Development of a Mission Abstraction Requirements Structure (MARS) and stochastic modelling for sensing service-driven mission performance prediction

Dave Thornley, Rob Young and Jim Richardson

International Technology Alliance, Technical Area 3, Project 7, Task 1

This work is part of the International Technology Alliance in Network and Information Science. References to projects and tasks indicate internal structure, as explained at <http://www.usukita.org> under Projects. Each project has a number of sub-tasks.

Research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

1. Overview

Research into the synthesis of functions which predict the quality of information from a given sensing service in the ITA is well under way. It is clear that understanding QoI is essential to sensor network design, but a means for creating utility or value functions to support design quantitatively has been missing up until this point. We currently prefer to use the term “value,” finding that the two words correspond to identical notions when tested in modelling. The word “utility” is perhaps too close to “utilization”, which has a specific meaning in performance analysis: the utilization of a service is the fraction of its capability that is consumed.

We now introduce an approach to generating utility functions by abstracting from the physics of sensing and the delivery of QoI, to models which provide the measures necessary for judging the value of information. This is focused by the introduction of a Mission Abstraction Requirements Structure (MARS) that holds all the details required to synthesize QoI delivery functions as already occurs within ITA, but also to build the stochastic models necessary to quantify mission utility. MARS will specify the factors necessary for QoI delivery function synthesis, and indicate how these relate to the aspects of QoI that we leverage in the modeling process. The original motivation for formulating MARS was to enable the synthesis of a concrete link between QoI delivery during sensing, and the value of that information to the achievement of goals. In analysing the structure of missions from an information-driven perspective, we currently consider delivered QoI to drive knowledge building, which enables synthesis of situational awareness, which in turn enables decision making and policy application and management.

This document introduces and combines viewpoints that have given rise to a viable approach to synthesizing mission specific utility, value and cost functions. We begin by summarizing some of the motivation and resolution of quality of information research in the ITA to date. We then describe the process commonly used in the military to specify the information requirements of a mission. This is followed by a discussion of the sensor network lifecycle, within which we introduce our view on the meanings of quality of information. We then run through an example chain of modeling steps which introduces the quantitative link that can be made between quality of information and utility, by abstraction away from physics to a stochastic description of the sensing activity, which drives knowledge building, situational awareness synthesis, activity selection and hence mission goal achievement. This represents a huge modeling structure, and a very significant research challenge. We are beginning to indicate what modeling for the purpose of resource utility assessment will be approachable within the ITA.

We conclude that the probability of attainment of goals, and the time taken to do so, provide handles on the utility of the sensing service that provided the information from which knowledge was constructed and so forth up to the plan and goals. One important outcome of this modeling activity is therefore that solving for goal probability and timing for competing sensing services will enable comparison of the utility of those services. This in turn will support decisions concerning allocation of sensing services to competing missions, and optimization of sensing service parameters.

2. Introduction

The concept of quality of information (QoI) acts as a focus for development of methods for characterizing sensor network output in Project 7 of the ITA. A range of powerful approaches have been developed for synthesizing functions which provide an estimate of the quality of information (QoI) which will be delivered by a given sensor network under specific circumstances *e.g.* [Kaplan 08]. Furthermore, more general frameworks for analyzing classes of sensing service have been proposed [Zahedi 08]. These enable a user to predict how a sensing network will report a given ground truth. We use the term “ground truth” to represent the current and previous facts which were sampled by the sensor network for the purpose of informing a mission, for example by detecting an event. There has been a long-felt need for a quantitative link between these *QoI delivery functions*, and suitable *design optimization objective functions* which would support sensor network and service design.

It has been proposed that the necessary objective functions will take the form of *utility functions*, which describe how useful a design will be. Frameworks for optimizing parsimonious designs in terms of asset commitment and energy usage [Pizzocaro 08, Johnson 08], and for addressing competition for resources across missions [Eswaran 08] have been proposed, and demonstrated using place-holder utility functions which reflect the form of the QoI delivery function directly. However, quantitative, normalized, mission-specific comparison of competing sensor network and service designs both in terms of parameterization of a given bundle, or in comparing competing bundles, requires predictive measures of their operational performance in the mission.

Any sensing service (combination of assets and approach to their use) will have strengths and weaknesses. A sensing service has one or more subjects that we intend it should measure in some way. A subject may act either consciously or coincidentally to exploit less capable aspects of the service (e.g. take a route over which lower QoI is delivered). Both cases must be fully described, ascribed relative likelihoods and importance, and their impact on mission outcomes measured (probability of success, time taken to achieve).

Prediction of the performance of a sensing service requires models of the subjects, sensing services, mission actors, orders and policy, knowledge construction, situational awareness synthesis and goals. Each model comprises states, transformation or transition functions, prior probabilities and reward measures. The states describe aspects of the ongoing situation which contribute or consume at a constant rate, such as a tracking service maintaining acceptable accuracy, using a given combination of assets and processing which consume battery life at an approximately constant rate.

This is an example of modeling which offers the required link between quality of information research, and the requirements for quantitative support of design and operation of sensor networks. The gamut of models necessary for modeling defense activities in general will of course be large, and we are just beginning to formulate the models that will populate that space. However, we can identify modeling features that are essential to the representation of a mission driven by a sensor network. In particular, these models will expose measures pertinent to the value/utility/benefit or cost of a proposed sensing

A *bundle* is a combination of sensing, platform, processing and communication assets [Preece 08]. Preece et al refer specifically to assets already deployed, supported by the middleware to which a request is to be made. Bundles are proposed initially in terms of satisfaction of the semantic aspects of the information requirement, *e.g.* a pair of suitably spaced acoustic arrays may perform spatial localization of a source, and the addition of a low-power radar can enhance resolution. To test such a bundle for the purpose of selecting operational parameters which balance energy consumption against the stability of the mission, and for comparison of effectiveness against another bundle also capable of the necessary semantics, it is necessary to model the services provided as it they could be expected to behave in the mission. This is a stochastic modeling task, since the subjects to be instrumented are not under our direct control (we may be able to constrain behaviours to some extent). In essence, the mission must be war-gamed with the proposed sensing service. Traditional Monte Carlo war gaming methods provide example event sequences to support qualitative introspection during planning. To measure the relative likelihoods of outcomes contingent on complex interactions in competing sensing services, however, requires an analytic treatment working directly with probability distributions, as opposed to a simulation treatment.

We introduce here the Mission Abstraction Requirements Structure (MARS) and illustrate the modeling activity supported by it. This acts as our focus for developing a generalized quantitative link between a sensing service's characteristics and its value to a mission. A specific, geometrically simple sensing service may be amenable to analysis against a simple mission behaviour projection to give a neat closed form effectiveness measure *e.g.* [Bisnik 06], but we require a general method which will cope with complex, novel sensing methods and information requirements, and which can be composed with other sensing activity for cueing or decision making. We require re-usable modeling components that are carefully characterized, and which may accrue metadata describing how well they functioned in example operations. This metadata may contribute to the functioning of the sensor mission matching stage, but will generally be followed by analysis against the specific mission being addressed. These modeling components encapsulate probabilistic descriptions of the behaviours and their interactions.

The motivating issue is assessment of information provision, and specifically the attainment of suitable information quality where and whenever it is needed, and balanced across the mission as a whole. We do this by abstracting from the physics of the sensing process, and the functional details of fusion algorithms, to give a behavioral model that can be analyzed in combination with mission activities and goals to enable assessment of the utility and value of its components.

2.1. Quality of information

Work in the ITA exploring the quality of information available from sensor networks began with an intuitive recognition that information of a higher quality would be of benefit to military operations by providing increased acuity and confidence in perception and management of engagements. The aim was then to improve this quality to the benefit of sensor driven activities carried out by US and UK forces. A common goal of such a scientific endeavour is to establish effective generalizations which enable formulation of beneficial results. In common with the substantial literature on information theory, we find that a tighter distribution of potential meanings of an item of information corresponds to a higher quality of information, since the response can be more confident, and hence has the potential for remaining pertinent and complementary to novel information through more of the subsequent evolution of a scenario.

This focus has driven work with results that comprise improvements to the quality of information delivered by some classes of sensor network functionality (predominantly using existing devices utilized through novel methodologies), and just as importantly, methods for specifying that quality to the customer in theatre such that they may respond appropriately. The form in which quality of information received is specified is dictated by the type of sensing function, and the application in which it is to be used.

The philosophy behind efforts in Project 7 has evolved from an initial desire for a general solution to the specification of QoI, to a recognition that measures which describe quality are highly application specific, but can be broadly taken to fall under headings of accuracy, confidence, time-dependency and trust.

While work to improve the *capabilities* of quality of information delivery by sensor networks continues, we have focused on finding a means by which those quality of information *capabilities* of a sensor network are linked to the *requirements* of command *per se,* and in the context of limited resources, both in terms of provision at the outset of a mission, and due to attrition through damage or predictable power failure. This constitutes research toward quantifying and predicting the *information provision performance* of a sensor network, which will then enable calculation of its *value* in the mission. This measure of individual value (for example inferior/ appropriate/superior) can then be balanced against the utility of competing candidates during the process of *designing* a deployment (perhaps to rank them in order of *preference*). Comparison of designs for competing missions then will enable calculation of measures of the *value* of a design in the missions competing for deployment. In these early stages we are looking at supporting requirements as resolved by decision makers in the command structure.

The sensing, communication and processing functions can be maintained in a suitable middleware, for example the ITA Sensor Fabric [Bergamaschi 08]. A semantic approach to selection of sensing bundles implemented in the fabric has been explored [Preece 08] such that suitable resources may be proposed, and requests placed with the fabric. The work we describe here introduces mission-specific quantitative estimation of the value of a sensing service to a mission. We begin with a proposed service description, whether this be use of existing infrastructure, or a proposal for a new or modified deployment. We then seek to synthesize models which enable comparison of the value of providing alternate services to the specific mission structure by predicting the consequences.

In this paper we consider information requirements about phenomena (objects or events, and environmental factors) in the physical world at levels of abstraction required for modeling of the sensor network *in service*. In the present introduction of our approach, we assume pre-defined, parameterized information requirements, and establish a mathematical statement of their meaning, e.g. detect gunfire or detect and track a vehicle in a specified area, over a specified timescale, with a projection of the potential behaviours of the mission elements suitable for stochastic modeling. This includes a description of how the subject may move and select its path depending on the mission circumstances, how the sensor network will respond to this, relevant elements of situational awareness, and in turn how situational awareness may affect what sensing services are required.

Rather than taking a philosophical approach, we propose a synthesis downward from command specifications toward the physics of the requested and necessary actions, identifying measures necessary to that process, many of which can be considered aspects of information quality. We briefly describe information quality at a small number of levels of abstraction marking stages of development of the quantitative link between command requirements and sensing physics.

Our contribution to the isolation and definition of QoI comes from a holistic analysis of the way missions are driven by intelligence to achieve command goals. Aspects of information and knowledge that can be called QoI are seen at many different levels, but when we consider stochastic modelling for performance prediction, it always boils down to a probability distribution over states, or transitions to alternative states. For example, the QoI of knowledge about a vehicle includes its class, which may be distributed among {tank, truck, SUV}. The QoI of an incoming report about a vehicle's can cause a transition from, for example, a state indicating that tracking is being maintained, but no class ascribed, to one which includes a description of the vehicle class. This target state will therefore embody knowledge selected from {tank, truck, SUV}. In an analytic stochastic model, there are a number of choices as to how this might be achieved. The state itself may encode the relative probabilities of the three types being the case, or we can have the relative probabilities made implicit in the model by providing a model state for each vehicle, and allowing the relative probabilities of each resolve when that piece of knowledge is used to select an outcome.

In Figure 0 on the next page, we sketch out the design and modelling progressions, the contents of which we introduce and begin to develop in detail later in the report. What we find is that QoI dictates the relative probability of alternate choices during the mission, and the timing and duration of those choices and resulting actions. Accuracy, confidence, timeliness and trust have so far proven suitable descriptions of the type of probabilistic values we need to work into the models. But, as is the case with many of the words used to describe what are in essence mathematical concepts, the interpretations are not always unique. This is one motivation for developing the Mission Abstraction Requirements Structures. It is intended as a repository for mathematical descriptions of mission elements. Translation to and from English or a coalition language is a challenge we address here only insofar as it is helpful in explaining why we are performing specific modeling steps.

