

AFRL-IF-RS-TR-2006-26
Final Technical Report
February 2006



VISITING FACULTY RESEARCH PROGRAM

The Research Foundation of State University of New York

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

STINFO FINAL REPORT

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2006-26 has been reviewed and is approved for publication

APPROVED: /s/

DUANE GILMOUR
Project Engineer

FOR THE DIRECTOR: /s/

JAMES A. COLLINS, Deputy Chief
Advanced Computing Division
Information Directorate

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 074-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this 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 Services, 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

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE FEBRUARY 2006	3. REPORT TYPE AND DATES COVERED Final Mar 03 – Mar 05	
4. TITLE AND SUBTITLE VISITING FACULTY RESEARCH PROGRAM			5. FUNDING NUMBERS C - F30602-03-2-0081 PE - 61102F, 62702F PR - SUNY TA - 03 WU - IT	
6. AUTHOR(S) Hongbin Li, Roger Chen, Nael Abu-Ghazaleh, Shuqun Zhang, Alistair Campbell, Peter Chen, Qiang Ji, Daniel Power, Lixin Gao, Kuo-Chi Lin, David Luginbuhl, Jingyuan Zhang, Abel-Aty-Zohdy, Edisanter Lo, Vira, Naren Adam Bojanczyk, Hongbin Li, James Haralambides, Laurence Merkle, Earl McKinney, Biao Chen, Peter Chen, Robert F. Erbacher, Sergey Lyshevski, Frank W. Moore, Jae C. Oh, Carla Purdy, Tarek Taha, Naren Vira, Shuqun Zhang, Thomas Hartrum, Joon Park, Yingrui Yang, Sibabrata Ray				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) The Research Foundation of State University of New York 35 State Street Albany New York 12207-2826			8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory/IFTC 525 Brooks Road Rome New York 13441-4505			10. SPONSORING / MONITORING AGENCY REPORT NUMBER AFRL-IF-RS-TR-2006-26	
11. SUPPLEMENTARY NOTES AFRL Project Engineer: Duane Gilmour/IFTC/(315) 330-2550/ Duane.Gilmour@rl.af.mil				
12a. DISTRIBUTION / AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.				12b. DISTRIBUTION CODE
13. ABSTRACT (Maximum 200 Words) The Research Foundation, for and on behalf of SUNY Institute of Technology (SUNYIT), has contributed significant research capability and capacity to the in-house program at the Air Force Research Laboratory. This was accomplished through the placement of highly motivated and accomplished Faculty members and Graduate Students pursuing advanced degrees in Engineering, Computer Science, Mathematics and other recognized technical disciplines critical to the advancement of information technologies. The program supported and enhanced the existing AFRL/Information Institute Summer Faculty Research Program and the Air Force Office of Scientific Research Summer Faculty Fellowship Program. SUNYIT worked closely with AFRL to help build, foster and nurture in-house research teams. Under this effort, SUNYIT recruited, placed, and supported administrative requirements of 31 faculty members and 10 graduate research assistants, and coordinated an additional 19 faculty extension efforts. This report contains abstracts of the research projects accomplished by the faculty members.				
14. SUBJECT TERMS Information Technology, Command & Control			15. NUMBER OF PAGES 26	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102

Table of Contents

1.0 Executive Summary	1
2.0 Faculty Researchers	1
2.1 The following sections describe the research that was pursued during fiscal year 2003.....	1
2.2 The following sections describe the research that was pursued during fiscal year 2004...	10
3.0 Graduate Students	21
4.0 Continuing Research Projects	21

1.0 Executive Summary

The Research Foundation, for and on behalf of SUNY Institute of Technology (SUNYIT), has contributed significant research capability and capacity to the in-house program at the Air Force Research Laboratory. This was accomplished through the placement of highly motivated and accomplished Faculty members and Graduate Students pursuing advanced degrees in Engineering, Computer Science, Mathematics and other recognized technical disciplines critical to the advancement of information technologies. The program supported and enhanced the existing AFRL/Information Institute Summer Faculty Research Program and the Air Force Office of Scientific Research Summer Faculty Fellowship Program.

SUNYIT worked closely with AFRL to help build, foster and nurture in-house research teams. Under this effort, SUNYIT recruited, placed, and supported administrative requirements of 31 faculty members and 10 graduate research assistants, and coordinated an additional 19 faculty extension efforts. This report contains abstracts of the research projects accomplished by the faculty members.

2.0 Faculty Researchers

2.1 The following sections describe the research that was pursued during fiscal year 2003.

Adam Bojanczyk/Cornell University

This research supports an ongoing effort at AFRL to design a decentralized software architecture controlling a distributed networked system of multi-model sensor nodes. The function of the system is to perform collaborative signal processing tasks like target detection, localization, identification and tracking.

The system should efficiently process computational tasks submitted by users who interact with the system from local or wide area networks. Coordination of all networked processing including the node execution environment, and optimizing resources to meet task objectives like latency and accuracy, is an important issue which this project addresses. The project involves the design of common computation and communication building blocks for sensor networking, and strategies for using heterogeneous sensor and network nodes to enhance performance of detections and decrease false alarms. The important assumption and feature of the system is that nodes can be reprogrammed remotely as they communicate with the rest of the system via wired or wireless links. It is reasonable to assume that nodes can function as html servers, can execute cgi, xml, and sql directives, and be equipped with a programming environment. Thus, it is possible to organize the network into a publish-subscribe system where publishers publish data, subscribers submit queries, and brokers match queries with data and possibly perform processing. The system sensor nodes need to be able to operate in several modes including: scanning, server interrupt and notify server only about unusual events. Sensor nodes maintain a database of locally sensed data, selected sensors' characteristics, and other knowledge. Varying degrees of communication bandwidth with respect to different classes of nodes influences the type of processing and communicated data. In the original prototype, the

data was post-processed off-line in a central computing facility. The goal of the current project is to collect and process data in real-time by the entire network. The signal processing problem addressed by the prototype system was that of distributed detection of the angle of arrival (AOA). In general, AOA systems are subject to multipath problems causing the system difficulty in accurately determining the point in the signal to be measured by all receivers. Two pairs of receivers can estimate the absolute location of the transmitter. Time difference of arrival systems operate by placing receivers at multiple sites, geographically dispersed in a wide area where each of the sites has an accurate timing source. The signal propagates to all of the antenna sites where the signal reception is time stamped. The differences in time stamps are then combined to produce intersecting hyperbolic lines from which the location is estimated.

Hongbin Li/Stevens Institute of Technology:

Hyperspectral imaging sensors (HIS) are capable of providing very fine spectral resolution that allows remote identification of ground targets smaller than a full pixel. Traditional approaches to the so-called subpixel target detection problem involve the estimation of the sample covariance matrix of the background from target-free training pixels. This entails at least two disadvantages: 1) large training requirement; and 2) high complexity, both due to the large size of the sample covariance matrix to be estimated.

