From Data to Actionable Knowledge and Decision

Katia Sycara
Robotics Institute
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213
katia@cs.cmu.edu

Michael Lewis
School of Information Sciences
University of Pittsburgh
Pittsburgh, PA 15260
ml@sis.pitt.edu

Abstract-*The Advanced Agent Technology Laboratory in the School of Computer Science at Carnegie Mellon University, the Human Computer Interaction Group at the University of Pittsburgh, the Munitions Directorate of Air Force Research Laboratory (MN/AFRL), Rome Labs and Northrup Grumman are partnering in early-stage research to address issues in high-level information fusion, including intent inferencing and retasking of sensors and munitions. To help translate information superiority to decision superiority (i.e. to rapidly arrive at better decisions than adversaries can respond to), many scientific, technical and technological challenges must be addressed. The most critical of those are information fusion and management at different levels, communication, planning and execution monitoring. The RETSINA multiagent infrastructure allows information producers and users to discover one another and establish direct links. The robust, decentralized Infosphere which results, can be stood up rapidly and ensures that information of the specified types will be delivered to the right users under the right conditions. We will present the overall approach for this research program, the role and contribution of each of the partners, the envisioned testbed, examples and research results to-date. In addition, we will discuss plans for the out-years.*

1 Introduction

We are embarking on a multidisciplinary research effort between computer scientists, engineers and cognitive psychologists from CMU, University of Pittsburgh, the Munitions Directorate of Air Force Research Laboratory (MN/AFRL), Rome Labs and Northrup Grumman to address issues in high-level information fusion. Current state-of the art in Information Fusion addresses primarily level 1 concerns, i.e. sensor fusion. To translate this information superiority into decision superiority (i.e. to rapidly arrive at better decisions than adversaries can respond to), requires addressing many scientific, technical and technological challenges. The most critical of these are information fusion and management at different levels, communication, planning and execution monitoring.

Simply bringing more raw information to decision makers in the battlefield cannot help due to cognitive overload while naïve approaches to computer-mediated information fusion and automated decision making may make things worse by hiding needed data or confusing situational awareness. For future information processing and planning systems, we need an in-depth understanding of the cognitive processes of the user being aided, coupled with innovative approaches for real-time information fusion at all levels, including multimedia and multimodal information from disparate distributed sources that include enormous amounts of uncertainty and noise. Such *cognitively congruent* systems will provide an intuitively understandable common operational picture for enhanced situation assessment and battle management along with planning guidance and monitoring functions in the uncertain and quickly evolving battlespace.

There are striking qualitative differences between the problems and viable approaches to level-1 and higher levels of fusion. In level-1 fusion, sophisticated sensing technologies and mathematics can be brought to bear on a static classification problem, e.g.;[4][5][6]. If the target can be detected and classified within time limits fusion has been successful. At subsequent levels in the chain success is not so easy to define. The basic problem of higher levels of fusion is directing the right information to the right decision maker at the right level at the right time, e.g.; it is essentially a data aggregation, categorization, and switching problem. In the modern battle arena, decision makers may be called upon to simultaneously monitor and control thousands of parameters. It is impossible for a human to do this task unaided. The most limited human resource in these situations is *attention*. The standard *supervisory control* solution from industry has been to automate most of the control functions and trigger alarms when monitored parameters exceed set points. A system, annunciated in this way becomes an *attentional filter* shifting the operator's focus from subsystem to subsystem as troubles are detected and corrected.

Unlike industrial systems that are relatively constrained, battlefield situations are unconstrained where information and need for attention are in constant flux. Infrared sources and acoustic signals indicative of a tank are only of archival interest if the tank is not in a position to pose a threat; however, if friendly infantry is moving into the

[1] This research was supported by AFOSR contract number F49640-01-1-0542.

area, the tank detection needs to be alarmed and responded to. Such *active annunciation* is needed to direct the commander's attention. This active annunciation relies on an underlying system that does in-context filtering, collection and organization of the different types of information needed; e.g.; information fusion. We have recently completed studies[7][8] on visual techniques for directing this attention. The problems that must be addressed to build these next generation systems are:

- Integration of heterogeneous information sources both internal and external to the Air Force (e.g. intelligence reports)
- Adaptation/Customization to rapidly changing threats and missions.
- Automation of information selection and dissemination to provide the right information to the right decision maker at the right time, to keep pace with the shrinking decision cycle.
- Scalability to accommodate ubiquitous sensing and other expanding information sources.

