Using CommonKADS to model the Air Campaign Planning process

John Kingston and Terri Lydiard
AIAI
University of Edinburgh

February 3, 1996

This document describes the content of a larger document with the same title. The larger document (currently 110 pages) is intended as a source of knowledge about the Air Campaign Planning process for ARPI researchers. It is based on knowledge acquisition and modelling performed between September 1995 and January 1996. The larger document is not yet considered to be complete; the authors invite comments, constructive criticism, and (in particular) requests from ARPI participants on useful additions to the larger document. The larger document is expected to be completed in late summer 1995.

The larger document is available on the World Wide Web at
`ftp://ftp.aiai.ed.ac.uk/pub/projects/oplan/ISAT/index.html`

1 Introduction to Air Campaign Planning

The knowledge described in this document focused on the processes used by the USAF to perform Air Campaign Planning. Air Campaign Planning, as the name implies, is the planning of air operations within the context of a military campaign. The key documents produced are the *Master Air Attack Plan*, which lists all the targets which should be attacked with associated priorities, and the *Air Tasking Order* (ATO), or battle plan, from which mission plans are derived.

When a crisis occurs, the Commander in Chief (CinC) will provide planning guidance to the Joint Forces Commander (JFC). This guidance is communicated to the component commanders (e.g. the Air Component commander), who will in turn communicate the guidance to the air campaign planners. Based on the guidance, the planning staff will take between 3 days and 1 week to build a plan that may be executed. Initial planning steps include assessing the situation and organizing the planning staff. Assessing the situation involves identifying enemy *centers of gravity* and air defences as well as determining a method (strategy), goals and targets. The concept of centers of gravity is a key to successful planning; these are points which, if attacked or otherwise disrupted, will disrupt some key part of the enemy's operations. Centres of gravity may be identified at different levels of abstraction; for example, the CinC might identify "transportation" or "national roadways" as a centre of gravity, whereas the planners may consider "choke points in the road network", "road bridges", or "the Forth road bridge" as centres of gravity, depending on which stage of planning they have reached.

Planning consists of three stages:

- identifying and prioritizing air objectives (e.g. "Disrupt road transport");
- identifying and prioritizing task objectives (e.g. "Damage or destroy road bridges near the front line");
- identifying and prioritizing targets (e.g. "Destroy the Forth road bridge").

It's important to note that the planning process does not always continue until each stage is fully completed. Often, the planning staff may only be 75% finished with some activities when it is time to move on to the next stage. However, the dynamic environment makes it difficult to complete a stage without the situation changing simultaneously. The time spent on any one task depends on the situation. The benefits of automated support in the planning process, which could enable certain stages of planning to be carried out more quickly and/or more thoroughly, are therefore clear.

2 Knowledge acquisition

The acquisition of knowledge about the Air Campaign Planning process was carried out in three ways:

1. Interviews with various experts – individuals with considerable knowledge of the air campaign planning process. These experts were primarily Doug Holmes and Joe Roberts of ISX, although some access was also obtained to members of Checkmate.
2. "Protocol analysis", in which a group of experts are presented with a real or hypothetical situation. The questions asked and actions performed by the experts are then recorded. In our exercise, the group of experts comprised three staff from Checkmate working on a hypothetical problem, involving an African country threatening to invade a nearby European country. One interesting observation that arose from this exercise was that the experts appear to follow a strategy of investigating all possible alternative ways of disrupting the threat before considering overt military intervention.

3. Filling in "repertory grids", which described key attributes of activities performed in air campaign planning. The repertory grid technique is normally used to acquire and investigate key attributes of objects or states; however, by applying it to activities, it was possible to acquire key information about these activities.

   There is much useful information to be gained from the repertory grid technique:

   - by obtaining and recording a complete set of information for each attribute of each activity;
   - by comparing grids between different experts to identify consensus and conflict;
   - by performing statistical analysis of the scores provided to identify underlying correlations;
   - by recording the 'asides' which the expert mentions as he is completing the grid.

3 Introduction to the modelling technique used

The modelling technique used to model the knowledge gathered about the ACP process was a variation of the CommonKADS methodology for knowledge modelling. CommonKADS [Wielinga *et al*, 1992] [Breuker & van de Velde, 1994] is a collection of structured methods for building knowledge based systems, analogous to methods such as SSADM for software engineering. It was developed between 1983 and 1994 on two projects funded by the European Community's ESPRIT program. CommonKADS views the development of a knowledge based system as a modelling activity, and so the heart of the method is the construction of a number of models which represent different views on problem solving behavior. The suggested set of models (any or all of which may be developed) is:

- An *organizational* model represents the processes, structure and resources within an organization;
- A *task* model shows the tasks carried out in the course of a particular process;
- An *agent* model represents the capabilities required of the agents who perform a process, and constraints on their performance;
- A *communication* model shows the communication required between agents during a process;

- An *expertise* model is a model of the expertise required to perform a particular activity. This model is divided into three components:
  - declarative knowledge about the domain;
  - the inference processes required during problem solving;
  - a hierarchical classification and ordering of the inference processes.
- A *design* model culminates in the design of a knowledge based system to perform all or part of the process under consideration.

For this project, organizational, task, communication and expertise models have been developed. It was decided that the modelling effort would take a "deep slice" approach; that is, the whole process would be modelled at a high level of detail (the Organizational Model), a smaller part of the process would be modelled at a greater level of detail (the Task and Communication models), and then one activity would be chosen for modelling at an even greater level of detail (the Expertise model). The intention of this approach was to demonstrate the potential of the modelling techniques of CommonKADS, in order to help in deciding how these techniques could usefully be applied.

4 Organizational model

The CommonKADS organizational model examines an organization from five major perspectives. These perspectives include:

- the activities[1] of the organization, linked together to form **processes**;
- the **structure** of the organization;
- the **power/authority** in the organization;
- the **resources** in the organization.

Each of these perspectives can be represented by a different diagram, or set of diagrams. An example of a diagram representing the process perspective can be seen in Figure 1.

