Ontology-Based Inference with Inferlets

Abstract

We propose to develop a massive ontology and use it as a framework for class-level inference for improved situation awareness. Our proposal is to conduct research of concepts and preliminary experiments about which we have written and presented conference papers at academic and DoD sensor and information fusion conferences. If successful, the ontology-based approach will leverage COTS database technologies and DARPA, ONR, AFRL, and other fusion, inference, and cognition technologies. It will enable massive fusion inference networks for weak evidence accumulation and long indirect inferencing to improve situation awareness. Our approach is simple and elegant, yet rigorous and comprehensive.

1. Background

The operational improvement potential for the Navy is significant for multiple transformational pillars, particularly Sea Shield, Sea Strike, and FORCEnet. If successful, the technology we propose will provide warfighters with more of the information they say they need, at the strategic, operational, and tactical levels. For example, for intelligence analysis, massive ontologic inference will transform evidence in the form of various pieces of input data into long and indirect accumulations of weak decision information that will signal and aid human analysis. In tactical warfare situations, greater information detail at a greater level of precision and accuracy, accompanied by accurate uncertainty estimates, will enable better resource pre-positioning and engagement optimization.

Currently, the ability to effectively achieve synergy among the research community's fusion, inference, and cognition technologies is limited because each project develops its own implicit ontology and typically has little to no capacity for standardized fusion algorithm interoperation. Limited algorithm interoperation and the difficulty in developing large, authoritative, logically coherent, and robust ontologies are major obstacles to large-scale and advanced levels of fusion. The approach we propose herein could overcome these obstacles by employing an authoritative and robust ontologic model for situation awareness that is responsive to warfighter information requirements. Further, utilizing the techniques we propose allows for autonomous but coordinated inference and fusion algorithms developed by multiple organizations.

The basic effort we propose is a modest effort to continue from our IR&D using a COTS real-time database for Multi-Sensor Integration. In the basic effort, we will conduct a proof-of-principle that the ontology implemented in the database can be a framework for autonomous yet coordinated "inferlets" (inter-object inference algorithms of many types, Bayesian, Kalman, etc.) In the basic effort we will operate against simulated inputs. If the proof-of-principle is positive and option year 1 is awarded, the next step will be to test against live recorded data and to expand the inference net for massiveness to, (1) validate the proof-of-principle's results with live rather than simulated data, and (2) prove the next principle, that weak evidence accumulation from long indirect inferences can provide meaningful and accurate results. If option year 1 has positive results, we propose an option year 2 to operate against actual live data as part of Sea Trial and to explore further integration and transition alternatives. An initial alternative we are contemplating is to utilize the system as an adjunct to in-place C4I systems such as GCCS. Incorporation into a next generation GCCS could be explored as opportunities for such were presented.

We are confident we can conduct this research and that it will be beneficial. Silver Bullet specializes in all the sciences required: ontology modeling, abstracted databases, fusion, inference, state estimation, sensors, intelligence, and operational information requirements. We have participated in OSD studies into warfighter information requirements, advanced fusion alternatives measures of performance, and knowledge-assisted fusion. Our personnel spearheaded knowledge-assisted fusion and Bayesian Networks for large-scale fusion in Navy combat systems and electronic warfare systems. The academic and DoD Government and industry R&D communities have favorably received our concept and preliminary experimentation papers.

2. Concept of Operation for the Navy

Complete and accurate situation awareness is important in almost all Naval operations. It is particularly critical for transformational capabilities such as Sea Shield and Sea Strike. For example, Sea Shield missile defense requires awareness of the potential threat situation and own-force status and activities for pre-positioning during the pre-engagement phase and for optimal resource (interceptors, sensors) allocation during the engagement phases. Situation awareness and related capabilities also appear in the FORCEnet capabilities taxonomy. Situation awareness applies to all levels of warfare, strategic, operational, tactical.

Situation awareness requires the integration of disparate sensors and sources of information and the transformation this information into actionable knowledge. Sensors and sources can range from high data rate radar and images of varying wavelengths and fields-of-view to processed reports and databases. It must be transformed to mission information about battlefield objects of varying level of composition, their activities, status, and intent, again at varying levels of detail and precision. Objects of interest can range from individual objects to formations and organizations, infrastructure, and political. In this infoscape, key knowledge is often difficult to perceive, due to ambiguities and weak evidence.

Table 1. Types of Information

Kinematics	Location, velocity, and trajectory (past and predicted), from detection to accuracy sufficient for PGMs
Identification	Broad type to specific unit and with varying certainty
Activity	General to specific plan and with varying certainty
Status	General to specific and with varying certainty
Intent	General to specific and with varying certainty

A recent DoD study [1][2] provides insight into the elements of situation awareness. The study exhaustively analyzed warfighter Essential Elements of Information (EEI) in order to quantitatively measure the benefits of improved multi-intelligence fusion. In this study, an information requirements model was developed to answer the question, *"What are the information needs of the* *warfighters that might be improved by advanced fusion architectures?"* Thousands of EEIs [3] were analyzed and categorized as to the required information. The relevant object types were categorized into object taxonomies. The high-level object types are shown in Table 2 along with summary narratives of the subtype objects. The information needed about those objects categorized into the information types shown at a high-level in Table 1. The information types had five levels of decomposition and constituted what in the study was called a "knowledge matrix". The object types with their levels of composition (Table 2) and the information types with their levels of precision and accuracy (Table 1) provide a spanning set of the situation awareness infospace.

Table 2. Object Types

Platforms and Facilities	Ships, aircraft, missiles, vehicles, SOF units, SAM sites, TELs, etc. from Company level up to Corps level.
Infrastructure	Communications networks, electrical networks/grids, transportation networks, etc.
Political	National organization, intent, internal conflicts, economic triggers and indicators, etc.

