

Development of an Autonomous Aerial Vehicle: A Case Study

NELSON C. BAKER

School of Civil Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0355

DOUGLAS C. MACKENZIE

College of Computing, Georgia Institute of Technology, Atlanta, GA 30332-0280

STEPHEN A. INGALLS

School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0150

Received December 2, 1991, Revised February 15, 1992

Abstract. This paper describes the design, development, and deployment of an unmanned autonomous aerial vehicle developed at the Georgia Institute of Technology during 1990–1991. The approach taken, the system architecture, and the embedded intelligence of the project as conceived by a team of students, faculty, and industrial affiliates is reported. The project focused on engineering a vehicle which performed an intended mission in the time, space, and weight restrictions specified as part of an AUVS 1991 Competition. This paper documents the system and its various components and also provides a discussion of integration issues.

The project demonstrated capabilities of existing and new technologies, but also highlighted many serious integration issues, particularly when using prototype components. The project also demonstrated the utility and mutual benefits of academic-industry projects. All members of the team benefited by working on a real and tangible project. Industrial participants gained first hand experience integrating their products with other components and many saw potential for their products and services in new markets.

Key words: Autonomous system, robot, aerial robot, concurrent engineering, case study

1. Introduction

Beginning in September, 1990, a team of Georgia Tech students and faculty along with several industrial partners set out to develop an unmanned autonomous flying vehicle. Vehicle specifications were dictated by the rules of a student competition organized by the Association for Unmanned Vehicle Systems (AUVS) held in July, 1991. The assembled team, multidisciplinary by necessity, engineered a vehicle which comes close to achieving the mission objectives. This paper describes the process by which the system was developed, the organization of the system, the individual components, and the software which makes the vehicle perform in an intelligent manner.

1.1. Overview

The AUVS Competition led to the creation of the Georgia Tech team and the collective interest in developing an autonomous aerial vehicle. Many of the design decisions discussed in this paper were made to meet requirements and mission objectives stated in the competition rules and guidelines [1].

The AUVS Competition prescribed that a team must use a vehicle which was completely autonomous and untethered. The vehicle was required to fly in an arena the size of a volleyball court containing two, 6 foot diameter bins. One of the bins contained three inch diameter orange metallic disks, randomly placed, which were to be retrieved and deposited in the second bin. The

objective was to fly a course which found a disk in the source bin, pick it up, and fly to the deposit bin where the disk was relinquished. This cycle was completed six times, transferring all of the disks from one bin to the other bin. The mission duration was three minutes with an additional six minutes allotted for setting up and starting the vehicle. All vehicles were required to provide a "kill" switch which transferred control to a human pilot in case of emergency.

A secondary mission for the development of the Georgia Tech vehicle was to provide a platform for future research and development. The experience gained via the development of the competition aircraft would provide an avenue for future research activities.

In order to perform the mission, the vehicle had to replace the capabilities of a human pilot. Thus, the measure of intelligence by any system or component was the degree to which it reduced the need for human control or augmented a human pilot's performance. The synergistic effects caused by the collection of components and their individual functionality developed the necessary intelligence for the system to completely replace the pilot for limited flight operations, i.e., those required by the AUVS mission.

1.2. Motivation

The AUVS mission, requiring an autonomous agent, is similar to missions proposed by various organizations and is one that has received increasing attention [2, 3]. These kinds of robotic agents, for example, are useful for surveillance operations, communication relays, and inspection duties for pipelines, large bridges, and hazardous environments. Any intelligence which can be added to a device which reduces the level of detailed control required by the operator, greatly increases the system's value and is beneficial. Therefore, each of the separate technologies and components on a vehicle become important and issues revolving around the integration of manipulation, mobility, real-time control, intelligent control, sensor fusion, and visualization technologies become a major undertaking and worthy of serious research. Thus, the motivation for our entry into the AUVS competition was more than just to compete in the event, but also

to develop a testbed for further research in various components needed by autonomous agents and the investigation of integration issues.

1.3. Paper Organization

This paper is organized so as to describe the various activities and components of the aerial vehicle system which lead to an intelligent agent. The discussion first begins with a description of the concurrent engineering activities that were performed. This phase identified key components and technologies required for successful completion of the project. Following the concurrent engineering activities, a well defined team was formed and its organization contributed in part to the solution developed. An overview of the system is provided and then descriptions of each of the major subsystems are presented. The paper then ties the subsystems together by discussing the integration issues faced by the project. At the conclusion of the paper is an evaluation of our approach to developing the system and some comments about the future capabilities.

2. Problem Statement

The team's initial commitment to compete in the competition defined the short term goals for the project. These can best be identified by discussing more specifically the competition and the goal's impact on formulating the scope of the robotic task.

2.1. Mission

The mission task to fly autonomously within a limited space and perform a complex mission in a very short period of time was a significant challenge. The team's goal was to retrieve all six of the disks within the three minute flight time. Therefore, design decisions were based on the fact that disk acquisition and flight time must fall within the total allotted mission time.

2.2. Competition Specifics

Constraints were placed on candidate vehicles by the competition rules. The vehicle chosen to per-

form the mission had to be capable of fitting inside a 6 foot cube. No tethers were allowed to the vehicle, so significant communication issues and discussions of on-board versus off-board equipment ensued. The flight path of the vehicle also had to contend with a three foot high opaque barrier which bisected the course (see Figure 1). The team decided early in the project that a vehicle satisfying all competition requirements was achievable, thus influencing many subsequent design decisions.

3. Concurrent Engineering Approach

Concurrent Engineering (CE) is the systematic application of scientific and engineering principles to the concurrent design of all aspects of the product and related processes. This includes meeting budget, schedule, and technical performance objectives while ensuring product functionality, maintainability, and manufacturability [4]. The large scale multidisciplinary development required to build a winning vehicle in the short time frame allotted, highlighted the benefits of applying concurrent engineering to this project. Application of CE techniques has been shown to reduce product development cycles by as much as 50% [4].

Five system documents were created in an effort to implement CE techniques. First, the Systems Engineering Management Plan (SEMP) was created, modeled after the SEMF issued by the Department of Defense's Joint Project Office for Unmanned Aerial Vehicles [5]. Common terminology to describe system components, team organization and responsibilities, a projected mile-

stone list and team management philosophy were included in the report. The Work Breakdown Structure shown in Figure 2 was modeled after the DoD's and developed to the component level (Level 4). It assisted in assigning team responsibilities. A Functional Task Decomposition was generated in conjunction with Quality Function Deployment (QFD) techniques [6]. Figure 3 shows the QFD planning matrix generated for the project. This matrix, sometimes called the "House of Quality", helped identify critical competition tasks and their relationship to customer (AUVS competition) requirements early in the design process. The matrix also provided an important awareness of coupled tasks as denoted by the upper triangular portion of the figure. The fifth document, an Integrated Schedule, presented explicit subsystem milestones necessary for timely completion of the system's integration and evaluation. Using the concurrent engineering process was intended to force earlier and better design decisions, resulting in fewer subsystem incompatibilities than would otherwise have occurred.

The most troubling problem in utilizing any of these tools was the inability to predict or meet schedules. This is an ongoing difficulty in industry, but academic projects suffer more due to the team members' conflicting time requirements. A complete discussion concerning the concurrent engineering activities associated with the project can be found in a related Master's thesis [7].

4. Organizational Structure

To properly address the design and development issues in a project of this magnitude and scope over such a short duration, some type of organizational structure had to be imposed. Initial goals for the project were to implement the concurrent engineering design environment through the use of the documents described above and then to start addressing critical design issues. Participants composing the team included students and faculty from the Schools of Aerospace Engineering, Civil Engineering, Electrical Engineering, Mechanical Engineering, and from the College of Computing. The team also had significant participation from industrial affiliates, Pacific RPV,




Fig. 1. Competition arena

Level 1	Level 2	Level 3	Level 4
Aerial Robot System, part a	Aerial Vehicle	Airframe	Fuselage Structure
			Alighting Gear
			Payload Platform
		Propulsion	Fairings
			Main Rotor
			Tail Rotor
			Engine
			Drive System
		Guidance & Flight Control	Fuel Control System
			Sensor Suite
	Stability Augmentation System		
	Power Generation & Distribution	Actuators	
		Electrical Bus	
		Overload/Surge Protection	
		Plugs/Connectors	
		Wires/Cables/Harnesses	
		Power Source	
	Ancillary Equipment	Power Regulation	
		Storage/Transportation Platform	
	Mission Equipment Package	Vehicle Positioning System	Camera/Conic device
External Vision Cues			
Computer Programs			
Disk Locating System		Camera/Optics	
		Strobe Light	
Disk Retrieval System		Computer Programs	
		Camera/Optics	
		Strobe Light	
		Computer Programs	

Level 1	Level 2	Level 3	Level 4
Aerial Robot System, part b	Mission Planning & Control	Ground Control Processing Station	Ground Control Subsystem
			Data Processing Equipment
			Communications Equipment
			Ground-Based Data Relay
			GCS Assembly and Integration
		Support Equipment	
	Computer Programs	Outerloop Controller	
		Mission Planning System	
		Database Management	
	Data Link	Airborne Data Link	Data Uplink Receiver
			Data Downlink Transmitter
			Downlink Postoffice
		Ground Data Link	Computer Programs
			Data Uplink Transmitter
	Logistics Support Subsystem	Maintenance Planning	Data Downlink Receiver
			Computer Programs
		Software Support	Provisioning
			Repair Parts
		Transportation and Handling	Training Spares
			System Transportability
Handling and Special Equipment			
Hardware Interface		System Storage	
		Facilities	
Facilities	Production/Assembly/Modification		
	Support Facilities and Equipment		
	Test/Demonstration Facilities		