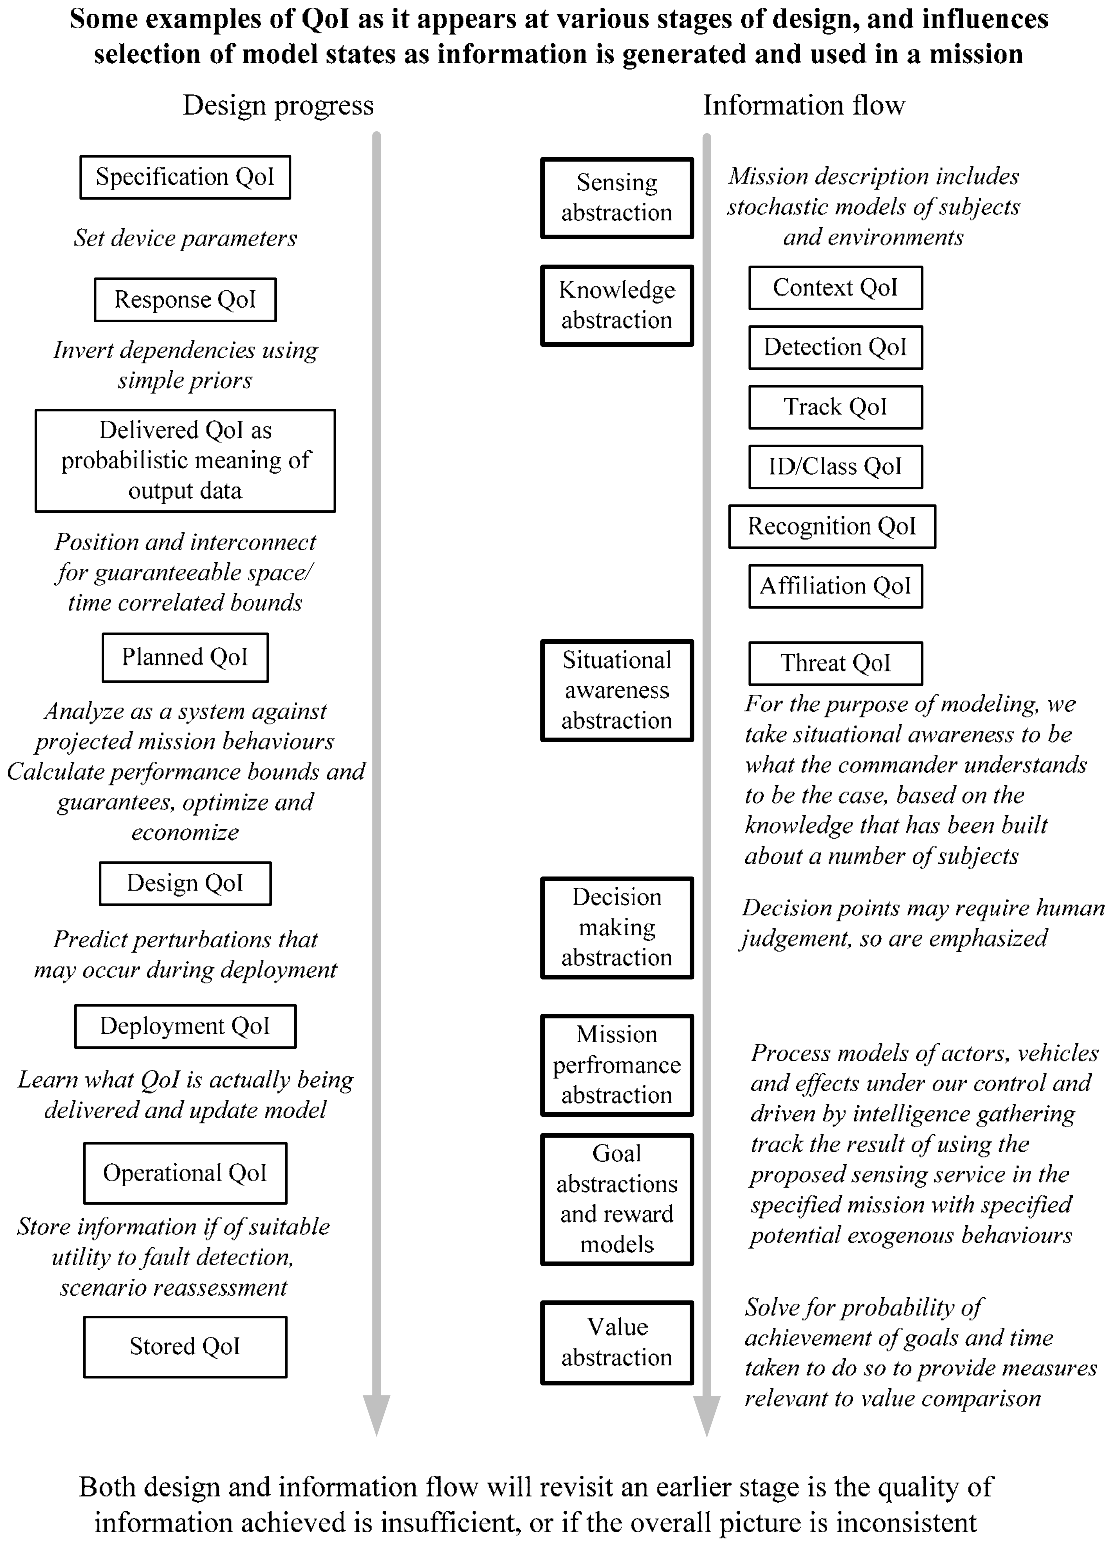

Figure 0: On the left, we suggest a breakdown of types of QoI to be assessed as design proceeds. On the right, we indicate the abstractions we use in our mission modelling to provide a concrete

link between information provision and outcomes in the mission, giving measures that can be ascribed comparative value.

3. Information requirements for sensor networks

To understand the context in which sensor networks are deployed on Intelligence, Surveillance and Reconnaissance (ISR) missions, it is necessary to be aware of the needs these missions are executed to meet, and of the background context of information gathering. This section of the report describes the background to how the UK military currently specify their information requirements and then go about meeting them.

The process for doing this is called Intelligence, Surveillance, Target Acquisition and Reconnaissance (ISTAR) that are so vital in an early stage of operations to set the framework for subsequent defeat of an enemy. ISTAR is defined as ‘The co-ordinated acquisition, processing and dissemination of timely, accurate, relevant and assured information and intelligence which supports the planning, and conduct of operations, targeting and the integration of effects’. It supports the commander and his decision making process to enable the commanders to achieve their goal throughout the Spectrum of Conflict.

ISTAR, when combined with an effective C4I (Command, Control, Communications, Computers, and Intelligence) capability, is known as C4ISTAR, and is the practical application of the Find Concept. The term Intelligence, Surveillance and Reconnaissance (ISR) is used at the Joint level. It follows the same basic concept as ISTAR but does not tie targeting and ISR to the same degree based on the perception that time-critical target engagement reduces at the operational level and above.

ISTAR is considered to be a “System of Systems” consisting of three principal elements: information, either raw or processed; the processes that enable information to be collected, collated and analysed into intelligence; and the physical architecture that encompasses the ISTAR collection systems, their organisations and the various staff cells. All three elements must be treated equally, because if any one element is neglected, there will be a reduction in the overall efficiency of the system.

To date there has been no single, simple methodology for assessing these three elements on a common basis. Indeed the emphasis of the intelligence is how it provides evidence in support of decisions made rather than act as a direct measure of Quality of Information. With the rise of pervasive sensor networks and the increasing frequency with which the sensor networks are deployed to make observations other than that it was fundamentally designed for the technology creates a strong desire for the definition of simple QoI metrics to support the design and deployment of ad hoc sensor networks. Such universally applicable metrics have not yet been found by analysing sensing approaches per se, and this motivates the investigation behind this technical report, which is working back from command requirements, through mission modelling, to the information sources, looking for opportunities to formulate suitable quality metrics.

3.1. Supporting the common operational picture and situation awareness

ISTAR supports the maintenance of the Common Operational Picture and the development of Situational Awareness (SA) for commanders, staffs and other users at all levels of command both inside and outside of the Land Component. Situational Awareness is the understanding of the operational environment in the context of a commander's (or staff officer's) mission (or task).

The Common Operational Picture is a snapshot in time of friendly and adversary forces and of the battle space environment. It is overlaid over a common geospatial framework and formed from the output from the database of information and intelligence which is common throughout a level of command and which is disseminated throughout that level of command. The Common Operational Picture informs the commander's Situational Awareness. ISTAR contributes to the Common Operational Picture by providing the database of information and intelligence on the adversary and, to a certain extent, on the environment. It is therefore a contributor to the Common Operational Picture and Situational Awareness but provides only part of the data.

ISTAR therefore requires skilled and/or automated coordination and combination of information sources to provide the potential battle-winning system of systems. It requires imaginative skill to achieve coherent ISTAR planning coupled with careful attention to detail and manipulation of the many platforms, systems and devices available to support a commander on the modern battlefield.

3.1.1. The Intelligence Cycle

The ISTAR process revolves around the four elements of the Intelligence Cycle – Direction, Collection, Processing and Dissemination. This cycle is adopted across the joint community and by many other countries. Traditionally the Cycle has focused on determining what information is required, collecting that information and then processing it to gain intelligence. The ISTAR process brings the Intelligence Cycle together with the integrated and co-ordinated use of collection assets.

The Intelligence Cycle (see Figure 1), is the fundamental process for developing intelligence, is central to the ISTAR Process. The Intelligence Cycle consists of Direction, Collection, Processing and Dissemination; each of these elements is a process in its own right, which include a number of sub-processes.

The sub-processes represent the sequence of activities whereby information is obtained, assembled, converted into intelligence and made available to users. This sequence comprises the four phases of Direction, Collection, Processing and Dissemination.

These sub-processes contribute to the production of intelligence products and are applicable at all levels of command. The impetus to the ISTAR process is the commander, whilst detailed requirements are provided through both the IPB and the Intelligence Estimate. These tools are the primary means of determining what intelligence is required and how the ISTAR collection assets can be used to gather the necessary information.

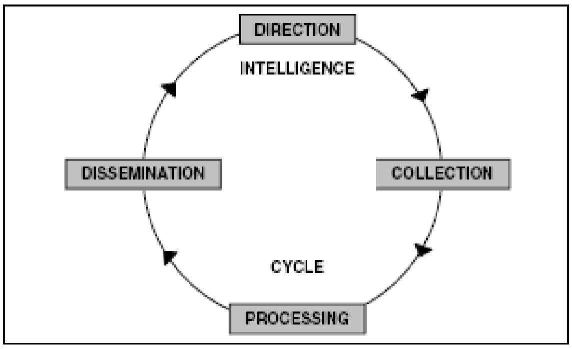

Figure 1. The Intelligence Cycle

Direction is the determination of intelligence requirements, planning the collection effort, issuance of orders and requests to collection agencies and maintenance of a continuous check on the productivity of such agencies.

The collection sub-process entails the exploitation of sources by collection agencies and the delivery of the information obtained to the appropriate processing unit for use in the production of intelligence.

Processing involves the conversion of information into intelligence through collation, evaluation, analysis, integration and interpretation.

Collation, in intelligence usage, is a step in the processing phase of the Intelligence Cycle in which the grouping together of related items of information or intelligence provides a record of events and facilitates further processing.

In intelligence usage, evaluation is the step in the processing phase of the intelligence cycle constituting appraisal of an item of information in respect of (a) the reliability of the source and (b) the credibility of the information.

Analysis is the step in the processing phase of the Intelligence Cycle in which information or intelligence is subjected to review in order to identify significant facts for subsequent interpretation.

Integration, in intelligence usage, is a step in the processing phase of the intelligence cycle whereby analysed information and/or intelligence is selected and combined into a pattern in the course of the production of further intelligence.

Interpretation is the final step in the processing phase of the intelligence cycle in which the significance of information and/or intelligence is judged in relation to the current body of knowledge. The term can also be used in its more usual sense of translating raw data into a more intelligible form - for example as in imagery interpretation.

Dissemination is the timely conveyance of intelligence, in an appropriate form and by any suitable means, to those who need it.

3.1.2. Information and Intelligence

The purpose of ISTAR is the production of intelligence from information. In order to understand the ISTAR system, the clear distinction between information, intelligence and data must be understood:

a. Information. Information is unprocessed data of every description that may be used in the production of intelligence. It is normally collected by individual sensors, systems or capabilities.

b. Intelligence. “The product resulting from the processing of information concerning foreign nations, hostile or potentially hostile forces or elements or areas of actual or potential operations”. Intelligence is the result of a process involving the evaluation, analysis, integration and interpretation of disparate pieces of information, usually in relation to existing information and intelligence, in order to clarify a situation and produce meaningful conclusions, assessments and predictions in response to the commander's intelligence needs.

There are three types of intelligence:

(1) Basic Intelligence. Basic intelligence is intelligence on any subject, which may be used as reference material for planning and as a basis for processing subsequent information or intelligence.

(2) Current Intelligence. Current Intelligence is intelligence which reflects the current situation at the strategic, operational, or tactical levels.

