

**REPORT DOCUMENTATION PAGE***Form Approved*  
**OMB No. 0704-0188**

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Service, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.

**PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.****1. REPORT DATE (DD-MM-YYYY)**  
OCTOBER 2009**2. REPORT TYPE**  
Conference Paper Postprint**3. DATES COVERED (From - To)**  
October 2008 – August 2009**4. TITLE AND SUBTITLE**  
N-CET: NETWORK – CENTRIC EXPLOITATION AND TRACKING**5a. CONTRACT NUMBER**  
IN HOUSE**5b. GRANT NUMBER**  
N/A**5c. PROGRAM ELEMENT NUMBER**  
N/A**6. AUTHOR(S)**  
James Metzler, Mark Linderman, Lee Seversky**5d. PROJECT NUMBER**  
558S**5e. TASK NUMBER**  
IH**5f. WORK UNIT NUMBER**  
NC**7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)**AFRL/RISD  
525 Brooks Road  
Rome NY 13441-4505**8. PERFORMING ORGANIZATION  
REPORT NUMBER**

N/A

**9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)**AFRL/RISD  
525 Brooks Road  
Rome NY 13441-4505**10. SPONSOR/MONITOR'S ACRONYM(S)**  
N/A**11. SPONSORING/MONITORING  
AGENCY REPORT NUMBER**  
AFRL-RI-RS-TP-2009-52**12. DISTRIBUTION AVAILABILITY STATEMENT***Approved for public release; distribution unlimited PA# 88ABW-2009-3733 Date Cleared: 19-August-2009***13. SUPPLEMENTARY NOTES**

This paper was published in the proceedings of MILCOM 2009 Conference, Boston World Trade Center, Boston, MA; 19-21 October 2009. This is a work of the United States Government and is not subject to copyright protection in the United States.

**14. ABSTRACT**

The Network-Centric Exploitation and Tracking (N-CET) program is a research effort to enhance intelligence exploitation in a tactical environment by cross-cueing sensors and fusing data from on-board sources with processed information from off-board platforms and sharing the resulting products in a net-centric manner. At the core of N-CET are information management services that decouple data producers and consumers, allowing reconfiguration to suit mission needs. Network centric algorithms utilize the availability of information from both homogeneous and complementary on-board and off-board sensors. Organic capabilities facilitate the extraction of actionable information from high bandwidth sensor data and ensure the necessary information arrives at other platforms and users in a timely manner. This paper provides an overview of the NCET architecture and the sensors and algorithms currently implemented upon it. The extent to which such algorithms are enhanced in a network-centric environment is discussed and the challenges of managing the resulting dynamic information space in a tactical publish/subscribe/query model are presented.

**15. SUBJECT TERMS**

Network-Centric, Tactical Information Management, High Performance Computing, Multi-Int Fusion

**16. SECURITY CLASSIFICATION OF:****a. REPORT**  
U**b. ABSTRACT**  
U**c. THIS PAGE**  
U**17. LIMITATION OF  
ABSTRACT**

UU

**18. NUMBER  
OF PAGES**

8

**19a. NAME OF RESPONSIBLE PERSON**

James M. Metzler

**19b. TELEPHONE NUMBER (Include area code)**

N/A

# N-CET: NETWORK-CENTRIC EXPLOITATION AND TRACKING

James M. Metzler, Mark H. Linderman, and Lee M. Seversky

Air Force Research Laboratory

Information Directorate

Rome, NY, 13441

james.metzler@rl.af.mil, mark.linderman@rl.af.mil, lee.seversky@rl.af.mil

**Abstract**—The Network-Centric Exploitation and Tracking (N-CET) program is a research effort to enhance intelligence exploitation in a tactical environment by cross-cueing sensors and fusing data from on-board sources with processed information from off-board platforms and sharing the resulting products in a net-centric manner. At the core of N-CET are information management services that decouple data producers and consumers, allowing reconfiguration to suit mission needs. Network-centric algorithms utilize the availability of information from both homogeneous and complementary on-board and off-board sensors. Organic capabilities facilitate the extraction of actionable information from high bandwidth sensor data and ensure the necessary information arrives at other platforms and users in a timely manner. This paper provides an overview of the N-CET architecture and the sensors and algorithms currently implemented upon it. The extent to which such algorithms are enhanced in a network-centric environment is discussed and the challenges of managing the resulting dynamic information space in a tactical publish/subscribe/query model are presented.

