Decision-Centred Visualisations for Tactical Decision Support on a Modern Frigate

Bruce A. Chalmers*
DREV, Decision Support Technologies Section
2459 Blvd. Pie XI (North)
Val-Bélair, Québec
Canada G3J 1X5
(418) 844-4000 (x 4698)
bruce.chalmers@drev.dnd.ca

James R. Easter**
Logica Carnegie Group
Advanced Decision Support Business Unit
5 PPG Place, Pittsburgh
PA 15222, U.S.A.

Scott S. Potter**
Logica Carnegie Group
Advanced Decision Support Business Unit
5 PPG Place, Pittsburgh
PA 15222, U.S.A.

Abstract

Decision support is a focus of attention for the mid-life upgrade of the Combat Control System (CCS) of the Canadian Navy's HALIFAX Class frigate. We are exploring concepts for developing decision aids to assist the ship's Command Team with tactical Command and Control. This paper proposes a nonlinear, empirical framework for the investigation. It then gives an overview of work aimed at assessing the feasibility and value for its analysis and design activities of a Cognitive Systems Engineering framework, known as Cognitive Work Analysis (CWA), for modeling intrinsic work demands and determining computer-based support interventions incorporating advanced decision aids to support these demands. It focuses on one aspect of the CWA feasibility study looking at the use of CWA for deriving decision-centred visualisations to enhance operators' situation awareness and action responses. This leads to preliminary requirements for an interface for the ship's tactical coordinator that have been captured in a storyboard.

1. Introduction

Technological developments in naval operations into the next century are expected to lead to increasing demands on information management and command decision making. Studies of the Canadian Navy's operational requirements for Maritime Command and Control in the 2010 time frame have recommended increased emphasis in the areas of data fusion and decision support to the shipboard Command Team as a means of increasing their battlespace awareness and decision speed and accuracy in the face of these challenges [Bertrand *et al.*, 1999]. In response, Command Decision Aid Technology (COMDAT), a pair of major R&D projects referred to as COMDAT I and COMDAT II, is being defined for the 2000-2006 time frame, to develop technology demonstrators to support requirements definition of the mid-life upgrade of the Combat Control System (CCS) of the Canadian Navy's HALIFAX Class frigate.

---

* Effective July, 2000, with Defence Research Establishment Atlantic, 9 Grove Street, P.O. Box 1012, Dartmouth, Nova Scotia, (902) 426-3100, bruce.chalmers@drea.dnd.ca.

** Now with MindSim Corp., 501 Grant St. #475, Pittsburgh, PA 15219, U.S.A., (412) 471-3456, jeaster@mindsim.com, spotter@mindsim.com.

Most tactical decision making in a modern frigate like the HALIFAX is performed within the ship's Operations Room. There, a team of operators interact with a CCS through consoles, aided by a number of other systems. The CCS lies at the heart of the combat system which provides the ship's warfighting functionality. It provides access to services for building and displaying a tactical picture and utilising the ship's capabilities for the basic purpose of threat neutralisation. Built around early 80's technology, the HALIFAX's CCS emphasises primarily the automatic control system, the so-called Threat Evaluation and Weapon Assignment (TEWA) system, which was, and still is, performing mostly numerical and simple rule-based calculations within the detect-to-engage cycle, while most of the higher-level cognitive activities are left to a team of operators.

COMDAT I will concentrate on enhancements to the CCS from incorporating multi-source data fusion. The aim is to improve the integration, at the level of the CCS, of organic and nonorganic data available to the ship and provide automated capabilities for developing a single integrated Maritime Tactical Picture. COMDAT II will demonstrate advanced decision support capabilities for enhancing operator performance in their situation and threat assessment (STA) and response identification, selection and management (RM) decision-making activities [McArthur *et al.*, 1999].

The purpose of this paper is twofold. First, we propose a nonlinear, empirical framework as the basis for exploring the development of advanced decision aids to be incorporated in a future decision support system (DSS) for augmenting the decision support capabilities of the HALIFAX's CCS. Conceptually, this DSS will provide interfaces and decision aids to support a variety of the Command Team's perceptual, cognitive and collaborative activities in conducting tactical Command and Control (C2) in its various warfare areas of air, surface and sub-surface warfare. Second, we present an overview of a study to assess the feasibility and benefits of a specific Cognitive Systems Engineering (CSE) approach to work analysis, known as Cognitive Work Analysis (CWA) [Rasmussen *et al*., 1994], [Vicente, 1999], for use in the analysis activity of the nonlinear framework. We emphasise one aspect of the feasibility study, concerned with using CWA to identify preliminary requirements for a threat management interface for the the ship's tactical coordinator, the Operations Room Officer (ORO). We also describe the decision-centred visualisations for the interface that were developed from these requirements.

CWA has been specifically developed to meet the challenges of complex sociotechnical systems, but it has so far received limited attention by the C2 R&D community [Vicente, 1999]; however, see [Rasmussen, 1998] and Chin *et al*. 1999] for some recent exceptions. While there are many reasons to believe that there is a good fit between the demands imposed by C2 and the characteristics of CWA, it is important to provide a basis for evaluating this degree of fit. There are a number of benefits from achieving this objective. First, we will have an improved understanding of the value of CWA to ongoing efforts in Canada to develop decision aids suitable for integration in the CCS. Second, the beginnings of a CWA for the HALIFAX Class ship will be developed. Third, the potential of CWA will be concretely illustrated by identifying preliminary requirements for an interface for the ship's ORO to display task-relevant domain semantics and support information extraction by the operator.

The paper is organised as follows. Section 2 argues for a holistic system perspective on design of computer-based decision support, accounting for characteristics of the workspace and its cognitive demands, the characteristics of human operators, and the opportunities afforded by the technological solution space for supporting operator demands. Section 3 discusses the aims of the R&D work described in this paper. Section 4 presents a nonlinear framework for exploring concepts suitable for implementation in computer-based artifacts that support operators in their cognitive work. Section 5 examines the Cognitive Work Analysis framework and motivates its consideration as a candidate for the analysis component of the nonlinear framework of Section 4 in the context of the HALIFAX's work environment. Section 6 gives an overview of the CWA feasibility study for the HALIFAX Class ship which was conducted in two parts. The method that was followed is described, along with the general nature of the models that were obtained in the first part. Section 7 looks in more detail at the second part of the feasibility study and describes how CWA models were used to first map the cognitive demands of the operational environment and then design an interface that supports information visualisation and enhances the ORO's battlespace awareness. Finally, Section 8 provides conclusions.

2. **Need for a Holistic System Perspective on Design of Decision Support**

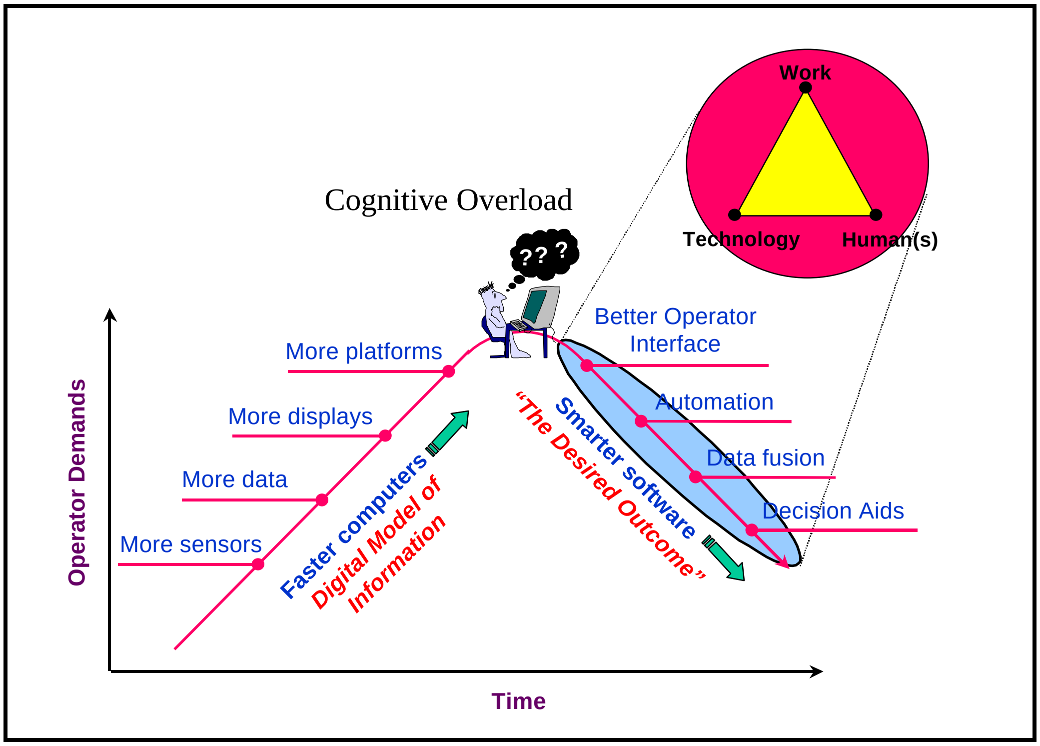

Figure 1. Need for a holistic perspective on design

The widely accepted technology-driven response to satisfying operators' information needs for decision making and problem solving in warfare is founded on the premise of a digital model of information that focuses essentially on amassing increasing amounts of data and processing this data into more information, as the basis of reducing battlespace uncertainty for commanders. While certainly providing important capabilities for increasing data coverage and accuracy in the dynamic representation of the battlespace, this effort alone, as characterised by the left portion of the curve in Fig. 1, can easily have the opposite effect to that desired by contributing instead to additional cognitive demands on operators, and, particularly in critical, high-tempo periods, even greater uncertainty as operators deal with the flood of increased data and making sense of all this data.

[Woods, 1985] draws attention to a host of potential errors in situations of data overload and limited human attentional resources, including tunnel vision, fixation, and the inability to differentiate the relative importance of data, which has the effect of reducing operator performance instead of enhancing it. [Schmitt and Klein, 1996] have noted that the view that technology will blow away the fog of war is a dangerous delusion. They argue that the digital model of information ignores the higher, less tangible levels of uncertainty that commanders have to deal with and may in fact increase their uncertainty about the battlespace. In Schmitt and Klein's view, the flaw in the digital information model is that it fails to come to grips with the decision maker's real needs for knowledge and understanding to be able to make judgements and decisions in the face of the inevitably significant amounts of uncertainty that are a natural product of the dynamic of war. Because there will always be uncertainty that cannot be removed in the model, they also argue that the human decision maker has ultimately to be able to cope in spite of uncertainty.

More recent technological efforts, based on developing data fusion and decision aiding technologies, promise significant opportunities for coping with the data explosion problem in modern warfare. They are shown on the right portion of the curve of Fig. 1. They offer the potential for improving data processing capabilities by enhancing data integration, and thereby reducing cognitve demands of operators and improving their situation awareness. However, realising the promised benefits of these newer software-based solutions is extremely challenging. While representing the desired outcome from algorithmic innovations, the benefits implied by right portion of the curve in Fig. 1 are certainly far from a given.

The literature offers numerous examples that attest to the difficulty of providing effective support in cognitively demanding work environments. In fact, there is increasing evidence that poorly engineered computer-based solutions can lead to substantial performance decrements of the joint human-machine system and even potentially catastrophic results. As examples, the burden associated with supervising automation as it performs an offloaded task can outweigh the benefits to improved performance [Kirlik, 1993]; performance decrements can result from automation-induced complacency (over-trust) [Parasuraman *et al*., 1993]; a partially automated system can induce more errors in cases where its knowledge is incompetent than if the operator is left in the loop and the system simply critiques the operator's performance based on recognising violations of intrinsic work constraints [Guerlain, 1993]. Examples of catastrophic results in the aviation and nuclear power plant industries can be found in [Endsley, 1996].

It has been suggested that when tools dominate, rather than constrain, the joint human-machine system, the designer runs a strong risk of solving the wrong problem, and of creating new problems and undermining existing work strategies [Woods and Roth, 1988]. [Woods and Hollnagel, 1987] emphasise that the key to the effective application of computer technology for supporting the human in complex problem-solving worlds is to conceive, model, design, and evaluate the *joint* human-machine cognitive system. [Klein, 1997] argues for the need to incorporate an understanding of the expertise of operators in information technology solutions so that such expertise can be productively employed and reinforced. Barriers to expertise that can be posed by information technology include excessive data, preprocessed data, excessive procedures, performing formal analyses, passive data handling, and limited ability for information seeking.

In summary, a broad, coherent consideration of the characteristics and demands of the sociotechnical work system of Command and Control is essential for success. An overarching goal must be to design the information and decision support to couple human operators and technologies into collaborative, adaptive control systems that can function effectively in the complex, dynamic work environment of the maritime battlespace. Considering the practical aspects of applying DSSs in these types of environments, Hollnagel [Hollnagel, 1988] suggests that a key consideration must be to provide the 'right' information (contents, meaning) at the 'right' time (in relation to the decision) and in the 'right' way (format, context) to achieve effective performance. This is not as easy as it may seem. Based on Wood's relational perspective [Woods, 1986] on the nature of information, this hinges on specifying what 'right' means in the dynamic context of the work environment, relative to the operators' expectations, intentions and interests. Designing displays and decision aids to support information extraction, instead of as a consequence of data avaliability, based on representations that reveal task-relevant domain semantics, emerges as a critical part of improving performance. Two other important dimensions of the joint human-machine performance problem, which Woods and Roth [Woods and Roth, 1988] refer to as the 'cognitive system triad', relate to the cognitive demands produced by the work domain and the characteristics of the cognitive agent(s) that must meet these demands.

A critical insight to be derived here is that the characteristics of the workspace and its cognitive demands, the characteristics of human operators, and the opportunities afforded by the technological solution space for supporting operator demands, are in fact so intricately intertwined that successful navigation of the solution space requires considering a number of complex, cross-disciplinary issues in a holistic manner, and at the outset, in dealing with the design of decision support. These are some of the significant features that distinguish the Cognitive Systems Engineering approach we are examining in this R&D effort.

3. **R&D Goals**

It is evident from our remarks in Section 2 that this is a highly multifaceted problem that cannot be resolved by focusing on technological advances alone, and that a variety of factors must be considered. For the designer, this translates into a need for greater understanding of the demands of the work environment and its complexity. Solutions need to take into account the human dimension and the technological dimension, in the context of the work that must be performed.

The general directions of the R&D effort described in this paper are firstly, in achieving a better understanding of what computer-based technology can offer as a means of augmenting the performance of the combined system consisting of operators and computer-based interventions in their workspace, and secondly, in turning such insights into viable computer-based solutions.