Our overall research hypothesis is that the way to address the above issues is through *adaptive and self-organizing collection of Intelligent Agents* who also possess models for discriminating and communicating situational distinctions salient to humans and the current mission.

2 Multiagent Architecture for Information Fusion

Our vision is that to effectively address the information fusion task at all levels in a uniform and *interoperable* way, each user is supported by a distributed collection of intelligent agents that coordinate to *access* and *fuse* information from distributed heterogeneous information sources, *monitor* information sources at many levels in a focused way, make *inferences* based on information access, determine *new information gathering* and fusion activities (e.g. new areas to search for targets) based on their inferences, and provide *customized views* and *alarms* to the commanders.

The software embodiment of the hypothesized system is the RETSINA multiagent infrastructure [13], which we have been developing over the past few years and which will be leveraged and adpated to fit the needs of the current research. RETSINA is a multiagent infrastructure that is designed to solve the problem of integration of information from distributed heterogeneous information sources and utilize the results in decision support in an open and dynamic environment. Information requests in RETSINA are specified generically, at a semantic level, independent of the structure, location, or even existence of the requested information. RETSINA processes these requests appropriately (e.g. decomposing a request) and flexibly matching them to underlying information sources that are relevant at the time the request is processed. In other words, RETSINA provides flexible, extensible means to locate information relevant to a task *during task execution*, and deals with incomplete, uncertain information and partial results.

RETSINA supports many different specialized agents that represent the user, the decision task that the fused information is meant to support (e.g. intent inferencing task), and the information resources. In RETSINA, these agents are of the following general types:

Interface agents receive the user requests. These agents use knowledge of the domain and the user role and functionality to help the user formulate and customize his/her information requests. They also plan appropriate interactions with other agents on the user's behalf. Interface agents can also learn user information needs and proactively pre-position information that the user has been observed in the past to need [1].

Task agents have knowledge of the task domain and also have planning abilities. Utilizing the task models and through planning, they are able to plan for a specific information gathering goal or an inferencing task. For example, they can *decompose* a high level information request (e.g. "find the state of readiness of a platoon") to lower level tasks and *form plans* for how to execute the information gathering subtasks, *find* (through middle agents) and *query* the appropriate information sources, and coordinate the *query execution* and *composition of the query results*. Task agents use planning mechanisms that combine planning and execution of the information gathering and inferencing tasks. RETSINA agents utilize the Hierarchical Task Network (HTN) formalism [11]. Planning is necessary, and it adaptively considers the current operational environment, for example different information sources may be present at one time but absent at a different time (e.g., the source or its communication link could be damaged during battle). We are identifying different sets of task agents for fusion tasks. For example, we are developing task agents that perform information fusion for different high level information goals, e.g. inferencing of readiness of particular units, inference of enemy intent, availability of friendly transport assets etc. These agents will have appropriate domain knowledge that guides the information gathering and fusing process to provide needed inferences.