## I. INTRODUCTION

This paper presents a prototype implementation of an architecture and algorithms to support net-centric exploitation and tracking in a tactical environment suitable for employment within a tactical pod or unmanned aerial vehicle. Tactical exploitation seeks to maximize the value of sensor data to achieve a set of missions. This involves the control of the sensor, the processing of its outputs and potentially the combining (fusion) of those outputs with other information to add context or to differentiate “interesting” information from “uninteresting” information. The tactical environment is characterized by unpredictability, scarce resources, and often a wealth of low quality sensor data that may be low resolution, outdated, or simply looking at the wrong things. Effective tactical exploitation must address each of these deficiencies.

If one’s objective is surveillance or intelligence gathering, tactical assets that operate in harm’s way enjoy proximity afforded to few other platforms. Different adversaries and different stages of conflict may require different sensors and exploitation approaches. Such diverse sensor requirements cannot be achieved with a fixed sensor suite. Furthermore, collections of forward-deployed sensors may need to communicate directly with one another (perhaps over low-bandwidth low-probability-of-intercept channels) to coordinate sensor tasking, perform real-time exploitation, and to prosecute targets that may only become evident as a result of the exploitation. In the early phases of a campaign, it may be necessary to combine reconnaissance and prosecution into a single mission

because loss of the element of surprise may imperil subsequent missions.

The idea that a sensor suite should be matched to the mission was one of the motivators of the concept that became the basis for N-CET. Initially termed PEAPod (Programmable, Extensible Architecture for Pods), the concept was conceived as an architecture for a tactical pod (similar to a Litening-AT targeting pod [14]) that could accommodate a variety of multiple front-end sensors and be placed on an aircraft to support specific missions. The pod would feature sufficient computational power to perform real-time sensor data processing and exploitation and be able to communicate with other pods in a net-centric manner.

Sensor data must be timely, of adequate resolution, and regarding targets of interest to be valuable. Given the limited bandwidth available in a tactical environment, if a sensor can collect data faster than it can be communicated, the sensor must record the data, process the data to reduce its size, selectively transmit the data, and/or discard it. Even though high-capacity storage may allow sensor data to be archived as it is collected, if it cannot be off-loaded until the completion of the mission, it may no longer be timely.

### A. Net-Centricity

Net-centricity is not simply about building communication networks; it is about collections of communicating (i.e., networked) entities synergistically collaborating to accomplish their respective missions better than would be possible alone. Ideally, a collection of net-centric platforms could comprise a cybernetic system in the original sense of the word that implicitly recognizes the essential roles of coordination, regulation, and control [3]. These three roles embody the idea of net-centricity and are fundamental to the design of the N-CET system.

N-CET coordinates the interaction of sensors, processing, and information sharing within a platform (and between platforms) to achieve maximal exploitation value. This may involve cuing one sensor off of another on-board sensor, or commanding electro-optical (EO) sensors on several platforms to capture images at a given point at the same time.

N-CET regulates the flow of information and processing to maximize mission effectiveness given finite resources. N-CET incorporates prioritized transmission of data over the radio network. The initial implementation presented here is simplified by adequate provisioning of computational resources for the tasks required. However, more stressing sensors and

exploitation algorithms are anticipated in follow-on efforts that will demand scheduling and prioritization of processing. In addition, context-based or demand-based exfiltration will play a larger role.

The heart of N-CET is the control software that performs sensor management and control (described in Section II-B3). Each platform operates under local control; managing collection modes for on-board sensors. It subscribes to internal and external data and commands, prioritizes tasks, and publishes commands to component subsystems. These in turn control the sensors to collect appropriate data that is published for subsequent processing.

### B. Motivation

The motivation of N-CET is based upon three observations: 1) there is no substitute for good sensor positioning, 2) modifying aircraft avionics is expensive and slow, and 3) most sensor systems are stove-piped and have unique architectures. Therefore N-CET must consider survivability to operate close-in, ease of fielding, and reconfigurability.

1) *Survivability*: N-CET does not directly address the survivability of a platform carrying it, but its net-centric attributes should assume that other nodes will come and go, either temporarily or permanently. Resilience (or fault tolerance) in such an environment may take many forms. At the networking layer, end-point discovery and maintenance are important. At the information management layer, all applications on an N-CET node rely upon publish/subscribe to communicate with off-board entities, so publish and subscribe-based protocols must be tolerant of lost peers. A focus on the connection level, however, skirts the real issue: implicit reliance upon external entities for tasking or information.