The work environment of the HALIFAX Class frigate is a tangible example of C2 decision support applications comprising complex human-machine systems where success depends heavily on the human operators' ability to efficiently cut through masses of data, visualise the state of the world and courses of action, cope in spite of uncertainty, and collaborate as needed in a fast-paced, highly adaptive multi-person environment. The critical, high-level computer-based decision support requirements to permit transforming the environment from an inefficient, data-intense, high cognitive demand situation to an efficient, information-rich, high-performance human-machine system can therefore be summarised as follows:

- provide for the information needs of operators and help them with understanding the situation and in deciding on their responses;
- support the communication and collaboration needs of the environment;
- provide flexibilty for adaptation and change; and
- support operators in coping in spite of uncertainty.

4. A Framework for Concept Exploration of Computer-Based Artifacts for Cognitive Support

4.1 *Motivation*

In some of our earlier work (e.g., [Chalmers, 1997]), we have alluded to the ill-structured nature of the problem of designing decision support systems for this work environment, likening it to solving a jigsaw puzzle consisting of uncertain pieces and an uncertain goal picture. [Alexander, 1964] describes design more generally as an effort to achieve fitness between two intangible entities, a form that has not yet been described and a context which cannot be properly described. Klein [Klein, 1996] has suggested the need for nonlinear problem-solving frameworks in dealing with ill-structured problems, to allow complete descriptions of problem detection, problem representation and problem-solving activities, and their interrelations. Compared with problem-solving approaches for well-structured or routine problems, this is intended to reflect a greater focus on explicitly dealing with the ill-structured nature of the problem by continually modifying goal definition and problem representation throughout the problem-solving process, rather than just at the beginning, based on continual feedback and insights gained from intermediate, tentative 'solutions'. According to Klein, nonlinear problem solving is:

- non-proportional: minor input changes can produce large output changes;
- interactive: problem representation may need to be continually changed throughout the problem solving process;
- open: problem solving may be highly opportunistic; and
- constructive: successful courses of action are constructed by recognising and exploiting critical, local features of the problem landscape rather than by searches in some predefined problem space.

We propose a view of concept identification, exploration and evaluation for this work environment that takes advantage of these insights. In particular, we suggest the need for a nonlinear, empirical framework to guide the exploration and development of advanced decision aiding concepts for incorporation in a computer-based system to support and augment the performance of the Command Team of the HALIFAX Class frigate. We formalise this in the next section.

4.2 *The Framework*

At a high level, we conceptualise an appropriate nonlinear framework as one that encompasses an opportunistic, dynamic evolution of multiple interacting activities, with continual goal and activity refinement, based on increasing our understanding of two key spaces. We refer to these as the problem space and the artifact space. The problem space is the work environment, which is the sociotechnical work system made up of the work domain, its workers or operators, their organisation, policies, doctrine, tactics and procedures. The work domain is the system being controlled, independent of any particular operator, automated controller, event, task, goal, or interface [Vicente, 1999]. The artifact space is our solution space. While it is certainly necessary to take a broad view of the solution space, we focus for the purposes of the discussion here only on solutions in the form of computer-based artifacts. However, since the introduction of any computer-based intervention into an existing work environment can have the effect of changing the work to be performed, impacting the nature of the work, its organisation, its demands, and the quality of its performance, both human-related and technology-related perspectives need to be considered throughout the process.

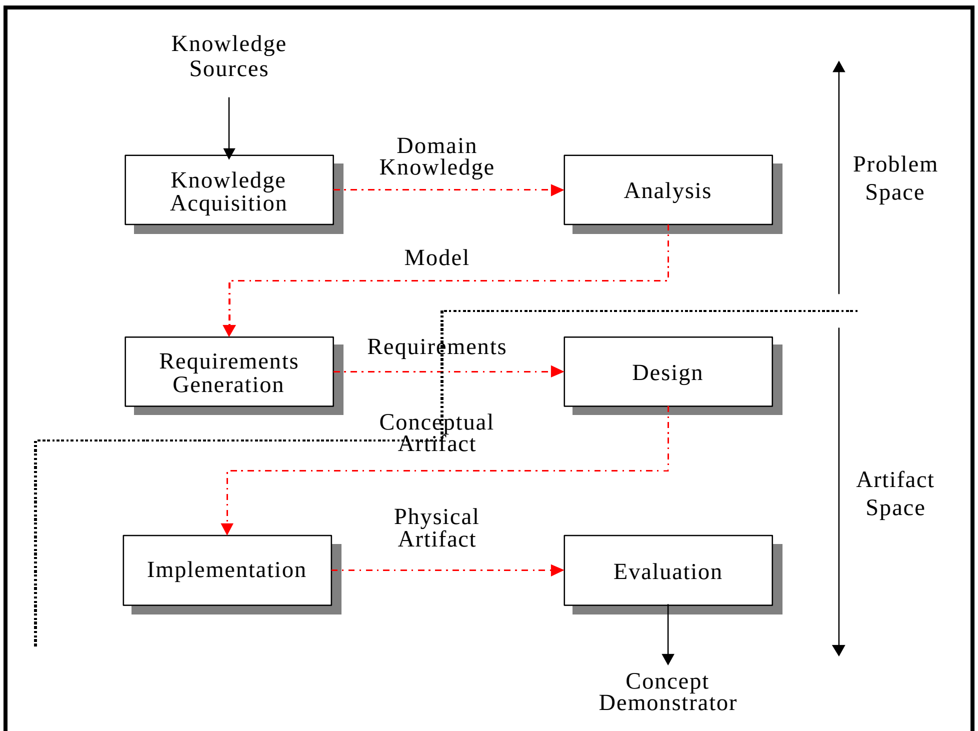

Figure 2. One possible trajectory in the technology demonstration 'life cycle'

A review of a variety of existing (essentially linear) systems engineering frameworks suggests that candidate activities in such a framework are: domain knowledge acquisition, analysis, requirements generation, design, implementation and evaluation. For concreteness, Fig. 2 illustrates just one activity trajectory within the framework (shown dashed in red), showing the activity nodes and their linkages in terms of inputs and outputs. In general, trajectories exploring the problem and artifact spaces will be considerably more chaotic than this.

A key output of the process, shown in Fig. 2, is a demonstration of some advanced computer-based decision support concept. In the sense we use the terminology here, a concept demonstrator is not an operational prototype of a support concept, but a means of capturing and instantiating design knowledge in a tangible, validated artifact that can be used as a basis for further exploration and development of an operational prototype. The process of producing an operational prototype should also benefit from adopting an analogous framework. However, if this current exploratory research phase serves its purpose, one would expect that producing an advanced prototype is going to be relatively more linear since by this time the concept demonstrator has contributed to structuring and representing both the specific problem being addressed and its design solution.

It is important to describe the dynamic, nonlinear aspects of the framework, since these are the critical features of what we are proposing here. Specifically, we refer to its opportunistic, bootstrapping, incremental and evolutionary nature. Several opportunistic activity sequences, formed from the nodes in Fig. 2, are conceivable (but not shown there). This is motivated by a number of factors, including where in the 'life cycle' a particular concept is under investigation for consideration in a technology demonstration, the intent of that demonstration, and the need to repeatedly shift attention between the two spaces determined by the problem space and the artifact space. This happens as design hypotheses, generated through analysis in the problem space, lead to detection and framing of the problem, fleshing out design goals for tentative solutions, testing hypotheses in the artifact space, feedback, and so on. The partitioning of the activity nodes between the problem space and the artifact space is illustrated in Fig. 2.

The bootstrapping characteristic of the proposed framework reflects the fact that each step of the process builds on the insights and products of earlier ones in an incremental manner. Furthermore, although these various activities are encapsulated in generically identified activity nodes in Fig. 2, their specific details (e.g., techniques used, strategies employed) on a given pass are linked to the actual approaches that need to be adopted on that pass and are not predetermined. This explains the evolutionary nature of the framework.

This framework is suggested as a basis for developing and instantiating advanced aiding concepts in a manner that does not lose sight of the domain's characteristics and the requirements of its operators, and for assessing, to the extent possible within the scope of the demonstration, the value of a concept and the merit of its further development or refinement.

The purpose of knowledge acquisition is to produce domain knowledge as needed for the other activities. This can involve a variety of techniques, including interviews or knowledge elicitation sessions with human subject matter experts (SMEs), verbal protocols, naturalistic observations of training or actual work situations, and consultation of procedural and system design documents (for the existing system), depending on the sources of the desired domain knowledge. Working with SMEs to construct work scenarios, or to elicit feedback using storyboards, mock-ups or other types of artifacts, or conducting a sea trial of a concept demonstrator, can involve aspects of both evaluation of an existing concept and an opportunity to pursue further knowledge acquisition (in this case, knowledge elicitation from human sources), leading to further analysis and problem representation, and so on. This is exemplary of the opportunistic and bootstrapping nature of the framework.

Separating knowledge acquisition from analysis in the framework emphasises the different natures of these two activities and makes the seldom made distinction between domain description, which is the aim of knowledge acquisition, and modeling, which is the purpose of analysis. However, the two are evidently closely related - in the reverse direction to their link shown in Fig. 2, the domain knowledge to be acquired, and hence the type of knowledge acquisition to be conducted at a given time, is determined by the specific needs of the analysis that is being pursued.

4.3 *Working out the Details*

We have focused on the activities in the framework, and their high-level characteristics. Establishing the specifics of each activity and the tool sets and infrastructures needed to pursue and support these activities requires a lot of empirical work, particularly given the web of potential interdependencies that exists between them. However, there is already a vast array of techniques one can expect to draw on, or adapt as needed, from the current knowledge acquisition, systems engineering and software engineering literatures, that should significantly help avoid reinventing the wheel. In addition, there has been a lot of recent focus in the cognitive psychology community on developing techniques for knowledge elicitation (e.g., see, [Cooke, 1994], [Hoffman *et al*., 1995]).

The remainder of this paper is concerned with a specific framework, and its application, for the analysis activity alone, that of Cognitive Work Analysis (CWA). CWA is a framework for modeling work demands in complex sociotechnical systems that arose from work first done in the nuclear power plant domain [Rasmussen *et al*., 1994], [Vicente, 1999]. As already remarked, the role of analysis is to generate design hypotheses which can be tested in the artifact space. This is akin to Woods' notion of design to support cognitive work as hypotheses about how artifacts shape cognition and collaboration [Woods, 1998]. More specifically, as Fig. 2 indicates, analysis is a modeling process whose outputs permit representing requirements as a basis for design. Importantly, while requirements shape design, in the end, design is usually underconstrained by statements of requirements. The process of turning requirements into design inevitably requires a lot of creativity on the part of the designer [Vicente, 1999].

5. **Cognitive Work Analysis**

It is accepted in the behavioural design disciplines that a good understanding of the work to be performed, and more specifically its demands, are important prerequisites for designing effective support intervention in a work environment [Vicente, 1999]. However, there is much less agreement on the nature of the analysis needed, what is the grain of that analysis, and what is a suitable framework for its conduct.

In this section, we motivate our consideration of CWA as a candidate for the analysis component of the empirical framework of Section 4 and briefly review some of its key ideas. We are not suggesting that this is the only analysis framework that should be considered for such a role, or that there may not be gains from employing some other approach, depending on the specific aspects of the work environment to be modeled, the characteristics of the portion of the workspace under consideration, or even the constraints on the design solution itself that are being imposed. In fact, it would be desirable to evaluate any proffered analysis framework for the capability of its modeling power to lead to effective design interventions in this work environment, and why, and to make comparisons between a variety of such frameworks for their strengths and weaknesses under a variety of scenarios. However, planning and conducting these types of experimental comparisons in a scientific manner would be technically challenging, expensive, and time consuming. As Miller and Vicente indicate in recent work comparing aspects of differing methods of analysis ('task' versus 'work') [Miller and Vicente, 1999], one also has to deal with the confounding factor arising from the impact of the designer's creativity in the transition from requirements to design.

There are also various pragmatic questions related to relative costs (based on time, required effort, risk, etc.) and the gains in using one approach over another. These remarks suggest the need to understand at the start the rationale for making interventions in the workspace and for identifying any constraints that can influence the choice of approach.

5.1 *Why Consider CWA in this Work Environment?*

Our reasons for considering CWA are related to the characteristics of the work environment being explored and the limitations of other work analysis approaches for modeling work demands in the presence of these characteristics. CWA's modeling layers are reviewed in Section 5.2.

Shipboard C2 possesses many features that usually characterise a complex sociotechnical work system. Such characteristics include (e.g., see [Rasmussen *et al*., 1994], [Vicente, 1999]): uncertainty; dynamism; team work; stress; risk; an open environment that imposes variable and unpredictable demands; large amounts of data to process and high potential for sensory overload; imperfect data; human interaction mediated via computers; and complex multi-component decision making. In practice, of course, complexity will vary, depending on the specific context and nature of the ship's mission. Complexity can also be expected to increase with the growing emphasis on littoral warfare and as the revolution in military affairs exerts pressure for increasing adaptation and agility in responding to rapid change.

As in any open work system, command personnel and their staff have to deal with a large variety of situations or events both internal and external to the ship, from familiar ones that they encounter routinely, to unfamiliar, but anticipated ones (i.e., anticipated by system designers, policy, doctrine, tactics, and procedures), to both unfamiliar and unanticipated ones. A designer aiming to support the cognitive work of the Command Team must therefore consider alleviating the demands posed by each of these event types. [Rasmussen *et al*., 1994] suggest that a three-level performance hierarchy, referred to as the SRK taxonomy, can be used to describe variations in the levels of human behaviour possible in these situations. These levels are defined by skill-based (SBB), rule-based (RBB), and knowledge-based behaviours (KBB), and they relate to the conceptually distinct ways in which constraints in the environment are represented and processed by the operator. In increasing order of cognitive processing demand on the operator, SBB corresponds to using an effortless, internal dynamic world model, RBB to implicit cue-action mappings, and KBB to explicit representation of relational structures and processes in the work environment. Knowledge-based reasoning, in particular, is associated with the level of reasoning required to handle unanticipated events, diagnose problems, identify deep features about a situation, adapt plans and responses to the momentary needs of the situation, and so on.

Most of the popularly touted methods for modeling work demands fall under the labels of Task Analysis (TA) and Cognitive Task Analysis (CTA) and essentially adopt the granularity of tasks for their modeling effort. Their basic approach consists of parsing the behavioural flow in the current work environment into events, and their associated decisions and actions. Activity elements in this event flow emerge as: predetermined goals; devices, tools or techniques for achieving these goals; tasks which are the activities necessary to achieve goals using a device; task components; and actions or simple tasks with no control structure [Preece, 1994].

Limitations in the modeling power of Task Analysis and Cognitive Task Analysis to work analysis include:

- Frequently, TA/CTA techniques only describe what currently happens in the work domain without offering any real analytic capability to support design [Preece, 1994].