(3) Applied Intelligence. Applied Intelligence is intelligence that has been analysed and refined to a degree.

Intelligence Requirements are those items of intelligence required by a commander in order to conduct current operations and to plan future ones.

Intelligence Preparation of the Battlefield/space, (IPB), is the systematic and continuous process of analysis of adversary/targeted force doctrine, order of battle, weather and terrain matched against the friendly commander’s mission in order to determine and evaluate the threat’s/targeted force’s capabilities, intentions and vulnerabilities.

Information Requirements (IR) are those items of information regarding the enemy and his environment which need to be collected and processed in order to meet the intelligence requirements of a commander. Given the Commander has limited resources he prioritises on how to deploy sensors by stating different levels of priority to the intelligence requirements. The most essential elements of information are those which contribute to the Priority Intelligence Requirements (PIR). These are those intelligence requirements for which a commander has an anticipated and stated priority in his task of planning and decision making.

3.2. How information requirements are specified

Typically information requirements that are questions about the physical world at some level of abstraction, usually limited to a specific geographic region and duration of time. It is currently possible for the commander to ask arbitrary questions. To enable computational processing of requests, it is desirable that the question should be expressed using a formal query language that can be mapped (in an accreditable manner) to the mathematical descriptions we require for comparison with modelling results.

For example [Preece 08] illustrate how high-level information requirements (IREQs) are defined informally by the use of natural language. They quote an example IREQ as asked by the Decision Maker: “Is there suspicious activity on the MSR road?”

The Mission Planner then takes this high-level IREQ and from it develops a set of scenario-specific information requirements (SSIRs), once more detailed in an inform manner using natural language. Using the example from Preece, the SSIRs could be:

- “Are there suspicious vehicles on the road?”
- “Is there suspicious pedestrian activity along the roadside?”
- “Are there suspicious objects located near the road?”

The SSIRs need to be broken down further to detail the phenomenology of a set of signatures, that if observed would provide sufficient information to answer the SSIR.

It is at this stage that the matching of sensor asset to mission can be facilitated. That is, the planning of sensor deployment can occur. Instances of predefined sets of intelligence questions quoted in the open literature include:

1) The Decision Support Centre database which contains 3500 specific questions or needs in support of the expected intelligence requirements expected to be evaluated in support of Precision engagement, Dominant manoeuvre and SEAD (Suppression of enemy air defence) missions ([Hall 01] or see www.dsc.osd.mil).

2) The Missions and Means Framework (MMF) [Gomez 07, Sheehan 03] provides a disciplined procedure to explicitly specify a mission, allocate means and assess mission accomplishment by:

- Unifying warfighter, engineer and comptroller understanding of missions and means
- Accounting for traditional T&E factors and traditional warfighter expertise factors that constitute mission success
- Being sufficiently credible, timely and affordable to make hard decisions that stay made
- Being consistent, concise, repeatable and scalable

3.3. The characteristics of effective intelligence products

Effective intelligence is intelligence that meets the commander's needs. In order to achieve this, intelligence products must have the following characteristics:

a. Relevance. Intelligence must support the commander's mission, concept of operations, and intelligence requirements.

b. Usability. Intelligence products must be in a format that can be easily used and they must highlight the significance of the information or intelligence they contain and include an assessment of the validity of the intelligence.

c. Timeliness. Intelligence products must be available in sufficient time to enable decisions to be made and implemented.

d. Objectivity. Intelligence must be unbiased, undistorted, and free from political influence or constraints. Intelligence methodology and product must not be directed or manipulated to conform to a desired result, preconceived views of a situation or adversary, or a predetermined objective or institutional position.

e. Availability. Intelligence must be readily available to those who need it, taking account of both source protection and the penalty of intelligence compromise. This may require pragmatic decisions to be made on the releasability of intelligence.

f. Completeness. Intelligence should be as complete as possible using all available information to answer customers' requirements and provide a full understanding of the situation whilst not slowing down the dissemination process.

3.4. The JDL model of data fusion

Intelligence products are derived from data gathered by sensors. The transformation of sensor data into intelligence is best described by the JDL fusion model. The JDL describes the sensor deployment and relates performance of deployment towards obtaining situational awareness.

The diverse and complex techniques researched in data and information fusion can appear initially daunting to comprehend and categorize against the higher-level military needs. As part of their mission to develop broad based, multi-service research and technology demonstrations in C3, the US Joint Directors of Laboratories (JDL) Data Fusion Group proposed a widely accepted functional model and lexicon for fusion. The current JDL model (revised in 1998) is the most widely used framework for categorizing data fusion related activities [Steinberg 99].

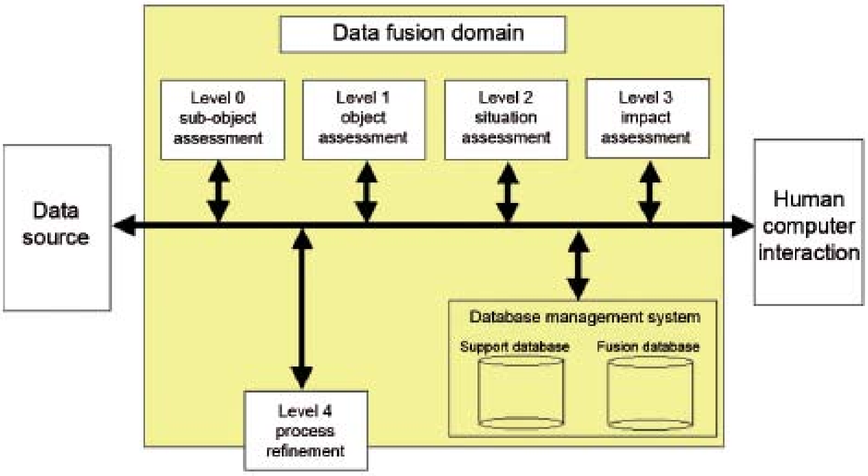

Figure 2. The JDL model of fusion levels

The JDL model proposal recognized that confusing data fusion terminology used in the fusion community had masked common research activity and reduced fusion technology reuse. The model presently consists of five levels (Figure 2), which address the refinement of sub-objects, objects, situation, impact and process. The levels are defined later in this section.

Proposed JDL level definitions are:

Level 0 Sub-object assessment—the estimation and prediction of signal / object observable states on the basis of pixel / signal level data association and characterization. This level of fusion is directed at improving the detection of entities against the environment.

Level 1 Object assessment—the estimation and prediction of entity states on the basis of observation-to-track association, continuous state estimation (e.g., kinematics) and discrete state estimation. This level of fusion activity is directed at integrating source data to improve the estimate of an entity's position, velocity, attributes and identity.

Level 2 Situation assessment—the estimation and prediction of relations among entities, to include force structure and cross force relations, communications and perceptual influences, physical context, etc. This level of fusion activity is directed at dynamically hypothesizing a description of the relationship between entities and events against the environment.

Level 3 Impact assessment—the estimation and prediction of effects on situations of planned or estimated / predicted actions by the participants; to include interactions between action plans of multiple players (e.g., assessing susceptibilities and vulnerabilities to estimated / predicted threat actions given one's own planned actions). This level of fusion activity is directed at projecting the current situation into the future and to implying enemy threats, blue and red vulnerabilities and operational opportunities.

Level 4 Process refinement (an element of resource management)—the adaptive data acquisition and processing to support mission objectives. This level of fusion activity is directed to monitoring and adapting the complete fusion system and managing the improvement of system performance.

In an attempt to clarify the relationship between the JDL model and more traditional frameworks of the intelligence process, DERA published a proposed mapping [Bedworth 99], which compares the functionality achieved at each JDL level against other traditional military intelligence models (see Figure 3).

<table>
<tr><th>Activity being undertaken</th><th>Waterfall model</th><th>JDL model</th><th>Boyd loop</th><th>Intelligence cycle</th></tr>
<tr><td>Command execution</td><td></td><td></td><td>Act</td><td rowspan="2">Disseminate</td></tr>
<tr><td>Decision-making process</td><td>Decision making</td><td>Level 4</td><td>Decide</td></tr>
<tr><td>Threat assessment</td><td></td><td>Level 3</td><td rowspan="5">Orient</td><td rowspan="2">Evaluate</td></tr>
<tr><td>Situation assessment</td><td>Situation assessment</td><td>Level 2</td></tr>
<tr><td rowspan="2">Information processing</td><td>Pattern processing</td><td rowspan="2">Level 1</td><td rowspan="3">Process / collate[2]</td></tr>
<tr><td>Feature extraction</td></tr>
<tr><td>Signal processing</td><td>Signal processing</td><td>Level 0</td></tr>
<tr><td>Source / sensor acquisition</td><td>Sensing</td><td></td><td>Observe</td><td>Collect</td></tr>
</table>

Figure 3. A proposed mapping of JDL to other intelligence models

A recent proposal [Bosse 07] has additionally suggested integrating a sixth level addressing cognitive refinement, which would include such functions as the modelling and compensation of human perception processes, limitations and biases through improved human computer interfaces (HCIs). This is likely to be incorporated into a future revision of the JDL model. It is important to recognize that the enumeration of the levels does not define the order of the data flow in a fusion processor. Instead a fusion process can dynamically apply functions in different levels in any order and in parallel to achieve its objective. The model simply provides a high-level taxonomy for categorizing the role of fusion technology components. A figure to show the iterative continuous nature of the inference process related to the JDL model is shown in figure x-3.

3.4.1. Situational awareness

Fusion at the higher JDL levels, i.e. information fusion, seeks to provide a machine counterpart to situation awareness. Situational awareness is commonly defined as the “perception of elements in a volume of space and time, the comprehension of their meaning and the projection of their status in the near future” [Endsley 95].

In Situation Assessment, hypotheses concerning relationships and situations are built and evaluated, whether implemented by people, automatic processes or some combination thereof involves inferences of the following types:

a) Inferring the presence and the states of entities on the basis of relationships in which they participate;

b) Inferring relationships on the basis of entity states and/or other relationships;

c) Recognizing and characterizing observed situations;

d) Projecting unobserved (e.g. future) situations. [Steinberg 08].

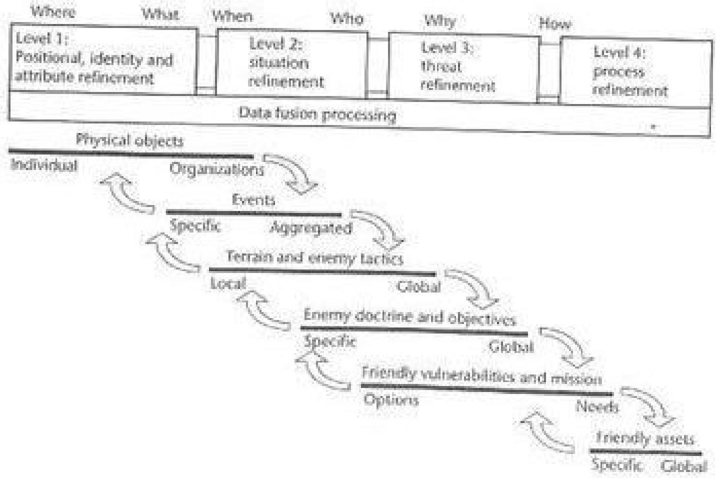

Figure 4. Linking Situational Awareness to the JDL model of fusion

Figure 4 is effectively a schematic to link situational awareness (how the military think about the utility of information) to a description of the information gathering process, the JDL model of fusion (which describes fusion in manner that facilitates building sensor networks). Surprisingly these views of how to gather information and then use the information have never been seamlessly linked up. The complete alignment of the JDL and SA model continues to be a key research topic.

3.5. Determining the requirements for QoI

The challenge for the sensor network designer is to determine a manner in which the quality of data collected by his deployed networks can be translated into a measure of the value of the information it will provide. A key stumbling point to date is that QoI at the datum level is application specific whereas the mission planner can task a variety of different applications to complete the desired mission.

Furthermore it is obvious that at the higher JDL levels concerned with situational awareness not all of the elements of interest to a decision maker are directly observable using the data and information sources available. Given the aspects of interest cannot be directly observed they must be inferred. Inference is perceived as being an essential act of information fusion. It indicates that the data fusion processes of L1 are very different to those needed for information fusion. Given the variety of fusion mechanisms included in the JDL model it is not surprising that there is no single, general metric is available to evaluate the performance of all the fusion algorithms.

4. Quality of information, the sensor network lifecycle and value of information

There are many ways of breaking down the space of possible QoI meanings, which has led to the apparently incoherent range of definitions used in research so far. This range arose naturally from exploration of measures of the output of a sensing service that reflect their capabilities. In this section, we discuss the faces of QoI that appear through the sensor network lifecycle.