[1] CommonKADS actually calls these "functions" rather than "activities", and also uses the term "tasks" elsewhere to describe similar activities. However, for reasons which are explained in [Uschold *et al*, 1995], I have chosen to use the term "activities" throughout this document.

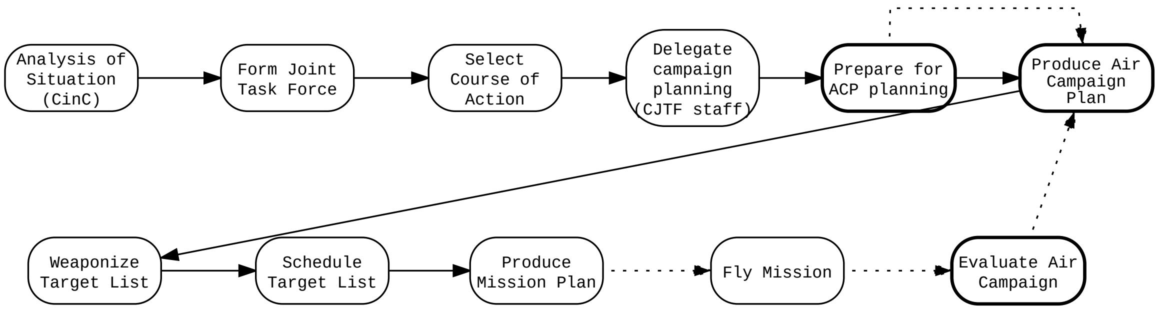

Figure 1: ACP Process Perspective

Key:		
	Solid link	the first activity *precedes* the second
	Dashed link	there is an *information flow* from the first activity to the second.

The process of domain familiarization also suggested that some aspects of the organization surrounding ACP – in particular, power/authority relationships – would benefit from being represented in more detail than the CommonKADS Organizational Model currently supports. It was decided that the multi-perspective approach to organizational modelling would be maintained, but that the perspectives would be altered to permit the representation of rights of access to resources and information, and the representation of who was responsible for performing which activities[2]. The result was that a set of perspectives were developed based on three basic entities: *activities*, *agents* and *resources*. Figure 2 shows the perspectives which were used, and how they were combined.

The **Dictionary** perspectives are simply sets of activities (or agents or resources) which are used as a basis for building the other perspectives. The **Network** perspectives show how entities of the same type relate to one another:

- Activities *precede* each other to produce a *process*;
- Agents *report to* each other and *belong to* departments to produce an *organizational structure*;
- Resources have a particular *status*, and the existence of one or more resources with a certain status may constitute a *state of affairs*. For example, "Air Campaign Plan completed" might be a state of affairs; it would exist if the status of the **Master Air Attack Plan** was *endorsed* and the status of the **Air Tasking Order** was *endorsed*.

[2]This distinction was derived from the ORDIT project which is intended to support organizational requirements definition [Dobson & Strens, 1994].

The **Cross Product** perspectives show how resources are produced, used, consumed or modified within the process (*Assets perspective*), what rights agents have to resources (*Rights perspective*), and the obligations and responsibilities to perform activities and to create states of affairs (*Responsibilities perspective*).

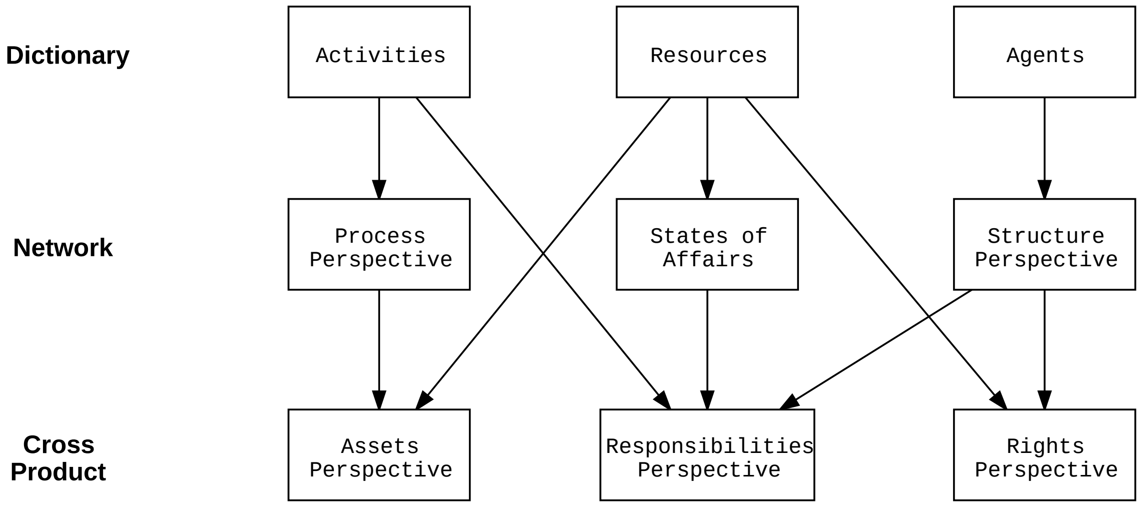

Figure 2: Enriched organizational model: perspectives and cross-products

An example of a cross product is shown in Figure 3. This shows a cross product of processes and resources, which is described as the "asset" perspective. This diagram helps to clarify the documents which are produced, and when they are produced; it also helps to identify the types of knowledge which are used by these processes, which is useful information in the later models.