In this effort, parametric adaptive modeling and detection techniques for HIS applications are investigated. Although such techniques have been successfully applied to solve radar detection problems, extensions for HIS applications are challenging, since HIS data are highly non-stationary. To deal with non-stationarity, we introduce a sliding-window based time-varying (TV) autoregressive (AR) modeling and detection technique, by which the spectral data is sliced into overlapping subvectors for parameter estimation and signal whitening. Experimental results using real HIS data show that the proposed parametric technique outperforms conventional non-parametric detection schemes, especially when the training size is small and/or in non-homogeneous environments.

Roger Chen/Syracuse University:

Distributed information management systems (DIMS) will include many sub-systems for multimedia manipulation/rendering, data/knowledge base management, graphical user interfaces, applications, security, and client/server interactions. The trend of software technologies in developing such systems (or subsystems) is to apply component-based techniques on top of object-oriented structures. Such a trend of componentization process will impose significant constraints, from the underlying software technologies, on system development, which makes interoperability between diverse software technologies a very challenging issue.

A preliminary feasibility investigation for a platform was performed, which allows the fusion of diverse software technologies in such a way that components from different technologies can seamlessly interact with each other to form a large system as if they were within a single technology.

The increasing complexity of knowledge/information storage, communications, processing and visualization as well as system/network infrastructures has made the conventional approaches to developing such information environments very inefficient.

To overcome such potential problems, one has to address the foundation of a DIMS system, such that it will provide the capability of allowing sub-systems and components to be pluggable into the environment and form a coherent system, where exchanges of events, data, and functions across the boundaries of sub-systems will be seamless and actions on all sub-systems will be globally coordinated and optimally controlled.

Unified modeling language (UML) is the standard notation for modeling software systems. UML class diagrams and sequence diagrams provide enough notations for software developers to describe component specifications and interaction specifications for component-based software architectures. With these foundations, a new type of component-based software synthesis tool should be able to: (1) transform the “interaction specification” in the description of component-based software architecture to “component instances”, which carries out these “requirements of interactions”; (2) realize the specification of component instances in the architecture using UML; and (3) generate the component composition of that architecture for a specific executable model.

Software technologies that were investigated include Java, COM, Active X, WSDL, Automation server, and NET. The focus was on the run-time behavior. A feasibility investigation of the functional requirements of a platform for software technology fusion was studied considering a wide range of software characteristics. A unified model is needed to represent a software system such that components, regardless of the origin of technologies, will be processed, manipulated, and accessed in the same way. The approach to distributed secret sharing using a non-uniform model is to place a secret entity on multiple distributed servers, such that each of them only carries a portion of the entity. Security compromises on a small number of servers will not allow attackers to gain enough information for the secret entity. Secret sharing alone does not provide a defense against mobile adversaries. The best defense against this is secret refreshing, wherein you create new secret sharing after a certain interval independent from the earlier ones and replace the old shares with the new shares.

Nael Abu-Ghazaleh/Binghamton University:

Parallel discrete event simulation (PDES) is an important application in use in many DoD projects; for example, PDES is used in large-scale wargaming, and in complex system design, analysis and verification. Improving PDES performance and capacity allows faster simulation times and more extensive analysis of more detailed models. These benefits are not application-specific: they should apply to any application that uses the improved simulation kernel. This research effort explored optimizing PDES in a Heterogeneous High Performance Computing (HHPC) environment. We profile a SPEEDES simulator and identify several opportunities. There were two major components: (1) optimizing the communication subsystem – a critical system for PDES since it is a fine-grained application; and (2) exploring the use of augmented field-programmable gate array or FPGA boards to accelerate simulation. While such approaches have been attempted for sequential and data path intensive applications, we believe that their use in clustered environments is novel. An example preliminary analysis result is that we observed an average of 20% performance improvement for one of our large benchmarks by removing the centralized communication server from the event message exchange path along with a number of other small improvements to the simulation cycle.

Shuqun Zhang/Binghamton University:

This research project addresses the tracking of objects in a video stream obtained from a moving airborne platform for annotation purposes. It requires developing a real-time tracking system that can track any Unmanned Aerial Vehicle (UAV) video objects indicated by a user mouse click. A general tracking framework is thus proposed, which is based on a spatio-temporal segmentation. The proposed algorithm compensates the image flow induced by the camera motion, and then detects and tracks object regions. The moving objects are detected using a temporal change detection algorithm. Change detection usually fails in detecting stationary objects. To overcome this problem, a simple method based on an image shrinking operation is used to make static objects “move” so they can be also tracked by the same algorithm as moving objects. Another problem with change detection is that many noise variations are also detected besides the object region, and if the object is small it can be very hard to separate out the moving target from the noise. An effective method of extracting object regions is proposed based on the assumption that the target should be close to the mouse click point. To extract more object regions, an edge detection-based segmentation is used. Segmenting the subsequent frames of video sequence and establishing a correspondence of moving objects between frames perform the final step of tracking. The main features of the algorithm are a low degree of computational complexity and generality suitable for objects of various types and sizes. Results were demonstrated on a few real video sequences.

Alistair Campbell/Hamilton College:

Many sorts of computer users have the need to use data provided by other computer users. As the amount of available data increases, so does the need for robust techniques for organizing, categorizing, describing, along with methods for finding and retrieving the data. One of the most promising techniques is to organize the data into a formal ontology, a machine-readable conceptual schema that describes the domain of the data. Ontologies consist of a controlled vocabulary together with relationships between the concepts denoted by that vocabulary, and axioms that further constrain the proper use and interpretation of that vocabulary. Usually an ontology defines a taxonomic (isa-relationship) hierarchy from more general concepts at the top, to more specific concepts in the lower levels. For data organization and classification, this taxonomic hierarchy is the primary organizational benefit of using a formal ontology.

One of the most appealing aspects of the formal ontology endeavor is achieving a controlled vocabulary. No matter what formalism or level of detail is employed, an ontology serves to ground the meaning of terms used in information interchange and reasoning.

Despite efforts to establish shared ontologies, much data remains uncategorized. Database schemata, XML Document Type Definition (DTD's) organizational personnel records, and similar data have not been linked to any of the established ontology efforts.

This work describes initial efforts in aligning flat data records described by glosses to the WordNet lexical database.