We can identify “specification QoI”, which is a property of the equipment as it comes out of the factory. This is distinct from “deployment QoI”, which is description of the capability of the sensor network when deployed. We also see “operational QoI”, which is a measure of what the sensor network delivers during the mission or set of missions. We are motivated by having the required operational QoI matched by the actual operational QoI at each stage.

Generally speaking, QoI depends on deployment parameters D, environmental parameters E and the actual behaviour B of the phenomenon to be sensed, then statements about “specification QoI” must make assumptions about D, E and B because all are unknown at that time; Statements about “deployment QoI”, both D and E are assumed known, and about the effects of B are are use only in terms of bounds or means. For “operational QoI”, we assume a stochastic description of B according to the mission description, and perhaps incorporating considerations used in its planning and design if such information is available, and predict measures of QoI that can be compared with command specifications.

4.1. The meaning of delivered information

The meaning of a report is a statement about the probability distribution of realities for the phenomenon of interest which could have led to that report. This may also include history, so may require reference to knowledge or situational awareness.

A commander wants to know ground truth, but in general a sensor network cannot provide that. A sensor network delivers information as reports, each of which has meaning and value. We take a stochastic approach, in which a report is interpreted as a probability distribution over the set of possible ground truths. For example, a gunfire detection system issues a report whenever it detects a shot. The meaning of this report might mean “the probability of a gunshot in the designated area in the past second is at least Pd”. The lack of a report (equivalent to a negative report) might mean “the probability of a gunshot in the designated area in the past second is no more than Pf”, with more detail if we have a suitable model of the opposition.

The "quality of information" of a report is a representation of the meaning of a report in terms of the locus of realities which could have led to that report being generated. For example, a report of location of a target at a range of 100m coming from a specified and characterized sensor network may actually imply that the target may be located according to a stochastic distribution which is Gaussian about the range of 100m with a standard deviation of 1m. The quality of information of that report might sensibly be summarized by "+/-1m". This is interpretable for a number of intended uses, including for example interception. However, this would also require an estimate of velocity, because we will have to predict where the subject and the interceptor will meer. The combination of location and an estimate or projection of the target dynamics may enable successful interception. Without the estimate of dynamics, or with too coarse an estimate, the location report is not of value to an interception mission . However, if the location report is to be used to cue or direct scouting, then it has value if there is a prior available on, for example, maximum speed of the target. Also, the accuracy of the location report can be significantly lower and still be of value. Assessing the value of a location report to a scouting operation is strongly terrain and target dependent, and the value may not vary smoothly with location, especially in an urban environment. Such analysis is likely to entail combinatorial exploration of potential paths between the current location of the scouting party and the potential target locations.

Priors are probability distributions which specify the potential state of a phenomenon before the sensor network begins operation, and "projections" are conditional stochastic descriptions of how the phenomenon will evolve over time. For example, the quality of information in a detection action depends on where a target is located while within the field of regard of the sensor from which we are to derive detection. A microphone used to detect tanks will emit a sample sequence, within which we must detect the signature of the tank (loud, low). We can compute beforehand where a tank of a specific class behaving in a specific way has to be located in order to be detected according to specific criteria. That locus constitutes one kind of prior on the location of that tank at an instant of detection, and is used in the calculation of quality of information delivery.

4.2. Quality of information and value of information

Informally, *quality of information* (QoI) is an expression of how well a particular body of information conveys the true state of the world, modeled at a particular level of abstraction. While different QoI metrics may be appropriate for different kinds of information delivered by different sensor networks, e.g. probability of detection vs. expected tracking error, but the same metric and QoI value are appropriate regardless of how the information supports the conduct of the mission. We consider QoI in sensor networks to have three dimensions: accuracy, latency, and trust. In this paper we focus on accuracy and latency, but also note that trust has the same *type* as the other two when modeled stochastically.

Informally, *value of information* (VoI) is the contribution that a body of information makes toward the overall success of the mission. Similar information, measured for example by exhibiting the same QoI (e.g. a NIIRS rating of 3), may have different value depending on how it supports the conduct of the mission (sufficient for foliage level estimation, insufficient for identification of guerilla activity), but also value as calculated against the ability of the imagery system to provide suitable information over the duration of the mission, and its timing characteristics both of image capture and processing facilitating timely knowledge building.

Some notion of VoI must be used when allocating constrained or even scarce sensor network resources to competing information requirements. The definition of precise VoI metrics deserves a study in its own right. The next section gives an example VoI development through stochastic modeling to illustrate the relationship between QoI and VoI. This requires the establishment of a body of information about the mission necessary for constructing the models and calculating the necessary measures. These measures generally appear in the parameters of models at higher levels of abstraction than QoI delivery. This will be seen, for example, in the expression of energy consumption states in the tracking vignette example in the next section.

4.2.1. The difference between QoI, QoS and VoI

Expressing the requirements placed on a mission in terms of required QoI can create confusion. We actually want to specify something like the *information provision performance* in the mission. Speaking in terms of performance naturally leads us to wonder if we are talking about some form of quality of service (QoS). QoS provides carefully calculated values of measures chosen as heuristic guarantors of performance, such as mean and maximum delays, or minimum bandwidth guarantees. This is formulated under the assumption that the load the system will be put under is unknown, but is comprised of a set of activities that the guarantees would *tend* to serve well. This is perhaps comparable to QoI delivery function development, in which one aims to maximize the capability of a sensing service, such that better guarantees can be specified. QoS guarantees, like QoI delivery functions are, however, calculated on a proposed specific deployment. This clearly indicates the opportunity for a feedback design loop in which the intention is to improve specific QoS parameters.

This may lead us to think that the QoS measures are now mission specific. In fact, they are not: they are deployment specific. To introduce mission specificity, we must construct service measures that satisfy command. This service is generally a cooperation between many sensing services, knowledge building and situational awareness synthesis. Modeling such a service begins with QoI delivery functions, includes aspects of all the mission elements, and ends with probability, timing and energy consumption measures for the goals of the mission. Further than this, we must drive the mission with representative mission traffic, i.e. subjects to be sensed and interacted with. The traffic (vehicles, persons, ballistic ordinance) is non-deterministic; therefore decisions about what happens “next” are encoded in stochastic models. Application of analytic stochastic modeling techniques and methodologies then enables us to go about linking the probability distributions that invoked each aspect of the mission in a manner that will enable synthesis of goal measurements.

Some of the necessary details may be behind caveats, in which case goals lower down the command or tactical levels are to be measured, with functions constructed to encode the importance of those goals. We describe a vision for the whole modeling picture, and give some example modeling constructs.

4.2.2. Use of stochastic models

There is considerable debate over the definition of quality of information, how it should be represented and what it can and does mean. We indicate here what QoI means in probabilistic modeling for sensing service design.

Firstly, however, we must be clear about why probabilistic modeling is necessary. Imaging sensors are currently ascribed QoI measures as distinct classes, such as in the NIIRS ratings. These provide a deterministic description – "this camera will enable differentiation of persons from livestock at a distance of X km". Generally, QoI is a probabilistic measure of what a sensor will supply, and therefore in and of itself motivates probabilistic mission modeling. Imaging is not regarded as probabilistic, and therefore suggests modeling in a logical style – "Object is present, therefore if it is a person, it will be identified as such by the camera, we cue an attempt to identify that person."

If we set aside any concern that there is a small probability that in fact the camera will not succeed in identification, then the action of the camera performing identification becomes a single deterministic transition which will set in motion the subsequent mission activities. The quality of information is essentially absolute in this case, or perhaps "sufficient". If however, we do include a modeling path for an unsuccessful identification, this enables us to explore the implications of an unidentified person at large, make prior estimates of the probability of this person having malicious intent, speculate as to which other parts of the mission the individual may seek to interact, and model these interactions. This opportunity clearly interacts strongly with questions of how decision making is to proceed in the command structure. A camera that has been sabotaged will also fail to identify a person, but we may know that this is the case, since no imagery will be captured. The decision makers may elect to assume that the camera will succeed, and not be sabotaged. This would simplify the modeling structure associated with the camera to a logical interaction, but our model of what is being imaged is still probabilistic, because there may or may not be a person present.

If we do allow for a failure in identification, then the transitions resulting from the imaging system are described as a probabilistic choice between success and failure. This is comparable with, for example, detection oriented sensing systems, which are ascribed a probability and latency of detection as their QoI [Zahedi 08]. Instead of connecting a sensing subject directly to the successful sensing state, this results in a probabilistic transition description of a suitable form. This can be modeled in a continuous time Markov chain model as a one or more Poisson processes to approximate the timing.

In general, then, and with an eye to supporting future sensing applications, we use probabilistic modeling. In the next section we introduce stochastic modeling, in which transitions are chosen according to a random distribution resulting in structured motion through different states of a model. In modeling deterministic imaging sensors, this connects the imaging subjects directly to their target sensing states (*e.g.* an identified person), rather than transitioning into a sub-model which describes the stochastic behaviour of a sensing service with random elements. The modeling of the impact of the sensed entity then proceeds according to its interaction with other aspects of the mission, which are not under our direct control, and therefore must be regarded as non-deterministic. Non-determinism is modeled using random variables, and this enables the construction of a stochastic model of the whole mission system, or the necessary parts to construct a model of the achievement of relevant goals.

5. Mission specific performance modeling

The quantitative analysis of a system that is so complex as to obviate detailed modeling, or with requires abstraction to a stochastic model. The Performance Analysis Process Algebra, or PEPA [Hillston 96], provides a suitable method for representing the interaction of the relevant aspects of a military activity, and was the first whose originating research group has focused on producing and fostering the development of practical tools for numerical solution of models written in it, and translated from it into more expressive languages. A demonstration of the ability to model spatiotemporal dependencies in a tactical sensing vignette [Thornley 08] is a basis for confidence that we should continue working with PEPA to build the richer structures necessary for mission modeling. To this end, we structure the illustrations we provide here in terms of cooperation between transitions between states in components of the overall model, rather than by message passing (as would have been the case in the Pi calculus or a simulation language). The representation of complex cooperations becomes cumbersome, but as we are investigating this issue in the formal event detection work associated with [Thornley 08], we continue with PEPA in this track (submission to SPIE 09), with confidence that a suitable extension or translation will emerge from our investigation of this or the event detection calculus investigations (submission to SPIE 09 modeling a simple mission with situational awareness feedback to sensing policy choices based on alert level).

We link activity in models using a *cooperation* over transitions. PEPA has two types of transition: passive and active. A passive transition will not occur unless it is included in a cooperation with one or more other transitions, at least one of which must be active. An active transition essentially drives passive transitions that cooperate with it. In PEPA, a system model is constructed of components cooperated together using the classic bow-tie operator. In this technical report, we will show this by using a common label (*e.g.* “tracking”) on cooperating transitions, and indicating which is passive (P) or active (A). The tracking transition we use in our example includes a “self-loop” passive transition that can serve to enable a dependent activity in the sensing service model, for example to cue observation by a UAV.

5.1. Example sensing vignette: track vehicles in an area of interest

Our introductory example derives from the common requirement for a sensor network to scan, cue and focus on a target. In this report we consider sufficient detail of the tracking process to illustrate essential modeling details necessary for mission-specific performance prediction.

We are to detect vehicles entering an area of interest, and track them to a specified accuracy and continuity while the subject is in that region (Figure 5). The action is described informally as a command to “Track a vehicle to within X metres’ standard error, and maintain that track for X% of the time the vehicle is in the area of interest.” This is example sketched in [Thornley 08b].

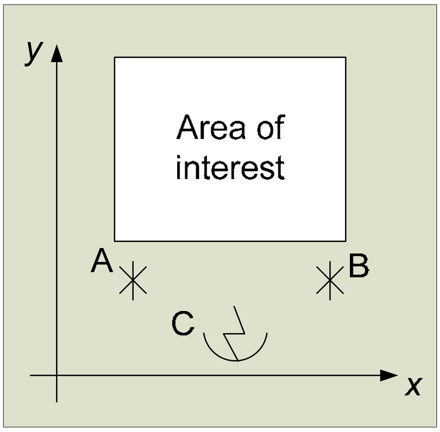

Figure 5. A sensing vignette in which an area of interest is to be monitored by two acoustic arrays (A and B) and a low-power radar. A and B are passive acoustic arrays, which we take to be either inactive, or active with the output of 3 or 6 elements being processed. The acoustic arrays provide direction of approach (DOA) estimates of an acoustic source. C is a low-power radar being used with a wide FOV to cover the area of interest, and is used to provide a range estimate.

During this tracking process, energy is consumed by the sensing, communication and processing assets which provide the service. The duration of a mission is specified, so we must model the battery usage of assets that run on batteries. The sensing service has a number of modes of operation, each of which has a predicable energy consumption rate.