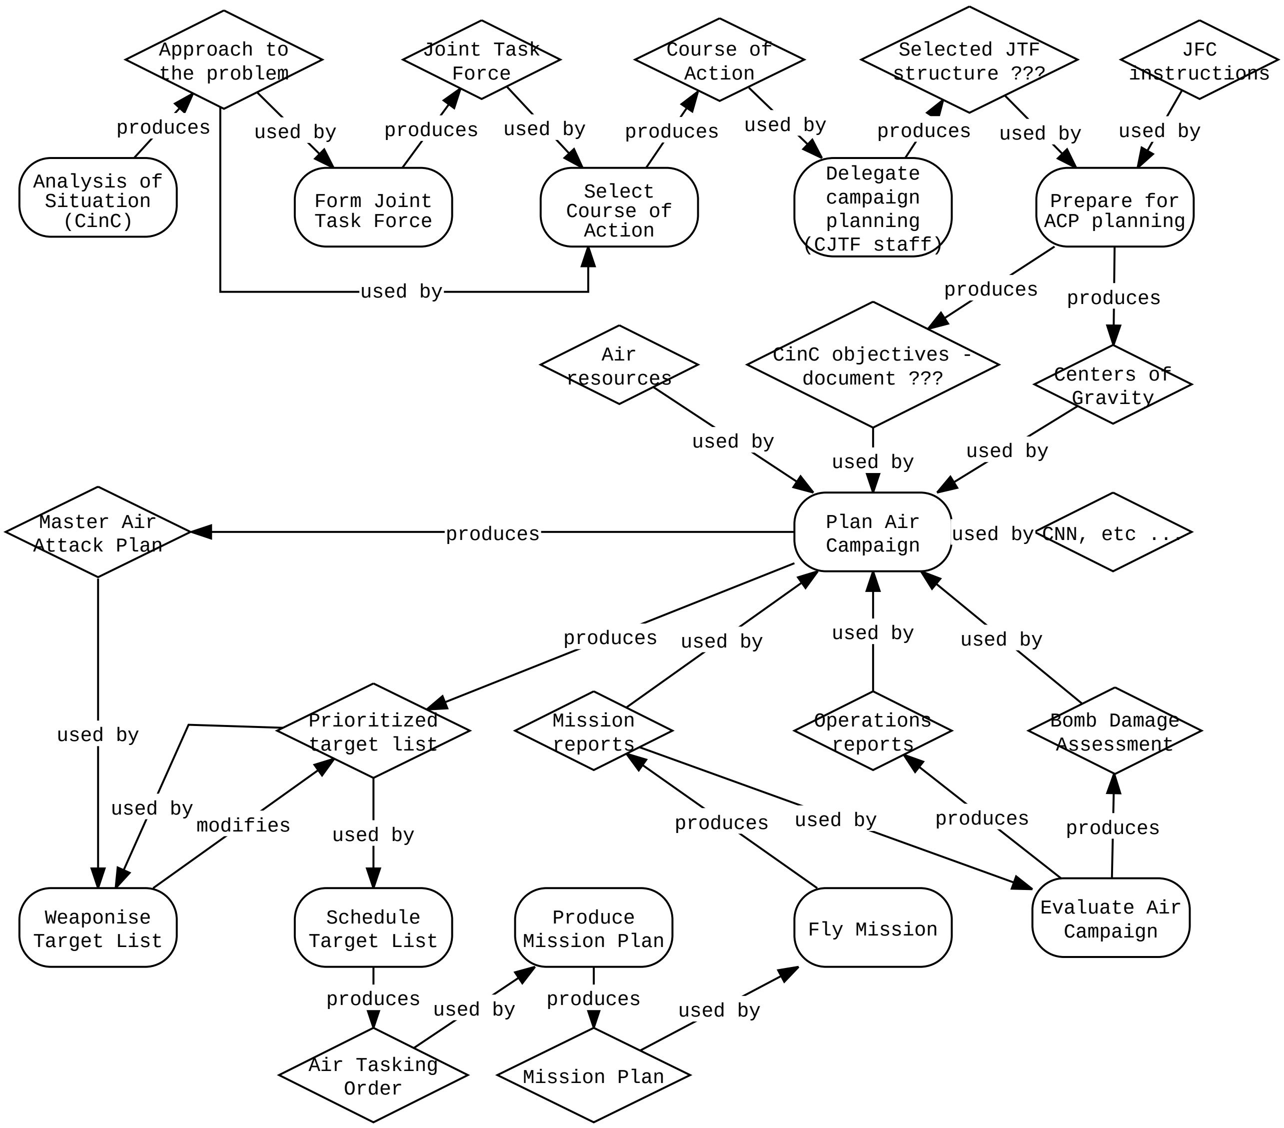

Figure 3: ACP Assets Perspective

	Labelled link	the activity *produces/consumes/uses/modifies* the resource
Key:	Diamond	Passive resource
	Rounded rectangle	Activity

Diagrams illustrating the Responsibilities and Rights perspectives can be found in the larger document. However, little information has been acquired on rights and responsibilities to date; the authors of this document would be happy to receive requests from ARPI participants who would find such information useful to their work.

4.1 The CommonKADS Task Model

Once the organizational model had been produced, a task model was produced which focused on the strategic planning of an Air Campaign Plan; that is, the part of the process which started with the CinC issuing policy and guidance, and finishes with the production of a Master Air Attack Plan (MAAP).

The CommonKADS Task Model is a more detailed version of the process perspective of the organizational model. The process perspective of the organizational model is intended to represent the overall workflow and/or control flow of an organization's primary process; the task model describes the activities which are executed in an organizational environment, with an emphasis on deciding which activities could be suitable candidates for computer software support [deHoog *et al*, 1993]. The extra detail includes the agents who can perform these activities; any temporal restrictions on activities; and the level of difficulty of these activities.

In addition to specifying all these attributes, it was decided that the IDEF3 technique for process modelling would provide useful diagrams of the processes involved, since it represents more detail about the exact ordering of processes than CommonKADS' graphical technique.[3] The two processes which are examined within the Task Model (**Prepare for Air Campaign Planning** and **Plan Air Campaign**) can be seen in Figures 4 and 5. "Asset perspective" diagrams were also produced (see larger document).

