The USC AFV-I: A Behavior-Based Entry in the 1994 International Aerial Robotics Competition

James F. Montgomery
monty@robotics.usc.edu

Andrew H. Fagg
ahfagg@robotics.usc.edu

George A. Bekey
bekey@pollux.usc.edu

Robotics Research Laboratory
Institute for Robotics and Intelligent Systems
University of Southern California
Los Angeles, CA 90089-0781

The International Aerial Robotics Competition is an annual event sponsored by the Association for Unmanned Vehicle Systems (AUVS). The competition requires flying robots to locate, manipulate, and transport objects from one location to another. These tasks are carried out under extreme conditions. A robot must operate without human guidance, ensuring system survival, while completing the goals of the competition within a fixed time limit.

Creating a flying robot with these capabilities presents many challenges. The robot has to make control decisions that achieve its goals based upon imperfect sensory data while adapting to unexpected situations such as gusts of wind or sensor failure. Additionally, these decisions need to be made in real-time to maintain the safety of the craft.

Teams of university students from around the world enter the competition, using a variety of approaches. The University of Southern California Robotics Research Laboratory's Autonomous Flying Vehicle-I (AFV-I)[1] finished first out of over twenty schools entered in this year's competition held on May 19. The robot uses a behavior-based approach as its control architecture. A set of behaviors, each responsible for achieving a specific task, interact to achieve the overall goals of the robot.

A behavior-based approach has many advantages over traditional methods of creating autonomous mobile robots[2]. Previous methods attacked the problem of robot control in a sequential manner, where a robot first senses, then perceives, models, plans, and acts in its environment. Since the world is information rich, the traditional method is prone to information overload, rendering it incapable of functioning in real-time, with possibly dire consequences. In addition, this method assumes that accurate global world models can be constructed from the incoming sensory information. A number of factors conspire to make this difficult, such as a rapidly changing world, limited computer processing power, and inaccurate, incomplete sensor models.

Conversely, a behavior-based approach solves the problem in a parallel fashion. Each behavior, acting concurrently, extracts from the environment only the information required to complete a given task at a given time. This, coupled with the elimination of a need for construction and maintenance of a global world model, greatly reduces the computational load on the robot.

Another advantage of the behavior-based approach is the ability to create layers of increasingly complex behaviors on top of simpler behaviors. If need be, the lower level behaviors are inhibited or modulated by higher level behaviors. In this way, a robot control system with increasing capabilities may be built and tested incrementally, without losing low-level capabilities already created.

However, the behavior-based approach has its own limitations. The interaction and possible couplings between behaviors are unknown a priori, but they may be crucial to the stability of the craft. It may be necessary to determine the couplings experimentally. Since no models are available (one of the strengths of the approach can also be a problem), this experimentation may be time consuming and potentially hazardous to the craft. This coupling problem only worsens as behaviors and layers increase, which creates problems when trying to expand system complexity.

The AUVS International Aerial Robotics Competition

The goal of the competition is to create an autonomous flying robot capable of carrying out a set of predefined tasks in a competition arena set up on a football field. The robots lift off from a designated starting area within the arena (Figure 1), and locate a black source ring that is 6′ in diameter with a 3″ high lip. Placed randomly within the ring are six 3″ diameter, day-glow orange disks. The goal is to transport the disks, one at a time, to an identical destination ring at the other end of the arena, 80′ away. A

3′ high central barrier separates the rings. The robot may only make contact with the ground in the starting area and in the rings. There is a six minute time limit in which to transfer the six disks.

Although no robot has completed all requirements of the competition held annually since 1991, a number of robots have achieved partial success. Autonomous lift off, hovering, navigation to a desired location, and landing has been shown by a Georgia Tech team. Other teams such as the U.S. Naval Academy, University of Texas at Arlington and ourselves have demonstrated a subset of these achievements. A variety of disk retrieval devices have shown the ability to locate and acquire disks after being moved to the source ring through human intervention, although not while attached to a flying robot. There are many challenges yet to be mastered, indicative of the overall difficulty of the task.