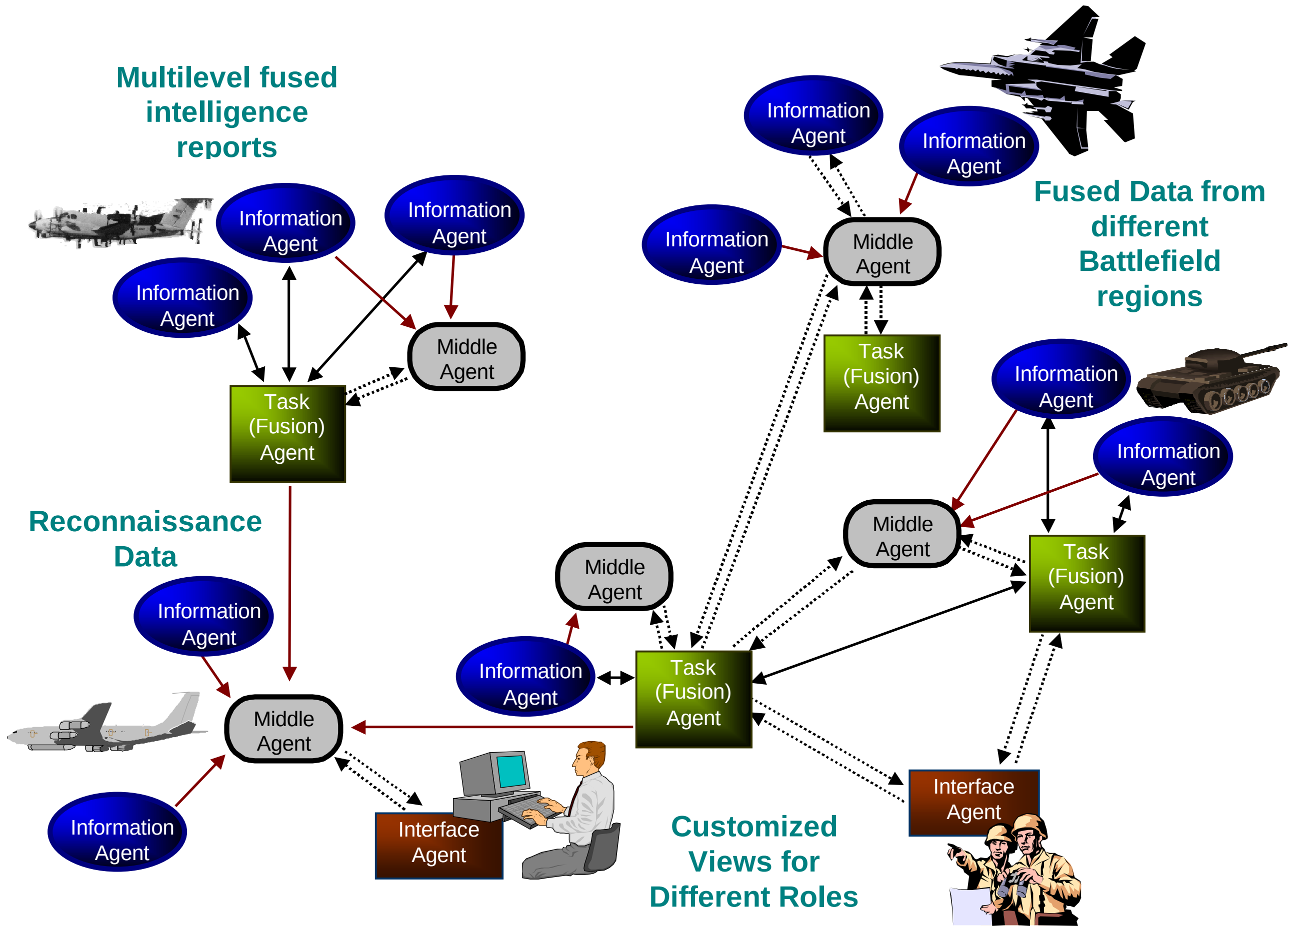

Figure 1. Agent mediated C4ISR

Information agents: each information agent *wraps an information source* and knows the particular details of how to interact with the source to answer a query. For example, besides one-time access to information, an information agent can be given event monitoring and notification triggers called *monitoring queries*. For example, "provide me the price of IBM every 5 minutes", or "provide me the price of IBM if it goes above a particular threshold" [3].

An information agent describes declaratively its underlying information source in a high level semantic description, its *capability advertisement*, which hides the details of query execution. When an information producing entity (an *information service provider*) is initiated, an information agent that wraps it is created. This is done both for new information sources, e.g. web pages, and also for legacy databases. This wrapping allows a passive information source to be "agentified" and interact with agents. An agent's capability advertisement is communicated to middle agents. Note that since all types of agents can also produce results that are potentially useful to others, all types of agents can advertise their capabilities with middle agents so they can be found.