2) *Ease of fielding*: One of the goals of N-CET is for the pod to rely upon the hosting platform as little as possible: minimally for transportation and power alone. Achieving net-centricity with minimal reliance on the platform requires that the N-CET node be largely self-sufficient; in particular, it may need organic communication abilities. Whether organic or host communications is used, it is essential that the node carefully manage information that is put into the pipe.

3) *Reconfigurability*: While there are many types of sensors of interest, all of them sense in the “analog” domain, and most of them digitize the data for subsequent processing. In the chain from sensing to digitization to exploitation, where is the appropriate boundary between the sensor front-end and the back-end? A goal of N-CET is to provide a considerable amount of general-purpose computational capability to allow it to subsume much of what has traditionally been considered part of the sensor for two reasons: 1) simplify the sensor, and 2) avoid losing data that might be unimportant in one context but important in another. Because sensors, like most hardware, are built with specific purposes in mind, it is tempting (indeed generally required) to build to those specific purposes. Also, like most hardware, it cannot be upgraded easily and certainly not quickly. N-CET anticipates a much more dynamic environment where suitable software is loaded

on a pod to match the specific sensor configuration and the mission.

### C. Design Considerations

The motivations described have been translated into design considerations for N-CET. It must be flexible, operate at the highest possible fidelity, save all information possible, and perform subject to policies that may be modified to suit mission needs.

Flexibility, the ability to adapt, is essential to the N-CET concept. From the design perspective, this means that the amount of hardcoded assumptions must be minimized. Architectural components such as the Joint Battlespace Infosphere (JBI) [8] information management services are inherently flexible. Others, such as the node controller, are not as flexible as ultimately intended to be. However, the use of publish/subscribe to route sensor data to available processors is a technique that provides flexibility as it did in RTMCARM [10] and Swathbuckler [9]. The use of general-purpose processors leaves open many algorithmic options.

N-CET's ability to archive data makes new applications possible. For example, another platform may detect interesting activity but by the time it gets to this node, the event may be several seconds old. Rather than hoping for the recurrence of a similar event, the archives may be searched for data from the time of the event, and additional exploitation may be possible. Of course, storage may be limited, requiring a lifecycle management policy for the data, but this is only one of several policies necessary for an N-CET system.

## II. SYSTEM DESCRIPTION

An N-CET node consists of hardware and network infrastructure, an instance of the 100X JBI [7] Information Management platform, the core clients, and supplemental exploitation and fusion clients. Currently, three mobile ground nodes are instantiated for development and field testing.

### A. Hardware Architecture

The hardware design of the initial N-CET prototype was based on the requirements for a flexible platform on which to host the various hardware components that make up the system, the majority of which are rack-mountable COTS. Because N-CET is targeted to be hosted inside an airborne pod or UAV, in future work will involve packaging the architecture for embedded environments.

The N-CET node comprises several hardware components and software modules. Figure 1 depicts the major components (power supplies and support components are omitted). On the left, the sensors, a gimbaled high-definition video camera and a directional RF antenna, are shown. The “head-node” is shown at center. It hosts the 100X JBI and the computationally less demanding components. Shown around the head-node are different pieces of hardware including the Sony PlayStation®3 (PS3) nodes used for computationally demanding tasks - principally video processing. The PS3 was chosen as a cost-effective Cell BE processor [6], with six available Synergistic




Fig. 1. N-CET node architecture

Processing Elements (SPEs), each capable of 25 GFLOPS [5]. The hardware components are connected with a 10 Gb/s Ethernet switch. The software components are shown as rounded boxes on the hardware they are implemented on.

The external communication and connectivity portion of each N-CET node consist primarily of several COTS IP (Internet protocol) radios including Freewave 900MHz and 2.4GHz IP radios and Microhard 900MHz IP radios. The Joint Capability for Airborne Networking (JCAN) [15] ties one or more of these radios together on each N-CET node and routes traffic through other nodes to compensate for link-loss that often occurs in airborne networks. For example, in the N-CET system with three nodes, when connectivity is lost between two nodes, traffic is transparently routed through the third node.