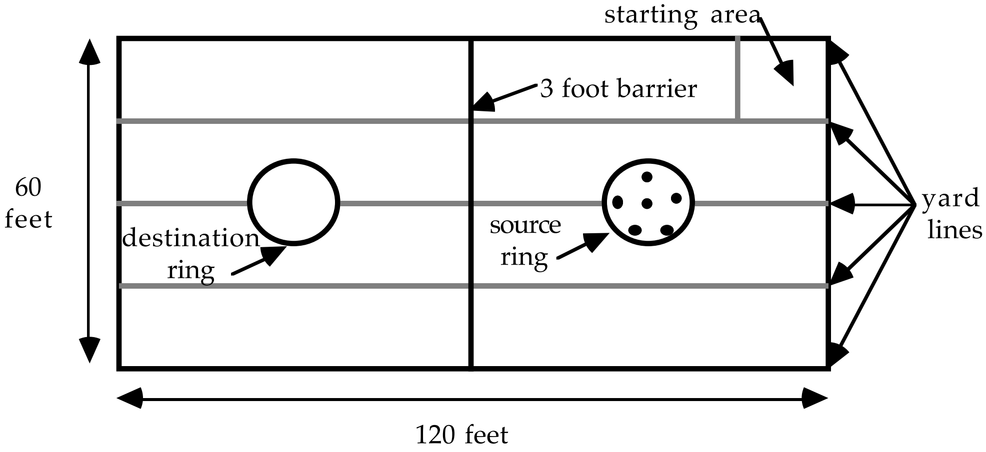

Figure 1. Competition arena layout

AFV-I SYSTEM DESIGN

A guiding design philosophy of the USC Robotics Research Laboratory is to create robots with ever expanding limits of autonomy. Although simulation is useful in examining issues of robot autonomy, it is not enough. In order to fully study autonomous agents, it must be done within the environment that they will ultimately operate. Therefore, they must be embodied and situated in the real world. These beliefs drive the design of AFV-I. Attributes that help define a robot’s level of autonomy are given:

(1) The degree of dependence on a human, ranging from total in the case of a teleoperated robot, to none in the case of a completely independent one.

(2) The degree of reliance on external resources (not part of the robot itself), such as sensory and computing devices.

(3) The difficulty of missions that are achievable. This can be measured by parameters such as duration of mission, number and variety of tasks to achieve, and degree of decision making capability required to achieve mission.

(4) The amount of structure needed in the environment to function. This includes natural (rivers) as well as artificial (roads) structure.

(5) The level of fault-tolerance of the system. The ability to recognize and overcome failures such as sensor or actuator loss.

(6) The degree of adaptability to unexpected situations such as getting lost or encountering an unknown obstacle.

(7) The learning level of the robot, measured by the ability to expand current capabilities with minimal or no human guidance.

These attributes are not complete by any means. However, they do ground research in building autonomous robots by providing measurable benchmarks of autonomy. These attributes are partially examined in the context of AFV-I and the competition requirements. Each attribute is attached a rating of the degree to which it has been satisfied of none, low, med, and high.

The AUVS rules state that a robot must function without aid from a human during its attempt to achieve the competition goals (1,high). A robot can make use of any external resources as long as they are outside the boundary of the arena (2,low). To carry out the competition goals, robots must complete a variety of tasks in a short amount of time (3,med). The arena has a fair amount of structure, with yard lines and rings in known locations, that the robots can use for navigation (4,low). No stipulations are made regarding the fault-tolerance, adaptability, or learning capabilities of the crafts (5,6,7,none).

For AFV-I, we further restrict two of the attributes of autonomy defined above, beyond the constraints of the competition. It is felt that the competition allowance of external resources is not restrictive enough. Approaches relying on external computing and sensing will have a difficult time scaling up to new situations where the field of operation is much larger or filled with obstructions. This makes it infeasible to do remote sensing and unless a robot’s radio link is robust, problematic to do remote computing for control. Therefore, we allow no external resources to be used; system power, sensors, and computers are all located on the robot itself. This has implications for the control system design. Weight-lifting limitations of the craft constrain the power consumption of the system as a whole, amount of computing power carried, and choice of sensors used. Also, the sensors are no longer positioned in a global frame of reference, and only supply egocentric-relative information (except for a compass, see Hardware below). These factors conspire to make building a control system much more difficult. We believe (and to some degree demonstrated) that a behavior-based approach is a feasible solution to this problem, and will scale to larger domains of operations, unlike the external resource based solutions. They will be unable to scale due to limitations in sensory information and computing power (2,high). The local environment AFV-I finds itself in is not further engineered beyond what is specified in the competition. For example, an external tracking mechanism is not used to help the robot locate its position. The robot is aware of the location of certain landmarks (yard lines, rings) in the arena it can make use of to determine its position (4,med). The remaining attributes have not been given much attention to date.