Transforming evidence (e.g., sensor or intelligence reports) into the types of actionable information described in warfighter EEI's requires massive inference networks that can compound and accumulate weak causalities, properties, and inter-object relationships over time and information space and that copes with uncertainty in a way that allows action but maintains awareness of alternatives. Inference is the general means by which information or evidence is transformed into actionable conclusions or decision information. For situation awareness, the inference processes are called "Level 2 Fusion" and "Level 3 Fusion" in the Joint Directors of Laboratories (JDL) fusion models [4]. Situation refinement (level 2) and prediction (level 3) involve inference and fusion techniques that have ways to cope with the inherent uncertainty in the information using a blend of artificial intelligence and statistical sciences. Because of the rigorous mathematical foundation of these inference techniques, the inherent uncertainties can be accurately estimated and used in decision logic and resource optimization. For example, a missile defense resource optimization algorithm (e.g., an auction algorithm) requires accurate uncertainty estimates of the target complex's kinematics (current and predicted) and classification alternatives (e.g., confusion matrix with uncertainty values) to determine how best to allocate and schedule radar dwells and intercept windows. The quantitative improvement in warfighter information and battle outcomes from improved level 2 and level 3 fusion was reported as part of the DoD study described above [1][2]

Progress in deploying large-scale fusion and mathematically rigorous inference systems has been slow in recent years despite substantial algorithmic developments in the fusion community. One reason is that there has not been a way to address large-scale evidence accumulation in a tractable manner that allows modular and collaborative evolution of fusion algorithms. Information and data modeling techniques have become quite mature over the past 20 years so that it is now possible to model a large-scale information domain tractably. Indeed, DoD has developed a very large model that can account for most of the EEIs cataloged in [1]. By extending information modeling constructs to semantic and inference nets, it is possible to use these information models as a framework for large-scale level 2 and level 3 fusion necessary for situation awareness. The framework provides a class-level ontology in which independently developed and evolving inference algorithms, what we call “inferlets”, can interoperate. The ontology provides the rules that allow the inferlets to be independent, yet coordinated. If successful, this framework could leverage other on-going developments in fusion, inference, and cognition such as DARPA’s Cognitive Information Processing Technology.

Initially, the end-state system resulting from this research could work alongside current C4I systems such as GCCS as an adjunct that provides warfighters with value-added situation awareness, that is, deeper analysis of patterns in the information. Transition into a next generation GCCS or other on-line C4I systems in support of FNC’s would be accomplished as opportunities present.

3. Operational Utility

The operational utility of improved situation awareness has three components: quantity (including detail and precision), accuracy, and relevancy. Inference increases the quantity of information by inferring additional information based upon a given set of information. The increase arises from the application of inference intelligence to given information so it is not just the same information formatted differently; it is genuinely increased information, by standard measures of information utility. We notice this in our everyday lives, that through the application of prior knowledge, a given set of sensor stimuli can be transformed into a large and diverse amount of information. So in order to measure the operational utility of the inference net, it is necessary to measure both the improvement of the intrinsic and end-use information: The following seven measures of operational utility are proposed.

a. Objects and instances inferred. Inference will result in objects/instances being inferred, e.g., the suspected organization of a vehicle. The count will be mitigated by the power of the inference, that is, the uncertainty so that weak inferences are given credit but in proportion to their strength. (Accuracy of the uncertainties will be measured also, see below.) The formula is:

$$\sum_{i}^{ontology} Property_i^{inferred} \cdot Mass_{property_i}^{90\%} \cdot MassAccuracy_{property_i}$$

Equation 1

where mass is computed in the usual way for numeric properties (e.g., speed) and for non-numeric properties (e.g., identity) using confusion matrices such as those shown in Figure 1 to define a measure on the space.

	H	S	N	AF	F
H	1.000	0.700	0.300	0.100	0.050
S		1.000	0.429	0.143	0.071
N			1.000	0.333	0.167
AF				1.000	0.500
F					1.000

Identity:
Hostile
Suspect
Neutral
Assumed Friend
Friend

	Sp	A	S	Lm	Lf	Sb
Sp	1.000	0.300	0.050	0.050	0.050	0.040
A		1.000	0.167	0.167	0.167	0.133
S			1.000	0.040	0.060	0.300
Lm				1.000	0.667	0.060
Lf					1.000	0.070
Sb						1.000

Environment / Cateogry:
Space
Air
Surface
Land-Mobile
Land-Fixed
Subsurface

Figure 1. Example Confusion Matrices with Distances Expressed as Relative Percentages

b. Inference distance. This will measure the contribution of the ontologic net in terms of its ability to support long-run inference, from evidence to intermediate hypotheses to final activations. Its measure is:

$$\sum_{i}^{onto \log y} InferenceRangeAccumulator_{property_i} MassAccuracy_{property_i}$$

Equation 2

where the InferenceRangeAccumulator is a node-by-node accumulation of a passed LinkCounter times the probability mass change that occurs in the iteration.

c. Decision information inferred. As in, this will measure the degree to which the spanning set (called "knowledge matrix" in [1]) is improved by the inference network. The inference network will be compared to a binary network that does not accumulate evidence and that concludes only when there is a preponderance of instantaneous evidence. That is, the control will be equivalent to a memoryless logical deduction system. Its measurement is not shown here in the interest of space but is similar to the "knowledge matrix satisfaction algorithm" in [1].

d. Object properties relative accuracy improvement. This measures the estimation refinement or accuracy improvement for properties of the objects such as kinematics, classification, and activity. The control is a "source selection" system such as typical of Navy Combat Direction Systems in which the "best" source properties are selected for the object (track). The measure is a comparison of the correct probability mass.

e. Object accuracy improvement. This measures how well the network estimates the objects in the environment. Although there are many detail types of measures in this area, they are fundamentally, (1) is every object in the true environment represented in the system; (2) are there extras; and (3) are there any contaminated objects. However, (3) is actually covered by the "object properties accuracy" measure since from a user or black-box point of view, it doesn't really matter what the "contamination" is as long as the property estimates are accurate (although one would expect contamination to result in loss of accuracy.) The control is an object for each associated report stream. In the interest of space, the measures are not shown here but are similar to those in [5].

This can be measured only in environments where truth is knowable such as simulations or range data.

f. Object properties absolute accuracy. These are like the "Object properties relative accuracy improvement" but the comparison is to known truth rather than to the source selector. The measures for numeric properties use a normalized distance (Mahalanobis) while the measures of accuracy for non-numeric attributes use the confusion matrices as a definition of distance from truth. This can be measured only in environments where truth is knowable such as simulations or range data.

g. Uncertainty estimation absolute accuracy. Uncertainty is important in higher levels of fusion, inference, and situation awareness, as discussed extensively in this proposal. Consequently, it is important to measure the accuracy of the uncertainty estimates from the inference network. The measure computes the network's error density over time as compared to the network's estimate. This can be measured only in environments where truth is knowable such as simulations or range data.