Fig. 2. Work breakdown structure

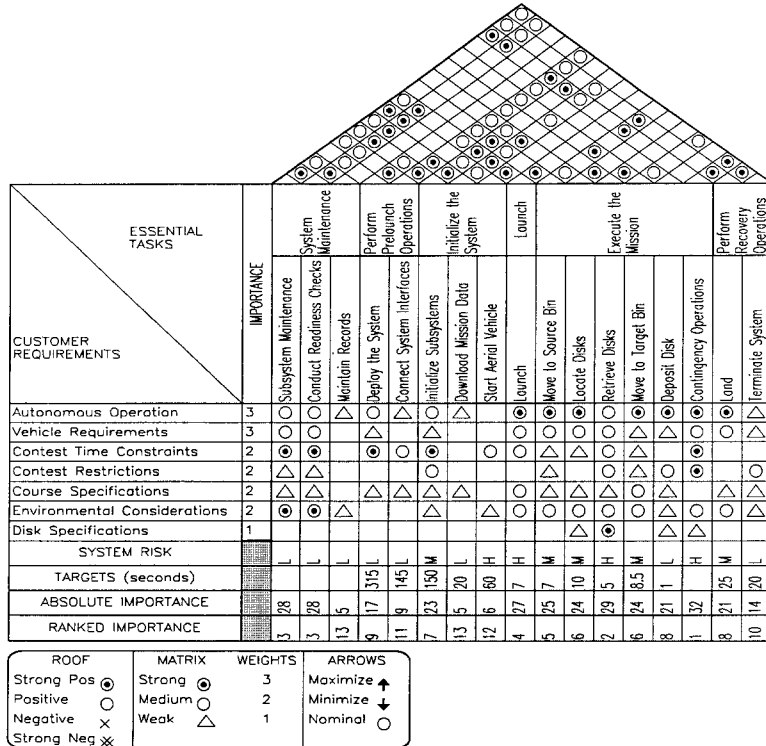


Fig. 3. Autonomous aerial robot quality function deployment planning matrix

The Mitre Corporation, Guided Systems Technologies, the U.S. Army Aerostructures Directorate, and TRON-Tek.

Initial design responsibilities of the various team members are shown in Figure 4. Towards the end of the project, however, many people moved between groups aiding in the integration of the various components. Perceived academic specialties of the represented engineering schools were aligned with technical requirements of prospective subsystems. The structure therefore allowed the groups to be uncoupled from other design activities with an underlying intent that system pieces would be integrated during the later portions of the project.

5. System Architecture

Having organized the team in a concurrent engineering design environment, development of a vehicle system began. The system was necessarily more than a collection of components, be-

cause each piece contributed, and integrated with, the rest of the package. This section of the paper provides an overview of the system as a whole while the next section will provide details concerning each of the components.

5.1. System Objectives

The driving objectives for development of an aerial vehicle system were the AUVS competition requirements described in Sections 1 and 2, coupled with the issues identified by the preliminary concurrent engineering study. These can be described using the high-level functional tasks of startup, takeoff, fly, navigate, retrieve a disk, deposit a disk, and land. Rather than using a human pilot and remotely controlled vehicle to perform these steps, the team needed to make the process autonomous within a ten month time frame. Additional objectives for the system included interchangeable components so that future research could be performed with the vehicle. A related benefit of the project was the opportunity to use

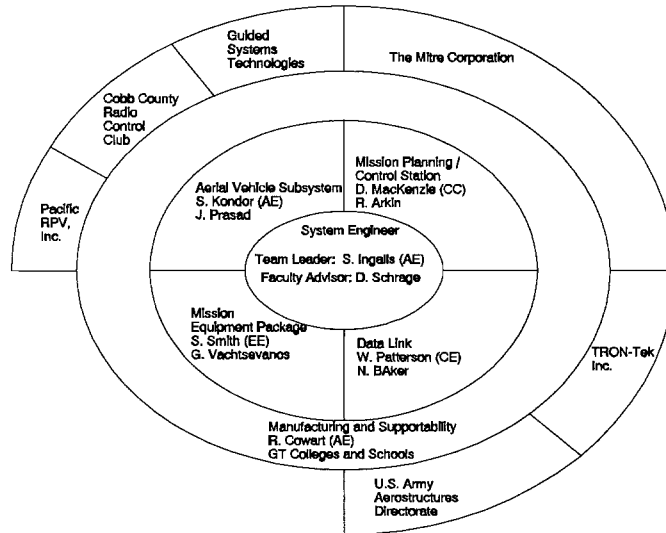


Fig. 4. Team organization

this entire design process as a concurrent engineering case-study.

5.2. System Performance Specifications

In order to achieve the mission objectives, several decisions were made early in the design scenario. These were mainly related to performance issues of the system, but influenced the design processes of most components.

First, as much as possible, reliability and functional redundancy was key to the design. Any one malfunctioning component could jeopardize the mission. Of course, redundancy issues had to be balanced against issues of aerial vehicle gross weight.

Secondly, the team's design objective during the entire process was to build a system capable of recovering all six disks. Therefore, timing issues for retrieving disks, for selecting flight velocities, and for decision making were directed toward this mission performance goal.

Flight performance of the vehicle was a third item. The ability of the aircraft to fly accurately to a desired location and hover at that location posed significant design requirements for the supporting subsystems. The retrieval mechanism could only be as good as the locational accuracy and hover stability of the vehicle. Vision systems needed to dynamically compensate for changes

in vehicle roll, pitch, yaw, and altitude. In order to achieve stable flight performance characteristics, an internal flight stability system was required. Flight control command update rates were determined to be on the order of ten times per second, a rate determined to be approximately twice that of a human pilot.

5.3. System Architecture Overview

Having defined the mission objectives and several key performance specifications, the research and development process was partitioned along functional lines and assigned to various portions of the team. Since weight and space were at a premium, initial design focused on investigating component weight and space capabilities that would result with a vehicle able to lift off the ground, and with sufficient power remaining for slow maneuvers. To reduce vehicle weight, some equipment would be placed at the ground station, relying on communication links to transmit data between the vehicle and the ground.

The system which evolved is depicted in Figure 5. The figure shows the ground-based activities in the center and aerial operations on the outer extremities. Four main airborne components existed. These were the Air Vehicle Sensor Suite, which assisted in stability augmentation; a Vehicle Positioning System which performed

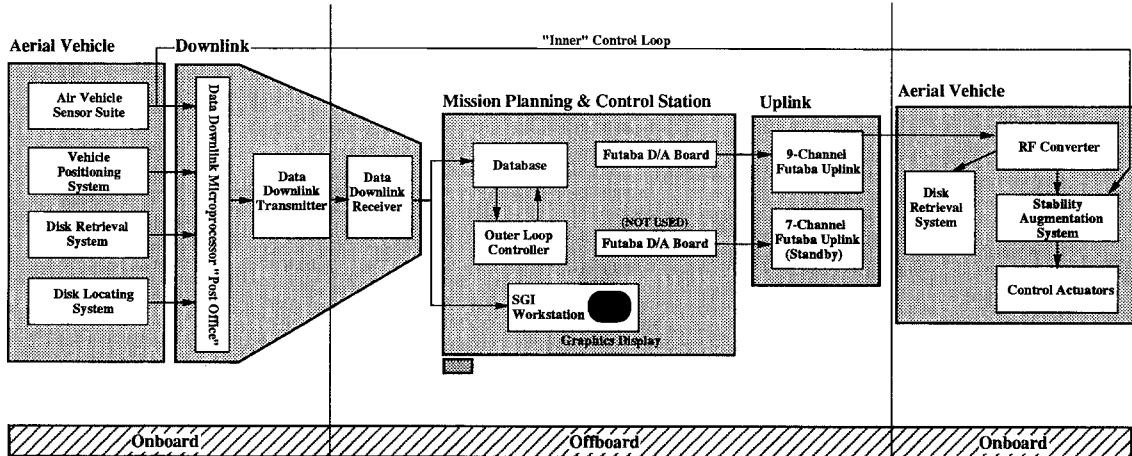


Fig. 5. Autonomous aerial vehicle system architecture

sensing and tasks related to navigation; a Disk Retrieval System for acquiring the disks; and, a Disk Locating System, which used a camera to determine disk locations in the source bin. Each of these subsystems transmitted data to the ground via a communication down-link. Onboard the vehicle, a “post office” collected the data from the four components, packaged it together, and sent the information to the ground. Once on the ground, the data was sent to the mission planning and control station for processing by the controller and for driving a ground-based visual “cockpit” display of the mission status. Vehicle command signals were then sent back to the aircraft through the uplink. These commands caused flight surfaces to move in intended directions and for the retrieval mechanism to respond in a desired manner. This path follows the processes which were to occur ten times per second during the mission.

The next section contains a description of each system component.

6. Subsystem Description

6.1. Aerial Vehicle

Competition requirements implied a vehicle capable of vertical takeoff with an ability to maintain stable hover at low altitude. The team de-

ecided that there was insufficient time to build a vehicle and that a commercially available airframe should be purchased. Three types of VTOL unmanned aerial vehicles were evaluated: coaxials, ducted fans and conventional helicopters. The commercial availability of conventional helicopters with the desired payload capability led to their selection. Further research determined that a modified version of Pacific RPV’s Bruiser helicopter, shown in Figure 6, would provide the best alternative.