The JCAN manages the airborne network that optimizes delivery over one or more radios. It also performs queue management and prioritization. Working with JCAN, the 100X JBI assigns a priority class to information based upon its type. In this way, for example, smaller high-priority messages may preempt or be sent simultaneously with larger messages.

### B. Core Components

While designed to be reconfigurable to specific missions and requirements, each N-CET node requires several core components that permit information management, federation of data among nodes, fusion of remote information with local information, and sensor control.

1) *Information Management:* The 100X JBI handles all information transfer within an N-CET node as well as among N-CET nodes. The 100X JBI is an implementation of the AFRL Information Management (IM) Core Services which provide publish, subscribe and query capabilities. In the pub/sub/query paradigm, information is encapsulated as a Managed Information Object (MIO) consisting of a payload and the metadata describing it.

The 100X JBI uses a client-server model in which the clients (being producers, consumers, or both) connect to one or more IM platforms (servers). Upon publication of an MIO

by a client, the 100X JBI brokers and disseminates the MIO to clients based on registered subscriptions and their corresponding predicates. The IM platform also provides archival capabilities, permitting clients to query MIOs based on type and filtered by predicates.

The 100X JBI server uses the YFILTER [4] XML predicate evaluation algorithm to enable the high throughput necessary to support the various N-CET clients. The server utilizes Oracle's DBXML in-process libraries for archiving data products when they are published to make them available for queries. The server also uses some Linux OS features for increased speed, including shared memory and direct disk writing that allow the JBI to sustain archival rates up to 240MB/sec to 4 SATA drives. This would support a frame rate of 50 HD frames per second. Dissemination rates have been tested up to 400MB/sec to multiple clients attached to the 10Gb/1Gb Ethernet switch.

2) *Data Federation:* Net-centric systems require information sharing between all nodes, both airborne and ground based. N-CET manages information locally on each node via the JBI platform, and in many instances, it is necessary to share information between nodes, i.e. platforms. To connect to JBIs on other nodes, a forwarding client is implemented on each node to share necessary MIOs with other JBIs. As a result, a client is able to subscribe only to its local JBI and also receive MIOs generated on other nodes. To accomplish this information sharing, the Forwarder establishes and manages connections to the JBI on any neighboring node that is within link range. The Forwarder subscribes to the necessary MIOs on its local JBI and upon receipt of an object publishes it to all connected JBIs. Threading and timeouts are used to properly manage connections, and predicates are used to avoid cyclical publishing of objects, i.e. the Forwarder only subscribes to and therefore publishes MIOs generated from its local platform.

3) *Controller:* The control of each N-CET node is performed locally through the direction of the Controller client. The Controller is responsible for correlating information from the exploitation clients, commanding the available sensors, and responding to user commands.

The Controller subscribes to system commands published by a user as ncetControl MIOs that are either *modes* or *tasks*. A mode is a persistent system command while a task is singular system command that is serviced and then the previously interrupted mode is resumed. The following modes have been implemented for the current sensors and algorithms.

- **trackTarget** The EO sensor is cued on a geo-location or bearing and all moving objects within the FOV are detected.
- **trackSurvey** The EO sensor repeatedly pans across a series of positions to capture a sequence of frames, and all moving objects within the FOV at each position are detected.
- **stop** All video processing halts.

The following tasks have been implemented for the current sensors and algorithms.

- **glance** The EO sensor is cued to a geo-location or bearing and captures an image.
- **panoramic** The EO sensor captures an image at a series of positions to create a panoramic view of the sensor Field of Regard (FOR).

In order to account for the asynchronous operation of N-CET, other clients do not subscribe to ncetControl MIOs, rather, the control mode is included in the resulting MIO's metadata. Clients may ignore a particular MIO based on the mode by which it was generated by specifying a subscription predicate. For example, client's in the video processing stream subscribe to videoFrame MIOs that were generated as a result of a trackTarget or trackSurvey command, ensuring that a task, such as a glance, that interrupted a sequence of video frames, does not affect the algorithm.

### C. Sensors and Algorithms

While the 100X JBI, Controller, and Forwarder are core components required on an N-CET node, the exploitation and fusion algorithms and their accompanying sensors are designed to be optional and interchangeable. This permits the rapid reconfiguration of an N-CET node to meet the mission needs, as specified in the design requirements. The current N-CET instantiation utilizes an ELINT sensor and an EO sensor. Several exploitation and fusion clients are implemented that make use of these sensors, and multiple others are currently being developed and integrated.