- The possibility with TA/CTA approaches exists, at least based on using only such techniques, of not abandoning existing designs [Preece, 1994], or breaking out of the assumptions underlying current designs of work processes, task structures, and work tools [Vicente, 1999]. This can lead to exploring design solutions related more directly to the 'what is' instead of the 'what could be'. The reason is that their potential for design guidance is based only on event-dependent descriptions of current work practice, in the form of tasks and actions, information cues, patterns and relationships workers perceive, the knowledge they use and strategies for processing information in performing these tasks.

There is therefore an inevitable preoccupation with representing existing trajectories in the behavioural flow of the work environment (to the extent that these trajectories can be sampled). It is not evident that this approach will lead directly to identifying intrinsic demands of the work that need to be supported (i.e., not workarounds that operators engage in because of deficiencies in currently provided tools).

- TA/CTA techniques lack flexibility to model work demands of operators that are a consequence of a great deal of behavioural variabilities [Vicente, 1995] and emergent behaviours in complex sociotechnical systems [O'Neill, 1996]. Examples of such behaviours include:

  - the need to adapt responses to unanticipated events arising from unforeseen disturbances that impact the state of the open work system (e.g., see [Vicente and Rasmussen, 1992]);

  - operators' discretionary need to resolve many degrees of freedom in particular situations, based on their own subjective performance criteria (e.g., see [Rasmussen, 1986, Rasmussen, 1997]); and

  - inadequacy of routine problem-solving methods for dealing with novel situations and the consequent need to reinterpret and reevaluate existing knowledge, usually involving almost ad hoc interactions among a variety of team members (e.g., see [O'Neill, 1996]).

  These demands can manifest themselves in several areas of the work, including the way it is shared among operators in high-load situations, the strategies employed in processing information depending on factors like expertise (e.g., use of RBB versus KBB based on experience), perceived work load and cognitive burden of using one approach over another, and the specific information that is needed for these strategies [Vicente, 1999]. O'Neill's reference in [O'Neill, 1996] to the evolving nature of artifacts that are produced and exchanged by a group of workers to represent their shared understanding of ill-structured problems that they confront is another interesting example.

  An important consideration in complex sociotechnical systems is that, due to their openess, dynamics, state uncertainty, and a variety of context dependent factors, the very same sets of actions can have different effects at different times, and different sets of actions may be required on different occasions to achieve the same goal. It will also not generally be possible to prescribe in advance a complete set of goals that need to be achieved and the potentially complex sets of interactions and dependencies between those goals. Work descriptions based on predetermined goals and specified tasks involving predetermined orderings of actions therefore offer limited expressiveness for representing these variabilities and hence for focusing design efforts to support the work demands they place on operators.

5.2 ***CWA Modeling***

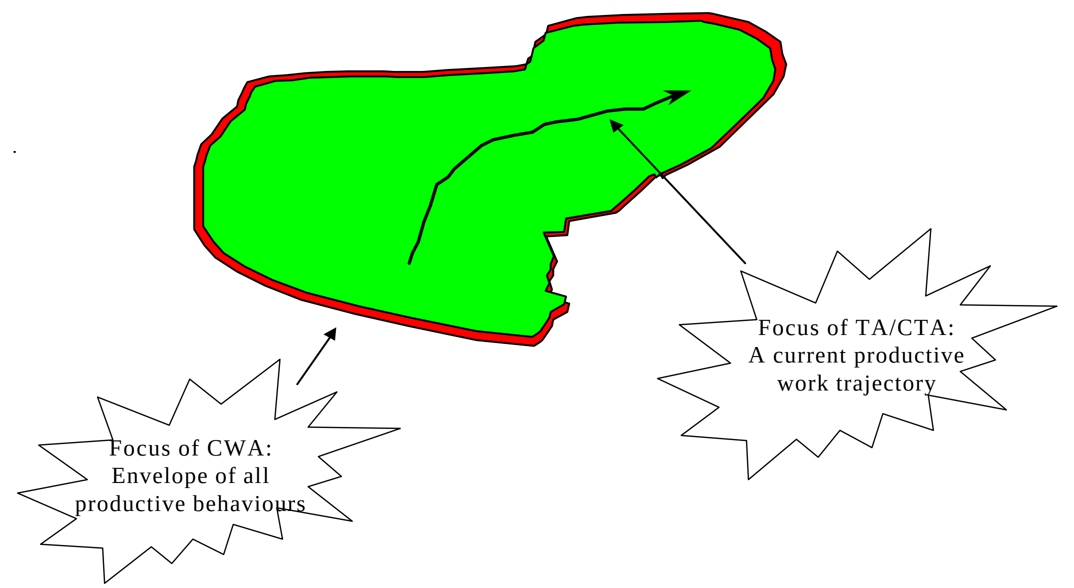

Figure 3 Different focuses of CWA and TA/CTA

Whereas TA/CTA focuses on current work trajectories in terms of stable task procedures, CWA broadens the scope of the analysis to the work place itself. CWA is based on the notion of intrinsic work constraints which are the behavioural constraints that need to be respected for effective performance in the work environment, independent of any particular device [Vicente, 1999]. Rasmussen refers to intrinsic work constraints as behaviour shaping constraints [Rasmussen *et al*., 1994] because they effectively delimit an envelope of behaviours within which all productive work can take place, independent of any operator, automated controller, or interface. The differing focuses of the two approaches are illustrated in Fig. 3. Not indicated in Fig. 3 is the possibility that TA/CTA methods, by their emphasis on parsing existing work practice, may also end up identifying trajectories related to workarounds that operators engage in owing to deficiencies in currently provided tools and which therefore lie outside the envelope of productive work trajectories that CWA focuses on.

Existing work devices such as currently available information and decision support systems, including their interfaces, are explicitly excluded from CWA modeling of the work domain to steer the analysis effort from the outset toward novel design solutions that are not constrained by current work practice. For this reason, it has been suggested that CWA is suited to a revolutionary approach to design, where radical leaps may be required simultaneously along several dimensions and a new design may fail if just one dimension of the change has not been adequately considered [Rasmussen *et al*., 1994]. However, Vicente [Vicente. 1999] is careful to point out that CWA can in fact be used in evolutionary design projects as well. In this case, some design decisions have effectively been made in advance and serve as inputs to the analysis. While this should lead as with TA/CTA methods to effecting design improvements in the existing system, the potential exists, however, for a less effective design result. This would be true of any revolutionary approach to design that is forced to follow an evolutionary design path, not just CWA.

In terms of the characteristics of the naval work environment described previously, the importance of CWA is that it provides the basis of an approach for eliminating the modeling limitations of TA/CTA methods identified in Section 5.1. We summarise the benefits of the CWA approach as follows:

- First, the models produced by CWA provide a conceptual means of describing those elements of the current work practice that lead to productive work (and which therefore need to be preserved, and even reinforced, in any future design). In this sense, CWA extends the modeling capabilities of TA/CTA methods.
- Second, CWA models aim to circumscribe currently unexplored possibilities for getting the job done. In this sense, therefore, CWA also provides a framework for delimiting new work trajectories and identifying design requirements for supporting the work demands of these trajectories.
- Third, since CWA models only identify work constraints instead of modeling predetermined work trajectories, this offers the modeler in situations where behavioural variabilities and emergent behaviours are an inevitable feature of the work environment the flexibility to represent those demands and subsequently generate design requirements for their support.

The CWA approach to modeling therefore has both descriptive and formative validity. Its descriptive capability permits understanding current behaviour in the C2 system. On the other hand, the formative aspect focuses on identifying requirements from a holistic view of the workspace that need to be satisfied for effective work. In a qualitative sense, a formative approach allows the designer to identify and anticipate the consequences of design interventions, without presupposing prescribed behavioural trajectories. [Vicente, 1999] contrasts this approach with other purely normative or descriptive approaches to work analysis. Normative approaches focus on shaping work as the designer perceives it *should be*, generally based on unrealistic or naive assumptions about the work. Descriptive approaches focus on the way things *are* and lean towards supporting current work practice. CWA focuses on the way things *could be* by developing work models that help identify novel possibilities for productive work.

For some specific examples of the anticipated benefits of CWA for improving HALIFAX Class ship operations, we note that CWA has the potential to contribute to the design of advanced interfaces and decision aids for the Command Team of the HALIFAX that:

- help them maximise ship effectiveness by understanding how their work domain behaves;
- support Command Team flexibility and adaptability in developing effective workarounds if ship components fail;
- help the Command Team develop novel solutions to unanticipated contact capabilities or behaviours;
- improve their situation awareness by exposing the breadth of the work domain; and
- support their quick adaptation to new technological capabilities added to the frigate by quickly associating these improvements with their potential to enhance mission accomplishment.

It is evident that the application of CWA hinges on identifying behaviour shaping constraints in the work environment (using appropriate knowledge acquisition methods, as Fig. 2 shows) that can account for the behavioural variabilities and emergent behaviours which are characteristic of complex sociotechnical systems. It will be difficult, if not impossible, to make the claim that all behaviour-shaping constraints have been modeled, or modeled with sufficient fidelity. The point is that modeling work demands has to be seen as a process of iterative refinement. An important consideration is that explicitly representing models of work constraints provides a structured way of recording what has been included and making design improvements in a consistent (i.e., not ad hoc) manner.

5.3 *Caveats*

As previously mentioned, CWA has been specifically developed for representing demands in complex sociotechnical systems. However, one should note that there are some characteristics that distinguish C2 from civilian sociotechnical work systems (e.g., nuclear power plants, process control) in which CWA has found its principal application to date. This includes having to deal with intelligent threats with the goal of denying information, or conveying false information. This one consideration alone means that there will inevitably be significant amounts of uncertainty that remain even with improvements in sensor technologies, data and information processing, and interface technologies.

In addition, areas of the shipboard work environment appear to differ in the extent of knowledge-based reasoning required (Rasmussen's KBB level). A simple example of this would be the contrast between air warfare and sub-surface warfare. Air warfare tends to be very procedural, whereas sub-surface warfare appears to require more deliberative reasoning. TA/CTA modeling may therefore suffice for the former, with little marginal gain from using CWA, whereas benefits could be substantial for the latter. This raises the possibility of segmenting the work environment into pockets of activity where a CWA approach to analysis is likely to produce high immediate returns versus other areas where the benefits of CWA may be less evident and the returns based on optimising the existing system may not justify the effort of a CWA approach. However, this remains for now only a possibility that merits future consideration.

Additional experimental work with CWA in this work environment is required to better understand the full extent and implications of the caveats we have outlined here.

5.4 *CWA Framework*

We have described the basic approach to modeling work demands adopted by CWA. In fact, CWA identifies, five layers of behaviour-shaping constraints in a sociotechnical work system, each layer related to analysing a different dimension of the set of work demands. They are illustrated in Fig. 4. The top three layers of analysis relate to what [Rasmussen *et al.*, 1994] refer to as the 'identification of activities' in the work system, and the bottom two to the 'identification of agents or actors' that perform the work. This separates characteristics of the work environment that must be satisfied from those of actors responsible for their satisfaction.

In the *work domain* layer, workspace constraints are identified, in a device-, event-, goal-, task- and actor-independent manner. These constraints define the work domain's purpose and content, and the structural relations that link purposes with the means the domain provides for achieving its purposes. Collectively, these constraints define its field of action possibilities.

Reducing degrees of freedom	CWA Layers	Typical Modeling Tools
↓	Work Domain Analysis	ADS
	Control Tasks Analysis	DL
	Strategies Analysis	IFM
	Social Organization and Cooperation Analysis	ADS, DL, IFM
	Competencies Analysis	SRK

Figure 4. Constraint layers in CWA

The *control tasks* layer models in a device- and actor-independent manner the goals to be accomplished, and their coordination constraints (e.g., temporal and logical constraints on workflows), using a product or input-output representation of the processing required to achieve goals. Prototypical work activities are identified, activities are represented structurally in decision-making terms, and processing short cuts in these decision-making activities are identified as a means of capturing potential processing variabilities among actors or as a result of the context of the situation. This modeling effectively provides constraint-based representations of domain tasks. The formative nature of these representations differs from one that prescribes the one ‘right-way’ to perform the task.

The *strategies* layer represents information processing in the system. It models the categories of generative mechanisms or processing procedures that can be employed by domain actors (humans or machines), as well as the constraints that govern strategy adoption and switching among these categories, depending on context, subjective task formulation and performance criteria.

The *social organisation and cooperation* layer models organisational constraints and content and form constraints (workload sharing, social organisation, etc.) underlying human-human and human-machine cooperation and communication.

Finally, the *competencies* layer models the constraints associated with the operators themselves, including their generic capabilities and performance limitations and more specialised competence requirements of their information-processing activities.

It is worth noting that the various constraint levels have far-reaching design implications for the workspace. While several directly impact computer-based DSS design, some are also related to other design interventions (e.g., training, selection, sensors; see [Vicente, 1999], [Reising and Sanderson, 1996]).

Table I. Constraint information of each layer of CWA

CWA Layers	Kinds of Constraint Information
Work Domain Analysis	Purpose and structure of work domain
Control Tasks Analysis	Goals to be satisfied; information, decisions, and cognitive processing required; coordination of control tasks
Strategies Analysis	Ways that control tasks can be executed
Social Organisation and Cooperation Analysis	Work distribution and sharing
Competencies Analysis	Levels of mental processing and information interpretation; required knowledge; performance limitations; novice-expert differences

As Fig. 4 indicates, the degrees of freedom in the design space (i.e., the space of feasible design interventions) are reduced as each new layer of constraint is added. The ordering of layers, from work domain to competencies, is related to a fundamental *ecological* orientation, as opposed to a *cognitive* one, underlying CWA. *The ecological approach* argues that in work environments that impose dynamic, external constraints on the goal-directed behaviours of its workers, it is necessary that environmental constraints (*ecological compatibility*) be considered *before* cognitive constraints (*cognitive compatibility*) [Vicente, 1999]. The point is that effective work performance in such domains rests on first satisfying the immutable constraints imposed by the work ecology. This orientation also has consequences for knowledge acquisition since the complexity of these work domains makes it difficult, if not impossible, for operators to have complete and correct models of their work domains. Starting the analysis by focusing on human mental models, using a CTA approach say, is therefore considered inappropriate.

The goal of achieving ecological validity has its roots in Brunswick's work [Brunswick, 1952] in ecological psychology. To understand the difference in focus between the ecological and cognitive approaches, we note that the ecological approach effectively places emphasis at design time on supporting operators in acquiring correct mental models of their domain, based on developing a structural description of its action possibilities for goal achievement, rather than on eliciting their potentially flawed mental models to support design. Following the ecological approach, the interface provides the means to expose work constraints by means of cognitively efficient representations (visualisations) on which the operator then 'overlays' his/her mental models in understanding and responding to the situation.