Several problems with the vehicle proved to be obstacles throughout the project. The engine was very sensitive to variations in ambient conditions and required tuning before every flight before it would generate the required performance. Second, the vehicle experienced pronounced vibration problems resulting from drivetrain modifications made to have the vehicle fit within the competition’s six foot cube size requirements. This vibration load, coupled with manufacturing deficiencies in the tail rotor blade grips and engine cooling fan, resulted in higher than normal component wear and failure rates. Long down times resulted in significant flight test schedule delays.

During flight testing, the Bruiser demonstrated the ability to fly with 18 pounds of payload, exceeding the target on-board payload by nearly four pounds, offering some margin for error in gross component weights.



Fig. 6. Photograph of aerial vehicle

6.2. Air Vehicle Sensor Suite

The on-board sensor suite sampled vehicle altitude, attitude (roll, pitch, yaw), and attitude rates of change. The sensors selected were required to be sufficiently responsive and accurate to allow stable flight. Small weight, size and power requirements were also important traits in on-board systems. Unfortunately, such sensors proved to be very expensive.

A commercial three-axis attitude and rate sensor package was loaned for use in the vehicle by Watson Industries, removing a large financial concern. However, failure analysis after a vehicle mishap pointed to the sensor unit as the crash's cause. Subsequent testing revealed one of the channels was malfunctioning. After three other units of the same model, from different sources, showed similar faults, a two-axis (roll, pitch) Watson system was combined with a flux gate compass (yaw position) and gyroscope (yaw rate) to provide the required three-axis functionality.

Commercial sensors capable of accurately registering vehicle altitude over the desired range of 0 to 10 feet proved unavailable. Therefore, the team undertook development of an ultrasonic ranging based altimeter. The system used three Polaroid sensors aimed at the ground, and aver-

aged the readings from each to reduce spurious errors. Lagging development schedules hampered the component's complete integration on the aircraft.

6.3. Flight Controller

6.3.1. Overview. The flight control subsystem was responsible for moving the vehicle from its current location to a target X,Y,Z position within the arena and then maintain that position in a stable hover. To accomplish this, two closed-loop controllers were used in a hierarchical arrangement [8]. The inner loop controller resided on the vehicle while the outer loop controller spanned the radio links to execute in the ground computer. Figure 7 presents a schematic of the flight controller subsystem.

The inner control loop executed on-board the aircraft in an Intel 80C196KC microcontroller based Stability Augmentation System (SAS). The SAS controller attached directly to the sensors and control actuators. It monitored the aircraft sensor suite and manipulated the pitch of the blades on the main rotor and the tail rotor in order to control the vehicle so that sensor readings matched target values provided by the outer control loop. The outer loop was simplified by using the SAS, since the controller could model




Fig. 7. Flight controller schematic

the vehicle as a stable system, unconcerned with power fluctuation, wind gust and changing load effects.

The outer loop controller, executing in the ground computer, used vehicle sensor data and position feedback inputs. It received inputs from mission planning that provided target X,Y,Z positions which might be many feet away from the current helicopter location. The outer loop controller then generated smooth velocity curves, within performance limits, to command the helicopter, through the SAS, to the desired location. Outputs of the outer loop provided inputs to the SAS to effect this move and were in the form of vehicle attitude targets (roll, pitch, yaw, etc.) moving the helicopter along the desired trajectory.

6.3.2. *Inner Loop.* The inner control loop executing on the SAS micro-computer operated helicopter control surfaces and read sensor data. It

was given target values for the vehicle attitude and it manipulated the control surfaces to drive the measured values to match the reference conditions. Figure 8 shows a typical block diagram for one of the four channels in the inner loop controller. Separate channels controlled roll, pitch, yaw, and altitude. The controller operated asynchronously at 20–30 cycles per second. Control gains were dependent on the vehicle configuration and desired performance characteristics, and were determined experimentally through flight testing.

The SAS, specifically designed to operate this helicopter, was offered to the team as a part of the purchased vehicle package. However, late delivery and difficulties with the vendor delayed final flight testing for over two months.

It was important that the SAS not be operated autonomously while the vehicle was in contact with the ground. Any attitude corrections attempted while the vehicle was constrained by the ground could easily tip the helicopter over. This was addressed by programming the SAS to operate in three modes: first, fully autonomous control; second, fully-remote control; and, thirdly operation in a partially autonomous mode. The latter state controlled the vehicle altitude without attempting attitude corrections. This mode was used for taking off. Once the helicopter cleared the ground, the SAS was switched to its normal operating state via a command from the outer loop controller. This process worked for taking off in calm air, but was not tested under windy conditions.

A final problem concerning the SAS was the switch from manual control to computer control and the switch back. The SAS was unaware of the transition and regarded any abrupt change in its target values as legal requests. Since the man-




Fig. 8. Inner loop flight controller schematic

ual controls could never be placed such that they matched the computer driven inputs, violent vehicle movements could occur after the switch. This problem somewhat limited the usefulness of the manual control system for testing the outer loop controller.

6.3.3. Outer Loop. The outer loop controller monitored the vehicle's position and velocity. It compared these values with target parameters and made corrections as needed to achieve the desired flight path. In addition, it received input commands in the form of target X,Y positions, as determined by the mission planning subsystem (described in section 6.4). Outputs from this controller were vehicle attitude commands. For example, to move the vehicle forward, the controller pitched the nose of the vehicle down, increasing forward velocity. When the vehicle achieved desired cruising speed, the pitch was reduced. The aircraft similarly decelerated as it neared the target location.

The outer loop controller was a closed loop linear controller with two independent channels operating in the X and Y planes. There were two copies of the control schematic shown in Figure 9, one controlling each channel. During takeoff and landing the altitude was increased or decreased linearly with no feedback. During the mission, altitude was held constant and therefore the Z axis did not have a controller in the outer loop. Control channels used actual vehicle position and velocity as feedback inputs. Gain terms were dependent on vehicle weight and engine performance, and were tuned experimentally.

Vehicle position was provided by the mission planner via the vehicle positioning subsystem. It is sampled over time to derive velocity. This method of determining velocity was chosen because commercial sensors were unavailable which could function at the vehicle's range of

mission airspeeds. The use of the visually determined position to infer velocity could cause problems, since position data was only expected to be accurate to ± 6 inches. A large amount of damping in the velocity term, to counterbalance velocity spikes in the samples, degraded vehicle responsiveness. This problem remains a concern to be evaluated through further testing.

The outer control loop relied on the radio links to control the vehicle. While reliance on radio links made interference a major weakness in the design, use of a ground station was chosen because it offered greater long term flexibility in system application. Methods for moving this computation on-board while maintaining some form of ground control station are being explored.

6.4. Mission Planning

The mission planning portion of the aerial vehicle system resided in the ground-based computer. This subsystem was predominantly software and represented the system's intelligent planning. Due to the critical nature of mission planning, most of the other subsystems interacted in some manner with its operations. A functional description, followed by a more technical description, of the mission planning activities follows.

6.4.1. Functional Description. The mission planning subsystem was the core of intelligence and decision-making for the vehicle. Without this portion of the system, a human could fly the vehicle, but with little support from the other on-board and ground based subsystems. As portions of the mission planning subsystem were added, the human pilot was provided with greater assistance in flying and performing the mission. Taken with all of the mission planning software in place, the vehicle was totally autonomous.




Fig. 9. Outer loop flight controller schematic

There were several important tasks which the mission planning subsystem performed. These composed the major pilot augmentation features, and were:

- World model maintenance;
- Visual cockpit display system;
- System monitoring; and,
- Motion planning.

The mission planning subsystem had the ability to visually replay previous missions and allowed simulation of missions. The capabilities of the mission planning system proved very valuable during debugging of the other subsystems.

Mission planning operations began when data was received from the aerial vehicle via the down-link. The data was evaluated and used to plan the next action for the vehicle. Actions consisted of both vehicle position commands and disk retrieval operations. Vehicle position commands were passed to the outer loop flight controller and disk retrieval commands were transmitted to the vehicle. Mission planning computations had to be performed in a timely manner to maintain the 10 Hz cycle time along the communication path identified for the mission.

6.4.2. Technical Description. The technical aspects of the mission planning activities revolved around the four tasks described above. Each of these will now be explained in detail.

6.4.2.1. WORLD MODEL MAINTENANCE. The world model consisted of a spatial and temporal representation of the vehicle, updated from the down-linked vehicle telemetry. Each data set was time stamped for temporal reference. The first activity of the world model was to interpret the data in the down-link. If a set of data was missing, or the error detection mechanisms signaled a problem with the arriving data, the entire data packet was not used. After verifying the completeness of the data, the data was added to the world model.

Prior to flight initiation, a priori knowledge of the environment was entered into the model. This information included parameters about the space allocated for the mission, the locations of source and depository bins, and locations of navigation aids. This information was considered

spatially exact and constant over time. As the vehicle began the mission, the position of the vehicle and objects which it “saw” were stored in the world model. Over time, refinement of unknown objects (those which were “discovered” as opposed to known a priori) took place. The vehicle could discern bad data by using the accumulated average of information in the world model. This gave rise to sensor validity issues and sensor fusion issues.

The world model did not directly augment a human pilot’s ability to fly the vehicle because there did not exist a mechanism for the pilot to see or visualize what was happening. However, all other pieces of mission planning depended on the information contained in the world model. The contents of the world model could be “replayed” to see what happened with various subsystems during the mission. The next portion of the mission planning subsystem added a visual interface to the world model, helping a pilot to fly the vehicle. It was not designed to be useful in such a small area as the volleyball court for this competition, but for missions encompassing greater distances.