The ELINT sensor is an RF detection and direction-finding sensor capable of intercepting an audio transmission from a Family Radio Service (FRS) hand-held radio and estimating the emitter's line of bearing (LOB) relative to the sensor. An Audio Exploitation client interfaces directly with the sensor, parses the proprietary messages being produced by the sensor, and publishes corresponding rfIntercept MIOs at the start of each transmission, intermittently during the transmission (to update a changing LOB), and at the end of the transmission. The rfIntercept includes transmission properties such as frequency, a unique identifier, and several measures of the quality of the intercept.

Speaker identification capabilities are currently being integrated into the Audio Exploitation client. An intercepted audio signal is streamed from the ELINT sensor to a speaker identification algorithm [17] implemented on a PS3. Speaker identification is made from a closed set database of speakers for which audio samples have previously been used to train a speaker model from features that parameterize the vocal tract of the speaker over short time segments. Upon interception of an audio signal, features are extracted and tested against each available speaker in the model and a decision is presented to the user as well as measures of certainty.

Social Network Analysis (SNA) [16] capabilities are also being integrated that will be used to determine the importance of the identified speaker, as well as identify and establish any groups the speaker is a member of or relationship the speaker may be a part of. The speaker's importance is measured by a Key Player Algorithm that assigns a rank to an entity in a

social network. The social network is stored in a database that will be updated as activity occurs, e.g., as multiple speakers are identified to be speaking on the same frequency within the same time period.

The rfIntercept MIO is federated to other nodes via the Forwarder. The Controller subscribes to the rfIntercept MIO type on its local JBI, thus receiving those generated both on-board and by remote nodes. Using timestamps and frequency data, rfIntercepts from different nodes are correlated. The node location and LOB of each rfIntercept is used to compute the geo-location of the emitter based on Brown's Least Squares Triangulation method [13]. This geo-location is published as an rfTarget MIO. The Controller autonomously cues the EO sensor onto these emitters, issuing a glance ncetControl command and restricting revisit rates for the same transmission based on time and difference in emitter position.

The Controller also subscribes to ncetControl MIOs issued by a user, such as an imagery request for a geographic location or a panoramic refresh. For videoFrame MIOs generated by these singular tasks, a Compression Proxy client compresses the raw image data using JPEG to reduce it to a size suitable for transmission to a ground user.

Frame capture and the EO sensor are controlled by the Video Capture client. The Video Capture client decomposes high-level instructions from the Controller into operations that control the camera hardware, the capture parameters, the meta-data generation for captured frames, and the publication of frames. Communication between the Video Capture client and the Controller is conducted through an API to allow alternate clients to command the camera. This allows reconfigurability of the system when a different control client is desired and establishes a common interface facilitating interchangeability of the EO sensor.

The EO sensor currently used is a SONY EV1-HD1 video-conferencing camera [1]. This camera has a resolution of 1920 x 1080 pixels and captures at a rate of 25 frames per second. Each uncompressed frame (YUV format) is read from a capture card on the head-node and published as a videoFrame MIO in both color and grayscale at a combined rate of 156 MB/s. The 100X JBI disseminates each frame type to the clients subscribing to it, and archives the data for future analysis. Video processing algorithms are used to extract the important information from this high resolution video, making possible high fidelity intelligence collection without the need for high bandwidth datalinks. Video processing involves three steps: 1) image registration, 2) motion estimation and object extraction, and 3) object tracking.

In order for the video tracking algorithms to function in an environment where the camera is moving, as in the case of a mobile node, a sequence of frames must be registered to a common coordinate system. That is, the translation, rotation and scaling of stationary objects that results from the changing position of the imaging sensor must be removed from successive frames so that moving objects may be detected within that sequence of frames.

Frame-to-frame registration is currently being implemented

on an NVIDIA Tesla C870 Computing Processor [12], a General-Purpose Graphics Processing Unit (GPGPU). The Registration client resides on the head-node and subscribes to grayscale videoFrame MIOs published by the Video Capture client. The client communicates with the Tesla C870 via a PCI-e x8 expansion card, and registration is performed on the GPGPU. A registeredFrame MIO is published for use by those clients requiring registered video frames.