5.1.1. Sensing service

We model the service in terms of its states of invariant consumption (energy, bandwidth) and output (quality, quantity) characteristics. To do this, we can specify contours in the system state space (characterized by a combination of the region in space covered, and parameterizations of the assets) where one or more modeling elements change state. For example, the tracking service enters different modes to maintain accuracy QoI. In this example, we also suggest a region in which the tracking accuracy cannot be achieved per se, but the overall mission specification may nevertheless be satisfied, since a certain percentage of track interruption is tolerated.

The scan, cue, focus process in this example comprises a passive acoustic detection service, a focusing service (cued by the detection process, and a tracking service. Acoustic detection (by either A or B with a small number of elements active) cues for initial location estimation, followed by an efficient tracking service. Details of the detection and cueing activities will be examined in detail in subsequent reports. Modes are labelled by specifying how many elements are active in each acoustic array involved, and whether the low power radar is active as follows:

CA6 = low power radar C and array A with 6 active elements operational for high resolution direction of approach (DOA)
CB6 = as CA6, B replaces A
CA3 = low power radar range and low resolution acoustic DOA
A6B6 = both acoustic arrays, each with 6 elements active
A6B3 = array A high resolution, B low

Scanning is implemented by each of the two acoustic arrays processing the output of perhaps one or two microphones for a signal which crosses some threshold of magnitude over a suitable period, or exhibits a suitable *signature*. This constitutes a detection service, whose detailed characteristics will be analysed in a subsequent report. This cues a process which serves to focus the next service appropriately. In the general case, this might involve anything from making a coarse estimate of location and causing a camera or radar to foveate, to dispatching a team to intercept. In our case, we must select the most appropriate mode of the tracking service.

If detected by A, activate A6. If no stable direction available, activate C and B6. If stable direction estimated, activate C if upper left quadrant, then rack back to A3 if close, or if other region, activate C and B6 and progress appropriately. A specification of acceptable delay in detection may require a simpler approach of activating all assets briefly, or may allow for a more careful progression of interrogation of the environment. That judgement must also be made in an analysis of the service response to subject projections.

For simplicity, we choose to divide space into contiguous, non-overlapping regions within which tracking accuracy is acceptable (Figure 6). The contours demarcating these regions are indicated in figure 6. We also include a region (8) where tracking accuracy is not feasible with the proposed service.

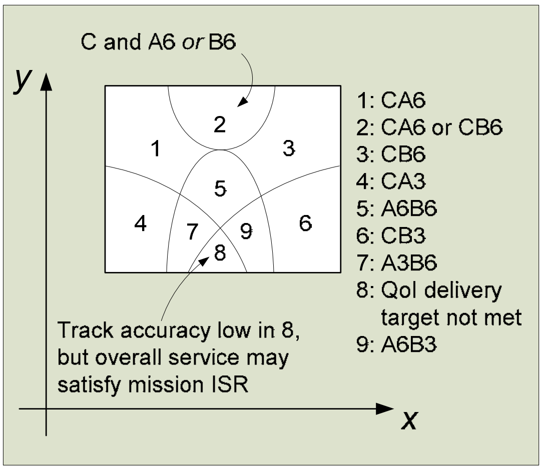

Figure 6. The area of interest has been partitioned into illustrative regions within which a specified sensing service can provide a class of tracking accuracy. The sensing service per region is shown to the right. “C” indicates that the low-power radar is in use. “A” or “B” followed by either “3” or “6” indicates that the indicated acoustic array is in use with 3 or 6 elements. In region 8, it is not possible to achieve sufficient tracking accuracy to meet the command specified minimum.

5.1.2. Subject projection

Projections of subject behaviours are necessary for analysis of sensing service performance in the context of a mission that has to deal with those behaviours. We use these projections to create the modeling states necessary to model the features which contribute to assessment of sensing service performance. We will be using sensing services to drive or enable transitions in other parts of the overall mission model. We may be more or less confident in our projections of subject behaviours, and in fact we suggest that any given solution be assessed for what subject behaviours might minimize its effectiveness.

To model spatial behaviour in an analytic state-based stochastic model, we could discretize that space in some manner. For example, our region could be split into a grid. However, we are not interested directly in modeling the location of a vehicle. Instead, we must effectively represent how effectively that vehicle is detected and tracked, specifically in relation to how we should respond to it according to policy, and how that information is to be used elsewhere. So, we do not use a grid based on the geometry of the area of interest directly, but calculate the transitions between distinct modes of the sensor network, and whether or not they satisfy aspects of the mission plan. In fact, we may also further sub-divide the state space on geometry of the sensing situation if it simplifies abstraction of the subject projections to drive transitions in the sensing service state space.

Here, we note that the tightest constraint on the sensing service – that of energy consumption – is a convenient byproduct of the abstraction approach. A given state of the sensing service will consume energy of given assets at predictable – most likely approximately constant – rates. Therefore, if we can solve for the equilibrium behaviour of the sensing service while responding to a projected subject or subjects, we may predict the energy consumption of the relevant assets, and hence the drain on a supply, or the lifetime of batteries.

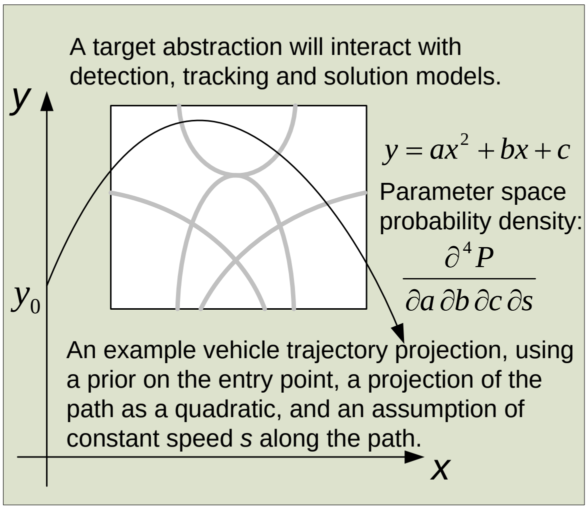

Figure 7. An illustration of target motion projection. This example represents the vehicle trajectory as a polynomial y(x) with an associated speed along the trajectory. The projection is given by specifying the joint probability density of the polynomial coefficients and the speed.

Figure 7 indicates how we relate a projection of a subject trajectory to the sensing states. From this, we generate traffic models which govern how transitions between sensing states will be driven when a subject is present, from the moment of incursion or discovery.

5.1.3. Traffic model

The projection used to generate the transition pattern and timing in the sensing abstraction pertains to a particular class of vehicle, perhaps with a specified intent (reach a location quickly, or patrol).

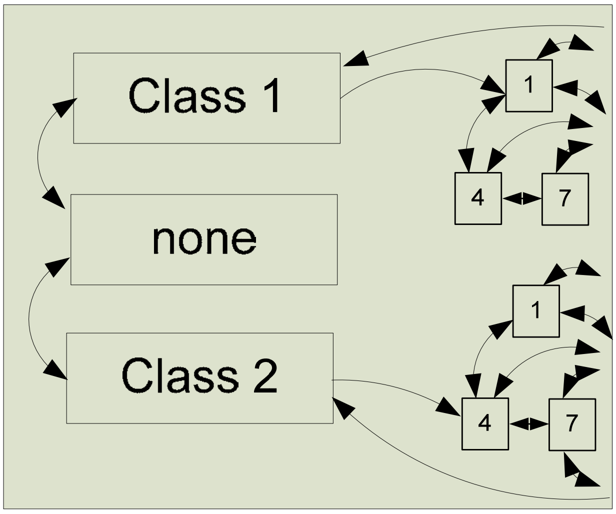

Figure 8. “TRAFFIC MODEL” The vehicle model is provided as a prior on location at detection (calculated against the detection service definition), and projections of trajectory. These are abstracted against the detailed sensor network model to produce partitions of space, each allocated a state. In each state, the target track description adheres to an invariant (*e.g.* accuracy better than +/-1m with 90% confidence, when the location is within area of interest). The vehicle abstraction includes states which guide the transition structure in the tracking trajectory abstraction.

The traffic model, illustrated in Figure 8, shows states which represents the selection of traffic to drive the mission model via the sensing service model. This example governs activity when only one subject is anticipated, selected from two possible classes. Each class has its own transition pattern and timing calculated to reflect the transition between the spatial regions that correspond to single, constant sensing modes.

In this example tracking vignette, we are to detect and track a vehicle for the purpose of enabling prosecution. This requires estimation of a number of guarantees on performance. Firstly, we must be sure the sensing service will last the mission. This could be loosely estimated using the tracking abstraction alone, on an assumption that the vehicle will generally complete a transit of the area of interest. This constitutes a pessimistic estimate, and with sufficient provision of materiel in practice, this might suffice. When competing for resources, however, more options will be successfully and correctly judged feasible if we can make closer estimates. In this example, successful prosecution before the subject completes transit of the area of interest will mean the sensing service is active for shorter periods, and hence last longer.

The traffic model therefore drives a sensing abstraction, part of which is provided by the tracking abstraction. A vehicle projection is selected by the traffic model, which in turn invokes a transition pattern in the tracking abstraction once detection is established. This traffic model therefore strongly affects measures of energy consumption, correctly reflecting the impact different traffic levels will have on reserves.

A more responsive overall mission will lead to the sensing service operating for shorter periods, but may require a mode selection policy which tends to use more energy (for example, if employing the low power radar more readily).

5.1.4. Tracking vignette abstraction

By using the subject projection to calculate representative branching probabilities and timing distributions of transitions between the sensing service modes, we create a stochastic model of tracking (

Figure 9) which enables us to calculate energy consumption and tracking interruption probability, and to connect the tracking service to those mission model components that require the tracking information.

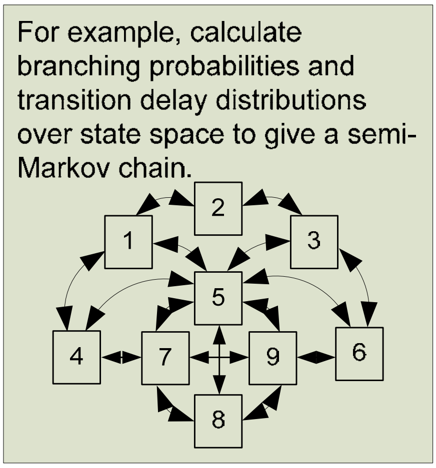

Figure 9. An indication of the transition structure between states of the sensing service. Consider state 1, which may be followed by 2, 4 or 5. We may model the system using branching probabilities pi,j and associated timing to generate a semi-Markov chain.

The state space may be further subdivided, or enabling transitions added where necessary to support the mission plan. For example, if the subject is not intercepted by a certain line in the area of interest, the plan may call for a neighbouring region to prepare for incursion of that subject. Crossing that line will be represented by a transition in the tracking model, which will be linked into the relevant response in the mission model.

5.1.5. Situational awareness modeling and feedback

Situational awareness (SA) modeling for utility estimation in the sense we describe is complementary to the situational awareness assessment in the JDL approach. In the JDL view, SA is identified and assessed according to importance and potential for impact. In the predictive models we describe, SA states identified as part of the JDL approach are tested explicitly, and their importance discovered. Our use of terms like “knowledge building” and “situational awareness” is under review, but in essence, they specify the use to which the model is to be put.

We briefly consider the simple sensing vignette in a slightly wider context to illustrate the importance of SA modeling. Consider monitoring an area of dead ground for vehicle presence, type and traffic levels. This area lies to the North of our base, which has other monitoring scenarios in play. It has been decided that opposing activity elsewhere should increase our sensitivity to local activity. Specifically, if the Situational Awareness states that there is opposing activity generally (which could either be a prior on the design, or tracked in the model for a sensor network design which operates during a changing scenario), then a single vehicle in our area is to be treated as suspicious, and hence be accurately tracked and identified as soon as possible. If the SA is quiet, then we track approximately, and intend to identify before the vehicle moves out of range. If, however, there are two vehicles in quick succession, we act as if suspicious. If a coherent track cannot be established, then the situation is treated as suspicious by default.

Does the likelihood of suspicious activity change the importance of information? Intuitively, yes. So we are modeling a larger system in which certain delivered information changes the SA, which then affects the importance of other information requirements and even spawns new information requirements.