A flat data set such as an XML tag set is described with formal syntax, but formal semantics are usually absent. We describe a system for associating explicit semantics with each element of a flat data set. A particular flat data set, the General Military Intelligence (GMI)

Namespace of the DoD Metadata Registry and Clearinghouse is explored for its assumed semantics. WordNet, a general-purpose ontology of everyday English, serves as the ontology to which the GMI namespace is mapped. Definition glosses of the GMI namespace are disambiguated using a lexical chaining technique. Tags themselves are then mapped to WordNet semantic concepts using three separate heuristics.

We explored issues in ascribing semantics to flat data sets by mapping the elements to ontologies. In particular, we presented a system that maps gloss-defined XML tags to the WordNet lexical ontology. By employing lexical chaining techniques, we are able to correctly disambiguate the words used in a tag gloss, and achieve 67% accuracy in mapping tags to semantics.

Peter Chen/Chen & Associates, Inc.:

We propose an effort to investigate how to make an architecture diagram compatible with the diagrams used in the information fusion group at AFRL and also to investigate the roles of knowledge-based techniques.

What is needed is the architecture of a comprehensive system that integrates different technologies including the technologies in these three important areas: information assurance, information fusion, and knowledge-based techniques.

We focused on building a decision support system for “information assurance;” however, the architecture proposed is very general and can be used for systems for other focuses or orientations.

This effort describes an architecture for decision support systems for information assurance using knowledge engineering and information fusion techniques. The architecture is divided into 3 levels: the bottom level consists of data sources and basic intrusion detection modules, the middle level consists of various modules performing knowledge engineering and information fusion techniques, and the top level consists of modules assisting decision makers and operators to analyze the situations and take appropriate courses of actions.

This architecture will be very useful in several applications: (1) identification of important software modules needed for a comprehensive decision support system for information assurance (or other functionalities); (2) a “roadmap” for planning and monitoring R&D activities in information technology; and (3) a big picture (bird-eye view) of various R&D program activities in major funding agencies (such as DARPA, AFOSR, ONR, ARO, NSF, etc.).

Qiang Ji/Rensselaer Polytechnic Institute:

A very difficult problem for the causal analysis tool (CAT), in particular, and for effects based operation (EBO), in general, is efficient selection of action plans. This is a problem that could significantly hinder the utility of CAT and the development of EBO. We investigated two methods to address this problem: one is a graph-theoretic approach and the other is a greedy approach. The graphical approach is aimed at developing a theoretical basis to reduce the search space and to more systematically and methodically search for the optimal or near-optimal solutions so that they can be identified in an efficient manner. The action ensemble based greedy approach recursively identifies an ensemble of actions to maximize the performance of an

ensemble on the goal. Preliminary empirical studies indicate the promise of both methods in both solution optimality and time complexity.

The idea is to first define a measure to quantify the utility of an action, and then use the utility measure to define pair-wise synergy between two actions. Synergy among multiple actions are then graphically inferred from the pairwise synergy matrix to have a large number of action combinations excluded from further consideration due to low synergy among them. Heuristic rules are then developed to identify candidates of optimal action combinations, based on their connection configurations in the synergy graph. The main contributions of our graphical approach include the proposed metric to quantify action utility and synergy, the introduction of a graphical representation of action synergies, and the development of the theorems and rules to efficiently parse the synergy graph. The graphical approach study, however, remains preliminary and needs further theoretical developments and experimental users to decide what actions to include initially, how many actions in the plan, and even the criterion used to evaluate the plan. We recommend it be implemented in CAT.

Daniel Power/University of Northern Iowa:

This project identified research and development needs associated with providing advanced military decision and planning support. The focus is on command and control in military operations with an emphasis on air operations. The report describes the research approach, defines key issues and terms, discusses the assumptions of military decision support and the current decision support “box”, briefly reviews how computerized decision support can assist in military decision tasks, discusses substitutes for computerized decision support in the military, and summarizes nine major recommendations that provide a “way ahead” for expanding the “envelope of possibilities” for advanced, computerized decision support for military command and control.

The major recommendations are:

1. Conduct a Decision Process Audit related to Command and Control in Operation Iraqi Freedom.
2. Develop a prototype for a Deployable Operations Command Center (DOCC).
3. Prepare an inventory of C2 decision support systems and software.
4. Prepare a literature review of military decision support research related to Command and Control.
5. Prepare senior military officers for design and use of decision support systems (DSS).
6. Create a multi-participant Air force game.
7. Investigate fifth generation agent-based simulation
8. Conduct more behavioral research on the impact of DSS on commanders and their staff.
9. Explore next generation Knowledge-driven DSS.

Lixin Gao/University of Massachusetts, Amherst:

Since the days of the Morris worm, the spread of malicious code has been the most imminent menace to the Internet. Worms use various scanning methods to spread rapidly. Worms that select scan destinations carefully can cause more damage than worms employing a

random scan. This effort analyzes various scan techniques. We then proposed a generic worm detection architecture that monitors malicious activities. We proposed and evaluated an algorithm to detect the spread of worms using real-time traces and simulations. We find that our solution can detect worm activities when only 4% of the vulnerable machines are infected. Our results bring insight on the future battle against worm attacks.

In the future, an efficient traffic monitoring infrastructure will be an important part of the global intrusion detection systems. A consequence of the worm detection method is that the attackers will have to use a limited number of IP addresses to scan the Internet. Therefore, the impact of worm scanning on the Internet traffic will be reduced.

In this effort, we discussed different types of scan methods and their effects on future worm propagation. We find that as the backbone link speeds and hosts of greater capacity are affordable to the attackers, it will be more difficult for us to detect worm scanning from the Internet traffic. However, the detection methods can still be useful in that it forces the attacker to use less traffic and scan more slowly and cautiously.

We designed two new scan techniques, routable scan and Divide-Conquer scan. Basically, they both use the idea of a routable IP address list as the destination base where the scan object is selected. A routable worm is easy to implement; it poses a big menace to the network security. We must keep in mind that anytime in the future the next worm incident may be worse.

For this strain of scan methods we designed a detection architecture. Specific detection methods were also designed. We find that using the victim number based algorithm, worms more serious than the Code Red and the Slammer can be detected when less than 4% of vulnerable machines are infected.

Kuo-Chi Lin/University of Central Florida:

A swarm is defined as a collection of autonomous individuals relying on local sensing and reactive behaviors interacting such that a global behavior emerges from the interactions. The advantages of using a swarm over a more deliberate team are robustness to disturbances, cost effectiveness in construction/operation, and mission scalability.

The UAV swarm is modeled as follows. The UAV swarm is homogeneous, except for a few specialists if needed. Each UAV only responds to local situations or threats based on the sensory inputs. The UAV's are controlled by a set of behavioral rules. Human controllers, either centralized or distributed, intervene only when necessary.

To reduce the solution space when designing a mission, the following propositions are made: if each UAV's low-level behaviors are properly designed, the swarm can exhibit proper collective low-level behaviors. The higher-level behaviors of the swarm can be the proper combination of sequences of low-level behaviors.