Hardware

A Kyosho Concept 60 RC helicopter powered by an Enya 80 nitro-methane fueled two-stroke engine serves as the robotic platform (Figure 2). It has five degrees of control: main rotor aileron and elevator cyclic pitch, tail rotor pitch, main rotor collective pitch, and throttle. The first three control the roll, pitch, and yaw of the craft, while the last two control its thrust. The aileron and elevator cyclic controls the roll and pitch of the helicopter by creating a lift differential in the main blades. This is accomplished by decreasing the lift of one blade and increasing the lift of the other, with these control changes occurring once per main rotor revolution at the same position in the revolution, thus the term cyclic. Where and by how much this change occurs in the revolution determines to what degree the craft will roll and/or pitch. The collective controls the thrust of the helicopter by changing the pitch of the main rotor blades by the same amount (collectively) creating either an increase or decrease in total lift of the craft.

A variety of sensors are mounted on the craft; a flux-gate compass for measuring heading, three downward facing ultrasonic sensors (two mounted on a crossbar on the front of the robot and one mounted on the tail boom at the rear) for determining the roll, pitch, and height of the craft, an RPM sensor mounted on the main rotor mast for measuring engine speed, and a gray-scale CCD camera to provide visual information. Three solid-state rate gyros are used to dampen the roll, pitch, and yaw of the robot. The disk retrieval device is based upon a ferromagnet that uses a tactile sensor which indicates when a disk has been successfully acquired.

There are two Motorola 68332 microcontroller-based custom-built computer boards on the robot; one to collect data from the sensors and control the actuators based upon the data, and the other a dedicated vision board to process the CCD camera visual information. There is a two-way communication link between the two boards. A set of nickel-metal-hydride batteries supplies power to the electronics.

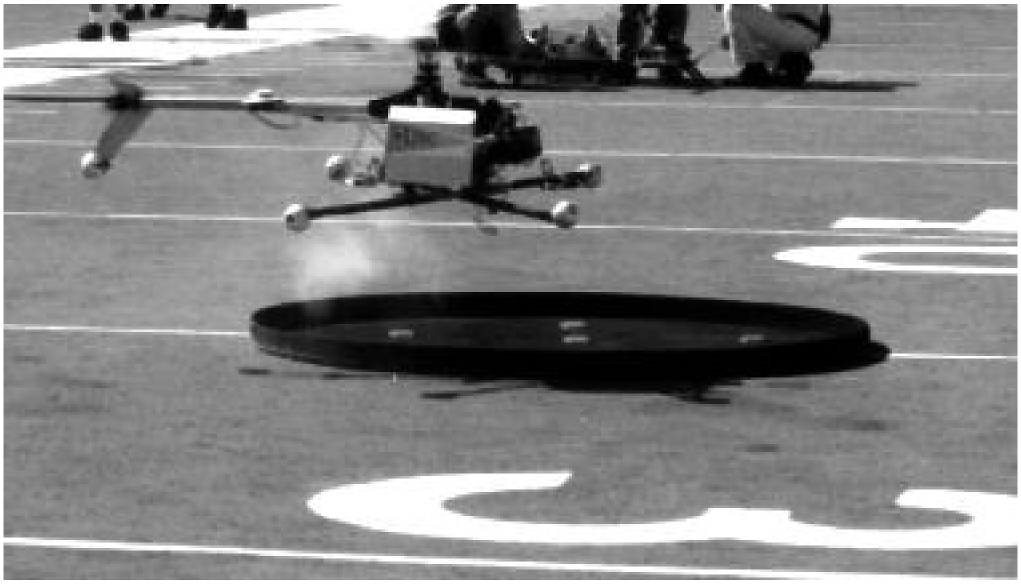

Figure 2. AFV-I hovering over the source ring during its winning flight

Control Architecture

Shown in Figure 3 is the AFV-I behavior-based control system architecture.