4. Technical Approach

The technical approach is to research massive inference using a class-based integrated ontology as the connectionist structure in which what we call 'inferlets' can then operate. In the following we will describe:

- A way to develop a robust massive ontology for situation awareness
- The means by which the ontology can provide the basis for inference
- Experiments indicating the feasibility of the ontology in fusion processes
- Proposed deliverables

4.1 Authoritative Basis for Massive Ontology for Situation Awareness

Fully-attributed Entity-Relationship (E-R) models, particularly ones with class-hierarchies, can form a basis for a class-level semantic net that can, in turn, be the basis for an inference net with distributed cooperating inference engines. Such an employment allows the use of large E-R models developed by DoD that cover all form of defense activities and objects. It brings the mathematical rigor of E-R modeling to bear along with the proven ability to collaborate and converge to large-scale ontologic concurrence. Both of these features are necessary if the ontology is to support automated and massive inference.

In paragraph 2, herein, we described a spanning set of information for situation awareness information based upon an exhaustive analysis of warfighter Essential Elements of Information (EEI). All of the entities shown in Table 2 are called, "object classes" in an object-oriented design, but in the more general Entity-Relationship (E-R) modeling might be called abstracted or generalized entities. There exists a robust, comprehensive, and rigorous model of much of this information domain, DoD's Command and Control (C2) Core data model. A high-level overview is shown in Figure 2. C2 Core employs powerful object class hierarchies to model very much the same types of objects, information, and levels as the EEI study.

This type of E-R model is equivalent to a class-level ontology. Class-level ontologies have the same sorts of benefits as class-object hierarchies – simplification of the model, permanence and stability of upper-class properties, consistency of object properties across like-classed objects. In class-level ontologies, it isn't necessary to explicate every object property and interrelationship because the class often does most the work. Only those exceptions to the class need to be modeled in the ontology. This simplifies the ontology and makes it easier to develop, validate, and maintain. For instance, ACTION can be cover all types of actions, from destroying a target to submitting a budget, via the TYPE CODE and the ACTION RESOURCE's and ACTION OBJECTIVE's, either of which can be the generalized objects FACILITY, PERSON, ORGANIZATION, MATERIEL, and FEATURE, all of which are superclasses. Class-level ontologies also are more stable than object-level ontologies because the upper-class properties tend to be fairly stable. For example, the C2 Core's ACTION, ACTION RESOURCE, and ACTION OBJECT can be expected to be fairly stable since they mirror the verb, subject, and object in many languages. Also, the class hierarchy prevents unintended inconsistencies between objects of the same class since they share the same class hierarchy. Since the upper class properties are grounded in well-reasoned invariants, they are not only more stable and consistent; they tend to have wide concurrence and validation.

The types (xxxTYPE in the figure) carry the general class behavior for the instances (xxxINSTANCES in the figure). Subclassing types and instances under the general concept

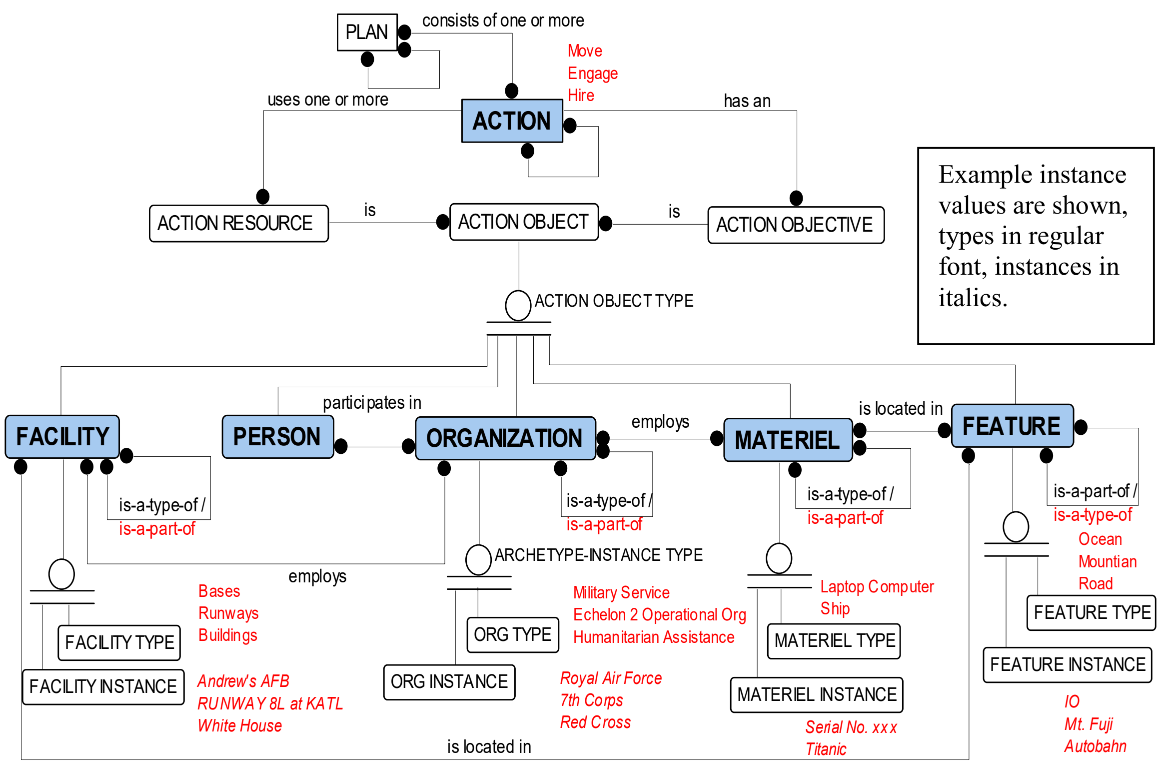

Figure 2. Major Entities in Command and Control Core Data Model

mirrors ordinary human discourse in which terms can interchangeably refer to types and instances. For example, I might speak of how fast a car can go, talking to the general type of car or a specific car. *In inference systems, having a general concept as a superclass of the type and instance allows transition from classification hypotheses at the type level to hypotheses at the* *actual real object level.*

Types generally have many levels of taxonomic typing. For example, targets may be categorized as Air, Surface, Subsurface, Ground, Space with Air subcategorizing as Bomber, Fighter, Transport, etc. and Fighter further subcategorizing as F-15, F-16, Viggen, etc. with yet another subcategory for F-16 such as F-16A, F-16B, etc. The purpose of typing is to allow for property inheritance. *In inference systems, this feature enables recognition of new instances.*

There are many benefits to this type of rigorous and mathematical ontology modeling. For more details, see [9].