Next, moving objects are detected in a sequence of frames by the Flux Tensor client, named for its use of the flux tensor algorithm [2]. The Flux Tensor subscribes to grayscale videoFrame MIOs and publishes blobList MIOs identifying the centroid, bounding area, and velocity of moving objects detected within a frame. The flux tensor algorithm was ported to the Cell architecture on a PS3 and redesigned to take full advantage of the specific architectural features of the SPE memory layout and communication patterns.

Two object tracking algorithms are currently being developed to track moving objects within a sequence of video frames, highlighting the ability to interchange clients in the system, and allowing each algorithm to be evaluated for strengths and weaknesses in varying environments. The first of the two, the Motion-Estimation Based Tracking (MEBT) client subscribes to blobList MIOs and processes these lists to identify track segments. The MEBT client uses an object association and multi-hypothesis tracking algorithm [2]. The algorithm supports near-line tracking and due to its low computation requirements, is implemented on the head-node in MATLAB®.

The second algorithm, Tracking Evasive Non-linear Targets (TENT), is an AFRL in-house program that focuses on tracking targets in a variety of domains including video and Moving Target Indicator (MTI) radar. In addition to blobList MIOs, the TENT client subscribes to color videoFrame MIOs so that feature tracking may be performed to augment moving object tracking. TENT tracks the objects identified by the Flux Tensor and uses the object properties to extract features of the object from the color video frame. Whereas objects disappear to the Flux Tensor when no longer in motion, feature detection allows TENT to maintain track on an stationary object that had been previously moving.

Both tracking algorithms generate a track MIO consisting of target identifiers and track information. This permits either or both of the tracking clients to be used depending on the mission requirements while being transparent to the subscribers of the information. This interchangeability is a design requirement for all clients in the N-CET system.

The final step in video processing is extracting the images of targets from the high-definition video frames. This reduces the bandwidth required to deliver high resolution imagery to the warfighter.

The Extraction client subscribes to color videoFrame and blobList MIOs, and using the timestamp in each object's metadata, extracts the image chip of each blob from the corresponding video frame. The client publishes the blob images of each frame as an imageChips MIO which includes




Fig. 2. Console user interface client

the metadata necessary for a subscriber to display the image chips relative to the position of the camera.

A user interface, the Console client, has been developed to display information generated by individual N-CET platforms and provide an interface for commanding the system from within a single environment for testing and evaluation purposes. The Console may subscribe to all MIO types or a subset depending on the user's requirements and the bandwidth available. The console can connect from any point within the network, however, it is typically remotely located and the bandwidth between it and mobile nodes is limited.

There are two main visualization components. A traditional bird's eye view visualization displays information that is in the geographical coordinate system, such as the position of the N-CET nodes and geo-located targets. The visualization of the information is overlaid on a geographic backdrop of the region of interest. Terrain data and satellite imagery is used to create a three-dimensional environment for the user to interact with. A second visualization component, the panoramic view, is used to display platform centric information represented in the camera coordinate system, such as the products of video tracking algorithms.

The Console is shown in Figure 2, with the panoramic view as a separate floating window at the top of the figure. In the birds eye view, two N-CET nodes are displayed and labeled. Also shown are rfIntercept MIOs published by each node, and the resulting rfTarget. The rfTargets for the same transmission are plotted over time to create a track of the emitter. The Console also displays User Target locations (center right) selected by the user and visualizes the images received for that location.

A User Target location is also shown on the map (center right), visualizing an imagery request at that location by a user. Captured imagery and other MIO metadata is shown in the MIO Browser on the right side of the bird's eye view.

The panoramic view is used to display information in the camera coordinate system, as defined by a focal length,




Fig. 3. Stockbridge Experiment Layout

azimuth, elevation, rotation, and imaging array size. This visualization provides a platform centric view and is used to fuse image-based information into a single common reference frame. A panoramic command is issued to create a background for the current FOR for each node's EO sensor, as shown in the top left of Figure 2. The products of the video tracking algorithms, such as tracks and imageChips, are described by these properties in their metadata, allowing them to be overlaid on the background at the proper location. For imageChips of the moving objects extracted from the HD video frames, the images are displayed on the screen and replaced at the arrival of a new imageChips MIO.