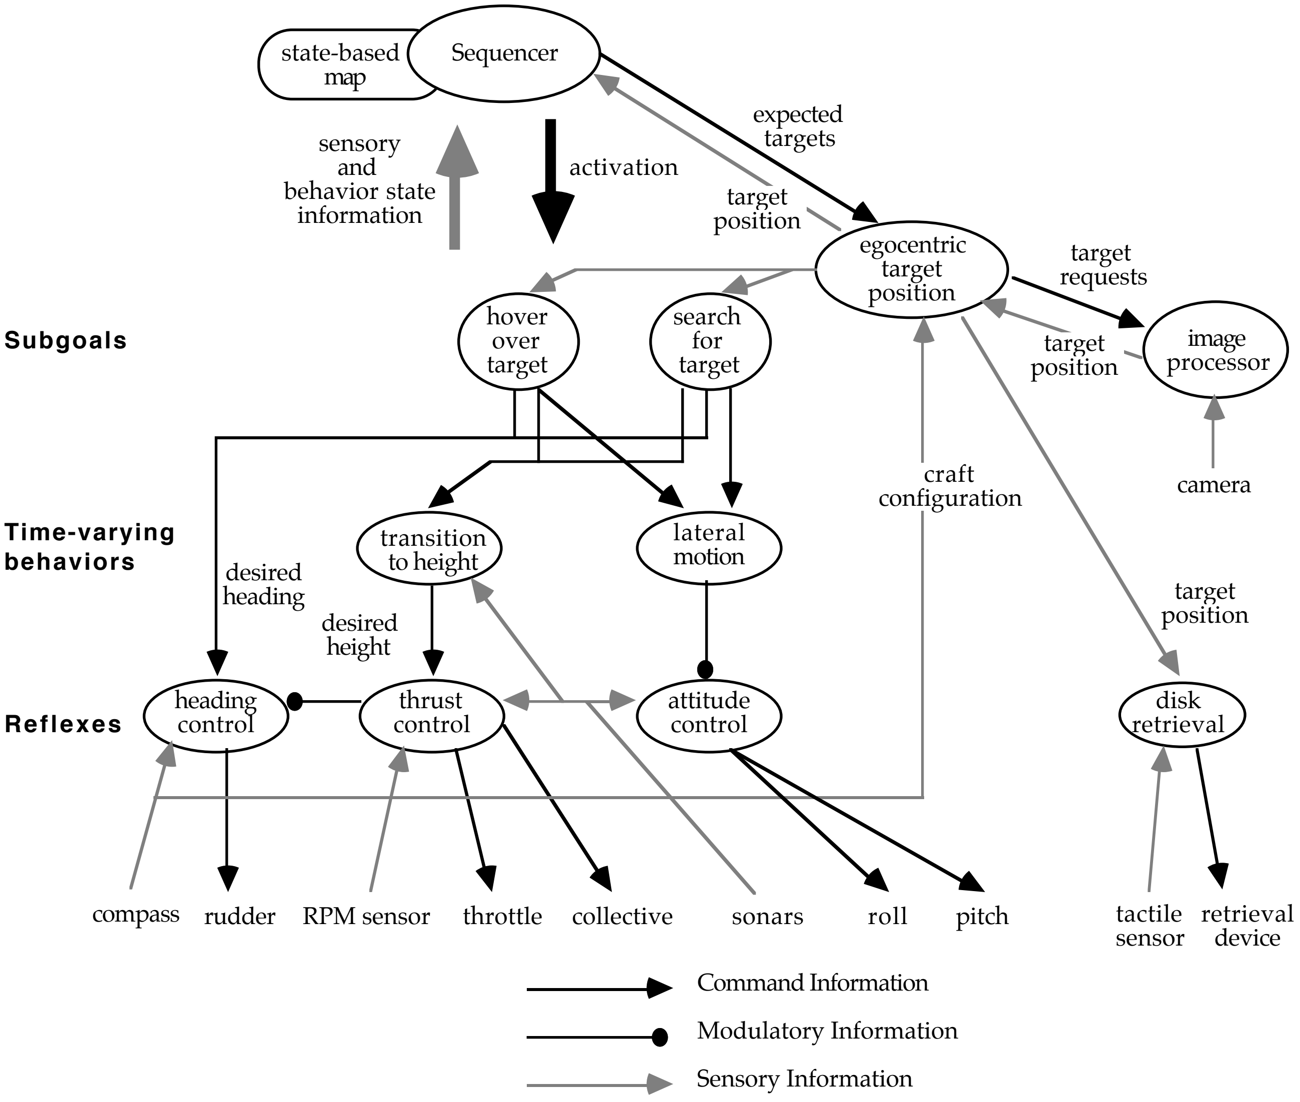

Figure 3. Behavior architecture

At the lowest level of control, survival is the main priority. To this end, the robot has a set of fast acting reflex behaviors that attempt to maintain system stability by holding the craft in a hover pointed in a desired heading. When the robot detects deviations the appropriate reflex returns the craft to its stable configuration.

The **heading control** behavior attempts to hold a desired directional heading by using the compass data to drive the tail rotor pitch. The **thrust control** behavior uses the sonar and RPM sensor data to control the collective pitch and throttle. This behavior is responsible for maintaining a desired height above the ground. The **attitude control** behavior tries to hold a stable hover (zero roll and pitch orientation and rate). It uses the sonars to determine attitude and then controls the aileron and elevator cyclic pitch to keep all three sonars the same calibrated distance from the ground. Assuming the robot is flying over level ground and all three sonars read the same distance from the ground, the robot will be in a zero roll and pitch orientation. A pid controller is used for the **thrust control** behavior, while a pd controller is used for the **heading control** and **attitude control** behaviors. Different levels of thrust will induce varying amounts of torque about the yaw axis of the robot. To counteract this, the **thrust control** modulates **heading control** in a feedforward manner. The **disk retrieval** behavior uses target information from the **egocentric target position** behavior (see below) and tactile information to control the retrieval device.

The two time-varying behaviors inject inputs into the reflex behaviors to alter their desired low-level goal of maintaining system survival. These inputs change dependent upon the current high-level goals of the robot. This demonstrates a key advantage of the behavior-based approach. Greater capabilities are created for the robot by layering more complex behaviors on top of previously constructed behaviors. This addition is transparent to the lower level behaviors, modulating, but not destroying, their underlying functionality.

The **transition to height** behavior inputs a desired height to **thrust control** to move the robot to a new goal height. The **lateral motion** behavior generates a sinusoidal modulatory signal input to **attitude control** causing the robot to roll and pitch out of its hover orientation in a specified direction for a given length of time.