6.4.2.2. VISUAL COCKPIT DISPLAY SYSTEM. Although not directly needed by the autonomous vehicle system, a visual display of the world model provided a very realistic and useful interface for debugging. This visual display could also be used to augment a human flying the vehicle. The design utilized for the project involved splitting the serial down-link communication so that the signals were sent to the ground-based computer for mission planning activities as well as to a Silicon Graphics Workstation. Thus, two world models actually existed, one on the ground-based computer, and the other on the Silicon Graphics Workstation. A vehicle cockpit was created to display, in real time, a graphical representation of the air vehicle within the competition space along with key pieces of mission data. There existed an altimeter gauge, a compass heading, pitch and roll indicators, and an indicator for the on-board battery status (see Figure 10). Also, when disks were located in the source bin, a separate image was displayed showing the placement of the disks within the bin (as known by the world model).




Fig. 10. Screen images of the visual cockpit system

This interface proved most valuable for debugging system components and for helping with the integration aspects of the project. Very easily, for example, one could visually see that one component was producing data in world coordinates with one type of unit measure and that another component was using either a different coordinate reference system or different units. The visual display also proved useful for replaying the mission for evaluation. For missions span-

ning larger spaces, this interface would provide better spatial resolution for a human pilot. It could also allow the human pilot to fly an "instrument" mission, that is, flying the vehicle where or when it can not be visually seen.

6.4.2.3. SYSTEM MONITORING. This part of mission planning monitored the status of various components in the system as well as the execution of

tasks initiated by the motion planning subsystem.

Before the vehicle was allowed to lift off the ground, the system verified that all components were operational. Items checked included the communication links, the vehicle systems, the ground systems, and the items set up by the team for navigation. This procedure was much like a human pilot going through a check list prior to takeoff. These same systems were monitored throughout the flight. If something went wrong, the mission was aborted. If the up-link portion of the communication path was lost, the vehicle was capable of maintaining a hover with only the on-board systems.

Sensing the successful completion of a mission was also a part of the system monitoring subsystem. When mission planning requested that a disk be retrieved, it monitored whether the retriever was able to attach to a disk. If a disk was not extracted, there was no need to fly to the depository bin. Rather, further investigation took place within the source bin.

This level of intelligence assisted the pilot when problems arose by showing what systems were working and which were not working. Also the pilot could determine if a disk retrieval had been successful, which was difficult to discern from a remote distance.

6.4.2.4. MOTION PLANNING. The final level of intelligence within the mission planning subsystem provided autonomous flight. This portion initiated the commands to takeoff, land, and move from one location in the competition arena to another location.

The mission planning system determined a flight path and identified this path by its end points, i.e., move from the start location, to the center of the source bin, to the center of the deposit bin, to the source bin, and continue until the disks were acquired at which point return to the start location for landing. The group decided that the vehicle motion would be as restricted as possible to allow for more simplistic motion planning. The flight plan dictated that the vehicle fly at constant altitude during the mission and with constant heading. To reverse direction, the vehicle would fly backwards. Therefore, upon deciding to commence flight, the first command was

to achieve the mission flight altitude. The motion planning facility commanded that the vehicle should reach this altitude with corresponding X and Y locations. The desired position was passed to the outer loop flight controller which handled how to achieve the position. Therefore, a hierarchical approach to the intelligence was developed. The mission planning subsystem was concerned with the higher level issues of what to do and when to do them. This information was then passed to other subsystems for execution.

6.4.3. Design Decisions. Several factors came to play in the development of the mission planning subsystem. One of the important issues was reliance on the communication subsystem. Without receipt of data to create the world model, the mission planner could do nothing. Along with the data was a need to store the data in the model and a system which could perform the necessary computations in a timely fashion. This consideration led to the use of two computer systems on the ground: a 386SX personal computer for planning and evaluation, and the Silicon Graphics Workstation for real-time visual displays.

The importance of distributed intelligence was key in the design [9]. During the design phase, the exact nature of the components was unknown. Therefore, from the mission planners point of view, a simplicity of design was needed whereby the mission planner would request a certain function or objective to be achieved. The corresponding subsystems could then act on these high level commands and perform the task. The mission planner did not need to be concerned with how the task was accomplished so long as the goal was realized.

6.4.4. Assessment. While the motion planning may seem somewhat simplistic, flying between two points, the development of the mission planning system was by no means trivial. The real-time issues and integration of the components was critical to keep the vehicle airborne. The world model and visualization portions provided tremendous capabilities for debugging. However, the enormous issues of integrating the system kept the mission planner from being fully operational at competition time. As will be discussed later with integration issues, the size of the team

and dynamic nature of the project sometimes resulted in various portions of code being duplicated in the system. This resulted in two versions of software in existence. The interfaces between the mission planner and the other subsystems could only be defined once the functionality of the other pieces was known. The concept of distributed intelligence is good, but the interfaces between the layers must be accurately identified and defined.

Due to difficulties in integration of the components, the system monitoring portion of mission planning was not present at the time of the competition. The mission planning subsystem assumed that requested actions were acted upon appropriately. Other ideas for improving the mission planning subsystem involve validating the data read from the down-link. Rather than throwing away entire data packets if one part was missing, perhaps some of the data could be recovered.

6.5. Vehicle Positioning Subsystem

6.5.1. Functional Description. The vehicle positioning subsystem was responsible for determining the X,Y location of the helicopter within the competition arena. This position was used by the outer loop controller for navigation. Sensors capable of reasonable accuracy at the low ground speeds traveled by the vehicle proved unavailable. Therefore, it was decided that a vision based triangulation system, using cues placed around the arena, would provide the necessary accuracy with the least monetary investment.

The visual triangulation system calculated the vehicle's position using lights mounted in known locations around the arena as landmarks. The apparent location of each light was determined using a vision system. The location of the helicopter was then determined by calculating the intersection point of lines drawn from each landmark to the image center. This was the vision-based equivalent of how aircraft pilots locate their position using radio transmitters as landmarks [10].

The deployed system did image processing on-board the helicopter to locate the lights in the image. These locations were then transmitted to the ground computer where the triangulation algo-

rithm determined the vehicle's position. This process executed 10 times per second with a predicted positional accuracy of better than ± 6 inches.

6.5.2. Technical Description. The Integrated Vision System (IVS) and associated optics were carried on-board the helicopter. The IVS consisted of a camera lens and Charge-Coupled Detector (CCD) attached directly to a single board Motorola 68000 based microcomputer. This arrangement allowed rapid image acquisition and digitization, typically on the order of 100 microseconds per frame. This integration reduced weight, volume and power requirements. Image processing was performed on-board, requiring only information describing the bright spots in the image to be transmitted to the ground. This greatly reduced the bandwidth requirements placed on the data link.

Externally mounted lights, used for navigation, were placed along the longitudinal and transverse axes of the source and deposit bins in

```

INPUT:   image[maxrow][maxcolumn]
OUTPUT:  list of blobs

BYTE queue[num_samples]
WORD queue_sum
BYTE queue_mean

BEGIN
  /* preload the queue from the diagonal*/
  FOR i = 1 TO num_samples DO
    queue[i] = image[i][i]
    queue_sum = queue_sum + queue[i]
  ENDDO

  /* now scan the image for blobs */
  FOR row = 1 TO num_rows DO
    FOR column = 1 TO num_columns DO
      IF (image[row][column] > (factor * queue_mean)) THEN
        /* pixel is within a blob */
        grow_blob(row, column)
      ENDIF
      /* note: pixels are set to 0 when grow_blob uses them */
      IF (image[row][column] > 0) THEN
        /* remove the oldest pixel from the queue */
        old_pixel = dequeue(queue)
        /* put the new one in */
        enqueue(queue, image[row][column])
        /* subtract the old value from sum */
        queue_sum = queue_sum - old_pixel
        /* add the new one in */
        queue_sum = queue_sum + image[row][column]
        /* compute the new mean */
        queue_mean = queue_sum/num_samples
      ENDIF
    ENDDO
  ENDDO
END

```

Fig. 11. Adaptive thresholding algorithm

the competition arena shown in Figure 1. To ensure that these six lights would be brighter than the ambient light, 500 watt incandescent light bulbs were used. The camera then used very short exposure times, a very small F-Stop and two polarizing filters turned out of phase to attenuate the image sufficiently to remove streaking and blooming. The lights were placed such that when the vehicle was hovering over a bin there would be three lights visible at right angles. This provides the most accurate positioning information when it was most needed.

To cope with the varying brightness of near and far lights as well as background intensity changes, an adaptive thresholding method was used, whose pseudocode is presented in Figure 11. The adaptive threshold was calculated based on the mean value of the intensities of the last *num_samples* pixels. The intensities of the last *num_samples* pixels were remembered in a queue. The sum and the mean value of the queued intensities were also saved. When each new pixel was sampled, it displaced a value from the queue. This displaced value was subtracted from *queue_sum* and the new value was added on and enqueued. The new *queue_mean* value was then computed. The *grow_blob* subroutine, not shown, started with an initial pixel and grouped surrounding pixels that exceeded a threshold intensity. It then calculated and re-

turned the center of intensity for the blob of light as well as region statistics that were used for exposure adjustment.

The camera lens was focused onto a mirrored right circular cone which provided a 360° horizontal field of view. The conic shown in Figures 12 and 13 were manufactured out of brass stock, polished until smooth and silver plated to create a mirror. The cone was mounted using a threaded hole in the base. A smaller hole drilled completely through the center of the cone allowed an LED to be mounted such that it was visible through the top of the cone and in the image.

