

An Autonomous Vision-Guided Helicopter

Omead Amidi

August 1996

Department of Electrical and Computer Engineering
Carnegie Mellon University
Pittsburgh, PA 15213

*Submitted to the
Department of Electrical and Computer Engineering
in partial fulfillment of the requirements for
the degree of Doctor of Philosophy*

© 1996 Omead Amidi

This research was partly supported by SECOM security company and Yamaha Motor Company. The views and conclusions contained in **this** document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of SECOM or Yamaha Motor Company.

Keywords: helicopter, autonomous, vision-based navigation, real-time image processing.

Abstract

Helicopters are indispensable air vehicles for many applications ranging from rescue and crime fighting to inspection and surveillance. They are most effective when flown at close proximity to objects of interest while performing tasks such as delivering critical supplies, rescuing stranded individuals, or inspecting damaged buildings. These tasks require dangerous flight patterns which risk human pilot safety. **An** unmanned helicopter which operates autonomously can carry out such tasks more effectively without risking human lives. The work presented in this dissertation develops an autonomous helicopter system for such applications. The system employs on-board vision for stability and guidance relative to objects of interest in the environment.

Developing a vision-based helicopter positioning and control system is challenging for several reasons. First, helicopters are inherently unstable and capable of exhibiting high acceleration rates. They are highly sensitive to control inputs and require high frequency feedback with minimum delay for stability. For stable hovering, for example, vision-based feedback rates must be at least 30-60 Hz with no more than 1/30 second latency. Second, since helicopters rotate at high angular rates to direct main rotor thrust for translational motion, it is difficult to disambiguate rotation from translation with vision alone to estimate helicopter 3D motion. Third, helicopters have limited on-board power and payload capacity. Vision and control systems must be compact, efficient, and light weight for effective on-board integration. Finally, helicopters are extremely dangerous and present major obstacles to safe and calibrated experimentation to design and evaluate on-board systems,

This dissertation addresses these issues by developing: a “visual odometer” for helicopter position estimation, a real-time and low latency vision machine architecture to implement an on-board visual odometer machine, and an array of innovative indoor testbeds for calibrated experimentation to design, build and demonstrate an airworthy vision-guided autonomous helicopter. The odometer visually locks on to ground objects viewed by a pair of on-board cameras. Using high-speed image template matching, it estimates helicopter motion by sensing object displacements in consecutive images. The visual odometer is implemented with a custom-designed real-time and low latency vision machine which modularly integrates field rate (60 Hz) template matching processors, synchronized attitude sensing and image tagging circuitry, and image acquisition, convolution, and display hardware. The visual odometer machine along with a carrier-phase differential Global Positioning System receiver, a classical PD control system, and human augmentation and safety systems are integrated on-board a mid-sized helicopter, the Yamaha R50, for vision-guided autonomous flight.

Acknowledgments

It has been a privilege to work with my advisor, Dr. Takeo Kanade. I am thankful for his support and teaching during my thesis work. He taught me how to build real working systems by persistently following every lead and attending to every detail with critical attention.

I would like to thank my committee members, Dr. Charles Thorpe, Dr. Charles Neuman, and Dr. Lee Weiss for their advice and technical insight. I'm thankful to Dr. Charles Thorpe for his guidance on a day to day basis and his generous sharing of his group's resources for my work. In particular, I am thankful for the use the Navlab autonomous vehicle for helicopter experiments.

I would like to thank, Mark Delouis, for assisting me through the span of my graduate work. He built all on-board helicopter mechanical and safety components and developed revolutionary hardware for the indoor helicopter testbeds. He mastered the challenging task of helicopter remote control and served as my safety pilot during every experiment. I am truly in debt of his diligence, patience, and kindness.

Finally, I would like to thank Keisuke Fujita and Yuji Mesaki of SECOM security company who supported my work during their two year stay at CMU Robotics Institute. They assisted me in the development of the on-board global positioning and image processing hardware.

Omead Amidi
August 9, 1996

Table of Contents

I. Chapter 1. Introduction	1
1.1 Challenges of Vision-Based Helicopter Flight	2
1.2 Contributions	3
1.3 Related Work	4
1.3.1 Helicopter Control	4
1.3.2 Controlling with Vision	5
1.3.3 Autonomous Systems	6
1.4 Approach	7
1.4.1 Vision-Based Position Estimation	7
1.4.2 Real-Time and Low Latency Vision	8
1.4.3 Experimental Testbeds	9
1.5 Dissertation Overview	10

II. Chapter 2. Vision-Based Helicopter Positioning for Autonomous Control	13
2.1 Positioning with a Visual Odometer	15
2.2 Definition of Coordinate Frames and Transformations	18
2.2.1 Helicopter and Local Ground Coordinate Frames and Transformations	18
2.2.2 On-board Camera Setup and Coordinate Frames	20
2.2.3 Camera Image Coordinate Frame and Transformation	21
2.3 Visual Odometer Tracking Algorithm	22
2.3.1 Overview	22
2.3.2 Position Sensing	25
2.3.3 Velocity Sensing	27
2.4 Visual Odometer Image Processing	29
2.4.1 Image Pixel Coordinates of the Target Template	29
2.4.2 Range Measurement	32
2.4.2.1 Image Center Range Estimation	32
2.4.2.2 Range Interpolation	34
2.4.3 Pixel Velocity at the Image Center	36

2.4.4 Template Matching	37
2.4.4.1 Matching Criteria	37
2.4.4.2 Coarse to Fine Search	38
2.4.4.3 Subpixel Interpolation of Matching Position	39
2.5 Position Estimation Experiments	40
2.5.1 Indoor testbed	40
2.5.2 Position Estimation Results	42
2.5.3 Velocity Estimation Results	46
2.6 Summary and Discussion	48
III. Chapter 3. A Real-Time and Low Latency Visual Odometer Machine	49
3.1 Visual Odometer Machine Specifications	50
3.2 Visual Odometer Architecture	52
3.2.1 Decentralized Communication	52
3.2.2 Modular Architecture	53
3.3 Components of the Visual Odometer Machine	55
3.3.1 Image Acquisition	55
3.3.2 Image Convolution	58
3.3.3 DSP Processing Module	59
3.3.4 Module Synchronization	60
3.3.5 Sensor Bridge Module	60
3.3.6 Image Display	60
3.4 Data Flow and Synchronization	62
3.4.1 Image Acquisition and Preprocessing	64
3.4.2 Image Transfer and Storage	64
3.4.3 Target Template Position Estimation Processing	65
3.4.4 Stereo Processing	66
3.4.5 Pixel Velocity Processing	67
3.4.6 Attitude Compensation	67
3.5 Summary and Discussion	67

IV. Chapter 4. Design and Evaluation of an On-Board Vision-Based Helicopter Control System	69
4.1 Yamaha R50	70
4.2 Indoor R50 Testbed	
4.2.2 Testbed R50 Helicopter	73
4.2.2.1 On-board Power	73
4.2.2.2 Attitude Sensors	73
4.2.2.3 Testbed Cameras, Scenery, and Lighting	75
4.2.2.4 On-board Computing	76
4.2.2.5 On-board Chassis Mounting	76
4.2.3 Testbed Safety	77
4.3 Helicopter Controller	78
4.3.1 Helicopter Control Inputs	79
4.3.2 PD Servo Control Loops	79
4.3.3 Controller Testing	81
4.4 Vision and Attitude Synchronization	83
4.5 Summary	85

V. Chapter 5. Outdoor Autonomous Flight	87
5.1 Secondary Positioning System	88
5.1.1 GPS Positioning Method	89
5.1.2 GPS Positioning Quality	90
5.1.3 GPS Evaluation Experiments	91
5.2 On-board Integrated System	94
5.2.1 Vision Processor	96
5.2.2 GPS Receiver	96
5.2.3 Compass	96
5.2.4 Laser Rangefinder	97
5.2.5 Real-Time Controller	97
5.2.6 Actuator Control	98
5.2.7 Safety Circuit and Human Interfaces	99
5.3 Autonomous Helicopter	100
5.3.1 Weight and Power	101

5.3.2 On-board Computing	101
5.4 Outdoor Flight Experiments	104
5.4.1 Experimental Setup	104
5.4.2 Experimental Approach	106
5.4.3 Experimental Results	110
5.4.3.1 Position Estimation	110
5.4.3.2 Velocity Estimation	115
5.4.3.3 Computer Control Trials	117
5.5 Summary	120
VI. Chapter 6. Conclusions and Future Work	121
6.1 Accomplishments	122
6.1.1 An Autonomous Vision-Guided Helicopter	122
6.1.2 Real-time and Low Latency Vision	124
6.2 Future Work	126
6.2.1 Object Tracking	126
6.2.2 Helicopter Positioning using Known Environments	127
6.3 Concluding Remarks	130
VII. Appendix A. Six Degree-Of-Freedom Testbed	131
VIII. Appendix B. Lessons Learned	137
IX. Bibliography	139

Chapter 1.

Introduction

Precise maneuverability of helicopters makes them useful for many critical tasks ranging from rescue and security to inspection and monitoring operations. Helicopters are indispensable air vehicles for finding and rescuing stranded individuals or transporting accident victims. Police departments use them to find and pursue criminals. Fire fighters use helicopters for precise delivery of fire extinguishing chemicals to forest fires. More and more electric power companies are using helicopters to inspect towers and transmission lines for corrosion and other defects and to subsequently make repairs. All of these applications demand dangerous close proximity flight patterns, risking human pilot safety. An unmanned autonomous helicopter will eliminate such risks and will increase the helicopter's effectiveness.

Typical missions of autonomous helicopters require flying at low speeds to follow a path or to hovering near an object of interest. Accurate position estimation of the helicopter relative to objects is necessary to perform such tasks. In general, such positioning equipment as inertial navigation systems or global positioning systems are well suited for long range, low precision helicopter flight and fall short for very precise, close proximity flight. Moreover these sensors estimate absolute position and cannot sense position in relation to task objects of interest. Visual sensing is the richest source of data for this relative position estimation.

The work presented in this dissertation demonstrates stable helicopter control based primarily on visual feedback. An “eye-in-the-sky” robot helicopter is developed which can perform missions outdoors while flying autonomously. The helicopter can fly precisely and at close proximity to ground objects by maintaining its relative location by on-board vision.

1.1 Challenges of Vision-Based Helicopter Flight

Helicopters are inherently unstable and require constant compensation for stable flight. The effectiveness of an autonomous helicopter is critically dependent on its accurate and stable positioning relative to objects in the environment. Estimating this relative position by on-board vision, a 3D object tracking problem, is difficult for several reasons.

- Helicopters can move quickly. Small and mid-sized helicopters can accelerate in the range of 0.5-0.7 g and can exhibit 40-60 degrees per second angular velocity under normal operating conditions. To keep up with the helicopter’s high degree of maneuverability, an on-board vision system must sample and process camera images at high frequency. On-board image processing must be performed at frame rate (30 Hz) or higher for effective vision-based object tracking. Higher rate image sampling also simplifies the tracking problem by limiting object displacements in successive images.
- Helicopters are highly sensitive to control inputs. Feedback latency is critical to stable helicopter flight. High throughput of image processing alone is not sufficient. Object tracking must be performed with minimum latency to provide adequate and timely feedback for stability. Small model helicopters require system latencies of no more than 1/30 to 1/60 seconds for stability.
- Helicopters move with typically significant attitude variations. A helicopter can bank 30 degrees as it transitions to forward flight. To maintain relative position, on-board vision must distinguish helicopter translation from rotation. Distinguishing rotation from translation in images under perspective projection can be difficult since small attitude variations can look virtually indistinguishable from small translational motion. This effect is exaggerated for the helicopter application since

tracked objects are frequently small relative to the helicopter altitude and cannot provide sufficient 3D clues for distinguishing rotation from translation.

- Helicopters have strictly limited payloads and available power. A vision system capable of meeting the above criteria must also be compact and efficient for practical on-board integration. Small (< 200 lbs) helicopter payloads range from 5-40 pounds.
- Helicopters are dangerous. The spinning rotor blades pose an immediate danger to nearby individuals. The responsive nature of helicopters makes them prone to out-of-control flight or crashes during experiments.

The research presented in this dissertation addresses these challenges by developing an autonomous vision-guided helicopter control system.

1.2 Contributions

The three contributions of this dissertation are:

1. The first autonomous robot helicopter stabilized and guided by an on-board “visual odometer” for position estimation: The odometer visually locks on to ground objects and maintains helicopter position at field rate (60 Hz) during flight. The helicopter integrates the visual odometer with sensors such as gyroscopes and a global positioning system (**GPS**) receiver, control and actuation, as well as safety and human augmentation systems.
2. A new vision machine architecture for real-time and low latency image processing: The architecture balances computational power and data bandwidth requirements to realize vision machines tailored to the applications at hand. Based on this architecture, a visual odometer machine is designed and realized on-board an autonomous helicopter.
3. Innovative testbeds for effective indoor experimentation with helicopters: Most significant is an indoor six-degree-of-freedom testbed built with light-weight composite material to support an electrical model helicopter. The testbed provides safety by preventing helicopter crashes and measures helicopter ground-truth position during flight for calibrated experiments.

1.3 Related Work

Building an autonomous helicopter system requires research in helicopter control as well as in helicopter position sensing. While this dissertation focuses on the position sensing aspect of the problem, it is important to recognize and employ existing work on helicopter control, visual servoing, and autonomous robotic systems to realize a working vision-guided robot helicopter.

1.3.1 Helicopter Control

The study of the helicopter control problem is not new. Helicopter dynamic modeling is well documented in the literature. In particular, Prouty [1] and Johnson [2] present excellent comprehensive studies of helicopter aerodynamical models and stability analysis. Overcoming the inherent instability of helicopters has been the focus of a large body of research, including detailed mathematical models (e.g., [3]) for control and Kalman filtering of multiple sensor data for state estimation (e.g., [4]). The controller design methods range from linear quadratic (LQ) design to H infinity design [5], [6], and [7] and predictive control [8]. For example, a stable closed loop control system has been formulated [4] by quadratic synthesis techniques for helicopter autolandings.

Incorporation of a human pilot model has been attempted based on quadratic optimal Cooperative Control Synthesis [9]. This model is used for control augmentation where the control system cooperates with the pilot to increase aircraft performance. The sophisticated pilot model developed by [10] attempts to describe the human's ability to look ahead. This ability is crucial to precise low-altitude helicopter control. While it is difficult to identify and verify these models, they provide a valuable basis for an intelligent helicopter controller, especially in the design of low-level control loops.

Manned flight tests of helicopter controllers have also been conducted. Notable implemented systems include those at NASA Ames Research Center [4], NASA Langley Research Center [9], and military aircraft manufacturers [11]. Fuzzy controllers have been successfully employed for helicopter flight experiments. In Japan, Sugeno's group at Tokyo Institute of Technology [12] has demonstrated helicopter control using fuzzy logic.

1.3.2 Controlling with Vision

The positioning feedback for the above helicopter control experiments is primarily provided by on-board INS/GPS or ground-based beacon systems instead of on-board computer vision. The computational complexity and the high data bandwidth requirements of vision have been major obstacles to practical and robust vision-based positioning and control systems. In spite of these drawbacks, promising results have been recently demonstrated in real-time vision processing, visual servoing of robotic manipulators, and accurate vision-based position estimation systems.

The development of low-cost special-purpose image correlation chips and multi-processor architectures capable of high communication rates has made a great impact on image processing. Examples of vision systems built from special-purpose hardware include transputer-based image hardware for two-dimensional object tracking [13], and real-time tracking and depth map generation using correlation chips [14].

High speed processors are making visual feedback increasingly more practical in many control applications. There has been significant development in visual control of manipulators carrying small cameras or eye-in-hand configuration. Researchers at Carnegie Mellon University's Robotics Institute have demonstrated real-time visual tracking of arbitrary 3D objects traveling at unknown 2D velocities using a direct-drive manipulator arm [15]. The Yale spatial robot juggler [16] has demonstrated transputer-based stereo vision for locating juggling balls in real time. Real-time tracking and interception of objects using a manipulator [17] have also been demonstrated based on fusion of the visual feedback and acoustic sensing.