This report analyzes the hierarchy of the UAV behaviors, and uses simulation to demonstrate those behaviors. Very simple bang-bang controls of the individual UAV motions are proven to make the swarm exhibit several higher-level behaviors.

David Luginbuhl/Western Carolina University:

As the military's reliance on information technology grows and as more and more information becomes available, there is a growing need for effective mechanisms of managing that information. To better understand specific needs for such a system and determine current capabilities gaps, we consider the broader landscape of Information Management (IM). The intent of this work investigated high-level technology challenges related to IM by developing a solid, meaningful definition of IM in the context of combat operations, as well as a detailed taxonomy of attributes and capabilities of IM; identify opportunities for technological improvement; and determine how these apply to current development of the Air force's vision of a Joint Battlespace Infosphere (JBI).

Jingyuan Zhang/University of Alabama:

Accurate weather prediction is extremely important to military command and control. The success of a military operation depends on the accuracy of weather prediction. Numerical weather prediction has proven to be the most accurate method today. A more accurate numerical weather prediction requires a finer resolution, which, in turn, requires more observation data to be collected and more observation stations to be placed. However, it is usually impossible to establish permanent weather stations in a battlefield. To collect necessary weather data, it is proposed to use an ad-hoc sensor network that can be set up quickly anywhere and anytime on an on-demand basis. In this project, a prototype is developed for local weather predictions based on observations from ad-hoc sensor networks. A protocol was developed that describes how to set up a sensor network for collecting temperature observations, how to synchronize the components in the prototype, and how to perform finer temperature predictions based on the observations from the sensor network and coarser predictions. The experimental results show the prototype is able to predict the temperature change more accurately by incorporating the local conditions.

Abel-Aty-Zohdy/Oakland University:

In this effort, experiments were completed to confirm the viability of Bacteriorhodopsin (bR) as a hopeful candidate for a novel memory device. More experiments must be done to measure a write efficiency that takes into account diffusive mechanical protein interactions. These measurements must be compared to the theoretical model and discrepancies must be accounted for. An initial model setup was created in the Java language. This program must be completed with the proper mathematical implementations and eventually built into a full 3-D model. This program is adaptable to add in new variables as new influences to the model behavior are discovered. If necessary, more protein interactions other than mechanical forces can be added into the program. Another key implementation that must eventually be made to the model is the specific time-dependent intermediate states of the write and page processes. As new strains of bR are being developed, the cumulative change in the intermediate state properties (intermediate state yields) is what causes the write-efficiency to change. As a future obstacle, these and other implementations to the model should be made. For the Java program, complicated mathematical implementations need to be added, and more coding would be required based on the way the math is going to work in the program. The possibility of

integrating this new memory device into a computational cognitive architecture is being explored.

Edisanter Lo/ Susquehanna University:

The objective of this research is to develop and implement an iterative auto-regression algorithm for detecting a target in hyperspectral images. The algorithm is based on the theory from a general linear model which is used to estimate the parameters. The modeling process is iterative. Initially, a spectral band (variable) is modeled as a simple linear regression of another spectral band. The resulting residual is then modeled as a simple linear regression of one of the remaining spectral bands. The iteration continues in this way and is terminated when there is a significant lack of fit in the model. The resulting model, which is called iterative auto-regression (ITAR), is no longer a simple regression model or an auto-regressive model. The test statistics used in the whitening process are the adaptive matched filter (AMF) and adaptive coherence estimator (ACE) test statistics for testing sub pixel targets assuming an unstructured background model. The algorithm was tested on real pixels from the hyperspectral images of the Purdue campus and Washington D.C. Mall to compare the relative performance between the parametric implementation of AMF and ACE using the whitening process with the covariance implementation of AMF and ACE.

Vira, Naren/Howard University:

The objective of this research is to develop a method to create three-dimensional models using a stereo triangulation technique. A Java-based algorithm was developed which can take a pair of digital images as input to recreate a texture mapped 3-dimensional (3D) model utilizing a general purpose, interactive JView based program engineered by AFRL. Though several procedures have been proposed in literature on how to generate 3D models from 2-dimensional (2D) stereo pairs by recovering the depth, they are limited in predicting an accurate pixel matching between image correspondences. Furthermore, the majority of approaches fail to reconstruct the 3D-image geometry to show how well their adopted techniques have been able to recover the scene depth at every pixel. We have employed a region-based block matching methodology for color image analysis. The reconstructed 3D texture mapped model can be viewed and maneuvered on a computer screen by moving up and down, rotating around, and zooming in and out with the help of mouse buttons. This interactive viewing capability is an aid to the visualization process and facilitates algorithm designers to test step-by-step their codes and make appropriate judgments as to the accuracy of the underlining reconstruction techniques (sort of a visual debugger). Three sets of examples are presented showing detailed disparity map and reconstructed geometries. Further research work is needed to improve the accuracy and the speed of the algorithm.

2.2 The following sections describe the research that was pursued during fiscal year 2004.

Adam Bojanczyk/Cornell University

The objective of this research is to study a distributed signal processing system for target location, tracking and identification. The system will be based on a novel beamforming technique developed by AFRL that accommodates signals of varying bandwidth and central frequency. Asynchronous computational methods will need to be designed which adjust computation to varying communication bandwidth among distributed nodes.

Hongbin Li/ Stevens Institute of Technology:

Hyperspectral imaging sensors (HIS) can provide very fine spectral resolution that allows remote identification of ground objects smaller than a full pixel in an HIS image. Traditional approaches to the so-called subpixel target signal detection problem are training-inefficient due to the need for an estimate of a large-size covariance matrix of the background from target-free training pixels. This imposes a training requirement that is often difficult to meet in a heterogeneous environment.

In this effort, a class of training-efficient adaptive signal detectors was examined by exploiting a parametric model that takes into account the non-stationarity of HIS data in the spectral dimension. A maximum likelihood (ML) estimator is developed to estimate the parameters associated with the proposed parametric model. Several important issues are discussed, including model order selection, training screening, and time-series based whitening and detection, which are intrinsic parts of the proposed parametric adaptive detectors. Experimental results using real HIS data reveal that the proposed parametric detectors are more training-efficient and out perform conventional covariance-matrix based detectors when the training size is limited.

James Haralambides/ Barry University:

In this effort an algorithm was developed which is our initial effort to attack the 3-D routing problem from a 3-D point of view rather than extend heuristic solutions for the 2-dimensional problem. We have constructed a multi-level placement and routing algorithm that produces 100% completion of nets. The design was modeled on a 3-D grid. Interconnections were allowed to overlap and nodes were treated as switching elements. We have incorporated a node-weight-based, shortest path algorithm to detect and form near-optimal cost paths between pairs of nodes at a geometric proximity. A probabilistic element is used to reduce the problem of producing local rather than global optima. The routing algorithm was modified to perform interconnections in a multi-level setting. The algorithm we have constructed works independently of the presence of obstacles (faulty nodes or groups of nodes on the grid.)

