

Agents for Net-Centric Warfare and Time Critical Targets

“Terrorists don’t think and move the way armies do. In Afghanistan they learned to hide when they suspected American Bombers were targeting them. To defeat groups like al-Qaeda, you have to attack as soon as you find them. With speed you can deny sanctuary to terrorists. In Afghanistan the Pentagon was able to reduce that time (the kill chain) to less than 10 minutes. But in those several minutes...at Tora Bora al-Qaeda fighters were able to escape due to delays in striking targets. In the future U.S. forces will have to locate targets and drop bombs in the time it takes to watch a couple of TV commercials. But speed has to be balanced with reliability (and accuracy), as at least a half-dozen accidental attacks on Afghan civilians and U.S. troops showed.” From a recent article in USA Today (“Afghanistan’s Lessons Shaping New Military”, Oct. 8, 2002).

Systems/Software Engineering Considerations

Time critical targets (TCTs) are time sensitive targets with an extremely limited time window of vulnerability, the timely attack of which is critical to ensure the successful execution of the global strike task force (GSTF) mission.

Allied experience with Iraqi mobile Scud theater ballistic missiles (TBMs) has led to increasing emphasis on coordinating systems and operational concepts for dealing with mobile, time critical targets. The need for increased precision is exacerbated by recent incidents of “friendly fire” (“collateral damage”, “fratricide”) against U.S., coalition and civilian personnel. As a result, many agencies are currently looking at multiple sensor programs and “network-centric collaborative targeting” information systems to increase the timeliness and accuracy of “engagement quality” information to the warfighter.

From a “system-of-systems” engineering perspective involving multiple sensor and weapons platforms, timeliness and accuracy requirements for supporting ISR operations clearly depend on the exposure times and changing signatures of targets, the accuracy and time demands of precision weapons, the need to synchronize non-lethal and lethal attack operations, and constraints imposed by differing levels of rules of engagement (ROE) designed to limit risks of collateral damage and/or fratricide.

Because no single defense element can provide the desired level of performance (near-zero leakage), especially against TBMs armed with WMD warheads, there has been substantial emphasis on developing a multi-layer defense capability (i.e., defense in depth) for TMD. For this analysis, we focus on Air Force counterforce operations against the TELs - protected by advanced surface-to-air missiles (SAMs) - and the necessary ISR

to support these operations. TEL strike/SEAD counterforce operations, to the extent that they can be successful, clearly benefit the other elements of the TMD defense-in-depth architecture by reducing the number of missiles and the salvo potential that other defense elements (BMD boost phase, midcourse, terminal, and passive defenses) will have to counter.

The combination of threat avoidance, judicious tactics, self-defense capabilities, and suppression packages (e.g., jamming aircraft, aircraft carrying anti-radiation missiles, air launched decoys), has been effective against older generation SAMs. However, these successes have not been without considerable costs and risks. Low-level air operations, in particular, have been hampered by an inadequate ability to detect, locate, identify, track, and target anti-aircraft artillery and man-portable infrared-guided and electro-optical-guided SAMs.

Successful air operations will be challenged further if future adversaries deploy more-capable air defense systems, such as the SA-10 and SA-20. The worrisome trends associated with these emergent threats include increased range, lethality, mobility, and netting (i.e., development of an integrated air defense system). The expected performance of these advanced air defense systems places current CONOPS and weapon systems at increased risk and will require improvements in all aspects of the engagement process.

Current studies show that today's ISR capabilities are inadequate to support pre- and post-launch counterforce operations against TBM TELs protected by advanced ("double-digit") SAMs. Improvements in six C2 functional areas are essential to ensure that the targeting process is completed in the tight, threat-driven timelines - with timelines approaching 3-5 minutes or less (not including the flight times of re-tasked sensors and strike aircraft), measured from the time of initial detection of a TCT to issuance of the engagement order to a strike aircraft. The six functional areas are as follows:

- Integrated tasking and rapid retasking of sensors and related processing, exploitation, and dissemination (PED).
- Timely integration (correlation and fusion) of information derived from multiple sources.
- Rapid target development and target nomination.
- Rapid weapon and target pairing.
- Timely decision and attack order dissemination.
- Rapid assessment of effects of weapon delivery.

The key categories of potential improvements to ISR performance are new sensors and sensor upgrades, and a dynamic, collaborative ISR environment supported by a high performance networked communications infrastructure.