CWA can be thought of as an *ecological* or *use-centred* approach to system design in a similar sense to that indicated in [Flach *et al.*, 1998] for interface design; it stands in contrast to other approaches classified as *technologically-centred*, *user-centred*, or *control-centred*. As [Flach *et al.*, 1998] indicates, these various approaches differ in terms of which constraints are given precedence, whether they be machine constraints, human performance or psychological constraints, stability constraints associated with human-machine integration, or workspace constraints in the context of the larger view of a work ecology. As is evident from Fig. 4, each of these four types of constraints would nonetheless be represented in a CWA.

Rasmussen has also pioneered a set of generic, conceptual modeling tools for representing constraints in each layer of the work analysis [Rasmussen *et al*, 1994], [Vicente, 1999]. Some of these are indicated in Fig. 4, along with their mapping to the various layers in which they have been found useful. ADS represents the abstraction-decomposition space, DL is the decision ladder, IFM are information flow maps, and SRK is the skills, rules, knowledge taxonomy that we have previously referred to in Section 5.1. The significant distinguishing characteristic of each of these tools is their ability to *formatively* represent behavioural, process, and task variabilities as needed by their respective constraint layers. Examples of some of the modeling tools used in the CWA feasibility study appear in Sections 6 and 7.

Table I summarises the typical constraint information provided by each layer of CWA.

6. **CWA Feasibility Study**

6.1 ***Overview***

The overarching objective of the feasibility study was to determine if the CWA framework can be effectively applied to the C2 work environment of a HALIFAX Class ship to explore concepts for computer-based support of the Command Team's tactical Command and Control activities.

Given the modest scope of the current effort, effectively over less than a one year time frame, with limited manpower and time resources, it was decided that a first assessment should look at:

- putting some of the upper layers of the framework into action by developing a preliminary set of CWA representations;
- observing pros and cons in an empirical manner, focusing on using CWA models to capture preliminary information and decision support requirements; and, finally,
- doing a conceptual design derived from these requirements, to be represented in the form of a dynamic storyboard, of an interface for this work environment.

Table II. Parts of the CWA feasibility study

Part I	Work Domain Analysis of HALIFAX Class Frigate
Part II	Control Task and Strategies Analysis of the ORO's Work Space; Development of Visualisation Concepts for a Dynamic Storyboard

The study was conducted in two parts as indicated in Table II. The shading in Table II corresponds to that in Table I and indicates the general kinds of modeling information that was captured in the two parts. However, there was some overlap in the modeling effort in passing from the broadly-based work domain analysis of the HALIFAX Class frigate conducted in Part I to the analysis in Part II specifically related to the ORO's tactical situation and threat assessment (STA) and response management (RM) activities. In particular, Part II built on the results of Part I as a means of focusing the efforts of that part of the study.

Part I

- Establish the feasibility of using a Work Domain Analysis (WDA) of the HALIFAX Class frigate for:
  - descriptive and formative modeling of the tactical C2 work domain; and
  - determining computer-based design interventions to support the decision-making activities of the ship's Command Team.
- Use the results of the WDA to make a preliminary determination of information display requirements for new interfaces for the HALIFAX Class Operations Room.

Part II

Demonstrate the utility of the CWA framework in generating novel computer-based decision support for operator demands on the HALIFAX Class ship by:

- conducting a control tasks and strategies analysis of the workspace of the ship's tactical coordinator, the Operations Room Officer (ORO);
- assessing current support tools for the ORO based on the CWA modeling results;
- using the CWA models to define preliminary visualisation and processing requirements for an advanced ORO interface.

This section gives an overview of the method that was followed in the two parts, along with the general nature of the models that were obtained in the first part. Preliminary results from the Part I study have previously appeared in [Chalmers and Burns, 1999]. More in-depth discussions of its analysis and findings are presented in [Burns *et al*., 2000], [Burns and Bryant, 2000]. Section 7 looks in more detail at the second part of the study.

6.2 *Method*

An iterative approach has been followed in both parts of the study. Preliminary knowledge acquisition sessions used a variety of operational level and training documents dealing with the principles and procedures of maritime warfare. Various quasi-structured interviews of SMEs in doctrine, training, and engineering were conducted, and a number of complete training exercises observed in a land-based training simulator, as well as on board a HALIFAX ship engaged in simulated air/missile attacks as part of harbour exercises. In the Part I study, this was augmented by a field trip on a HALIFAX ship involved in (primarily) simulated air defence exercises in a task group setting over the course of a three day sea trip between Norfolk, Virginia, and Halifax, Nova Scotia.

CWA models constructed at each stage based on these knowledge acquisition sessions were evaluated in further interviews with SMEs to correct misconceptions, fill in gaps in the models, and generally test their validity. In the Part I study, this also involved walking SMEs through an existing training scenario to make contact with the elements of the model and identify model deficiencies, and, in a final phase, having them complete detailed questionnaires testing out various aspects of the final versions of the models, including model completeness and frequency of consideration of model elements across a variety of mission types.

SME inputs were also used in the Part II study to develop the details of a tactical scenario that included a variety of air, surface, and sub-surface threats, as a means of supporting the development of the dynamic storyboard for a portion of the ORO's workspace. This scenario and the validated CWA models were used to develop visualisation concepts for the ORO's interface which were then captured in the storyboard. The storyboard was subsequently evaluated by walking SMEs through its operation and by obtaining their feedback both verbally and in the form of responses to a questionnaire.

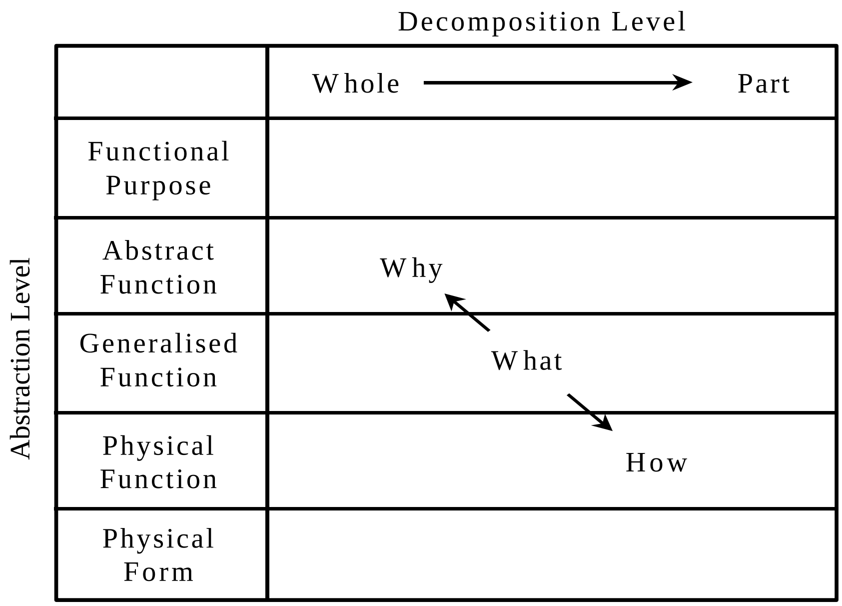

Figure 5. Abstraction-decomposition space

6.3 ***Part I Models***

As previously reported in [Chalmers and Burns, 1999], the WDA for the HALIFAX ship involved developing an abstraction hierarchy for tactical C2 (see also Section 5.4). A general abstraction-decomposition space (ADS) models functional relationships between work domain elements and their purposes, as well as part-whole relationships. Figure 5 shows the five levels of abstraction that have been used: functional purpose (frigate purposes), abstract function (principles or conservation laws), generalised function (processes), physical function (capabilities), and physical form. It shows the 'why', 'what', and 'how' structural means-ends links that exist between consecutive levels of an ADS: at the level of 'what', the operator determines the 'why' of a model, its goal, by looking at the level above; and determines the means, the 'how', to achieve the goal from the model level below.

WDA models were elaborated as two separate domains, one for ownship and the other for contacts with which ownship interacts. These two domains interact through the natural environment which also impacts the capabilities and opportunities for action in each domain. This leads to the general 3-part WDA model shown in Fig. 6, where only the frigate portion of the model is expanded for simplicity. Modeling the environment separately provides a non-redundant way of exposing the relational dependencies with the ownship and contact domains. The various model entities, along with their means-ends links for the frigate model are shown enlarged in Fig. 7. The differences in shading of model elements highlight three classes of model constraints that were identified: physical constraints, information-gathering constraints, and social-organisational constraints [Burns *et al.*, 2000]. Physical constraints represent the structural boundaries due to physical laws. Information-gathering constraints are concerned with maximising the amount of information available to the ship. Social-organisational constraints define a boundary between what it is physically possible for the HALIFAX Class frigate to do and what is incompatible within the task group and naval environment within which the frigate operates.

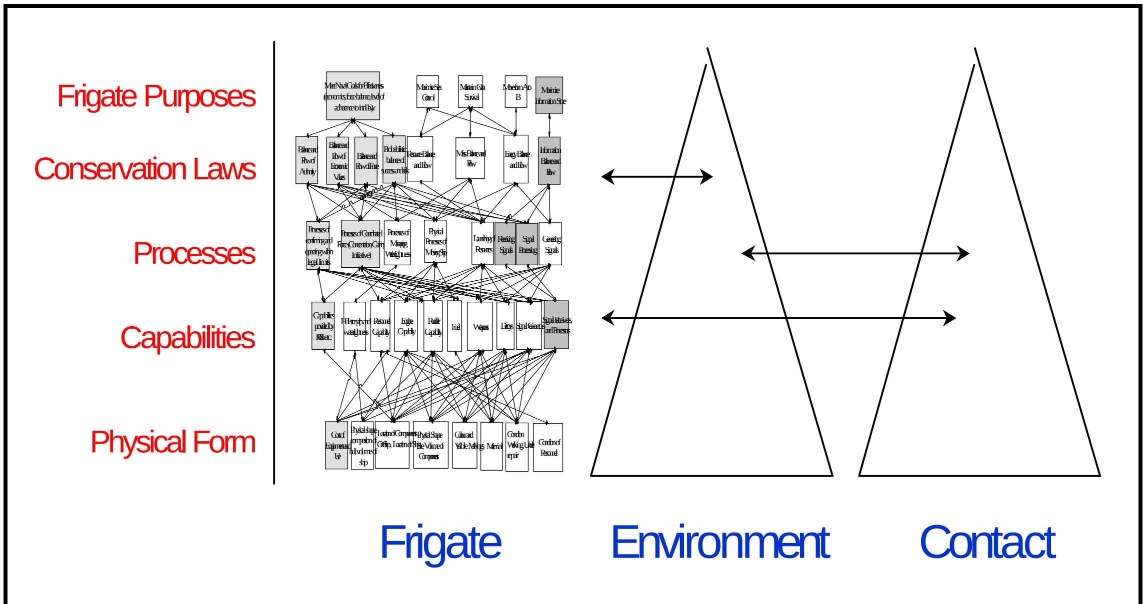

Figure 6. General 3-part WDA model

7. CWA Feasibility Study: Part II

This part of the study used as its starting point results from the general 3-part WDA model developed in Part I (Fig. 6). The focus was on modeling decision and information requirements, derived by CWA, to permit creating decision-centred visualisations to support the ORO in his/her tactical situation and threat assessment (STA) and response management (RM) activities. We concentrate on describing here the modeling process followed and highlighting the key steps along the way from CWA modeling to the development of the visualisation concepts. Additional discussions of the results appear in [Easter and Potter, 2000]. [Potter *et al.*, 2000] presents another case study where a similar approach was followed in a strategic decision-making environment.

The modeling effort crossed over several of the nodes in the abstraction hierarchy in Fig. 6. Within the domain of HALIFAX Class frigate resources, it addressed primarily the ‘Launching of Resources’ node at the generalised function level, as well as the supporting nodes of ‘Weapons’ and ‘Signal Generators’ at the physical function level. Within the domain of contact resources, it addressed the ‘Generating Signals’ node at the generalised function level, as well as, similarly, the supporting nodes of ‘Weapon’ and ‘Signal Generators’ at the physical function level.

7.1 *Decision Requirements Assessment*

To help narrow the focus of the visualisation design, a decision requirements assessment of the current support tools for the ORO’s role was first conducted based on interviews with SMEs and on observations of a number of training exercises in a land-based training simulator and on board a HALIFAX ship. Three critical roles of the ORO emerged from this knowledge acquisition exercise.

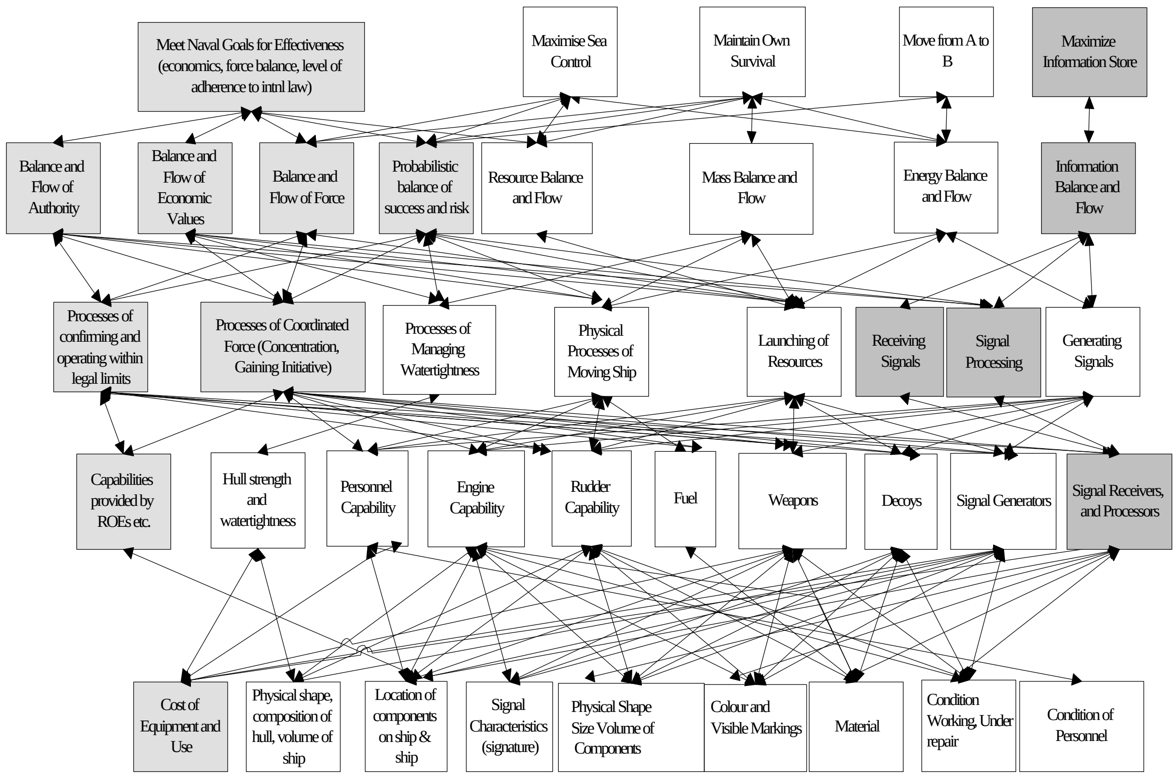

Figure 7. Domain of HALIFAX Class frigate resources [Burns *et al.*, 2000]

1. Integration of warfare areas:

   - The ORO serves as the 'integration point' of three warfare areas, air, surface, and sub-surface, as well as of both internal and external events relative to ownship, and must keep one foot in each world to coordinate actions and events.