Middle agents provide a level of indirection in information processing. They act as *intelligent registries* of agents. They allow service/agent discovery and lookup by *semantically matching* information needs of *requester* agents (or humans) with available information resources so that the requests can be routed to the available resources. Note that industrial lookup services such as SUN's JINI can be viewed as middle agents with very restrictive matching abilities. We have identified different types of middle agents with different performance tradeoffs [2][16]. Naturally, there are many distributed middle agents in the system. Through distributed discovery and query propagation, service providers can be found and queried regardless of which middle agent they have registered with. The battlefield is an "open" and dynamic information environment where information sources, communication links and computational processes can dynamically appear and disappear. The intelligent process and information registries (the middle agents) provide ways for information providers to find information requesters, allows automated discovery of services, provide for replication and *functional substitutability* of information sources (if a source goes down, another can be found through an information broker), and provide *customized subscription services*, so that information has *targeted dissemination*, thus conserving precious network bandwidth and minimizing risk of enemy eavesdropping on communications. We have already implemented and demonstrated the middle agents *interoperability architecture* in many different scenarios (e.g. a NEO scenario, various joint mission planning scenarios using ModSAF [12]; electronic commerce applications [15] and financial portfolio management [14]).

Realization of an infosphere depends on being able to create, track & maintain a huge collection of information at a variety of levels of abstraction. This can be done within a RETSINA multi agent system by creating agents to track and maintain information about newly identified entities. So, every time a new "object" is added through level-1 fusion it would register with its geographic, alarm, and view appropriate middle agents. These middle agents, in turn, would notify any higher-level entity that was registered to be informed of new entities of this type. One type of interface agents are *presentation agents*: these agents coordinate presentation of different information views and alarms to the user interface agents and user display.

In our agent-based architecture, besides the automated information fusion done by information and task agents, humans also serve as explicit information fusers by entering and/or annotating their assessments of infosphere entities. These entries are then maintained by the system and made accessible to human and software agents alike. The architecture uses advertisement and subscription mechanisms to provide precise filtering of information by directing attention to problem situations (*alarms*) and role and task relevant information (*views*). While users will retain access to any information available from registered entities by querying middle agents, automation of information selection and distribution is the key to reducing cognitive overload. In our architecture presentation agents automate access to information in two ways: by pushing an alert when alarm conditions are met (*active annunciation*) and by providing indexed access to pre-selected typed information views to perform standard tasks. Subscription to both types of information sources is based on a user's role. So, for example, the role of a platoon commander might entail subscriptions to a variety of geographically near information sources making up views appropriate for reconnaissance, re-fueling, and other activities. The commander's alarms would likely be for nearby entities as well but might include aircraft at a substantial distance if within attack range. An officer at the command center might have substantially different views and alarms for the same battlespace.

Unlike annunciator systems for airplanes and powerplants which rely on static set points, active annunciation requires specifying complex relations between a potentially threatening entity and a role. The alarms are registered when the entity's agent is created and continuously monitored thereafter. Views by contrast, are role driven and therefore maintained by the presentation agent. Additional standard views can be defined over Infospace parameters such as level of fusion/aggregation or certainty to allow viewing the battlefield at different levels of strength or to separate facts from conjecture.

The control of the overall system is asynchronous, *bottom up* and *top down*. User requests and task agent plans provide a top-down context and guide information retrieval and fusion; on the other hand, bottom up, data driven control is provided by automatically triggered updates and notification. The whole system composed of humans, agents and entities in the world (e.g. missiles, targets, etc) can be thought of as an asynchronous, distributed, virtual blackboard. The system closes a high level feedback control loop by *continuously adapting* several key aspects of its behavior to its dynamic situation: its perceptual strategy (what information to get at each time, what areas to search next for critical entities, e.g. targets), its choice of reasoning tasks to perform, its choice of what actions to suggest, and what output information to display.

2.1 RETSINA and the JBI

Much of the technology necessary for the JBI infosphere has striking correspondences with elements of the RETSINA architecture.

The RETSINA infrastructure can be leveraged to rapidly build a JBI prototype to support:

- Insertion of data from sensors, databases, ATR (Automatic Target Recognition) systems and other sources through *information* agents.
- Information exchange through subscription of interest with *middle* agents; this facilitates discovery and delivery of information as it becomes available.
- The development of *task* agents or fuselets, that generate higher-level information by fusing existing data with military doctrine, commander instructions, or other available data.
- The development of *entity* agents to represent *updatable knowledge objects* that correspond to, and model, entities within the battlespace. Such agents can be updated as new information arrives at the JBI, and queried to enhance situational awareness.
- Agents that represent higher echelons, by collaborating with entity agents (representing lower level echelon entities), and maintaining higher level information.
- Customizable interface agents that allow presentation of multi-modal data on a variety of displays and devices.