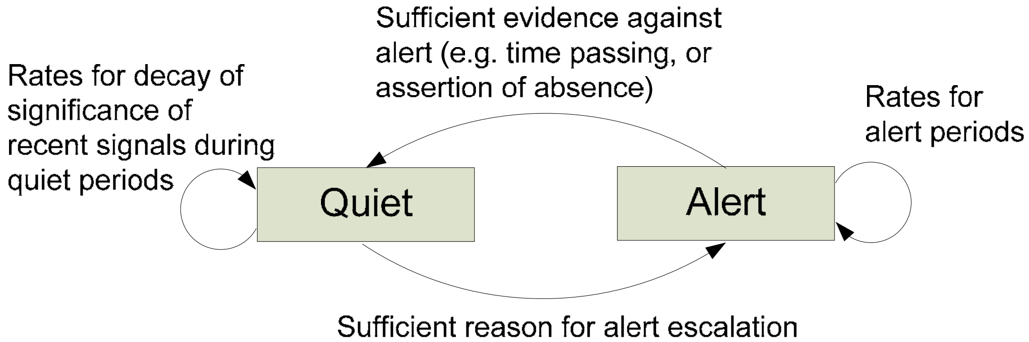

Figure 10. “SITUATIONAL AWARENESS”. Situational awareness influence on interpretation of individual reports or groups of reports. Alert rates dictate the sensor network response to recent vehicle numbers. Creating this situational awareness is a separate construct, involving recognition of combinations of the states/outputs of combinations of monitoring scenarios.

The component of situational awareness we are interested in occupies one of two states: “Quiet” or “Alert” (Figure 10). We maintain a counter of the number of vehicles seen recently, with states “Recent0”, “Recent1” or “Recent2” (Figure 11).

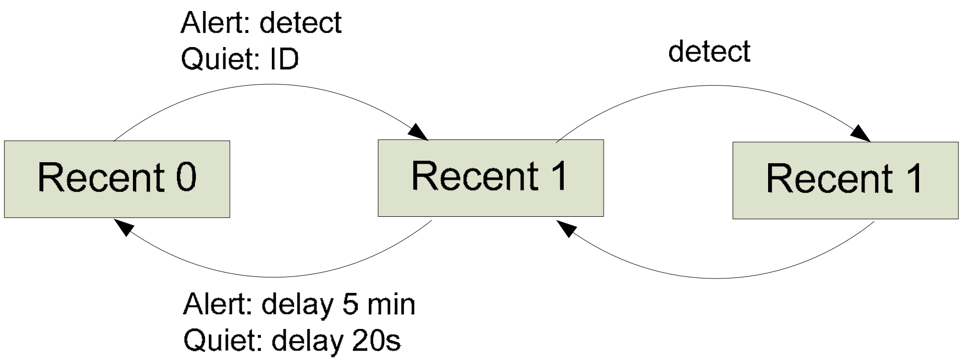

Figure 11. "HISTORY MODEL" Counting vehicles can influence on monitoring policy by driving the achievement of different alert levels. These alert levels themselves may then increase sensitivity to the number of vehicles sensed.

Sufficient evidence for transition to Alert, or return to Quiet comprises patterns in the states of the monitoring component and other scenarios. Implementing this use of interaction patterns in a stochastic model is a research issue in its own right, and is of value at a wide range of levels of abstraction from command down to signal fusion.

A tracking accuracy has been stipulated for periods of alert, and a track accuracy sufficient for recognition and ID can be derived from performance models of the relevant infrastructure. Command has specified that a vehicle be tracked for *[a number]* seconds with *[X%]* continuity. This will enable a third party to identify the vehicle *[keyhole satellite?]*.

5.1.6. Detection model

Entering the region corresponding to tracking state 1 does not mean that tracking state 1 has been established. Another process describing detection is required.

Consider a scenario in which a number of these monitoring components are operational. To create situational awareness in the form of an Alert level, we combine the outputs of a number of monitoring scenarios in what boils down to an event detection; the correlation of suspicious activities, possibly diffuse over time, sufficient to require an alert. This is the subject of on-going research.

5.1.7. Knowledge model

Keithley [Keithley 00] proposed a Knowledge Matrix for the analysis of the value of fusion approaches. This was designed with a rigid, relatively simple sequential structure of information interactions in mind. We are developing a more versatile, re-configurable component-based modeling approach, in which information flows may interact in a complex manner with the mission plan, but also with on-going information gathering tasks. The knowledge matrix therefore becomes a more general knowledge model, which takes the form in common with the other modeling aspects supported by MARS of a set of states of knowledge, with transitions between them that are driven by the models of the information sources, which in turn respond to models of the subjects in the mission.

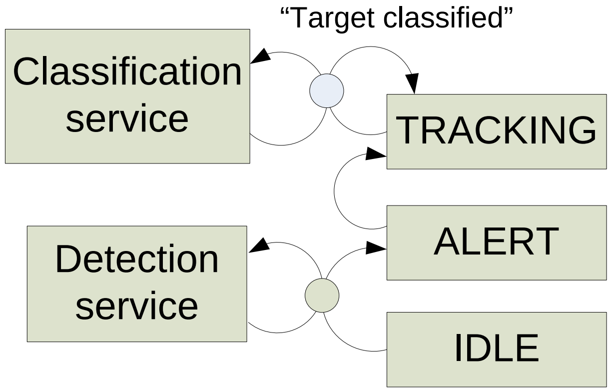

Figure 12. The output of sensor network services drives progress through the knowledge model. Figure 11 shows vehicle counting knowledge, and we illustrate building of resolution of object knowledge here. This in turn enables synthesis of situational awareness (SA). Instrumentation of the time taken to achieve certain aspects of SA will enable detailed assessment of synchronization matrices used in planning, for example.

While tracking is established, we specify a confidence that other mission activities such as identification and interception may be feasible. This is modeled by providing cooperative enabling transitions associated with such activities. A service intended to classify a vehicle may do so by focusing a camera on the vehicle, which requires location information. This is modeled as a self loop in the tracking model, and the transition associated with the required foveation step emerging from the classification service model in Figure 12.

5.1.8. Situational awareness (SA) model

Situational Awareness (SA) is a view of mission circumstances at an abstract level which supports the application of elements of the plan, and dictates how each element of the mission under our command should be acting. For example, a situational awareness "In harms way" associated with personnel motivates a suitable defensive and/or supportive response. SA is formulated from current knowledge through interpretation against the Decision Makers' prior analysis of the mission, and quantified through fusion of elements of knowledge, which are themselves fused from information available from applicable sources of intelligence. These sources include the sensor networks whose design and maintenance we are to support. Information requirements are imposed by command, and sensor networks or services are to be designed to satisfy them.

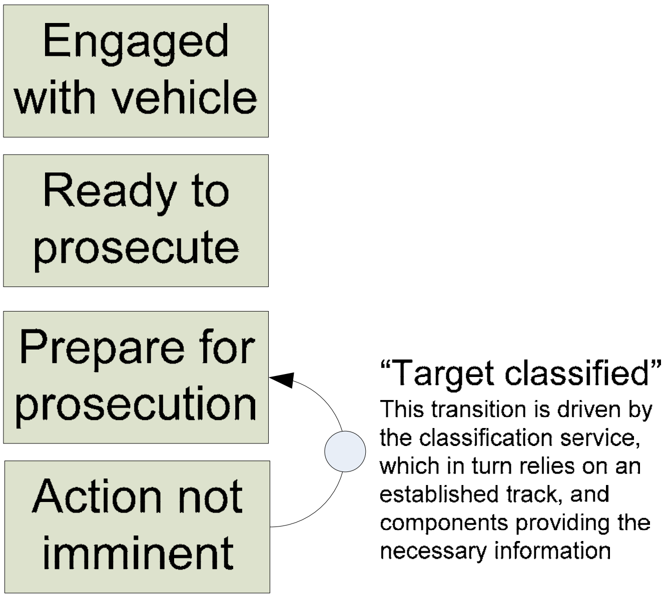

Figure 13. Aspects of situational awareness which concern the personnel involved in the mission. When action is not imminent, they may be involved in other activities, which will require that other activity to cooperate with the transition to preparing for prosecution. That preparation will take some period of time, after which the personnel model will be able to cooperate with information that motivates readiness for action. This readiness may involve exposure to danger, which we may wish to measure and optimize.

Figure 13 shows an informal set of situational awareness states, and an example transition between the awareness that action is not imminent and an awareness that preparations should be made to prosecute a potential adversary. A change in situational awareness may generally occur when a combination of a suitable context being established, and the necessary transition into a state which completes a context description that requires a change in situational awareness.

Actions may be selected or their rates modified according to situational awareness, so probabilities, timings and linkage patterns between modeling components may be selected by the situational awareness. The structure and dynamics of the mission (see for example [Allen 08, Poltrock 08]) are mimicked in sufficient detail to track the outcomes we wish to measure. State space size is controlled through abstractions which retain essential detail (e.g. [Gillies 08]).

We have illustrated the abstraction of the tracking process. Detection abstraction requires a similar analysis of a detection service QoI response (e.g. [Zahedi 08]) with combined information sources and history as required, partitioning space on changes in service modes, and on regions of acceptable detection delay and clarity.

The classification service may comprise processing of acoustic array output (e.g. [Guo 08]), or imagery capture, feature processing and database lookup relating to the images and the mission context (e.g. [Bent 08]), for which we estimate timing distributions.

5.2. From sensor network QoI delivery to mission performance

The purpose of a sensor network is to meet one or more specific information requirements, using one or more sensors whose outputs are communicated and typically processed or fused to generate the information delivered to the consumer. The information is delivered as a sequence of “reports”, each of which must be susceptible to interpretation as information relating to the mission.

Development of QoI capability predictions during normal functioning, and more recently during failure or misbehaviour modes [Szczodrak 08] enables the prediction of the way a sensing system, which may include communications and fusion processes, will respond in a specific circumstance. Thus is enables us interpret the meaning of a given report, which gives us the information quality, which is then used to guide the use of that report as either a new item of knowledge, or in fusion with other reports to generate higher level knowledge.

To test the effectiveness of a sensor network in satisfying command information requirements, its overall performance as an information production and provision service must be assessed. The issue of quality of service is in general fraught with overlapping meanings. In Project 7, which supports the needs of the ITA in assessing the information coming out of sensor networks, QoS relates to the communication network or networks which carry the sensor data and fused information. Briefly: the physics of sensing is abstracted to a measure of error in the result of a sensing act. The effect of QoS on QoI is being investigated in Task 7.3. The effect of data losses, which may be a QoS or a sensor malfunction, on QoI is part of the work in Tast 7.2. The effect of communications QoS is therefore subsumed in QoI prediction. In Task 7.1, delivered report QoI measurements are subsumed in a prediction of the extent to which command requirements are satisfied in service.

5.3. Modeling challenges

5.3.1. Modeling space: location and motion

Conditioning on projections enables us to partition space meaningfully. A state-based model is not suited to modeling spatial location per se. However, we are interested in modeling the sensing process, not the explicit motion of targets. The motion of targets is made implicit in the model by calculating an abstraction of the target motion, which is encoded in the state transition structure, probabilities and timing.

The quality of delivered information provided by a report is a measure of the “tightness” of the probability distribution associated with the report. This is subsumed in the sensing abstraction.

5.3.2. The state space explosion

The most significant barrier to exact modeling of real systems is the size of the computational load. The state space explosion is the Achilles heel of stochastic modeling, in which we must represent a number of joint states with an upper bound of the product of the number of states in the full set of interacting components. This is an upper bound, because many potential states are simply not encountered (*e.g.* trivially, active radar and a discharged battery). The solutions to this practical limitation are twofold: distributed processing and stochastic abstraction. We can consider the use of processing in the large for research, experimentation, and off-line computation of sensor network design policy (we use the term policy in this instance to mean a list of prescribed combinations of sensor network designs to be allocated in specified combinations of circumstances). For very large systems, or on-the-fly calculation of design requirement responses, we formulate approximations, in which the state space of components, or combinations of components (sub-spaces of the whole model) are replaced with approximate equivalent components with a smaller number of states, preserving the statistics necessary for reproduction of the behaviours of interest with quantified error. For example, the acoustic array model may be pared down, or the models of a group of acoustic arrays, or the complete sensing package for one are of interest may be replaced with a simpler summary. This process of replacement involves calculation of an equivalent behaviour to the necessary resolution, which may be performed on the original physical details, or on an existing abstract model.

5.3.3. Making importance estimation implicit, reducing subjectivity

Information is of value when it enables progress in a part of the mission. The value may be higher if progress can be faster, for example if the information arrives sooner, or if the availability of a key piece of information means vigilance/caution is not necessary during movement of personnel or equipment. Progress in a mission may be measured in a stochastic model by integrating time spent in a state representing maintenance of desirable states, *e.g.* percentage accurate track of a vehicle calculated as a percentage of the time the vehicle is present. Using these to judge value, however, would require knowledge of the *importance* of those states in a suitable sense.

Achievement may also be measured as a percentage of on-going goals achieved, for example, ID of opposition with less than a given delay: this may be summarized as a percentage of occasions of the presence of opposition correctly ascribed as such within an acceptable period. Again, the importance and value of these factors must be assessed by the decision makers. One way of supporting these decisions is to provide characteristics of the tradeoffs between expenditures and benefits.