RAPiD and DROID [18], developed by Roke Manor Research Limited, are systems designed for vision-based position estimation in unknown environments. RAPiD is a model-based tracker capable of extracting the position and orientation of known objects in the scene. DROID is a feature-based system which uses the structure-from-motion principle for extracting scene structure using image sequences. Real-time implementations of these systems have been demonstrated using dedicated hardware.

1.3.3 Autonomous Systems

Integrating efficient model-based and learning techniques with powerful hardware architectures has produced an array of autonomous land and air vehicles. Significant advances in autonomous automobiles have demonstrated vision-based control at highway speeds. Most notable are Carnegie Mellon's autonomous ground vehicle projects (Navlab [19][20], Automated Highway System [21], Unmanned Ground Vehicle project [22][23]) and the work of Dickmanns involved with the European PROMETHEUS project [24] at the University of Bundeswer, Munich.

Dickmanns applies an approach exploiting spatio-temporal models of objects in the world to control autonomous land and air vehicles [25]. He has demonstrated autonomous position estimation for an aircraft in landing approach using a video camera, inertial gyros and an air velocity meter. Vision-based state estimation is also pursued at NASA Ames Research Center [26][27][28] using parallel implementation of multi-sensor range estimation for helicopter flight.

An aerial robot competition sponsored by Association for Unmanned Vehicle Systems, described in [29], has recently encouraged the development of a number of small model vertical takeoff autonomous robots. The competition task requires flying robots to autonomously carry small objects from one location to another. Recently, tasks requiring on-board vision, such as object identification, are being added to the competition requirements.

Most notable competitors are teams from Stanford University, University of Southern California (USC) [30] and [31], Georgia Institute of Technology [32], and University of Texas at Arlington (UTA)[33]. Researchers at Stanford are concentrating on carrier-based **GPS** technology for helicopter position and attitude estimation. The USC team is approaching the helicopter control problem based on a behavioral paradigm and low complexity vision to aid in helicopter navigation. The Georgia Tech robot supports on-board sensors and flight control systems. Mission planning and vision tracking are performed off-board using a ground station. UTA researchers have developed a vertical take-off aircraft which uses a responsive rigid propeller instead of the traditional articulated helicopter rotor blade designs. They have integrated control, navigation, and communication systems on-board their aircraft for autonomous operation.

1.4 Approach

The main goal of this dissertation is the development of an airworthy autonomous helicopter system which employs vision as its primary source of guidance and control. This goal is pursued by developing: a high-level position estimation algorithm through a visual odometer, a real-time and low latency vision machine architecture for on-board system implementation, and an array of experimental test-beds to incrementally design and evaluate each system component.

1.4.1 Vision-Based Position Estimation

A visual odometer locks on to and tracks feature rich objects to sense helicopter motion. The odometer maintains this visual lock using high speed image correlators which estimate helicopter range and motion relative to the ground objects. The odometer closely integrates on-board attitude sensors with the image correlators to resolve 3D helicopter translation which is key to accurate helicopter control.

The visual odometer implements an object tracking algorithm. Viewing the ground through a pair of on-board cameras, the algorithm locks on and tracks feature-rich objects appearing in image windows or templates. As shown in Figure 1-1, the algorithm initially locks on to objects appearing at the image center and maintains this lock while the objects are in the field of view. As the objects leave the image, the algorithm selects another image template to lock on to and continues positioning the helicopter. Image templates are tracked by high-speed image correlation or template matching. For full 3D motion estimation, the algorithm matches templates in both camera images simultaneously for stereo range detection, and in successive images, for velocity estimation.

Since helicopter attitude and height variations can significantly affect the appearance of tracked objects in successive images, the algorithm must actively update image templates for robust tracking. Furthermore, the algorithm must sense helicopter translation for accurate control which requires eliminating the effects of rotation on image displacements. The algorithm accomplishes these difficult tasks by tracking multiple templates and by measuring helicopter attitude with on-board angular sensors. The relative motion of two tracked templates in images determines height and heading changes




Figure 1-1 Vision-based Positioning

which are used for scaling or rotating tracked templates for consistent matches. The effects of helicopter roll and pitch variations are determined by tagging each image with helicopter attitude during camera shutter exposure interval. Using this synchronized attitude data, the algorithm estimates the effects of rotation on image displacements based on camera lens parameters. This dissertation develops custom-designed hardware to filter and tag the camera images for the tracking algorithm.

1.4.2 Real-Time and Low Latency Vision

Controlling a highly unstable plant such as a small helicopter requires frequent state feedback with minimum latency. Providing this feedback with vision at suitably high rates and with small delay can be very challenging. This is especially true for computationally complex tasks such as image correlation or template matching. Processing rates of 30-60 Hz with 1/30 second latency are experimentally determined to be sufficient for stable model helicopter control.

In spite of the growing commercial development of high-speed vision systems, many are designed for high image processing throughput with little regard to processing latency. Powerful general-purpose vision systems capable of low latency processing are too bulky and expensive for on-board integration. Furthermore, most commercial vision systems are incapable of precisely synchronizing external sensor data acquisition within the image processing pipeline. This capability is espe-

cially important to vision-based helicopter control which requires synchronized data acquisition from a variety of on-board sensors. These factors motivated the development of a new real-time and low latency image processing architecture presented in this dissertation.

The architecture's design incorporates processing capabilities modularly. A uniform communication scheme keeps all processing and sensing components compatible. The system's processing pipeline can be easily expanded horizontally, to incorporate more processing capabilities, or vertically, to increase data bandwidth capacity. Images flow through the system via a network of high-speed point-to-point communication links with consistent latency making up a 2D pipelined machine architecture. The links interconnect modules, such as image convolvers or template matchers, which also have predictable performance and latency matched with other modules and the helicopter control system.

1.4.3 Experimental Testbeds

Building an autonomous robot helicopter requires careful and calibrated experimentation with real helicopters. This is difficult because helicopters are unstable and typically exhibit undesirable oscillatory behavior without active and frequent control compensation. There is also the danger of the helicopter fuselage or the spinning rotor blades crashing into nearby individuals and causing major harm during experiments.

Repeated failures due to faulty sensors, computer algorithms, or helicopter mechanics are likely during development, and careful safety measures must be in place to protect the researchers and the experimental aircraft. Outfitting a helicopter for outdoor free flight from the start can be risky since each unavoidable malfunction is a major loss of time, effort, and resources. With these concerns, this thesis develops a series of innovative testbeds for safe indoor experimentation.

Each testbed supports an electrical or a gas powered model helicopter for inexpensive and logistically manageable experiments indoors. Model helicopters are faithful reproductions of full size helicopters with respect to the crucial rotor controls and the techniques developed to control them directly apply to larger scale helicopters. In most cases, the smaller models are more agile and usually more

difficult to control than the larger and less responsive helicopters. The testbeds limit allowable helicopter travel area and attainable velocity using mechanical links and dampers, therefore reducing the risk of mechanical failure from hard landings and violent flight patterns. Minimizing the effects of these mechanical linkages on helicopter dynamics is the main challenge in designing such testbeds.

A chronological progression of indoor testbeds to outdoor flight is shown in Figure 1-2. Experiments with small model helicopters were started with an attitude control testbed (a) limiting helicopter motion to only one rotational axis and continued with a six-degree-of-freedom testbed (b) which allowed full helicopter motion in a semi-spherical area. The experiments with the small model helicopters were expanded to another indoor testbed (c) housing a mid-sized helicopter, the Yamaha R50, which is also employed for the outdoor prototype autonomous helicopter (d).

Beyond addressing the safety issues, the indoor testbeds are designed to provide several other important capabilities. Each testbed is outfitted with non-intrusive sensors for accurate measurement of helicopter ground-truth position and attitude for quantitative performance evaluation. With this ground-truth feedback, experiments on key components such as position estimation algorithms or sensor platforms can be performed independently allowing precise analysis of different components under controlled conditions.

1.5 Dissertation Overview

The dissertation is divided into six chapters. Chapter 2 provides a high-level view of the dissertation's vision-based position estimation by describing a visual odometer positioning algorithm. The chapter describes the issues involved in helicopter positioning and the techniques employed to address these issues.

Building on vision-based techniques, Chapter 3 addresses design issues in building vision systems for high speed applications such as helicopter control. The chapter introduces a reconfigurable vision machine architecture designed for low latency and real-time image processing and describes the architecture by analyzing a visual odometer machine developed for on-board helicopter position estimation.




Figure 1-2 Experimental Approach

Chapter **4** and Chapter 5 present indoor and outdoor experiments to build a prototype vision-guided autonomous helicopter. Chapter **4** presents indoor design and evaluation of an autonomous helicopter system. The chapter presents the development of an indoor testbed for the Yamaha R50 helicopter which is employed to verify a PD-based helicopter controller and to prove the effectiveness of the visual odometer machine in positioning the helicopter. Chapter 5 presents the outdoor flight experiments using a fully integrated autonomous helicopter supporting on-board vision, GPS, and real-time control systems. Results **of** outdoor free flight tests are presented and vision-based helicopter positioning and control performance are evaluated using accurate carrier-phase **GPS** receivers.

Finally, Chapter **6** presents conclusions and presents preliminary experiments to probe future research directions of the work presented in this dissertation. The future research directions presented include vision-based object tracking, helicopter positioning in known environments, and vision machine designs for other applications such as factory inspection and industrial robotics.

Chapter 2.

Vision-Based Helicopter Positioning for Autonomous Control

The control of an autonomous helicopter is only as good as its positioning. Control response and accuracy is, in essence, dictated by how frequently and how promptly the helicopter's position is determined during flight. Positioning the helicopter can be either global or relative depending on the task at hand. Global positioning is necessary for long distance flight where the helicopter must reach a predetermined destination. Relative positioning, on the other hand, is necessary for precise flight in relation to objects of interest in the environment. Vision is particularly well-suited for this type of relative positioning and is the main focus of this chapter.

On-board vision can estimate helicopter motion by tracking stationary objects in the surrounding environment. Objects are displaced in consecutive images as the helicopter moves, and this displacement can be accurately measured by image processing to determine helicopter motion.

A key concern is the trackability of objects in the field of view. Tracking is possible only if visible objects possess distinguishing features which can be consistently identified in image sequences. Highly contrasting and randomly textured scenery is common in outdoor environments and can provide feature-rich imagery for vision-based motion sensing. On-board vision can take advantage of the abundant natural features to “lock” on to arbitrary objects and track them to sense helicopter motion.

It is difficult, however, to sense helicopter translation, which is essential for autonomous control, with vision alone since image displacements also occur with helicopter rotation. Distinguishing between rotation and translation in a sequence of images under perspective projection is extremely difficult. For instance, helicopter rolling motion can appear very similar to lateral translation in consecutive images. This ambiguity can be resolved by accurately fusing data from angular sensors with image displacement measures. New generations of light-weight gyroscopes and angular rate sensors available today can provide reliable measurement of angular change between images to isolate rotational effects on image displacement.

This chapter sets the stage for the dissertation by presenting the algorithm for vision-based helicopter position tracking developed for autonomous flight. The algorithm maintains helicopter position and altitude by capturing images from a pair of ground-pointing cameras and visually locking on to ground objects. The algorithm is built upon fast template matching engines and is implemented by a reconfigurable vision machine on-board a prototype autonomous helicopter. The vision machine and the prototype helicopter are presented in Chapter 3 and Chapter 4.

This chapter begins by outlining the vision-based position estimation approach which implements a visual odometer for tracking helicopter position. Subsequent sections analyze the visual odometer by defining system coordinate frames and transformations and describing the odometer's position tracking algorithm. The tracking algorithm is further analyzed by discussing each of its image processing components, including position and velocity trackers and stereo image processing. The chapter concludes by presenting experimental results of indoor helicopter test flights demonstrating algorithm performance.

2.1 Positioning with a Visual Odometer

A major contribution of the work presented by this dissertation is a visual odometer which tracks helicopter position relative to an initial known location with on-board vision. The odometer determines the position of objects appearing in the camera field of view relative to the initial helicopter location, and thereafter tracks the objects visually to maintain helicopter position. As the helicopter moves, older objects may leave field of view, but new objects entering the scene are localized to continue tracking helicopter position.

The visual odometer relies on a “target” template initially taken from the center of an on-board camera image. The location of the object appearing in the target template is determined by the sensing of camera range to the object using the current helicopter position and attitude. With the sensed object location, helicopter position is updated as the odometer tracks the object in the images with template matching.

Template matching between consecutive images provides lateral and longitudinal image displacement which may result from both helicopter translation and rotation. Template matching in two images, taken simultaneously by a stereo pair of cameras, measures helicopter range. Three dimensional helicopter motion is then estimated by combining the lateral and longitudinal image displacements and range estimates with helicopter attitude, measured by on-board angular sensors.

Several important observations are in order regarding this position tracking approach:

Effects of Rotation: Helicopter translation is a direct result of its change in attitude, often causing large image displacement. Figure 2-1 shows the significance of this effect while the helicopter flares for reducing forward speed or stopping. The effects of rotation must be eliminated from the measured image displacement to determine the change in helicopter position. The visual odometer determines these effects by precisely measuring the variation in helicopter attitude between images. This correction is only valid provided that attitude data is captured in precise synchronization with the camera shutter opening.




Figure 2-1 Effects of Rotation

Range Estimation: The tracked objects must be visible by both stereo cameras for range estimation. The odometer must guarantee range measurement during all anticipated helicopter flight maneuvers for robust position estimation. Assuming locally flat ground, the odometer estimates template range using current helicopter attitude, measured by angular sensors, and one reference range point at the center of the image. Although not ideal, this approach simplifies object range estimation and allows easy integration of other range sensors, such as a laser rangefinder, to the system.

Template Matching Accuracy: Helicopters can move rapidly relative to tracked objects. As a result, the template matching process must be consistent and robust to accommodate for the quick rotation and distance variations. Accurate template matching to retain visual "lock" on objects requires anticipating how the appearance of the objects changes in future images before performing the matching operation. As shown in Figure 2-2, objects at the same ground location can significantly change in size and orientation in the field of view as the helicopter turns and varies its altitude. Failing to compensate for these variations will result in poor matches which reduce the quality of helicopter position updates.




Figure 2-2 Changing Template Appearance with Rotation and Distance

For accurate and consistent matches, templates must be rotated, scaled, and normalized in intensity. The visual odometer determines incremental template rotation and scaling factors by tracking multiple templates concurrently. By tracking an auxiliary template in parallel with the primary or main template the odometer estimates the effects of rotation and height variations directly from the image. Observing the direction and magnitude of a vector connecting the centers of the templates determines the rotation angle and the scale factor necessary to prepare templates for subsequent accurate matches.

Template Matching Speed: The computationally complex matching operations must be performed frequently to ensure accurate template matching and to provide sufficient rate of feedback for helicopter stabilization. Searching for a matching position in the entire image is not always necessary. To reduce computational requirements, the odometer searches for templates in a small window surrounding the previous match. The search window size is chosen based on the matching frequency, helicopter proximity to the objects appearing in the template, and anticipated helicopter movement based on the previous match displacement.

It is worth noting that the search area can be reduced with increasing processing frequency and that high processing frequency may be achievable only if the search area is smaller. Therefore, it is beneficial to perform the matching operation as fast as possible limited only by the camera image acquisition frequency.

2.2 Definition of Coordinate Frames and Transformations

This section defines a number of coordinate frames and transformations necessary to analyze the visual odometer. To track helicopter position using observed image features, coordinate frames must be defined for the helicopter environment, the helicopter body, the on-board cameras, and the on-board camera images along with their respective coordinate transforms.

A local ground frame is aligned with the earth's magnetic North, determined by a magnetic reference compass, and horizontally leveled using the gravity vector, measured by inclinometers. The helicopter's center of mass is chosen as the origin of the helicopter body coordinate frame and each camera's focal point is chosen to be the origin of its camera coordinate frame. Finally, the camera image coordinates are defined by the 2D image pixel coordinates and the camera range of the objects appearing at the pixel coordinates.