4.2 Implementation of the Ontology in a COTS Database Management System (DBMS)

Expressing an ontology model in E-R model also has an important practical benefit, that is, they can be directly implemented in COTS Data Base Management Systems (DBMS) such as Oracle and SQL Server. Nearly all formal database design is now done using this technique, thus there is a large body of developed models from which to draw. DBMS's provide robust built-in features that can improve many fusion system performance attributes such as:

- System Reliability. COTS DBMS' have very high reliability so system crashes and hangs caused by data corruption can be reduced
- Presentation Accuracy. Data integrity problems in fusion systems are common and sometimes crippling to operations. These problems can include systems crashing, hanging, outputting incorrect data, and/or confusing operators. DBMS' enable improvements in data integrity compared to flat, unrelated, and unmanaged data structures in several ways including "normalization" which prevents duplicate values and referential integrity and data validation rules.
- Object Fidelity. The objects of fusion can be treated with more accurate and complete object semantics. Employing this proven powerful technique in fusion system database design would accrue the same benefits it does for all the other systems for which it is applied: more complete and correct expression of the domains objects, their properties, and their inter-relationships
- Inference Execution Accuracy. Inference can operate with greater completeness and faithful object inter-relationships.
- Backtracking and Auditing. Input, state, and output history can be maintained to a very large degree.
- Inference Design Accuracy. Inference design can be made more logically coherent and complete.
- Tractable Data Management. Reference, track file, sensor file, etc. can be managed in a systematic manner.

DBMS features that enable these types of benefits are described in more detail in [10].

4.3 Using the Ontologic Connectionist Structure as a Basis for Inference

Data models are expressive and can have great fidelity to the ontology of the modeled domain. However, they model all beliefs in binary -- true or not true.

There are four types of augmentations to models like C2 Core that are necessary to express uncertainty. The following will generalize from prior work to cover more types of uncertainty representation and to allow for all forms of inference, including estimation techniques. These are described in detail in [9]; one type is described herein for convenience. Other work dealing with uncertainty in databases includes attribute conflict resolution in heterogeneous databases [6]. In this work the Dempster-Shafer evidential technique was investigated as a means to resolve conflicts in values between attributes of the same entity in different databases. Lee's [7] work in this area also showed how uncertain updates could be encoded in a relational model.

The simplest case is that of a model relating two objects via an associative entity. This model can support many-many relations and/or information about the relation that is not specific to the individual objects, as in the left side of Figure 3. This associative entity provides a many-to-many relationship between PERSON and ORGANIZATION, allowing a person to be a part of many organizations, conversely, allowing an organization to have many people. Common-sensical but powerful, it is nevertheless binary -- you're either a part or not; no maybe's. To represent uncertainty in this case, merely add a confidence value to the associative entity, as shown in right side of Figure 3. An instance in the associative entity represents each hypothesis.

It is fairly easy to "attach" a measurement model to the situation awareness model using DoD models because they all interlock on common entities. An example of such interlocking is shown in Figure 4. The measurement model is a transmitter model, appropriate for electronics signals measurements. The tie-in to C2 Core is the entity MATERIEL-ITEM. This model bears some resemblance to Electronic Warfare classification databases but has all the benefits of formal modeling and DBMS implementability we briefly described in paragraphs 4.1 and 4.2. Most importantly, it can be used as a framework for autonomous yet coordinated inferlets, as described in the following paragraph.

