Knowledge Level Planning in the Search and Rescue Domain

Version 9.

Hugh Cottam†, Nigel Shadbolt†, John Kingston‡,
Howard Beck‡, Austin Tate‡.

† University of Nottingham, AI Group, Dept of Psychology,
University Park, Nottingham NG7 2RD.
‡Artificial Intelligence Applications Institute, University of Edinburgh, 80 South Bridge, Edinburgh EH1 1HN.

Abstract

The increased use of intelligent decision support systems has created a demand for efficient acquisition, implementation and maintenance of the knowledge required by such systems. The field of knowledge level modelling has developed as a means to this end. This has led to the construction of methodologies for KBS development that facilitate a generic approach to knowledge acquisition. Such generic approaches have achieved great success when applied to various domains, yet have thus far largely neglected the generic areas of planning, scheduling and resource allocation. In this paper we outline the development of such a generic approach within the domain of planning for Search and Rescue. Our generic approach makes a distinction between domain derived knowledge level models and those derived from systems. We describe how the combination of these two types of model can achieve definite benefits within the course of KBS development.

Acknowledgements:

Thanks to RCC Edinburgh for their co-operation and support, in particular Squadron Leader W. Gault. The work described has been done under contract to the Defence Research Agency Flight Systems Division, Farnborough.

1. Introduction

This paper describes work undertaken and insights derived from a project commissioned by the Defence Research Agency Flight Systems Division. The project is entitled “Acquiring and Using Planning Knowledge for Search and Rescue” and was motivated by three concerns. Firstly there was the clear requirement for intelligent decision aids within the domain of search and rescue (SAR) planning. Second was the requirement by the Flight Systems Division of DRA Farnborough to find ways in which their knowledge engineering work within planning system development could be made more efficient and reusable. Finally there was the mutual interest of the AI Group at Nottingham and AIAI at Edinburgh, in the development of generic approaches to knowledge acquisition for planning systems.

1.1 The need for generic approaches to KA for planning

The increased use of intelligent decision support systems has created a demand for efficient acquisition, implementation and maintenance of the knowledge required by such systems. The knowledge level concept [Newell, 1982] and the field of model based knowledge acquisition have developed as a means to this end. This has led to the construction of methodologies for KBS development that facilitate a generic approach to knowledge acquisition. e.g. KADS [Breuker *et al.*, 1987.] or VITAL [O'Hara *et al.*, 1992]. These methodologies rely upon the concept of a generic problem-solving model (PSM). Generic PSMs resulted from the discovery that when a certain number of PSMs were purged of their domain-specific content, the resulting structures seemed invariant over various domains. Users of knowledge level methodologies have thus built up extensive libraries of generic PSMs, aimed at facilitating the reuse of both knowledge engineering effort and system software itself.

Such generic approaches have achieved great success when applied to various domains, yet have thus far largely neglected the areas of planning, scheduling and resource allocation. This point should be clarified as it might well be argued that there are in fact a number of existing generic PSMs for planning. e.g. *CommonKADS Library for Expertise Modelling* [Breuker & van de Velde, 1994]. The important observation to be made about these existing PSMs for planning is that they are often system derived. This is in fact an observation that holds for many of the PSMs in various libraries such as KADS [Breuker *et al.*, 1987.], VITAL [O'Hara *et al.*, 1992], PROTEGE [Musen *et al.*, 1995.]. Because such planning PSMs model how computers plan rather than how humans plan, their efficacy for human expert knowledge acquisition is debatable, and they may enforce an unsuitable system architecture upon the domain.

An important discussion point concerns the origin of the generic PSMs associated with a methodology. The originators of a methodology may deduce the ontological elements for use in such models, yet the structure of generic models must be inferred or validated inductively. There are several established generic PSMs that have originated as a result of system analysis, as opposed to human expert analysis. e.g. the heuristic classification model of Mycin [Clancey, 1985]. However the strength of a model such as heuristic classification is not simply the initial system analysis, but its proven efficacy for knowledge acquisition. This observation has important implications for the construction of generic PSMs within the generic area of planning. Existing planning PSMs have resulted as an attempt to extend the use of knowledge level methodologies, yet are system derived and await validation through their use and refinement in the context of knowledge acquisition. However, there is a danger when applying system derived PSMs to knowledge acquisition, in that the knowledge engineer may tend to force the characterisation of the domain to fit the "off the shelf" model.

Our observations on the nature of existing PSMs for planning, led us to the goal of constructing an explicit PSM for the SAR domain from a combination of domain analysis and wider ontological considerations. Thus the structure of the PSM was domain driven, which we considered to be of great importance in order to avoid imposing a pre-conceived PSM upon the domain. Once shorn of its domain specific content this model could be compared to generic PSMs such as those mentioned earlier from the CommonKADS library. The applicability of the resulting generic PSM could then be considered for other domains. Initially the ontological issues regarding the entities that would constitute a PSM for planning were addressed, and possible structures of PSMs for planning in the SAR domain were then considered.

1.2 Ontologies for planning

There is widespread interest in ontologies to support knowledge sharing across a range of domains and, at this time, a rapidly growing interest in the development of ontologies for planning [Tate, 1994]. These ontologies establish a set of consistent terms to describe the entities that constitute a plan and the relationships between such entities. The ontologies thus represent "what" is reasoned about in planning, but do not explicitly represent "how" planning is performed. It was our objective to merge generic planning ontologies with a knowledge level modelling approach in order to formulate a planning PSM for the SAR domain.

An example of a generic planning ontology is that proposed by [Tate, 1994] which has been adopted as a base for the US ARPA working group on Plan Ontology, and as a component of Enterprise modelling ontologies [Fraser & Tate, 1995]. These efforts are being undertaken for a variety of purposes, one of which is to assist the knowledge acquisition process. Other developments that have much in common with the Plan Ontology work are the Workflow Management Coalition glossary [1994] and the Process Interchange Format standards [Lee, 1994].