However, if we can complete the quantitative modeling links between sensing service capability through mission behaviour right up to the achievement of the goals specified by command *per se*, then estimation of importance of sub-goals is made implicit in the model. Through experimentation with such complete models, we suggest that it will be possible to recognize patterns in the importance of certain types of infrastructure, and machine-learn or synthesize service utility as functions of mission and subject description parameters. This will flag unsuitable proposals earlier in the process, and hence contribute to opportunities for more agile planning, as the magnitude of the necessary search is reduced.

5.3.4. Timing accuracy

We have begun this research using continuous time Markov chain models generated from PEPA timed stochastic process algebra models. These use Poisson processes to generate transitions, which express mean timing characteristics well, maintaining orders of events, and through modern analysis techniques, enable calculation of transient responses. However, delay measurements are rendered somewhat diffuse, and tightening these will assist the automated selection of sensing services and methods as a pattern matching, rather than an abduction process. CTMC modeling enables us to synthesize utility functions given definitions of the system components, and to test alternative mission components from the ground up.

Improving the tightness of timing estimates will assist in the creation of viable libraries of sensing service descriptions by supporting the synthesis of functions which map directly from mission descriptions to normalized utility functions. The research to achieve this begins with the use of local state-space enrichments to approximate more accurate timing distributions, and analyzing the effect of these on outcomes. We mention the use of a semi-Markov model to describe the sensing abstraction. The difference between this and an enriched Markov chain state space is in essence down to which solution tools we explore, and optimization methods we develop.

6. The Mission Abstraction Requirements Structure (MARS)

The term “Mission Abstraction Requirements Structure” is chosen to give a neat acronym for an entity which lays out the information required to synthesize the models necessary for mission abstraction to a form which enables calculation of the utility, value or cost of certain components.

To design and optimize a sensor network, especially where resources are scarce, requires detailed investigation of the functioning of that sensor network as it would during the intended mission or missions. This exploration may be carried out either as simulations or analytic calculations synthesized using a stochastic description of the mission. Further, to select which missions should receive which assets when these are scarce requires careful exploration of the limits of operation of those devices.

At the beginning of the information gathering process, physical facts are transformed by a sensor into items of information provided as reports with associated estimates of their quality. These reports are used to build a state of knowledge about mission status and activity (*e.g.* there is a tank in our field of interest). From this, situational awareness is inferred (*e.g.* we should consider prosecution, and be suspicious of other activity). This situational awareness can invoke other information gathering activities (e.g. attempt to recognize the tank), and influence the selection of parameters for on-going activity (e.g. increase the sensitivity of nearby detection processes). This activity can include activating other sensing modalities, or sending personnel to reconnoiter, intercept or prosecute. Sensing, formulating situational awareness, reacting and assessing outcomes must all be modeled to enable estimation of the relevance, utility and value of information from a sensing deployment design. We begin by detailing the issues to be resolved in sensor network design, clarify the place of QoI in the design process, and how it is abstracted to mission specific performance. We propose MARS as a structured collection of the details required to carry out this modeling, which comprises an abstraction of the information pervading military actions up to a level at which it can be compared quantitatively against command ISR requirements.

6.1. Mission plans are dependency-constrained, timed state-transition systems

The synchronization matrices commonly produced to maintain the formulation and integration of sub-plans are timed transition structures with interdependencies.

A key motivating factor for the formulation of the MARS approach is to enable quantitative calculation of the compliance of information services with command requirements, and further to enable assessment of the value of information from variant candidate sensing packages/services in the context of missions that compete for the assignment of potentially scarce sensing and network equipment. This requires the surveying of the performance of the candidate sensing package or service in all the potential circumstances of the mission. These potentials are represented as priors and projections on phenomena, which result in sensor information provision, which guides progress through further information gathering. This requires a description of how knowledge is to be built, and how situation awareness is assembled to enable the satisfaction of command to be quantified. These descriptions are produced by the planning process.

The simplest aspect of a mission to describe and calculate is its potential duration, since energy consumption rates emerge early in the modeling. Testing satisfaction of duration requirements, given an analysis of the sensing service for satisfactory operation, requires a statement of intended duration (e.g. 60 days), and perhaps an indication of the potential value of being effective for longer.

6.1.1. Outcome evaluation

If we were put in a position to model the entirety of the design of a mission, being party to decision makers' assessments of the components of a mission, we might define measures of the worth of outcomes of aspects of the mission. We begin, however, by supporting sensor network design as a fulfillment of the explicit demands of the command structure at the level of decisions having been taken.

This places a step function on the value of an outcome according to fulfillment of an individual information specification. However, if there is more than one information requirement at the highest level, these step functions would require some form of weighting to enable comparison of alternate sensor network designs. Therefore we supplement the tracking of information gathering activity with calculations of the potential for tradeoff between costs and benefits.

For example, energy expenditure rate – which dictates the duration of a mission – and tracking fidelity or confidence comprise the classic tradeoff. This can be calculated at the level of QoI delivery in situations where tracking fidelity or detection confidence does not vary significantly across an area of interest (for example when large numbers of sensors are used). However, in general when using scarce resources, or attempting to set minimal duty cycles to maximize mission duration, we test the limits of functioning of the equipment, and this must be assessed in relation to the specifics of a mission.

6.2. Information requirements

Information requirements have been introduced in considerable detail in section 3. For the purpose of stochastic model, these become probabilistic constraints on acceptable quality of information in terms of accuracy, confidence and latency during the progress of relevant parts of the mission. This requires a model of the mission which generates and uses that information to enable the calculation of the necessary probabilities and latencies. This constitutes a statement of the performance of the sensor network in the mission. High level information requirements, such as a demand for the recognition of a tank within 20 seconds of its detection in an area of interest, are broken down into intermediate requirements by decision makers, with at least one proposed action for gathering it.

For example, to transition from a state of no knowledge to one in which the presence of a tank has been established, we may find that this is established by the presence of acoustic noise above a certain threshold over some period. This does not constitute an information requirement at the command level: it is an information requirement at the decision level.

6.3. Sensor QoI delivery functions

At the highest level description, a sensor gives a report type, and the delivery of that report is intended contribute to changing the state of knowledge of some element of the mission. This could be a fusion node, or a more general intermediate state governing, for example, cueing another sensor, or acting as a placeholder while other reports are delivered, which may together constitute the representation of an event, recognized as a state *pattern*.

Reports are ascribed a predicted delivered QoI. This prediction requires specification of the phenomenon signature, the signal path characteristics and the mapping from signal structure to report content to be made.

For example, an acoustic sensor (microphone) emits a sequence of samples, which may be put to a number of uses.

- In an acoustic array, the output of a number of microphones is fused through beamforming to give an angle of approach, which has the type “radians”, and the associated QoI provides information about the potential error in that measure, which has the same type.
- Alone, a microphone may be used for basic detection by continuously processing the sample sequence to identify the signature of a specific phenomenon, such as the passing of a tank (noise levels over some threshold, perhaps with some spectral analysis), a gunshot (transient event detection). Each of these may be combined in an array to provide beamforming in addition, or integrated over time if the acoustic field pattern has suitable structure (for example Kaplan’s combination of passive gunfire detection and location estimation). Detection commonly requires a defined signal to noise ratio for acceptable detection.

The general factors required for modeling the quality of information from a sensor include: field of regard (width of radar output beam, field of view of camera, directionality of microphone), spatial resolution (pixel density of camera), sampling rate (frame rate of camera, rate of sampling of microphone), latency (includes any processing and communication steps, depending on the level of abstraction), responsiveness, and an environmental point measurement to specify the sensing and communication channels.

For modeling the quality of service of a sensor, we take the QoI delivery function, and analyse it with respect to the space of input parameters to that function, and how that space is exercised within the mission. We have shown how the physical space parameters of a sensing scenario drive the definition of a number of tracking states, each of which constitutes a region within which certain factors remain constant. Those factors in the example are the equipment usage profiles, which give the energy consumption rates. In the example we use this to formulate mission endurance *given satisfactory information provision during the mission*.

QoI delivery functions provide a means for interpreting a report from a sensor network in terms of the locus of realities which may have led to that report, annotated with estimates of their probability given the provision of that report.

6.3.1. Phenomenon signatures

The signature of a phenomenon is the expression of that phenomenon in the operation of the sensor under analysis. A given phenomenon may have different signatures in variant sensors. For example, the signature of a tank in an acoustic array is a frequency spectrum in the microphone data which can be tallied between the microphones to give a direction. The radar signature of a tank comprises a mapping from the incident radiation from the radar transmission via the radar cross-section, the absorption and emission efficiencies. Both the acoustic and radar signatures vary with the angle to the principal axis of the vehicle. We have some control over the radar sensing act, in that we may select the output power, beam width and emission mode. The acoustic detection approach relies on sufficient acoustic output.

Therefore, the MARS entries covering detection of a tank must include details of the sensing geometry (location of sensors and target), parameters of the equipment, environmental conditions in the signal path (discussed below) by which the sensing act may be fully described.

Each (sensor, mode, target) tuple has a characteristic signature *e.g.*

- (radar, wide beam range, SUV) has a signature derived from the radar cross section, returning a signal which is processed to give a range estimate.
- (microphone, omnidirectional high gain, SUV) spends some fraction of the time at a higher signal amplitude than some threshold.

Essentially, we require a description of key phenomena attributes in terms of sensors deployed to observe signature (optical, thermal, acoustic, magnetic etc). The form of these is highly specific to the application, and may be different according to the analysis, fusion and approximation methods.

6.3.2. Environment

We must specify the character of potential factors that affect sensing and communication. For example, the presence of smoke may entirely obscure the image in a camera, rendering that sensing mode inoperable during some periods of a mission. Air density affects the speed of acoustic propagation, affection the QoI of acoustic sensors, such as beam forming in acoustic arrays. Signals may also be attenuated, be subject to additive noise, or suffer non-linear distortion to an extent dictated by the environment. The environmental factors modeled may include temperature, humidity, air pressure, visibility and rainfall.

Other, more statically defined effects may need to be modeled, such as those arising in areas where the topology of the area of interest may affect sensing. We may find radar or acoustic sensing is subject to multi-path scatter or occlusion, due to ground formation or usage (e.g. buildings, large equipment/vehicle storage).

In summary, we require sufficient description of the physical parameters of the environment to calculate the effect of ambient conditions and relevant aspects of the area on observation on the phenomenon signature, and hence on the QoI prediction.

6.3.3. Event signatures

Reports from sensors – or sensor network subsystems – may be further processed, and fused with other reports, to generate a higher level report, used to modify knowledge. The choice of subsuming such analyses into a state-transition model, or explicitly maintaining the fusion patterns in a state-transition model is made to balance detail of reproduction and feasibility of model analysis. If a number of sensors provide information which may be interpreted as a number of concurrent services, then the fusion process may be made explicit as a response to precalculated combinations of states to correctly reflect the correlations and orderings between fused reports delivered by those services. An event signature here is a pattern of report deliveries, and hence a pattern of transitions in the stateful model.

6.3.4. QoI delivery function retrieval or synthesis

All the above combine to enable the synthesis of QoI delivery functions. Where a common combination of phenomenon and sensing package have previously been analysed for the QoI of their report output, the function which maps from a specific set of circumstances to report and QoI annotation will be inserted directly. Where this has not been previously calculated, the details in the MARS will be used to synthesize a suitable QoI delivery function with sufficient generality for the mission or missions under analysis. Research to develop QoI delivery functions for candidate sensor networks in terms of accuracy, confidence and latency will populate the physics face of MARS, and provide library functions ready to be abstracted to mission specific performance measures. Thus MARS holds all the information required to synthesize QoI delivery functions, and may include those QoI delivery functions which do not require preconditioning with mission priors. For example, with large numbers of sensors in an isolated monitoring mission, the QoI delivered may vary relatively little over an area of interest, thus enabling mission agnostic measures of the relationships between energy consumption and information quality output. However, in situations where sensing resources are limited, and competing missions interact, we will be exploring the limits of operation of the systems under design. Inclusion of the details of the mission program and especially the anticipated behaviours of phenomena of interest at the root of the stochastic estimate synthesis will generally be necessary.

6.3.5. Trust compared to accuracy and timeliness

Trust may be a difficult concept to measure directly, but is relatively simple to implement or represent in a stochastic model. We can propose an estimate of whether an item of intelligence is true, and this might be implemented in a model as a probability of performing a transition associated with a trustworthy report with the given meaning. More generally, we may propose a set of meanings of any given report whose trustworthiness is in question, and ascribe a probability estimate to each. In the resulting stochastic model, the effect of the trust we predict we ought to place in that source of information can be estimated in terms of the set of potential outcomes, with ascribed probabilities. Trust is therefore congruent with accuracy and latency as an aspect of QoI, since each is a set or continuum of possibilities, annotated with probability functions.