Table 1: Parallels between JBI and RETSINA infrastructure

Joint Battlespace Infosphere	Retsina Architecture
Browsing capabilities	Several different interface agents have been developed, for customized display of geospacial data (MokSAF, Tandem [9]), presentation through web-based interfaces (A-Match, Coala [15]), and for any-time, anywhere delivery of device independent data. Current work on DARPA Agent Markup Language (DAML) will lead to the creation, location and presentation of semantic data on the web.
Interaction	The RETSINA architecture currently supports multi-modal interaction and delivery of information, from large displays and electronic whiteboards, to PDAs, limited interaction devices (e.g. WAP-enabled cell phones) or audio-only devices.
Fusion	RETSINA task agents have been used to fuse data supplied by other agents for information verification aggregation and decision-making prior to presentation (Warren [14]).
Objects; authoring and subscription	Middle agents facilitate discovery and communication of other agents through the use of a capability description language. This describes the functional capabilities of the agent, and how it may be contacted. Existing information agents' posses information (or acquire it dynamically) that can be requested. Monitor queries actively check subscribed data sources until new information is available.
Structured Common Representation	The DARPA Agent Markup Language (DAML) facilitates a semantic markup of information based on a structured common representation. RETSINA agents utilize this markup to describe and reason about concepts, thus supporting interoperability with other RETSINA agents, and other agents that conform to the DAML specification.
Automatic Data Capture	Several RETSINA agents support automatic data capture for proactive information retrieval and user support. JoCCASTA and Agent-Storm RETSINA-based applications [12] to monitor communication (both text-based and voice) between individuals, and issue requests for information to other agents based on the context of the discussion.
Tailoring Information to meet user needs	The RETSINA environment supports personalized information delivery on different hardware platforms, and tailors the information when possible to meet the limitations of the device used.

2.2 Proactive Fusion and Intent Inferencing across multiple layers

Local models in the form of plan fragments and local inference models are used to localize and control information gathering and fusion. These plan fragments model the pieces of information that must be collected and reasoned about, assumptions and dependencies among them. We are in the process of developing *inference-specific task agents* associated with different types of intent inferencing for different categories of battlefield entities. For example, an agent is associated with inferring the state of readiness of platoon-1, another one of platoon-2. Each one of these agents has associated model fragments that control and trigger its information gathering plans and evaluation of evidence for that type of inference. Agents at higher level of abstraction (note that abstraction layers do not necessarily imply hierarchical system control) are responsible for combining results from more specific inference agents, plus planning for and gathering additional information that may be needed for the higher level inference. In turn, such inference may be used to focus future information gathering activities (e.g. where to search for a target).

Intent/Threat inferencing can be done *proactively* through planning by a task agent *in the context* of an agent's or user's task. Associated with planning goals and sub-goals are the information requirements that could determine whether a goal is satisfied. The planning agent then, instantiates the information-gathering requirements, finds through middle agents the information providers that can fulfill the information requirements and interacts with them. The result would be a collection of entities of varying levels of aggregation and certainty. Information will be filtered in this way to provide displays of information of differing degrees of certainty. In addition, views would be filtered by level of aggregation to provide high-level summary or low-level displays of entities. The planning mechanism we have implemented in the RETSINA task agents uses a Hierarchical Task Network (HTN) formalism [11] that integrates information gathering, planning and execution monitoring. Such planning makes use of appropriate domain knowledge and guides information access and monitoring of information resources. For example, inferring the threat intent of enemy tanks could require:

1. Information retrieved from sensors that can recognize tanks,
2. Identifying the formation of tank tracks over time, and
3. Recognizing that the tanks are in attack "wedge" formation. The recognition and tracking tasks could be done by different agents that are coordinated through the planning agent.

This scheme is more flexible than conventional evidence propagation schemes where every variable of interest must have been included in the model statically from the start. For example, determining the state of readiness of the unit comprised of platoon-1 & platoon-2 may include additional pieces of information that must be dynamically incorporated into the assessment, than simply the union of the information needed to infer readiness of platoon-1 and readiness of platoon-2. Second, the task-based distributed control of the overall multiagent system does not require all of these task inferencing models be active at the same time, (which would unnecessarily congest the network) but only when they are in the inferential chain of higher level inference processes that have been activated either top down or bottom up by some observation of interest. Therefore, this *coordination scheme promotes scalability and robustness* of the overall system.