2. Keep the 'big picture':

   - The ORO must keep an overview of the entire situation and avoid becoming locked onto a particular problem at the expense of others. The ORO may need to alert or request the assistance of others on significant events outside of their current field of view.

3. Maintain an understanding of 'functional impact' of events:

   - The ORO must synthesise information about the different events and activities in the world into a coherent representation of the state of affairs. In other words, the ORO must maintain a high-level situation awareness.

Some of the conclusions from the tools assessment are:

- The Operations Room displays are, fundamentally, copies of drawings that, not surprisingly, are artifacts of the way the problem space was dealt with by practitioners prior to the development/availability of computer technology or as a consequence of the technical limitations of the era in which they were designed. This implies that the computer is being used primarily to speed up the work process (mostly in the tight detect-to-engage loop), without attempting to simplify or improve the human decision-making process, particularly in conditions of high stress, uncertainty and time pressure.

- There is limited capability in the current CCS for synthesising incoming data to provide the operator with new levels of understanding. This problem is compounded by the large number of sources of data, with differing quality and timeliness characteristics, that an ORO needs to integrate and make sense of in order to perform his/her tasks. Often these sources are independent of each other, i.e., they have separate interfaces with different protocols for accessing the necessary data. Examples are the data available on the current display, text and verbal messages, data from other information systems, and status boards.
- The ORO's job as the 'threat manager' is reactive with often very little time for completion.
- The alarm system presents all abnormalities/problems as if they were of equal importance.
- There is little or no computerised support for the ORO in supervising the Operations Room personnel or in tracking the general status of the Operations Room (short of communicating directly with other key personnel) and getting a summary of the implications when one or more elements are operating in degraded mode.
- Existing displays tend to show only the present situation, i.e., what is going on now. For example, there is a need, particularly in the sub-surface warfare area, to indicate where the submarine cannot be. In addition, the displays do not provide insights into how the computer system arrived at the conclusions presented.

A number of specific problems that fall in the general area of threat management were also observed. They are consistent with the general lack of data synthesis and the consequent negative effects in critical, high-tempo situations (e.g., missing the 'information' for the 'noise' which is characteristic of data overload). The threat management area of the ORO's tasks was therefore identified as a promising first area for using CWA modeling to derive improvements over existing CCS interfaces.

7.2 *Functional Abstraction Hierarchy Modeling*

As a means of integrating the previous WDA models and with the focus on modeling decisions and their information requirements, a hybrid CWA model, referred to as a Functional Abstraction Hierarchy (FAH), was built. We describe first the general framework of a FAH and then give an overview of the specific model that was developed for the ORO's workspace.

7.2.1 *Model Framework*

A Functional Abstraction Hierarchy is a formative structure for mapping cognitive demands created by the work domain in a fairly similar manner to that of Rasmussen's abstraction-decomposition space. It provides a means of representing goal-means or functional interrelationships between work domain entities. It represents both functions, which provide the means for achieving domain goals (functional commodities), and the set of processes that make up a function and provides the commodity and knowledge about how they work to achieve the function's purpose [Woods and Hollnagel, 1987]. It can be thought of as a hybrid modeling tool because it incorporates characteristics of multilevel flow modeling [Lind, 1991] and abstraction-decomposition [Rasmussen *et al*. 1994].

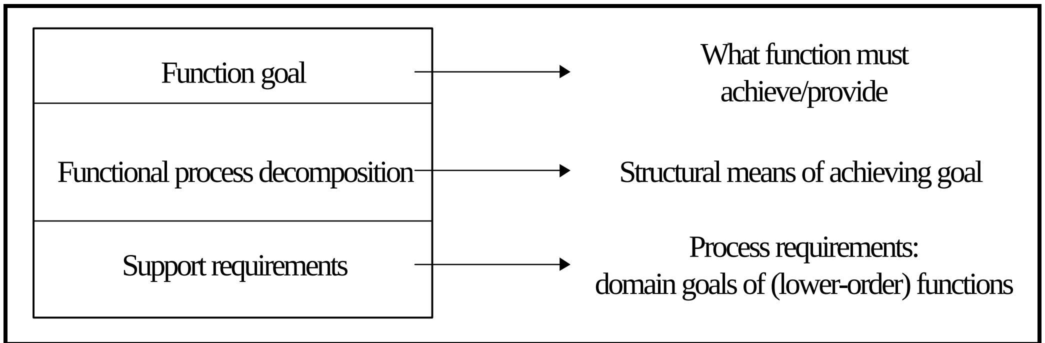

Figure 8. A function node in a goal-means network

In more detail, a FAH represents work domain objectives and the functions that must be available and satisfied in order to achieve goals. In turn, these functions may be abstract entities that need to have other, less abstract or less aggregated functions available and satisfied in order that they might be achieved. The representational framework is that of a network of functions, also referred to as a goal-means network [Woods and Hollnagel, 1987], linking abstract aggregates of work domain entities to their concrete details. Nodes in the FAH network represent domain functions that can be linked with other domain functions by support requirements relationships. This arises when a component process has a requirement that serves as a domain subgoal of some other, usually lower-order, function. These links can also be labelled by performance criteria governing goal achievement. This is summarised in Fig. 8. We have also used a second type of link, a decomposition link, that provides a means of mapping or expanding a process function into its constituent subgoals. Various conventions may also be used to represent function combination or aggregation, alternative functions or processes, pre-conditions, side effects, preferred orders, etc.

The FAH provides a framework for making explicit the goals to be achieved in the domain, the choice points and decisions that arise, the information needed to support those decisions, and the cognitive and collaborative activities entailed [Woods and Hollnagel, 1987], [Roth and Mumaw, 1995]. The FAH enables the designer to determine where decision making is likely to be difficult due to the nature of the problem space (e.g., due to lack of information, multiple interacting goals, or time pressure). For example, the FAH helps identify places in the problem space where objectives compete with each other (e.g., where choices have to be made that require some level of sacrificing of one objective in order to achieve another, perhaps more heavily weighted, objective), or otherwise constrain each other (e.g., where the satisfaction of multiple goals need to be considered in determining the best course of action). Further, the FAH helps to reveal the context of a decision, what the necessary inputs are, and the complicating factors that can arise (e.g., sensor failures, equipment malfunctions) to increase decision difficulty.

7.2.2 *FAH Model*

We used the FAH to represent constraints in the top two CWA levels (see Fig. 4) of the ORO's workspace. We focused on identifying the hierarchy of goal/process nodes related to the STA/RM work area. A high-level (partially expanded) version of the FAH is presented in Fig. 9. The target area of threat management selected for the subsequent visualisation effort (see Section

7.1) is shown in red. As implied there, the visualisation can also be expected to address some of the cognitive demands of the decomposition nodes in the FAH related to the threat management function.

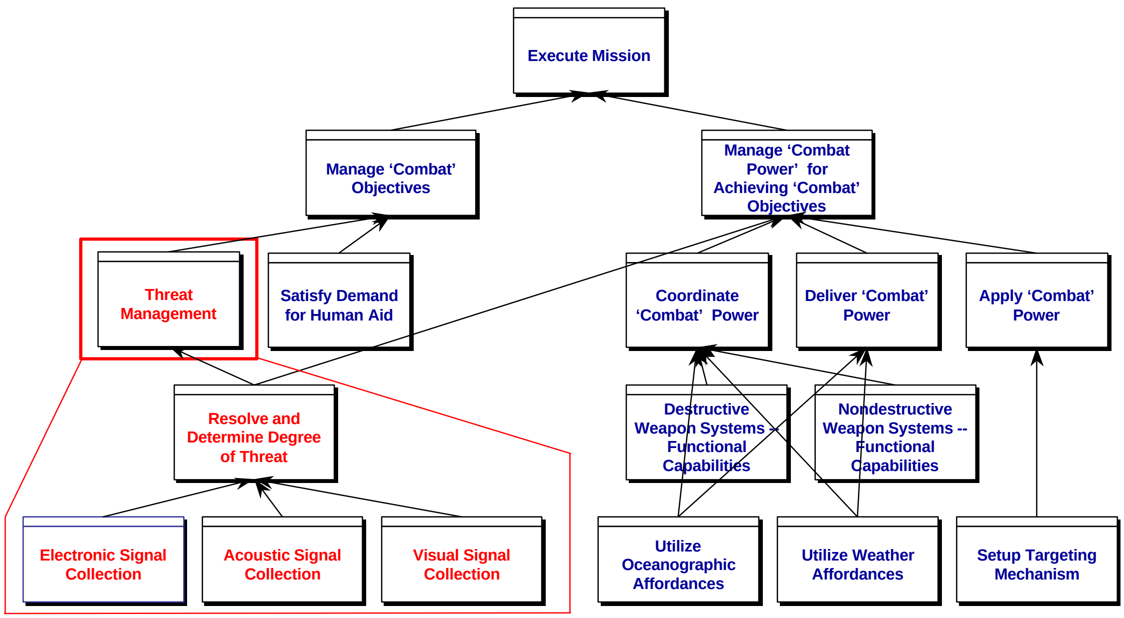

Figure 9. High-level version of the Functional Abstraction Hierarchy

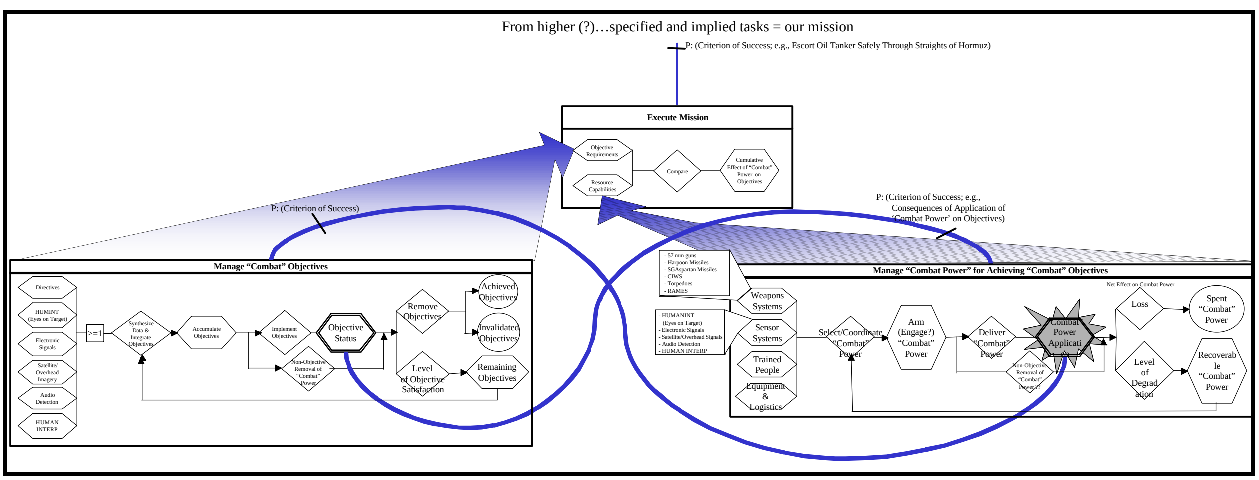

Figure 10. Top three nodes of the Functional Abstraction Hierarchy

At the top of this FAH is the primary goal, which is to “Execute Mission”. This domain goal is decomposable into two subgoals, namely “Manage ‘Combat’ Objectives” and “Manage ‘Combat Power’ for Achieving Combat Objectives”. The first goal, related to managing combat objectives, allows the ORO to monitor the situation as his/her decision(s) is/are implemented and to adjust moment-to-moment objectives as the situation dictates. The second element of the decomposition has to do with assembling and applying available resources to accomplish the mission, based on the ORO’s decision-making abilities to drive the process. These three goals, executing a mission by extracting one or more objectives from it and assembling and applying available resources in an attempt to accomplish the mission, serve as the core of the FAH. They form a continuously on-going process to which all other goals present themselves as interrupts as the situational context demands. These three nodes are illustrated in more detail in Fig. 10.

The FAH helped to expose key functional commodities of the decision-making process that the ORO engages in during STA/RM reasoning. These can be summarised as follows, with typical examples shown bulleted.

Functional Role:

- What is the ship's role within the Task Group?
- What are the mission objectives?
- What are the Rules of Engagement (ROEs)?

Functional Capabilities:

- What is the range and quality of the ship's sensor system?
- What is the range and kill probability of the weapon systems?
- What are the sensor and weapon ranges of the threats?

Functional Predictions (impact of environment on functional capabilities):

- Where are the known threats?
- Where can the enemy go? Where can't they go?
- Which threats can the ship respond to? When can it respond to them?
- What is the impact of current conditions (onboard and environmental) on resource capabilities (for ownship as well as the threat)?

7.3 ***Decision Modeling***

7.3.1 ***Underlying Framework***

As indicated earlier, the goal-means network may be used to derive the critical decisions and associated cognitive and collaborative activities required to achieve domain goals. These cognitive activities and decisions centre around goal-directed behavior, such as monitoring for goal satisfaction and resource availability, planning and selection among alternative means to achieve goals, and controlling activities (initiating, tuning, and terminating) to achieve goals [Roth & Mumaw, 1995]. By organising the specification of operator information, decision, and control requirements around nodes in the goal-means structure, rather than around predefined task sequences (as in traditional task analysis), the representation helps insure that resulting displays and decision aids reflect a decision-centred perspective.

The critical decisions and associated cognitive and collaborative demands derived from the FAH constitute a second type of design artifact – *decision requirements*. Decisions can be tied directly to nodes in the FAH and provide an intermediate artifact that forms part of the design thread. The result is an end-to-end connection from goal nodes in the FAH to supporting visualisations and decision aids. The critical decisions identified in Section 7.3.3 for the target region of the

FAH were derived from the knowledge elicitation activities (interviews, observations) as well as analytically by applying a number of questions, adapted from [Roth and Mumaw, 1995], at each pertinent node of the FAH.

Monitoring/Situation Awareness

Goal Monitoring:

- What are the goals? (e.g., win the battle vs. sustain minimal damage?)
- Goal satisfaction: Are function-related goals satisfied under current conditions?
- Margin to dissatisfaction: Are goal limits/restrictions being approached?
- Consequences of goal dissatisfaction: Which higher-level goals are supported by this function, and what are the consequences if this function is not achieved?

Process Monitoring:

- Active processes: Which resources are currently active? What is the relative contribution of each active resource to goal achievement? Are the resources performing correctly?
- Process element monitoring: Are the individual resources and their components working as they are supposed to?
- Automation monitoring: Are the automated support systems functioning properly? Which goals are these systems attempting to achieve? Are these appropriate goals?