Laurence Merkle/ Rose-Hulman Institute of Technology:

The AFRL in-house project "Hybrid Architectures for Evolutionary Computing Methods" investigates novel computing architectures that facilitate evolutionary computing (EC)

methods such as genetic algorithms (GA's) and genetic programming. This class of methods is inspired by the processes of natural evolution, and the methods are most frequently applied to the solution of difficult optimization problems. For example, this project applies them to the problem of parameterizing models consisting of sets of coupled, non-linear, ordinary differential equations to fit experimental data.

The first area of significant progress accomplished in this effort is the development of parallel implementations of the GA using both the "farming model" and the "island model" on a 48-node cluster. The second area of progress is the development of the first known general purpose architecture for reconfigurable hardware implementations of GA's, and the application of that architecture in this project. The status of related "evolutionary hardware" design efforts was reviewed, including visits to coordinate with AFIT and Wright-State University.

Two areas of future research were investigated. The first combines the progress made in both areas by using the built in high speed FPGA to FPGA board communications channels in the heterogeneous cluster to implement the first-ever FPGA-per-island parallel GA. The second investigates the automatic application of these tools to new optimization problems, taking into consideration area and timing requirements for hardware implementations of problem specific function evaluation, mutation, and recombination units.

Earl McKinney/ Bowling Green State University:

Understanding how to support teams that face crisis is essential. To support their teams in crisis, military organizations seek to leverage advances in information technology. These advances include automation support to the warfighting team (e.g. an electronic checklist for a flight crew), as well as collaboration support, such as linking engaged combat troops to intelligence services. While automated support is rapidly developing, very little consideration has been given to enhancing the collaboration support for teams that face crisis. Collaborative support for these teams in the past was limited by the available technology. Now, with advances in network capacity and sensors, IT has enabled teams that face crisis to obtain collaboration support from others in the organization. Crisis by their uniqueness reduce the utility of automated support. The challenges are discovering what is happening, and thinking through irrevocable decisions. As a result, automated support, while valuable, should not be the only available support for teams that face crisis. Collaboration with other human experts is necessary to aid problem discovery and to consider ramifications of responses.

This project suggests a preliminary set of IT system attributes to support collaboration for teams that face crisis. These attributes are based on two frameworks that have been developed to mitigate the effects of crisis. One is an organization approach called the High Reliability Organization (HRO), the other, a team approach known as Crew Resource Management (CRM).

Biao Chen/Syracuse University:

This research effort studied the potential of multiple-input multiple-output (MIMO) communications in ad hoc networks and developed transceiver design methodologies to realize their potential. For an ad hoc network where multiple peer-to-peer communications need to be simultaneously facilitated, the following three frameworks appear to be logical choices.

- Orthogonal communication through channel division.

- Overlay transmission: all transceiver pairs communicate in the same time-frequency channel. This overlay transmission is enabled due to the spatial multiplex capability of each MIMO node.
- Contention based channel access, such as (slotted) Aloha.

Previously, we showed the performance limitation of the overlay transmission in large scale ad hoc networks. Specifically, we have shown that:

1. Pure overlay transmission has an asymptotically limited network capacity. If the number of transceiver pairs grows, the total network capacity is limited for blind transmitters; while for the informed transmitter case, the network capacity is limited by $t + r + 2ptr$. In the above, t and r are respectively the transmitter and receiver antenna numbers.

2. For orthogonal transmissions, the network capacity grows logarithmically as the number of users grows for orthogonal transmissions.

However, it has been observed that with a finite number of users (on the same order of magnitude as the transceiver degrees of freedom), overlay transmission does enjoy advantage over orthogonal transmission in achievable capacity. In particular, there seems to be a ‘sweet’ region for the MIMO interference transmission where the per user achievable rate appears to stabilize before dropping further; numerical results indicate that the location of this region is strongly correlated to the number of antennas. This is tantalizing: it points to a very promising networking protocol where a newly active user can be added to an already occupied channel without compromising each individual user’s Quality of Service.

The above observation, while interesting, is based on empirical results. Furthermore, the results are obtained based on a single user detection assumption which is a performance bottleneck in a multiuser system. An interesting question is that: if advanced transceiver structure is implemented, such as those that use successive interference cancellations, how much extra performance gain can be obtained over orthogonal transmission?

This is the focus of this study. The preliminary results are summarized below:

1. Adopting the classical superposition code idea, we are able to obtain a lower bound on the sum capacity for a two user MIMO overlay transmission system.

2. For an orthogonal transmission, such as frequency division multiple access (FDMA), the sum capacity is a strictly concave function of the bandwidth allocation factor. The same results hold for time division multiple access (TDMA) with an average power constraint.

3. The low bound for sum capacity of the MIMO overlay transmission significantly outperforms that of the sum capacity of the orthogonal transmission.

Peter Chen/Louisiana State University:

Validity assessment of information from various data sources is a crucial problem in information fusion. This effort identified several key issues and proposes a framework to solve this problem. It also describes a prototype of a decision support system to support the estimation of the composite data values from heterogeneous databases with different validity assessment values. We also address several issues closely related to information validity assessment, such as meta data modeling and reverse engineering of existing database schemas into conceptual models.

Intelligent extraction and validation of information from multiple database systems that are of different models (e.g. relational, Object-Oriented, flat files, etc.) are the key prerequisites to efficient knowledge integration for decision support.

Our effort attempts to provide an active information management framework by allowing applications or users to query or extract information from databases of different models. In addition, the active information system cooperates with the users by automatically informing the applications or users of the necessary information whenever critical situations occur. One of the key aspects in our effort is to develop a framework and techniques for information validity assessment. How can we handle this problem? We think there are at least three things that will be very useful:

- A framework for analysis.
- Computational formulae or algorithms for conflict resolution.
- A software program that helps people make decisions.

We outlined some of the important steps in a framework for information validity assessment. We also described some algorithms for conflict resolution. We implemented a prototype of a decision support system for helping people to make decisions under conflicting data situations.

Robert F. Erbacher/Utah State University:

The need for visualization-based intrusion detection and analysis techniques has been shown due to the complexity of the underlying data and the inability of purely algorithm techniques, such as data mining to effectively analyze intrusion data. We have developed several visualization mockups which have great potential for aiding in the identification and analysis of intrusions. Our goal with these visualization techniques was to focus on solutions for competent attacks. The goal was to ignore or filter out script kiddies and the like as they can be easily blocked by firewalls and should not consume the analyst's time. These techniques rely on glyph-based techniques which have proven effective for the representation of multi-parametric data. These visualization techniques in conjunction with effective interaction techniques will create a complete environment.