During initialization of the vision system, the LED in the conic was energized. The location of this light in the image marked the center of the cone and was used as the optical center by the triangulation algorithm, Figure 14, when it converted image locations into polar coordinates. This method facilitated finding the center of the cone automatically and gave some resiliency to rough handling. The X,Y location of the vehicle in the starting position was input into the system to complete the bootstrapping process.

The vehicle positioning system operated in a feed-forward mode during flight. The last position of the vehicle was used to predict where the vehicle currently was located. This prediction was used by the triangulation algorithm when it matched blobs to lights. With the system cycling




Fig. 12. Vehicle positioning camera conic device




Fig. 13. Vehicle positioning camera conic assembly.

10 times per second, the vehicle only moved a few inches between images at cruising speed.

6.5.3. *Design Decisions.* Several other methods for determining vehicle position were explored. A vision-based system that used cameras on the ground to track the helicopter was proposed. However, it was decided that the 3 minutes allowed to set up equipment would not allow sufficient time to position and calibrate the cameras. A commercial radio locating system was ex-

plored but found to cost much more than this project could afford.

For the 1992 competition, the team is evaluating designing a custom radio locating system that uses commercial hobby transmitters and receivers to reduce development costs to acceptable levels. A differential Global Positioning System (GPS) is also being evaluated. This system would utilize a ground-based receiver and telemetry unit to reduce spatial error sufficient for use during the competition. A vision-based system using laser beams for landmarks instead of light bulbs also shows promise.

6.5.4. *Assessment.* The camera was originally designed to be mounted under the keel plate of the helicopter. While constructing the solid model of the vehicle and placing components, it was determined that a less obstructed field of view would be achieved by mounting the camera on top of the helicopter, just forward of the rotor mast. This required turning the conic assembly upside down since the CCD had to be within 6 inches of the IVS board. This reorientation resulted in significant concerns about the positional accuracy. With the camera right side up and the vehicle on the ground, the lights were reflected near the apex of the cone. The accuracy in this position was limited by manufacturing precision and camera resolution since the lights were very close together in the image. As the vehicle achieved flight altitude, the lights moved toward the base of the cone, which afforded the greatest accuracy while the vehicle was performing the mission. In the inverted state as deployed, the lights were at the base of the cone when the vehicle was on the ground and moved towards the point as it climbed. This problem was a major concern. A possible solution was to use two sets of lights at different heights. The vehicle would use the low lights when on the ground and switch to high lights during operation to alleviate some of the accuracy problems. This option was never fully developed due to insufficient design and implementation time.

By matching the angle of the conic to the camera lens, the vertical field of view was expected to be $\pm 7^\circ$ from the horizontal axis with 15° pin-hole optics. Unfortunately, the assembled cam-

```

INPUT: blob locations in image coordinates
OUTPUT: vehicle X,Y position

BEGIN
  FOR each blob DO
    Create a unit vector from the center of the image to the blob
    Rotate the vector to compensate for vehicle yaw
  ENDDO

  FOR each light being tracked DO
    Match light to blob closest to the location in the last image
  ENDDO

  FOR each remaining light DO
    Match light to blob closest to the predicted location
  ENDDO

  FOR each pair of lights DO
    /* Don't intersect lines that are nearly parallel */
    IF (30 ≤ angle of intersection ≤ 150) THEN
      /* Lines are along unit vectors */
      /* describing the blob location */
      calculate_intersection(line1,line2)
    ENDIF
  ENDDO

  Calculate mean of intersection points
  Discard points farther than 12 inches from the point nearest the mean value
  Recompute mean with remaining points

  Return mean as vehicle X,Y position
END

```

Fig. 14. Triangulation algorithm

era optics did not meet expectations and the delivered vertical field of view was closer to $+5^\circ$ to -10° . This lack of vertical field of view and resulting sensitivity to roll and pitch remained an unresolved issue.

The positioning system in static testing demonstrated accuracies better than the design goal of ± 6 inches. The system was not tested on the vehicle due to delays in system integration.

6.6. Disk Locating Subsystem

6.6.1. Functional Description. The disk locating subsystem was responsible for determining the location of disks in the source bin. An IVS identical to that used by the vehicle positioning subsystem was deployed to look down into the bin and locate the disks. A camera flash illuminated the scene and reduced the effects of ambient light changes. To compensate for a long strobe recharge time, the system used knowledge of the position of the vehicle to only take pictures when the disks should have been visible.

The on-board system determined the centers of the disks that were discernible in the image and transmitted that information to the ground. The ground system then used the vehicle's altitude, heading and position to translate the image coordinates to world coordinates. These disk locations were then combined with the beliefs of where the disks were located in the bin to improve the system's knowledge and accuracy. The disk location with the highest certainty associated with it was then selected as the next target for retrieval. This method reduced the effects of transient reflections that were erroneously marked as disks.

6.6.2. Technical Description. The camera optics gave a $\pm 30^\circ$ field of view. This resulted in an image of a 5 foot by 5 foot area of the ground when taken from the 5 foot cruising altitude. This was not large enough to include the complete 6 foot diameter bin, but a lens with a wider field of view was too thick to use with the camera.

An orange filter was used to highlight the orange painted disks in the image. A polarizing filter was used to attenuate glare and reflections from both the sun and camera strobe. The strobe

exposure used very short integration times, resulting in steady images taken from the vibrating helicopter.

There was a matching problem involved with the ground-based system. It received disk locations from the helicopter in an egocentric frame of reference. A ground-based landmark was not present in the images that could be used to reduce the ambiguity, so the vehicle's altitude, yaw and position were used to convert the disk locations to a world coordinate frame of reference. Therefore, the disk locations could be determined no better than the accuracy of the vehicle positioning system. The system kept track of disk locations and attempted to match new reports with disks previously reported. The number of times that a disk location was matched was termed the *Hit* count. Each match was used to reduce the spatial uncertainty of the disk locations by averaging the new reading with past readings. Six slots were allocated to hold the disk locations. If the new disk location could not be matched to an existing location, the new disk location was inserted into a slot. If all the slots were in use, the slot with the lowest *Hit* count was replaced. When a target was required for retrieval, the disk with the highest *Hit* count was selected. The algorithm used is shown in Figure 15, in pseudocode.

6.6.3. Design Decisions. It was necessary to use a strobe with the vision system to allow reduced exposure times and thereby reduce image blur-

```

FOR each disk location received DO
  Correct disk location using vehicle yaw
  Convert pixel coordinates to world coordinates
  FOR each disk we are tracking DO
    IF this disk is within 6" of tracked disk THEN
      Average this location with tracked location
      Increase tracked disks HIT count
    ENDIF
  ENDDO
  IF did not match a received disk location THEN
    IF disks being tracked < 6 THEN
      Add disk location to list
    ELSE
      Replace disk with least number of HITs
    ENDIF
  ENDDO

```

Fig. 15. Disk matching pseudo-code

ring. The choice of strobe was driven by the need for interfacing it to the IVS. Conventional 35mm camera flash units were decided on because of their light weight, self-contained construction and ease of interface to the camera. A unit with a variable power setting was chosen to allow optimization of the power to recharge time tradeoff.

6.6.4. Assessment. When the system was tested in direct sunlight over the competition bins, specular reflections from the sun were stronger than the flash illuminated disks. This necessitated increasing the power setting on the flash, as well as adding the polarizer and reducing the camera F-Stop to attenuate the light back to usable levels. This overcame the reflection problem but added the encumbrance of a 7 second strobe recharge time. This problem was surmounted by utilizing knowledge of the vehicle's position. The camera only took one picture per trip, each time immediately after the vehicle arrived over the source bin. This allowed sufficient time for the strobe to recharge between images, but made no allowance for unusable images.

6.7. Disk Retrieval System

Mission requirements specified that six disks would be located in a bin on one side of the competition arena. The vehicle needed a way to capture the disks one at a time and move them to the depository bin. Several mechanisms were investigated by the team and the paragraphs below describe these activities.

6.7.1. Functional Description. The essential elements of the retrieval device were predicated on several mission determined objectives.

First, the device had to retrieve only one disk at a time. Therefore, if two disks were in close proximity, the device needed to determine which to acquire. The skills involved in acquiring a disk could be any one, or combination of: magnetic, suction, grasping via holes in the disk, or grasping via a fork lift type maneuver between the top and bottom portions of the disk (see Figure 16). The ability to acquire the disk involved a certain level of intelligence in order to detect the disk's presence, to possibly touch or come in contact




Fig. 16. Target disk geometry

with the disk, and to acquire a selected disk, possibly choosing from more than one.

The chosen method, secondly, needed to be able to perform the first objective within the geometries presented by the vehicle's positioning subsystem and stability characteristics. Studies with the vehicle positioning system and flight controls, determined a spatial tolerance of plus or minus six inches on the surface of the court, combined with a vertical tolerance of plus or minus three inches were to be expected. Therefore, the device needed to be able to sense and acquire disks, subject to some movement and unknown spatial positioning of the helicopter.

Third, when positioned over a disk, the retrieval device needed to be able to pick the disk up within time tolerances dictated by flight trajectory limitations. The intentions of the mission were to acquire all six disks within the three minute flight time. This translated into a time requirement for each retrieval operation of approximately ten seconds or less.

Therefore, the retriever needed to possess intelligence sufficient to differentiate between materials in contact with the retriever's surface and to select the disk, given a choice, resulting in the highest probability of retriever success.