New Sensors and Sensor Upgrades

Because of the depth and operational flexibility inherent in long-range TBM and advanced SAM threats, we will require new sensors and sensor platforms. These must address the serious challenges of providing affordable and militarily useful capabilities,

including deep-look, long-dwell, all-weather/day-night operations, and acceptable survivability in the face of advanced air defenses (e.g., Global Hawk and possibly high altitude balloons). Improved signals intelligence systems are needed that can operate across the entire signal spectrum and quickly identify and accurately locate emitters that operate intermittently, as well as perform specific target identification (STI) and tracking. There will be increased emphasis on, and requirements for, ISR sensors that can surveil mobile targets until weapon systems achieve target destruction. For post-launch operations, focused-look surveillance capabilities (e.g., space-based IR) are sufficient, whereas for pre-launch counterforce operations, broad-area search with much higher search rates at higher imagery resolution (relative to existing capabilities) is needed (e.g. SBR). These advanced capabilities will be needed in peacetime, day-to-day, for intelligence and strategic warning, in short-notice crises, and in war.

Sensor upgrades may also be appropriate for the fighters, as well, depending on the specific operational concept (e.g., the extent to which the fighters can rely on off-board targeting, such as coordinates for JDAMs employment, versus having to find, track, identify, and target the TELs themselves based on relatively crude off-board cues). For example, improving the resolution of the F-15E's APG-70 in the synthetic aperture radar (SAR) mode, adding a GMTI capability, providing the ability in the air-to-air mode to track (and hence, backtrack) in-flight ballistic missiles, and the addition of sensor management aids to improve area search and fusion with off-board sensors, may all be worthwhile improvements for the TMD counterforce mission.

The Need for Dynamic, Collaborative ISR Processes and Tools

Numerous factors drive the need for more responsive, accurate ISR processes in support of military operations: the fleeting exposure times and changing signatures of TCTs; the accuracy and time demands of precision weapons, particularly against mobile targets under collateral damage constraints; and the need to synchronize non-lethal and lethal attack operations.

Two factors determine the precision required from the location and identification functions of the ISR process: the type of weapons to be employed against the target (e.g., unguided or GPS-guided weapons such as JDAM vs. target-signature guided weapons such as HARM); and the theater ROE. For example, in an aggressive ROE situation, employed when political or military repercussions for mistakes would be lenient, commanders may use procedural methods to identify a hostile target. Kill boxes, free fire zones, no-fly zones, guilt-by-association, and point of origin guidance may not require much sensor accuracy to identify a contact as a hostile target. In other cases, moderate or even restrictive ROE may be necessary to minimize the risk of collateral damage or fratricide. Thus, as the risk of unintended consequences increases, the precision required increases.

Our doctrine of information extraction (target detection, identification and location) from ISR assets depends upon exchange of multi-spectral information between sensors. Allowing multiple sensors to "talk" with each other ("cross-cueing") creates a synergy that takes advantage of specific sensor strengths (FOV, resolution, revisit cycle,

day/night, all-weather, etc.) and overcomes system weaknesses. For example SIGINT systems generally can detect and identify potential targets, but initially can provide only coarse positional information until numerous intercepts have been made. Imagery systems can provide detailed positional information, but have difficulty surveying large areas. Providing identity and rough positional information from SIGINT sources to IMINT systems managers electronically, in real-time, could greatly enhance the speed and accuracy of the overall targeting process as well as provide a measure of confidence against deception and/or risk of collateral damage.

More-dynamic ISR systems, procedures, and CONOPS are needed to shorten response times so that mobile TCTs can be killed before they flee. Automated tools and procedures are needed for rapid retasking of sensors and associated PED, and for processing, correlation and fusion of data from multiple intelligence disciplines (i.e., multi-INT fusion). Decision aids are also needed to more effectively support the dynamic management of future integrated TEL strike and SEAD packages. SEAD CONOPS must be updated to reflect the need for more-dynamic ISR procedures. Furthermore, these CONOPS must be flexible enough to respond appropriately to variations in enemy air defense strategies ranging from maximum engagement and concomitant friendly-force exposure to episodic engagements with minimal exposure.

The fielding of stealth aircraft and longer-range standoff weapons allows the target to be engaged with minimum exposure time inside the lethal range of enemy air defense systems. But these weapons must be cued. Cueing calls for enhanced ISR capabilities that support rapid retargeting of new standoff weapons. To achieve kill chain timelines of a few minutes or less, “targeting constellations” may have to be predetermined.

Current Operational Situation

A major shortfall in today’s optimal allocation of scarce ISR resources and to cross-cued and/or simultaneous collections of TCTs is the lack of agreed-upon (by DoD and the intelligence community) CONOPS, a system-wide interoperability architecture, and automated tools for integrated tasking and battle management of (1) sensors from multiple intelligence disciplines (multi-INT), (2) sensors from multiple platform domains (cross-domain), and (3) the associated PED to support military monitoring, assessment, planning, and execution processes and timelines.