Generic planning entities defined in the Plan Ontology work include the following [Tate, 1995]:

Activity -- is a behaviour performed by one or more agents.
Agent -- is an entity that can perform or participate in behaviour, or hold some purpose.
Issue -- an implied or pending constraint upon a plan. Requirements remaining to be addressed in the plan.
Activity Decomposition -- is a set of sub-activities or sub-activity constraints. There are normally multiple ways the activity can be decomposed.

Constraints -- there are three main types of constraint:

> **Implied Constraints** -- represent future constraints that will be added to the plan as a result of handling unsatisfied requirements.
> **Plan Level Constraints** -- actions in the plan associated with begin and end time points.
> **Detailed Constraints** -- of two main sub types:
>> **Ordering Constraints** -- define temporal relationships between actions.
>> **Variable Constraints** -- co-designation and non-co-designation constraints on plan objects.

Additional detailed constraints may include authority constraints, that define agent to agent relationships such as the authority of one agent over another, or delegation.

These generic planning entities represent abstractions of the possible items that may exist within a particular planning domain. The range of generic entities defined in the Plan Ontology will in general be a superset of the entities required for a specific domain. They therefore also define a superset of the knowledge roles or information types that will be contained within our planning PSM.

2. The SAR Domain

Within this project the course of knowledge acquisition largely followed the lifecycles of KBS described in methodologies such as KADS [Tansley & Hayball, 1993] and VITAL [Motta *et al*., 1994]. The lifecycle does however vary from the norm at certain stages owing to the lack of a generic PSM on which to base the modelling process. Knowledge acquisition (KA) commenced with extensive tutorial KA involving basic documentation of the domain.

Planning for search and rescue is based at the Rescue Co-ordination Centre (RCC) at Pitreavie near Edinburgh. The RCC have responsibility for the support of military flying and to provide cover for Royal Flights, yet their most common role is in support and co-ordination of civilian emergencies. The RCC's geographical area of responsibility extends from a line South of Birmingham northwards as far as Iceland, and out into the North Atlantic and North Sea. The Southern part of the UK is covered by an equivalent RCC based at Plymouth. The RCC have direct responsibility for the allocation, application and co-ordination of military assets for SAR (this includes SAR helicopters, RAF Nimrods and RAF mountain rescue teams). They may however have to co-ordinate with a number of civilian emergency authorities such as fire, police, ambulance, coastguard and civilian mountain rescue teams. They might also take responsibility for overall co-ordination of a rescue incident that includes the allocation and application of these civilian rescue assets. A rescue incident can vary in scale from retrieving a walker with a sprained ankle to handling a large aircrash.

Tutorial KA continued with domain experts giving in depth descriptions of individual rescue incidents and the decisions associated with the handling of those incidents.

These interviews were taped and transcribed for analysis. A number of video recordings were made of actual incidents in progress and the RCC's handling of them. These recordings proved invaluable in eliciting an explicit structure describing problem solving for SAR.

The video medium proved particularly effective at capturing the temporal relationships prevalent within planning at the RCC. The videos enabled us to rigorously document both the RCC's information sources and the manner in which they interacted with these during the course of their workflow. This was documented in an information stores glossary, which is a listing with in depth explanation, of all the stores used by the RCC for information within the course of their SAR work. Most of the RCC's information sources are either paper based or magnetic boards. These types of representation tend to be used for those information items within the RCC that are frequently changing. A computer database is used containing information of a more static nature; e.g. geographical positions of things such as hospitals, landing sites, and decompression chambers.

3. The Context of SAR Planning

Fig 1 shows the high level tasks identified in the RCC's workflow along with the dependencies that exist between them. Arrows represent a temporal dependency. In order to appreciate the factors affecting SAR planning (task 4) it is important to gain a wider picture of the RCC's work flow, and the interaction and nature of the five high level tasks outlined in the task decomposition. Tasks 2, 3 & 4 in the diagram were considered to be knowledge intensive. Task 1 involves initial data acquisition and is not a significantly knowledge intensive component.

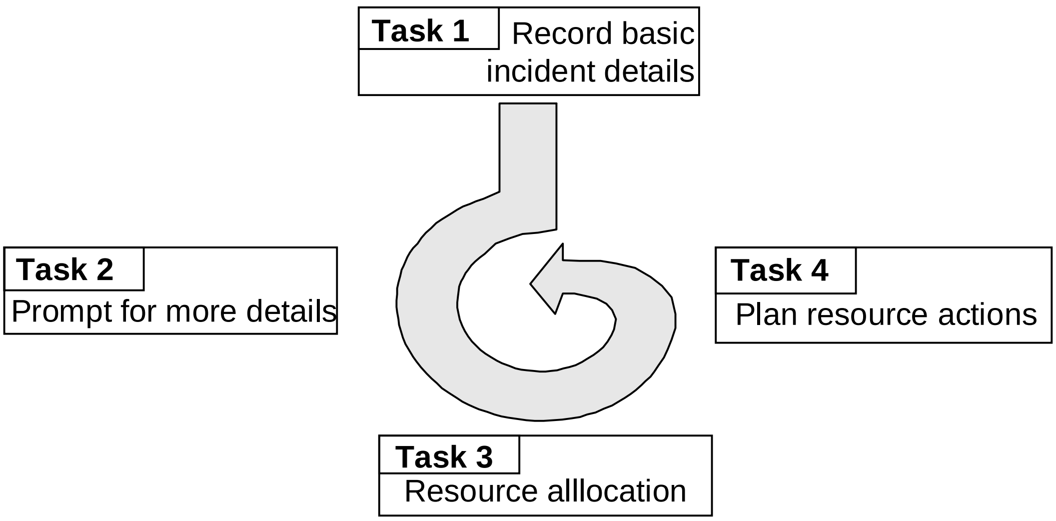

Fig 1: High Level Task Decomposition

Task 2 arises when the RCC actively seek more information about the present world state. It can be initiated from tasks 1, 3 or 4 due to the realisation that important facts are missing. These are facts that will be necessary for resource allocation and planning decisions to be made. When initiated from task 1 the RCC will make rapid decisions as to the nature and probable future course of the incident. This represents the anticipation of information that will be necessary for the RCC's problem solving associated with that incident. When initiated from task 3 or 4, task 2 represents the discovery that vital information is missing for the successful completion of problem solving within those tasks. It can be seen that if the entry point to task 2 is from task 3 or 4, then the nature of the missing facts has already been ascertained. Task 2 is therefore either straightforward or very knowledge intensive. In the latter case the reasoning is dependent upon experience and intuition. It is also likely to be a heavily interactive process, involving conversation with an individual from another authority (or resource).