### III. LESSONS LEARNED

#### A. Experimentation

In addition to laboratory testing, outdoor experimentation was conducted in June 2009 at an AFRL test site in Stockbridge, NY. The Stockbridge site was originally built to evaluate the radar cross-section of aircraft and provides towers upon which to mount the N-CET sensors and terrain for both vehicle and dismount targets. Two N-CET nodes, ncet1 and ncet2, were fielded as shown in Figure 3. While the sensors for ncet1 were mounted at ground level, the sensors for ncet2 were mounted approximately 80 feet above ground on a tower providing downward-looking viewing angles. The tree lines shown in Figure 3 limited the visible range of the sensors and did not allow multiple viewpoints to any one area, however this created an interesting scenario of passing visual tracking from one node to the other as targets moved along the routes shown in Figure 3.

The objective of the June experiment was to determine if the N-CET system could detect multiple RF emitters (targets), identify the targets (visually and, in the case of audio RF communications, by name and social associations), locate the targets, and track the targets in a sequence of video frames. To supplement the ELINT sensor and demonstrate the flexibility of the system, a Ground Moving Target Indicator (GMTI)



Fig. 4. Example still Image from imageChips sequence

sensor was also added to the system. The FRS two-way radios used as emitters also transmit GPS data which provides ground truth for the targets. A GMTI simulator used this data in real time to publish GMTI detections for the targets and a generic GMTI tracker was used to create persistent tracks and publish groundTrack MIOs. In addition to RF detections, the Controller subscribed to the groundTrack MIOs, and in the same manner as for an RF target, cross-cued the EO sensor onto the GMTI target.

The N-CET system was successful in cross-cueing the EO sensor to visually identify targets. The groundTrack MIOs were displayed on the geospatial view of the Console and the images captured for each track were displayed in the MIO Browser. The RF direction finding sensor was less accurate, however when quality rfIntercepts were exchanged between nodes, an rfTarget was computed and published along with the associated imagery.

Motion estimation and object extraction were also tested for stationary cameras. Figure 4 shows the extracted images of moving objects contained in a single imageChip MIO (from a sequence) overlayed onto a static compressed background image. In this scenario, two walking subjects rendezvous with two vehicles and get inside. At the point in the sequence shown here, the subjects have entered the vehicles, however they still appear in the static background image at their location when the background was captured. The complete sequence lasts approximately 100 seconds and the raw color HD video frames consists of 7.4 GigaBytes of data. The resulting processed data, the imageChips publications, are approximately 4.4 MegaBytes for the same 100 second scenario. In this case, the data downlink requirement has been decreased from 74 MBps to 44 Kbps while maintaining a high resolution image of the targets of interest.

#### B. Information Management

A major effort in the N-CET program has been the structuring of the information that is generated and used by the clients in the system. During the initial development of the system, the integrators of each technology created the schema's that defined the metadata for the MIO that specific technology would generate. Although the schemas had been formalized

well in advance, problems occurred during integration that were the result of differences in the interpretation of the metadata, both elements and their values. The Information Management Services are intended to provide the tools for information users to better deliver and receive information and the user must decide how to structure the information so that it is valuable. This requires an integrator with a true system wide perspective who understands how clients are making use of the information and the assumptions and assertions made when processing and producing the information. The 100X JBI's requirement for strictly structured metadata defined by an XML Schema Definition (XSD) ensures that information adheres to these formats once established.

Inherent in this system wide perspective is the ability to recognize where pedigree information is required in order to trace the lineage of certain MIOs. For example, when constructing the operating picture for the user, the Console visualizes the relationship between MIOs, such as the rfIntercept used to create an rfTarget and any resulting image captured by the EO sensor. In addition, multiple rfTarget MIOs generated for a transmission are visualized as a track. To accomplish this, pedigree information has been added to the necessary schemas. For example, the rfTarget MIO has sourceInfo elements identifying the source, transmission identifier and intercept identifier of the rfIntercept MIOs used in the triangulation. Likewise, videoFrame MIOs are given a requestID element to identify the intercept, target or user command that initiated the image capture.

#### IV. FUTURE WORK

The N-CET architecture and initial system implementation currently provides a platform for experimentation of existing and new net-centric capable sensors and algorithms. However, a number of challenges and open problems associated with both the system implementation and the architecture of the N-CET concept still remain. Future work should investigate alternative system implementations with pod form factors such as the LITENING (AN/AAQ-28) [14] or LANTIRN (AN/AAQ-14) [11] in order to begin transitioning from a ground-based to an air-based experimentation platform. In addition, there exists a tradeoff between the decoupled design of the publish/subscribe/query model and the need for more global design and control of information management services for tactical dynamic environments. The current underlying information management services decouple data producers and consumers. As a result, the current N-CET implementation must couple consumers and producers in a local fashion manually in order to satisfy specific static objectives and resource constraints. On the one hand, the decoupled design allows for a modular information management system. On the other hand, more global control and management is needed for the better understanding and utilization of the net-centric resources. Future work should investigate the interplay between net-centric components and develop general strategies for managing information which may include both local and more global techniques.