The present process for tasking and retasking of sensors relies on tools developed for individual intelligence discipline (INT) sensors; usually, the tools also depend on the sensor platform. Even for conducting cross-domain, multi-INT tasking simply to allocate the best sensor (regardless of domain or INT) to a target, the current collection management process is manual and time-consuming.

If cross-domain, multi-INT tasking is desired for more-sophisticated collection operations (i.e., cross-cueing and simultaneous collections), further time-consuming manual interventions - mostly by personnel remote from the AOC - are necessary. Moreover, once tasking is completed, there are no management tools for PED tasking to ensure that collected information is processed and analyzed to support user timelines. In addition, the

current PED tasking and retasking processes are not well integrated with the monitoring, assessment, planning, and execution processes and timelines of weapon systems at various levels of command (e.g., unit, wing, air operations center, joint task force).

Typically, the outputs from a single-INT PED are geared more toward the support of specific standing requests for information or general updates to an overall geographical area of interest rather than to the support of highly dynamic operations. The lack of adequate rapid integration of cross-domain multi-INT PED outputs (products and services) and lack of rapid delivery of these outputs to the various levels of command often impede the rapid development of alternative courses of action and the execution of TCT operations. For the purpose of this discussion, we define this dynamic ISR function as the timely control and management of:

- Access to multi-source, cross-domain, multi-INT data and related context information
- Correlation/fusion of data and information into data streams, information products, and services
- Integrated display of these diverse forms of data and information to support military users

Currently, extensive ad hoc manual interventions are necessary to correlate, fuse and display cross-domain multi-INT data and information for use in the monitoring, assessment, planning, and execution of military operations. These interventions are necessary because of the following:

- Data are collected by multiple, often stove-piped, sensor systems that have diverse phenomenology (e.g., electro-optical, infrared, SAR, GMTI, acoustic).
- Different organizations and activities are responsible for transforming data and information into products and services to support military operations. Moreover, they may apply different data correlation tools, techniques, or algorithms to derive products and services (e.g. spatial registration, change detection).
- Data streams, products, or services may not include accurate time stamps, and their spatial reference frames and accuracy may vary.
- Numerous displays have emerged to help decision-makers visualize correlated information.

As a result TCT is currently executed with human mental fusion of the available (often limited) information. It has been reported that in Operation Allied Force, ISR and combat operations personnel had direct real-time feeds and displays of data from individual sensors (e.g., Predator, JSTARS), as well as intelligence products and services from other organizations. However, they had to manually integrate all this information to prosecute ground TCTs.

As we gain the ability to detect and track moving targets, using multiple GMTI information sources (i.e., national, theater and tactical systems), it must be determined how best to integrate the data from these sensors and convert the data to useful

information on targets of interest. Also important will be the development of techniques and procedures for integrating GMTI and video information and other sources of data (e.g., hyperspectral imagery) that can help find, fix and track moving TCTs. The emergence of fighter and bomber aircraft (e.g., F-22, Joint Strike Fighter) with improved sensors and improved data processors presents a further opportunity to improve situational awareness across all echelons of command. However, to realize the potential of these new sensing capabilities, data from fighters and bombers have to be backlinked to data correlation nodes (e.g., the control and reporting center) and integrated with relevant data from other sources (e.g., national reconnaissance and airborne surveillance sensors).

Emergence of the Network-Centric Warfare Model

Existing battle management systems will need to be enhanced to support such pre-launch and post-launch counterforce operations in very short timelines (approaching 3-5 minutes or less). The limited numbers of individual sensors are deficient in ability to provide autonomous detection, combat identification, and precise targeting across a theater of operations. Furthermore ISR planning and execution management capability and processing is slow to adjust the collections plan or to capitalize upon opportunities in dynamic situations for efficient and effective support to TCT. C2 nodes do not have the automated capacity to correlate/fuse data from the vast array of ISR platforms for timely target and identification development, nor the automated means to assess weather effects and aid effective target-to-weapon pairing.

Advances in individual sensors alone cannot meet this challenge. The geometric improvement in accuracy and timeliness required can only be achieved by system-wide network-centric collaboration of the ISR fleet. This concept envisions collaborative operations among the various ISR fleet elements by the automated horizontal integration of sensors at the front-end (“upstream”) to coherently capture fleeting events via synchronized sensing and fusion operations so as to create composite TCT strike packages containing actionable information in real time.