We also provide an initial architecture to guide the future implementation of the environment. As a first step in this implementation, a Java native interface based wrapper has been created for Simpcap, an engine for accessing large-scale libpcap files. The architecture incorporates support for the tap+bridge interface as well as other database formats. This will provide support for many of the interfaces in use by AFRL and other organizations. The architecture also incorporates an advanced user interface and support for multiple visualization displays.

Multiple visualization displays are critical as no single visualization can show all aspects of a database. Thus, these difference visualization techniques allow for different views of the data, either by examining the same data parameters or completely separate ones. This provides for a more complete representation and analysis of the database than is otherwise possible.

Finally, we incorporate a discussion of several other related topics to information awareness. This includes cyber command and control, computer forensics, trending, and intrusion detection. While we believe the developed techniques will prove immensely effective,

we have begun to identify future directions of research to continue the process of refinement and improvement and engender future systems which provide complete solutions to the problems at hand.

Sergey Lyshevski/Rochester Institute of Technology:

Carbon molecules are examined as multi-terminal electronic nanodevices to be utilized in the envisioned 3D computing architectures. The specific attention is focused on functional molecules that exhibit the desired electronic characteristics. These multi-terminal carbon-based 1-nm-size transistors should be aggregated within neuronal topologies. We propose a **Molecular Multi-Terminal Electron-Tunneling Fullerene-Transistor (MEF^{3T})** that consists of C_{xx} molecule (endohedral or doped fullerene derivative) covalently bonded to organic ligand. Modeling, simulation and analysis of the fullerene-ligand complex, as a MEF^{3T} prototype, are performed and reported.

There is a need for development of high performance computing architectures that meet requirements and specifications of future Air Force Command, Control, Communications, Computers and Intelligence (C4I) combat management systems. The importance of super high performance 3D multi-terminal transistors, as a bottom-up primitive, is a fact that these transistors lead to design of aggregated 3D neuronal-based circuitry topologies. Correspondingly, novel super high performance 3D computing architectures become a reality. With the major focus on the device level research, we propose to utilize MEF^{3T} as the basic bottom-up device in envisioned computing architectures. Three-dimensional computing platforms, designed from the super high performance MEF^{3T} bottom-up primitives, will lead to:

- Enormous military advantages guarantying information processing preeminence and computing superiority;
- Strong commercial potential with immediate applications in new generations of preeminent processors and memories;
- Novel technology development with a sound technology transfer feasibility to future Air Force C4I systems.

Molecular carbon-based electronic nanodevices will revolutionize electronics and information technologies due to enormous performance enhancement. It is found that MEF^{3T} guarantees super high bandwidth (switching frequency), functionality and incredible device density.

Frank W. Moore/ University of Alaska at Anchorage:

This effort developed a genetic algorithm that evolves optimized sets of coefficients for signal reconstruction under lossy conditions due to quantization. Beginning with a population of mutated copies of the set of coefficients describing a wavelet-based inverse transform, our genetic algorithm systemically evolves a new set of coefficients that significantly reduces mean squared error (relative to the performance of the selected wavelet) for various classes of 1-D and 2-D signals.

Collectively, the results of this study suggested that the number of coefficient sets capable of producing high-fidelity signal reconstructions under lossy conditions may be much larger than previously believed. The novel coefficient sets evolved during this study violated

wavelet properties required for perfect reconstruction, such as invertibility and non-redundancy. Nevertheless, the corresponding inverse transforms consistently outperformed the DauM inverse transform, often producing significantly higher fidelity reconstructions of periodic signals and images, as measured by the percentage reduction in the mean-squared error (MSE) of each reconstructed signal. The results of this study strongly encouraged the identification and use of evolved inverse transforms for signal reconstruction under lossy conditions. In particular, our results conclusively demonstrated our GA's ability to automatically identify novel sets of coefficients for inverse transforms that successfully reconstruct various classes of periodic signals and images under low conditions subject to quantization.

Jae C. Oh/Syracuse University:

This project continued an ongoing long-term research effort in developing theory and application of a formal model for resource sharing and allocation in massively distributed multi-agent environments in which agents are selfishly rational. We are particularly interested in the application of the formal model to the coalition formation among allied forces in an extremely complex global wargame situation consisting of military forces from multiple divisions and countries. Often, a unit engaging in a conflict may require assistance from other units in order to successfully resolve the conflict. The question is, when such assistance is needed by a unit, unit i , which other coalition unit should i request from for assistance? Furthermore, when a request is received, the receiving unit must decide whether to send help or not. When the number of units is large and the units are deployed in a distributed fashion in a large area, a centralized method for unit deployment may not work.

Each unit tries to maximize its utility throughout its deployment period. Since units are rational, we need some kind of reputation management similar to a peer-to-peer (P2P) computing environment. Reputation management in most P2P systems depends on the reputations of nodes circulated in the environment. Because reputations are circulated, the mechanism suffers potentially serious problems.

We present a new mechanism that attempts to address the above problem and derive the utility equation used internal to the agents. The experimental results show that the distributed algorithm deploys resources slightly better than centralized schemes.

In this effort, we worked on: (1) a reputation inference system for rational multi-agent environments; (2) a utility equation that is used by rational agents in interacting with others; and (3) the application of the model to coalition wargame scenarios. Experimental results show that the new model is suitable for large distributed resource allocation and management problem where centralized methods may fail. In numerous experiments, the new model outperforms a centralized non-game theoretic model, in resource sharing and allocation, as well as task completion rate and cost.

Carla Purdy/University of Cincinnati:

Two major projects were involved in the effort for this time period:

1. Development of a robust, well-tested library of VHDL-AMS and Verilog components, along with at least one complete architecture, for efficient application-specific implementations of genetic algorithms in hardware.

2. Development of preliminary designs and simulation models for biomolecular computation devices to be used in place of or in addition to traditional computing devices for eventual implementation in a wide range of sensor and actuator applications. The interpretation of the sensor inputs and the controllers for the actuators to be developed eventually will use heuristic techniques such as the genetic algorithms; along with other techniques such as neural nets, fuzzy measures, and wavelet-based smoothing and approximation. We have completed high-level models of the biochemical switch in VHDL-AMS and have validated them through simulation and we are continuing to work on the models to interface the two technologies.
3. In addition, we continued development of methodologies, tools, and educational modules for robust verification and validation of hardware / software components which employ heuristic techniques and / or exhibit stochastic behavior were considered.