6.7.2. Technical Description. The retriever developed for the AUVS mission was an electromagnetic array that would be dropped onto the competition surface. Once in the lowered position, appropriate magnets on the array were energized to acquire a disk. The array and attached

disk were retracted back up to the vehicle and flown over the deposit bin, where the electromagnets would be turned off to drop the disk. A Motorola MC68HC811 microcontroller operated the retrieval subsystem. Three main components comprised the retriever device: a retraction system, a magnet array, and a connection mechanism for the assembly to the aerial vehicle (see Figure 17 and 18).

6.7.2.1. RETRACTION SYSTEM. The retraction system consisted of a cable and reel type mechanism. When the vehicle was positioned over a disk, a signal initiated by the mission planning subsystem was sent via the up-link to deploy the retriever. An electromagnetic brake on the reel was released and a gearmotor lowered the magnetic array to the ground. Due to friction in the system, free fall caused by the weight of the assembly was not sufficiently fast enough for reaching the ground in the required time. Deployment speeds of three feet per second could be obtained using the motor. When the array sensed and grasped a disk, the retriever was activated to retract the mechanism. The array's descent was monitored by an optical shaft encoder. This information provided the retriever with knowledge of how far the magnetic array was extended below the vehicle, and subsequently allowed the re-

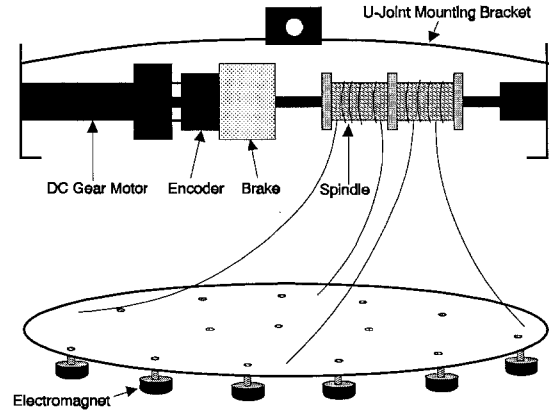


Fig. 18. Disk retrieval system schematic

triever to know how far the array must be retracted. Retraction speeds of two feet per second were obtained. When the array reached the fully retracted position, the motor was de-energized and the brake activated to hold the assembly in the retracted position. A cable system between the magnetic array and vehicle allowed flexibility and tolerance in the vertical and horizontal directions.

6.7.2.2. MAGNET ARRAY. Located on the ground side of the retraction system was an array of electromagnets. Due to the spatial positioning




Fig. 17. Disk retrieval system

tolerance of ± 6 inches, multiple electromagnets were attached to a one foot diameter piece of plexiglass. Positioning twelve electromagnets within the one foot diameter circle allowed any disk location to be in contact with at least one magnet of the array. Testing with disks in various locations under the circle provided electrical inductance values which could be detected and measured across each magnet, even when magnets were only partially on a disk.

A Motorola MC68HC811 microcontroller controlled the selection process. When the magnet array reached the ground, a signal from the retrieval subsystem initiated a search for disks. Each of the magnets was individually energized to measure the presence or absence of a disk, calculated from the inductance value measured across each magnet. Any measurement above the predetermined zero value (no disk) assigned a probability of disk presence. After all magnets had been tested, the best region (high probability) for disk retrieval was activated. This activation region may have been one magnet, if a disk was directly under a magnet, or more than one when a disk was located between discrete magnet positions. Having successfully activated magnet(s) to acquire a disk, the serial communication link between the microprocessor on the array and the microprocessor in the retrieval subsystem caused the array to be retracted. If a disk was not found with this method, the vehicle could be repositioned by mission planning and a new search initiated.

6.7.2.3. CONNECTION MECHANISM. The entire retrieval mechanism was attached to the aerial vehicle using a U-joint assembly. The “dropping” and retrieving of the magnet array was determined to be critical to the flight stability of the aerial vehicle. The amount of movement (both translational and vertical) had to be isolated as much as possible from the vehicle. Therefore, a U-joint assembly connecting the retrieval mechanism to the vehicle was developed. This resulted in a single point load on the vehicle. The location of this connection was chosen to be the vehicle’s center of mass in order to reduce any moments imparted to the vehicle and initiated by the retrieval mechanism. With the U-joint allowing pivot in two directions, the array when de-

ployed, could remain stationary while the vehicle hovered within some spatial tolerance.

6.7.3. Design Decisions. The retrieval portion of the mission was separated from other design activities in order to study a variety of disk acquisition methods. Tactile arrangements were investigated for “scooping” or for grasping the disks, in addition to the method described above. A study using permanent magnets, as opposed to electromagnets, was also investigated. Various methods for attaching the retriever to the vehicle were studied; including a telescoping arm. At one point, a small autonomous ground vehicle was considered for deployment into the source bin to acquire disks. This vehicle would be tethered to the aerial vehicle and, after capturing a disk, could be retracted much like the magnetic array.

In all scenarios, issues of weight, power requirement, and total time for acquisition were compared. Design complexities were identified with the method of choice seeming to be use of the electromagnet array. The team continues to identify and refine these, and other, methods for acquiring the disks.

6.7.4. Assessment. Reviewing the selected retriever mechanism still raises several unsolved issues. The current mechanism (or any other mechanism) has been designed for the one foot diameter spatial tolerance. If the value could be reduced, other retrieval methods might be incorporated. Obviously, if the one foot tolerance is not sufficient, the potential exists for deploying the retriever and not finding a disk. Secondly, the array’s deployment location is dependent on the ability to accurately locate disks from the aerial vehicle. These two issues significantly impact retriever design.

Evaluating the existing retriever also identifies several problem areas. At irregular intervals, the flexible cabling systems becomes tangled, resulting in the array’s inability to deploy and/or retract. This problem is very similar to knotting of fishing line around a spinner. If too much line is unreeled prior to braking the spinner when the array contacts the ground, too much slack can exist. The “swinging” of the array while being retracted, in addition, must be fully understood in terms of its relation to vehicle dynamics. Time

critical issues of making stability corrections to the vehicle while retracting the retriever require additional study.

6.8. Communications

Early in the design process, the team decided to place some of the computing resources on the ground. This decision provided less weight on the vehicle in terms of batteries to drive computer components and provided for a ground station to monitor vehicle activities. However, this also dictated that a communication link be established which could reliably move data between the vehicle and the ground at a suitable speed. Therefore, the communications system became a critical component of the vehicle system design.

6.8.1. Overview. The communication subsystem involved several subparts: an on-board “post office”, a down-link, a ground station, and an up-link. These four components provided a one-way communication path from the vehicle to the ground and back to the vehicle. Communications connected the aircraft with the ground, connecting the sensors and actuators to the planners and evaluators.

The system could best be visualized as a one-way path (see Figure 5). The “post office” collected information on the vehicle and packaged that data into one cohesive bundle for transmission to the ground via the down-link. The down-link provided a one-way corridor for information transfer to the ground. Once on the ground, the information was interpreted and evaluated before planning the next set of vehicle actions. These actions were encoded and sent via the up-link to provide actuator control on the vehicle.

Competition rules specified that a “kill” switch had to be present so that a safety pilot could override the autonomous actions of the vehicle. The issue presented then, was one of how to provide a communication system with the reliability and capacity to handle the information exchange [11]. The following paragraphs describe three of the major portions of the communication system, the “post office”, the down-link, and the up-link. Information about the ground station has been provided in the Mission Planning portion of the paper (section 6.4).

6.8.2. Post Office.

6.8.2.1. FUNCTIONAL DESCRIPTION. The post office collected various kinds of information on-board the aerial vehicle for transmission to the ground. During the design, the team realized that with different components on the vehicle, each would want to send information to the ground. The team also determined that having multiple transmitters, or multiple channels on a single transmitter would be expensive in terms of weight and power consumption. Therefore, a system which collected information from the components on the vehicle and packaged them into one piece of data for transmission would be necessary. This system of collection and distribution was dubbed the “post office”.

The current vehicle configuration supported four suppliers of information to the ground. These were the aerial vehicle sensor suite, the vehicle positioning system, the disk locating system, and the disk retrieval system. Each of these systems produced telemetry data which the post office collected and passed to the down-link.

6.8.2.2. TECHNICAL DESCRIPTION. Each of the four on-board systems which sent data had their own microprocessor. These systems transmitted data to the post office via serial communication. The post office which received the data was designed around a Motorola MC68HC811 controlling an octal Universal Asynchronous Receiver/Transmitter (UART) as shown in Figure 19. Software was developed to read each of the input component data ports and store the data into memory in a specified format. Each of the input ports on

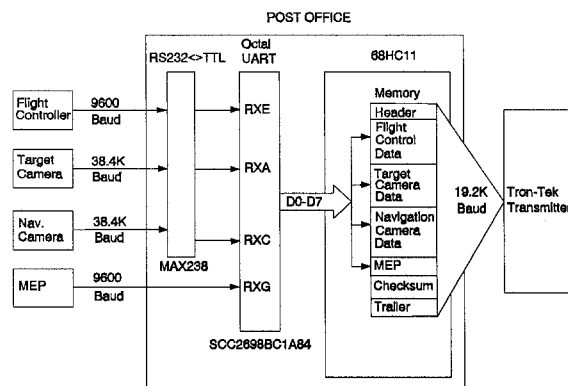


Fig. 19. Communications post office schematic layout

the UART had sufficient memory to store its own data until an interrupt directed the HC11 to read the data. After reading each of the four incoming data sources, the post office memory was transmitted through an output port connected to the down-link. The post office software was developed to read and send complete data packets down the down-link at the system desired rate of 10 Hz.