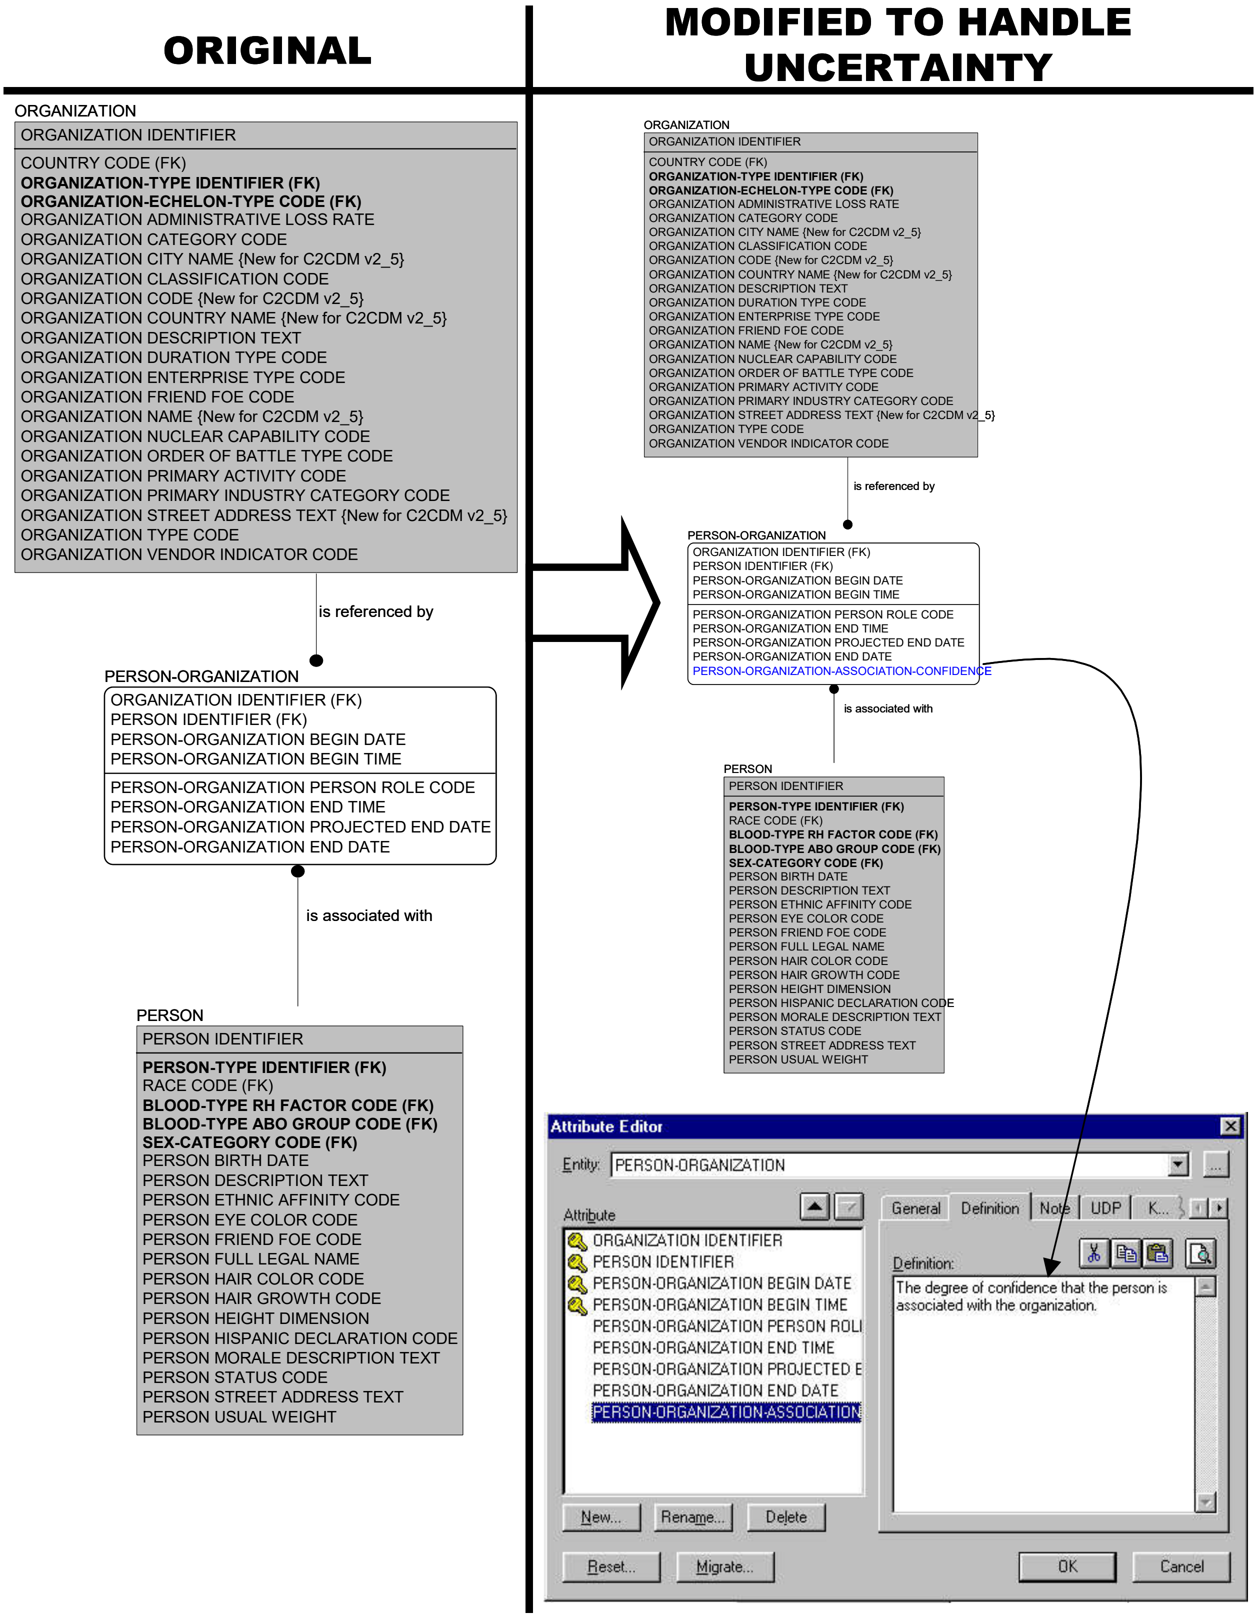

Figure 3. Example Associative Entity Modified to Handle Uncertainty

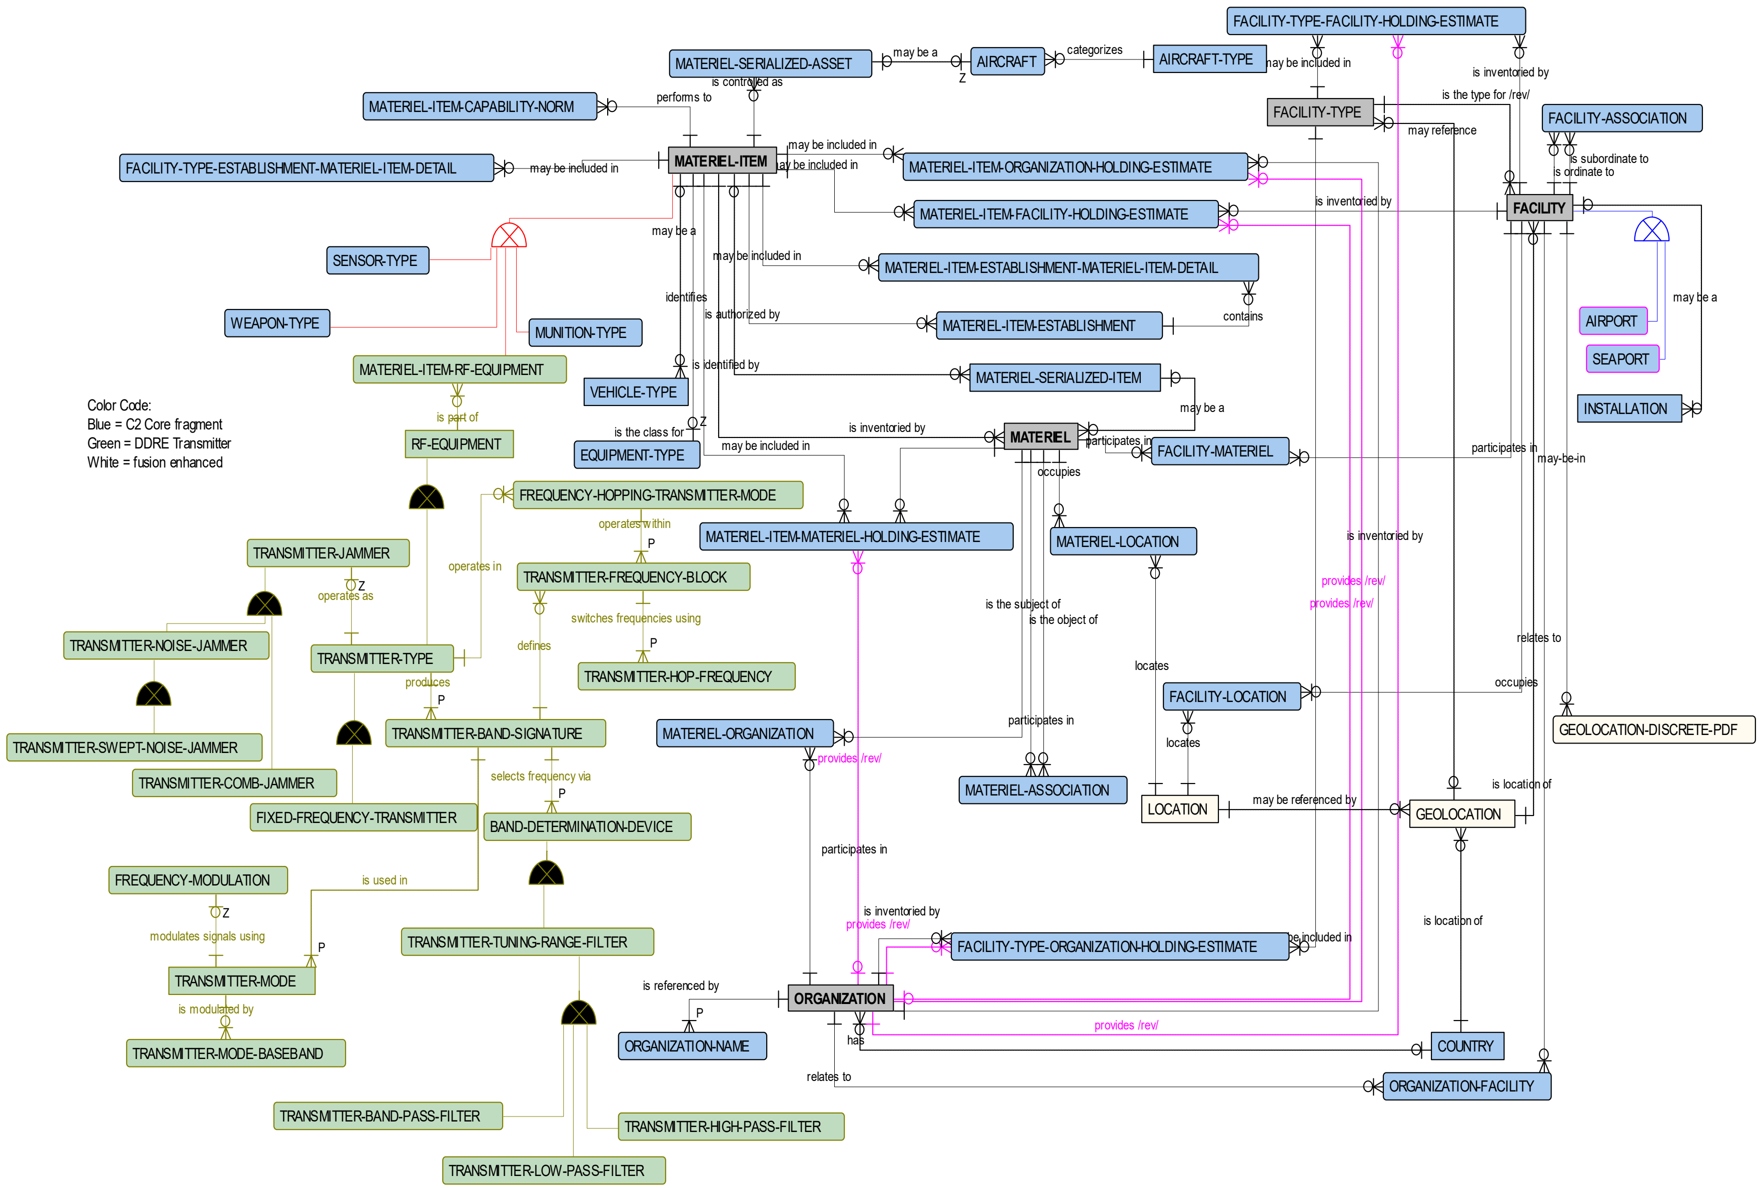

Figure 4. Class-Level Ontology Model Relating Sensor Measurements to Order-of-Battle and Force Structure

4.4 Inferlets

We now address computation issues in executing the instrumented model so that it becomes an inference network. To see how a class-level ontology like the one shown in Figure 4, imagine an update to the frequency and modulation type tables, either in the form of a new instance or an update to an existing one. If new, association hypotheses to prior measurements would be formulated, along with activations to transmitter type, and perhaps even transmitter instance hypotheses. These would then spread to the host equipments, platforms, and facilities, and then to the possible organizations, missions, and actions.

The structure provides a flexible structure for storing various inference algorithms, in a manner analogous to methods in OO databases; in a relational DBMS, they would be "stored procedures". These inference elements, resident with their object, are called "inferlets". The inferlets reside at a "type" level and will be applied to instances depending on the type-instance hypotheses. For example, if aircraft "tail number" A123 could be a MIG-29, then the type MIG-29 (or a higher superclass) would contain the inferlet applicable to the instantiated, and perhaps new proposed, propagations from the instance A123. The propagation paths will be created from the set of path types specified in the superset of associations from the class object AIRCRAFT and its superclases. "Propagation" here is being used in the broad sense, from Bayesian, Dempster-Shafer, and neural-network, to state-estimation algorithms, i.e., any sort of method that would influence belief regarding the state(s) of adjoining nodes.

There are four challenges upon which we propose to focus our research, as described in the following paragraphs. These are:

- Tasking model for inferlet governance
- Prioritization by utility estimation
- Activation completion
- Controlling feedback

Some issues, applicable to spreading activation networks in general, are discussed extensively in the literature [8], and are beyond the scope of this research. We will, however, show how it is possible to execute the afore-discussed model and that problems are no different than those encountered in inference networks in general.

The tasking model allows the inferlet activations from any belief update (or new belief information) about an object to spread in a non-ordered (asynchronous) manner. DBMS "triggers' are an example of how inferlets activation could spread. We will build a task executive to parcel computing resources to the inferlets that request activation The tasking model will retrieve and execute the inferlet. Inferlets could act like automata. However, we believe governance is required for three reasons

First, it is necessary because the computing demands of activation networks can be large so it is necessary to prioritize the activations in the most profitable manner. Many of the belief updates will be of low immediate consequence. The tasking request from the activated nodes might include the degree of change of the input. but the automata cannot estimate utility in this manner because they have only local information. What is at once the great strength of the Bayesian Net and asynchronous automata is a weakness in estimating the system benefit of an update. We propose an associate to the task executive that will discern if any sets of requested tasks seemed to be leading in productive directions, perhaps using information utility measures along the lines in [1][2] and the Operational Utility measures we describe in paragraph 3. That is, the Operational Utility measures will be more than just testing measures, but will be used online for task control.