As can be seen from these brief descriptions, all three components of this project are interrelated both through potential applications and through the ubiquitousness of nondeterministic phenomena which must be dealt with. Thus the work in (3) above can be seen as a foundation for allowing us to use the technologies described in (1) and (2) with increased confidence and understanding. As circuit technologies shrink to the nanoscale, we must develop the ability to routinely build robust, reliable hardware/software systems which incorporate stochastic behavior. In addition, we must continue to exploit the power of bio-inspired computing strategies, such as genetic algorithms, in order to solve the increasingly complex information interpretation and decision problems which arise in modern computational systems. This work sets standards for problem definition, benchmark data set content and representation, randomization, algorithm/system definition, specifications of the experimental environment, performance measures, statistical analyses of results, visualization and archiving of results, and is also designed to support collaborative experiments.

Tarek Taha/Clemson University:

For space based embedded processors, low energy consumption is a major design criterion. Reducing the energy consumption lowers both packaging and overall system costs. Within an embedded processing system, the memory hierarchy is one of the main sources of power consumption. The SDRAM (synchronous dynamic random access memory) is expected to be one of the main sources of energy consumption within the memory hierarchy, because it is a large off chip structure to which significant amounts of data are transported along wide off chip buses.

This effort investigated the concept of reducing the energy consumption associated with accesses to the SDRAM. It developed techniques to measure the energy consumption of SDRAM accesses by incorporating activity counters within the memory controller. These counters measure the activity of the SDRAM and also monitor the activity on the off-chip bus to the memory. Using these counters, the investigation also looked into whether a wide or narrow data bus to the SDRAM would be more energy efficient. The two bus sizes investigated were 32 bits and 128 bits (resulting in a 256 bit transfer taking 4 cycles versus 1 cycle respectively).

The results of the investigation show that activity counters are able to track SDRAM activity and monitor changes on the bus. Internal DRAM activity did not appear to be affected significantly by data bus width. However, the change on the bus energy by changing bus width was more noticeable. The result of increasing the data bus width generates different results based on how energy consumption is calculated. At present, there are two approaches to calculating the bus energy consumption. One method is to count only wire transitions to 1 since transitions to 0 do not need to charge a wire. An alternate (and predominant) method is to account for all wire transitions (i.e., both 0 to 1 and 1 to 0). Based on the first method of measurement, the wider data bus produces about a 30 % reduction in bus energy consumption. The second method of measurement, on the other hand, shows a narrower bus produces a 30% reduction in bus energy consumption. This work also looked into intermediate cases where transitions to 0 do require energy, but are weighted less than transitions to 1.

Naren Vira/Howard University:

The use of hand gestures provides an attractive alternative to cumbersome interface devices for human computer interaction. In particular, visual interpretation of hand gestures can help in achieving the ease and naturalness desired for human computer interaction. This has motivated a very active research area concern with computer vision-based analysis and interpretation of hand gestures. In order to enhance multimedia capabilities for the interactive DataWall, we are interested in exploring this technology. Furthermore, since the Datawall room has a large display screen size (12' x 3'), oftentimes it is difficult to precisely identify the information presented on screen that someone in the audience is pointing to. On the same token, presenters pointing at information on the display by hand during the presentation cannot clearly be visualized by the audience. Both scenarios result into loss of effective communication.

We initially concentrated on investigating the feasibility of utilizing an image triangulation technique for accurately positioning and tracking a virtual pointer pointing towards the DataWall. The modeling and simulation task was carried out in which synthetic images of the pointer (generated using Studio 3D max) were input to a MS Visual C++ code. The projected locations on the DataWall of a virtual pointer were compared with the known locations retrieved from the Studio 3D max. The results were promising and the pointing actuary of the pointer on the DataWall is in the neighborhood of 0.06 feet. This actuary is regarded to be well within our specified goals. Further development on the work is warranted.

Preliminary analysis of the present work reveals that the image triangulation method works reasonably well for locating the pointer's pointing projection on the DataWall. It is now worthwhile to carry on the project to demonstrate using real camera systems on the DataWall.

Shuqun Zhang/College of Staten Island /City University of New York:

This project aims at developing robust tracking algorithms for dealing with complex motion and shape dynamics in UAV videos, where real-time processing and minimal human intervention (a mouse click) are required. Earlier research on the same topic shows that the development of general tracking algorithms can deal with large variability of object shapes and sizes, but the segmentation may not be accurate enough for all types of objects. In this effort, we studied the use of switching/mixing algorithms for improved segmentation and tracking

performance. A combined tracking method is proposed to combine the power of spatio-temporal segmentation and deformable models. The spatio-temporal segmentation is used to determine the size of the object being tracked and to track small objects, while larger objects are tracked using a modified snake model. The proposed statistical snake method uses both edge and region information by modeling image gradient as a Rayleigh distribution. The image gradient is enhanced by either color or Gabor filtering. Several strategies are used to overcome the problems in complexity, and removing manual parameter tuning. The proposed tracker handles clutter and occlusion problems by first detecting them using the size of object bounding box and then automatically correcting the bounding box by adding more forces to the snake model. By removing motion computation, the proposed snake-based tracker can easily achieve real-time implementation. Promising experimental results are provided using real UAV video sequences.

Thomas Hartrum/Wright State University

An integrated demonstration system for JBI publishing and subscribing by embedded assets was developed by adapting the Guardian Agent code from Boeing's Insertion of Embedded Infosphere Support Technology (IEIST) simulation system. That effort allowed the real-time demonstration of publishing from a Joint Unmanned Combat Air System (J-UCAS) simulator using Mobile Code sent in a Force Template simultaneously subscribing to the information. When ported to the multi-processor system with the real JBI platform, a run-time error occurred for which there was insufficient time to locate and correct.

The run-time error involving publishing to the real JBI platform was located and corrected. This allows full publishing of and subscribing to the J-UCAS in real time. In order to demonstrate an integrated publishing environment, a simulator was developed to allow an Air Force Synthetic Environment for Reconnaissance and Surveillance (AFSERS) Predator simulation to publish data to the JBI. Based on operational considerations, this was implemented to run on the AFSERS platform as a stand-alone publisher without the need for an Airborne Warning and Control System (AWACS) or Guardian Agent using a two-process architecture.

Finally, a subscriber was developed to subscribe to the Predator data. This is conceptually similar to the subscriber previously developed for the J-UCAS data, but uses a completely revised architecture to provide better modification and extension.

Overall this effort met its goals. The IEIST-based demonstration is running on a multi-processor network, successfully publishing and subscribing to J-UCAS data from the Boeing simulator in real-time. The AFSERS/Predator demonstration is also running, successfully publishing and subscribing to Predator data in real-time. These can be run simultaneously, demonstrating the integrated capability of publishing and subscribing through JBI from embedded tactical platforms.

