**AFRL-IF-RS-TR-2005-366**
**Final Technical Report**
**October 2005**

# AGENT-BASED SYSTEMS ENGINEERING

**Dartmouth College**

**Sponsored by**
**Defense Advanced Research Projects Agency**
**DARPA Order No. K544**



**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-2005-366 has been reviewed and is approved for publication

APPROVED: /s/

FRANK H. BORN
Project Engineer

FOR THE DIRECTOR: /s/

WARREN H. DEBANY, JR., Technical Advisor
Information Grid Division
Information Directorate

# REPORT DOCUMENTATION PAGE

*Form Approved*
*OMB No. 074-0188*



**1. AGENCY USE ONLY (Leave blank)**

**2. REPORT DATE**
OCTOBER 2005

**3. REPORT TYPE AND DATES COVERED**
Final Aug 00 – Mar 05

**4. TITLE AND SUBTITLE**
AGENT-BASED SYSTEMS ENGINEERING

**5. FUNDING NUMBERS**
C - F30602-00-2-0585
PE - 62301E
PR - TASK
TA - 00
WU - 07

**6. AUTHOR(S)**
George Cybenko and
Daniella Rus

**7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)**
Dartmouth College
Thayer School of Engineering
Hanover New Hampshire 03755

**8. PERFORMING ORGANIZATION REPORT NUMBER**

**9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)**
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington Virginia 22203-1714

AFRL/IFGA
525 Brooks Road
Rome New York 13441-4505

**10. SPONSORING / MONITORING AGENCY REPORT NUMBER**
AFRL-IF-RS-TR-2005-366

**11. SUPPLEMENTARY NOTES**

AFRL Project Engineer: Frank H. Born/IFGA/(315) 330-4726/ Frank.Born@rl.af.mil

**12a. DISTRIBUTION / AVAILABILITY STATEMENT**
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

**12b. DISTRIBUTION CODE**

**13. ABSTRACT *(Maximum 200 Words)***
This project combines robust and proven concepts from traditional mathematical systems engineering with the technology of web-based agent systems, leading to new modeling paradigms and technical results for agent-based computing. The main goal of this project is to develop a scientific approach to agent-based computing using concepts and paradigms from classical systems and computer engineering. The techniques are applied to multi-agent system coordination and adaptation problems that arise in UAV swarm surveillance and scheduling applications.

**14. SUBJECT TERMS**
Agent-Based Systems, Surveillance, Multi-Agent Systems

**15. NUMBER OF PAGES**
14

**16. PRICE CODE**

| 17. SECURITY CLASSIFICATION OF REPORT | 18. SECURITY CLASSIFICATION OF THIS PAGE | 19. SECURITY CLASSIFICATION OF ABSTRACT | 20. LIMITATION OF ABSTRACT |
|---|---|---|---|
| UNCLASSIFIED | UNCLASSIFIED | UNCLASSIFIED | UL |

## Table of Contents

**Technical Summary**

The DARPA TASK program's Agent-Based Systems Engineering project, performed at Dartmouth College, has made significant progress towards developing a quantifiable, scientific theory of agent and multi-agent based systems. The project was framed as a collection of four hypotheses which are recapped below along with the major findings and accomplishments associated with each.

**Hypothesis 1:** An agent taxonomy based on systems engineering concepts such as observation, modeling and prediction together with the new concept of brokering is a comprehensive foundation for agent-based systems engineering.

**Proposed Evaluation:** We will complete a detailed taxonomy of agent functionality as outlined below. It will be demonstrated that a variety of software agent systems can be decomposed and categorized using that taxonomy.

**Project Findings and Accomplishments:** We successfully developed this taxonomy and implemented several software systems to demonstrate its effectiveness. The taxonomy is summarized below.

*Information agents:* These agents are designed to monitor databases, locate and extract information from web and internet sites, filter raw sensor data feeds for relevant information and so on. Information agents obtain observations of the real world system state.

*Modeling agents:* These agents use observations of the real world state to estimate the real world state. Such agents perform, for example, tracking, data fusion, estimation and interpolation. A weather prediction system, wrapped as an agent, would fall into this category. Agents that correlate data from different sources use a model of the world based on some model from a knowledge base as the basis for the correlation.