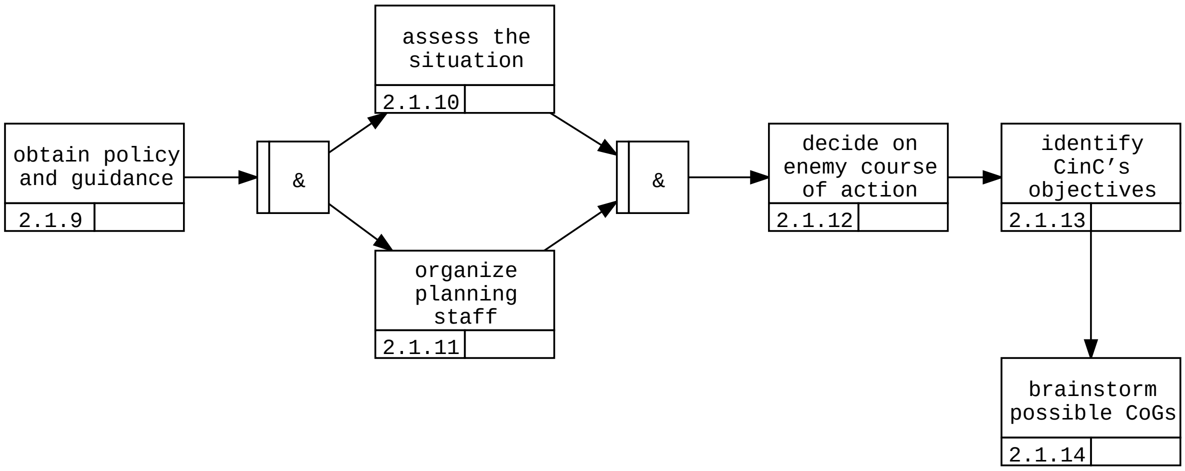

Figure 4: Task Model: Prepare for Air Campaign planning

[3]The numbering scheme of IDEF3 diagrams is as follows: 2.1.9 indicates "the ninth node defined in all diagrams, which is part of the first decomposition of node number 2"

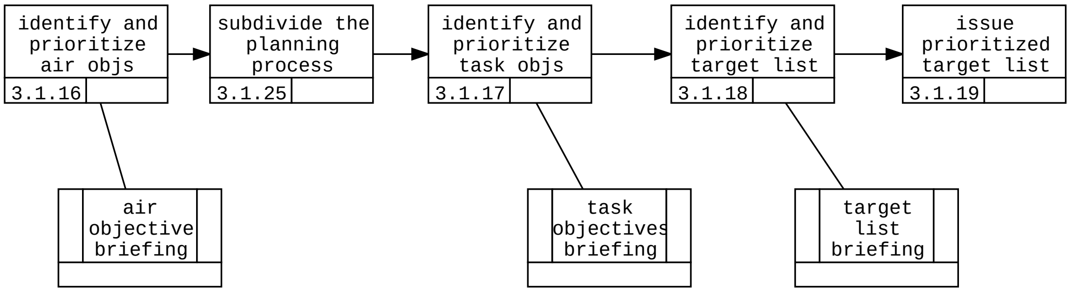

Figure 5: Task Model: Plan Air Campaign

Key:

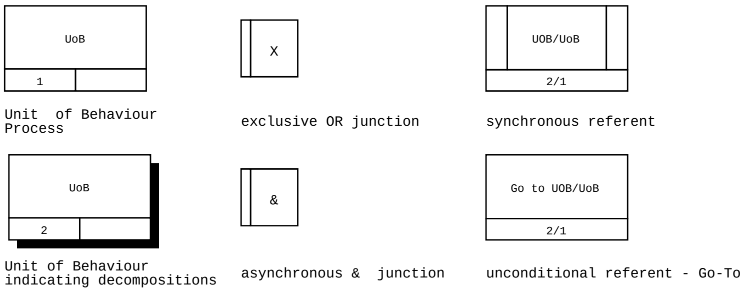

The information on the level of difficulty of the activities, which was acquired using the repertory grid technique, is the most useful for determining which activities would benefit most from computer support. Activities which take considerable amounts of time are good candidates for computer support; if the activity requires an experienced officer to perform it and/or has a high likelihood of making a mistake, then it would probably benefit from support using a knowledge based system.[4] An examination of the repertory grid shown in Figure 6, which shows the average repertory grid scores from three members of Checkmate, suggests that the following activities would benefit from support knowledge based support:

- Assessing the situation;
- Deciding on the enemy's most likely course of action;

[4]Seriousness of mistake is also an important factor to consider, but statistical analysis of the repertory grids shows that, for air campaign planning, seriousness of a mistake is strongly correlated with the expwerience required to perform an activity.

- Identifying the CinC's objectives;
- Identifying and prioritizing air objectives;
- Producing and sequencing task objectives

At this level of detail, no activities were required as needing conventional computing support. However, the repertory grid technique was also applied at a lower level of detail to examine the various sub-activities involved in identifying objectives and targets (see larger document). The results of this analysis were that at least two activities – the identification of targets and the prioritization of targets – would benefit from conventional computing support.

Scale	Plan Air Campaign	Obtain policy and guidance from CinC/JFC	Organize planning staff	Assess Situation	Identify campaign objectives and military objectives	Subdivide the planning problem	Decide on enemys most likely course of action	Review enemy course of action	Communicate plan to subordinate	Brief CJTF	Produce prioritize air objectives	Produce prioritize task objectives	Produce prioritize target list
1 = Little 5 = Much	Experience required	3	4	4	4	3	4	4	3	4	5	4	4
1 = Trivial 5 = Very serious	Seriousness of mistake	4	3	5	5	3	5	4	4	4	4	4	4
1 = Rare 5 = Frequent	Likelihood of mistake	3	2	3	3	2	3	2	3	2	3	3	3
1 = Short 5 = Long	Time required	2	2	4	4	2	4	3	2	2	4	4	3
1 = Analytic 5 = Synthetic	Analytic vs Synthetic	4	4	3	3	3	2	3	4	3	3	3	2
1 = Little 5 = Much	Evaluation required	4	2	4	5	2	4	3	2	3	4	4	4

Figure 6: Repertory grid: preparing & planning an air campaign.