6.8.2.3. DESIGN DECISIONS. Beyond the initial design decision requiring a device to collect the data on the vehicle, several other design parameters influenced the post office. The current design assumed that all input data from the four on-board vehicle components were correct. No data checking was performed at this level. Secondly, the post office assumed that the data read was most recent. There were no temporal ties between the four input channels other than the sequence of reading the ports. The underlying assumption with this method was one which said that any related change between the four input channels would occur at time intervals greater than one-tenth of a second. The method for reading the data also assumed that the four input devices were equally important and no priority was given to any particular input channel.

Following from these design assumptions, the device which resulted had to be small and of light weight. It had to be able to gather information from four different sources, package the information (approximately 100 bytes when assembled) and send it out another port within the 10 Hz design specification. The power to drive this component had to be minimal to reduce battery weight and the system should be programmable to accommodate changes in both the number of input ports and the amount of data coming from any one input port.

6.8.2.4. ASSESSMENT. The configured device successfully meets all of the design parameters. If funding would have allowed, a more robust product could have been purchased rather than student built. This would have made the process of code development on a personal computer and then transfer of code to the post office microprocessor easier.

6.8.3. Down-link.

6.8.3.1. FUNCTIONAL DESCRIPTION. The down-link received the data package from the post office and transmitted the information to the ground. Associated with each packet of data (the grouping of information obtained by one collection sweep of the post office) was error checking information. Header bytes were added so as to help decode the transmission. Following the header bytes were four groups of data, the aerial vehicle sensor data set, the vehicle position data set, the disk location data set, and the retrieval mechanism data set. By defining the total length of the data package as a constant, the ground station could decode the transmission data package appropriately. The header bytes help realign the software in case of transmission errors. Figure 20 shows the data sent by each of the four on-board component systems.

System Component	Data	Representation	
Aerial Vehicle Sensor Suite	Pitch Attitude	2 signed bytes	
	Roll Attitude	2 signed bytes	
	Yaw Attitude	2 signed bytes	
	Pitch Rate	2 signed bytes	
	Roll Rate	2 signed bytes	
	Yaw Rate	2 signed bytes	
	Altitude	2 unsigned bytes	
	Unused Data	14 signed bytes	
	Vehicle Position	Major Fault Status	1 unsigned byte
		Minor Fault Status	1 unsigned byte
Light 1 - X Pixel		2 unsigned bytes	
Light 1 - Y Pixel		2 unsigned bytes	
Light 2 - X Pixel		2 unsigned bytes	
Light 2 - Y Pixel		2 unsigned bytes	
Light 3 - X Pixel		2 unsigned bytes	
Light 3 - Y Pixel		2 unsigned bytes	
Light 4 - X Pixel		2 unsigned bytes	
Light 4 - Y Pixel		2 unsigned bytes	
Light 5 - X Pixel		2 unsigned bytes	
Light 5 - Y Pixel		2 unsigned bytes	
Light 6 - X Pixel		2 unsigned bytes	
Light 6 - Y Pixel		2 unsigned bytes	
Disk Location	Minor Fault Status	1 unsigned byte	
	Major Fault Status	1 unsigned byte	
	Disk 1 - X Pixel	2 unsigned bytes	
	Disk 1 - Y Pixel	2 unsigned bytes	
	Disk 2 - X Pixel	2 unsigned bytes	
	Disk 2 - Y Pixel	2 unsigned bytes	
	Disk 3 - X Pixel	2 unsigned bytes	
	Disk 3 - Y Pixel	2 unsigned bytes	
	Disk 4 - X Pixel	2 unsigned bytes	
	Disk 4 - Y Pixel	2 unsigned bytes	
Retrieval Mechanism	Disk 5 - X Pixel	2 unsigned bytes	
	Disk 5 - Y Pixel	2 unsigned bytes	
	Disk 6 - X Pixel	2 unsigned bytes	
	Disk 6 - Y Pixel	2 unsigned bytes	
	Carriage Status	1 unsigned byte	
	Array Status	1 unsigned byte	

Fig. 20. Down-link data packet

6.8.3.2. TECHNICAL DESCRIPTION. The down-link transmission device was a 1 watt video and audio transmitter operating at a frequency of 460 MHz. The digital data was sent as an analog approximation via the video channel. A frequency of 460 MHz provided sufficient bandwidth for reliable approximation of the digital signal. The output voltage of the post office signal was reduced to a value of 1.4 volts peak-to-peak, normally accepted by the transmitter, as compared to the nearly 24 volts peak-to-peak of standard RS232 communications, using a standard diode clamping circuit [12].

The signal was detected by the ground station receiver using a zero crossing circuit detector. As the signal crossed the zero reference voltage, the output of a comparator changed and the signal was boosted to RS232 levels for transmission to the ground computer via a serial data communications port. The ground computer then unpackaged the data for processing as explained within mission planning.

As part of the assembly and unpackaging of information sent via the down-link, software was required to represent all information as byte components. Some of the data sources send 1 byte signed data or 1 byte unsigned data versus other sources sending 2 byte signed and unsigned data. Therefore, software was necessary to split 2 byte data into representative 1 byte blocks. At the ground station, the data was reassembled knowing what format was expected. Those operations were necessary to allow a general interrupt function to write byte data from the serial port to memory at the ground station computer without losing any information (i.e., inappropriate timing of an interrupt fetching data at the port and writing to memory). This problem was associated with the one way communication path, therefore no flow control, such as XON/XOFF, could be utilized.

6.8.3.3. DESIGN DECISIONS. During the design phase of the project, the team studied the communication throughput required to pass data packets containing approximately 100 bytes of information at the specified timing of 10 Hz. From these studies, the team concluded that the amount of data had to be transmitted faster than

9600 baud to allow the ground-based computer to do other processing besides reading the serial port. However, most "wireless" modems operate at transmission speeds of 9600 baud or less. Therefore, the solution presented above was developed. This communication system operated effectively at 19200 baud and could be extended to 38400 baud. The speed of the down-link gave the design added flexibility in the size of the data packet sent to the ground.

Other design considerations concerning the down-link, involved the battery power required to drive the transmitter resulting in miniature rechargeable nickel-cadmium (NiCad) batteries being used. The output from the batteries was passed through a voltage regulator. The overall communication system weight was thus reduced using these batteries coupled with the small size of the transmitter. The frequency of transmission was another design decision. The frequency had to be sufficiently high to allow reliable recognition of the digital data over the analog signal.

6.8.3.4. ASSESSMENT. The selection of the 460 MHz frequency was one of the poorer choices in the system. This frequency is heavily utilized for many other communication services (e.g., taxi cab dispatch and other 2-way radio communications). The system worked well and provided the adequate throughput, but was prone to being unreliable due to other signals stronger than the transmitter on the aerial vehicle. This problem is the suspected reason for interference associated with the communications system around the arena during the competition. The same components can be utilized with a better frequency selection and possibly a stronger signal. However, each of the options requires additional power resources for the transmitter which adds weight to the vehicle.

6.8.4. *Up-link.*

6.8.4.1. FUNCTIONAL DESCRIPTION. Once the ground station had interpreted the information from the down-link, a planning and evaluation phase was undertaken to determine what action the vehicle should perform next. The actions involved commands such as those to increase the engine speed, lower the retrieval mechanism, etc. These

commands maneuvered the vehicle and operated the retrieval functions.

When a human pilot flew the vehicle, a radio control transmitter was used. This transmitter had two joy sticks and several other digital and proportional channels for causing actuator movement on the vehicle. The up-link subsystem provided a mechanism to mimic the human commands. Each of the commands on the radio control transmitter was sent via an analog signal to a receiver on the vehicle. If an actuator was to move a certain direction, a human pilot moved a joy stick some corresponding direction. The up-link subsystem allowed the computer to "move" the same joy stick an appropriate amount, thus sending the same command to the vehicle. Therefore, the actions a human operator provided to the radio control transmitter were simulated by computer, but in some cases, the stick function was altered from that traditionally used by radio controlled pilots.

6.8.4.2. TECHNICAL DESCRIPTION. The radio control transmitting unit used by human pilots to control the helicopter was retained in the design. However, a "bypass" switch was utilized so that the signals produced by the computer were spliced into the transmitter. The computer signals were the same signals as those produced by a human operator for controlling the vehicle. In this manner, whenever the "bypass" switch was closed, the computer was operating the helicopter. If a human operator sensed trouble with the vehicle, the "bypass" switch was opened and the human pilot regained control.

The up-link involved creating software inside the ground computer which produced a range of digital data values for each of the transmitter channels. A midrange value represented the "neutral" position of the transmitter joy-stick movement. When the computer software determined that an actuator should be moved in a particular direction, the digital data value was adjusted accordingly. The digital values were processed through an analog output (D/A) board for proper conformance with the signals and voltages in the radio control transmitter.

6.8.4.3. DESIGN DECISIONS. The design scenario for the up-link provided two key requirements for

the vehicle. One, the human override or "kill" switch required by the competition specifications was provided. Even more importantly, however, the vehicle could be flown by a human with all of the subsystems in place for testing and verification. A recording of human operator commands could also be captured and then used to develop the software in the ground station.

Significant amounts of data analysis were performed to measure the full spectrum of signals and accompanying voltages produced by the transmitter which the D/A board had to simulate. These signals were tuned on the computer board to match as best as possible the same control surface motions on the vehicle with the human pilot and the computer driven software.

6.8.4.4. ASSESSMENT. In the course of the project, the up-link D/A board was manufactured by our student team to save on costs. However, a more standardized D/A board would have been more reliable. One of the problems realized with the board was due to temperature variation and its influence on accurately matching the signal voltages.

