Carrier Phase GPS and Computer Vision for Control of an Autonomous Helicopter

Bruce R. Woodley Henry L. Jones II Edward A. LeMaster Eric W. Frew
Prof. Stephen M. Rock
Aerospace Robotics Laboratory, Stanford University

September 18, 1996

BIOGRAPHY

Bruce Woodley is a Ph.D. candidate in the Department of Electrical Engineering at Stanford University. He received his MSEE from Stanford in 1996 and a Bachelor of Applied Science in Engineering Physics from Simon Fraser University in 1993. Email: woodley@sun-valley.stanford.edu

Hank Jones is a Ph.D. student in Aerospace Engineering. He received his MSAA from Stanford in 1996 and his BSME from the University of Mississippi in 1995. Email: hlj@sun-valley.stanford.edu

Ed LeMaster is a Ph.D. student in Aerospace Engineering. He received his MSAA from Stanford in 1996 and a BSAAE from the University of Washington in 1995. Email: lemaster@sun-valley.stanford.edu

Eric Frew is a Ph.D. student in Aerospace Engineering. He recieved his MSAA from Stanford in 1996 and a BSME from Cornell in 1995. Email: ewf@sun-valley.stanford.edu

Prof. Stephen Rock is an Associate Professor in the Department of Aeronautics and Astronautics at Stanford University. Email: rock@sun-valley.stanford.edu

ABSTRACT

The Stanford University Aerospace Robotics Lab (ARL) has developed a fully autonomous hoverable unmanned air vehicle (HUAV) which uses differential carrier phase GPS (DCPGPS) as the only sensor for both control and navigation. Precision autonomous operation in a structured environment has been demonstrated at the 1995 Aerial Robotics Competition, where the ARL helicopter demonstrated retrieval and transportation of 4 oz objects around a small field.

The ARL has experimentally demonstrated the advantages of combining differential carrier phase GPS for high-bandwidth vehicle control with computer vision to enable the operation of an unmanned, autonomous helicopter in a dynamic unstructured environment. DCPGPS techniques and conventional control methods are used to close a short-period inner loop, stabilizing the helicopter's dynamics and providing a global navigation system. An outer loop is closed using feedback from the vision system to determine the precise location of the ground and other objects in the environment.

This paper examines the combination of differential carrier phase GPS and computer vision techniques to enable an autonomous vehicle to operate in unstructured environments. Experimental results are presented and future work is discussed.

INTRODUCTION

Helicopters have found broad use in many diverse applications, including search and rescue, fire fighting, fine-scale terrain mapping, and agricultural operations. The operation costs of manned helicopters are quite high. It is expected that the operational use of unmanned helicopters could greatly reduce the cost of many of these tasks.

The remote piloting of hoverable unmanned air vehicles is a very difficult task, requiring great operator skill and attention as well as a high-bandwidth, low-delay data link between the pilot and vehicle. Furthermore, most applications require that the pilot maintain visual contact with the HUAV at all times. These requirements limit the applications of HUAVs and prohibit large scale deployment of these systems.

Many airborne applications require operation in dynamic, unstructured environments. Imaging is the most natural method of obtaining information about an unknown environment, and has broad applications in military reconnaissance, agricultural inspection, and mapping. Imaging plays a critical role in the operation of HUAVs, providing information to ground based operators and enabling operation in an unstructured, dynamic environment.

The Stanford Aerospace Robotics Laboratory has been working towards developing prototype HUAV systems which address these issues. This paper examines the combination of differential carrier phase GPS and computer vision techniques to enable an autonomous vehicle to operate in unstructured environments.

EXPERIMENTAL SYSTEM

The HUMMINGBIRD autonomous helicopter was originally constructed for research in the control of an unstable flying machine using Global Positioning System measurements as the only sensor. The hardware system, shown in Figure 1, has now been extended to deal with the combination of GPS measurements with vision system measurements. The basic airframe is a heavily modified Schluter Futura hobbyist model helicopter. The airframe is capable of carrying twenty-five pounds of experimental equipment. Nominal gross takeoff weight is forty-six pounds.

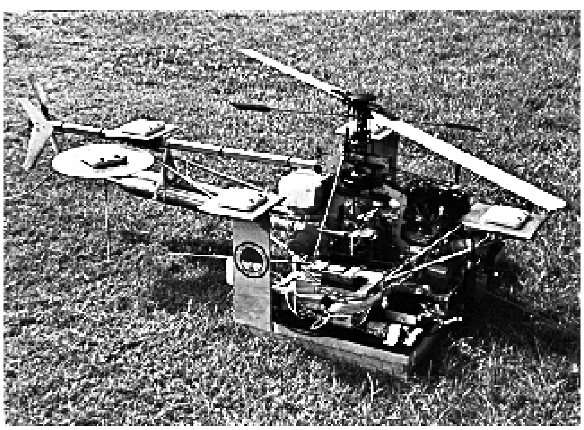

Figure 1: **Current configuration of the Stanford HUMMINGBIRD.**

A plywood plate, clamped to the bottom of the helicopter landing gear, supports the onboard electronics and their power supplies. Four GPS antennae are mounted on the helicopter—the port and starboard antennae are mounted on wooden towers at the extreme aft corners of the plywood plate, the nose antenna is supported by small aluminum tubes, and the tail antenna is mounted directly onto the helicopter tail boom. The electronics consist of a 486-based computer, a wireless ethernet modem, an RS232 radio modem receiver, two GPS receivers (Trimble TANS-Vectors), two color cameras, and one microwave wireless video transmitter. All electronics are shock mounted on closed-cell foam mounts.