2.2.1 Helicopter and Local Ground Coordinate Frames and Transformations

The local ground coordinate frame is the principal reference for helicopter position tracking. It is local in the sense that its origin is at an arbitrary location in a bounded and level indoor or outdoor area. As shown in Figure 2-3, the local ground frame's x and y axes are directed towards East and North and the z axis is pointed away and orthogonal to the horizontal plane. References to the "ground frame" in this dissertation are to this local ground coordinate system.

The origin of the helicopter body coordinate frame is chosen at the helicopter's center of mass which is along the main rotor axis. The body frame's x , y , and z axes point forward, left, and upward, respectively, as shown in Figure 2-3.




Figure 2-3 Helicopter and Ground Frames and Transformations

Let $P^H = [x_p^H \ y_p^H \ z_p^H]^T$ denote the location of a point P in the helicopter coordinate frame. The point's location in the ground coordinate frame $P^G = [x_p^G \ y_p^G \ z_p^G]^T$ is:

$$P^G = R_H^G P^H + T_H^G \quad (2-1)$$

where R_H^G is the (3×3) Euler rotation matrix [34] which is a function of the three Euler angles ϕ , θ , and ψ in (2-2) and $T_H^G = [x_h^G \ y_h^G \ z_h^G]^T$ is the helicopter's coordinate in the ground frame:

$$R(\phi, \theta, \psi) = \begin{bmatrix} \cos\theta\cos\psi & \sin\phi\sin\theta\cos\psi & \cos\phi\sin\theta\cos\psi \\ & -\cos\phi\sin\psi & +\sin\phi\sin\psi \\ \cos\theta\sin\psi & \sin\phi\sin\theta\sin\psi & \cos\phi\sin\theta\sin\psi \\ & +\cos\phi\cos\psi & -\sin\phi\cos\psi \\ -\sin\theta & \sin\phi\cos\theta & \cos\phi\cos\theta \end{bmatrix} \quad (2-2)$$

2.2.2 On-board Camera Setup and Coordinate Frames

The two on-board cameras are mounted side by side and approximately parallel to the x axis of the helicopter body frame as shown in Figure 2-4. The front camera is chosen to be the “main” camera providing images for lateral and longitudinal image displacement measurement and stereo matching. The rear camera is used for stereo matching of main camera image templates. The origin of the camera coordinate frame is chosen at the focal point of the main camera and the axes are directed as shown in Figure 2-4. The two cameras are accurately aligned so that the main camera x axis passes through the rear camera image center horizontally, dividing it into two equal rectangles.




Figure 2-4 Camera to Helicopter Transformation

Since the cameras are affixed to the helicopter, the main camera position and attitude do not vary in the helicopter body frame and are represented by the Euler rotation matrix R_C^H and translation vector T_C^H . The rotation matrix, R^H , is a function of the three camera Euler angles, $(\phi_c^H, \theta_c^H, \psi_c^H)$, in the helicopter frame and the translation vector, T_C^H , is the coordinates of the main camera in the heli-

copter frame, (x_c^H, y_c^H, z_c^H) . The location of a point P in the camera frame, $P^C = [x_p^C \ y_p^C \ z_p^C]^T$, can be transformed to the helicopter frame according to:

$$P^H = R_C^H P^C + T_C^H \quad (2-3)$$

2.2.3 Camera Image Coordinate Frame and Transformation

The camera image coordinate frame is defined as a combination of the 2-D pixel coordinates of the main camera image, (x^{im}, y^{im}) , and the camera range, z^{im} , of objects appearing at those pixel coordinates as shown in Figure 2-5.




Figure 2-5 Image Coordinate System

The location of a point P in the camera image frame, $P^{im} = [x_p^{im} \ y_p^{im} \ z_p^{im}]^T$, can be transformed to the camera frame, P^C , according to:

$$P^C = \begin{bmatrix} x_P^C \\ y_P^C \\ z_P^C \end{bmatrix} = z_P^{im} \begin{bmatrix} \frac{x_P^{im}}{f} \\ \frac{y_P^{im}}{f} \\ 1 \end{bmatrix} \quad (2-4)$$

where f is the focal length of the camera lens.

2.3 Visual Odometer Tracking Algorithm

The visual odometer maintains helicopter position by a tracking algorithm. The algorithm senses image displacement in subsequent images by template matching and determines helicopter motion between images to incrementally update helicopter position and velocity. This section presents the tracking algorithm of the visual odometer, starting with a high level discussion of the algorithm followed by a presentation of its underlying components.

2.3.1 Overview

The algorithm updates helicopter position by locking onto objects that initially appear at the main camera image center and tracking them in subsequent images. For simplicity, the tracked objects will be referred to as “the target” hereafter. “Locking on” refers to the algorithm’s instantaneous sensing of the target’s ground frame position and its subsequent tracking by on-board vision.

Relying on the target’s ground location, the algorithm continuously senses the target’s location in the helicopter frame to estimate helicopter position. The algorithm computes the helicopter’s ground frame position using the two sets of target coordinates, in the ground and helicopter frames, together with current helicopter attitude, measured by on-board angular sensors. Therefore, the algorithm must first sense the target’s ground position and then track it while the target is in the field of view.

Since it is unclear if a tracked target is about to leave the field of view before each processing cycle, the algorithm must constantly maintain a potential new target replacement while it tracks the current target. To accomplish this, the algorithm tracks the current target and prepares a new target concurrently using two threads of execution as shown by the algorithm's high-level flow chart in Figure 2-6. The two execution threads are bootstrapped by an arbitrary initial two dimensional (x,y) helicopter position and commence by capturing camera images and current helicopter attitude from on-board cameras and angular sensors.

If a target is currently available and localized in the ground frame, the primary thread, labeled in Figure 2-6, senses the current target's image coordinates by image processing. The primary thread then transforms the target's image coordinates to the helicopter frame to compute the helicopter's position.

While the primary thread is estimating the helicopter's position, the secondary thread maintains new potential targets. This thread captures a new target from the main camera image center and estimates its position in the ground frame by first estimating its image coordinates by image processing, followed by coordinate transformations once the current helicopter position is determined by the primary thread.

Once the primary thread has estimated current helicopter position, it decides to keep the current target or discard it for a new one from the secondary thread. For instance, the primary thread may discard the current target if it is near the image border. How close the target can travel to the image border is determined by current altitude and anticipated helicopter motion from image to image. This topic is discussed in Section 2.4.1 in detail.

In addition to preparing new potential templates, the secondary thread also estimates helicopter velocity while the helicopter is being localized by the primary thread. Searching for the previous potential target in the current image, the secondary thread detects target displacement between images which it then transforms to the ground frame for velocity estimation. The sensed velocity has lower latency than velocity derived from differentiating the estimated position. Lower latency velocity is beneficial for helicopter stabilization and provides another source of data for redundant helicopter motion estimation.

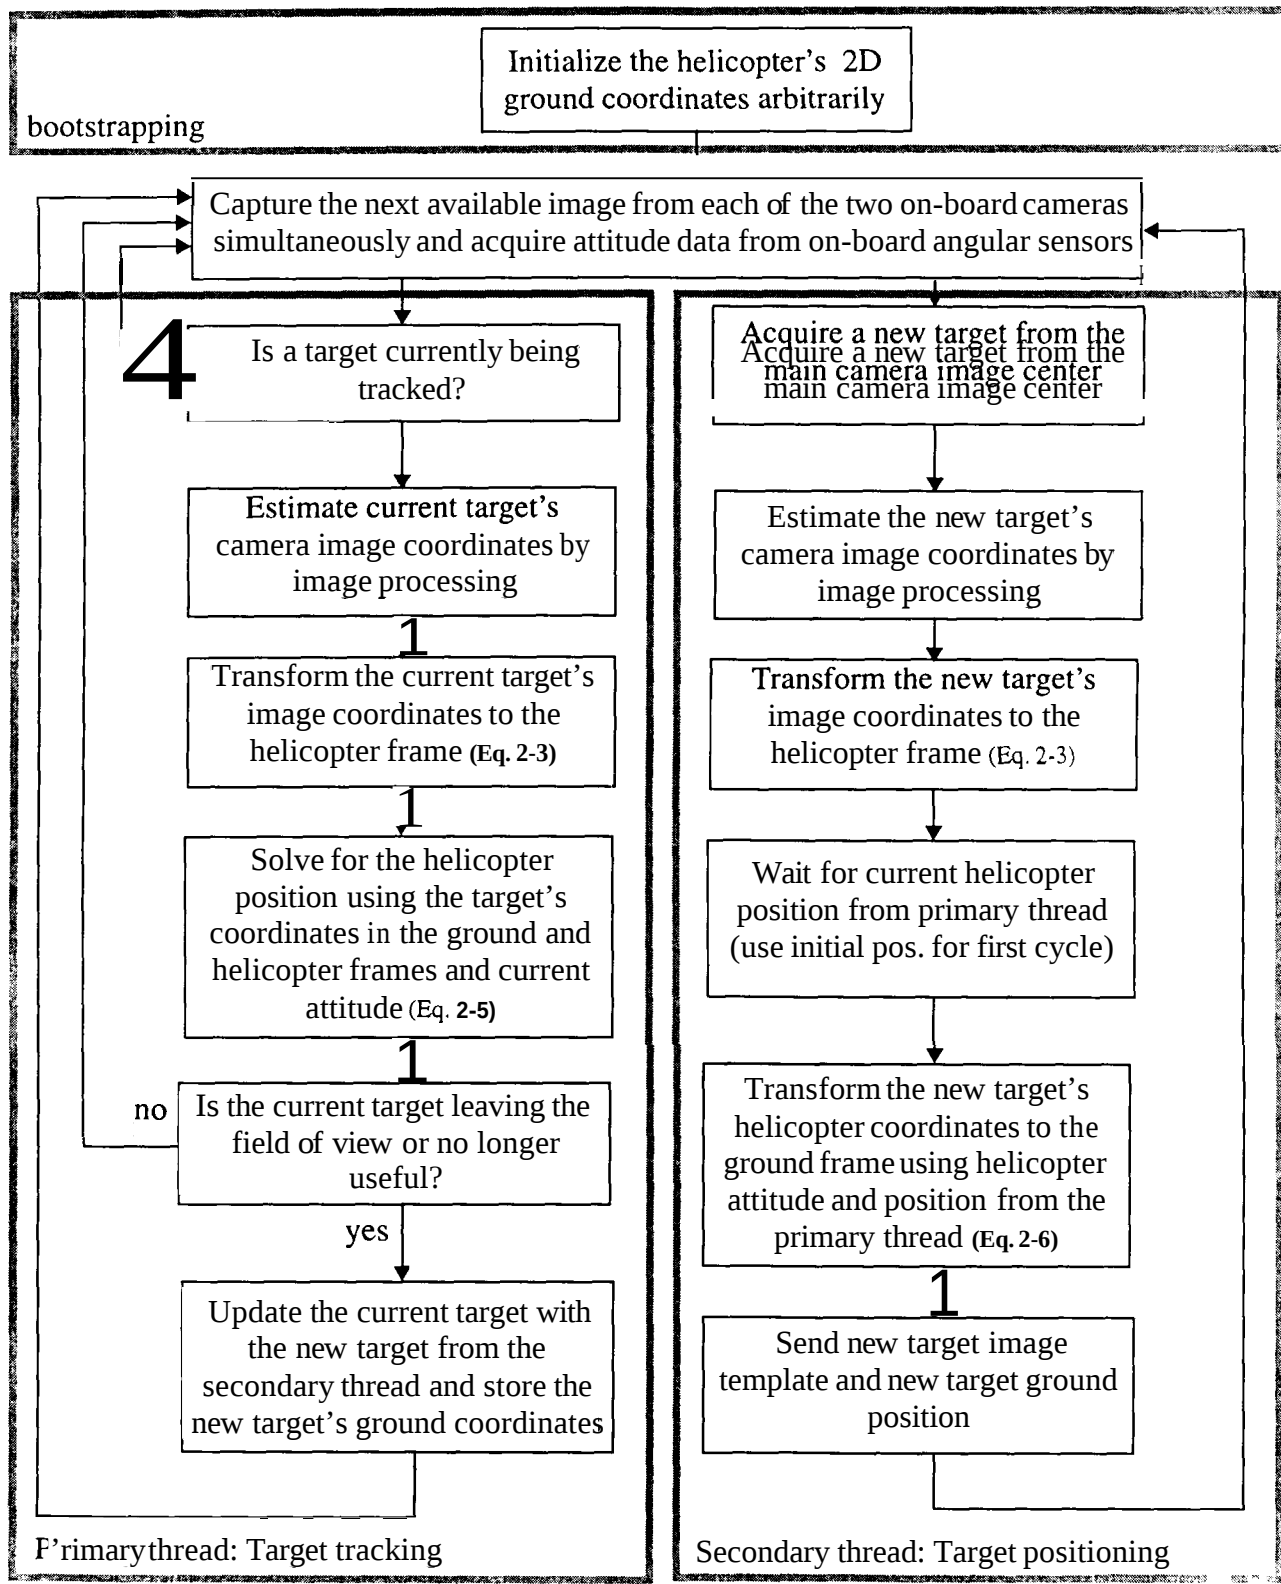


Figure 2-6 Visual Odometer Tracking Algorithm Flow Chart for Position Estimation

2.3.2 Position Sensing




Figure 2-7 Tracking Algorithm Position Sensing

To analyze the tracking algorithm's position sensing, let us examine a situation where the helicopter moved from ground frame position P_0 to P_1 as depicted in Figure 2-7. (Note that the associated vectors for each position have the corresponding 0 or 1 subscripts.)

From the current estimated position, p , the algorithm estimates the new position, P_1 , by first localizing the ground target position, P_T , and sensing the target view vectors V_0 and V_1 , and cam-

era vectors C_0 and C_1 , in the ground coordinate frame. By vector arithmetic the new position is computed from (2-5). The superscripts denote the coordinate frame in which vectors are represented.

$$P_1^G = P_T^G - V_1^G - C_1^G \quad (2-5)$$

where:

$$P_T^G = V_0^G + C_0^G + P_0^G \quad (2-6)$$

$$V_1^G = R_{H_1}^G R_C^H V_1^C \quad (2-7)$$

$$C_1^G = R_{H_1}^G C^H \quad (2-8)$$

and:

$$V_0^G = R_{H_0}^G R_C^H V_0^C \quad (2-9)$$

$$C_0^G = R_{H_0}^G C^H \quad (2-10)$$

$R_{H_1}^G$ and $R_{H_0}^G$ are the helicopter-to-ground rotation matrices at the two positions, R_C^H is the constant camera to helicopter rotation matrix, and C^H is the constant main camera translation vector in the helicopter frame.

The algorithm commences with the secondary thread's localization of a ground target taken from the image center using (2-6). This localization requires sensing the ground target location in the camera frame which defines the view vector, V_0^C , in (2-9) and transforming the camera translation vector to the ground frame by constructing the rotation matrix, $R_{H_0}^G$, in (2-10).

The secondary thread determines V_0^C by first sensing the target's image coordinates and applying (2-4). The two dimensional pixel coordinates are simply zero as the target is acquired from the image

center and the secondary thread estimates the target's range by stereo image processing. The details of the range detection are presented later in Section 2.4.3.

The secondary thread constructs the rotation matrix $R_{H_0}^G$ using (2-2) based on helicopter attitude which is sensed by on-board sensors. It then transforms the camera translation vector to the ground frame using (2-10) to localize the ground target by (2-6).

With a localized template to lock on to, the primary thread estimates the helicopter's new position, P_1^G , using (2-5). To apply this equation, the primary thread determines the view vector V_1^G by image processing and transforms the camera translation vector, C_1^G , to the ground frame using measured helicopter attitude.

To determine the view vector, V_1^G , the primary thread senses target location in the camera frame by image processing and transforms its location vector to the ground frame using (2-7). The image processing steps to locate the target include target location in the image by template matching, and target range detection by stereo image processing. Sections 2.4.4 and 2.4.5 present the details of target image location and range detection.