One of these activities – the identification and prioritization of air objectives – was chosen for further analysis using the CommonKADS Expertise Model. A major reason for choosing this activity was the similarity of the sub-activities required for identification and prioritization of air objectives, producing and sequencing task objectives, and identification and prioritization of targets.

4.2 Communication Model

The task model provides considerable detail about the individual activities within a process, but it does not show many details about the communication between different individuals or groups that is required as part of the process. This is the role of the communication model. Figure 7 shows a communication model for all the activities identified in Figure 2; it has been drawn using the Role Activity Diagram (RAD) technique.

Figure 7: Communication Model for Air Campaign Planning

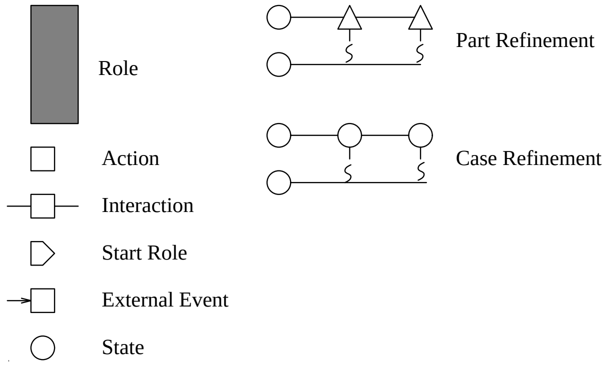

Key for Role Activity Diagram

4.3 The CommonKADS Expertise Model

The Expertise Model is a model of the expertise required to perform a particular activity. The Expertise model is divided into three components, each of which is represented by one or more diagrams:

- *domain knowledge:* declarative knowledge about the domain;
- *inference knowledge:* the inference steps required during problem solving;
- *task knowledge:* control information, largely concerned with the ordering of inference steps.

The Expertise model can be developed bottom-up (by acquiring domain knowledge and identifying relevant processes), top-down (by modelling task & inference knowledge and identifying relevant domain knowledge) or by a mixture of the two, with each component being partially developed and then used to support the development of other components. In this case, the task knowledge was developed first, and then the inference knowledge and domain knowledge were developed concurrently.

4.4 Task knowledge

The task knowledge within the expertise model is in a similar format to the task model, and the process perspective of the organisational model. It consists of a diagram showing precedence relationships between the various sub-activities of identifying and prioritizing air objectives, and also gives a detailed description of the individual activities.

N.B. The diagram format used for the task knowledge is currently the same as that used for the organisational model; a different format may be used in future.

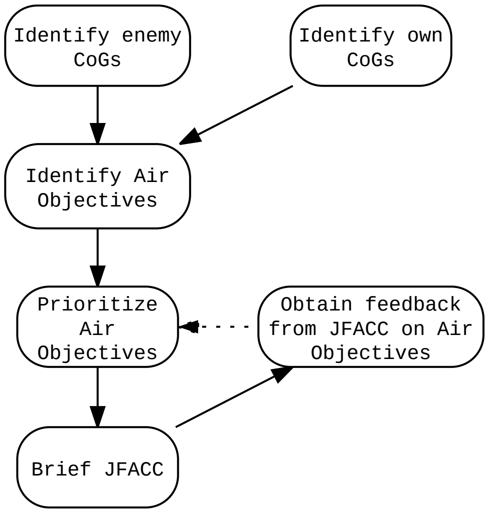

Key:		
	Solid link	the first activity *precedes* the second
	Dashed link	there is an *information flow* from the first activity to the second.

Figure 8: Task structure for identifying and prioritizing air objectives

N.B. The attributes of activities are not provided in this shortened document, but can be found in the larger document.

4.5 Inference knowledge

The inference knowledge in CommonKADS provides a "bridge" between the task knowledge and the domain knowledge. It does this by representing each activity asx an "inference step", using one of a standard set of verbs to describe the operation being performed by that activity[5]; if necessary, the activity is broken down into several inference steps. The inputs and outputs of each inference step are then identified; each of these should map to one or more items in the domain knowledge.

The key to the inference knowledge for identifying and prioritizing air objectives came from two observations made in the domain knowledge. The first observation was that all types of objectives (CinC objectives, air objectives and task objectives) follow a similar format. This format is:

Achieve a GOAL on a `center of gravity` with an optional *constraint*.

[5]A second standard set of verbs is also available in CommonKADS, for those who prefer to identify activities by what they are trying to achieve rather than the type of operation which is being carried out.

The second observation was that centers of gravity can be conceptualized as being structured in a hierarchy, the levels of which match the different types of objectives. On this basis, each center of gravity identified in a CinC objective should have a one or more centers of gravity specified at the next level down the hierarchy, which will become the centers of gravity for air objectives. The inference knowledge therefore shows that identifying air objectives requires:

- identifying the goal and the center of gravity in a CinC objective;
- identifying corresponding centers of gravity at the air objectives level of the hierarchy;
- selecting a goal which, when applied to these centers of gravity, will fulfil the effect required by the CinC objective.
- composing these into an air objective.

Once the air objective has been identified, its priority can be calculated.

The four-step process described above appears quite simple; in practice, the activity is (not surprisingly) somewhat more complex. The "inference structures" which display the inference knowledge for identifying and prioritizing air objectives can be found in Figure 9 (the top level), Figure 10 (the inference steps required in selecting a goal), and Figure 11 (the inference steps required in calculating a priority). Ellipses represent inference steps, and rectangles represent "knowledge roles" (domain knowledge which plays a part in the inference process).[6]

[6] The numbering scheme used in these diagrams is the same as that used by the IDEF3 notation.