6.9. *Power Distribution*

The systems both on the vehicle and on the ground required electricity to operate. The ground computer operated from a traditional 110 volt AC outlet. The on-board systems were all powered by batteries. Differing battery types were explored and the low cost of traditional NiCad rechargeable batteries was found to outweigh the weight penalty they imposed. Standard radio control battery packs were used in the vehicle where possible.

The vehicle did not have a homogeneous power distribution system since many subsystems were developed separately and provided their own power system. A central power regulator provided clean power for the down-link transmitter, but most systems were just connected directly to battery packs. There were over a half dozen different battery packs on the vehicle, each requiring recharging before the system could fly. A central umbilical cord that could be plugged into a wall socket was needed to allow

extended bench testing and evaluation during the integration phase.

A second problem was the duration of the battery power supply. The IVS boards consumed a large amount of energy. The battery packs allowed less than four minutes of operation before the cameras failed. This was not long enough for the three minute flight and potential six minutes of setup time. The use of a generator instead of batteries is being evaluated. It probably will weigh more than the batteries it replaces, but operational costs should be less and mission duration would not be limited by battery life. This central power source would also provide a convenient connection for an umbilical cord.

7. Integration Issues

Having discussed each of the components, attention again is focused on the system and integration efforts. Unfortunately, this is the area which, in large part, prevented the Georgia Tech team entry from performing on the competition day. All of the components were tested, and for the most part, were working, but integration was not complete. The following paragraphs describe some cursory issues concerning the component integration in this project.

Perhaps, some of the more important integration issues were based on the group organization and composition. Team participants were organized by discipline and located in many different labs across campus. The group structure was one which focused on subsystem performance, not on mission performance, resulting in inadequate attention to integration. Additionally, whenever attempting to integrate two components, someone had to deal with different lab locations, unfamiliarity with the tools available, and usually no authority in the integration location to obtain tools. If one common area had been available for the team, integration would have been easier and quicker.

A second major factor influencing the integration, as well as the design, was the amount of time to research and develop the prototype. From the beginning of the academic year (late September for Georgia Tech) to the competition date was approximately ten months. To assemble

a team, bring on industrial affiliates, design and produce the components, and then integrate resulted in an extremely tight schedule. Significant to this process is that most of the student team members were volunteering their time as an extra-curricular activity, already limited due to course activities for graduate and undergraduate degrees.

A better sense of the magnitude of integration efforts would have helped from the beginning. However, with a prototype vehicle, containing parts almost all of which were themselves prototypes, many unexpected events were bound to arise. Perhaps one of the biggest design decisions which influenced the project was one of mission timing. From the outset, the team strove for a three minute flight with six minutes of total setup time. This decision drove many other design decisions and was ultimately, not of tremendous importance to the sponsors during the competition.

Equally important was the decision to obtain a proven vehicle. Significant time and energy was expended in the direction, only to find that no single vehicle, without modifications, met our requirements. Therefore, a vehicle was selected and modified, but was not mature enough for other system mounting and testing until very late in the design phase.

More effort should have been given to address the integration from a planning perspective. Information about interfacing with other components was more of an afterthought than part of the design. Most of the design time was spent on getting components to function independently. Insufficient detail was given to self-monitoring, and communication of the status to other layers of intelligence. This would have allowed the mission planning component to function more robustly.

8. Evaluations

This project has demonstrated the need for, and benefits of, joint academic-industry endeavors. The students and faculty learned much by working on a real and tangible project. Industrial affiliates gathered experience for improving their own work and/or products, acquired research

and testing of new equipment, and saw potentials for products/services in new markets.

The project provided an excellent avenue for stimulating interdisciplinary work between individuals, schools, and university-industrial participants. All of the participants learned a great deal about the new and exciting technology of autonomous robotics. They also learned that hurdles still remain before this kind of vehicle becomes a common occurrence.

The authors' believe that the single biggest obstacle preventing realization of this technology is integration. Developing standard interfaces between components, further maturing the prototypes so that components are more reliable, and developing a system architecture which accommodates these issues is necessary.

The result of the Georgia Tech entry has formulated a working model for an intelligent aerial vehicle. This attempt includes a structured system architecture, particularly in terms of multiple levels of intelligence and the requisite communication needed. However, significant refinement is needed before an adequate analysis or evaluation of the total system can be made.

Through projects such as this, individuals have learned and experienced a great deal. We applaud the AUVS for creating the contest and eagerly look forward to participating again.

Acknowledgments

The authors would like to acknowledge the support and contributions from the many individuals who provided the necessary assistance on the project.

We would like to thank the Association for Unmanned Vehicle Systems, and its representative, Rob Michelson for sponsoring the competition. The success of the event was measured in terms of knowledge gathered by many players, not only on the Georgia Tech entry, but also by other contestants.

We wish to acknowledge those students on the Georgia Tech team who spent many hours working on this project: Todd Cargle, Brian Clark, Russ Clark, Chris Coughlin, Robbie Cowart, Brian Eversmann, Mark Gordon, Steve Ingalls, Shayne Kondor, Andreas Lipp, Doug Mac-

Kenzie, Victor Nelson, Walter Patterson, Dave Paulli, Dennis Rossman, Steve Smith, and Mike Tinski.

We would also like to acknowledge the faculty advisors on the project: R. Arkin, N. Baker, S. Dickerson, B. Heck, J.V.R. Prasad, D. Schrage, and G. Vachtsevanos.

Without the financial and equipment support of the following individuals and organizations, the project would have never succeeded. These are: Gary Pochlein, Vice President, Georgia Tech Interdisciplinary Programs Office; Greg Walker and Art Phelps of the Army Aerostructures Directorate colocated at NASA Langley Research Center; George Makawiec of Vigyan Incorporated; J. Eric Corban of Guided Systems Technologies; John Smith and John Moore of Pacific RPV, Incorporated; Roger Creech, Bill Grass, and Dr. Kim of TRON-Tek; Marc Slack of Mitre Corporation; Watson Sensors; and the Digital Equipment Corporation.

Finally, the following individuals and organizations provided assistance in many ways to the project. C. McKeithan, Aerospace Engineering; J. Craig, Aerospace Engineering; each of the home Schools and Colleges of the participating students and faculty; the Aerospace Engineering shop; Richard Daniel and Joe Hulsey of the Cobb County Radio Control Club; and, Ken Mauragas and Howard Cooley of the Roswell Air Force.

References

1. Association for Unmanned Vehicle Systems, First International Aerial Robotics Competition, 1991.
2. Department of Energy, Program Research and Development Announcement (PRDA) Number DE-RA21-92MC29245, "Research and development for environmental restoration and waste management technologies," *Commerce Business Daily*, August 20, 1991.
3. R. Bajcsy, Report from the NSF-DARPA Sponsored Workshop on: *Identification of Research Needs for a Mobile Platform with Material Handling Capabilities*, National Science Foundation, March, 1991.
4. R.I. Winner, et al., *The Role of Concurrent Engineering in Weapons Systems Acquisition*, Report R-338, Institute for Defense Analyses, Alexandria, VA, 1988.
5. Unmanned Aerial Vehicle Joint Project Office, Department of Defense, Washington, D.C. 20361-1014, Draft Unmanned Aerial Vehicle System Engineering Management Plan, 1989.

6. J.R. Hauser and D. Clausing, "The house of quality," *Harvard Business Review*, May/June, 1988.
7. S.A. Ingalls, "Application of concurrent engineering methods to the design of an autonomous aerial robot," Master's thesis, School of Aerospace Engineering, Georgia Institute of Technology, November, 1991.
8. B.C. Kuo, *Automatic Control Systems*, Prentice-Hall, fourth edition, 1982.
9. J.S. Albus, "A theory of intelligent systems," *Proceedings of the Fifth IEEE International Symposium on Intelligent Control*, September, 1990.
10. W.K. Kershner, *The Student Pilot's Flight Manual*, Iowa State University Press, 1987.
11. J. Arditi, "Design procedure for a UAV data link," *Unmanned Systems*, 7(2), pp. 32-37, Spring 1989.
12. N.C. Baker and W.M. Patterson IV, *Project Report: Digital Transmission Using the TRON-Tek Video Transmitter*, Georgia Institute of Technology, School of Civil Engineering, 1991.



Nelson C. Baker obtained a BS in Civil Engineering from Georgia Tech in 1980, an MS in Civil Engineering from Carnegie Mellon University in 1985, and a PhD in Civil Engineering from Carnegie Mellon University in 1989. Dr. Baker's research focuses on the representation and use of spatial information for quantitative and qualitative reasoning for robotics and engineering design. These research activities are in conjunction with several projects including a robotic excavation project, an autonomous aerial robotic project, and a structural layout design advisor. As faculty advisor for the aerial robotic project, described in this paper, Dr. Baker supervised the data link and world modeling activities.



Douglas C. MacKenzie received the B.S. degree in Electrical Engineering in 1985 and the M.S. degree in Computer Science in 1989, both from Michigan Technological University. He was employed as a Software Engineer in Allen-Bradley's Programmable Controller Division in Cleveland, Ohio from 1988 until 1990. He has been a Ph.D. student in the College of Computing at Georgia Tech since 1990.

Captain Stephen A. Ingalls is an officer in the United States Army. He received his master's degree in Aerospace Engineering from Georgia Tech in March 1992 with a special emphasis concerning concurrent engineering techniques as applied to an autonomous aerial vehicle development project. Mr. Ingalls received his Bachelor of Engineering from the United States Military Academy at West Point in 1982. He is an Army pilot with nearly 800 hours of flight time rated in a variety of military helicopters. Captain Ingalls is currently assigned as a Civil and Mechanical Engineering faculty member at West Point.