The challenge is to achieve the technical, cultural, and organizational changes that will make this improvement possible. Perhaps the most important developments in this area are the network centric collaborative targeting (NCCT) CONOPS and the associated TCT functionality (TCTF) developed by AC2ISRC (USAF) that includes the capability to assist C2 operators with real-/near-real time data access, visualization, and manipulation to support target development/decision-making and asset tasking (sensors, weapons).

“The primary purpose of NCCT is to rapidly synchronize ISR sensors so that they can collaboratively focus on common targets. This process dramatically improves target location accuracy, timeliness, and completeness. NCCT accomplishes this by establishing a wideband network between ISR participants and using a common set of rules of interaction carried with each participant. The NCCT network allows ISR participants to selectively exchange raw sensor information about specific targets very rapidly with high update rate and very low latency.” (NCCT CONOPS)

Effective performance of the NCCT network will depend on factors such as:

- A system-wide protocol which contains “published” capabilities of individual component sensors, location/attributes of sensors and objects of interest (e.g., targets) that are active in a common reference frame (e.g., datum, time reference, TLE, etc.) and expressed in a standard “vocabulary”, and technical standards for system-wide data exchange, correlation/fusion algorithms and dissemination; (for example NCCT envisions the production of a “smart track” data package which would be propagated within the NCCT infostructure, with multiple sensor “operators” collaboratively populating the fields required to support a targeting decision and engagement sequence);
- Dynamic data fusion capabilities based on data from similar and dissimilar sensors, archived data, target templates, and shared track identities; tools that currently support data fusion will need to be restructured to support in-line dynamic processing under variable rules of association;
- Dynamic collection and reporting tasking capabilities based on real-time visualization/amplification rules and cooperative cross-cueing; these functions are accomplished by direct, automated sensor-to-sensor interaction and data correlation at the front-end of the collection process; autonomous real-time retasking will be enabled by predictive templates of effective sensor combinations against multiple target types under specific combat conditions accounting for unexpected, potential enemy courses of action and corresponding reactive sensor coverage options;
- A rule set enabling the network to perform collaborative, decentralized, concurrent operations without human intervention in order to ensure assets are properly synchronized on time-critical events; the rule set must be sufficiently flexible/reconfigurable to enable updates in reaction to the realities of the battlefield environment.

Network Centric Operational Performance Characteristics

In the fundamental shift to network-centric operations, distributed sensor data fusion at multiple operational sites, supported by secure information networks, emerges as a key enabler of increased combat power. The operational value or benefit of sensor networks is derived from their enhanced ability to generate more complete, accurate, and timely information than can be generated by sensors operating in stand-alone mode. Sensor networks provide significant performance advantages over stand-alone sensors in key mission spaces by overcoming the fundamental performance limitations (e.g., coverage and accuracy) of individual stand-alone sensors. The value-added processes of data fusion and sensor tasking can at least in part overcome these limitations.

For example, employment of sensor networks allows us to overcome line-of-sight obscuration by terrain. The combination of sensor tasking and data fusion enables multiple sensors, based in space, the air, or on the ground, to effectively increase the amount and quality of information available. Certain classes of objects cannot be tracked, located, or identified with sufficient accuracy using a single type of sensor or sensor phenomenology. This deficiency can sometimes be overcome by linking sensors of

different types to achieve an all-source capability. This is of particular value in detecting and identifying high-value targets, such as mobile surface-to-air or surface-to-surface missile launchers, as well as surface-to-surface missiles in flight.

Dynamic sensor fusion and tasking provides two things. First, it permits the optimum orchestration of sensor entities to increase battlespace awareness by providing multi-source information. Second, it provides warfighters with the operational flexibility to synchronize observations with the timing and tempo of operations. Sensor tasking may be either pre-planned or real-time. When operating in the pre-planned mode, active and passive sensors can be tasked to collect and provide information required to support pre-planned operations. For example, the collection of battle damage information could be accomplished with pre-planned sensor taskings. The ability to rapidly transition between pre-planned and real-time modes enables warfighters to task sensor entities in real-time to generate needed information on demand. This operational capability will enable synchronization of collection and analysis with rapidly changing timing, operating tempo, and priorities of joint operations.