Using the RETSINA infrastructure, we have demonstrated prototypes for automated coordination of threat inferencing and sensor information integration through automated interaction of our agents with the ModSAF simulation environment. When such threats are inferred, alarm notifications are sent to the user(s). We have provided such information aggregation capabilities in the TANDEM task [9][10], which is a low fidelity simulation of Aegis cruiser target identification and threat intent identification by a team.

3 High Level Information Fusion Testbed

Unlike level-1 fusion research which can validly rely on archived sensor traces, the active and temporally sensitive character of higher level fusion requires an active battlespace with many disparate forms of data. Because we cannot gain access to timeslices of radio traffic, sensor signals, and computer messages from real operations we must generate such data in order to study higher levels of fusion. We are using the OTBSAF simulation as our testbed to investigate *realtime information fusion.*

OTBSAF is a transitional computer generated forces (CGF) simulation bridging the gap between ModSAF, the CGF simulation which grew out of SimNet, and OneSAF, an HLA compliant next generation successor. OneSAF is slated for wide use in both training an research so our adoption of this simulation benefits both us by providing a large scale detailed military model and the larger community of simulation users by adding more sophisticated sensor models to a widely used tool.

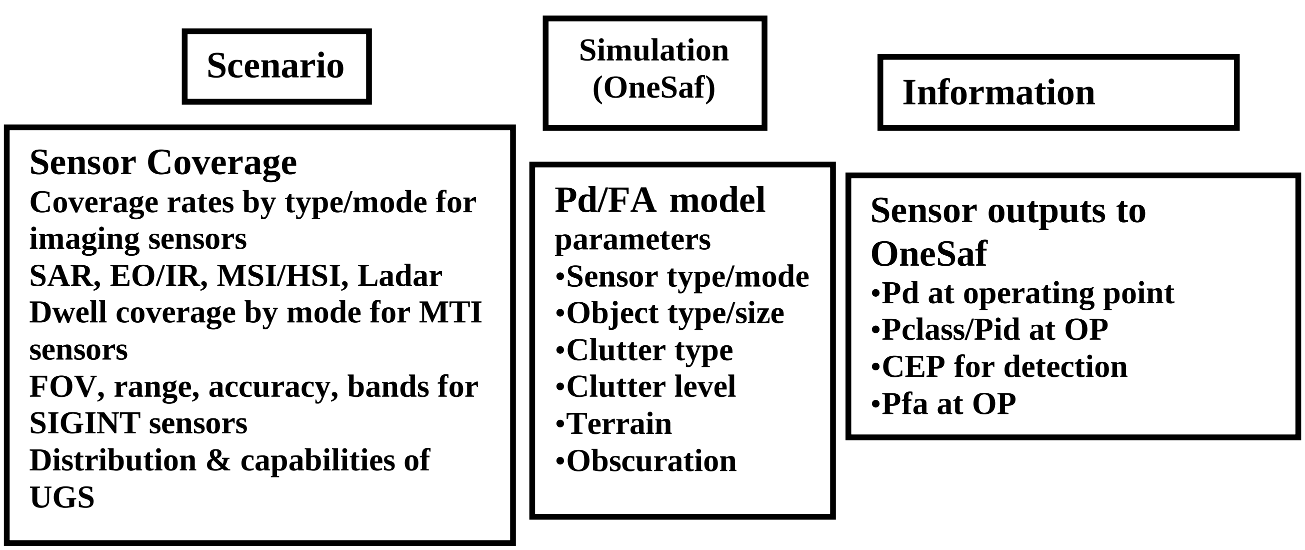

Figure 2. Sensor Model

Working in conjunction with researchers from Northrup Grumman we are currently developing sensor models for SAR, Ladar and other advanced instruments for OTBSAF that will take into consideration factors such as multiple resolutions, resolution/coverage tradeoff, and the effect of clutter/confusability on identification ROCs. These models along with existing models of conventional sensors will be used to distort the ground truth from the simulation so as to provide a realistic approximation of the sorts of input likely to be available to higher level fusion. Figure 2. shows the basic categories of inputs and outputs planned for the sensor models.