The next layer of behaviors is responsible for achieving the subgoals of the robot. The **search for target** behavior moves the robot in the direction of the desired target to be acquired. This direction parameter is instantiated by the **sequencer** (see below). The **hover over target** behavior is activated once a target has been visually located and it commands the craft to maintain a position over the target. Both of these behaviors get target input from the information providing **egocentric target position** behavior. This behavior reports whether a target has been found and at what location. **Search for target** is active when a target is not found, while **hover over target** is active after a target is found by **image processor.**

The **image processor** behavior is another information provider, responsible for locating targets in the CCD camera imagery. This behavior receives target requests from the **egocentric target position** behavior depending upon which part of the mission goals are next to be achieved and returns positions of located targets back to it. Targets are locations of vertical and horizontal lines and centroids of blobs (disks or lines).

The **sequencer** at the highest level of control determines which behaviors to activate and what parameters to instantiate to achieve a desired subgoal. It makes these decisions using the **state-based map** which contains the information necessary to carry out the competition mission: what targets to look for, the order in which to look for them, the behaviors to activate and the parameters to instantiate. After the **sequencer** determines the subgoal is reached based upon sensory and behavior state information, it repeats the activation and instantiation process for the next subgoal. This continues until the overall mission goals are completed.

COMPETITION STRATEGY

The **sequencer** determines how the underlying behaviors are activated to achieve a particular mission. If the current suite of behaviors is insufficient to carry out a desired task, a new behavior can be formulated to fill the hole in the robot’s abilities. The current set of behaviors as shown in Figure 3 should be sufficient to fulfill the competition goals.

The sensors other than the CCD camera are used to achieve attitude, height, and directional control but are useless for driving lateral movement of the robot or finding targets. Thus, AFV-I depends heavily upon vision to interact within its environment. Since the competition takes place on a football field, there are some very convenient landmarks (yard lines) in the competition arena that the robot knows about and can use.