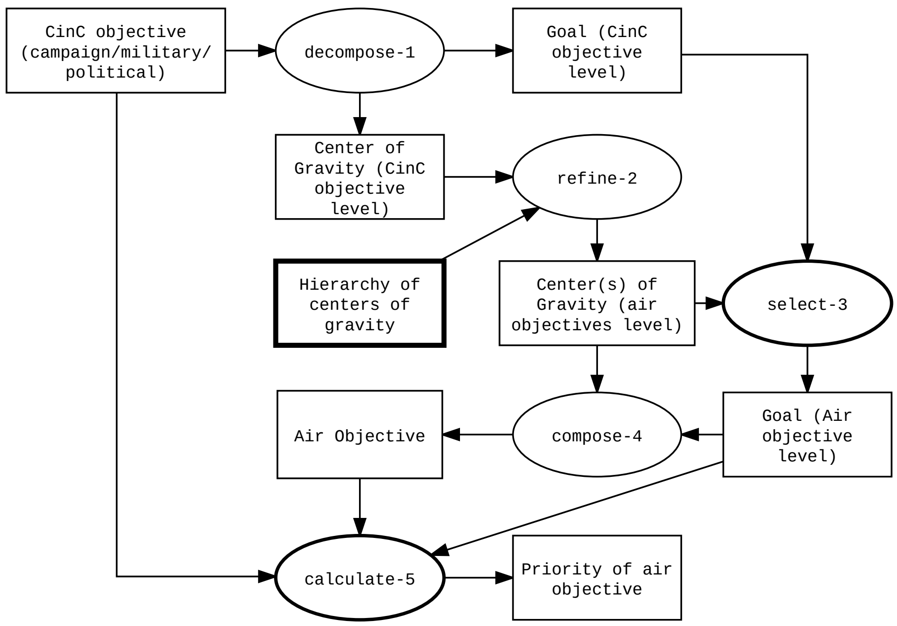

Figure 9: Inference Structure: Top level

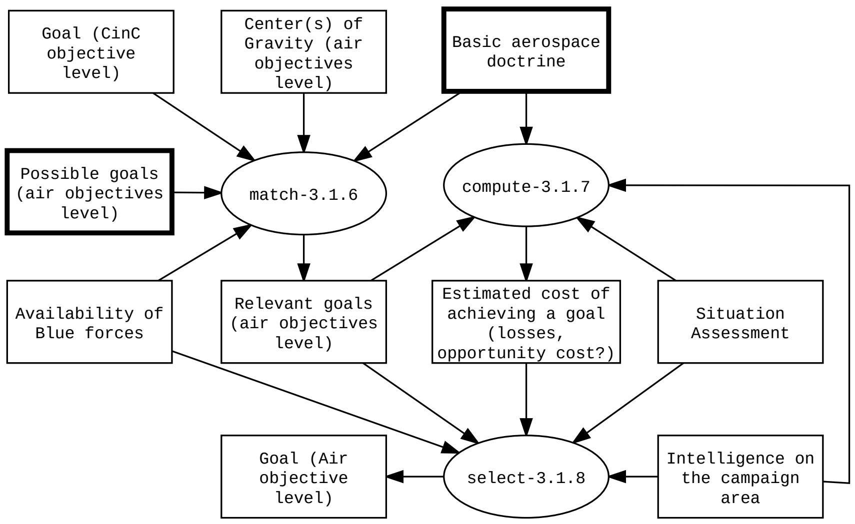

Figure 10: Inference Structure: Selecting a Goal

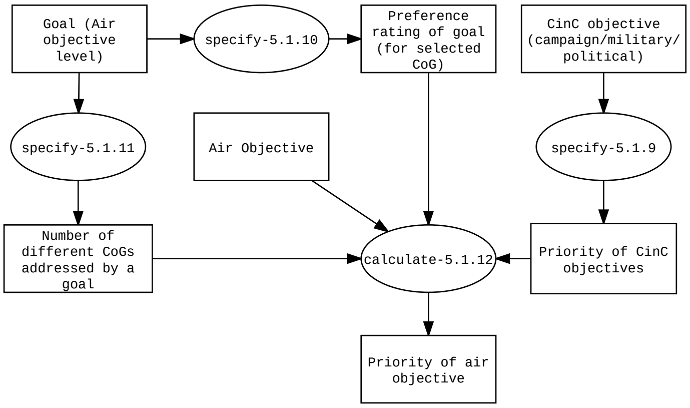

Figure 11: Inference Structure: Calculate Priority

4.6 Domain knowledge

Domain knowledge for this activity has been acquired by analyzing a list of CinC objectives, air objectives and task objectives which was provided by Checkmate. As stated above, it was observed that all objectives consist of three major concepts:

- Centers of gravity (e.g. **national roadways**);
- Goals for particular centers of gravity (e.g. **Disrupt** or **Destroy** a CoG);
- Constraints limiting the goals, such as time constraints (e.g. Destroy the bridge **before noon tomorrow**) or geographical constraints (e.g. Disrupt all radar sites **within 100km of the capital city**). Other types of constraints include categorical constraints (e.g. differentiating weapons of mass destruction from other weapons).

These three concepts comprise the key domain knowledge for this activity.

Centers of gravity: The centers of gravity which were acquired have been structured into a hierarchy, which is sufficiently detailed that it needs to be spread over several diagrams. One of these diagrams is shown in Figure 12.