*Planning agents:* These agents take a current estimate of the world state and produce a plan or course of action. Such agents can task other agents for example to find additional information or data, can alert a human user to a certain situation, can operate a fire control unit and so on. In the DARPA CoABS program, planning agents take logistics and geographical data and produce evacuation plans for a region, for instance. In the commercial domain, planning agents can suggest a purchase to a user or even execute that purchase on the user's behalf.

*Brokering agents:* This type of agent is unique to the software agent domain and has no analogue in classical systems engineering. Brokering agents locate different types of agents dynamically, create a possibly unique system realization and facilitate interoperability between the different constituent agents. In classical control and systems engineering, everything is "hardwired" so that resource discovery and interoperability are not an issue.

We implemented a variety of Information Agents, including sensors based on acoustics, network security systems, video imagery and computer process logs. The corresponding Modeling Agent work has been consolidated into a single generic software technology which we are calling *Process Query Systems* about which more is said below. The Planning Agent activity has been focused on a top-down agent design methodology in which desired collective behaviors of a multi-agent system are defined and mathematical principles are used to derive decentralized control strategies which each agent implements to achieve the desired emergent behavior or effect. This has been demonstrated in a variety of settings including multiple UAV motion planning and coordination, distributed surveillance, and sensor registration.

Initially, the project studied the IVRS REF in detail and documented findings in a report that summarizes previous work on intelligent vehicles and roadways with a focus on what different components and relationships have been proposed for a centralized, hierarchical solution. The report categorizes the different components as agents of different types within our taxonomy, as outlined above, and demonstrates that solutions to the IVRS problem can in principle be decomposed into information agents, modeling agents and planning agents quite naturally.

We initiated a review of decentralized algorithms that produce desired global, centralized solutions to optimization problems especially as they might occur in the various REFS. We identified several major categories of techniques. Among them was a method based on approximations to centralized gradients by decentralized information. The idea is quite simple but appears to be relatively powerful and generally applicable to several of the problems being considered in the REFS.

Specifically, consider a globally defined objective function for a distributed agent based system that requires global information to compute. Each coordinate of the underlying variable space is a feature (location, velocity, etc) of an individual agent within the system. A gradient descent method for minimizing this objective function requires computation of partial derivatives of the objective function with respect to the various agent variables. This requires global information and communication.

We observed that in some problems within the Intelligent Vehicle and Road System (IVRS) REF, the global objective function has terms which involve strong local interactions and weaker interactions at a distance. Using the strong interactions, which can be communicated using only local information exchanges, an approximation to the true gradient can be computed. That approximation can be shown to have a positive projection onto the true gradient and therefore serves as a descent direction for minimizing the global objective function. The approximation property and associated noisy computation of the gradient can be folded into a powerful analytic machinery developed by Bertzekas and Tsitsiklis during the 1990"s for stochastic optimization. The key concept is that local interactions are strong and interaction strength diminishes with distance. Distance can be a conceptual metric, not necessarily spatial distance.

Another direction that has been pursued actively is the coordination of multiple agent systems using decentralized, asynchronously communicated information. Examples of such coordination include resource control and apportioning of fixed assets. In the IVRS problem for example, this can be equal spacing of a platoon of vehicles using purely local information sharing. This circle of ideas appears to be a completely new approach in decentralized computation using ideas from matrix iterative analysis, linear programming and graph theory in remarkable tight and powerful ways.

The methodology was demonstrated on a 2D distributed agent coordination problem which we called the "2D Opera Problem" which required the distributed control of mobile physical agents with the goal of exiting a region while maintaining inter-agent separation. The 2D Opera Problem however allows vehicle dynamics that are not feasible in 3D. We identified a promising 3D version that involves a hybrid approach which involves fine and coarse levels of control. Coarse control is based on a cellular partition of 3D space. Control at this level uses idealized dynamics. At the coarse level, 3D dynamics are similar to 2D dynamics in that objects can 'stop' and maneuver without regard to 'fine' dynamic constraints. Vehicles/aircraft within such a coarse cell are subject to specific dynamic constraints and use classical control techniques. Control at the coarse level is governed by a 3D Opera Problem approach, namely coarse controls to move between cells as based on the descent algorithms we used in 2D. Within a cell and maneuvers between cells are governed by specific, realistic vehicle dynamics. The stability of this hybrid control scheme was explored and documented.