Feedback Monitoring:

- Procedure adequacy: Is the current procedure achieving the desired goals?
- Control action feedback: Are actions achieving their desired goals?

Planning

- Operational and tactical priorities: Which goal has the highest priority?
- Availability of warfighting resources: What alternative resources are available for achieving mission goals?
- Choices among alternatives: Can an alternative resource be deployed?
- Consequences and side-effects of actions: What other warfighting resources and functions are affected by the current actions?

Control

- Resource control: How is a resource controlled for deployment, tuned for optimum performance, terminated?

7.3.2 *Decision Ladders*

The decision ladder template ([Rasmussen *et al*., 1994], [Vicente, 1999]) is a generic representation of the steps that may be involved in decision making. It is a 'template' because it can be used to show how a particular step is instantiated in a specific case, and how, for that particular decision, some steps may not be needed. It is a formative modeling tool that represents:

- information processing activities;
- resulting states of knowledge;
- multiple entry and exit points; and
- processing shortcuts via ‘leaps’ and ‘shunts’.

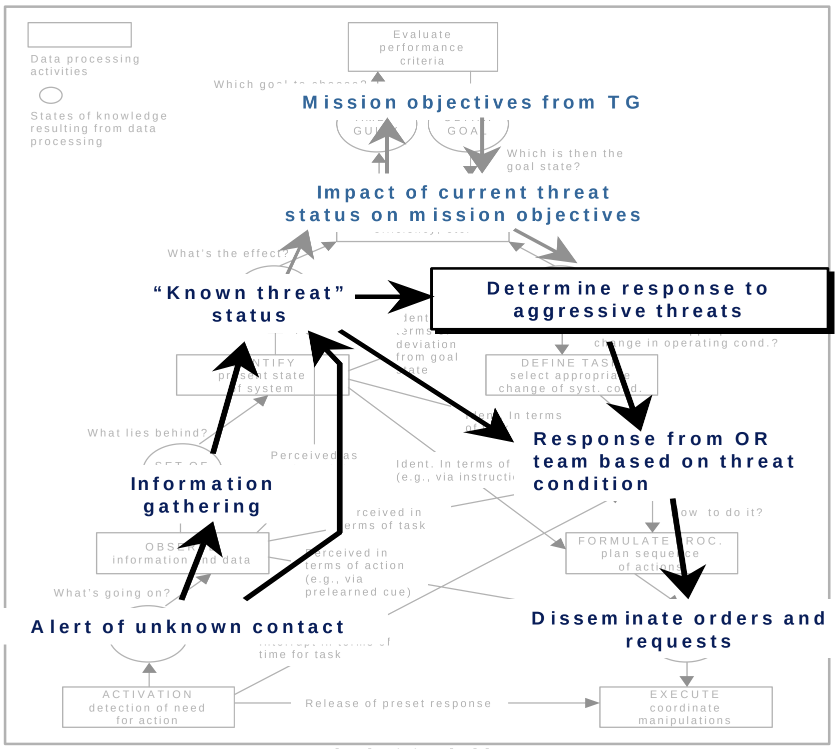

The decision ladder

Figure 11. Decision ladder for threat management

We illustrate the use of the decision ladder for the target region of the FAH (see Fig. 9). The decision ladder is one in which a current contact must be evaluated to determine its degree of threat, threats prioritised within the context of overall mission objectives, and a response determined based on the prevailing threat condition. Initial decisions relate therefore to resolving the identity of a contact and determining the degree of threat posed by potentially aggressive contacts. Emphasis then shifts to response determination and dissemination of orders.

This decision model implies two potential shortcuts in the decision-making process. First, there may not be a need for additional information gathering after the initial alert to a threat condition is received (if the initial alert is sufficiently well-defined as in a ZIPPO call[1]). Second, there may not be a need for information processing related to determining the appropriate response to the threat (if the threat is sufficiently lethal).

[1] Planned antiship missile defence reactions are effected as ZIPPO calls. They allow procedures for rapid and coordinated response to missile attacks to be executed.

7.3.3 *Critical Decisions*

From this decision ladder framework, four critical decisions in the target region of the FAH were identified (decision number in parentheses):

Process Monitoring:

- Monitor all contacts for ones that are potentially hostile (D1).

Goal Monitoring:

- Determine the degree of threat (based on identity determination and proximity) for all contacts (D2).
- Determine available response time for all aggressive threats (D3).

Control:

- Determine appropriate response, including priority and order of response, to aggressive threats (D4).

Feedback monitoring (e.g., kill assessment) is not explicitly part of this set of critical decisions, as this decision has been included in the "Manage 'Combat Power' for Achieving Combat Objectives" node in the FAH (Fig. 9).

7.3.4 *Information Requirements*

The next step is to identify and document the information requirements to support the cognitive tasks and critical decisions. The FAH representation can help in this step because it indicates the factors that need to be considered in making decisions. A designer can see what input data is necessary to provide the user with insights to be able to quickly and accurately make decisions. In cases where the required data is not directly available it provides a rationale for obtaining that data (e.g., by providing additional sensors) or creating 'synthetic' values.

To support the critical decisions listed above, the following information requirements were identified from the FAH (associated decision in parentheses) as a basis for the visualisation design in the target region:

- availability and functionality of contact sensing equipment (D1);
- contact characteristics (e.g., type, location, capability) (D2);
- contact behavior (e.g., current and projected course, response to warnings) (D2, D4);
- for incoming aggressive weapons, time until impact (D3);
- for incoming aggressive platforms, time until likely weapon release (e.g., closest point of approach (CPA) or based on weapon ranges) (D2, D3); and
- contact characteristics vs. degree of threat vs. characteristics of onboard weapons, including both static and dynamic weapon characteristics (e.g., current status of a weapon system vs. its prescribed capability) (D2,D4).

7.4 *Visualisation Concepts*

These various modeling efforts culminated with the development of visualisation concepts for a storyboard for the ORO's threat management interface. The focus at this point is on developing the mapping between domain semantics, that is information on the state and behavior of the domain (e.g., critical decision and information requirements uncovered), and the syntax and dynamics of the visualisation. The general design goal is to reveal the critical information requirements and constraints of the decision task through the operator interface in such a way as to capitalise on the characteristics of human perception and cognition. The approach is consistent with Cognitive Systems Engineering principles in representational aiding [Woods *et al.*, 1988] and Ecological Interface Design [Vicente and Rasmussen, 1992].

The display concept and how it supports the cognitive tasks and decisions is first captured in a Display Task Description. This description explicitly identifies the goals of the display in terms of the cognitive tasks and critical decisions it is intended to support. It also specifies supporting information and graphic elements required to support the cognitive tasks and decisions. A display task description is an intermediate design artifact that links the decisions and related cognitive and collaborative tasks revealed by the domain analysis (as represented in the FAH) with the visualisation and decision support concepts intended to support those decisions.

7.4.1 *High-Level Goals*

Consistent with the critical roles of the ORO identified in Section 7.1, the following high-level goals for the development of the visualisation concepts were identified:

- Maintain the 'big picture':
  - Provide the ORO with an overview of the tactical situation to avoid cognitive tunnel vision on a particular aspect of the problem. This may include alerts with respect to significant events outside of the ORO's current field of view. This was based on the observation (noted in the field work) that OROs were prone to focus attention on specific threats, potentially failing to manage responses to the entire situation.
- Provide coordinated functional and physical representations of the state of the work domain:
  - Provide the ORO with a combination of functional and physical (geographical) information that will, in combination, provide a coherent representation of the tactical situation. One key is to allow the ORO to effortlessly transition between the physical and the functional information. Virtually all of the ORO's existing displays are, from a CWA perspective, physical displays; that is, they represent physical aspects of the problem space at the level of physical entities or objects (e.g., their physical location and direction of movement). However, the information requirements identified in this work tended to be of a functional nature (e.g., time until likely weapons release, projected course information). Since the intention was not to simply add additional information to an already crowded tactical display, this led us to consider two integrated displays, consisting of a functional one and a physical one, rather than just one display, on which to map these various representations.
- Integrate the three worlds of air, surface and sub-surface warfare:

– Provide the ORO with a mechanism to integrate the three warfare areas with respect to threats and resources. This will support the ORO as the 'integration point' of these three worlds, consistent with the ORO's roles identified in Section 7.1.

7.4.2 *Display Task Description*

We concentrate here only on the functional component of the display representation. The focus of the current version of this display is a functional visualisation, organised around functional distance (i.e., distance represented in terms of time as a means of capturing notions of risk and imminence of tactically significant events) and the state of those contacts that have been identified as potentially aggressive threats. Although the following items were not fully represented in the final visualisation (due to a lack of project resources), the key elements of this functional organisation can be described as follows.

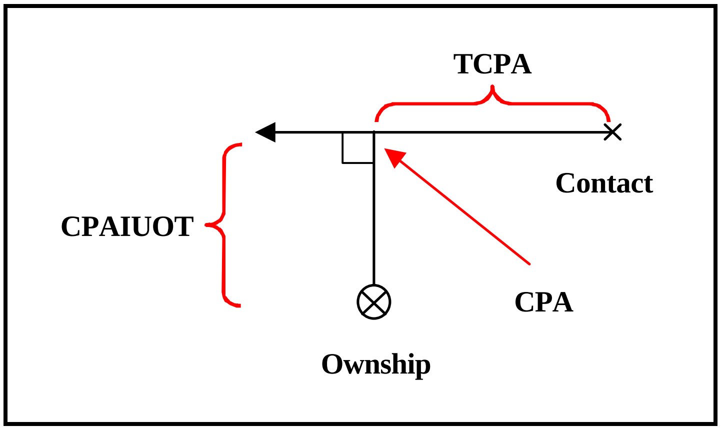

Figure 12. Computing TCPA and CPAIOUT

- Functional distance to a contact defined as a function of (see Fig. 12):
  - time to its CPA (TCPA); and
  - CPA distance transformed into units of time (CPAIUOT) (i.e., contact's CPA distance divided by its speed).
- Contact state defined as a function of:
  - contact classification (i.e., unresolved/suspect/hostile);
  - threat identification (i.e., missile/aircraft/torpedo);
  - threat ranking from the automated TEWA system or based on future threat evaluation algorithms;
  - planned response (i.e., weapon of choice)
  - engagement status (i.e., weapons assigned/engaged);
  - availability of response (i.e., time until threat is in arc of fire of weapon of choice).

Based on these parameters, the critical contacts to focus on include:

- unresolved contact functionally close to ship;

- hostile/unengaged contact close to ship; and
- hostile contact outside ship’s arc of fire.

The specific cognitive demands on the ORO we aimed to support with the functionally-organised visualisation include:

- determining the time it will take a threatening contact to reach its CPA with respect to ownship;
- prioritising threatening contacts based on their temporal ‘distance’ to the ship, ‘distance’ to achieving their weapon launch range (in the case of aircraft and submarine threats), and their lethality;
- ensuring that threat detection, evaluation and response activities are being executed effectively and in a timely manner, including the direction and operation of the ship’s sensors, weapon and decoy systems (i.e., that the Operations Room threat detection and response mechanisms are acting appropriately); and
- supporting assessments across all the warfare areas.

A display representation, based on the time-based transformations involving the closest point of approach, was selected as the basis for building a first useful visualisation to support the ORO’s threat management activities.

7.4.3 *Display Description*

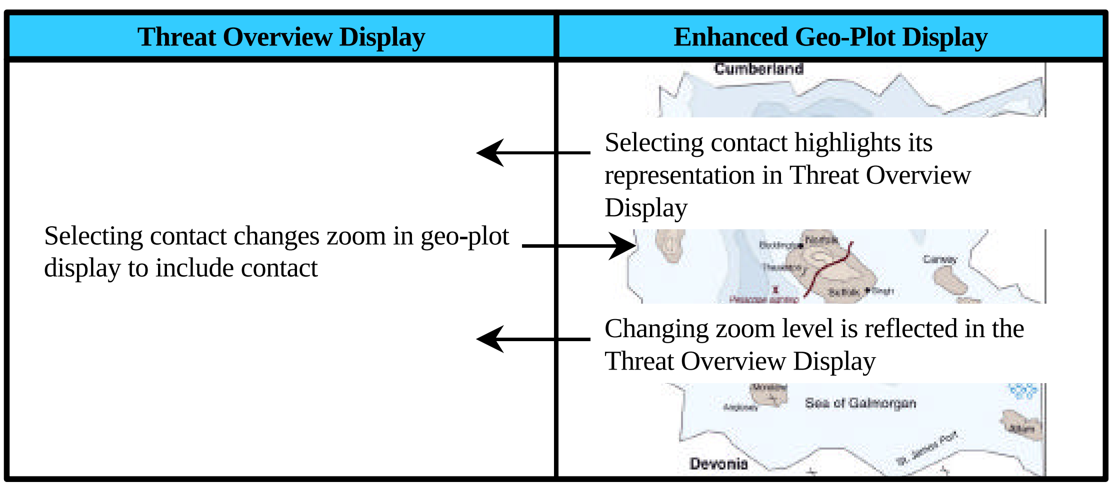

Figure 13. The two displays: Threat Overview Display and Enhanced Geo-Plot Display

The concept of coordinated functional and physical displays is illustrated in Fig. 13. It shows the functional display, the Threat Overview Display, on the left, and the more usual physical display, the Geo-Plot Display, on the right, along with their interactions. The contents of these displays are shown in more detail in Fig. 14. We identified various enhancements for incorporation in the physical display (e.g., route planning, combat systems status and performance predictions); however, this was not the principal thrust of the design effort that we describe here.