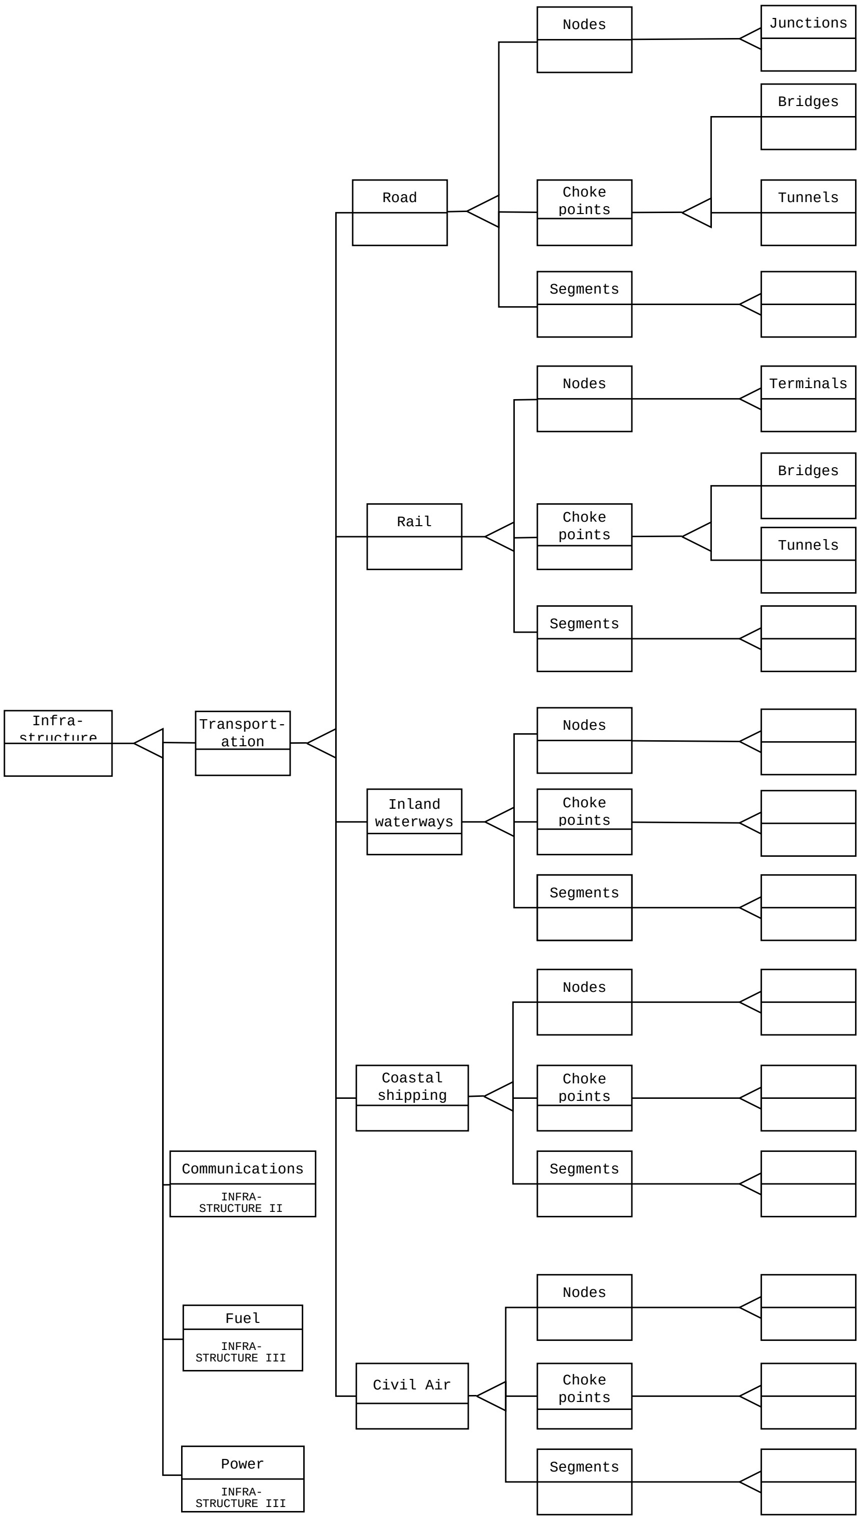

Figure 12: Centers of Gravity related to Transportation

The hierarchy in Figure 12 should be read from left to right. The top level (on the left) is **National Infrastructure**, which is one of the 5 categories of centers of gravity identified by Col. Warden. The next level is **Transportation**, since national infrastructure consists of so many centers of gravity that the authors of this document have subdivided them under the following categories:

- transportation;
- communication;
- fuel & power.

The third level consists of different categories of transportation (national roadways, national railways, coastal shipping, etc.). This is typically the level of centers of gravity which are included in the CinC objectives, and so this document refers to the third level of the hierarchy as the "CinC objectives level". The fourth level consists of *categories* of centers of gravity related to each type of transportation (e.g. **choke points**); these often appear in air objectives, and so the fourth level is referred to as the "air objectives level". The fifth level is known as the "task objectives level"; it specifies types of targets which could be attacked. There is implicitly a sixth level (the actual targets), but this cannot be drawn in without reference to a particular situation.

N.B. The blank boxes are place-holders in parts of the hierarchy as a reminder of where further knowledge acquisition is required to complete the hierarchy.

The key observations which lead to this hierarchical structuring were:

- Many centers of gravity appear to be decomposable into a lower level of abstraction. For example, **national roadways** can be decomposed into **nodes**, **choke points** and **segments**; similarly, **choke points** can be decomposed into **bridges** and **tunnels**. These 3 levels of decomposition correspond to the centers of gravity named in different types of objectives (CinC objectives, air objectives and task objectives).

- In many cases, centers of gravity at the same level of abstraction have a similar decomposition. For example, both **national roadways** and **national railways** are sub-CoGs of **transportation**, so they are at the same level of abstraction; and both of them decompose into **nodes**, **choke points** and **segments**. The reason for the similar decomposition is that both national roadways and national railways are networks. The power of this observation becomes clear when it is applied to other types of network, such as the electrical network; these, too, have nodes (e.g. power stations), choke points (e.g. certain transformers) and segments (e.g. overhead power lines).

The second observation above is indeed a powerful one, for it suggests that a top-down approach to identifying centers of gravity may come close to exhaustively identifying all possible centers of gravity. The key to the approach is to specify centers of gravity which **should be** present by drawing analogies with other centers of gravity at the same level of the hierarchy. For example, if the planners were asked to identify all centers of gravity related to transportation in a country, they would consider nodes, choke points and segments for all types of transportation. This should lead to the consideration of possible centers of gravity which might otherwise be overlooked; perhaps consideration of choke points for coastal shipping might lead to an investigation of narrow safe shipping channels as possible centers of gravity.