Vision driven behavior sequencing

The vision system serves two purposes. It augments **attitude control** indirectly through **lateral motion** when **hover over target** is active. It is also used by the **sequencer** to determine when a particular subgoal has been met and it is time to work toward the next subgoal by activating the appropriate set of behaviors.

The vision processor utilizes an image subsampling strategy to reduce the camera information from 250K pixels to 2K pixels. The processor has three modes of operation:

- Track a horizontal and a vertical line
- Track a blob

- Track a blob while looking for another blob in the center of the camera’s field-of-view (FOV).

The robot is commanded to headings either perpendicular or parallel to yard lines simplifying the image processing task of finding vertical and horizontal lines. With the limited amount of information for each image and the simple features, we are able to process 8-10 frames/second.

Figure 4 illustrates the strategy for carrying out the entire competition sequence. Each square annotated by a letter represents the camera's FOV at a given time. The squares are used for discussion purposes only and are not to scale.

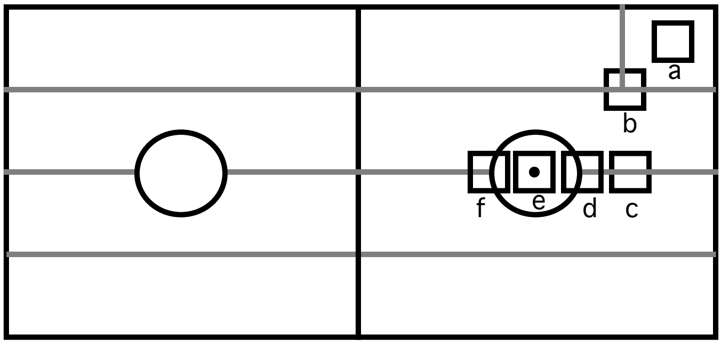

Figure 4. Visual targets in the arena

The three reflex behaviors are activated by the **sequencer** and stay active throughout the mission. The remaining behaviors are activated by the **sequencer** as required. Before lift off from the starting area (a), the robot will see nothing but green field. The first landmark to locate is the ‘T’ created by the intersection of a yard line and the border of the 15’x15’ starting area. The robot’s yaw axis is oriented parallel to the yard lines, as a result they will appear as vertical lines in the FOV, while the starting area border perpendicular to the yard lines will appear as a horizontal line.

The **sequencer** instantiates both vertical and horizontal lines as targets, giving that information to **egocentric target position**. The **image processor** is given these target requests, while **search for target** is activated and **hover over target** is primed in preparation for when the targets come into view. Only one of these two behaviors is active for each target, depending upon whether a particular target is visible. In this case there are two targets, a vertical and a horizontal line. It is possible to see the vertical line and have **hover over target** be active for it, while still in active **search for target** mode for the horizontal line. At the start, neither target is visible, so **search for target** causes **lateral motion** to modulate **attitude control** to move the robot toward the 'T’, while giving **heading control** a desired heading and commanding **transition to height** to input a desired height to **thrust control.**

The remaining steps are examined in less detail, but the same flow of control occurs, with behaviors acting in parallel and the higher level behaviors affecting the lower level ones. When the ‘T’ is located, the visual target information is used to center the ‘T’ in the FOV as shown in (b). Once (b) is reached, the robot will stay until a stable hover, height, and heading is demonstrated. The **sequencer** instantiates a vertical line as the next target and the robot begins lateral translation in the appropriate direction until (c) is achieved, with the vertical line centered in the camera image.

The next goal is to move toward the source ring. This is accomplished by shifting the desired location of the “center of mass” of the visible part of the vertical line down in the FOV. This is a perceptual trick, causing a lateral movement of the robot forward and toward the ring. The vertical line is tracked while blobs (disks) are searched for. When (d) is reached, and a disk located, **disk retrieval** is activated to drive the lateral motion of the robot to center the retrieval device over the disk (e). Retrieval attempts are made until a tactile sensor on the device indicates that a disk is acquired. The vertical line again becomes the target and, if necessary, height is increased, until (f) comes into view. The robot uses the line to traverse toward the destination ring, looking for the 3' high barrier along the way. When crossing over the barrier, the expected shorter distance indicated by the sonars will be ignored until the robot has reached the other side. Lateral movement forward continues until the destination ring is located using the same approach as when locating the source ring. The disk is released into the ring, and the robot returns to the source ring for another disk following the yard line as before. This procedure continues until all six disks are retrieved at which point the **sequencer** commands the robot to land and terminate operation.