The second reason we believe the inferlets require governance is knowing when the spreading is done. This is a common problem in many connectionist structures and is a well-studied problem in neural networks. Examples of techniques to end the spreading are energy functions (e.g., Liapunov) and activation path thresholds. In the former, a cost function computes the energy of the spreaded activation and considers it to be complete if it reaches local minima or an established threshold. In the latter, a threshold is established (or computed dynamically) for each activation. If the activation value is too low, it does not fire. These generally accumulate at the node so that successive updates at the source node can build up energy that may then overcome the threshold. It could also be that there is no "done", but that activation requests are continuous but are effectively rejected by low prioritization by the Operational Utility estimator.

The third reason we believe the inferlets require governance is reduction, if not elimination, of uncontrolled feedback to prior nodes. Directedness alone can alleviate some of this problem, but not all, as there are cases where two nodes can mutually influence each other. Proper structure of the inter-node influencing can also help; by making sure node influence is by truly causal factors. Unfortunately, this may run counter to independence architecture of the network. Cessation or de-prioritization of insignificant activations, as part of the tasking and activation cessation controllers just described, we believe will help. Evidence tracking is insufficient, because a node's updated state could be based upon multiple neighbor-node updates. Similarly, pedigree tracking is generally impractical because the pedigree could include all the node states, back to the start of the system that influenced current state. An approximation we will research is to track only significant and current pedigree. Maintaining with the each belief hypothesis an accumulation of the inputs that most contribute to current belief could do this.

4.5 Tools and Early Experimentation Results

Silver Bullet has explored the ability to use these large E-R models implemented in a "real-time" DBMS for fusion [10]. Two experiments were conducted for Multi-Source Integration (MSI) of a scale appropriate for the advanced E-2C MSI, the other for ESM/ELINT Similar Source Integration (SSI) of a scale appropriate for the now-canceled Advanced Integrated Electronic Warfare System (AIEWS). The results of the preliminary experiments provide good evidence that, as would be expected, high-performance embedded DBMS's have much higher data access times than application dependent data structures. However, the results do show that even under intense scenarios and with massive fusion on a general-purpose medium performance off-the-shelf computer using a non-real-time operating system, the embedded DBMS can perform adequately. The MSI experiment had average sensor updates of 473/sec resulting in requirements to access and average of 3,368 complex object structures per sec.

Similarly, the ESM/ELINT experiment had average 112 updates per sec. resulting in the requirement to access 4,512 complex objects per sec. These accesses were against a track, situation awareness, and intelligence database with information on over 130,000 complex objects. With more powerful processors, a real-time operating system, and a distributed computing environment, the embedded DBMS could be expected to perform even better. A key enabler in the experiments was the associative memory. This kind of layering, of application-dependent associative memories, with high-performance embedded DBMS's, followed a conventional off-line DBMS for amplifying display and historical data, can enable the types of fusion benefits described in [10] and summarized in paragraph **Error! Reference source not found.**, herein.

5. Plan of Work

5.1 Proof of Concept

The concept for ontology-based fusion and inference has been described in [9] and preliminary experiments have been conducted and reported in [10]. Conduct further experimentation to determine whether basic inference, equivalent to that of 'hard-coded' systems, can be duplicated in the formal ontologic framework and general inference technique described in the referenced papers.

The results targeted are to determine how Entity-Relationship (E-R) models, as mathematically formal ontologic models, can serve as a blueprint for a multi-level object class inference network, that COTS "real-time" DBMS's such as TimesTen [11] can implement the ontology with performance that supports fusion and inference, and that multiple inference types, based upon inter-class or inter-object inference properties, can be executed modularly in controlled spreading activation. If successful, this approach could provide a foundation for the large-scale fusion required for improved situation awareness. This will involve the following tasks:

a. Ontology Design Refinement. Build a small experiment using a portion of the C2 Core model and a portion of the DDRE model for Transmitters as a starting point. These models interlock via the common object class, "MATERIEL-ITEM". Figure 5 is an overview of the ontology E-R model and some inferlets and object instances that will be subject to experimentation and demonstration in the basic effort. The subclasses have specific inference rules that would be maintained at the superclass level and are activated based upon the superclass relationships.
b. Ontology Generation. Implement the model in a "real-time" DBMS.
c. Initial Knowledge Population. Populate the transmitter structures with a portion of an ELINT library and Order Of Battle (OOB) database. Upon creation of a new sensor report, the inferlets, stored in the object class and object structure, will be invoked in a spreading activation to infer target presence, kinematics, classification, organization, activity, status, and affiliation with other targets. The inferlets shall be Bayesian and Kalman; they could be other types as well.
d. Inferlet Adaptations. Inferlets shall be adapted from [12], modularized and conformed to the generalized ontology represented by the E-R model implemented in the DBMS. The inferlets will be triggered by updates to associated entities (objects).

e. Situation Awareness Experiment Design. Design an experiment to show conventional inference from signal parameters to object type and organization.