An excellent example of the power of dynamic sensor fusion:

“USAF officials at Eglin AFB, Fla., redirected a falling bomb to hit within 25 ft. of a moving ground target. Late last month a 1-ton, inert JDAM--modified by adding a data link--was launched by an F-16 flying at 20,000 ft. and 5.3 nautical mi. from the target, which was the second vehicle in a convoy of six moving at 18 mph. A surrogate fourth-generation JSF radar was triangulated with an E-8 Joint-STARS ground surveillance radar to track that vehicle. It was further refined with a NIMA terrain database. The updated location was relayed to the modified JDAM in flight. The test, part of DARPA's affordable moving surface target engagement (AMSTE) program, will be followed by three more conducted at China Lake, Calif., this month. Researchers will drop JDAMs from F-14Bs and a C-model Joint Stand Off Weapon from an F/A-18 using Link 16 at a range of about 20 mi. The radar triangulation will be done with a surrogate Navy Global Hawk radar and a Joint-STARS. In October, additional tests will focus on the ability to track a target for 20-30 min. New algorithms can fingerprint a single target to aid tracking of individual vehicles across a crowded battlefield.” Aviation Week September 23, 2002

Key Enabling Technology: “Intelligent” Software (a.k.a. “Agents”)

As the military moves to network centric targeting and operations, the TCT challenge is driving the movement from client/server and web-based computing to network centric computing. In this evolution, the computing, communication, and content all converge, thereby rendering the network itself as the computer.

Several key trends are becoming increasingly important to end-users in this environment:

- Faster access to large volumes of ISR strategic and tactical data, as well as improvements in the collection ability of sensors and the sophistication of weapon systems, have resulted in a major increase in complexity.

- The volume of information available is so great it is difficult to process and disseminate, resulting in dramatic increases in discarded or ignored data containing potentially valuable information for time-critical operations. Although handling large volumes of data is an important capability, *situationally efficient* handling of data from heterogeneous sources is more important.
- Improvements in both packaging and wireless communications technologies have enabled an increase in the mobility of users, enabling real-time connectivity for sensor managers and “shooters”.
- Distributed computing is challenged to integrate heterogeneous, largely autonomous/“stove-piped” legacy computer components as part of “real-time” *collaborative* environments.

The assignment of well-understood functionality to automated processes is an important response to these trends. For example, software can take on responsibility for information finding, retrieval and filtering, can personalize human-computer interaction, and can enable tasks to be carried out on these behalf of users whether they are present or not, and do all this with guidance from the users, rather than by direct control. From the perspective of network integration, agents can be applied as interaction entities to mediate differences among legacy network components, while providing a syntactically uniform and semantically consistent intermediary role.

The idea is to build computer surrogates that possess a body of knowledge about the process of information extraction from sensed data and about the user/operator in relation to the process (preferences, decision-making style, etc.) The problem of combining information from a rapidly growing number of distributed, heterogeneous information sources is addressed by organizing these sources into a network of information agents. The goal of each agent is to provide information and expertise on a specific topic (eg, sensor type) by drawing on relevant information from other agents.

A multi-agent system is a loosely coupled network of problem-solver entities that work *collaboratively* to find answers to problems that are beyond the individual capabilities or knowledge of each entity.

The goal is to design flexible systems that allow users to exploit the patterns and knowledge they understand when they are useful without getting in the way when they are not. Such systems have the following abstract properties: (1) they represent and automatically process certain information in formally specified ways (e.g. access and display data from conventional databases); (2) they represent and make it easy for end users to process the same or other information in ways that are not formally specified (e.g. personally configured conventional electronic mail); and (3) they allow the boundary between formal processing by computers and information processing by operators (not skilled programmers) to be radically and easily changed at runtime. They provide end users with powerful building block tools (at the “right” level of abstraction for the domain in question) enabling the creation of a wide range of different applications by progressively modifying a working system.

This is the case with systems of intelligent agents, where agents are seen as entities that emulate mental processes or simulate rational behavior: (1) personal assistant agents, where agents are entities that assist users to perform a task; (2) mobile agents, where entities are able to roam networking environments to fulfill their goals; and (3) information agents, where agents filter and coherently organize unrelated and scattered data.

Intelligent Agent Dimensions

Intelligent agents are software entities that carry out some set of operations on behalf of a user or another program with some degree of independence or autonomy, and in so doing, employ some knowledge or representation of the user's goals or desires. Intelligent agents can then be described in terms of a space defined by the three dimensions of **agency, intelligence, and mobility**.

First, **agency** is the degree of autonomy and authority vested in the agent, and can be measured at least qualitatively by the nature of the interaction between the agent and other entities in the system. The degree of agency is enhanced if an agent represents a user in some way. This is one of the key values of agents. A more advanced agent can interact with other entities such as data, applications, or services. Further advanced agents collaborate and negotiate with other agents.