Task 3 represents the RCC’s decision as to which resource to apply to an incident. It is initiated from task 1 or 4, and it is normally obvious to the RCC what type of resource to deploy and which particular resource to select. The type is governed by the classification of the incident. The particular resource selection is governed by the geographical position of the incident and the availability of resources. This only becomes reasoning intensive when resources are not freely available.
Task 4 is the RCC making decisions about the actions that resources should take when applied to an incident. This is where detailed planning takes place as is evidenced by the description of this task in the next section.

4. Knowledge Level Problem Solving Models

Within our discussion of planning for SAR (task 4), we shall refer to outline plan templates, executable plans, goals and actions. Goals are considered as being distinct from the actions that satisfy those goals. Planning within the RCC can be viewed as the process of moving from an outline plan to an executable plan. An outline plan template consists of a set of partially ordered outline goals that define the requirements of an incident. An executable plan is a set of ordered physical actions to be taken at the RCC. The reasoning process of the RCC enables this transition between outline plan template and executable plan.

4.1 A Knowledge level model for asset utilisation in SAR (task4)

Fig 2 shows a knowledge level model representing the inference layer for problem solving within task 4, from the perspective of a single search and rescue incident. The model represents the inference types and domain roles that exist within task 4 reasoning. The boxes represent domain roles for knowledge, and the ellipses represent inference steps. The shaded boxes represent support knowledge for a particular inference. The model is based on the KADS methodology [Breuker *et al.*, 1987.] and is expressed in domain terminology. The PSM was constructed through a lengthy process of second stage KA involving taped structured interviews, video tape analysis, protocol analysis, incident documentation and structured analysis of specific incident cases.

The model represents the process of reasoning from an outline plan to an executable plan; converting outline goals into an ordered set of physical actions to be carried out by the RCC. The detailed description of Fig 2 is given in section 4.3.

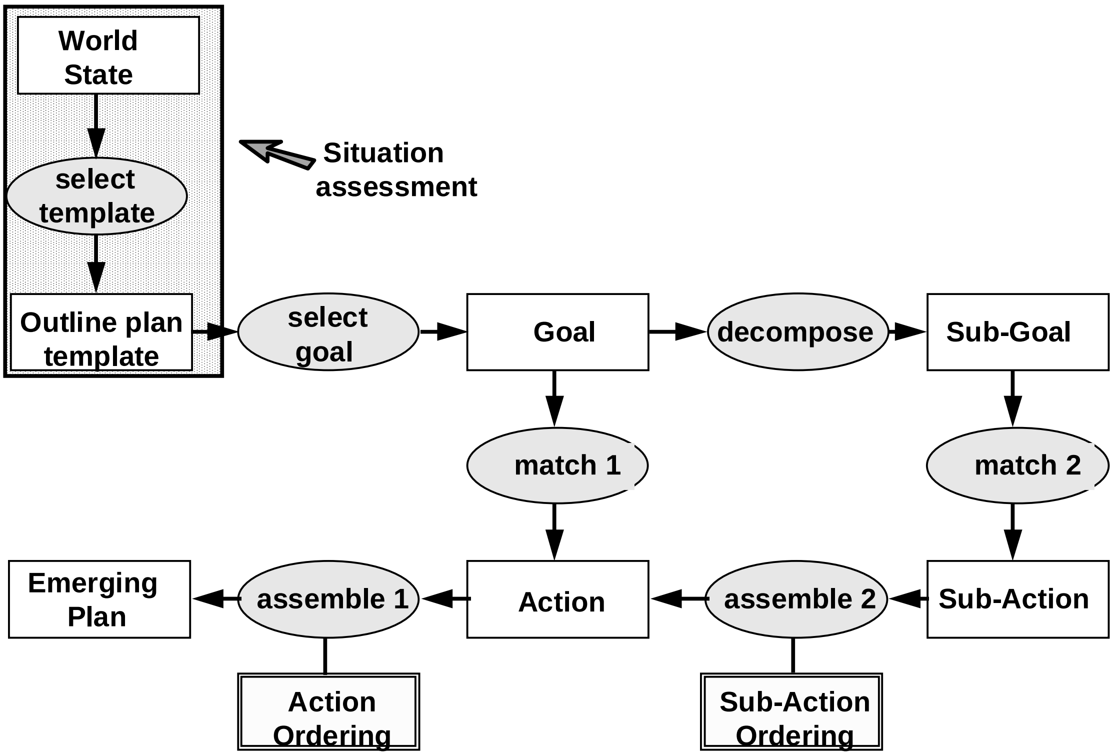

Fig 2: Domain Problem Solving Model

Planning for SAR is a progressive task that spans the temporal duration of a particular resource's application to an incident. This process involves the use of heuristic expert knowledge in order to make planning decisions in a domain where future data and constraints on planning are unpredictable. Due to this unpredictability, the decomposition of outline goals, and instantiation into planned physical actions, is usually not performed until the situation demands it. The RCC often resort to this least commitment strategy when planning. In this manner, the maximum amount of factual knowledge about the current situation is gathered before decision making. Oftentimes the RCC must hold back from taking physical actions, because if they wait for a small amount of time their factual knowledge of the situation will have increased so as to make a more effective decision possible.