A key component of the hierarchies which contributes to the goal of identifying all relevant centers of gravity is the identification of *intra-hierarchy links*. These show where a center of gravity in a different part of the hierarchy is relevant to the current center of gravity. The name of the related hierarchy is displayed in the bottom half of the CoG node. For example, one form of communication is by newspapers and printed media; the distribution of these requires transportation. A planner examining the **segments** of the newspaper communication network would therefore be directed to look at the centers of gravity for the appropriate type(s) of **transportation**.

Following these intra-hierarchy links may lead to the identification of unexpectedly relevant centers of gravity. To take an example based on real life, a planner who was attempting to prevent a large army of tanks from invading would examine the hierarchy for **Fielded Forces**, and would discover "Logistics" as a center of gravity. In this category, he would be invited to consider **Ammunition**, **Fuel**, **Spare Parts** and **General Munitions**. On examining Fuel, he might decide that fuel **storage** is a potential center of gravity; but not only is it difficult to attack, it is also an inadequate supply for a large army of tanks. This inadequacy suggests that the other possible source of fuel – a supply line – should be considered.[7] Here an intra-hierarchy link applies; the fuel is either supplied by transport vehicles (in which case a link must be made to **National Infrastructure: Transportation**) or, as in this case, it is supplied by pipeline, in which case the appropriate starting point is **National Infrastructure: Fuel**. Since fuel pipelines form a network, then **nodes**, **choke points** and **segments** must be considered. The segments in a fuel network are pipelines; it turns out that there are only 5 pipelines used for pumping fuel to these storage areas, which makes the pipelines a possible center of gravity. However, fuel pipelines are hard to attack and easy to repair. The nodes in the network are the input fuel plants and the storage areas, which are either too heavily defended or too far away to consider. The consideration of choke points, however, leads to the identification of three (undefended) pumping stations

[7] In fact, there are two other possible sources of fuel; direct access to a fuel production plant, and siphoning it from other vehicles. The former is unlikely to avoid notice, and the latter is infeasible for a large army, even if they were not using specialized fuel.

at a point where the gauges of the fuel pipes change, thus ensuring that all fuel is routed through these three stations.

Goals: The goals which were derived from the objectives were analyzed and are shown in tables in the larger document, along with suggested definitions of each goal. The goals can be classified two dimensions: whether they imply *temporary* or *permanent* effects on the center of gravity, and whether they are *partial* or *complete* effects. This four-way categorization corresponds loosely to the standard set of goals used by the planners when considering interdiction: DELAY, DISRUPT, ATTRIT and DESTROY. However, table 4.6 suggests that "delay" could helpfully be renamed to "deny use".

Goal	Definition
Disrupt	Temporarily reduce effectiveness
Attrit	Permanently reduce effectiveness
Delay/Deny use	Temporarily reduce effectiveness to zero
Destroy	Permanently reduce effectiveness to zero

Table 1: Standard goals for interdiction, with definitions

Constraints: Many CinC objectives and Air objectives, and some task objectives, have associated constraints. These constraints may be:

- geographical (e.g. "Destroy SAMs **beyond 100km of the DMZ**");
- temporal (e.g. "Sustain (Blue) forces **for a minimum of 45 days**");
- categorical (e.g. "Deny use of national roadways for the transport of **military equipment**");
- capability (e.g. "Deploy (Blue) forces **capable of deterring YOC aggression**").

Analysis of the constraints shows that the constraints are diverse, but a few patterns can be seen:

- The distinction between **military** and **civilian** is frequently drawn. Typically, the planners will avoid targetting the civilian population unless there is a military end state. An example of an objective which does target population is:

  DENY USE of `national television systems` to influence population (of the hostile and adjoining country).

- Another distinction which appears to be important is the distinction between **weapons of mass destruction** and other weapons. This is reflected in the hierarchies of centers of gravity, where "special weapons" (i.e. nuclear/ biological/ chemical weapons and ballistic missiles) are treated separately from other weapons.

- In the higher level objectives, the majority of geographical constraints are concerned with fielded forces, and their locations in relation to battlegrounds and demilitarized zones. However, there are also some concerns related to the whole region, typically connected with evacuation of U.S. civilians and other non-combatants.

5 Conclusion

By modelling a "deep slice" of the air campaign planning process, the uses of the CommonKADS modelling technique have been illustrated. In the process of modelling, a number of discoveries have been made which shed light on the process of air campaign planning:

- the representation of air campaign planning in terms of **activities**, **agents** and **resources**;

- the identification of activities within air campaign planning which would benefit from knowledge based support and/or conventional computing support;

- the possibility of exhaustive identification of centers of gravity using top-down modelling;

- the possibility of finding unexpectedly relevant centers of gravity by following intra-hierarchy links.

References

[Breuker & van de Velde, 1994] Breuker, J. and van de Velde, W. (1994). *The CommonKADS Library: reusable components for artificial problem solving.* IOS Press, Amsterdam, Tokyo.

[deHoog *et al*, 1993] de Hoog, R., Martil, R., Wielinga, B., Taylor, R., Bright, C. and van de Velde, W. (1993). The Common KADS model set. ESPRIT Project P5248 KADS-II KADS-II/M1/DM1.1b/UvA/018/ 5.0, University of Amsterdam and others.

[Dobson & Strens, 1994] Dobson, J. and Strens, R. (1994). Organisational Requirements Definition for Information Technology Systems. In *ACM International Conference on Requirements Engineering, Denver, USA, 1994*. ACM.

[Uschold *et al*, 1995] Uschold, M., King, M., Moralee, S. and Zorgios, Y. (1995). *The Enterprise Ontology.* AIAI, University of Edinburgh, Also available on the World Wide Web: http://www.aiai.ed.ac.uk/ entprise/enterprise/ontology.html.

[Wielinga *et al*, 1992] Wielinga, B., Van de Velde, W., Schreiber, G. and Akkermans, H. (1992). The KADS Knowledge Modelling Approach. In *Proceedings of the Japanese Knowledge Acquisition Workshop (JKAW'92).*