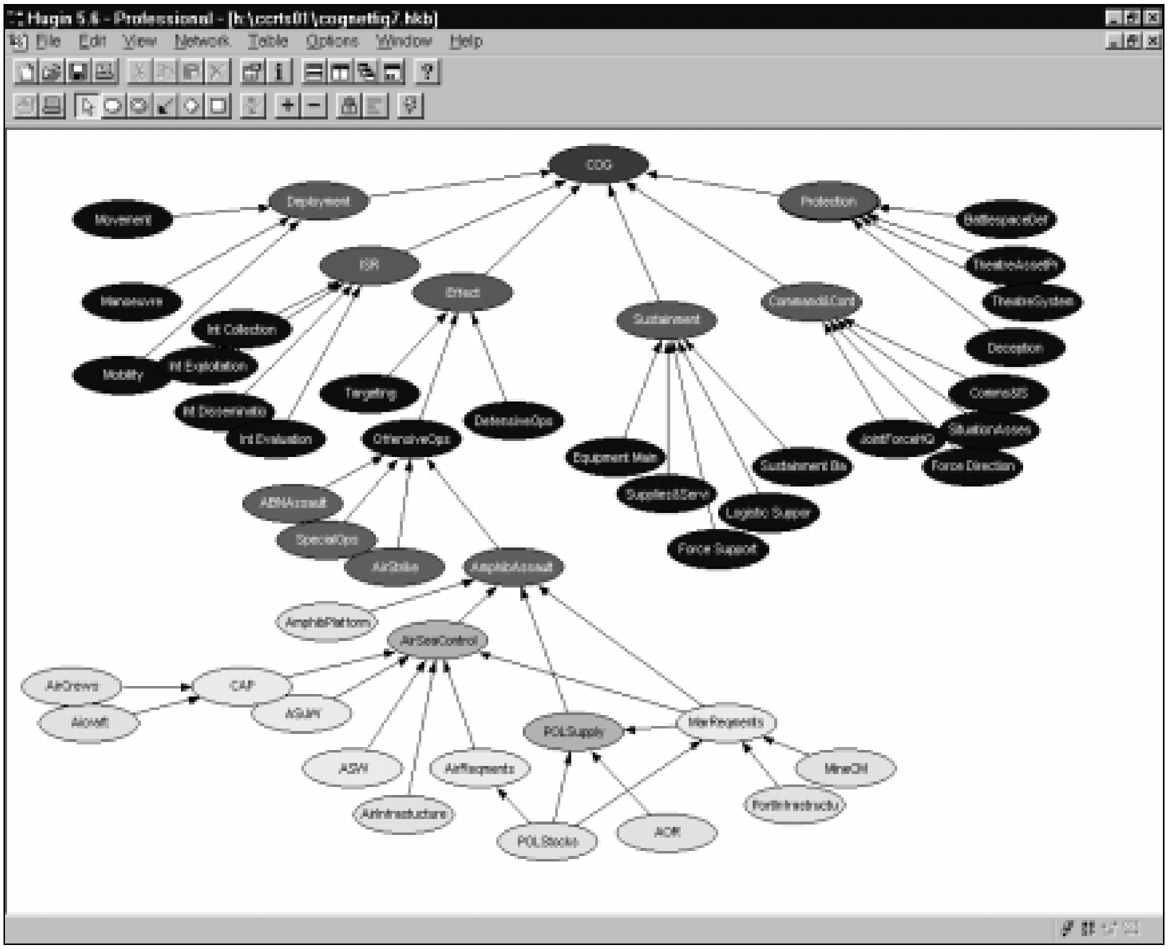

Figure 8: Designing a generic model

5. Conclusions and Future Work

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 described in this paper goes a long way to facilitate this task and it provides an effects-based analysis capability. The graphical representation allows a specification of direct dependencies resulting in a network structure that reflects induced dependencies. Furthermore it provides a framework, which can serve as a knowledge base representing generic causal relationships to aid knowledge reusability and knowledge transfer. Future research will include investigating modularisation techniques using object-oriented Bayesian networks.

Another promising avenue for research is to link the critical capability analysis described in this paper to the next step in the planning process in which a line of operations is derived from a sequence of tasks. Operational architectures philosophy advocates the adoption of a standard language (such as a Joint Task List) for describing a defence force’s capabilities. Once a JTL has been defined a set of Joint Mission Essential Tasks (JMET) considered essential to the assigned mission is derived from the high level concept of operations and the mission objective. However, the logical link between the essential tasks and the defined objective is not always clear. By definition the mission objective for a friendly force is to negate the threat COG while maintaining our own. Thus both COGs have a direct relationship to the mission objective. We propose, therefore, that a JMET list (a subset of the JTL) can be produced for each mission based on the results produced by COGNET. The relationship between the generic model and each mission-specific sub model would then be equivalent to the relationship between the UJTL and the JMET.

6. Acknowledgements

The authors would like to thank Moira Chin and Gina Kingston from Joint Systems Branch, DSTO, for discussions on the Australian approach to the categorisation of operational-level military capabilities.

7. References

[1] ADFP 9 Joint Planning,
http://defweb.cbr.defence.gov.au/home/documents/adfdocs/ADFP/adfp9.htm

[2] Chairman of the Joint Chiefs of Staff Manual, CJCSM 3500.04A. *Universal Joint Task List Version 3.0.* 13 September 1996.

[3] L. Falzon, L. Zhang and M. Davies. A policy analysis approach to operational level course of action analysis, Proceedings of the 5th International Command and Control Research and Technology Symposium, Australia War Memorial, Canberra ACT, Australia, 24-26 October 2000.

[4] J. V. Jensen. *An Introduction to Bayesian Networks*. London UCL Press, 1996.

[5] D. Koller and A. Pfeffer. Object-Oriented Bayesian Networks. Proceedings of the Thirteenth International Conference on Uncertainty in Artificial Intelligence, San Francisco California, Morgan Kaufmann, 1997.

[6] K.B. Laskey and S.M.Mahoney. Network fragments: representing knowledge for constructing probabilistic models. Proceedings of the Thirteenth International Conference on Uncertainty in Artificial Intelligence, San Francisco California, Morgan Kaufmann, 1997.

[7] K. Matthews and M. Davies, C2-led simulation and synthetic environments for air operations,_Proceedings of the 6th International Command and Control Research and Technology Symposium, US Naval Academy, Annapolis MD, 19-21 June 2001.

[8] J. Pearl. *Probabilistic Reasoning in Intelligent Systems*. California, Morgan Kaufmann, 1988.

[9] S.Srinvas. A probabilistic approach to hierarchical model-based diagnosis. Proceedings of the Tenth International Conference on Uncertainty in Artificial Intelligence, San Francisco California, Morgan Kaufmann, 1994.

[10] L. Zhang, L. Falzon, M. Davies and I. Fuss. On relationships between key concepts of operational level planning, Proceedings of the 5th International Command and Control Research and Technology Symposium, Australia War Memorial, Canberra ACT, Australia, 24-26 October 2000.

[11] L. Zhang, B. Mitchell, L. Falzon, M. Davies, L.M. Kristensen and J. Billington, Model-based planning using coloured Petri nets, Proceedings of the 6th International Command and Control Research and Technology Symposium, US Naval Academy, Annapolis MD, 19-21 June 2001.