Critical to the safe flight testing of the helicopter is a fully manual flight mode. An independently powered radio receiver, with computer override capability, can be used to take manual control of the helicopter at any time. Flight testing is always performed under the direct supervision of a skilled RC pilot.

Ground station equipment consists of a dual-processor Pentium computer (coupled to the wireless ethernet network), a single antenna GPS receiver, an RS232 radio modem transmitter, and a microwave wireless video receiver. The GPS receiver on the ground receives signals from up to six satellites. Carrier phase measurements are uplinked to the helicopter directly through the single direction dedicated RS232 link. No GPS processing is performed on the ground.

The helicopter receives GPS signals from three sources—the ground station (via the RS232 link), from a local “position” GPS receiver which produces raw carrier phase signals from one antenna, and from a local “attitude” GPS receiver which produces differential measurements between a master antenna and each of three slave antennae. The helicopter’s on board 486 then applies algorithms as described in Conway [1] to determine the vehicle attitude, attitude rate, position, and velocity. No real-time computation using the video information is currently done onboard the helicopter.

Before the helicopter begins flight, the planned trajectory is determined and uploaded to the onboard computer. During flight, the ground station computer receives and stores video images marked with the corresponding GPS-determined helicopter position and attitude. After landing, the ground station downloads all data taken by the onboard computer during flight and begins a post-flight analysis. The video data is then tested for validity and the final processing is performed.

The wireless ethernet communciation between the helicopter and the ground-based Pentium system allows the helicopter to utilize these off-vehicle computer resources. In the near future, it is envisioned that real-time vision algorithms could be applied to direct the control of the helicopter at a high level.

DIFFERENTIAL CARRIER PHASE GPS

Key to the success of the HUMMINGBIRD platform has been the use of GPS as a sensing technology for closed loop control of an unstable air vehicle. GPS attitude and position sensing has several distinct advantages, including:

- no moving parts
- very little calibration required
- drift-free in both position and attitude
- future promise of reduced size, power consumption and increased reliability associated with electronic parts.

A major achievement of this project is the fully autonomous flight of a helicopter using GPS as the *only* sensor for both attitude and position stabilization and control. A 10 Hz update rate for all 6 DOF is sufficient to ensure stable flight of our test vehicle. No other rate gyros, compasses, et cetera were used. For a complete description of the algorithms used in this work see Conway [1].

The ability to fly pre-programmed paths was demonstrated in the 1995 International Aerial Robotics Competition, where our helicopter was the first—and to date only—air vehicle to pick up and move a small ferromagnetic disk fully autonomously. The trajectory flown is shown in Figure 2.

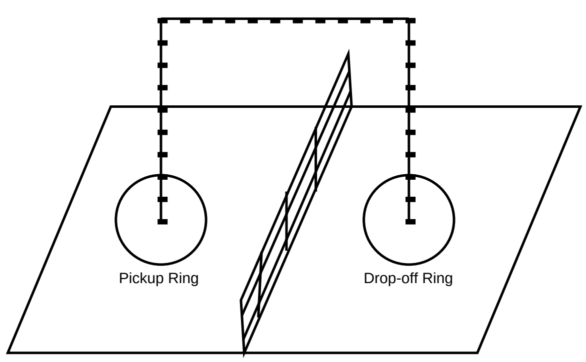

Figure 2: **Competition flight trajectory**

The helicopter autonomously retrieved a disk in the Pickup Ring, crossed the 1-m barrier at an altitude of six meters, and delivered the disk to the Drop-off Ring.

Further work has begun to extend the GPS sensing system to become an integral part of a computer vision system. The mixing of the two technologies seems appropriate since GPS provides an excellent sensor for position, while vision provides a large amount of local information about the surrounding environment but provides little information about the global position of various objects.

GPS-AUGMENTED VISION

As previously stated, DCPGPS techniques and conventional control methods are used to close an inner loop, stabilizing the helicopter's dynamics and providing a global navigation system. An outer loop is closed using feedback from the vision system to determine the precise location of the ground and other objects in the environment.

The precision available using DCPGPS is useful when it is coupled with traditional computer vision techniques. Many such techniques exist for localizing objects within a scene relative to the camera. If one can directly measure the position and orientation of a camera, one can then determine the location of sensed objects in global coordinates.