**Hypothesis 2:** This framework will allow the quantification of agent and agent-system performance in a scientifically rigorous manner.

**Proposed Evaluation:** We will demonstrate that information, modeling and planning agent performance can be characterized quantitatively by introducing metrics analogous to traditional systems engineering but appropriate for agent-based computing domains. We will continue development of agent computing performance, already initiated in the DARPA CoABS Program, and explore joint semantic-computing performance metrics and tradeoffs.

**Project Findings and Accomplishments:** We have successfully quantified agent situational awareness through the Process Query System technology which was started in this project and has continued to be applied on other efforts at Dartmouth and elsewhere. The basic idea behind Process Query Systems is that data streams being collected from distributed, noisy sensors are fused according to multiple process models of the environment in which the system is operating. That is, the designer of a Process Query Systems application articulates a collection of process models that describe the environment. For example, in a battlefield situational awareness setting, this would include kinematic process models of vehicles and or other dynamic objects in the region. Similarly, in the network security domain, the process models are models of network and host attacks, infrastructure failures and so on. The reader is encouraged to consult the publications about Process Query Systems listed in the Publications section for further details.

A generic situational awareness metric, called (alpha,beta)-currency was defined, developed, implemented and evaluated in several simulation scenarios. The two parameters of (alpha,beta)-currency are defined as the probability alpha that the state estimate correct to within a temporal grace period of beta. For example, as beta increases to infinity, this metric means we have alpha probability of knowing the situation state a long time ago, which is easier as information accrues and ambiguities are resolved. As beta gets smaller, our temporal requirement increases in that alpha probability of correctness is demanded of recent states. A similar analysis of the alpha parameter shows that alpha close to one means high accuracy is required or possible while alpha close to zero means that confidence in the beta current state estimate is relaxed.

We also investigated problems of target surveillance with the aim of building a general framework for the quantitative evaluation of the performance of a system of autonomous agents. To this end, we designed a class of semi-distributed stochastic navigation algorithms, that drive swarms of autonomous scouts to the surveillance of grounded targets, and we developed a novel approach for performance estimation based on analyzing sequential observations of the system's state with information theoretical techniques. We studied the interrelations between randomness, resource consumption (fuel) and ergodicity of a decentralized control system in which the decision-making process is partially stochastic.

We investigated the role of very simple and noisy sensors for tracking problems. We proposed a simple binary sensor model, where each sensor's value is converted reliably to one bit of information only: the object is moving toward the sensor or away from the sensor. We showed that a network of binary sensors has geometric properties that can be used to develop a solution for tracking with binary sensors and developed algorithms and simulation experiments. We developed two classes of algorithms: one that assumes the sensors have no range limits and another that limits the sensor range. Our extensive simulations showed low error that decreases with sensor density.

**Hypothesis 3:** The taxonomy introduced above will lead to more efficient, scalable and robust brokering so that agent discovery, reuse and interoperability will be enhanced.

**Proposed Evaluation:** We will develop agents in accordance with this taxonomy and disseminate the taxonomy and agent-based systems engineering paradigm to other DARPA Agent-Based Computing contractors. The effectiveness of our approach to agent-based system engineering will be measured in DAML, TASK and CoABS integration experiments.

**Project Findings and Accomplishments:** The project developed and evaluated a software system called a Semantic Message Oriented Middleware (Semantic MOM) that uses semantic information to index, locate and route messages between participating agents. Work on context-based routing has also been performed and published.

The publish/subscribe paradigm of Message Oriented Middleware provides a loosely coupled communication model between distributed applications. Traditional publish/subscribe middleware uses keywords to match advertisements and subscriptions and does not support deep semantic matching. To this end, we designed and implemented a Semantic Message Oriented Middleware system to provide such capabilities for semantic description and matching. We adopted the DARPA Agent Markup Language and Ontology Inference Layer, a formal knowledge representation language for expressing sophisticated classifications and enabling automated inference, as the topic description language in our middleware system. A simple description logic inference system was implemented to handle the matching process between the subscriptions of subscribers and the advertisements of publishers. Moreover our middleware system also has a security architecture to support secure communication and user privilege control.