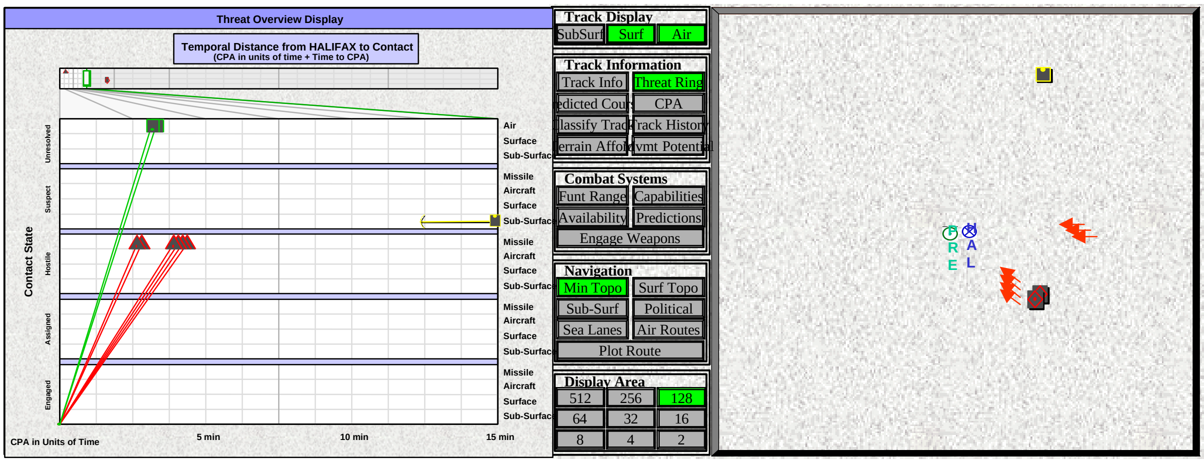

Figure 14. A more detailed view of the two displays[2]

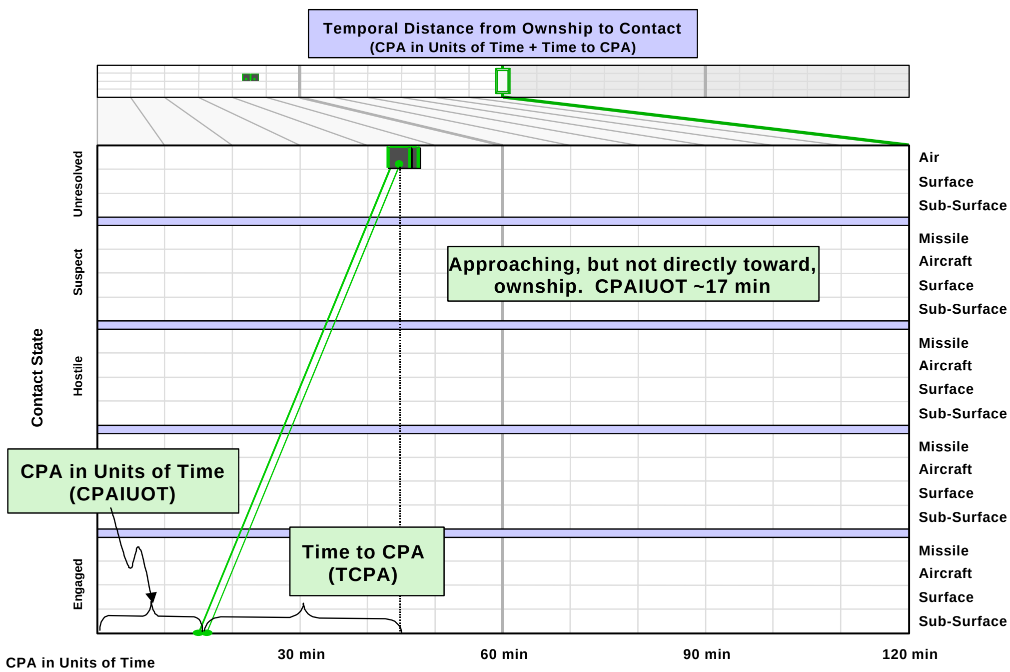

Figure 15. TCPA and CPAIUOT for two aircraft flying near ownship

As indicated in Section 7.4.2, a contact's functional distance is intended to capture notions of risk and imminence of tactically significant events arising from its observed or likely threatening actions. The current version of the Threat Overview Display uses two parameters, TCPA and CPAIUOT, for this purpose, associated with events surrounding the contact reaching its CPA[3]. Moreover, since these two parameters are represented in a common currency (i.e., time), we can represent both on the same plot and get a sense of the direction, relative to the ownship, of a contact's course (explained below). Such a representation is shown in Fig. 15.

---

[2] Display symbology and colour coding was based on [NATO STANAG 4420, Edition 2].

[3] TCPA is a natural choice for this purpose. CPAIUOT measures the time to reach ownship if the contact, on reaching its CPA, simply turns ninety degrees to ownship and maintains its current speed.

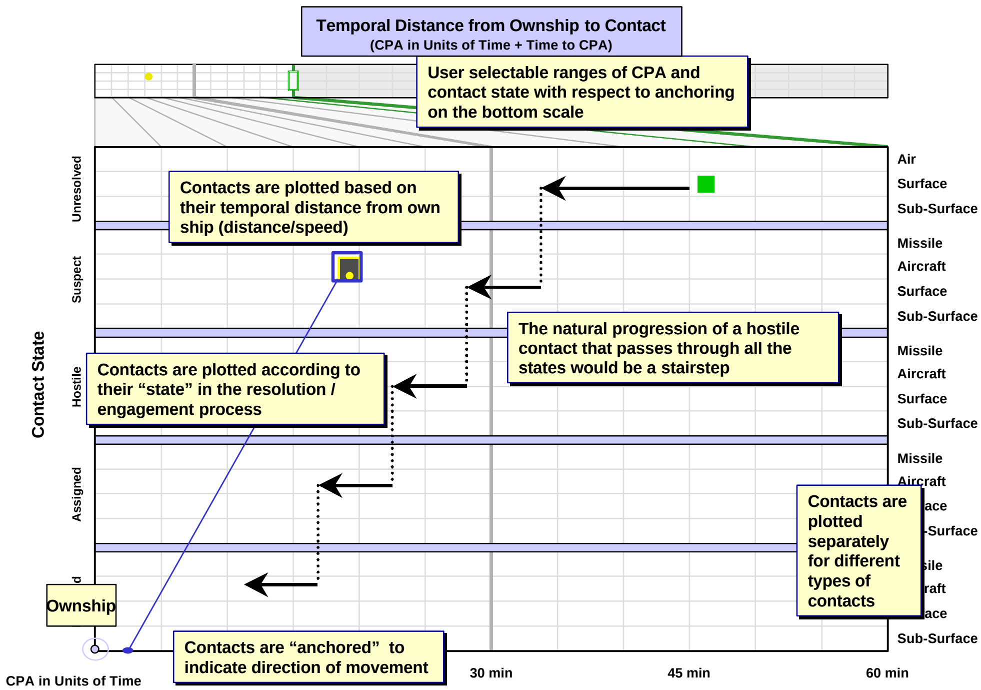

Figure 16. Attributes of the Threat Overview Display

The Functional Display is divided into two major windows (see Fig. 16). The first is the long narrow scale at the top, which has four rows. This is the 'overview' or 'big picture' portion of the display. It shows the set of contacts whose identification process is incomplete or which have some degree of threat attributed to them. Its time scale is large, typically two hours or more. This means that any object that ownship has sensed within a two hour range, given its current speed, appears in this space, located at its appropriate range (in units of time) from ownship. The four rows in this overview portion correspond to the categories of missile, aircraft, surface, and sub-surface. Within a row (or category), contacts may occupy the same temporal space. In this case, contacts are plotted top to bottom in a manner consistent with the detailed portion of the display (i.e., contacts from the top (unknown) may be obscured by hostile contacts). In this manner, higher priority contacts take precedence in this portion of the display space.

The second and larger window provides the detailed view of those contacts that have been classified as threats to the ship. The coordinates of this display space are Time-to CPA plus CPA in Units of Time (i.e., total time from the ship), shown on the X-axis, and the current state of a contact with respect to its classification and the ship's weapons engagement sequence, shown on the left-hand Y-axis. Each step in the classification and engagement sequence is divided into the appropriate number of parts to represent the contact's location in the three warfare areas.

The CPAIUOT is shown by a contact's anchor on the lower X-axis and will move to the right or to the left depending upon changes in the contact's course. If the contact does not change course and speed, then the location of the CPAIUOT will remain constant. The contact icon is 'anchored' to the CPAIUOT by connecting the two points with a straight line. Such anchoring helps to indicate the direction in which the contact is moving relative to ownship, i.e., a positively sloped line indicates that the contact is moving closer to the ship. The greater the slope (but not vertical), the closer the contact is to reaching its CPA point. If the CPAIUOT anchor is located on top of the ownship point (at the origin), then the contact is heading directly toward ownship. As the contact approaches its CPA, the anchoring line becomes vertical, and as it moves away from ownship, we have adopted the convention, based on SME inputs[4], that both the location of the contact icon and its CPAIUOT anchor move together and to the right, with the anchoring value on the X-axis at any given moment equal to the temporal distance of the contact from ownship.

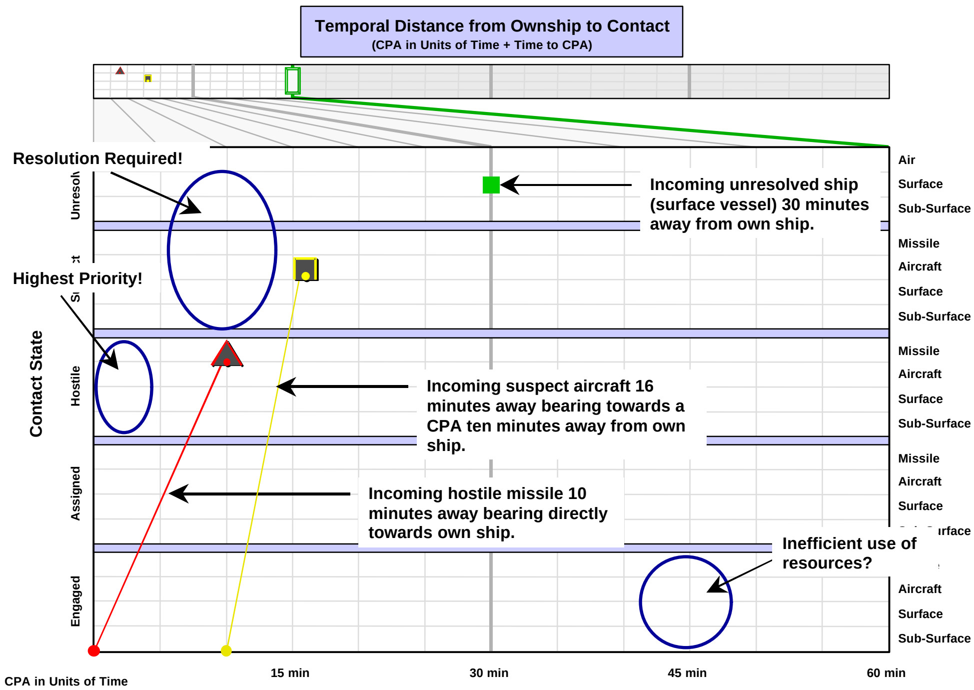

Figure 17. Additional attributes of the Threat Overview Display.

The coordinates of the display space help the ORO to:

- 'integrate' the three warfare areas;
- understand where contacts are with respect to their classification and identification;
- which threatening contacts are going to arrive at ownship first and, by implication, the priority of response that should be attached to each contact); and
- where each of these contacts is with regard to the engagement by onboard weapons.

This last point addresses the issue of providing the ORO with a means for assessing how well the resources under his/her command (the Operations Room staff and the weapons' automation) are performing with respect to preparing to deal with the threat. The 'natural' progression of a hostile contact in this window is shown in Fig. 16.

The time scale of this second window can be adjusted by the operator by using the cursor to slide the green-framed slider shown in Fig. 15 – to the right or left, depending upon his/her needs or

[4] This was related to SME concerns with an early version of this feature about losing track of the direction of movement of a contact, toward or away from ownship, as the ORO's attention shifts at a critical moment away from the Threat Overview Display to other matters and then back to this display.

desires. This permits temporally focusing in or out on the threat picture. Figure 17 shows some additional attributes about this display space which are a result of the choice of the coordinate system.

7.4.4 *Storyboard and SME Feedback*

SME inputs helped develop the details of a tactical scenario to be used as the framework for the animation sequences of a dynamic storyboard incorporating the visualisations described in Section 7.4.3. This knowledge acquisition also led to further refinements in existing CWA models. There were two cycles of evaluation and SME feedback of the visualisation concepts, the first during the early stages of development of these concepts, and the second as part of getting more specific feedback from a wider ORO community to the visualisation design itself. In the second cycle, the storyboard was evaluated by walking the SMEs through its operation in the context of the tactical scenario and obtaining their feedback both verbally and in the form of responses to a questionnaire.

The critical aspects of the tactical scenario included:

- a semi-realistic context in which to present and evaluate the visualisation concepts;
- inclusion of a variety of threats – air, surface, and sub-surface; and
- multiple, simultaneous threats of air, surface, and sub-surface origin.

In general, the SMEs found the learning curve required to make effective use of the new display concepts quite straightforward, and accomplished the task in a very short period of time. They were also able to provide many useful recommendations for further refinements and improvements of the visualisations. Some examples include:

- other parameters to consider in further developing the notion of functional distance to a contact (e.g., 'launch range' from ownship and time to it for a hostile platform prior to launching its weapons);
- the need to reflect in the display the greater locational uncertainties in the sub-surface domain relative to the above-water warfare; and
- a clearer separation of the aspects of the display concerned with the ORO's monitoring of the resolution process from those that are specifically concerned with the assignment and engagement of weapon systems to counter threats.

8. **Conclusions**

This paper has proposed an empirical, nonlinear framework as the basis for exploring the development of advanced decision aids to be incorporated in a decision support system within the CCS architecture of the HALIFAX Class frigate. It also reviewed a two-part project aimed at examining the feasibility and benefits of a specific Cognitive Systems Engineering approach to work analysis, known as Cognitive Work Analysis (CWA), as a component of this empirical framework. A comparison with other methods of work analysis, such as Task Analysis and Cognitive Task Analysis, was provided to help situate and motivate the need to look at CWA for determining support requirements, driven by the current and anticipated nature of the cognitive demands operators will have to deal with in this work environment. The second part of the feasibility study was concerned with providing, in the form of a case study, with results captured in a dynamic storyboard, a demonstration of the utility of the CWA framework in generating novel computer-based decision support to satisfy operator demands on the HALIFAX Class ship.

The study found that CWA is certainly applicable to the work environment of the HALIFAX. Focusing on a specific portion of the ORO’s workspace, as was done in the second part of the study, turned out, in hindsight, to be an advantageous aspect of the work. It provided a significant opportunity to demonstrate the potential of the approach, in an end-to-end manner, spanning knowledge acquisition, modeling and analysis, requirements generation, and design activities in the proposed, exploratory framework. This quickly exposed the framework’s nonlinearities. At the same time, it represented what we think is an important consideration, that should be followed as a guide in follow-on CWA projects, viz. that the complexity of this type of work environment is such that it is best attacked by making a strategic choice of the portion of the workspace to concentrate on for exploiting the strengths of the CWA framework. For example, we identified that the characteristics of sub-surface warfare, with its apparent emphasis on deliberation over reaction, are such that this domain may well offer a target area for promising returns from a CWA analysis. However, future work is needed to confirm this.