The primary thread positions the helicopter by this method while it is in view and trackable. If the primary thread must switch to a new target, it acquires the new target's position from the secondary thread and follows the same procedure for uninterrupted helicopter positioning.

2.3.3 Velocity Sensing

Along with helicopter position estimation, the visual odometer's tracking algorithm estimates helicopter velocity. While the primary thread is estimating current helicopter position, the secondary thread estimates the pixel velocity, referred to as optical flow, at the image center to estimate helicopter velocity.

The relationship between image center optical flow and helicopter velocity can be determined by differentiating (2-5). If the stationary target position is known, (2-5) can be rewritten for position P_0^G as follows:

$$P_0^G = P_T^G - V_0^G - C_0^G \quad (2-11)$$

Rewriting (2-6) and (2-7) for V_0^G and C_0^G and substituting them into (2-11) yields:

$$P_0^G = P_T^G - R_{H_0}^G (R_C^H V_0^C + C^H) \quad (2-12)$$

which represents current helicopter position in terms of sensed view vector and camera translation vector. Differentiating (2-12) yields:

$$\dot{P}_0^G = \dot{P}_T^G - \dot{R}_{H_0}^G (R_C^H V_0^C + C^H) - R_{H_0}^G \dot{R}_C^H V_0^C \quad (2-13)$$

The ground target is assumed to be stationary; therefore, $\dot{P}_T^G$ is zero and can be eliminated from (2-

13). Since the target is taken from the image center, the view vector V_0^C is simply $\begin{bmatrix} 0 & 0 & z_T^{im} \end{bmatrix}^T$

where z_T^{im} is the ground target's range to the camera. Substitution of the image coordinate expressions from (2-4) into (2-13) yields:

$$\dot{P}_0^G = -\dot{R}_{H_0}^G \left(R_C^H \begin{bmatrix} 0 & 0 & z_T^{im} \end{bmatrix}^T + C^H \right) - \frac{1}{f} R_{H_0}^G R_C^H \begin{bmatrix} z_T^{im} x_T^{im} & z_T^{im} y_T^{im} & f z_T^{im} \end{bmatrix}^T \quad (2-14)$$

Equation (2-14) describes how the secondary thread estimates helicopter velocity by sensing the

image optical flow vector, $\begin{bmatrix} x_T^{im} & y_T^{im} \end{bmatrix}^T$, and range, z_T^{im} , at the image center. The helicopter to ground rotation matrix and its time derivative are determined based on sensed helicopter attitude from on-board sensors.

2.4 Visual Odometer Image Processing

The visual odometer tracking algorithm is built on three elements estimated by image processing:

- Target image pixel coordinates: x_T^{im}, y_T^{im}
- Target range: z_T^{im}
- Image center optical flow: $\dot{x}_T^{im}, \dot{y}_T^{im}$

The odometer estimates these quantities by image template matching and synchronized helicopter attitude measurement. The following subsections examine the odometer image processing algorithms for estimating these quantities and the template matching method upon which these algorithms are built.

2.4.1 Image Pixel Coordinates of the Target Template

The odometer initially acquires a target template from the image center and thereafter tracks it to maintain helicopter position. Therefore, the odometer's helicopter positioning accuracy is directly affected by the accuracy and robustness of this tracking operation. Templates must be tracked consistently as images rotate with changing helicopter attitude, vary in size as the helicopter ascends or descends, and vary in intensity with changing lighting conditions. Therefore, for robust and accurate matches, the odometer must calibrate the target template before matching by rotating, scaling, and normalizing intensity.

To aid in calibration of templates, the odometer tracks an auxiliary template in parallel with the main target template. The auxiliary template provides another anchor point in the image to measure image rotation and scaling. The change in angle of the baseline between the templates measures image rotation, and the change in their baseline measures image scaling.

Figure 2-8 depicts the odometer's image processing steps for detecting template pixel coordinates. The odometer acquires an auxiliary template near the main template at the image center. The auxiliary template is offset by a nominal (20 pixel) distance in the x direction and has the same y pixel coordinate as the main template. This x direction offset provides the initial horizontal base line representing zero image rotation.

The odometer stores the two templates from the initial image as locked on targets and commences matching them to incoming images. The templates are calibrated using the baseline and image intensities of the previous match. The first match does not require template calibration since the image rotation, scaling, and intensity variation in one cycle is assumed to be insignificant relative to the algorithm tracking frequency. The validity of this assumption is later demonstrated experimentally.

To accommodate for changes in template intensity, the template pixel intensities are normalized to correspond to the most recent image match. In addition, a scale factor is determined by comparing the intensity within the calibrated templates and the intensity within the matched areas. Finally, the stored templates are rotated by the baseline angle before locating the templates in next image. The primary thread of the algorithm then employs the pixel coordinates of the main template to estimate helicopter position.

The templates leave the image from time to time with helicopter motion and must be reacquired from the image center. The secondary thread always captures new templates and the primary thread locks on to them under the following conditions:

- Either one of the current templates is about to go out of view. **As** is presented Section 3.4.3, templates are searched for in the neighboring area of the last successful match. **A** template is replaced if the neighboring search area reaches an image border.

The current template baseline length, D , is above a threshold value. **A** large baseline indicates a significant reduction in helicopter altitude or large attitude changes which degrade the resolution of




Figure 2-8 Detecting Target Image Pixel Coordinates

calibrated templates, and leads to poor matches. Large template separations also signal a potential mismatch of one of the templates; switching both templates improves the matching for upcoming images.

- The current baseline length, D , is below a threshold value. A short baseline is a result of a significant gain in altitude or a large change in helicopter attitude. The angular resolution of image rotation angle reduces with baseline length and new templates are necessary restore this resolution.

2.4.2 Range Measurement

Helicopter position and velocity estimation by object tracking requires knowledge of the object range to the camera. Object range along with lens focal length are the two variables necessary to transform object image coordinates to the camera frame as shown by (2-4). The odometer only needs to estimate object range during flight as lens focal length can be determined off-line using a number of lens models, see [35] for comprehensive work on lens calibration.

The odometer measures object range by stereo image processing. It interpolates one range estimate at the image center to detect the range of objects appearing in different parts of the image. The odometer interpolates the one range estimate using current helicopter attitude by assuming the ground below the helicopter is locally level. This approach guarantees range detection for a predetermined helicopter altitude range by choosing the camera baseline so that the center template of one camera is always visible by the other. Matching individual templates as they move through images is more accurate, but significantly limits the allowable template travel region to guarantee visibility by both cameras.

2.4.2.1 Image Center Range Estimation

The odometer estimates the image center range by locating the new potential templates captured from the main camera image center in the rear camera image. The two observations of the same object along with known camera baseline and focal lengths allow the odometer to solve for camera range. Figure 2-9 shows the processing steps and the associated variables of this process.




Figure 2-9 Stereo Camera Range Estimation

The odometer normalizes the main camera template intensity to match the intensity of the rear camera template before the matching operation. The intensity of the previous success full match is used for this normalization. The normalized main camera template is then located in the rear camera which is centered vertically (y coordinate is zero) and offset in the x direction since the cameras are assumed to be accurately aligned.

Denoted by x_p^{imr} , the location of the match in the rear camera gives camera range by (2-15) in which D and f are camera baseline and focal length, respectively.

$$\frac{im}{z p} = -\frac{Df}{x_p^{imr}} \quad (2-15)$$

2.4.2.2 Range Interpolation

Relying on a locally flat ground assumption, the odometer interpolates the image center range to estimate the range of templates in arbitrary image locations by using measured helicopter attitude. Figure 2-10 depicts the relevant variables and vectors used for this range interpolation.




Figure 2-10 Range Determination

Referring to Figure 2-10, the odometer interpolates the image center range, R , to find template range, T , using current helicopter attitude in the ground frame. Using similar triangles, we find the following relationship between T and R :

$$T = \frac{E}{P} R \quad (2-16)$$

where $E = \frac{\bar{v}^G \cdot e_z}{\bar{w}^G \cdot e_z} R$, $P = \frac{R}{w^H \cdot e_z}$, and $e_z = [0 \ 0 \ 1]^T$

$\bar{v}$ and $\bar{w}$ are unit vectors pointing to the current template object position and along the main camera z axis, respectively, and are defined as follows:

$$\bar{v}^G = R_H^G R_C^H e, \quad (2-17)$$

$$\bar{w}^G = R_H^G R_C^H \bar{w}^C \quad (2-18)$$

The C , G , and H superscripts denote camera, ground, and helicopter coordinate frames, respectively, and the matrices R_H^G and R_C^H define helicopter to ground and camera to helicopter rotations previously described in (2-2). The odometer estimates the only unknown variable, $\bar{w}^C$, based on current template image coordinates using (2-19) where (x_T^{im}, y_T^{im}) are the template image pixel coordinates and f denotes main camera focal length:

$$\bar{w}^C = \frac{1}{\sqrt{(x_T^{im})^2 + (y_T^{im})^2 + f^2}} \begin{bmatrix} x^{im} \\ y^{im} \\ f \end{bmatrix} \quad (2-19)$$

2.4.3 Pixel Velocity at the Image Center

The odometer senses the pixel velocity or optical flow at the image center to estimate helicopter velocity. It senses the optical flow by locating the center template of the current image, which is continuously captured to serve as a potentially new template, in the successive image. This operation is outlined in Figure 2-11.




Figure 2-11 Template Matching to Estimate Pixel Velocity

The odometer matches the previous center template to the current image. The found location, (x_P^{im}, y_P^{im}) , indicates the image displacement in one algorithm period which determines pixel velocity at the image center, (x_T^{im}, y_T^{im}) , using (2-20), where F represents the template tracking frequency in this equation. Template calibration is unnecessary since a fresh template is acquired in each cycle.

$$\begin{bmatrix} \dot{x}_P^{im} \\ \dot{y}_P^{im} \end{bmatrix} = F \cdot \begin{bmatrix} x_P^{im} \\ y_P^{im} \end{bmatrix} \quad (2-20)$$

2.4.4 Template Matching

The visual odometer tracking algorithm is built upon image displacement measurement by template matching; therefore, positioning accuracy and robustness directly depends on the odometer's template matching capability.

2.4.4.1 Matching Criteria

The visual odometer employs the *Sum of the Squared Differences (SSD)* matching criteria to locate templates in incoming camera images. The **SSD** criteria, one of a large class of image comparison strategies [36], is the traditional choice because of its proven effectiveness in many object tracking applications. In each image, the odometer searches for the location yielding the minimum **SSD** of image and template pixels to locate a matching area. Therefore the odometer must compute the following:

$$ssd(x, y) = \sum_{n, n} [I(x + D_x, y + D_y) - T(x, y)]^2 \quad (2-21)$$

for each examined image location where $I(x, y)$ represents the image intensities, $T(x, y)$ represents the template intensities, (D_x, D_y) represent the image location being examined, and (n, m) represent the template dimensions in pixels.

Evaluating the **SSD** criteria is computationally complex. Examining each image location requires nm multiplications and subtractions. Computational cost for finding the template can be reduced by restricting the search area to a small window around the previous successful template match. The area size linearly affects the computational complexity.

The odometer examines a search area around templates based on an experimentally determined maximum change in template location within one processing period. As the helicopter altitude decreases, the same translational motion causes a larger displacement in the image. The entire image may need to be searched for to locate the template match candidates. The odometer's search implementation is described in Section 3.6.

2.4.4.2 Coarse to Fine Search

The odometer employs a coarse-to-fine strategy to further improve the processing complexity. Initially, the template and image are subsampled by calculating the **SSD** of every fourth pixel to narrow the search to a 9x9 pixel area as shown in Figure 2-12. The subsampled match is then improved by computing the **SSD** at the unexamined pixels within the subsampling neighborhood.

Image subsampling can be susceptible to mismatches, especially in images with highly contrasting intensities. Typically an image pyramid [37] is constructed by interpolating adjacent pixels for multi-resolution searches. The interpolation is computationally complex, but necessary for consistent template matching to high contrast images. However, for the helicopter application, images must be filtered due to the significant inherent noise of the power plant and on-board electronics which significantly lowers image contrast and eliminates the need for pixel interpolation. In fact, by smoothing images with an 8x8 Gaussian convolution mask, the odometer produced consistent matches of high contrasting natural vegetation by subsampling alone. There was no need for pixel interpolation.




Figure 2-12 Coarse-to-fine search

2.4.4.3 Subpixel Interpolation of Matching Position

The odometer improves the template match location to subpixel accuracy by fitting a two dimensional parabolic surface to the SSD error of the pixel match candidates using the following equation:

$$ssd(x, y) = ax^2 + by^2 + cxy + dx + ey + f \quad (2-22)$$

where (x, y) are the match candidate locations and (a, b, c, d, e, f) represent the parabola coefficients.

Typically there are more match candidates than the six parabola coefficients, and the odometer employs a least squares parabolic fit to determine the subpixel match location. The least square parabola coefficients are determined by:

$$\begin{bmatrix} a & b & c & d & e & f \end{bmatrix}^T = (A^T A)^{-1} A^T e \quad (2-23)$$

where A is an $n \times 6$ matrix with rows representing one of n match candidates and e is a vector of SSD errors for the corresponding pixel. Each row of A consists of the parabola variables,

$\begin{bmatrix} x^2 & y^2 & xy & x & y & 1 \end{bmatrix}$, evaluated at the particular integer pixel coordinates. The matrix A can be stored as a constant provided the subsampled match candidates are always translated to the image center before fitting the parabola. With this approach, the right hand side of (2-23) can be calculated by one matrix vector multiplication at run-time. Figure 2-13 shows an example of a fitted parabola to an (8x8) pixel grid of match candidates.

In addition to subpixel accuracy, the fitted parabola provides match uncertainty information. A steep parabola versus a shallower one signals a more accurate match. For instance, the fitted parabola of Figure 2-13 is shallow, indicating a poor match. The parabola coefficients which determine the shape of the parabola provide valuable information about the match quality. They can be used to construct a covariance matrix describing template match uncertainty in two dimensions. This uncertainty measure is essential for data fusion from multiple template matches or external sensors to improve image displacement estimation.




Figure 2-13 Parabola Fit to SSD Error

2.5 Position Estimation Experiments

Before performing potentially dangerous and logistically difficult experiments outdoors, the visual odometer was tested and verified indoors using smaller and easier to manage model helicopters. This section presents the indoor experimental setup and position estimation results of a working visual odometer prototype.

2.5.1 Indoor testbed

A six-degree-of-freedom (6-DOF) testbed was developed for evaluating various position estimation and control systems developed in the span of the work presented in this dissertation. As shown in Figure 2-14, the 6-DOF testbed (described in Appendix A) supports an electrical model helicopter attached to poles by graphite rods for safety and helicopter ground-truth position estimation. The testbed allows unobtrusive helicopter free flight in a cone shaped area with mechanical stops preventing the helicopter from crashing or flying away.




Figure 2-14 6-DOF Testbed for Indoor Position Estimation Experiments

The testbed helicopter is outfitted with two CCD cameras for image acquisition and three small gyroscopes for attitude measurement. The level area underneath the helicopter is covered with gravel to simulate outdoor rough terrain. At times, various objects of interest are placed on the gravel for object detection experiments. In addition, adjustable lighting installed around the testbed evaluates system performance under changing lighting conditions.

With this experimental setup, experiments with the visual odometer experiments were performed by off-board computing (not shown) before system integration on-board a larger and more capable helicopter for outdoor flight.

2.5.2 Position Estimation Results

Figures 2-15 and 2-16 compare ground truth lateral and longitudinal position, and helicopter height¹ (dashed lines) measured by the testbed, with vision-based estimates (solid lines), which were collected during flight tests of the 6-DOF testbed helicopter. To accurately compare the positioning, the odometer's position estimates are delayed by the 1/60 second system latency to match the ground truth data. The graph at the bottom of each figure shows the absolute value of the positioning error.