Additionally, we examined the role of information in markets that allocate computation to software agents. A comparison of two types of markets illuminated the importance of information and the incentives for buyer agents and seller agents to share their preferences with each other. In our comparison, the distinguishing feature of the two markets types was the alignment of agents' interests. We define a closed-interest market as one where resources are collectively owned among the agents. An open-interest market makes no assumptions on the interests of agents or resource owners.

The incentives of agents in the two markets drastically differ. The open-interest model motivates agents to be less trusting and to not share information. This aspect stems from the model's greater applicability to resource allocation, but has a deep impact on system efficiency. The project summarized some economic theory and allegorical evidence from these models and system implementations that support the claim, and developed with guidelines for system development.

**Hypothesis 4:** Hybrid systems can be built using both traditional systems engineering and agent-based computing components. Such systems will significantly enhance U.S. capabilities in countering asymmetric threats.

**Proposed Evaluation:** We will demonstrate a hybrid system using both traditional systems engineering techniques and agent-based computing. A candidate domain could involve combining satellite tracking (based on traditional signal processing and target tracking techniques) with human intelligence reporting.

**Project Findings and Accomplishments:**

Several demonstrations, simulations and experiments were performed to support this part of the project.

We investigated a problem arising in decentralized registration of sensors. The application we consider involves a heterogeneous collection of sensors - some sensors have on-board Global Positioning System (GPS) capabilities while others do not. All sensors have wireless communications capability but the wireless communication has limited effective range. Sensors can communicate only with other sensors that are within a fixed distance of each other. Sensors with GPS capability are self-registering. Sensors without GPS capability are less expensive and smaller but they must compute estimates of their location using estimates of the distances between themselves and other sensors within their radio range. GPS-less sensors may be several radio hops away from GPS-capable sensors so registration must be inferred transitively. Our approach to solving this registration problem involves minimizing a global potential or penalty function by using only local information, determined by the radio range, available to each sensor. The algorithm we derive is a special case of a more general methodology we have developed called 'Emergence Engineering' that was described above already.

Dartmouth demonstrated several unique capabilities combining all of these findings. They involved decentralized sensing, situational awareness estimation and UAV scheduling based on the technologies developed in this effort. One such effort is shown on the last quad chart developed for the project in 2004, copied on the final page of this report.

One such experiment involved unmanned aerial vehicles (UAVs) that were employed in aerial reconnaissance military missions. Traditionally, these flying vehicles are operated by a fully centralized navigation control. The mode of deployment is to launch them remotely, and fly them towards the target region, passing over geographic barriers. A natural evolution of this tactical employment of UAVs is to deploy a large number of UAVs to carry out missions of surveillance and target tracking. Our work proposed a fully decentralized control system for the navigation of UAVs on patrol of a given area with the task of surveillance and tracking of unknown (possibly mobile) targets. The obtained solution was intended to provide a large fault-tolerant autonomous system composed of fully automated, unmanned, easy-to-deploy, and cheap (expendable) vehicles.

We presented a nontrivial robotic implementation of the distributed algorithm previously introduced for the following:

**Instance**: A battlefield consisting of a terrain and a set of potential targets that move along unknown trajectories.

**Problem**: Deploy a number of scouts (flying or grounded vehicles) and a number of fixed sensors having limited sensing and communication range and bandwidth with the following objective and optimization requirements:

• Objective: Know number and position of targets over time

• Optimization/Trade-off: Steer UAVs along unpredictable trajectories minimizing the consumption of fuel per cell visit.

Our initial solution was devised for a scenario in which all the UAVs were flying at different altitudes to prevent collisions and there were no fixed sensors. In the actual demonstration, UAVs were substituted with land-based agents (wheeled iRobots), and there were also fixed sensors that were treated as stationary UGSs. In other words we assumed that UAVs and fixed sensors possess similar computing and sensing capabilities.