4.2 Critiquing the domain PSM using a system derived PSM

The domain PSM (Fig 2) was validated by the expert, which involved lengthy discussion with multiple domain experts. This discussion was based upon relating the model back to specific incidents in SAR, in order to confirm that all cases could be characterised accurately within the model. A knowledge level methodology advises that the next stage in the development lifecycle should be the population and refinement of the PSM with domain knowledge. This then leads into the system design stage. At this point we departed from the suggested course of development. Rather than launching directly into the domain knowledge acquisition, we wanted first to consider operationalisation issues.

The reason for this is a common problem occurring in KBS development that relies upon a domain inferred PSM. Either during the design process or the actual implementation, it often becomes apparent that there are vital knowledge elements missing from the original PSM. This potential incompleteness of the PSM is a recognised problem in domain driven knowledge acquisition [Ford & Bradshaw, 1993]. Our proposed method of overcoming this was to select a generic system based planning PSM (as mentioned above), and attempt to establish a mapping into this from the domain based PSM. The rationale behind this was that the system PSM would be complete and sufficient, because it represented an operationalised architecture. If a clear mapping existed between the elements of the PSMs, then the system PSM may aid us in anticipating any omissions in the domain PSM, that would compromise its ability to produce decisions.

The system PSM that we selected for this purpose was derived from the Open Planning Architecture (O-Plan) system [Currie & Tate, 1991]. The O-Plan system is a generic computational planning architecture. The generic nature of the system and the fact that initial CommonKADS models of O-Plan had already been proposed [Kingston, 1995], made this an attractive option for our purpose. We established the existence of a clear mapping between the main ontological elements in the domain PSM and those in the O-Plan PSM (Fig 3). This approach successfully aided us in the identification of knowledge omissions in our domain PSM, prior to the completion of KA. This two model approach presents a number of advantages:

- comparison between domain and system based PSMs facilitates the early identification of omissions in the domain PSM, without enforcing a system structure upon it.
- the explicit mapping between the two types of PSM provides a mapping between domain-specific terminology and generic planning terminology.

4.3 A Knowledge level model for the O-Plan system

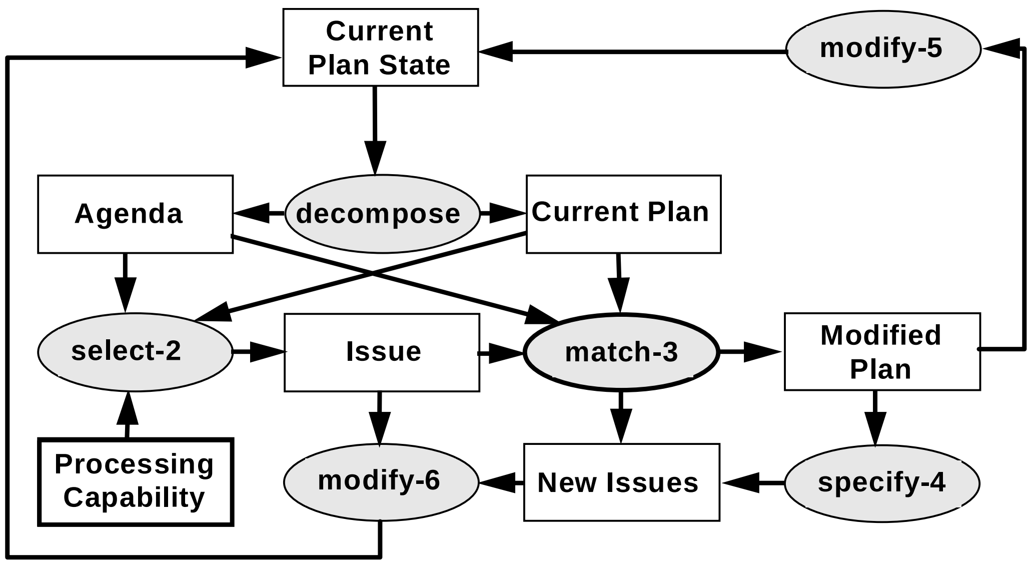

Fig 3: System Problem Solving Model

Both the system PSM (Fig 3) and the domain inferred PSM (Fig 2) possess inference steps which allow the transition from one form of knowledge to another. The different forms of knowledge are known as knowledge roles. In this section we describe the inference steps and knowledge roles in the domain based PSM, and how these items map to the O-Plan based PSM. Although the two models initially bear little resemblance to each other, there is in actual fact a clear mapping between both the ontological elements and inference steps that are depicted in the domain PSM to corresponding items within the O-Plan PSM. Both structures represent the matching of goals to possible actions. Both PSMs also facilitate the decomposition of goals and identify the selection of the next goal as an important inference step in the planning process. The comparison showed that the O-Plan PSM had a richer representation for the selection of goals, highlighting the necessity for knowledge that supports this inference step in the architecture of any intended system. The comparison had therefore successfully identified weaknesses in the domain based PSM.

The following is an explanation of the knowledge roles and inference steps in the domain PSM, accompanied by their mappings into the O-Plan PSM:

4.3.1 Goal selection

Select template -- The input to this is the world state, and the output is an outline plan template consisting of outline goals that correspond to a generic type of incident. They have a partial temporal ordering, yet it is only when the RCC plan the application of resources to an incident that these goals are more fully defined and ordered. In complex situations it may be the case that no outline plan template exists and one will have to be constructed. This then requires a knowledge of temporal constraints on the goals that make up an outline plan template. This template maps to the initial "Current Plan State" in the O-Plan PSM. The outline plan template is a set of outline goals to be resolved, and the agenda consists of a set of issues to be resolved. In this case, the goals in the domain map to outstanding issues in the O-Plan PSM. The outline plan template in the domain PSM is selected at the commencement of planning for an incident.