In the first set of experiments performed using the HUMMINGBIRD helicopter, the UAV was maneuvered over a grass field upon which many distinctive objects were placed. Assuming the field is level, and given the camera's height above the ground (measured with GPS), a single image is sufficient to determine the location of a target in global coordinates.

It is possible to extend this technique to provide a crude "mosaicking" capability by capturing successive images, each with corresponding DCPGPS position and attitude. By appropriate image transformations, a large image can be constructed from many, possibly overlapping, images.

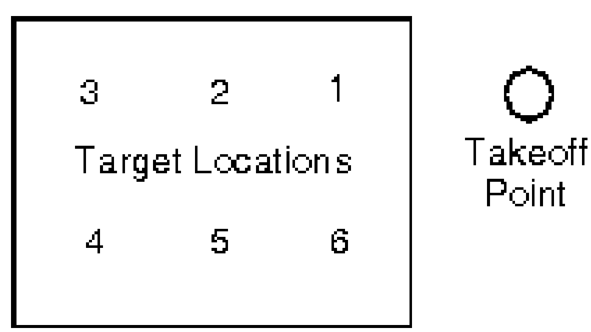

Figure 3: **Test Field**

Figure 3 shows a typical test setup for the object location algorithm. The test field measures 15 meters by 9 meters, with a known take-off point five meters outside the surveyed area. Targets, which were black plastic barrels or 4-inch metal disks painted orange, were placed at presurveyed target locations within the field.

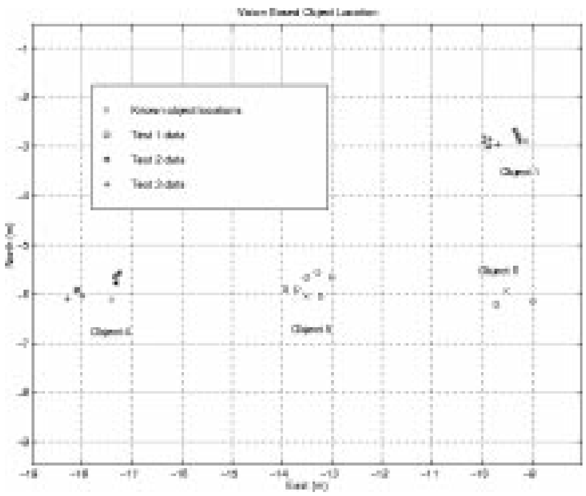

Figure 4: **GPS-Augmented Vision - Experimental Results**

Figure 4 presents the results from three test flights over the field. In each test, objects were placed on two of the six possible target locations, and HUMMINGBIRD then proceeded to survey the disk locations from an altitude of about three meters. Each data point corresponds to a single photograph of the field.

Table 1: **Vision Based Object Location**

Test Number	*Target Number*	*Position Error (m)*	
		Mean	*Standard Deviation*
1	5	0.27	0.39
1	6	0.31	0.52
2	1	0.21	0.12
2	4	0.46	0.11
3	1	0.75	0.13
3	4	0.73	0.14

As Table 1 shows, our accuracy at locating the disks is quite good. The mean location is generally within 0.50 meters, and the standard deviation usually under 15 centimeters. The data clearly indicates that the greatest error sources are constant offsets, which will be eliminated in the future with better algorithms and improved equipment calibration. Closer synchronization between the GPS attitude solutions and the camera images should also help to reduce the small amount of data scatter present.

FUTURE RESEARCH

The ARL is currently working on several additional projects to further implement vision with the HUMMINGBIRD helicopter. These include stereo imaging, motion based stereo, servoing off of camera inputs, and mobile object tracking and following.

In many computer vision applications, target objects are not located at known distances from the camera. In these situations a single image is not sufficient to determine the location of the object. HUMMINGBIRD employs a pair of stereo cameras to enable it to determine the range to target objects. This will help facilitate the retrieval of objects from on top of obstacles, or when the altitude of the target object is not known beforehand.

Conventional stereo vision is only effective out to a limited range determined by the baseline distance between cameras. Because HUMMINGBIRD is capable of determining the locations and pointing angles of its onboard cameras quite accurately using GPS, it is able to perform motion based, or long-baseline, stereo. This involves taking two images from different locations and then triangulating to determine the global location of the object. Eventually, this could provide the capability to perform accurate stereo mapping of objects kilometers away.

In the future, HUMMINGBIRD will be able to servo off of its vision system by working to keep a target object centered in its field of view and at a constant range. This will greatly facilitate object retrieval or manipulation and allow for greater resolution when performing aerial surveying tasks. HUMMINGBIRD will also be able to track and follow moving objects, improving its ability to work in dynamic environments.

References

[1] Andrew Conway. *Autonomous Control of an Unstable Helicopter Using Carrier Phase GPS Only.* PhD thesis, Stanford University, 1994.