Helicopter maneuvers were performed under computer control (refer to Section 4.4 for control system description) during the flight tests to observe the positioning accuracy under abrupt (1-3 Hz, 5-10 degrees amplitude) attitude oscillations. Helicopter attitude variations during data logging is shown in Figure 2-17. In spite of the constant attitude oscillation and camera vibration, the lateral and longitudinal position estimates are accurate within 1.5 cm with 1/60 second latency. Errors in the longitudinal direction are 50-60% larger due to the lower image resolution² in this direction which reduces the template match location accuracy.

The errors stem from small position displacements introduced each time a target templates is reinitialized from the image center. The lens imperfections, which were not carefully modeled, are a significant source of these position variations. Templates leaving the image are near the border and

1. Helicopter height is plotted in relation to stand rod height of 1.1 m above the ground.

2. Image sampling frequency of 6 MHz was used producing an aspect ratio of 2.3.

Position Estimation Experiments




Figure 2-15 Ground Truth Lateral and Longitudinal Position vs. Vision-Based Estimates




Figure 2-16 Ground Truth Height vs. Vision-based Estimate

largely affected by lens distortion compared to new templates taken from the image center which have little distortion effects. A discrepancy of up to 3 image rows was observed near the image border using a 6 mm wide-angle lens. This pixel error translated to 0.54 cm error in helicopter lateral positioning for the helicopter height (~ 1.1 m) in the experiments. The same 3 pixel error produced a maximum of 1.2 cm error in the longitudinal direction each time a template was reinitialized. During hovering or oscillatory maneuvers, templates may leave and enter the image from the opposite image boundaries and, in effect, cancel some of the error over time. The drift rates of 2-3 cm were observed for 1 minute hovering flight tests.

The discrepancy in the odometer's height estimation versus the testbed ground truth was partially caused by the 1-3 cm high stones placed under the helicopter. The testbed estimated height with respect to the flat ground under the stones. This estimate does not match the true helicopter distance from the stones. The height estimate did not drift since no integration was performed. Other sources of error, such as camera alignment and lens distortions, were not modeled and the 1-3 cm accuracy was found sufficient for helicopter control applications.

Position Estimation Experiments




Figure 2-17 Helicopter Attitude During Vision-based Position Estimation

2.5.3 Velocity Estimation Results

Helicopter velocity measurement based on image center pixel velocity is shown below in Figure 2-18. The estimated velocity is within 15cm/s of the stand velocity estimate. The velocity discrepancy stems from the stand's damping of helicopter vibration before it is measured by the testbed sensing elements. The vibration has a fundamental 12.5 Hz from the main rotor spinning frequency.




Figure 2-18 Ground truth lateral and longitudinal velocity vs. vision-based estimates

As a measure of consistency, the lateral and longitudinal position estimates presented earlier are differentiated and compared with the optical flow based velocity results. Figure 2- 19 shows the differentiated position (dotted) versus the measured velocity (solid). The two velocity measures agree within a 5 cm/s margin. This suggests that what appears as measurement noise is, in fact, the actual helicopter vibration. This velocity estimate accuracy was deemed sufficient for helicopter control.




Figure 2- 19 Differentiated Position vs. Measured Velocity

2.6 Summary and Discussion

This chapter presented a visual odometer for helicopter positioning, one of the major contributions of the work presented in this dissertation. The odometer incrementally maintains helicopter position by sensing image displacements. It senses these displacements by visually locking onto ground objects by image template matching. The odometer eliminates the effects of helicopter rotation from sensed image displacement by measuring changes in the helicopter attitude with each camera image capture. The disambiguated image displacement is then transformed to determine helicopter motion. The odometer relies on two main assumptions. The first assumes that the helicopter flies over locally flat ground and the second assumes that the objects appearing in the field of view are rich in features.

The locally flat ground assumption simplifies range measurement by local interpolation of one range estimate at the image center using current helicopter attitude. Extending the algorithm's capability to handle non-flat areas requires matching of the main target template in two cameras. This approach would provide accurate range to the target objects regardless of the ground shape. This extension of the algorithm is not pursued in building the first autonomous vision-guided helicopter prototype. Flight experiments presented in this dissertation were performed over locally flat or gently sloping farm land.

The assumption regarding the availability of image features can be made less restrictive by selecting high texture image regions. The statistical distribution of template matches can be applied to the entire image to determine areas of high contrast suitable for target template selection. This approach requires a ten fold increase in the computational power currently realizable on-board the helicopter and was not pursued.

This chapter introduced the visual odometer by describing its underlying components and provided evidence of its motion estimation ability through indoor flight tests. An off-board implementation of the odometer, integrating on-board angular sensors through a tether, positioned an electrical model helicopter within 1-3cm accuracy as evaluated by the six degree-of-freedom testbed. The odometer was implemented by a new real-time and low latency vision machine architecture. The vision machine architecture and specific implementation details of the odometer are the focus of Chapter 3.

Chapter 3.

A Real-Time and Low Latency Visual Odometer Machine

A real-time vision machine must perform high-bandwidth, and versatile operations on an uninterrupted, overwhelming volume of image data. Observing some of the characteristics of successful and unsuccessful vision machine architectures provides insight into vision machine design. An add-on function box, such as a convolver, is an inexpensive way to obtain video-rate performance for certain types of operations, but it is single functioned and lacks programmability. In a rigid synchronous pipelined system, an image processing pipeline must be synchronized with incoming images. Therefore a minor computational deficiency due to additional requirements in one of the stages causes major latency and throughput penalties.

While a generic computer is completely “programmable”, it is hard to perform even a relatively small window correlation at video rate. A massively parallel SIMD architecture with a grid of processors can perform uniform local image operations, such as a small-kernel convolution or graphical image warping, at lightening speed. It is, however, common experience that as soon as the processing becomes global requiring information from remote pixels, or as the control becomes data dependent, such machines are miserably slow.

What is observed is that an image processing function can be local or global, its operations can be uniform or non-uniform, **and** its control flow can be data-dependent or data-independent. Depend-

ing on its properties and applications, the most suitable architecture varies. Input bandwidth (i.e., data access requirements), processor bandwidth, and output bandwidth must be properly balanced to achieve the highest performance.

The vision machine architecture must stress modularity, expandability, and simplicity in configuring a target machine, rather than blind “generality.” It is not important how “general” a fixed machine is, but how quickly a specific machine can be configured for an application at hand.

This chapter presents a new reconfigurable low latency vision machine architecture developed to implement the helicopter visual odometer. The odometer machine specifications, architecture, and components are presented along with an analysis of the machine’s data flow and synchronization with external sensing.

3.1 Visual Odometer Machine Specifications

The helicopter visual odometer poses a number of difficult demands on an on-board vision system. The system must process images in real time with minimum latency. It must possess high computational power and data throughput, and must allow close integration of external sensing with vision. Furthermore, the machine must be compact enough in size, and efficient in power usage to fit on-board a small helicopter.

Processing latency: Stable helicopter flight with vision requires low latency image processing. Helicopters are unstable and require frequent state feedback for stable control compensation. Feedback latency rapidly degrades the stability of this control compensation. Therefore, to be effective, the helicopter visual odometer must be realized by a low latency vision processing system. Indoor testbed experiments using classical PD controllers demonstrated that feedback latency of 1/30second is adequate for stabilizing small and mid-sized helicopters. This latency figure is employed as the visual odometer machine base latency specification.

Low processing latency also improves positioning accuracy and, in fact, reduces computational complexity of the visual odometer. Low latency increases the chance of finding a template in a smaller search area around the previous successful match, therefore reducing the necessary computational power. Long processing delays require searching larger image areas using more complex algorithms which, in turn, require more computational power.

Computational power and data throughput: A vision machine must balance computational power and data bandwidth for low latency processing since captured images must first travel to the processor's local memory before the processing can begin.

The visual odometer maintains helicopter position, velocity, and height by tracking multiple templates simultaneously in two filtered camera images. The visual odometer machine must capture images, filter them by convolution, and search for the templates all at video field rate (60 Hz) to keep within the short latency window necessary for stable helicopter flight. High data bandwidth for image transfer and high computational power for image processing must be accurately matched in performance to implement this system. An implementation of the visual odometer tracking algorithm capturing (256 by 256) pixel images, (40 by 40) tracked templates, and (8 by 8) Gaussian image convolution at field rate requires close to 540 million multiplications per second and 8 MBytes per second data throughput.

Close integration with other sensors: The helicopter attitude variations result in large apparent image displacements, which impose difficulty in precise operation of the visual odometer. This effect can be eliminated by measuring these variations with on-board navigational sensors and compensated beforehand. The attitude compensation is effective only if helicopter attitude is measured in perfect synchronization with each camera shutter opening.

To provide this level of sensor integration for the visual odometer, a helicopter system including vision, navigational sensors, and control must be equipped with accurate event synchronization. Most critical is a tagging system to label incoming images with synchronized sensor data, and flexible external interfaces to capture data from different sensing devices.

Physical compactness: An on-board implementation of the visual odometer must be compact in size and weight, and use power efficiently to be carried on-board a small helicopter. Computational power and data throughput resources must be tailored to the needs of the odometer to optimize the use of available on-board space, payload, and power.

3.2 Visual Odometer Architecture

The lack of commercially available vision systems capable of meeting the requirements presented earlier motivated the development of a new architecture for real-time and low latency vision to implement the visual odometer machine. This section describes this architecture by focusing on its two distinguishing features: a decentralized communication scheme and modular structure.

3.2.1 Decentralized Communication

The architecture uses a network of decentralized, high-speed, asynchronous communication links which serve as the system's arteries instead of a shared global bus. The same links carry system control packets for initial boot-strapping, monitoring, and diagnostics.

There are a number of advantages to this approach. Communication rates among system modules are consistently predictable since the links are independent and can operate without interruptions. Furthermore, module additions or deletions do not affect the communication bandwidth of other system modules. In fact, different modules can be tested individually or bypassed in the processing pipeline to pinpoint trouble spots.

The communication scheme also reduces latency by eliminating large synchronous frame stores typically present in vision systems. Images can flow to all processing elements which internally store and process only relevant image segments as early as possible. This feature is critical to reducing processing latency for matching operations. For instance, as images are digitized line by line, a processor need not wait for the arrival of the entire image before locating a template which was previously observed near the top of the image.

Processing incoming images without a frame store requires carefully balancing incoming image traffic with module processing capabilities. Modules must asynchronously keep up with the large volume of data which is continuously sampled from the synchronous camera image signal. The communication scheme of the architecture addresses this issue by employing intelligent communication port (comm-port) interfaces and data broadcasters.

Each module supports a communication interface at each connection site or port. These interfaces support small queues to eliminate the effects of uneven input/output data rates of the asynchronous links which must continuously accept synchronous image data. The queues provide temporary storage to even out data transfer surges, thus allowing modules to receive data at constant predictable rates. The size of these port queues depends on the data transfer variation and must be carefully selected by the system designer.

The communication scheme also supports a data broadcasting capability to cope with applications demanding higher communication bandwidth and/or computational power than any single module can provide. Data broadcasters transfer multiple copies of the same data from one comm-port in parallel to multiple processors to minimize processing latency. The broadcasters can support their own port interfaces with data queues to even out transfers to each receiving module. The prototype vision machine shown in Figure 3-1 employs one broadcaster to divide the velocity and position estimation tasks between two DSP modules.

3.2.2 Modular Architecture

The decentralized communication scheme works hand-in-hand with a decentralized modular processing architecture. Interconnected via high-speed links, system modules incorporate local intelligence to perform complex tasks orchestrated by one external real-time controller. Each module is treated as a raw source of data or computation with timing, data flow, and synchronization predetermined before machine operation begins. System supervision by one central controller reduces the complexity of individual modules and allows compact and low cost implementations of most system modules. The

controller captures complicated non-vision tasks such as external communication and user interfaces which are typical sources of processing uncertainty.

The system architecture relies on predictable vision processing latency and timing of each module. Each module is rated for its computational power and bandwidth to perform a specific vision task. Following this rating system, existing modules of varying throughput and computational power can be employed in the system, or new modules can be developed to optimize systems for different applications.

Using all available modules as their tool-box, system designers can build systems with varying throughput and latency by expanding the processing flow vertically or horizontally as shown in Figure 3-2. If latency is not important, high throughput can be achieved by a long horizontal chain of modules connected as a pipeline, with each stage performing an image processing step. On the other hand, if latency is critical to the application, modules can be arranged vertically to operate in parallel.




Figure 3-1 Horizontal and Vertical System Expansion

The vision machine supports four types of modules which include: processing modules, interface or bridge modules, synchronization or timing modules, and broadcast modules. Processor modules provide raw computation for image processing. Interface or bridge modules connect the machine to external sources such as cameras or sensors. In addition, bridge modules allow communication with global busses or networks for standard communication with commercially available systems. Synchronization or timing modules generate timing signals for machine event scheduling. Finally, using the decentralized communication scheme, the broadcast modules carry out the data communication fan-out described earlier.

3.3 Components of the Visual Odometer Machine

The visual odometer machine is composed of a number of modules including: image A/D and D/A converters, image convolvers, powerful digital signal processing (DSP) elements, an image tagging and synchronization module, and external communication bridge modules. Figure 3-2 shows how these modules are interconnected to realize a prototype visual odometer machine. This section presents the underlying structure and the implementation details of each these module.

3.3.1 Image Acquisition

Image acquisition is fundamental to the operation of vision machines. The visual odometer machine acquires images from two cameras through two independent A/D converter modules. The modules sample the analog camera signals and output images digitally through their output comm-ports. The structure of the A/D module is shown in Figure 3-3.

The A/D module provides a generic image digitization facility with a few non-standard features: programmable image sampling and synchronization, real-time configurable image blanking, and high-speed communication ports. These features were found to be extremely useful for high-speed image processing.




Figure 3-2 Visual Odometer Machine Block Diagram




Figure 3-3 A/D Block Diagram

The module supports custom-designed circuitry to generate sampling clocks and control image blanking in addition to an A/D converter and comm-port interfaces. The clock generation circuitry provides programmable image sampling and synchronization frequencies. Programmable image sampling can dramatically reduce image data traffic by proper frequency matching to camera CCD array resolution. This can provide virtually the same image content captured in significantly smaller images. On the other hand, all available pixels can be used when digitizing video signals with longer rows as with high resolution line cameras as well as image capture from non-standard video sources such as variable frequency cameras. These are important considerations, as image capturing synchronized with rotor blade revolutions is a potential outdoor requirement of the system. The A/D module also supports a configurable image blanking controller circuit that allows the processing elements to select regions of interest in the image in real time to further help reduce image data traffic.

The output comm-port interface incorporates a small storage queue to even out data output traffic. The status of this queue is used as a means of image synchronization. A full queue indicates that the receiving module is not capturing images and data is simply thrown away. If the receiving module commences reading data, the A/D module blocks transfers until the start of the next valid image field to properly synchronize image transfers without explicit hardware connections.

It is assumed that the receiving module can keep up with the image data rate and the queues will never overflow during machine operation. The size of the output queue is chosen carefully to equalize the variable image traffic and processor input data rates during valid pixel and blanking intervals.

The implemented A/D module design incorporated an 8-bit BrookTree A/D converter, supporting a built-in image look-up table (LUT), clock generator chips, and custom designed state machines implementing the comm-port interface.

3.3.2 Image Convolution

Fast convolution is essential for image processing. In addition to edge detection and smoothing, matching and feature extraction can be performed using special convolution masks. As previously presented, the visual odometer relies on fast image smoothing to subsample images for efficient template matching. In addition, image smoothing by convolution reduces the significant noise from the helicopter power plant and electronics which corrupts the camera signals.

The visual odometer machine filters images from the A/D modules using a real-time image convolver module. An application specific integrated circuit or ASIC is employed for low latency image convolution. The convolver ASIC, GEC Plessey 16488, can perform (8x8) convolutions at 10MHz per second input image data traffic which, in effect, delivers 640 MOPS. The convolution ASIC internally stores the (8x8) convolution mask and provides dedicated external expansion signals to increase mask size.

For compact implementation, the visual odometer machine simply includes the convolution ASIC within the A/D module. To provide valid data near image borders, raw digitized images are transmitted to the convolver before image window blanking. The image convolution latency is 22 pixel clocks as the convolver chip operation is internally pipelined without waiting for the first 8 lines of the image to fill the pipe stages which are filled up by image lines above the region of interest.