Select Goal -- This represents the selection of an outline goal from the outline plan template. The selection step is often simple, corresponding to the default ordering of goals defined in the outline plan template. However, in a complex or rapidly changing situation, the selection of goals becomes more knowledge intensive. It is here that our model comparison was informative, as it suggested the existence of certain types of control knowledge affecting goal selection. It is clear that in the O-Plan PSM there is a much richer representation of the knowledge affecting which goal/issue to resolve next. These are described in the O-Plan PSM as three possible expansions of the **match-3** inference step. The three expansions represent three different ways in which O-Plan can attempt the resolving of issues. The expansions are depicted in the O-Plan PSM as three separate inference structures. Two of these structures have clear mappings to the domain PSM (the decompose expansion is described in the next section). The third does not and represents knowledge about goal selection that is not accounted for in the domain PSM. In O-Plan this knowledge drives a backward chaining search process that decides which issues to resolve when the present issue's conditions are not satisfied. Issues are selected in order to achieve actions that will satisfy the original unsatisfied conditions. This backward chaining process caused by interaction between issues had not been considered in the domain PSM. All that had been considered was a knowledge of basic dependencies between goals, of the form Goal A must be satisfied before Goal B can be considered. The comparison between PSMs suggested a deeper form of knowledge about interdependencies between conditions necessary to resolve goals and the actions that satisfy them. This knowledge will provide important support to the user when incidents become complex and the default ordering of goals described by the outline plan template may not be applicable. The SAR domain has examples of the need for this "backward chaining" e.g. if the RCC are in charge of an incident involving mountain rescue in poor visibility the default ordering of the outline goals may not apply. The outline plan template for this incident places "scramble resources" before "ascertain weather". The outline goal "scramble resources" decomposes to "scramble helicopter", "scramble mountain rescue team" or both, depending upon the world state. A condition of this decomposition will be that visibility at the incident scene is sufficient for a helicopter to operate. If the visibility condition is unknown then the outline goal "ascertain weather" will be initiated in order to effect an action that will satisfy or negate this condition.

As mentioned the two other expansions of the **match-3** inference step have clear mappings to the domain PSM:

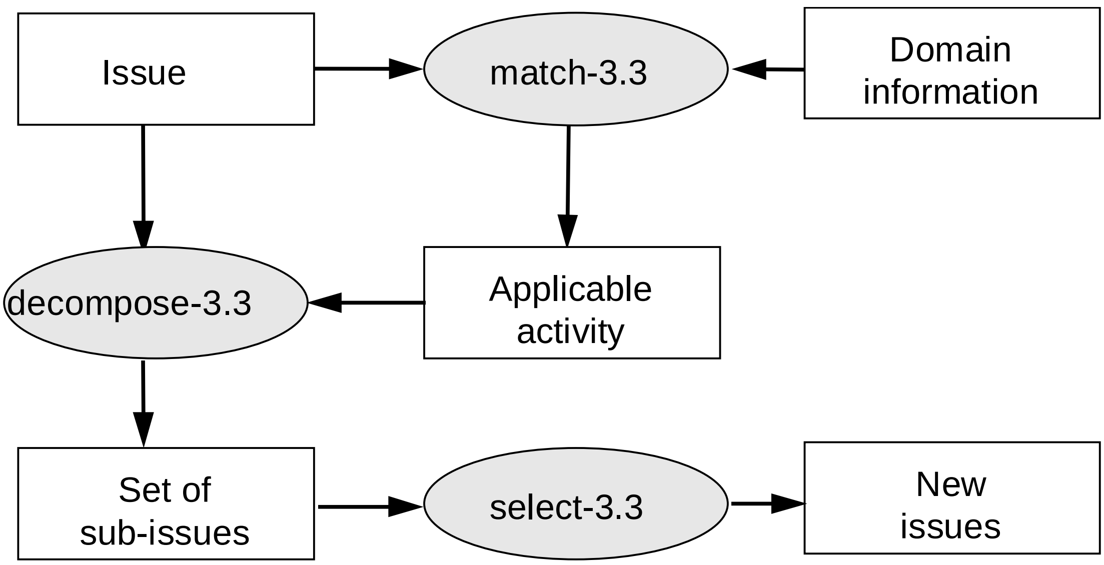

Fig 4: Expansion of match-3 for issue decomposition

4.3.2 Goal decomposition

Decompose -- This inference step decomposes a goal into a sub-goal. Similarly in O-plan issues may be decomposed into sub-issues. There is a clear correspondence with one expansion of the **match-3** inference step in the O-Plan PSM (Fig 4). The degree to which goals are decomposed by RCC varies. Some high level goals match to high level actions that consist of an invariant ordered set of physical actions. The existence of such invariant high level actions and their suitability to satisfy high level goals are factors affecting the degree to which goals must be decomposed. In some cases goals will have to be decomposed to the granularity of physical actions. It would seem that decomposition increases when an incident and its associated goals are out of the ordinary. Intuitively the level of reasoning increases in the exceptional cases.

4.3.3 Matching goals to actions

These inference steps involve finding actions that can fulfil (or help to fulfil) goals. There will generally be multiple actions that can fulfil a particular goal. The match therefore depends upon the present world state. There is a mapping between these two inference types and the third method of expanding **match-3** in the O-Plan PSM.

Match 1 -- The input to this is a high level goal. The output is a high level action. This step will require a knowledge of high level actions that satisfy high level goals, and the conditions that make the match valid.

Match 2 -- The input to this is a low level goal, corresponding to a physical action. The output is a lower level action. This is similar to the Match 1 inference type, though it is actually matching to a physical action.

4.3.4 Assembling the executable plan

Assemble 1 -- The input to this is a high level action, corresponding to an outline goal. This high level action consists of a set of ordered physical actions or of sub-actions. The output is the current executable plan. As discussed earlier the executable plan is gradually formulated throughout the course of the resources application to the incident. This is supported by a knowledge of action ordering, though a lot of this ordering will have been decided at the goal ordering stage.

Assemble 2 -- The input to this is a set of lower level actions; the lowest level being physical actions. The output is a higher level action; the highest being those that actually make up the executable plan (i.e. those that correspond to the high level goals).