Next, **intelligence** is the degree of reasoning and learned behavior, ie, the agent's ability to accept the user's statement of goals and carry out the task delegated to it. At a minimum, there is a statement of preferences, perhaps in the form of rules, with an inference engine or other reasoning mechanism to act on these preferences (an Expert or Knowledge Based System). Higher levels of intelligence include a user model or some other form of understanding and reasoning about user demands, and planning the means to achieve this goal. Higher on the intelligence scale are systems that learn and adapt to their environment, both in terms of the user's objectives, and in terms of the resources available to the agent. Such a system might, like a human assistant, discover new relationships, connections, or concepts, and exploit these in anticipating and satisfying user needs.

Last, networked agent applications, such as in NCCT, add a third dimension to the picture called **mobility**. This is the degree to which agents travel through the network. Some agents may be static, either residing on the client machine (to manage a user interface, for instance) or instantiated at the server designed to “intelligently” process raw sensor data at the request of other agents via remote procedure calls (RPC). Mobile scripts may be composed on one machine and shipped to another for execution in a suitably secure environment; in this case, the program travels before execution, so no state data need be attached. Finally, agents may be mobile with state, transported from machine to machine in the middle of execution, and carrying accumulated state data with them. Such agents may be viewed as mobile objects, which travel to agencies at which they can present their credentials and obtain access to services and data managed by the agencies. Agencies may also serve as brokers or matchmakers, bringing together agents

with similar interests and compatible goals, and providing a "meeting point" at which they can interact safely.

The capability for relocating an agent to another agent context, such as the NCCT "smart track" described above, is important because it allows encapsulated, self-adjusting, autonomous data processing operations to move into close network proximity to the data source. Upon arrival, remote agents can perform the necessary data reductions and then either (1) "carry the results with them" to the "next" appropriate agent context, or (2) send the reduced data to the appropriate sending agent via its messaging functionality.

Agent messaging infrastructures provide the regulations that agents follow to communicate and to understand each other, thereby enabling knowledge sharing across the network. These infrastructures deal with the following aspects:

- *Ontologies*: they allow agents to agree about the meaning of concepts. This layer ensures that a term and indeed, even an object or entity, will have a uniform meaning amongst all agents involved in interaction even if different names are used for them. In short, ontology semantically unifies agent communication.
- *Communication Protocols*: they describe languages for agent communication.
- *Communication Infrastructures*: they specify channels for agent communication.
- *Interaction Protocols*: they describe conventions for agent interactions.

In a complex application domain such as TCT, the cooperative decision-making process may require frequent agent-to-agent messages rendering the time-efficiency of the decision-making process dependent upon the efficiency of the messaging layer. If there is a large amount of remotely stored data to be processed (e.g., imagery), there can be significant time saving advantages to transmit the processing method to the location where the data is stored ("function shipping"), than to transfer the data to the client ("data shipping"). Mobile agents could also improve the support for mobile computing, i.e., mobile devices that are only intermittently connected (e.g., cockpit radios). A user might choose to put together a query (a "task" for the agent), connect to the network, launch the agent with the respective task, disconnect from the network, and at a later time retrieve the agent along with the results it has collected.

One might argue that many of these capabilities could be achieved through stored procedures (RPC) or similar mechanisms. Although this may be the case, RPC methods are unable to respond to situations requiring remote processing methods that are not predeterminable, but rather depend the current state of "evidence accumulation" processes and on the real-time data itself. Time-critical targeting multisensor data fusion algorithms under "variable rules of association", like suggested in the NCCT CONOPS, may present such a challenge. Furthermore, RPC requires that the server "owners" statically install the server procedure components in advance, while mobile agents are dynamically installed by the application itself. This can be a significant advantage in the adaptation of new functionality (e.g., TCTF) in existing ISR networks whose servers are owned and operated by independent agencies.

Even so, mobility may not be the panacea just yet. It brings a host of security, privacy, and management challenges. Initially, applications will likely be built around static agents; mobility will appear gradually over time, as the infrastructure for agents matures.

Conceptual Multi-sensor Fusion Agent Architecture

Figure 1 below depicts a logical view of a multi-agent system designed to support TCT operations by supporting the automation of information extraction from multiple, distributed heterogeneous sensor sources. Unique agents are envisioned to represent the expertise at each sensor source and correlation/fusion processing level: raw data processing/filtering, registration, correlation, change detection and multi-sensor, object-level fusion. Facilitator and mediation agents are used to coordinate communication among agent classes and archived/historical track and context data sources, following rules of interaction established by TCT operations personnel.