3.3.3 DSP Processing Module

High speed processing of images with low latency requires fast computing capable of acquiring and processing images at high frequencies. There are a number of compact CPU platforms with such capabilities including: SGS-Thomson Inmos T9000 Transputer [38][39], Intel i860 [40][41], and Texas Instruments TMS320C40 Digital Signal Processor (C40)[42][43].

The C40 is an ideal platform for image processing and is extensively used to implement the visual odometer machine. It is a powerful image processor for several reasons. The most significant is its high communication bandwidth though versatile communication-ports (comm-ports) well-suited for high-speed image transfers. The C40 supports six asynchronous comm-ports, each rated at 20 MBytes per second (MB/s) transfer rate; 14-16MB/s rates have been observed to be more typical. (See [44] for a detailed analysis of the C40 comm-ports.)

The C40 supports six DMA channels for high-speed data transfer. Each DMA channel has its own dedicated comm-port connection for high speed external data transfer. These dedicated connections help reduce the data traffic of the two main 100Mb/s external 32 bit memory interfaces of the C40. The DMA channels can perform non-stop complex data transfers using their own set of programmable instructions stored as “link pointers.” Images can be split in pieces and transferred to other C40s without any CPU intervention.

The processor also supports on-board high-speed memory for instruction cache and critical data storage. With careful resource management, incoming images from comm-ports can be stored in independent SRAM banks allowing the processor uninterrupted access to the image during image processing operations. For fast processing, the DMA channels can store image portions of interest in zero-wait-state SRAMS (20 ns access time) instead of the more traditional slow VRAM-based frame buffers.

Most CPU floating point operations such as 32-bit multiplication are single cycle instructions (two clock cycles) and the current C40s are clocked at 60MHz with plans for 80 MHz versions by the end of 1996. The C40 is rated at 275 MOPS and 320 MB/s data throughput.

3.3.4 Module Synchronization

Real-time image processing requires accurate event synchronization. The visual odometer machine relies on accurate synchronization to schedule image processing operations and to coordinate image acquisition with helicopter attitude measurement. The machine supports a central synchronization generator (sync generator) module to govern the processing and data acquisition operations with each camera image capture.

The modules configure cameras using the **NTSC** format in non-interlaced mode with a special setting to provide the same image field updated at 60 Hz frequency. The sync generator also incorporates counters to count each horizontal image line. The line numbers from this counter are used by other modules to trigger and tag incoming attitude data from gyroscopes. In effect, this tagging mechanism synchronizes attitude data acquisition with each camera shutter opening and image video line.

3.3.5 Sensor Bridge Module

Triggered by the sync generator, a sensor bridge module tags camera images with sensor data for the visual odometer machine. The sensor bridge can capture data from a variety of sources, including four independent A/D converters, ten quadrature decoder circuits, and digital input lines. In addition, the sensor bridge can output data through four **D/A** converters and a number of digital output lines. The block diagram of this module is shown in Figure 3-4.

A distinguishing feature of this module is its configurable state machine, which can arbitrarily define the type and sequence of data to be acquired or outputted. By encoding all input/output data into packets, this module reduces system **I/O** complexity and provides a simple interface to a large number of sensors and external control outputs through two communication ports.

3.3.6 Image Display

Image display is not critical for the visual odometer machine operation, but is necessary for viewing processed images. The visual odometer machine employs a **D/A** module for image display. The mod-




Figure 3-4 Sensor Bridge

ule is similar to the A/D in several respects. As shown in Figure 3-5, the module supports the same variable clocking and blanking control circuitry as does the A/D. These components enable image display on a range of monitors with different horizontal frequencies and resolution. For proper image synchronization on the monitor, an image synchronization handshaking procedure similar to the A/D module is performed at the input ports.




Figure 3-5 D/A Module Structure

To help eliminate redundant image transfers for simple display purposes, the display module has separate input datapaths for image and overlay. For example, the image received from the A/D could be displayed while the overlay plane is updated by external processing elements.

The D/A module is implemented with a BrookTree D/A chip set and clock generator circuits. The module supports triple LUTs for 256 pseudo-color display and 16 overlay colors, and can receive data asynchronously through its input ports and refresh an NTSC monitor in real-time solely from its input data received from comm-ports. The image transmitting modules need simply to send image data at predefined sizes and the module automatically produces the proper synchronization for display, dramatically reducing image display system complexity.

3.4 Data Flow and Synchronization

The tracking algorithm, implemented by the visual odometer machine, requires two main execution threads. The primary thread estimates helicopter position, while the secondary thread estimates helicopter velocity and prepares new potential templates for future tracking. Both threads estimate template range by stereo image processing and transform current template position and image center pixel velocity to the ground coordinate frame using the helicopter attitude which is measured by on-board angular sensors.

The visual odometer machine implements the above two execution threads of the algorithm using four **C40** DSP engines. To minimize latency and optimize processor utilization, the machine schedules as many image processing operations as possible as early as possible. The machine allocates two C40s to estimate current template position and image center pixel velocity, another to measure range by stereo vision, and a fourth to integrate the image processing results and compensate for helicopter attitude changes. An external real-time controller, communicating via a bridge module, accesses external data storage for system initialization and logging tasks. Similarly, gyroscope attitude data is acquired through another external bridge interface.

Figure 3-6 shows the processing event scheduling of the vision machine by a data flow and synchronization time-line. The horizontal rectangles represent data transfer and processing events and the vertical lines represent the camera shutter openings. The following sections provide detailed descriptions of the vision machine data flow and synchronization.

Data Flow and Synchronization




Figure 3-6 Data Flow and Synchronization

3.4.1 Image Acquisition and Preprocessing

The visual odometer machine acquires images from two NTSC cameras. The cameras are synchronized with one central sync generator and operate in a special non-interlaced mode which provides the same image field after each camera shutter opening. This mode guarantees a fresh image field at every sixtieth of a second.

The machine's A/D module samples the camera signals at 6 MHz which provides a maximum of 360 pixels per NTSC video line which is 64 microseconds in duration (60 for the video signal + 4 for the horizontal frequency interval). The central 268 pixels are chosen as the line area of interest using A/D blanking offsets. Similarly, 236 of the available 260 video field lines are centered and chosen to compose the vertical image area of interest.

The (268x236) field images are preprocessed by the convolution module with an (8x8) Gaussian convolution mask. For the implemented machine, the latency of the digitization and preprocessing operations were close to 4 microseconds and considered negligible.

3.4.2 Image Transfer and Storage

Rows 1,3, and 8 of Figure 3-6 show the image data transfer to C40s. They consist of periods of activity during the valid image window and inactivity during blanking intervals where no transfers are performed.

C40 DMA coprocessors transfer and store images in high-speed static memory freeing the main CPU to perform only image processing. As relevant image areas arrive, the DMA coprocessors signal the main CPU to commence processing. Since most transfers are not simple periodic operations, complex instructions using C40 link pointers are loaded during system initialization so that the DMAs can be controlled without main processor intervention. In some cases, the data packets themselves are tagged with DMA instructions for variable length transfers. The DMA channels operate in the higher speed "split mode" which allows direct connection to the normally memory-mapped comm-ports, therefore reducing the memory bus data traffic of the C40s.

Since the C40 has a 32 bit data bus structure, the incoming 8 bit image data from the comm-ports can only be stored in four pixel data elements using conventional C40 hardware implementations. The data must be unpacked before image processing can start. Unpacking an entire (268x256) image requires 3 millisecond by the main C40 processor.¹ Similarly, outgoing data for display must be packed which also consumes valuable computational power.

3.4.3 Target Template Position Estimation Processing

C40-1 tracks the two target templates of the visual odometer. The processing events of C40-1 are shown by row 2 of Figure 3-6. Processing begins after a predetermined central image region is transferred to local memory by a DMA channel. This region encompasses two (40x40) partially overlapping position templates and a 20 pixel border as shown in Figure 3-7. In case the currently tracked templates leave camera view, image pixels in this region are captured at each cycle for possible initialization of templates in the next cycle. The template pixels are stored along with their borders to provide the extra image area for rotating and scaling templates as the helicopter moves. The pixels are unpacked and prepared by the main processor in parallel while the rest of the image is arriving.

The main processor begins searching for the two target templates upon image transfer completion. The search area for each template encompasses the last matched template surrounded by a 16 pixels wide border. Immediately following the coarse to fine template match process described in Chapter 2, the match locations are transmitted to C40-2 via a different DMA channel, and the processor starts template preparation for the next cycle.

If the templates must be updated, the processing is terminated and previous template locations are simply recorded for integration. More frequently, the templates must be rotated, scaled, and adjusted for image intensity variations for the next cycle. The processor updates the templates while the next image is being transferred to one of its memory banks.

1. For efficient image packing/unpacking for the visual odometer machine, the C40s are programmed in assembly language to take advantage of parallel load/store instructions and all independent global, local, and on-chip data storage to maintain the processing pipeline as full as possible. In addition, costly extra static **RAM** is incorporated in each C40 module to store packed and unpacked images. Intelligent hardware image packing and unpacking circuits are designed but not implemented.




Figure 3-7. Position Estimation Templates

3.4.4 Stereo Processing

C40-2 matches templates for stereo range measurement. Row 6 of Figure 3-6 shows its main processor activity. Image segments from both cameras are transferred to local memory by the **DMA** coprocessors. The stereo matching locates the (40x40) center template of the front camera image, received from C40-3, in a portion of the right camera image, acquired from the A/D module. To reduce latency, the main processor unpacks the front camera image template while the search area in the rear camera image is being received. The cameras are accurately calibrated to limit the search area to a horizontal rectangle in the rear camera image. To minimize the search area off-line stereo matching is used to align the cameras.

The matching process starts by a slow match to initially locate the template in a (72x 100) search area and continues by fast matches searching for the template displaced by ± 16 pixels from the last successful match. The range measurement result is combined with the position estimates from C40-1 and transferred to C40-3 by a DMA channel.

3.4.5 Pixel Velocity Processing

C40-3 matches one template for velocity estimation based on image center pixel optical flow. Its processing events are shown by Row 10 of Figure 3-6. Filtered main camera images are transferred through the broadcast module by a DMA coprocessor. A predetermined central image region encompassing a (40x40) template and a 16 pixel wide search area is used for the search area. Template matching begins immediately after this portion of the image is received by a DMA channel. While the matching process is being performed by the main processor, a DMA channel transfers the center template to C40-2 for stereo matching. The template transfer time is quite short in comparison with the A/D transfers since the **C40** comm-ports have significantly higher (16-20 MB/s) transfer rates. Position, range, and velocity template matching results are combined in one packet and sent to C40-4 by a **DMA** channel.

3.4.6 Attitude Compensation

C40-4 compensates for the effects of helicopter attitude variations on image displacement. Row 10 of Figure 3-6 shows C40-4's processor activity. C40-4 receives sensor data sampled and tagged with each video line video line (~15KHz) by the sensor bridge. The main processor filters the sensor data, with a latency of 100 video lines or 6.4 milliseconds, before compensating for the attitude variations. To compensate for filter latency, sensor data acquisition begins before each shutter opening to ensure precise attitude measurement at each shutter opening.

3.5 Summary and Discussion

This chapter presented the second contribution of the work presented by this dissertation, a configurable vision machine architecture for real-time and low latency image processing. The architecture's versatile communications scheme and the modular architecture are instrumental in its configurability to different applications. The architecture can be modified for low latency or high throughput to best meet application requirements at optimal system size and cost. Yet this configurability has some drawbacks. Machine event scheduling and programming can be difficult. The system designer must keep track of each module's capabilities and limitations to arrange an optimal execution order and

“just-in-time” arrival of various data at different modules. The designer may have to experiment with many different configurations until the optimal system is realized.

The flexibility and power of the architecture is demonstrated by a visual odometer machine which integrates image acquisition and display, **DSP** processing elements, central synchronization, image convolution, and external sensing. The visual odometer machine successfully maintained helicopter position and velocity at a 60 Hz update rate with **26** millisecond latency.

Chapter 4.

Design and Evaluation of an On-Board Vision-Based Helicopter Control System

The design and implementation of the visual odometer for helicopter positioning were the first steps in building an autonomous vision-guided helicopter. Indoor flight tests using the six degree-of-freedom testbed demonstrated promising results in vision-based helicopter positioning and opened the path to more ambitious outdoor free flight experiments.

On-board integration for autonomous outdoor flight raised a number of important concerns. It was not clear that the visual odometer machine and other control and power systems could, in fact, be integrated on-board a small helicopter capable of performing a useful mission. Furthermore, since the indoor flight tests employed off-board computing and power, vibration and noise effects of the helicopter power plant on the on-board systems were never investigated. Of major concern were the effects of engine noise and vibration on camera image quality, high bandwidth image transfer integrity, and on-board navigational sensor data.

This chapter describes the experimental approach to develop and verify an integrated vision-based helicopter control system on-board a mid-sized model craft, the Yamaha R50. The chapter presents the integration of different system components such as vision computing, attitude sensing, low-level control, and power systems. Each component was individually tested using indoor testbeds, and was designed to physically fit and weigh within the small available space and payload of the helicopter.

4.6 Yamaha R50



Figure 4-1. Yamaha **R50** Helicopter

Shown in Figure 4-1 under human pilot remote control, the **R50** is a commercial product of Yamaha Motor Company for aerial agricultural pest control. Designed for spraying fields in hard to access areas, the **R50**, unlike airplane crop dusters, can fly close to the ground at slow speeds. This ability significantly improves pesticide effectiveness and reduces the undesirable overspray and dispersion of chemicals in the atmosphere.

Powered by a 98 cc water cooled 2-stroke engine which produces 12 HP, the **R50** has a payload of 20 Kg with a maximum takeoff weight of 67 Kg and can continuously operate for 60 minutes. Its overall length is 3.5 meters, with a main rotor diameter of 3.1 meters and body length of 2.7 meters.

4.7 Indoor R50 Testbed

Several factors prompted the development of a tethered testbed for the R50 as an intermediate step towards autonomous operation. Initial experiments with the R50 would be significantly safer with protective tethers since the R50 is quite massive and can accelerate to dangerous speeds rather quickly. In addition, it is easier to investigate and resolve different system integration issues indoors using proper test equipment. In particular, effects of vibration on sensors and computing enclosures can be quite significant and requires careful investigation.




Figure 4-2. Indoor R50 Testbed

4.7.1 Testbed Design

As shown in Figure 4-2, the testbed design allows relatively large (1.5 meter) longitudinal travel while severely limiting helicopter travel laterally and vertically. Yet, despite this limitation, the one-axis travel arrangement can provide great insight regarding the behavior of the R50 under total computer control outdoors, and can test the integrity of the vision-based positioning and control systems.

An off-the-ground, heavily reinforced platform is built to provide a level area for helicopter longitudinal movement. To prevent tail collisions during large helicopter pitching which may occur during undesirable oscillations, the platform is designed to be shorter than the R50.

The R50 is tethered with ropes which are fastened to the ground and two poles positioned on either side of the platform as shown in Figure 4-3. A steel rod with hooks on either end connects the ropes to the R50. The rod is secured to the helicopter's center of gravity to eliminate any torques from restraining forces which could cause dangerous rotations.




Figure 4-3. Indoor Test Flight of the R50

Different kinds of rope restrain the lateral and vertical axes of the R50. High strength and rigid Kevlar rope is used for rigid vertical restraint and limit the possibility of large rotations. High strength flexible nylon rope is used for lateral restraint to ease impacts to the helicopter when it reaches the travel extremes. This flexibility significantly reduces the magnitude of impacts experienced by the main rotor hub which absorbs forces generated by changes in the momentum of the spinning rotor blades as the helicopter hits travel extremes. To further limit potentially dangerous helicopter rotation, longer cylindrical skids with smooth plastic ends to reduce friction are installed on the helicopter.

4.7.2 Testbed R50 Helicopter

The testbed R50 helicopter is used to design and evaluate the power system, sensing, and computing for on-board integration. An on-board chassis supports all on-board systems for indoor flight tests. The chassis aligns cameras and attitude sensors, as well as isolating harsh vibration from system components. Figure 4-4 shows the different components of the testbed R50 helicopter.