There is no replication of these assembly inference steps in the O-Plan PSM, although typically any set of actions in O-Plan have ordering constraints instantiated during planning and are therefore assembled implicitly within the plan. The explicit assembly of actions is however important to the RCC, as it serves to summarise what has been done (or what is intended to be done), in order to achieve goals. This reflects the nature of the RCC's planning, which proceeds in small chunks corresponding to the outline goals in the template. There may be activity in several chunks at once, though this tends to be the exception rather than the norm.

There are clearly three inference steps in Fig 3 (besides the "backward chaining" expansion of **match-3**) that are not represented in the domain PSM. The **modify-6** step describes how the actions carried out in the plan modify the world state (world state is part of current plan state as it includes constraints). This could be included in the domain PSM as a step from emerging plan back to world state. The **modify-5** and **specify-4** steps describe how intended actions in the plan cause new issues to arise, and this therefore modifies the current plan state. These are not represented in the domain PSM, and this form of knowledge has not been observed in the domain. This is probably due to the previously mentioned least commitment strategy of the RCC decreasing the amount of intended actions that exist within their plans. This type of knowledge may be important in forms of SAR incident that we have not yet witnessed, and we regard this as an important area for future KA.

4.4 A Simple Problem Solving Example from the SAR domain

Figure 5 shows a typical outline plan template for a “routine or commonplace” incident (from the RCC point of view). It outlines the application of a single helicopter to an incident.

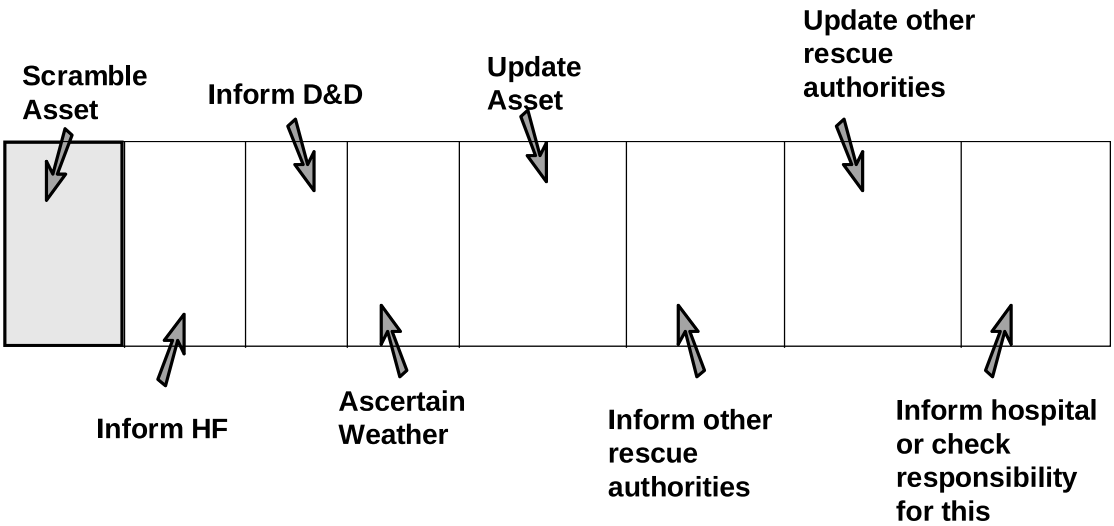

Fig 5: Outline Plan Template for a Mundane Incident

Scramble asset -- has been carried out in task 3 and need not be considered here.

Inform HF (high frequency radio operator) -- is quite straight forward and maps to a single physical action.

Inform D & D (Distress and Diversion) -- is also straightforward, merely requiring the decision as to whether this should be London Distress & Diversion, Scottish Distress & Diversion or both. The decision whether or not to set up a Temporary Danger Area, would considerably increase the complexity of this action.

Ascertain weather -- may seem a low level goal that easily translates into physical actions. However this can potentially become a complex knowledge intensive task. In an incident involving the use of helicopters, weather information may be crucial as it determines the operability of these assets. There are many aspects to weather information; the general weather situation, the weather at the incident scene, the weather on route to the incident, or at the helicopter base. How will these factors affect the assets potential utilisation? Obtaining accurate and up-to-date weather details is a non-trivial task. The complexity of the situation will reflect the number of sub-goals and the level to which these sub-goals must be decomposed before a set of physical actions are calculated to satisfy the goal.

Update asset (i.e. keep helicopters etc, informed of present situation) -- may also be complex as there are several possibilities for communication with resources, and communication may often be intermittent and unreliable. However there will be some default physical actions such as contacting the base or other authorities if communication cannot be secured.

Inform other rescue authorities -- involves reasoning about the likely progress and possible expansion of an incident, in order to consider all authorities who could potentially become involved and should therefore be informed.

Update other rescue authorities -- this is to some extent a continuation of the previous goal. It does however include updating authorities that are already involved, such as those that originally informed RCC of the incident.

Inform hospital or check responsibility for this -- is an example of a high level goal that translates to a high level action. For example if it has already been defined which hospital the casualty should be taken to, the high level action corresponding to this goal would consist of three physical actions:
(i). check responsibility for informing hospital by contacting the callout authority.
(ii). if RCC's responsibility, inform hospital of incident.
(iii). later inform hospital of casualty ETA at hospital and details of any injuries.

This is an example of a high level goal mapping to a standard set of physical actions. In other situations, where the hospital has not been decided, more decomposition and reasoning may occur. It can be seen in this simple example that even a mundane incident may require substantial problem solving in order to handle successfully.

5. The Modality of KBS support

Three possible modes of KBS support for task 4 were identified:
(i). **Visualisation and representation** -- this involves computerising the information stores and representations used by RCC, as well as the development of additional representations and visualisations for entities that are not currently stored. The most important of these would be plans themselves.

(ii). **Analysis** -- this involves representing the plans that the RCC follow, and using knowledge-based planning techniques to keep the computerised plans up to date with the present state of plan intention and execution. This state is supplied by the users of the system yet advice can be offered about plan formulation, potential interaction between plans, and weaknesses of individual plans.