In this architecture the agent hierarchy can be described as follows:

Filtering Agents: based on a pre-specified goal architecture, raw data is filtered/parsed into subsets at source; intelligent routing for early fusion

Facilitator Agents: provide intelligent communication and coordination services supporting effective agent-to-agent interactions

Data Cleaning/Formatting Agents: routers feed high value data into a distributed data cleaning architecture based on pre-specified event interests; render data into common knowledge architecture (e.g., temporal/spatial registration)

Data Aggregation/Data Mining/Event Classification Agents: 1st line of early threat identification (change detection/template association, etc.); agents may interact with relevant data cleaning agents, obtain real-time updates, and classify threats

Mediation/Mining/Fusion Agents: manage lower level agents, obtain classification data (e.g., templates, signatures), and apply data fusion, data mining algorithms to database and associated context and sensor-derived information

User Interface Agents: eliminate redundancies and non-critical data collection/analysis; focus resource utilization on high priority targets; may trigger automated rule-based responses such as new sensor tasking or initiate countermeasures automatically based on pre-specified goals and data quality/fusion confidence levels

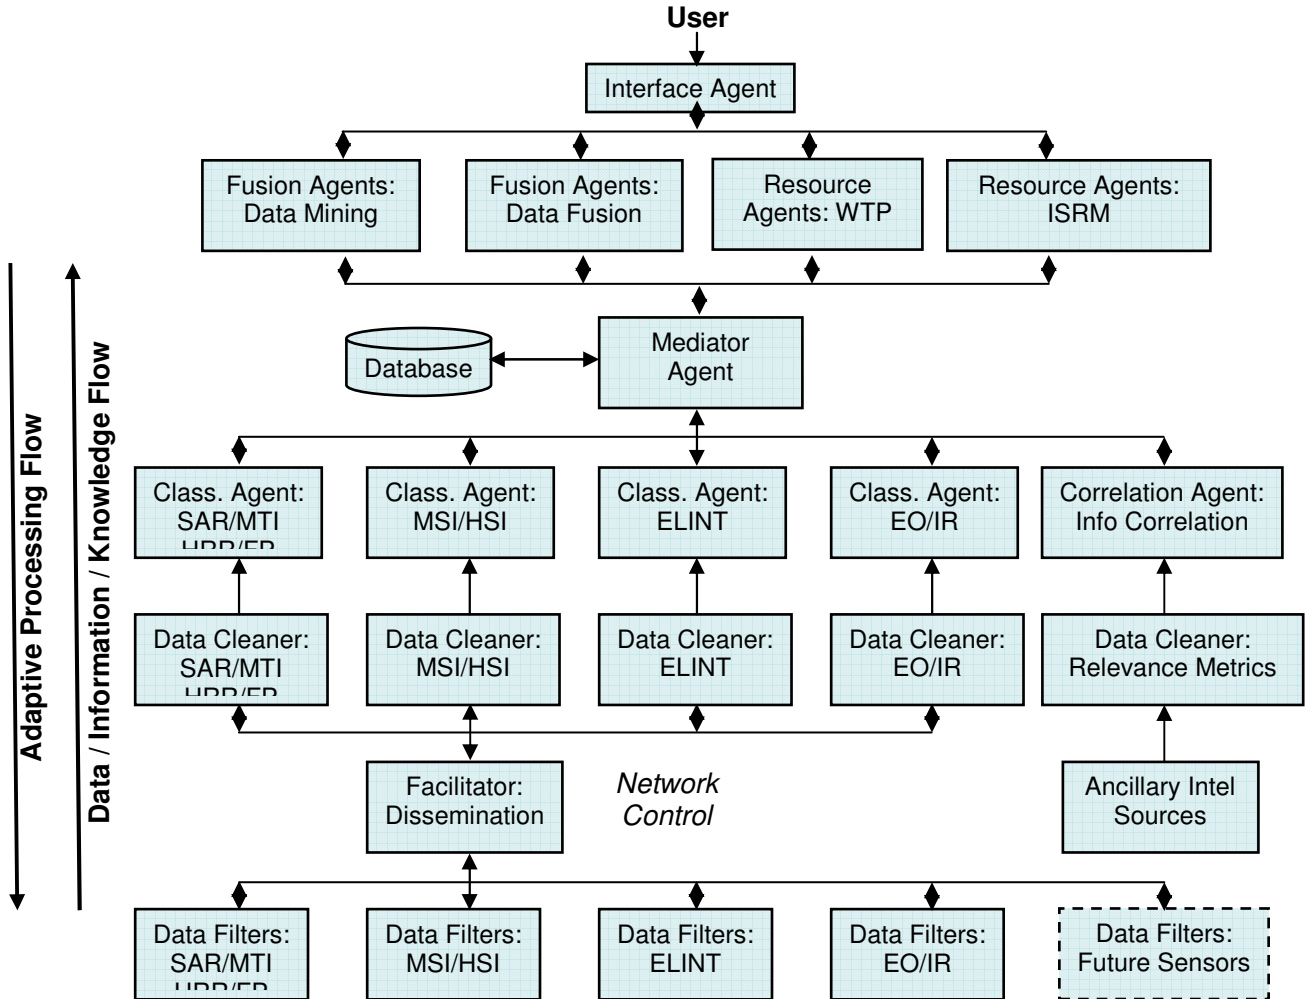


Figure 1: Logical View of a layered, intelligent multisensor processing agent architecture