PROJECT STATUS

At this year's competition, the University of Southern California, University of Texas at Arlington, and Southern College of Technology finished in first, second, and third place, respectively.

During the competition, AFV-I completed a number of autonomous flights, with the longest duration being 29 seconds. The craft demonstrated the implementation of the **heading control**, **thrust control**, **attitude control**, **transition to height**, and **lateral motion** behaviors. These behaviors performed well, maintaining robot stability even in the presence of substantial winds that shifted direction and speed constantly.

A number of accomplishments have been achieved. We have built an autonomous flying robot, dependent upon only its internal resources, capable of sustained and safe flight. We have demonstrated the ability of the behavior-based approach to fly a model helicopter autonomously for short missions. This was accomplished using imperfect sensory information, making effective control decisions in real-time, and flying a craft with complex flight dynamics without crashing.

There are many issues remaining to be resolved. Due to atmospheric conditions in Atlanta on the day of the competition that reduced the lift capability of the craft, we had to remove the vision system (CCD camera and vision board) to allow the robot to rise out of ground effect. This meant it was flying blind; able to control its height, heading, and attitude, but not its lateral position. We are investigating the use of a new helicopter capable of lifting all robot subsystems. Integration of the vision system into the robot control system is of high priority since it is necessary to carry out the rest of the competition requirements.

More research into the attributes of autonomy defined earlier need to be addressed. We have primarily concentrated on the first four with limited attention to the last three. Using the numbering scheme defined earlier, the state of our exploration is summarized. The robot is operates totally independent of human intervention and external resources (1,2,high). Mission difficulty is defined by the competition requirements (3,med). We reject further modification of the structure of the competition environment (4,low). Fault tolerance issues remain, with many single point failures that need to be addressed. The craft has to detect faults and recover, if possible. Failures range from the benign such as the robot becoming lost, expecting to find a yard line but does not, to the catastrophic such as a sonar failure Detecting and handling failures fits well with the behavior-based approach. Behaviors can be created that will monitor the system for specific failures and take the appropriate action when recognized (5,none). It has limited ability to adapt to uncertain situations such as gusts of wind or imprecise knowledge of the exact location of the disks inside the source ring. However, shifting the source-ring over one yard line or increasing the height of the barrier to 10 feet would be beyond the adaptability of the robot (6,low). A variety of control parameters, such as the actuator center values necessary for a stable hover, are currently adjusted through human trial and error. Instead, a rapid, on-line learning of these parameters is desirable (7,none).

FUTURE RESEARCH

Our research to date has been driven by the needs of the competition and our own desire to exceed these needs. We have examined the behavior-based approach to autonomous robot control, experimented with a variety of sensors, explored vision based navigation, and probed various dimensions of autonomy. In order to create robots capable of even more complex applications, additional research is required.

Applications for autonomous flying robots include those that are too dangerous or monotonous for humans; such as performing surveillance behind enemy lines during wartime, cleaning up toxic waste, fighting forest fires, delivering food to refugees in areas of civil unrest, inspecting miles of remote pipe and power lines, or crop-dusting agricultural fields. Robots capable of these applications will require additional abilities such as highly maneuverable, long-range flight and increased manipulatory power. This drives the need for more accurate and varied sensory information with increasingly sophisticated control architectures capable of interpreting and acting upon this information.

To date, our current control architecture has been hand designed. As was described earlier, it can be difficult increasing the complexity of a behavior-based system due to the possible coupling between behaviors. Therefore, it is desirable that methods be developed to overcome this weakness. Due to a variety of previous research described below, it is believed that integration of a fuzzy rule based system with the behavior-based controller is a promising approach. The strengths of the behavior-based approach can be maintained while the weaknesses reduced or eliminated.

Fuzzy systems have demonstrated the capability of dealing with the uncertainty in unstructured, real-world environments for a variety of applications, autonomous robots included. For a number of systems, the fuzzy based approach has been shown superior to other conventional approaches when comparing performance results, the Sendai subway in Japan being an example[3]. Although research has been limited in the area of fuzzy logic controlled autonomous flying vehicles, promising findings do exist[4]. For these reasons, implementing an AFV using a fuzzy/behavior-based system capable of achieving the goals of the AUVS IARC is a specific goal.