Also both their sensing and communication range are assumed to be limited. Our approach is based on a top-down methodology developed to design provably performant agent-based control systems. This was a first response to the problem of supplying agent-based technology with a general unified model for performance estimation and is based on a three-stage process:

1. The first stage of the design is characterized by the assumption that each autonomous subcomponent of the system has full knowledge of the global state. This means that we looked for a centralized solution. Specifically, we produce a navigation algorithm that is run locally by each UAV but with global knowledge of all the events that occur in the large area being monitored. One way to simulate this is by assuming the existence of an omniscient super monitor (e.g. a satellite with unbounded transmission range and bandwidth) that records all the events occurring locally and allows all the UAVs and sensors to learn about them.
2. The second stage is to remove the supermonitor and replace global resources with local resources compatible with the constraints of the given distributed environment. This introduces communication issues due to the fact that now vehicles must acquire information through interactions with other locally reachable components in order to infer the most likely global state of the system. The result of this stage is a navigation algorithm that runs locally and uses local information. It records all the events that occur within range and maintains a representation (hypothesis) of the global state. When two UAVs/sensors come into contact, an exchange of knowledge and experience takes place on a peer-to-peer basis.
3. The final stage consists of calibrating the obtained distributed solution. We want to determine the minimum amount of information that needs to be exchanged between peers in order to keep the uncertainty on the global state (relative to the ideal condition of complete centralization) below a certain bound, which in turn depends on a global objective that the entire monitoring system needs to fulfill. An example of global objective is a function that quantifies how much the system knows about the position of the unknown targets over time. As exact knowledge of the position of the targets at any time is unrealistic, given that we deploy a finite number of scouts with limited radius of visibility, we are forced to accept a margin of uncertainty that needs to be properly quantified. This is attained through the notion of (alpha,beta)-currency that was already successfully employed in the framework of Web search engines. (A crucial problem in that context would be the determination of the time needed for a search engine to re-index the entire Web in order to maintain an "up-to-date" database of pages.) Alternatively we defined a risk function that instead quantifies the cost of not knowing enough about the position of the targets over time.

The above concepts were demonstrated and the TASK Program final demonstration day in Washington, DC in the summer of 2004.

**Personnel Supported (all personnel were at Dartmouth originally):**

| Name | Role | Present Position |
|---|---|---|
| George Cybenko | Co-PI | Professor, Dartmouth |
| Daniela Rus | Co-PI | Associate Professor, MIT |
| Valentino Crespi | Postdoc | Assistant Professor, Cal State LA |
| Jonathon Bredin | Ph.D. student | Assistant Professor, Colorado College |
| Guofei Jiang | Postdoc | Group leader, NEC Research Princeton |
| Daniel Bilar | Ph.D. student | Assistant Professor, Colby College |
| Han Li | MS student | Software engineer |
| Alex Jordan | MS student | Research scientist, BAE Systems |
| Robert Gray | Postdoc | Research scientist, BAE Systems |
| Wayne Chung | MS student | Ph.D. student, Dartmouth College |
| Jay Aslam | Faculty | Associate professor, Northeastern |
| Zack Butler | Ph.D. student | Assistant professor, RIT |
| Ron Peterson | Research staff | Dartmouth College |

**Publications:**


J. Aslam, Z. Butler, F. Constantin, V. Crespi, G. Cybenko, D. Rus. Tracking a moving object with a binary sensor network. Proceedings of the ACM SenSys'03 Conference, November 5-7, 2003, Los Angeles.

V. Berk, W. Chung, V. Crespi, G. Cybenko, R. Gray, D. Hernando, G. Jiang, H. Li and Y. Sheng. Process Query Systems for Surveillance and Awareness. Proceedings of the 7th World Multiconference on Systemics, Cybernetics and Informatics (SCI 2003), July 27-30 2003, Orlando, Florida.

Luis Caffarelli, Valentino Crespi, George Cybenko, Irene Gamba and Daniela Rus. Stochastic Distributed Algorithms for Target Surveillance. Proceedings of the conference: Intelligent Systems and Design Applications (ISDA2003), to be held in Tulsa, Oklahoma, August, 2003.

Valentino Crespi and George Cybenko. Decentralized Algorithms for Sensor Registration. 2003 International Joint Conference on Neural Networks (IJCNN2003), Portland, Oregon, July 2003.

A. Das, G. Kantor, V. Kumar, G. Pereira, R. Peterson, D. Rus, S. Singh, J. Spletzer. Distributed Search and Rescue with Robot and Sensor Teams. In Field and Service Robotics 2003.

Q. Li, M. deRosa, D. Rus. Distributed algorithms for guiding navigation across sensor networks. Mobicom 2003.

V. Crespi, G. Cybenko, D. Rus, M. Santini. Decentralized Control for Coordinated flow of Multi-Agent Systems. Proceedings of the 2002 World Congress on Computational Intelligence. Honolulu, Hawaii, May 12--17, 2002.