In a parallel effort specialized RETSINA agents are being developed to perform the alarm and view generation functions. Together, the sensor models, agents which access their data, and agents which distribute alarms and views realize a simulated JBI based on the OTBSAF simulation. The resulting High Level Fusion Testbed will allow us to conduct realistic experiments in accuracy and scalability of complex fusion tasks such as threat inferencing. The development of a High Level Information Fusion Testbed based on OTBSAF is also a necessary bridge to our collaboration with AFRL/MN. As shown in Figure 3, OneSAF serves as the ground truth reference simulation for the LOCAAS End-game Analysis Program Simulation (LEAPS) at AFRL/MN. The LOCAAS (Low Cost Autonomous Attack System) simulation of cooperating wide area search munitions which will complete the kill cycle enables us to study information acquisition and use in a dynamic realtime environment.

4 Conclusions

In this paper, we have described our multiagent approach to the problem of information fusion that spans all levels.The approach uses a community of software agents to perform fusion at levels 2-4 and an agent-based publish/subscribe mechanism to distribute the information. We are in the process of developing sensor models and scenarios and incorporating them in the OTBSAF (OneSAF Testbed Baseline) battlefield simulation to test the approach. We have made substantial progess in three additional research subareas of our program: (1) we developed and experimentally tested initial models of directing human attention in open environments that contain hundreds of potential attentional targets [7][8], (2) we built models of optimal allocation of munitions to targets in dynamic environments [17] and (3) developed and tested models of automated classification of financial news stories and automated topic tracking [www.cs.cmu.edu/~softagents/text-miner.htm].