(iii). **Synthesis** -- this involves the actual production of plans.
We decided that for the SAR domain our efforts would be most effectively applied in support of the first two of these modes.

6. Implementation

The structure of the acquired PSM was then carried through into the design and system implementation. The SAR system was developed using CLIPS and HARDY. HARDY is AIAI's hypertext and diagram tool, which is integrated with CLIPS (a rule based programming language developed by NASA). HARDY was used for the front end of the system, while CLIPS was used to implement the actual knowledge based planning functionality.

When the system was designed, it became apparent that there were two options for implementing the planner. The first option was to implement the acquired goal/action matching rules directly in CLIPS; the second option, was to encode these acquired rules and their conditions as possible actions. These "declarative rules" would then be activated by a set of "meta" rules. In the latter approach, planning consists of matching the conditions of a possible action with the current state of the world; if the conditions are matched, then that action can be introduced to the plan. If the conditions are not matched, further processing can be carried out to see if the relevant goal can be decomposed, or if another action can be introduced to the plan as an issue in order to change the state of the world (this is the "backward chaining" described earlier).

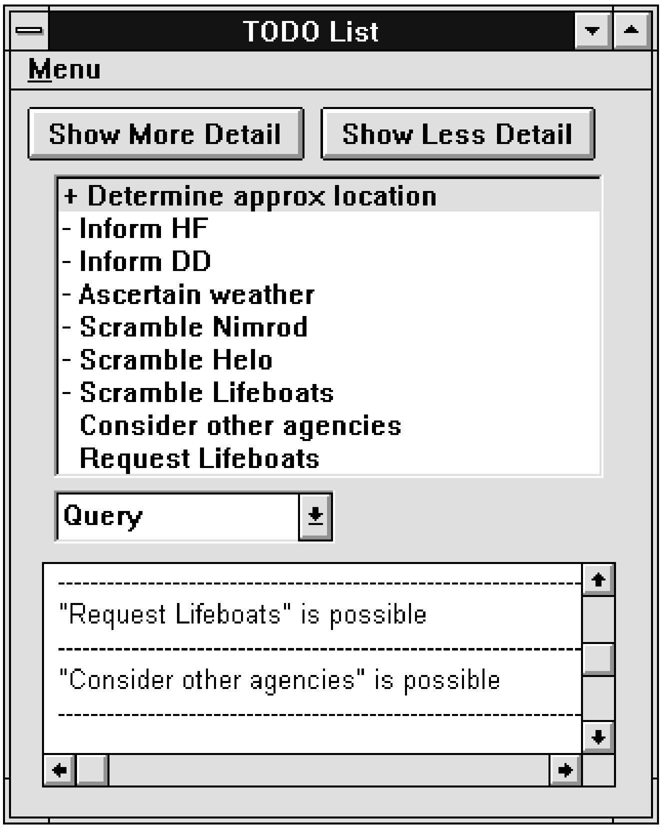

Fig 6: Rescue System TODO List

The second approach was chosen. A set of "meta" rules were written in CLIPS, which match a set of "possible actions" (encoded as CLIPS instances) against data about the search and rescue incident and the availability of resources (also represented as CLIPS instances). This “meta” rule approach enables a virtual planning architecture to be imposed upon the CLIPS language. Another set of CLIPS functions maintains a set of HARDY nodes and arcs which provide a number of different views on the plan and the underlying data. These include a “TODO list” (Fig 6) and a PERT chart (Fig 7) which provide the user with access to and views of the (partially developed) plan. The system facilitates fast information input and plan navigation as well as interactive plan generation with the user.

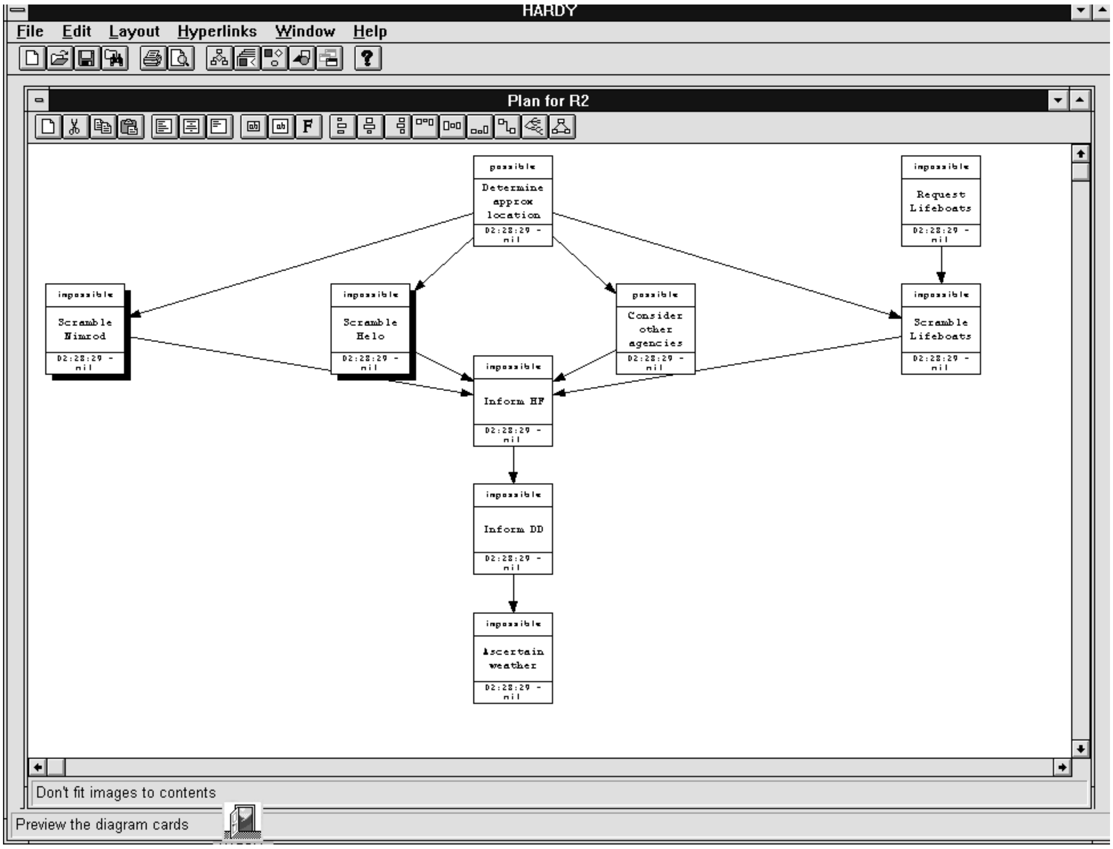

Fig 7: Rescue System Pert Chart (Plan Representation)

These are linked to a ”status board” and an “incident form” which are both duplications of representations currently used by the RCC. All of these views, apart from the PERT chart, allow input to the knowledge base as well as displaying output.