#### V. CONCLUSION

The core design requirements for N-CET have been accomplished. While interfaces and control process will need to be created for different sensors, N-CET is capable of accommodating any sensor that can be connected via serial, Ethernet, or coaxial cable (video). The JCAN and radios provide organic communication to other nodes and ground users. An N-CET node is not currently packaged in a size, weight and power envelope conducive to a pod form factor or UAV, however focus will shift to this requirement as development continues. Several exploitation and fusion algorithms have been implemented in the N-CET system and more will be incorporated. The components have been designed in a net-centric manner, taking advantage of the variety of sources and modalities of information, and to allow reconfiguration of the clients to suit mission needs.

#### REFERENCES

- [1] Sony Broadcast and Business Solutions Company. Overview available at <http://bssc.sel.sony.com/BroadcastandBusiness/markets/10005/EVIHD1.shtml>.
- [2] Filiz Bunyak, Kannappan Palaniappan, Sumit Kumar Nath, and Gunasekaran Seetharaman. Flux tensor constrained geodesic active contours with sensor fusion for persistent object tracking. *JOURNAL OF MULTIMEDIA*, 2(4):20–33, August 2007.
- [3] <http://everything2.com/index.pl?node=cybernetics>.
- [4] Y. Diao and M. Franklin. Query processing for high-volume xml message brokering, 2003.
- [5] Michael Gschwind, H. Peter Hofstee, Brian Flachs, Martin Hopkins, Yukio Watanabe, and Takeshi Yamazaki. Synergistic processing in cell's multicore architecture. *IEEE Micro*, 26(2):10–24, 2006.
- [6] IBM. Overview available at <http://www.research.ibm.com/cell/>.
- [7] L. Kwong-Yan, R. Linderman, and G. Ramseyer. 100X Joint Battlespace Infosphere (JBI). *AFRL/IFTC, Rome, NY, Tech. Rep. AFRL-IF-RS-TR-2006-265*, 2006.
- [8] M. Linderman, J. Brichacek, S. Haines, B. Seigel, G. Chase, D. Ouellet, and J. O'May. A reference model for information management to support coalition information sharing needs. *The 10th International Command and Control Technology Symposium (ICCRTS)*, 2005.
- [9] Richard W. Linderman, Joshua Corner, and Scot Tucker. Swathbuckler: Real-time wide swath synthetic aperture radar image formation using embedded hpc. In *HPCMP-UGC '06: Proceedings of the HPCMP Users Group Conference*, pages 244–251, Washington, DC, USA, 2006. IEEE Computer Society.
- [10] Michael O. Little and William P. Berry. Real-time multichannel airborne radar measurements. *IEEE National Radar Conference*, 1997.
- [11] Lockheed Martin Missile and Fire Control. Overview available at <http://www.lockheedmartin.com/data/assets/3212.pdf>.
- [12] NVIDIA. [http://www.nvidia.com/docs/IO/43395/C870-BoardSpec\\_BD-03399-001\\_v04.pdf](http://www.nvidia.com/docs/IO/43395/C870-BoardSpec_BD-03399-001_v04.pdf).
- [13] Richard A. Poisel. *Electronic Warfare Target Location Methods*. Artech House, Boston, 2005.
- [14] Northrop Grumman Targeting and Surveillance. Overview available at <http://www.dsd.es.northropgrumman.com/litening/litening.html>.
- [15] Northrop Grumman Information Technology. Overview available at <http://www.it.northropgrumman.com/offer/mission/jcan.html>.
- [16] Todd Waskiewicz and Peter LaMonica. Developing an intelligence analysis process through social network analysis. *Proceedings of SPIE, Evolutionary and Bio-Inspired Computation: Theory and Applications II*, 6964, May 2008.
- [17] R. Yantom, D. Benincasa, and S. Wenndt. Effects of co-channel speech on speaker identification. *Proceedings of SPIE, the International Society for Optical Engineering*, 4232:258–261, 2001.