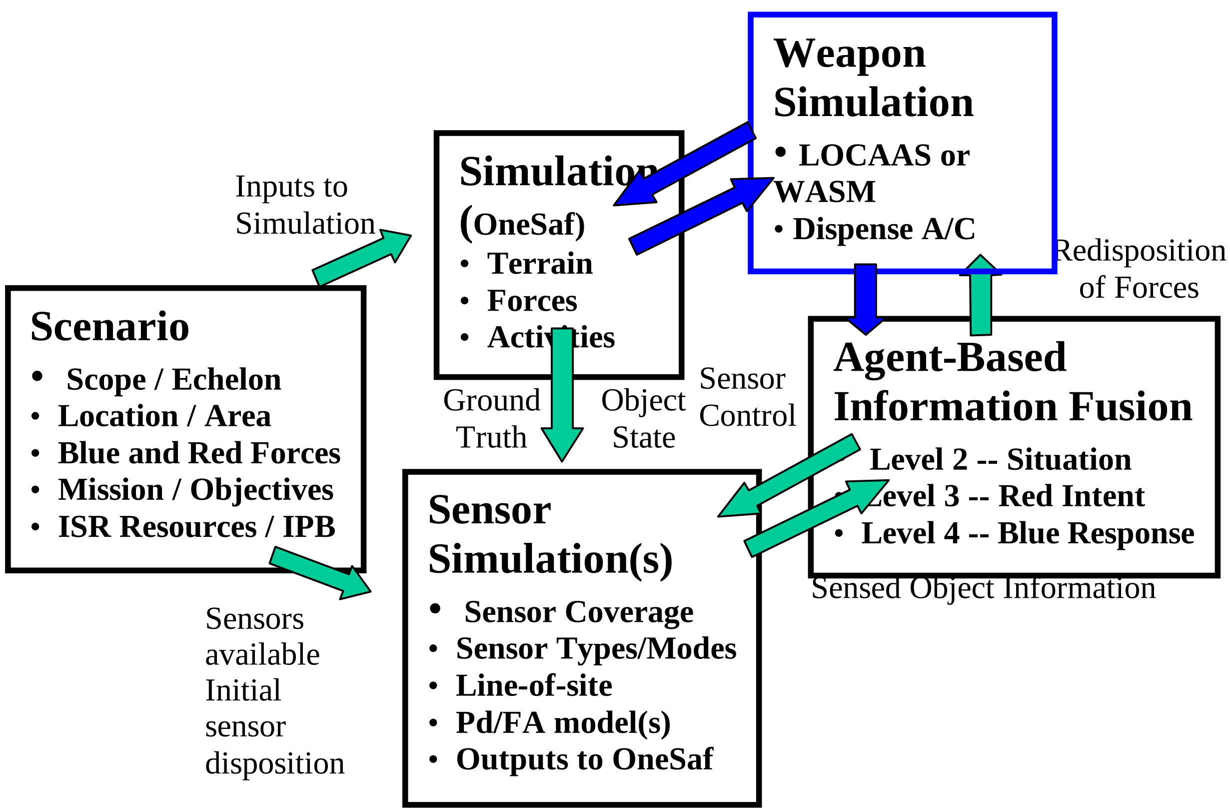

Figure 3. Information Fusion Testbed and LOCAAS simulation

References

[1] Chen, L, Sycara, K. "WebMate: A Personal Agent for Browsing and Searching", *Proceedings of the 2nd International Conference on Autonomous Agents*, Minneapolis, MN, May 10-13, 1998.

[2] Decker, K., Sycara, K. "Middle-Agents for the Internet", *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence* (IJCAI-97), Nagoya, Japan, 1997.

[3] Decker, K. and Sycara, K. "Intelligent Adaptive Information Agents". *Journal of Intelligent Information Systems*, vol.9, pp. 239-260, 1997.

[4] Geiger D., and Heckerman, D. "Knowledge representation and inference in similarity networks and Bayesian multinets". *Artificial Intelligence*, 82:45-72, 1996.

[5] Jensen, F.V. "An introduction to Bayesian Networks". UCL Press 1996.

[6] Kjaerulff, U. "Nested Junction Trees". *Proceedings of the 13th Conference on Uncertainty in Artificial Intelligence*, Providence, R.I. 1997.

[7] Hughes, S. and Lewis, M. "Attentive interaction techniques for searching virtual environments." (to appear) Proceedings of the 46th Annual Meeting of the Human Factors and Ergonomics Society; Sept. 30-Oct. 4, Baltimore MD 2002.

[8] Hughes, S. and Lewis, M. "Directing attention in open scenes." (to appear) Proceedings of the 46th Annual Meeting of the Human Factors and Ergonomics Society; Sept. 30-Oct. 4, Baltimore MD 2002.

[9] Lenox, T., Lewis, M., Roth, E., Roberts, L., Shern, R. and Rafalski, T. "Support of Teamwork in Human-Agent Teams", *Proceedings of the 1998 IEEE International Conference on Systems, Man, and Cybernetics*, Oct 12-14, San Diego, CA 1341-1346, 1998.

[10] Lenox, T., Hahn, S., Lewis, M., Payne, T. and Sycara, K. "Agent-based aiding for individual and team planning tasks" IEA 2000/HFES 2000 Congress, 2000

[11] Paolucci, M., Sycara, K. and Shehory, O. "A Planner for Agents in Open, Dynamic Multiagent Systems", *Proceedings of the IJCAI-99 workshop on Scheduling and Planning Meet Real-time Monitoring in a Dynamic and Uncertain World*, Stockholm, Sweden, August 2, 1999.

[12] Payne, T., Sycara, K., Lewis, M., "Varying the User Interaction within Multi-Agent Systems", *Proceedings of Autonomous Agents 2000*, Barcelona, Spain, pp. 412-418, 2000.

[13] Sycara K., Decker, K., Pannu A., Williamson M., and Zeng D., "Distributed Intelligent Agents", *IEEE Expert: Intelligent Systems and their Applications*, Vol. 11, No. 6, December 1996, pp. 36-46.

[14] Sycara, K., Decker, K, and Zeng, D. "Intelligent Agents in Portfolio Management". *Agent Technology: Foundations, Applications, and Markets*, N. Jennings and M. Woolridge (eds). Chapter 14, Springer Fall 1998, pp. 267-283.

[15] Tsvetovat, M. and Sycara, K. "Customer Coalitions in the Electronic Marketplace". *Proceedings of the Fourth International Conference on Autonomous Agents*, June 3-7, Barcelona, Spain, 2000. pp 263-264.

[16] Wong, H.C. and Sycara, K. "A Taxonomy of Middle Agents for the Internet" *Proceedings of the Fourth International Conference on Multiagent Systems*, July 10-12, Boston MA., 2000 pp. 465-466.

[17] Yang, W. and Sycara, K. "Dynamic Weapon Target Assignment Problem" (working paper).