The system interfaces with another software package (AutoRoute™), which provides a computerised map display of the location of an incident and of resources. The communication between the planning system and AutoRoute™ is handled using ODBC™ functions which are built into HARDY.

7. Conclusion

The work described represents the construction and demonstrated use of a domain driven knowledge level modelling approach to KBS development in a planning domain. We make a definite distinction between problem solving models that are inferred from the domain and those that have been derived from systems. Such domain derived models do not presently exist for the support of system development within the generic area of planning. Our work merges a knowledge level modelling approach with the work that has been done on ontologies for planning, in order to formulate a generic approach to the acquisition and utilisation of knowledge for planning systems. The approach was tested and refined through the development of a knowledge based system for the support of planning for search and rescue.

The distinction that we have made between domain and system based models, led us to investigate possible benefits that could be gained by exploring the mappings between these models. In the context of the search and rescue development, we discovered that the comparison of these models enabled us to identify omissions in the domain model without biasing its structure to that of the system model. It also enabled us to identify specific areas for future KA. We believe this twin model approach may have more general applicability for KBS development, particularly in situations where no generic model is available.

During a KBS development lifecycle there must be iteration between the developmental stages. Later stages of development iteratively inform and revise the earlier stages. It is expected that this iterative cycle will improve the final product. The price paid is that a large amount of iteration in the development lifecycle increases the effort expended. For this reason approaches that enable the detection of shortcomings in the earlier stages represent a potentially large saving in development effort.

The two model approach described demonstrates such a saving in the case of domain driven knowledge level modelling for planning. It should be stressed that this approach is aimed at model validation when no suitable generic PSM exists, and the intention would be to arrive at a single PSM that could then be added to the library.

REFERENCES

Breuker, J. and van de Velde, W.(Eds.). (1994)
CommonKADS Library for Expertise Modelling.
IOS Press, Amsterdam, September 1994.

Breuker, J., Wielinga, B., Van Someren, M., De Hoog, R., Schreiber, G., De Greef, P., Bredeweg, B., Wielemaker, L., Billault, J-P., Davoodi, M. and Hayward, S. (1987)
Model driven knowledge acquisition: interpretation models.
KADS-I Project Deliverable, University of Amsterdam, Holland 1987.

Clancey, W.J. (1985)
Heuristic classification.
Artificial Intelligence, 27: 289-350, 1985.

Currie, K.W. and Tate, A. (1991)
O-Plan: The Open Planning Architecture.
Artificial Intelligence, 52 (1): 49-86, 1991. Also available as AIAI-TR-67.

Ford, K.M. and Bradshaw, J.M.(Eds.). (1993) *Knowledge acquisition as modelling.*
New York: Wiley. 1993.

Fraser, J. and Tate, A. (1995)
The Enterprise Tool Set - An Open Enterprise Architecture

Proceedings of the Workshop in Intelligent Manufacturing Systems, International Joint Conference on Artificial Intelligence, Montreal, Canada. 1995.

Kingston, J.K. (1995) *CommonKADS Models for Planning tasks.* forthcoming AIAI technical report, 1995.

Lee, J. (1994) *"Process Interchange Format Proposal" Technical Report,* ICS-TR-94-27, University of Hawaii, Dept. Information & Comp Sci. 1994.

Musen, M., Gennari, J., Eriksson, H., Tu, S., Puerta, A. (1995) PROTOGE-II: *Computer Support for Development of Intelligent Systems From Libraries of Components.* Proceedings of the Eighth World Congress on Medical Informatics, Vancouver, Canada. 1995.

Motta, E., O'Hara, K., Shadbolt, N., Stutt, A. & Zdrahal, Z. (1994). *A VITAL Solution to the Sisyphus II Elevator Design Problem.* Proceedings of the Eighth Banff Knowledge Acquisition for Knowledge Based Systems Workshop, Banff, Alberta, Canada. 1994.

Newell, A. (1982) *The knowledge level.* Artificial Intelligence, 18: 87-127, 1982.

O'Hara, K., Shadbolt, N. R., Laublet, P., Zacklad, M. & Le Roux, B. (1992). *Knowledge acquisition methodology.* VITAL deliverable DD212. Nottingham University, UK. 1992.

Tansley, D., and Hayball, C. (1993) *Knowledge Based Systems Analysis and Design: A KADS Developer's Handbook.* BCS Practitioner Series, Prentice Hall. 1993.

Tate, A. (1994) *"Plan Ontology", a paper to the Workshop on Ontology Development and Use.* San Diego, California, USA, November 1994. (Also available as AIAI Technical Report AIAI-TR-161, and through URL: ftp://ftp.aiai.ed.ac.uk/pub/document/1994/94-krsl-plans.html)

Tate, A. (1995) *Characterising Plans as a Set of Constraints - the <I-N-OVA> model - a Framework for Comparative Analysis* ACM SIGART Bulletin, 6 (1), January 1995.

Workflow Management Coalition Members. (1994) *Glossary: A Workflow Management Specification*. The Workflow Management Coalition, Belgium:Brussels, November 1994. (Also available through URL: http://www.aiai.ed.ac.uk/WfMC/)