4.7.2.1 On-board Power

A 12 V (7 AH) battery supplies all the power for on-board computing and sensors. The power dissipated is about 150W. DC-DC regulated converters and custom-made filtering circuits provide clean +/-12 and +5 V power signals from the 12V battery even when it is drained to as low as 9 Volts. The on-board computer monitors the battery voltage by an A/D converter and produces warnings if the voltage drops below 10.5 Volts.

4.7.2.2 Attitude Sensors

Two light-weight (40 g) and low cost gyroscopes, made by Gyration Inc., are mounted close to the center of gravity for attitude measurement. One gyroscope is directional for measuring helicopter heading and the other is vertical to measure roll and pitch. Both gyroscopes have two nested, optically encoded gimbals with 0.2 degree angular resolution. The gyroscopes are mechanical and incorporate a mass spun at high speeds by a DC motor. The motor speed is regulated by the input voltage and is set at its maximum (15,000 RPM) for best performance.




Figure 4-4. Testbed R50 helicopter

The directional gyroscope is quite drift prone and a few **(2-4)** degree per minute accumulating drift rates are not uncommon. The drift rate of the vertical gyroscope's roll and pitch angles is similar but does not accumulate over time since it incorporates a pendulous inner gimbal with a 10 minute time constant. Using the gravity vector to eliminate long-term drift, the pendulum levels the gimbals. This leveling scheme assumes zero lateral acceleration over long time periods the order of 15-20 minutes.

Both gyroscopes provide relative angular measurement by generating digital pulses as the helicopter changes its attitude. The pulses are integrated to estimate attitude relative to starting gimbal positions. The vertical gyroscope levels itself in a **20** minute period and the roll and pitch values are initialized using the gravity vector measured by three accelerometers. The heading from the directional gyro is simply initialized to zero in the forward direction of the testbed platform.

An accurate navigational quality vertical gyroscope, on loan from Humphrey Gyroscopes, is used to evaluate the performance of the small inexpensive vertical gyroscope during indoor flight experiments.

4.7.2.3 Testbed Cameras, Scenery, and Lighting

As shown in Figure 4-5, two ground pointing CCD cameras (Sony XC-75) are mounted on the side of the R50. The cameras are fitted with wide angle (6 mm focal length) lenses to provide the large view angle necessary for low altitude image processing.

A non-repeating stone pattern is printed on the testbed platform to provide a feature rich scene for the visual odometer to lock on to. The pattern is covered with plexiglass for protection from the helicopter skids. The plexiglass is matted with sandpaper to reduce undesirable reflections.

The cameras are synchronized by the sync generator to produce non-interlaced fresh image fields at approximately 60 Hz. The shutter speed is set at open at 1 millisecond intervals to provide clear images under the harsh vibration as well as proper synchronization with attitude sensors.

The testbed is lit by multiple light sources from different out-of-phase circuits to reduce image flickering due to 60 Hz AC power, exaggerated by the short shutter opening interval. To further reduce

the flickering effects caused by the indoor light sources, the cameras are synchronized with 59 Hz vertical frequency, governed by a synchronization generator, instead of exactly 60 Hz.

4.7.2.4 On-board Computing

A seven slot VME cage houses all computing hardware (only six slots are required). The cage is not used on-board for indoor experiments and is removed during actual flight experiments to protect it from the unknown vibration characteristics of the R50. A tether transmits camera signals and sensor data to the cage and an RF transmitter sends control signals to the helicopter actuators.

4.7.2.5 On-board Chassis Mounting

For easy integration on multiple helicopters, the sensors and the computing enclosure are integrated into one detachable package using a chassis shown in Figure 4-5. During the indoor experiments, all the sensors are used during indoor experiments except the laser rangefinder which is mounted on the chassis for future outdoor experiments.




Figure 4-5. On-board Chassis

The chassis could not be mounted on the helicopter until vibration characteristics of the R50 are determined. Since the chassis itself can affect the vibration, all chassis components are weighed and replaced by one large metallic plate of equivalent weight to approximate the inertial characteristics without using actual components. For vibration analysis, three accelerometers, aligned with helicopter axes, are mounted on the metal plate.

A variety of mounts are tested, starting with solid rubber at the chassis's four mounting points. With rubber mounts, the accelerometers sensed a 15 g wave at frequencies as high as 167 Hz. To reduce this vibration, cylindrical rubber mounts with varying density are tested under real flight conditions. Commercial mounting materials could not withstand the vibration amplitudes and caused large chassis oscillations which are undesirable for image processing. Different rubber mounts are manufactured by drilling holes in different densities in the rubber and individually tested until the maximum observed vibration of each axis is within 0.5 g.

4.7.3 Testbed Safety

Careful tethering of the helicopter does not guarantee the safety of nearby individuals. The rotor blades may shatter or tear off the helicopter, thereby causing major injury. For safety, a control room is constructed for indoor flight experiments. To stop flying debris, two small windows are covered with wooden barricades supporting multiple layers of bullet-proof Lexan and steel mesh reinforcements.

To prevent violent helicopter oscillations, a safety control system is developed as a backup for vision-based flight experiments. In case of vision system failure, the backup control system measures helicopter position using string potentiometers. The back up system automatically takes over helicopter control if large longitudinal velocities are observed. In case of backup control system failure, an experienced safety pilot is also present in the control room to take over control using an RF transmitter.



Figure 4-6. Helicopter Control Room

4.8 Helicopter Controller

An effective helicopter control strategy is as important as the positioning estimation method. Helicopter control is inherently challenging due to dynamic coupling, and nonlinearities. Despite this, classical PD control can be quite effective for stable hovering and low-speed point-to-point maneuvering.

Employing indoor testbeds, a helicopter control system composed of a number of PD servo loops was developed. The system is capable of both hovering and slow ($< 10\text{m/s}$ or 20mph) speed helicopter flight. The success of the simple PD control approach is attributed to the high quality of vision-based positioning feedback.

4.8.1 Helicopter Control Inputs

Similar to full-sized helicopters, the remotely piloted models have four primary control inputs which are: the collective pitch angle, the lateral and longitudinal cyclic pitch inputs, and the tail rotor pitch angle. The collective pitch angle changes the angle of attack of the main rotor blades producing vertical lift which regulates vertical ascending or descending. The two cyclic pitch inputs, lateral and longitudinal, vary the rotor blade pitch sinusoidally within each revolution thus causing the rotor plane attitude changes which produce lateral and longitudinal forces to accelerate the helicopter back and forth or from side to side. The tail rotor input controls helicopter yaw angle. An additional control input to the helicopter is the engine throttle position which regulates rotor speed. This input is typically controlled by a governor or is varied with the collective pitch input to keep the rotor RPM as constant as possible.

4.8.2 PD Servo Control Loops

Figure 4-7 shows the block diagram of the control system implemented for helicopter stabilization.

The controller inputs are ground-based helicopter position and attitude, $\begin{bmatrix} x_h^G & y_h^G & z_h^G & \phi_h^G & \theta_h^G & \psi_h^G \end{bmatrix}$, and reference or desired position and heading, $\begin{bmatrix} x_{ref}^G & y_{ref}^G & z_{ref}^G & \psi_{ref}^G \end{bmatrix}$ and the outputs are the four primary helicopter control inputs and throttle position. The controller is composed of several PD servo control loops for controlling helicopter attitude, lateral/longitudinal position and velocity, and height. The blocks labeled PD symbolize a linear combination of measured position and measured velocity.

A PD controller is biased and tuned for height control. The output of the controller is limited to an operating range determined by actual flight experiments. The limit also prevents sudden loss of altitude should the height sensors malfunction. In addition, the collective pitch input is used to set the engine throttle opening to maintain constant rotor RPM. A piece-wise linear function of collective pitch, recommended by the helicopter manufacturer, consisting of three operating points is used to bias the throttle input based on the collective pitch. Similarly, the tail rotor input is offset based on the collective pitch input to decouple helicopter yawing with varying loads on the main rotor. This offset is added to the output of the PD yaw control loop.




Figure 4-7. PD servo loops

Helicopter position and velocity are also controlled by PD loops which output desired accelerations in the form of helicopter roll and pitch angles. The helicopter position and velocity, measured in the ground frame, are transformed to the helicopter frame to determine errors in the correct control axes. Position and velocity control loops are not combined as it was desired to operate the system in velocity mode alone. To reduce noise, the measured velocity is employed in the PD position loop instead of differentiating measured position. The reference roll and pitch angles from the position and velocity loops are range limited before being sent to PD controllers maintaining helicopter attitude. This range limiting is necessary to maintain helicopter control within a the operating range of the linear PD servo loops and to prevent large helicopter attitude changes in case on-board sensors malfunction.

4.8.3 Controller Testing

As accurate attitude control is central to helicopter control, an attitude control testbed, shown in Figure 4-8, is employed to test and tune the PD attitude control loops using an electric model helicopter. The testbed consists of an electrical model helicopter mounted on a swiveling arm platform. An optical encoder mounted with a frictionless bearing measures ground-truth angles in real-time. The model helicopter supports a detachable sensor package for sensor calibration using the optical encoder.




Figure 4-8. Attitude Control Testbed

Following the attitude control experiments, the entire controller is implemented and tested on the 6-DOF testbed (See Appendix A). The controller stabilized the testbed helicopter with hovering accuracy of 15 cm by using 60 Hz vision-based position feedback.

Similarly, the controller is integrated with vision-based feedback for flying the R50 indoors. The controller hovered the R50 within 30 cm of the reference position. Figure 4-9 shows the R50 under computer control about 3 inches off the platform. Longitudinal and yaw control are also tested and tuned during these experiments.



Figure 4-9. Indoor R50 Computer Control

4.9 Vision and Attitude Synchronization

Accurate synchronization is key in eliminating the effects of attitude on helicopter positioning with vision. The visual odometer relies heavily on helicopter attitude, measured precisely at each camera shutter opening, to compensate for image displacements caused by helicopter rotation. The transformations in (2-7) to (2-10) are based on accurate helicopter attitude measurement with each image capture.




Figure 4-10. Attitude Synchronization Testbed

Ineffective attitude compensation can produce completely inaccurate helicopter positioning feedback which can be catastrophic during free flight. For this reason, the odometer's attitude compensation is tuned and evaluated experimentally using an attitude synchronization testbed. As shown in Figure 4-10, the testbed restricts camera movement to one dimension and incorporates a string poten-

tiometer and a shaft encoder to measure ground truth camera translation and rotation. A calibrated stereo pair of cameras is mounted on the testbed, along with a gyroscope to measure attitude. Gravel is placed under the cameras to provide features for vision.

While keeping the camera stationary, the correlation of measured attitude with image-based displacement while cameras rotate reveals the exact shutter timing in relation to filtered attitude data. This timing is employed to trigger attitude data acquisition by the sensor bridge module in the compensation system described in Chapter 3. Figure 4-11 displays the significant correction observed when measured precise timing was employed for attitude compensation. The dashed graph represents odometer output without attitude compensation, while the solid line represents the compensated position estimates which closely match the ground truth measurements shown by dots. Without correct synchronization, the compensated data is oscillatory and unusable.




Figure 4-11. Attitude Compensation

4.10 Summary

This chapter presented the experimental approach to developing an autonomous helicopter for outdoor free flight. System components such as vision computing, power, and sensing were individually tested and integrated on-board a mid-sized model helicopter, the Yamaha R50.

The helicopter's on-board systems are tested indoors using a tethered testbed. The testbed restricts movement in the helicopter's longitudinal direction and incorporated ground-truth positioning sensors for experiment evaluation and safety. Another testbed for attitude synchronization verifies vision-based positioning under severe attitude variation.

A control system is developed to stabilize the helicopter using vision-based positioning feedback. The control system is made up of a series of nested PD servo loops for attitude, velocity, and position control. The controller is integrated with on-board vision to hover the R50 indoors using simulated natural scenes.

Summary

Chapter 5.

Outdoor Autonomous Flight

The indoor experiments with the six-degree-of-freedom testbed and the R50 helicopter demonstrated the effectiveness of the visual odometer in helicopter positioning under controlled laboratory environments. Implemented by a low latency and real-time vision matching, the odometer was repeatedly tested and verified off-board and then controlled the testbed R50 helicopter while connected to on-board power and sensing through tethers. The indoor tests proved that the helicopter system is air worthy for outdoor autonomous flight. However, outdoor untethered flight raises two critical issues not present during indoor experiments.

First, without protective tethers, the helicopter can fly out of control if the on-board vision system malfunctions. Loss of control can cause serious damage to the helicopter and possible human injury. To avoid this, a secondary system is necessary to maintain helicopter position and stabilize the helicopter in case of malfunctions. In addition to the safety issue, during less precise high altitude flight, a secondary positioning system can guide the helicopter to a predetermined destination where the visual odometer can then start accurately positioning the helicopter for high precision maneuvers. The secondary positioning system can also measure the performance of the visual odometer and the on-board helicopter controller during free flight experiments.

Second, outdoor autonomous flight requires an integrated system that combines the vision system with actuation and the secondary positioning system on-board the helicopter. The system must support a safety mechanism which allows a human operator to switch system positioning modes or take over the helicopter controls in case of system failure. Since it is difficult to pinpoint problems if the system is totally under computer control, human interfaces are especially critical in system development. An interface capable of human and computer control augmentation is developed for effective system performance evaluation.

This chapter presents the approach to developing an on-board helicopter system which integrates the visual odometer machine, secondary positioning using Global Positioning System, navigational sensors, actuator controls, and a safety system with human interfaces. The chapter concludes with the presentation and analysis of the helicopter positioning and control data collected during outdoor flight tests.

5.1 Secondary Positioning System

It is difficult to rely on vision alone for helicopter positioning throughout free flight in natural environments. The visual odometer relies on trackable ground features, but they may not be always available. For example, a bare and snow covered field does not have many features. When vision-based positioning encounters such featureless environments, a secondary system is necessary to stabilize the helicopter. Furthermore, a prototype visual odometer machine is prone to malfunctions during initial testing and requires a backup system for reliable positioning for helicopter control.

A global positioning system (GPS) receiver is an ideal secondary source of positioning to assist vision. A GPS receiver senses the range of multiple GPS satellites orbiting the earth and estimates global position by triangulation. GPS is especially well-suited for positioning aircraft since at higher altitudes, satellites are not usually obstructed by objects. An autonomous helicopter can rely on GPS for high altitude flight to a destination and then switch to vision-based positioning as it flies in close proximity of objects of interest.

5.1.1 GPS Positioning Method

The Global Positioning System is a ranging system which uses known positions of satellites in space to estimate unknown positions on land, sea, air and space. Satellite signals are continually tagged with the time they were transmitted so that when received, the signal transit period can be measured by a synchronized receiver. Apart from the determination of a vehicle's instantaneous position and velocity, GPS can precisely coordinate events in time using the satellite signals.

GPS uses “pseudoranges” derived from the broadcast satellite signals. The pseudorange represents the distance to a satellite and is derived by measuring the travel time of a coded signal from the satellite and multiplying it by its velocity. The clocks of the receiver and the satellite are employed to measure the signal travel time; since these clocks are never perfectly synchronized, instead of true ranges, “pseudoranges” are obtained where the synchronization error (denoted as clock error) is taken into account.




Figure 5-1. GPS Principle of Operation

Four unknown quantities must be determined to estimate the GPS receiver position. They include: the three coordinates of the desired position based on true satellite range, and the satellite and receiver clock error. Therefore, at least four satellites are necessary to compute the unknowns from equations of the form:

$$(x_i - x_0)^2 + (y_i - y_0)^2 + (z_i - z_0)^2 = (R_i - c\Delta T)^2 \quad (5-1)$$

where, as shown in Figure 5-1, (x_0, y_0, z_0) represent the vehicle position, (x_i, y_i, z_i) and R_i represent the known i th satellite position and measured pseudo-range, c is the speed of light, and ΔT represents clock error.