Joon Park/Syracuse University

As information systems become ever more complex and the interdependence of these systems increased, the survivability picture becomes more and more complicated. The need for survivability is most pressing for mission-critical systems, especially when they are integrated with other Commercial Off-the-Shelf (COTS) products or services. When components are exported from a remote system to a local system under different administration and deployed in

different environments, we cannot guarantee the proper execution of those remote components in the currently working environment. Therefore, in the runtime, we should consider the component failures (in particular, remote components) that may either occur genuinely due to poor implementation or the failures that occurred during the integration with other components in the system. Also, we should protect the component against cyber attacks.

Although advanced technologies and system architectures improve the capability of today's systems, we cannot completely avoid threats to them. This becomes more serious when the systems are integrated with COTS products and services, which usually have both known and unknown flaws that may cause unexpected problems and that can be exploited by attackers to disrupt mission-critical services.

Traditional approaches for ensuring survivability do not meet the challenges of providing assured survivability in systems that must rely on commercial services and products in a distributed computing environment.

We identified the following generic problems of the component-sharing services in large, complex, distributed systems that span multiple organizations.

- *Problem 1:* An autonomous mechanism to support trusted component-sharing services between different organizations or systems is needed due to the fact that there is no single administrator who can control every aspect of the various systems used in an enterprise.
- *Problem 2:* Testing software components before deployment cannot detect or anticipate all of the possible failures that manifest themselves during runtime, especially when external components are integrated. Some failures are detected only when the components are deployed and integrated with other components in the operational environment(s).
- *Problem 3:* The currently available redundancy-based static approaches cannot solve the problem completely. If one component has failed because of reason R1, the rest of the redundant components will fail for the same reason.
- *Problem 4:* Even if we know the reasons for and the locations of the software failures, in most currently available recovery approaches in distributed computing environments, changing the component's capability (e.g., for immunization) in runtime is not possible, especially when the source code is not available (which is not an uncommon situation).

In this work we have applied our ideas to real systems whose components have internal failures and are vulnerable to cyber attacks. We addressed the importance of autonomous administration of components where single administration is not feasible. In a mission critical system, when integrating local components with other COTS components and dynamically downloaded components from another administration, we must ensure that the remote components are safe from failures and cyber attacks. Our approach combines and extends the techniques of fault injection analysis and runtime-code instrumentation to provide component recovery and immunization. We have introduced the system architectures in detail and developed new mechanisms that can be used in many different mission critical systems. Finally, we have developed separate prototypes in two most popular platforms (i.e., Java and .Net) to prove the feasibility of our ideas.

Yingrui Yang/Rensselaer Polytechnic Institute

Human decision making is really a two-stage process: to form an appropriate decision problem and then to solve it, and often needs to work iteratively. Most current decision theories focus on Stage 2 of the decision process but neglect Stage 1 (Joyce, 1999); consequently, the so-called “small-grand world” problem (SGW) has remained an open question since Savage (1954). This effort investigated mental decision logic (MDL) of the SGW problem, which models reasoning processes underlying Stage 1 decisions. We discovered how MDL works in modeling the SGW problem; the idea is to use domain-specific mental predicate-argument structures in transforming between act-state structures. We show this treatment explains framing effects by an example. We developed an arithmetization of MDL by using the Gödel number method. It shows that the SGW functions are consistent but asymmetric. Also, we defined a game-theoretic version of the SGW problem, and formulated illusory mathematical expectation and illusory equilibrium based on the idea of MDL. The notion of field of game is introduced, which allows pseudo-collaboration in non-collaborative games, and pseudo-competition in cooperative games under what we call weak reasoning force. The tension of a weak game is based on the distance between illusory equilibrium and illusion solution. As an example of a possible application, we investigated the idea of “cognitive economics” and “cognitive finance” (in a sense different from behavioral ones).

Sibabrata Ray/University of Alabama

We researched the problem of server replication and placement against a resource-conscious single adversary. We want to expand this research to the case of multiple adversaries (with varying degrees of animosity against each other) and possible shared resources.

Servers, in the context of our effort, are softwares running on a network of computers and providing important services. In an adversarial environment, the adversary is capable of compromising computers on the network and causing the servers to malfunction. The adversary is assumed to introduce different types of faults. Two fault models are generally considered important, namely, the *benign fault* model and the *Byzantine fault* model. Under the benign fault model, a compromised computer ceases to function. Under Byzantine fault model, a compromised computer runs Trojan horses instead of correct servers.

The general technique for protecting a system is *voting with replication and placement* of servers. Replication means creating many copies of a server and running the copies (replicas) on many different computers in the expectation that the adversary will find it hard to compromise all/many of the replicas and therefore the server will continue to function correctly. Placement means running some or all of the replicas of the critical tasks on relatively well protected machines. Voting means collecting the output from all or many replicas and accepting the majority agreement as a correct result.

Current research on fault tolerance primarily concentrates on voting protocols and provides little guidance for replication and placement beyond: 1) create enough replicas so that every replica may not be compromised (for benign fault); 2) create enough replicas so that only less than half of them may be compromised and then take majority vote (Byzantine faults with perfectly secured vote taker); or (3) create enough replicas so that only less than a third of them may be compromised and then take majority vote (Byzantine faults with no perfectly secured vote taker). While such guidelines are good when isolated servers are considered against well-known/weak/simple adversaries, those are not enough for protecting a set of servers against a stronger and intelligent adversary.

3.0 Graduate Students

The students supported under this effort were all enrolled in a graduate education program, demonstrating excellent academic accomplishment, they included:

2003 Grad Students

- Ronny Lewin
- Gavin S. Page
- Joshua M. Surman
- Joseph Patrick Dib

2004 Grad Students

- Jason Gibbs
- Jennifer Vitalbo
- Joel D. Landis
- James M. Metzler
- Gavin S. Page
- Joshua M. Surman

4.0 Continuing Research Projects

This initiative was intended to allow AFRL scientist and engineers to identify and support the continuation of outstanding faculty research projects begun during the summer. Due to limited funding, the continuation efforts proved to be very competitive and sought after. Eight continuation projects were supported following summer 2003 and eleven projects were supported following summer 2004, they included efforts by:

2003 Extension Grants

- Peter Chen
- Hongbin, Li
- Lixin Gao
- Qiang Ji
- Sibabarta Ray
- Jingyuan Zhang
- Nael Abu-Ghazaleh
- Shuqun Zhang

2004 Extension Grants

- Biao Chen
- Hongbin Li
- Joon Park
- Peter Chen
- Robert Erbacher

- Sibrabrata Ray
- Thomas Hartrum
- Tarek Taha
- Yingrui Yang
- Frank Moore
- Shuqun Zhang