Technical Challenges

Data fusion from distributed, heterogeneous data sources

Developing a knowledge-building framework for a network-centric environment is a large undertaking, but quite doable (current and past DARPA agent framework, ontological, ATR and dynamic multi-sensor fusion and tasking initiatives in the IW domain will be leveraged to the extent feasible). The two principal issues are ensuring software design scalability and runtime scalability. To satisfy both scalability issues, we must develop a data fusion framework that addresses all of the concerns raised in preceding sections, as well as others that arise. We must carefully choose an agent framework with the proper support for robust time-critical decision-making and efficient communications layering.

Effectively utilizing domain information/knowledge

The two principal challenges/tasks are (1) scalable management of large volumes of data,

and (2) collecting ISR-related domain information with proper strategies for data reduction and interpretation. Task (1) is extensive, but it has been done successfully in many environments. Task (2) is harder to estimate, because it requires access to human resources that can enumerate, as well as verify, the many domain rules and conditions, for example, military techniques, tactics, and procedures (TTP).

Integrated semantic/spatial database algorithms

This task is quite challenging, but the rapid growth in XML technologies enhances our ability to structure the collected knowledge, making it more amenable for efficient algorithmic analysis/processing.

Secure operations

Many agent frameworks provide no (or inadequate) agent-level security. For now, secure operations can be achieved at the level of the data fusion framework with a combination of access list-based control of network and runtime processes and encrypted over-the-wire communications. For the long term, the viability of software encryption strategies is suspect for any complex environment, due to runtime scalability concerns, and hardware encryption may be necessary.

In agent-based, network-centric IW environments there are two principal security concerns:

Secure over-the-wire data transport among/across network nodes for

- Data/object persistence via a networked database
- Inter-agent messaging and data exchange
- Agent mobility among Java processes

Secure external access to distributed software components including

- Static data residing in database systems
- Dynamic mobile agents executing within and across distributed Java processes

Category (1) security requires either (a) guaranteed-safe encryption strategies over a clear (open) wire (or radio/satellite transmission), or (b) guaranteed-safe protection of the wire or transmission (optimistically) allowing unencrypted operation. Category (2) security, especially for mobility concerns in Java processes, is questionable at present. Several Java-based mobile agent frameworks have provided distributed access list-based security for agents, but this issue remains unsolved due to Java's security model, which may be revamped in the next 2-3 years.

Scalability

In the military domain where electronic information warfare operations tasks are network-centric, compute-intensive, and network-complex, scalability issues become

very important. In general, monolithic (even though object-oriented) software architectures are insufficient.

For large-scale, network-centric applications, scalability issues include:

Avoiding poor execution performance due to localized or ineffective distributed architectures

- Per-host CPU bottlenecks
- Network traffic bottlenecks
- Failure to transport and/or process data efficiently
- Data loss from input devices

Avoiding ongoing design complexities due to ineffective software infrastructure

- Monolithic and/or complex OO designs that are not easily extensible
- Domain-incorrect, hard-coded, distributed component architecture

Avoiding the inability to grow the system due to ineffective domain modeling

- "Dumb" domain solutions that scale linearly (at best) as the domain grows
- Improper factoring/separation of the domain implementation from the application core

Agent frameworks, on the other hand, provide mechanisms for avoiding/conquering these architectural issues. For example, agent-based distributed architectures provide mechanisms such as mobility for dynamically adjusting to compute-intensive and network-complex execution environments. Agent frameworks enable software design and development strategies that handle complex domains using extensive encapsulation and modular (pluggable) components. Agent-based solutions can avoid the linear-scalability issues by using intelligent planning strategies to adapt to a dynamically changing domain: sporadic sensor failure, newly active domain components, redistribution of data processing tasks based on device proximity and current data volume, and so on. Unfortunately these advantages are accompanied by technological risks that must be addressed on a case-by-case basis.

Above caveats notwithstanding, the risk of failure of our ISR systems to address the threats posed by time-critical targets clearly makes a strong case for the measured application of powerful, proven agent technology capabilities.