For a system as complex as an autonomous flying helicopter, determining and organizing the rules is not trivial. Automatically generating these rules is a desirable approach. With the current AFV-I control system, various time-varying inputs and outputs can be gathered. Using this training data with techniques such as those described in[5,6], patterns can be recognized and rules generated. It is thought, based upon this previous research, that the rules generated will be sufficient to assure the survival of the robot. Then these rules can be automatically tuned through techniques such as reinforcement-learning to optimize system performance.

Additional fuzzy research could be in the augmentation of a traditional controller, where rules are specified for regions of the state space where the conventional control system did not work well. This can be implemented by creating a hybrid control system, where a mixture of the fuzzy and traditional outputs are merged to varying degrees. Proportions are varied dependent upon which controller worked better for a specific region of the state space. This could be useful in overcoming the traditional controller’s difficulty in dealing with the difference in thrust required for altitude control in and out of ground-effect.

Another direction of research being pursued is that of navigation. The current system makes use of a state-based map. At each state, the control system stores the appropriate action to take and the sensory input that will signal the next state transition. When the environment changes, however, this will require a complete reconfiguration of the state-based map. A more general approach being pursued stores the environmental information in a spatial map, from which the action/state-transition pairs are automatically derived given a general specification of the path to be taken by the craft. This approach can be made more robust by tracking multiple landmarks at any given time, allowing for cases where some landmarks are not visible or are in different positions than predicted by the map.

Navigation for larger-scaled missions is also of primary interest. Here, sensory systems such as GPS (Global Positioning System) can give the control system general positioning information within a large map. However, the control system will continue to make use of existing sensory sources to maintain the general safety of the craft and to localize the craft's position on a finer scale using visual landmarks. This level of research opens up a whole host of other problems, including representation of the maps, how to combine the coarse map information with the locally available information to make control decisions, and how to update maps based upon local information.

ACKNOWLEDGMENTS

We thank the rest of the current and past USC AFV-I team members; Joel Hoff, Scott Cozy, Scott Souter, John Gullickson, Steve Goldberg, Tony Lewis, Gerry Kim, Gaurav

Sukhatme, Parag Batavia, Mike McHenry, and Dave Lotspeich. We are grateful to our sponsors; MTA Hobbies of Torrance, CA, GP Batteries of Carlsbad, CA, Sundance Model Products of Newbury Park, CA, and the Gordon Marshall Fund at USC. Thank you to Olaf Kubitz for the picture of AFV-I hovering over the source ring.

REFERENCES

1. A. H. Fagg, M. A. Lewis, J. F. Montgomery, G. A. Bekey, “The USC Autonomous Flying Vehicle: An Experiment in Real-Time Behavior-Based Control”, Proceedings of the 1993 IEEE/RSJ International Conference on Intelligent Robots and Systems, IEEE, Piscataway, NJ, 1993, pp. 1,173-1,180.

2. R. A. Brooks, “A robust layered control system for a mobile robot.”, IEEE Journal Robotics and Automation, 1986. **RA-2**(RA-2): pp. 14-23.

3. S. Yasunobu and S. Miyamoto, “Automatic train operation by predictive fuzzy control,” In Industrial Applications of Fuzzy Control, M. Sugeno, editor, pages 1-18. North-Holland, Amsterdam, 1985.

4. M. Sugeno, M.F. Griffin, A. Bastian, “Fuzzy hierarchical control of an unmanned helicopter,”, presented at IEEE International Conference on Fuzzy Systems,

(FUZZ-IEEE ‘93). The paper did not make the proceedings but it is available via anonymous ftp from host cs.arizona.edu in the directory /japan/kahaner.reports with filename copt-fuz.93.

5. J. A. Dickerson, B. Kosko, “Fuzzy function approximation with supervised ellipsoidal learning,” Proceedings of the 1993 World Congress on Neural Networks. Volume II, pp. 9-17.

6. Y. Nakamori, M. Ryoke, “Identification of fuzzy prediction models through hyperellipsoidal clustering,” IEEE Transactions on Systems, Man, and Cybernetics, Vol. 24, No. 8, August 1994.