Q. Li, J. Aslam, D. Rus. Distributed Energy-Conserving Routing Protocols for Sensor Networks. IEEE Hawaii International Conference on System Science, 2002.
R. Peterson and D. Rus. Interacting with a sensor network. ACRA 2002 (November 2002).

P. Corke, R. Peterson and D. Rus. Flying robot navigation using a sensor net. 2003 International Symposium on Experimental Robotics.

Valentino Crespi, Wayne Chung and Alex Jordan. Decentralized Sensing and Tracking for UAV Scheduling. Proceedings of the SPIE Defense and Security Symposium (conference 5403-62), Orlando, FL, April, 2004.

George Cybenko, Vincent Berk, Valentino Crespi, Robert S. Gray and Guofei Jiang. An Overview of Process Query Systems. Proceedings of the SPIE Defense and Security Symposium (conference 5403-129), Orlando, FL, April, 2004.

Diego Hernando and Valentino Crespi. Sampling for Process Detection with Applications to Surveillance and Tracking. Proceedings of the SPIE Defense and Security Symposium (conference 5403-25, Session 5), Orlando, FL, April, 2004.

Valentino Crespi, Wayne Chung, George Cybenko, Alex Jordan.  Distributed Sensing and UAV Scheduling for Surveillance and Tracking of Unidentifiable Targets, SPIE Defense and Security Symposium, March 2005, Orlando, FL.

DARPA
TASK

# Agent-Based Systems Engineering

8/2004

**GOAL: Distributed Coordination and Control of Autonomous Multi-Agent Systems**

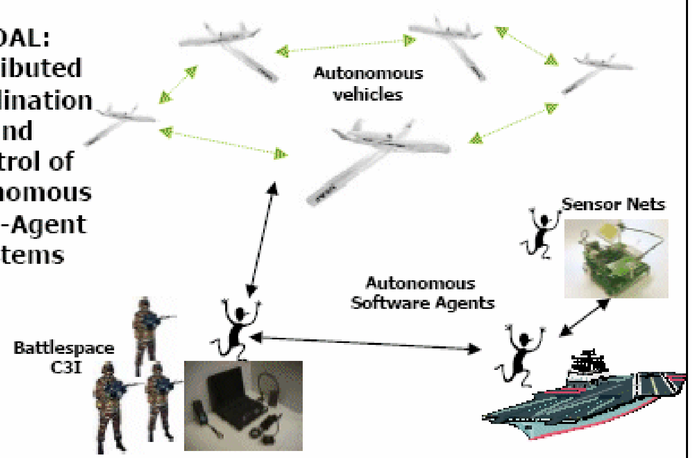


## New ideas

- Process Query System for tracking
- Engineering systems theory applied to agents
- Top-down theory, bottom-up implementation
- Tractable and Automated Performance Analysis
- Semantic services and agents
- Use of powerful stochastic optimization results

## Revolutionary Impact

- On large-scale, autonomous software systems
  - Resource control, stability and performance
- On large-scale autonomous physical systems
  - Decentralized, effective coordination and control
- On technology for decentralized system operation
  - Mathematically sound, computationally efficient

## Schedule

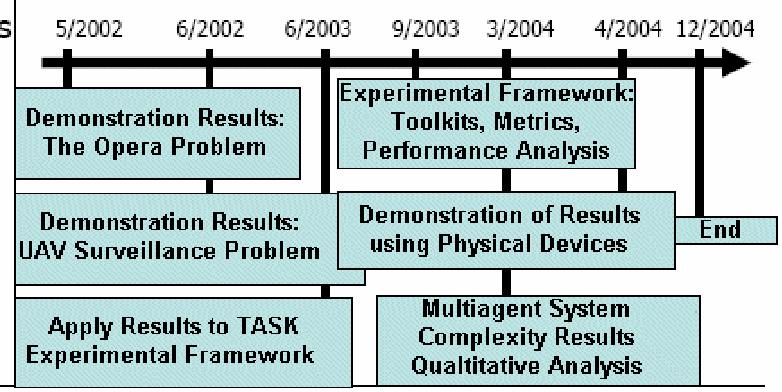


Dartmout College PI: George Cybenko