The FAH provided an underlying framework for making explicit the goals to be achieved in the domain, the supporting processes to achieve these goals, and the relationships between these goals and processes. This design artifact enabled determining where decision making is likely to be difficult due to the nature of the problem space (e.g., due to multiple interacting goals). Further, the FAH helped reveal the context of the decision, what the necessary inputs are, and the complicating factors that can arise (e.g., sensor failures, equipment malfunctions) to increase decision difficulty. From the FAH representation it was possible to derive the critical decisions required to achieve the domain goals. These cognitive activities and decisions centered around goal and process monitoring, planning, and control to achieve the identified goals. By organising the specification of operator information, decision, and control requirements around nodes in the goal-means structure, rather than organising requirements around predefined task sequences (as in traditional approaches to task analysis), the representation helped insure that the resulting visualisation reflects a decision-centred perspective.

Based on the results of the CWA, it was possible to design a visualisation to support the situation and threat assessment process by organising potentially aggressive threats according to two critical elements in the problem space – temporal “distance” from ownship and state in the resolution and engagement processes. This approach resulted in a transformation of the currently used notion of CPA into a unified, temporal perspective for managing potential threats.

It is important to note that the visualisation concepts developed in this project represent, nonetheless, an initial, untested concept. While the feedback received from the operational community was extremely positive, there needs to be follow-on development work to turn this into a robust prototype, based on taking full advantage of the feedback received so far. To support this, more realistic evaluations could be performed by having operators interact with higher fidelity prototypes. This type of investigation could range from possibly exploring whether or not there is added workload associated with adding a second display to the ORO's workstation during times of high threat activity to more thorough, iterative evaluations of the effectiveness of the concepts in the Threat Overview Display.

It is also important to note that the current visualisation concept addresses only a portion of the ORO's workspace as defined by the FAH, and that the current version of the FAH only covers a portion of the entire ORO workspace. Thus, additional CWA and visualisation design would need to be performed to add additional displays to cover other aspects of the ORO's workspace, besides threat management. A natural extension to the current target area of the case study discussed in this paper would be more explicit emphasis on the response management aspect of STA/RM (including the impact of ownship's position and manoeuvering capabilities). The functional display, when not being used to help prioritise threats and manage the ship's responses to threats, could be used for almost any other aspect of the ORO's assigned tasks. These tasks need to be explored and appropriate interfaces designed and validated.

Finally, we note that the strength of CWA lies in its formative modeling of work demands (constraints) as the basis for designing decision support. However, attempts to turn this into a rigorous, codified, and readily accessible engineering process (a complete end-to-end design methodology) are still of a preliminary nature. As examples, work is needed to develop further knowledge elicitation techniques that are applicable in building CWA models; and the expertise required to conduct a CWA is not currently widely available. The preliminary success of the case study in this paper, aimed at developing an advanced ORO interface, certainly lends support to arguments for further experimental CWA work and for advancing this process in naval C2. Related to this is the need for productivity enhancing tools as CWA is a very human intensive effort. The reader can find some recent attempts to address these challenges in the work of [Skilton *et al*., 1998] and [Potter *et al*., 1998]. Even more generally, it is worth noting that the characteristics of situations where CWA is likely to be very useful for design are, in general, exemplary of the Revolution in Military Affairs. This suggests the potential for good use of CWA not only in the tactical C2 arena, but also, perhaps even more so, at the strategic and operational levels of Command and Control.

9. **Acknowledgements**

The authors wish to thank LCdr Dean Morrissey who in his role as Naval Liaison Officer helped with arranging access to SMEs during the work on CWA. On several occasions, he also served as an on-site SME himself and helped provide many useful insights into the naval domain, as well as important inputs in developing a scenario for the ORO's storyboard. The authors also gratefully acknowledge the invaluable support provided by personnel of Canadian Forces Naval Operations School and Sea Training (Atlantic), HMCS Fredericton, HMCS Montreal, and Lockheed Martin Canada, who acted as enthusiastic SMEs at various stages of the CWA study.

10. **References**

[Alexander, 1964] C. Alexander. *Notes on the Synthesis of Form*. Harvard University Press, Cambridge, Masschusetts, 1964.

[Bertrand *et al*., 1999] S. Bertrand et al. *The Naval C4I Way Ahead: A Blueprint for 2010*. Draft Report, Canada, 1999.

[Brunswick, 1952] E. Brunswick. *Conceptual Framework of Psychology*. University of Chicago Press, Chicago, 1952.

[Burns and Bryant, 2000] C.M. Burns and D.J. Bryant. *Feasibility Study of Cognitive Work Analysis for the Command and Control Work Environment of a HALIFAX Class Ship*. Final Contract Report, PWGSC Contract No. W7701-8-1427/001/XSK, 2000.

[Burns *et al*., 2000] C.M. Burns, D.J. Bryant, and B.A. Chalmers. *A Work Domain Model to Support Shipboard Command and Control*. In preparation, 2000.

[Chalmers and Burns, 1999] B.A. Chalmers and C.M. Burns. *A Model-Based Approach to Decision Support Design for a Modern Frigate*. Proceedings of the TTCP Symposium on Coordinated Battlespace Management, Space and Naval Warfare Systems Center, San Diego, CA, USA, 1999.

[Chalmers, 1997] B.A. Chalmers. *Design Issues for a Decision Support System for a Modern Frigate*. Proceedings of the 2nd Annual Symposium and Exhibition on Situational Awareness in the Tactical Air Environment, Naval Air Warfare Center, Patuxent River, Maryland, U.S.A., 1997, pp. 127-149.

[Chin *et al.*, 1999] M. Chin, P. Sanderson, and M. Watson. *Cognitive Work Analysis of the Command and Control Work Domain*. Proceedings of the 1999 Command and Control Research and Technology Symposium, Vol. 1, 1999, pp. 233-248.

[Cooke, 1994] N.J. Cooke. *Varieties of Knowledge Elicitation Techniques*. International Journal of Human-Computer Studies, Vol. 41, 1994, pp. 801-849.

[Easter and Potter, 2000] J.R. Easter and S.S. Potter. *Decision-Centred Visualizations for the Command and Control Work Environment of a HALIFAX Class Ship*. Final Contract Report, PWGSC Contract No. W7701-8-3667/001/XSK, 2000.

[Endsley, 1996] M.R. Endsley. *Automation and Situation Awareness*. In R. Parasuraman and M. Mouloua (eds), “Automation and Human Performance: Theory and Applications”, Lawrence Erlbaum Associates, Inc., 1996, pp. 163-181.

[Flach *et al*., 1998] J.M. Flach, K.J. Vicente, F. Tanabe, K. Monta, and J. Rasmussen. *An Ecological Approach to Interface Design*. Proceedings of the Human Factors and Ergonomics Society 42nd Annual Meeting, 1998, pp. 295-299.

[Gordon and Gill, 1997] S.E. Gordon and R.T. Gill. *Cognitive Task Analysis*. In C.E. Zsambok and G. Klein (eds.), “Naturalistic Decision Making”, Lawrence Erlbaum Associates, Publishers, Hillsdale, Mahwah, NJ, 1997, pp. 131-140.

[Guerlain , 1993] S.A.E. Guerlain. *Factors Influencing the Cooperative Problem-Solving of People and Computers*. Proceedings of the Human Factors and Ergonomics Society 37th Annual Meeting, 1993, pp. 387-391.

[Hoffman *et al*., 1995] R.R. Hoffman, N.R. Shadbolt, A.M. Burton, and G. Klein. *Eliciting Knowledge from Experts: A Methodological Analysis*. Organizational Behavior and Human Decision Processes, Vol. 62, No. 2, 1995, pp. 129-158.

[Hoffman *et al.*, 1998] R.R. Hoffman, B. Crandall, and N. Shadbolt. *Use of the Critical Decision Method to Elicit Expert Knowledge: A Case Study in the Methodology of Cognitive Task Analysis*. Human Factors, Vol. 40, No. 2, 1998, pp. 254-276.

[Hollnagel, 1988] E. Hollnagel. *Information and Reasoning in Intelligent Decision Support Systems*. In E. Hollnagel, G. Mancini, and D.D. Woods (eds.), "Cognitive Engineering in Complex Dynamic Worlds", Academic Press, London, 1988, pp. 215-228.

[Kirlik, 1993] A. Kirlik. *Modeling Strategic Behavior in Human-Automation Interaction: Why an "Aid" Can (and Should) Go Unused*. Human Factors, Vol. 35, No. 2, 1993, pp. 221-242.

[Klein, 1996] G. Klein. *Nonlinear Aspects of Problem Solving*. Information and Systems Engineering, Vol. 2, 1996, pp. 195-204.

[Klein, 1997] G. Klein. *Implications of the Naturalistic Decision Making Framework for Information Dominance*. Report No. AL/CF-TR-1997-0155, Wright-Patterson AFB, OH, Armstrong Laboaratory, Human Engineering Division, 1997.

[Klein, *et al*., 1997] G. Klein, G.L. Kaempf, S. Wolf, M. Thorsden, M., and T. Miller. *Applying Decision Requirements to User-Centered Design*. International Journal of Human-Computer Studies, Vol. 46, 1997, pp. 1-15.

[Lind, 1991] M. Lind. *Representations and Abstractions for Interface Design Using Multilevel Flow Modelling*. In G.R.S. Weir and J.L. Alty (eds.), "Human-Computer Interaction and Complex Systems", Academic Press, 1991, pp. 223-243.

[McArthur *et al*., 1999] B. McArthur, A. Ashley, É. Bossé, and D. Beevis, *The Command Decision Aid Technology (COMDAT) Project*. Proceedings of the TTCP Symposium on Coordinated Battlespace Management, Space and Naval Warfare Systems Center, San Diego, CA, U.S.A., 24-28 May, 1999.

[Miller and Vicente, 1999] C.A. Miller and K.J. Vicente. *Task "Versus" Work Domain Analysis Techniques: A Comparison*. Proceedings of the Human Factors and Ergonomics Society 43rd Annual Meeting, 1999, pp. 328-332.

[NATO STANAG 4420, Edition 2] *Display Symbology and Colours for NATO Maritime Units*. Military Agency for Standardization, NATO.

[O'Neill, 1996] J. O'Neill. *The Role of Problem-Solving in C3I Systems*. Proceedings of the 2nd International Symposium on Command and Control Research and technology, Market Bosworth, Warwickshire, UK, pp. 348-355, 1996.

[Parasuraman *et al*., 1993] R. Parasuraman, R. Molloy, and I.L. Singh. *Performance Consequences of Automation-Induced "Complacency"*. The International Journal of Aviation Psychology, Vol, 3, No. 1, 1993, pp. 1-23.

[Potter *et al*., 1998] S.S. Potter, E.M. Roth, D.D. Woods, and W.C. Elm. *Toward the Development of a Computer-Aided Cognitive Engineering Tool to Facilitate the Development of Advanced Decision Support Systems for Information Warfare Domains* United States Air Force Research Laboratory Technical Report, AFRL-HE-WP-TR-1998-0004, 1998.

[Potter *et al*., 2000] S.S. Potter, E.M. Roth, D.D. Woods, and W.C. Elm. *Bootstrapping Multiple Converging Cognitive Task Analysis Techniques for System Design*. In J.M. Schraagen, S.F. Chipman, and V.J. Shalin (eds), "Cognitive Task Analysis", Lawrence Erlbaum Associates, Inc., 2000.

[Preece, 1994] J. Preece. *Human-Computer Interaction*. Addison-Wesley Publishing Company, Reading, Massachusetts, 1994.

[Rasmussen *et al*, 1994] J. Rasmussen, A.M. Pejtersen, and L.P. Goodstein. *Cognitive Systems Engineering*. Wiley, New York, 1994.

[Rasmussen, 1986] J. Rasmussen. *Information Processing and Human-Machine Interaction: An Approach to Cognitive Engineering*. North-Holland, New York, 1986.

[Rasmussen, 1997] J. Rasmussen. *Merging Paradigms: Decision Making, Management, and Cognitive Control*. In R. Flin, E. Salas, M. Strub, and L. Martin (eds.), "Decision Making Under Stress: Emerging Themes and Applications", Ashgate, Hants, England, 1997.

[Rasmussen, 1998] J. Rasmussen. *Ecological Interface Design for Complex systems: An Example: SEAD-UAV Systems*. Final Report, EOARD-Contract:F61708-97-W0211, 1998.

[Reising and Sanderson, 1996] D.V. Reising and P.M. Sanderson. *Work Domain Analysis of a Pasteurization Plant: Building an Abstraction Hierarchy Representation*. Proceedings of the Human Factors and Ergonomics Society 40th Annual Meeting, Santa Monica, CA: Human Factors and Ergonomics Society, 1996, pp. 293-297.

[Schmitt and Klein, 1996] J.F. Schmitt, and G.A. Klein. *Fighting in the Fog: Dealing with Battlefield Uncertainty*. Marine Corp Gazette, 1996, pp. 62-69.

[Skilton *et al*., 1998] W. Skilton, S. Cameron, and P. Sanderson. *Supporting Cognitive Work Analysis with the Work Domain Analysis Workbench*. Proceedings of the Australasian Computer Human Interaction Conference, Adelaide, SA, 1998.

[Vicente and Rasmussen, 1992] K.J. Vicente and J. Rasmussen. *Ecological Interface Design*. IEEE Transactions on Systems, Man, and Cybernetics, Vol. 22, No. 4, 1992, pp. 589-606.

[Vicente, 1995] K.J. Vicente. *Task Analysis, Cognitive Task Analysis, Cognitive Work Analysis: What's the Difference?*. Proceedings of the Human Factors and Ergonomics Society 39th Annual Meeting, 1995, pp. 534-537.

[Vicente, 1999] K.J. Vicente. *Cognitive Work Analysis: Towards Safe, Productive, and Healthy Computer-Based Work.* Lawrence Erlbaum Associates, Publishers, Hillsdale, Mahwah, NJ, 1999.

[Woods and Hollnagel, 1987] D.D. Woods and E. Hollnagel. *Mapping Cognitive Demands in Complex Problem-Solving Worlds*. International Journal of Man-Machine Studies, Vol. 26, 1987, pp. 257-275.

[Woods *et al*., 1988] D.D. Woods and E.M. Roth. *Cognitive Systems Engineering*. In M. Helander (ed.), "Handbook of Human-Computer Interaction", Elsevier Science Publishers B.V., North-Holland, 1988, pp. 3-43.

[Woods, 1985] D.D. Woods. *Knowledge Based Development of Graphic Display Systems*. Proceedings of the Human Factors Society 29th Annual Meeting, 1985, pp. 325-329.

[Woods, 1986] D.D. Woods. *Paradigms for Intelligent Decision Support*. In E. Hollnagel, G. Mancini, and D.D. Woods (eds), "Intelligent Decision Support in Process Environments". NATO ASI Series, Vol. F21, Springer-Verlag, Berlin, 1986, pp. 153-173.

[Woods, 1998] D.D. Woods. *Designs are Hypotheses About How Artifacts Shape Cognition and Collaboration*. Ergonomics, Vol. 41, 1998, pp. 168-173.