6.4. Sensing abstraction

As described in our example modeling sequence in the previous section, we require a structure of transitions between states that describe the modeled mission system during a period when its elements are invariant in the measures of interest. This includes sensing service mode, data delivery rate and latency, energy consumption rate, and perhaps percentage interruption of QoI delivery over time with a specified statistical definition, which may itself be a sub-model.

The example model neglected environmental conditions for simplicity. Environment is modeled in a similar manner to traffic for selection of an appropriate class, and hence response. See [Thornley 08a] for a simplified environmental response model.

6.4.1. Phenomenon projections

Weather **–** weather can affect mission activity at a gross scale, and may therefore constitute part of situational awareness, e.g. “We must act according the restrictions on visibility and progress speed imposed by monsoon conditions”.

The importance of information about a phenomenon is strongly mission-specific. For example, an acoustic array’s performance will vary with temperature, but to a relatively small extent. But if air temperature is at body heat, infrared sensing of personnel may be compromised, so that a negative response should not be taken as an indication of absence of personnel.

Targets **–** If we are to detect and track a vehicle in a given area of interest, we must model its potential behaviours, as in the example.

A phenomenon projection will affect the selection of parameters of QoI delivery functions and lead to the synthesis of abstractions over space, time and sensing mode that give measures of interest, and drive higher level abstractions in the model.

6.5. Knowledge building

Knowledge progresses through information gathering and fusion, as described in section 3. For the purpose of modeling, we keep track of states of invariant levels of mission-pertinent knowledge Our analysis of this began with reference to Keithley's Knowledge Matrix [Keithley 00], which describes the assembly of high level knowledge according to a pre-defined sequence of resolutions. This gave the essence of the progressive nature of mission information gathering, but with insufficient detail of the meaning and interactions of each element of knowledge with respect to the mission. We therefore refer to a knowledge model, in which each element corresponds to a state in a transition sequence, with the transitions corresponding to transformations of the state of knowledge to the next higher resolution.

Each transformation from one state of knowledge to another is achieved by maintaining specified QoI levels while monitoring mission behaviours. This may take place over a number of measurements and component actions, each of which is modeled in a manner which enables calculation of the cost and benefit.

6.6. Situational awareness mapping

The body of knowledge currently held (maintained in the Knowledge Model) maps to a number of aspects of situational awareness (SA). This may use long term knowledge about a political situation or short term knowledge about the location of a Red force. Emerging patterns in the mission knowledge base lead to new situational awareness (there is now a tank in each of two locations, so raise alert level). Situational awareness elements and merging knowledge may interact to create new SA (there is opposing activity, and now there is an SUV crossing the dead terrain, so we should raise alert levels in the regions it is heading for).

We model the knowledge base as a set of components, each of which tracks transition through alternative or mutually exclusive states of knowledge of a part of the mission. These model take the general form of a column of Keithley's knowledge matrix [Keithley 00], but without the assumed structure of a single matrix. Instead, the relationships between sequences in the knowledge model are found in the links between the attainment of patterns in states of knowledge states, and transitions in situational awareness sequences. These drive progress through the mission plans.

6.7. MARS enables sensing service value assessment

We define the MARS structure as providing the information necessary to link physics to command via mission, personnel and equipment models, and aspects of command decomposition as appropriate.

MARS is not a rigid, predefined structure, with rigid parameter types tied to specific types of mission or equipment. Instead, it is a flexible structure that is synthesized from the requirements imposed by command, and the equipment available. We are beginning to formally define the synthesis of a MARS for a mission; downward from command through decision making and action assignment; and upward from physics through sensing, fusion and dissemination mechanisms.

Our remit is to support the design of sensor networks through *in silico* testing of variant designs against a mission program. MARS collates and structures all the factors required for modeling this. The output of the modeling comprises, for example, measures of the probability of achieving the information requirements, and structures over which we may calculate tradeoffs to support the decision makers' process. We may also eventually be able to detect correlations between elements of information in the mission which may enable further optimizations of the structure of the mission. The most important contribution we make with the introduction of MARS is to perform the jump between quality of information as researched over the IPP and much of the BPP, and the assessment of value (and related aspects) to specific missions, which will enable proposed sensing bundles in the sensor mission matching work to be quantitatively compared in terms of their *performance* in the mission, and hence enable suitable selections to be made based on the relative value of those performances across the competing missions. Further, MARS models mission aspects necessary for the assessment and comparison of any information and intelligence sources, and choices of action for which we can specify probabilistic models.

An important strength of this approach is that it explicitly captures the inherent uncertainty in aspects of military activities, and provides a structure within which the probabilities of not achieving alternative sub-goals are quantified.

7. Future work

We have illustrated that it is possible to abstract sensor network QoI capability to scenario-specific performance models, and indicated how these are compiled into action/mission models. This will enable us to construct the necessary formal models for representing states of the system which will then enable solution for measures which can be compared against mission information requirements. We have laid out illustrative modeling structures for a sensing, knowledge building and situational awareness sequence that demonstrates that we can link from QoI delivery to mission-specific performance using cooperating stochastic state transition models.

MARS is a combination of: the information requirements, states of knowledge to attain, command specified methods for transition between these, the information required to synthesize mathematical models of those states and transitions (which includes the physics of sensing/other collection method in time and space, any structures such as cueing &c., and crucially priors on target behaviours in the mission, and descriptions of the *importance* of each knowledge state (insofar as this is available), and these will enable the calculation of the value of the output of a sensor deployment in a mission by solving for the equilibrium and transient (time-dependent) behaviours of all or part of the mission model as necessary,

Command requirements and decision maker's action specifications are the motivation for this research. In future work we will select or collaboratively define a suitable query formalism for posing information requirements. This may resemble a temporal logic, but must comprise a representation of intended causality. In the present phase of research, we use expert knowledge of the information types in our research team, which will be formalized in on-going collaborative research. We work with mathematical formulations which map directly to, and will eventually be automatically derived from, established military command methodology.

The theoretical modeling work proceeds with fleshing out the example, enriching the mission context model, alignment with realistic planning structures (these provide the structure over which we drive progress). The process algebraic modeling proceeds with encoding these models in PEPA, formulating the complex interaction patterns with a goal of identifying common structures and requirements and formalizing these, introducing them to the language as appropriate. Some extensions to PEPA are necessary for this, and the development of these is under way. Initial numerical results will be produced using the standard Markov chain translation of PEPA models, to be enhanced by research into the creation of better timing approximations through state-space enrichment or migration to semi-Markov or general models as appropriate

8. Summary

In summary, to date, we have provided an example information about, and models of, mission behaviours which describe the transformations and modeling required to estimate sensing service performance against command requirements. The example modeling elements we have provided illustrate the modeling opportunities, challenges and research directions we will pursue to provide the functionality required to quantitatively support sensor network and service design, deployment, operation, update, post-hoc analysis and procurement.

This work is producing methods for supporting sensor network and sensing service design, and also broadly applicable and valuable computer science results in formal modeling methods and numerical analysis of stochastically timed models.

9. References

[Allen 08] J. A. Allen, D. Mott, A. Bahrami, J. Yuan, C. Giammanco and J. Patel, “A Framework for Supporting Human Military Planning” , ACITA 2008

[Bedworth 99] Bedworth M and O’Brian J. The Omnibus model: a new model for data fusion. Proc 2nd International Conference on Information Fusion (1999).

[Bent 08] G. Bent, P. Dantressangle, D. Vyvyan, A. Mowshowitz and V. Mitsou, “A Dynamic Distributed Federated Database” , ACITA 2008

[Bisnik 06] Nabhendra Bisnik, Alhussein Abouzeid, Volkan Isler, “Stochastic Event Capture Using Mobile Sensors Subject to a Quality Metric”MobiCom’06, September 23–26, 2006, Los Angeles, California, USA

[Bergamaschi 08] Flávio Bergamaschi, Dave Conway-Jones, Christopher Gibson, and Andy Stanford-Clark, IBM UK, Dinesh Verma, Seraphin Calo, IBM US “Policy Enabled ITA Sensor Fabric” in proceedings Policy 2008, 2-4 June 2008 in Palisades, New York, USA.

[Bosse 07] E Bosse, J Roy & S Wark, Concepts, models & tools for information Fusion. Artech House 2007.ISBN 13 978-1-59693-081-0.

[Endsley 95] Endsley MR. Towards a theory of situation awareness in dynamic systems. Human factors Journal. Vol 37 No 1 Pg 32-64 March 1995.

[Eswaran 08] S. Eswaran, A. Misra and T. La Porta, “Utility-Based Adaptation in Mission-oriented Wireless Sensor Networks” , ACITA 2008

[Funk 00] H.B.Funk, C.A.Miller, C.L.Johnson, and J.P.Richardson, “Applying intent-sensitive policy to automated resource allocation: Command, control, and, most importantly, communication,” in Proc. The Fifth Annual Conference on Human Interaction with Complex Systems (HICS 2000).

[Gillies 08] D. Gillies, D. J. Thornley and C. Bisdikian, “Using Probabilistic Inference to Estimate the Quality of Information in Military Sensor Networks” , ACITA 2008

[Gomez 07] Gomez, Preece & de Mel, Towards Semantic Matchmaking in Sensor-Mission Assignment: Analysis of the Missions and Means Framework. Version 1.3. Technical Report Task 2, Project 8, Technical Area 3. 23 March 2007.

[Guo 08] B. Guo, M. S. Nixon and T. Damarla, “Acoustic Information Fusion for Ground Vehicle Classification” , ACITA 2008

[Hall 01] Hall D and Llinas J. Handbook of Multisensor Data Fusion. CRC Press, ISBN 0-8493-2379-7 (2001).

[Hillston 96] J. Hillston, A Compositional Approach to Performance Modeling, Cambridge University Press, 1996.

[Johnson 08] M. Johnson, H. Rowaihy, D. Pizzocaro, A. Bar-Noy, S. Chalmers, T. La Porta and A. Preece, “Frugal Sensor Assignment” , ACITA 2008

[Kaplan 08] L. M. Kaplan, T. Damarla, and T. Pham, “QoI for passive acoustic gunfire localization,” in First IEEE Workshop on Quality of Information for Sensor Networks (QOISN’08).

[Keithley 00] H. Keithley, “Multi-INT Fusion Performance Study,” Joint C4ISR Decision Support Center, DCS-00-02, 2000.

[Poltrock 08] S. E. Poltrock, M. Handel, H. Bowyer and P. Waggett, “A Dynamic Model of Mission Context”, ACITA 2008

[Pizzocaro 08] D. Pizzocaro, M. P. Johnson, H. Rowaihy, S. Chalmers, A. D. Preece, A. Bar-Noy and T. La Porta, “Maximizing Utility of Sensor-Mission Assignment with Uncertain Demands”, ACITA 2008

[Preece 08] A. Preece et al, “Reasoning and resource allocation for sensor-mission assignment in a coalition context,” in sub. to Milcom 08, July 2008.

[Sheehan 03] J. H. Sheehan, P. H. Deitz, B. E. Bray, B. A. Harris, and A. B. H. Wong, The Military Missions and Means Framework, in Proc. Interservice/Industry Training, Simulation & Education Conference, pp 655-663, 2003.

[Steinberg 99] Steinberg A, Bowman C and White F. Revisions to the JDL Data Fusion Model. Proc of SPIE AeroSense, Orlando Florida, USA 430 (1999).

[Steinberg 08] Alan N. Steinberg & Galina Rogova. Situation and Context in Data Fusion and Natural Language Understanding. In proceedings of Information Fusion 2008.

[Szczodrak 08] M. Szczodrak, S. Zahedi, P. Ji, D. Mylaraswamy, M. Srivastava and R. Young, "Simulation Framework for QoI Characterization of Sensor Networks in the Presence of Faults" , ACITA 2008

[Thornley 08a] David J. Thornley, Chatschik Bisdikian, Duncan F. Gillies, "Using stochastic process algebra models to estimate the quality of information in military sensor networks", Modeling and Simulation for Military Operations III. Edited by Trevisani, Dawn A. Proceedings of the SPIE, Volume 6965, pp. 696507-696507-10 (2008)

[Thornley 08b] David J. Thornley, Robert I. Young, James P. Richardson "From mission specification to quality of information measures – closing the loop in military sensor", ACITA 2008

[Verma 08] D. C. Verma, C. W. Wu, T. Brown, A. Bar-Noy, S. Shamoun and M. Nixon, "Location Dependent Heuristics for Sensor Coverage Planning" , ACITA 2008

[Zahedi 08] S. Zahedi, M. B. Srivastava, and C. Bisdikian, "Framework for quality of information analysis for detection-oriented sensor network deployments," in 2nd Annual Conference of ITA (ACITA'08), Sept. 2008.

End of references.