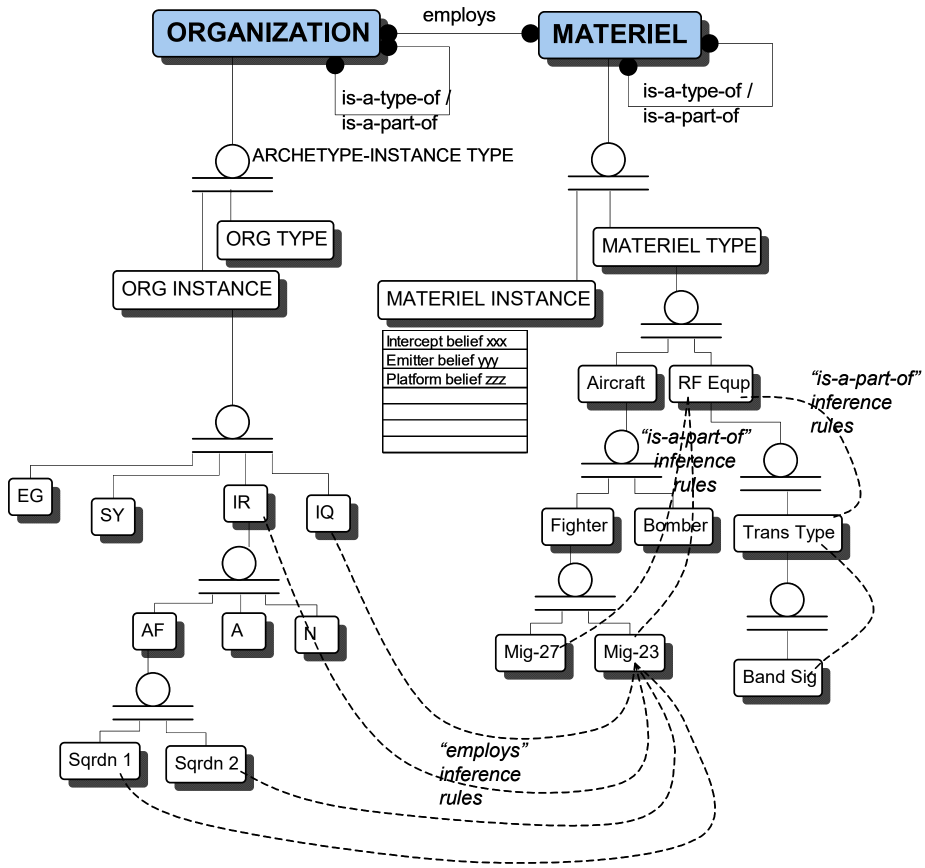

Figure 5. Basic Effort Experiment

f. Evidence Generator (Simulator). Generate simulated ELINT/ESM measurements to be input to the system

g. Measures of Performance Software. Develop the instrument the software for measures of performance for Operational Utility.

h. Controls Setup. Setup switches and add logic to the network so it can be used as a control for the Operational Utility measures.

i. Conduct. Record and analyze results from comparison to conventional systems and for weak evidence accumulation, iterating and experimenting with alternate inferlets and ontologies according to experimentation results.

5.2 Step One – Live Data Feeds and Long-Run Inference

Conduct live data testing using recorded ELINT data and other data sets provided by the Government to validate the Basic Effort results, which used simulated data. Expand from the proof-of-principle in the Basic Effort to research long-run inference in this phases. The experiments shall include long-reach evidence accumulation such as to infer activity or inter-object mission associations (e.g., operating base of an aircraft or launcher of a missile). The following subtasks shall be performed:

a. Live Data Testing, Injection translator. Develop a translator to translate the live data formats for injection through the simulator used in the proof-of-principle.
b. Live Data Testing, Conduct. Conduct testing with the live data sets, measure results, and iterate the network to improve results. The simulator shall also be updated to be in better calibration with live data.
c. Long-Run Inference, Additional Ontology Knowledge Population. Add additional knowledge population to the ontology to support long-run inference.
d. Long-Run Inference, Inferlets. Develop additional inferlets for long-run inference.
e. Long-Run Inference, Controlled Experimentation. Experiment with the long-run inference network to see if decisionable information results.
f. Long-Run Inference, Situation Awareness Experiment Design. Design an experiment to test long-run inference, wherein the decision information inferred is very remote from the raw evidence input to the system.
g. Long-Run Inference, Evidence Generator (Simulator). Modify the simulator and input scenarios as necessary to provide the base evidence.
h. Long-Run Inference, Measures of Performance Software. Develop the instrument the software for measures of performance for Operational Utility.
i. Long-Run Inference, Controls Setup. Setup switches and add logic to the network so it can be used as a control for the Operational Utility measures.
j. Long-Run Inference, Test Conduct. Record and analyze results from comparison to conventional systems and for weak evidence accumulation, iterating and experimenting with alternate inferlets and ontologies according to experimentation results.

5.3 Step Two – Sea Trial and Transition Planning

In this phase, Participate in Sea Trial of the inference network, iterating the design as a result of experimentation results. Also formulate and analyze alternatives for transition and integration into on-line systems.

a. Sea Trial, Planned Exercise Candidates Analysis. Work the Government to determine a suitable Sea Trial appropriate for experimentation with the inference network.
b. Sea Trial, Data Feeds Integration. Integrate the network into the data feeds for the Sea Trial, ashore, afloat, and/or airborne, as specified by the selection of the Sea Trial in subtask (a).
c. Sea Trial, HCI Options and Integration. Integrate the network for Human-Computer Interface (HCI) for the Sea Trial, ashore, afloat, and/or airborne, as specified by the selection of the Sea Trial in subtask (a).
d. Sea Trial, Test Conduct. Support Sea Trial and testing and experimentation with the inference network.
e. Integration Alternatives, Architect. Formulate alternatives for transition and integration of the inference network for ashore, afloat, and/or airborne alternatives.

f. Integration Alternatives, Assess Benefits and Impacts (Cost, Risk, Performance). Model cost, risk, and performance for each of the alternatives, assess the alternatives, and present decision information to the Government.

6. References

[1] *Multi-INT Fusion Performance,* Joint C4ISR Decision Support Center, OASD (C3I), Washington, D.C., 2001

[2] Keithley, H., "An Evaluation Methodology for Fusion Processes Based on Information Needs", in *Handbook of Multisensor Data Fusion*, ed. By Hall, D. and Llinas, J., CRC Press, New York, 2001

[3] Community Imagery Needs Forecast (CINF), Air Sovereignty Operations Center (ASOC), and others

[4] Waltz, E., Llinas, J., Multisensor Data Fusion, Artech House, Inc., Norwood, MA, 1990

[5] Fox, M., Mills, G, Thode, S, Crespo, A.; Navy Tactical Command System-Afloat (NTCS-A) 1990-91 Correlator/Tracker Evaluation; Technical Report 1475; Naval Ocean Systems Center; November 1991.

[6] Ee-Peng Lim; Srivastava, J.; Shekhar, S., "An evidential reasoning approach to attribute value conflict resolution in database integration", *IEEE Transactions on Knowledge and Data Engineering*, Oct. 1996, Volume: 8 Issue: 5

[7] Lee, S.K., "Imprecise and Uncertain Information in Databases: An Evidential Approach", in *Proceedings of the 8th International Conference on Data Engineering*, IEEE, 1992

[8] Liu, J., Maluf, D, Desmarais, M., "A New Uncertainty Measure for Belief Networks with Applications to Optimal Evidential Inferencing", *IEEE Transactions on Knowledge and Data Engineering,* May/June 2001, Vol. 13, No., 3, IEEE

[9] McDaniel, D., "Multi-Hypothesis Database for Large-Scale Data Fusion", *Proceedings of the Fifth International Conference on Information Fusion*, International Society of Information Fusion, Sunnyvale, CA, 2002

[10] McDaniel, D., and Schaefer, G.,"Real-Time DBMS for Data Fusion", *Proceedings of the National Symposium on Sensor and Data Fusion*, Infrared Information Analysis Center, 2003

[11] http://www.timesten.com

[12] McDaniel, D., *Electronic Warfare IDentification (EWID) Small Business Innovative Research Final Report*, Space and Naval Warfare Systems Command, 1996.