The GPS receiver constantly maintains the position of all visible satellites by using the navigation data provided by each satellite. Satellite ephemeris data provides parameters which describe each satellite's path based on orbital mathematics and Keplerian equations. (See [46] for a summary of equations and the satellite navigation message details.)

5.1.2 GPS Positioning Quality

GPS is a rapidly evolving technology for outdoor navigation and guidance. Currently, positioning accuracy using a single GPS receiver is within a hundred meter radius which is not suitable for precise helicopter control. This relatively low positioning accuracy is mainly caused by undesirable atmospheric effects and purposely degraded satellite signals. Recently, differential GPS correlation methods use ground station receivers to improve the positioning accuracy to about 1% of distance traveled by light in 1 microsecond or a 3 meter radius. The ground stations provide a correction reference for the atmospheric effects and helps in eliminating the effects of signal degradation.

New algorithms further improve positioning accuracy by taking advantage of the carrier phase of satellite signals. The carrier phase position resolutions were typically performed by off-line computations with powerful computers; however, high-end units on the market can perform the phase resolution "on-the-fly."

Carrier-phase methods employ a double differencing method to remove systematic errors ranging from signal degradation by the defense department to clock errors and atmospheric effects. These methods maintain the range differences of a number of satellites with respect to one satellite, referred to as the pivot satellite, from both a pre-surveyed ground station and the mobile GPS receiver to eliminate errors. Different strategies employ brute force search algorithms to Kalman filtering to determine the satellite signal carrier phase ambiguity in real-time, and take advantage of the phase of the carrier to estimate receiver position to within 10-20cm accuracy. (See [45] for a comprehensive presentation of GPS principles and operation and [46] for satellite navigation message format.)

5.1.3 GPS Evaluation Experiments

A carrier-phase differential GPS receiver, the NovAtel Rt20 [47], is employed for outdoor flight experiments. The unit is capable of 10-20cm positioning accuracy at 5 Hz provided the differential ground station is nearby (within 10 miles) and the helicopter is flown in an area with an unobstructed view of available GPS satellites.

Experiments are conducted to evaluate positioning accuracy and latency of the NovAtel RT20 carrier phase differential GPS receivers. A nearby (30 meter) differential station global position is surveyed by the Omnistar positioning system to provide corrections to the mobile receiver. The ground station is set up in an open area with no large, nearby obstructions. A choke-ring antenna is employed to help reduce multipath effects from reflected satellite signals.

The ground truth position is measured by an instrumented table, pictured in Figure 5-2, allowing GPS antenna movement in one direction and measuring ground-through position with an accurate string potentiometer. Experiments were conducted to compare the GPS data to ground-truth position at various rates of antenna oscillatory motion. As the graphs in Figure 5-3 illustrate, the experiments proved the performance of the GPS receiver. GPS positioning was accurate to 10 cm with 0.1 second average latency. The better than anticipated performance was due, in part, to the proximity of the ground differential station and the large number of available satellite signals. The GPS receiver had seven satellites in view during these experiments.




Figure 5-2. Instrumented Table for GPS Evaluation

On the average, a five minute initialization time was required for receiver phase ambiguity resolution down to the 10 cm positioning accuracy whenever satellite signals were interrupted. Therefore, to be effective, the GPS receiver must have an uninterrupted view of satellites during test flights away from large satellite-occluding objects such as buildings or hills.

In summary, the calibrated tests with the instrumented table proved that the GPS performance is adequate as a backup positioning system and for evaluating the vision-based positioning using a nearby ground station. To investigate GPS performance in flight, additional tests were performed under real flight conditions to determine if the main rotor blades occlude satellite signals. The GPS receiver and antenna were mounted on the helicopter as shown in Figure 5-4. Repeated human piloted flight tests proved that the R50 fiberglass and metal rotor blades did not occlude satellite signals. In addition, the receiver maintained satellite tracking even with the quick and significant attitude variations commonly exhibited by the helicopter during flight maneuvers.

Secondary Positioning System




Figure 5-3. Carrier Phase DGPS Performance



Figure 5-4. On-board GPS Receiver and Antenna

5.2 On-board Integrated System

Figure 5-5 shows a block diagram of the on-board integrated system built for outdoor autonomous helicopter flight experiments. The system is an extension of the indoor prototype visual odometer machine and retains the same modular point-to-point comm-port architecture.

The vision system is augmented by a GPS receiver, a flux-gate magnetic compass, and a laser rangefinder for redundant helicopter position estimation. Comm-port interfaces were developed for each of these sensors for easy integration and compatibility with all system components. The system's real-time controller implements the PD servo loops, and commands the helicopter actuators through a safety circuit. Human pilot interfaces are incorporated into the safety circuitry for computer augmented helicopter control. This section describes the operation of the on-board integrated system components.

On-board Integrated System




Figure 5-5. On-board system block diagram

5.2.1 Vision Processor

The visual odometer machine, calibrated and tested using the indoor testbeds, is integrated on-board with few small modifications. Taking advantage of the modular system architecture, more DSP elements are incorporated into the system for additional sensor data acquisition and integration with vision. The tasks of attitude compensation and coordinate transforms are divided into two **DSP C40s**. One simply performs synchronized filtering of attitude data while the other receives position data from the remaining sensors, including a **GPS** receiver, a laser rangefinder, and a magnetic compass. In addition, a video transmitter is integrated on-board to send processed images to the ground for monitoring.

5.2.2 GPS Receiver

As described in Section 5.1, experiments verified that the carrier phase differential GPS receiver, NovAtel RT20, is an accurate and reliable source of positioning for free flight experiments. The GPS receiver performs satellite tracking and carrier phase ambiguity resolution using transputers which provide high-speed external links for system expansion. With cooperation from NovAtel Communication, Ltd, these links are internally accessed to develop a low-latency comm-port interface to the C40s and the remainder of the system. For synchronized operation, the sync generator triggers the GPS receiver every 12 images or 5 **Hz**. For performance evaluation, on-board wireless modems receive differential corrections and transmit position logs during helicopter flight to the ground station.

5.2.3 Compass

To remedy the accumulating drift of the directional gyroscope, a North seeker digital magnetic compass is integrated into the system. The compass uses a toroidal flux-gate sensing element free floating in an inert fluid to keep the sensing element horizontal. The compass manufacturer, KVH Industries, provided a more viscous fluid filled sensing element to help reduce vibration effects on heading measurements. In spite of this precaution, the compass heading data is quite noisy and requires extensive filtering.

The digital output from the compass is sampled with the same external trigger as the GPS receiver, and a hardware interface is developed to transmit the heading data to other system components via a comm-port.

The difference between the compass heading and the yaw gyroscope is filtered to correct the drift of the yaw gyroscope. The filter is a low-pass with time constant of 5 seconds. **A C40** acquires compass data, performs low-pass filtering, and transmits corrected heading data to other system modules.

5.2.4 Laser Rangefinder

For redundant height measurement, a laser rangefinder is integrated on-board. The laser rangefinder, manufactured by Yamaha Motor Company, has a 20 meter range with a 20 Hz measurement frequency. Preliminary experiments demonstrated small (1-2%) range variations with reflective surface color which are not explicitly modeled for helicopter height estimation. Data are clocked out of the sensor serially and a comm-port bridge that includes shift registers is developed to transmit range data to other system modules.

5.2.5 Real-Time Controller

The real-time controller, integrating an MC68040 microprocessor, implements the helicopter control system. Composed of several PD servo loops, the controller receives position and attitude estimates from the integrated vision system at field rate (60 Hz) and controls the helicopter by transmitting commands to the on-board actuators and augmentation systems.

The real-time controller also bootstraps and configures the **C40** network as well as performs several other functions during system operation. The network interface of the controller provides access to mass storage for data logs and initialization. The controller provides user interfaces for run-time system configuration and supports 128 digital I/O lines for interfaces to safety circuits, the actuation system, and two system comm-ports.

5.2.6 Actuator Control

The Yamaha R50 is designed to be remotely controlled with an RF transmitter during flight experiments indoors. The transmitter provides control sticks and trims for remote control by human pilots and incorporates decoupling curves to reduce cross coupling of helicopter control inputs. Off-board computing generated analog signals in place of the stick potentiometers of the transmitter to remotely control the R50 for the indoor test flights.

To control the R50, the four stick inputs of the transmitter are sampled and converted to five actuator positions. Three actuators control the main rotor collective and cyclic pitch as shown in Figure 5-6. Two side actuators move in opposite directions to produce lateral pitch, while the middle actuator controls longitudinal pitch. All three actuators move in parallel for regulating the collective pitch. The remaining two actuators (not shown) control tail rotor pitch and engine throttle.




Figure 5-6. Main Rotor Actuators

The transmitter digitally encodes the desired actuator locations which it sends to an on-board receiver. The receiver produces pulse width modulated signals at 50 Hz to motor controllers which move each actuator. For autonomous operation, the integrated system must directly generate pulse width modulated signals to control the five actuators. In addition, the actuator coupling terms must be resolved on-board.

Pulse width decoders and encoders were developed using AMD Mach **435** complex PLDs for on-board helicopter actuator control. To command helicopter actuators, the real-time controller was interfaced to the PLDs. The controller internally stores the actuator mixing tables to convert desired collective and cyclic pitch terms to the three main rotor actuation positions. Changes in throttle setting and tail rotor are also made by the controller using an internally stored piece-wise linear tables.

For augmented control, the five actuator positions sent from the remote transmitter are decoded using inverted actuator mixing tables to determine human collective, cyclic, and tail rotor commanded positions. These inputs are then normalized and used as augmented control inputs to regulate helicopter position or velocity using the on-board control system.

5.2.7 Safety Circuit and Human Interfaces

The indoor testbeds provide protection and safety by limiting helicopter travel and measuring ground truth helicopter position for stable recovery from out-of-control flight. The absence of this protection for outdoor free flight requires on-board safety circuits and human interfaces to minimize accidents.

A “heartbeat” mechanism is developed to detect system failures during autonomous operation. The heartbeat is a periodic signal that indicates system health. The real-time controller is configured to periodically monitor all on-board systems, including the vision system, **GPS**, actuator locations, and battery voltage, and to generate a heartbeat by outputting a pulse on one of its external digital lines. In case of a malfunction, the heartbeat signal remains unchanged for an extended period. During normal system operation, the heartbeat signal reports system processing duration and frequency.

The safety circuit times the heartbeat intervals; if the heartbeat signal is stuck in one state, the circuit detects a system fault. In addition, the safety circuit includes multiplexors for switching actuator control between the on-board RF receiver, carrying out human pilot’s actions, and the on-board computer. An extra channel of the receiver, set by the remote transmitter, switches between the control modes.

5.3 Autonomous Helicopter

For outdoor flight, the chassis for supporting the cameras and attitude sensors for the indoor R50 is integrated on-board another R50 helicopter. Figure 5-7 shows the R50 on-board system components.




Figure 5-7. R50 On-board System

The on-board visual odometer machine and real-time controller are housed in the computing cage which is mounted below the helicopter fuselage. To the immediate right of the cage are the laser rangefinder and bridge assembly, and the video transmitter. Small hardware circuits implementing the sync generator and sensor bridge modules are mounted in the front of the cage. On either rear skid, there is a wireless modem for GPS differential correction and ground telemetry. The main system battery is mounted in the front of the helicopter to balance out the weight of the GPS receiver mounted

on the tail. The two ground pointing cameras are mounted on the main system chassis as in the indoor R50 helicopter. Critical components such as the attitude sensors and the safety circuit are well protected inside the helicopter frame and are not visible in the figure.

The following subsections describe the major components of the on-board implemented vision-based helicopter control system.

5.3.1 Weight and Power

The weight of all on-board equipment, about 18 Kg, is less than the 20 Kg payload of the R50. The power dissipated by the system is on the order of 180W and the computer operation is possible for 12-14 minutes using inexpensive (\$20) 7 Ah lead-acid batteries. Silver-zinc batteries of the same weight could power the system for 1 hour but are not used because they are an order of magnitude more expensive. For redundancy, a separate lower capacity battery powers the receiver, actuators, and the safety circuit for reliable remote piloting in case the computing battery power is drained in free flight.

5.3.2 On-board Computing

Figure 5-8 shows the on-board VME computing cage and the hardware modules. The ribbon cable connections are identical comm-port links between different components. Image A/D, D/A, and convolution modules are merged into one image capture and preprocessor printed circuit board, shown in Figure 5-9. Using comm-port connections, the board can transmit digitized images and receive processed images for display. Image synchronization is performed with intelligent state machines to eliminate additional external connections. The comm-port interfaces provide internal buffering and termination to guard against data corruption from signal reflection or other noise sources; transfers are rated at 10 Megabytes per second. Images can be digitized from up to 4 multiplexed NTSC/PAL camera inputs by the integrated A/D module. The module incorporates all programmable image sampling and blanking features described in Chapter 3. In addition, the A/D module provides access to an internal synchronized video data bus for external synchronization and can pass images through an 8x8




Figure 5-8. On-board Computing

image convolver for image preprocessing. In addition, processed images can be displayed through an RGB pseudo-color display driver. An input comm-port interface transfers the images to the display driver and provides the proper synchronization for video screen refreshing. Processed images can be overlaid on top of captured images from the A/D module. A VME bus bridge module is also incorporated for system initialization and external communication.




Figure 5-9. Computing and hardware modules

5.4 Outdoor Flight Experiments

The next step following the development and integration of system components was outdoor autonomous flight experiments. The experiments were conducted in an isolated flight site with grassy terrain and an open view of GPS satellites.¹ The site's terrain provided enough features for establishing visual lock by the visual odometer, but was not locally flat. The helicopter tests were conducted at the summit of gently sloping hill.

5.4.1 Experimental Setup

The Navlab I [48] autonomous land vehicle was modified to house the helicopter for transportation and to provide a mobile ground computing platform for system development and evaluation. The interior of the vehicle is shown below in Figure 5-10.




Figure 5-10. Helicopter Vehicle Interior

1. The flight site is the property of William Wittey who kindly permitted the use of his farm in Zelienople, PA for helicopter flight tests.

The vehicle also serves as a mobile GPS differential base station. The global location of the experiment site is measured by another GPS differential station each time the vehicle is driven to the field and initialized as the local differential station for helicopter positioning. The vehicle has an on-board GPS receiver to transmit corrections to the helicopter during experiments. Figure 5-11 shows the vehicle setup during helicopter flight.

A pair of wireless radio modems transmits GPS corrections to the helicopter and receives helicopter position and status during flight tests. A video receiver captures processed images transmitted from the helicopter for viewing inside the vehicle.

The vehicle also incorporates an on-board power system, mass data storage, and a local area network to power and bootstrap the helicopter computers. While the helicopter is on the ground, two cables between it and the vehicle provide power and communication. The power system keeps the helicopter computes on-line and charges its batteries before each flight.




Figure 5-11. Vehicle setup for helicopter flight

5.4.2 Experimental Approach




Figure 5-12. Ground testing

The helicopter was tested on the ground before computer controlled flight experiments were undertaken. As shown in Figure 5-12, the helicopter **was** strapped on to a wheeled platform and moved around on the field to test position and velocity sensing and actuator compensation.

The grassy terrain proved rich enough in features to lock on to by the visual odometer machine and the GPS system was operational, provided its antenna was clear of obstructions. Running all systems in parallel, the vision-based lateral and longitudinal positioning were compared with estimates from the GPS receiver to verify accurate position estimation. At times, the platform was lifted and rotated to ensure proper attitude compensation and height sensing by the visual odometer machine.

Two warning lights, made up of LED grids as pictured in Figure 5-13, are mounted on the helicopter to indicate computer control and status to a safety pilot on the ground. The lights are controlled by the safety system which detects failures using the computer heartbeat. Different patterns are generated on the LED grid to indicate power, vision, or GPS failures. Smaller indicator lights to the right of the warning lights show battery status while the helicopter is on the ground.

The safety system also regulates graceful switching in and out of computer control flight while the human pilot flies the helicopter. An “auto” switch on the pilot’s control transmitter switches control between the human and computer. Current human stick and trim locations from the transmitter are decoded and sent to the control system to properly bias the PD controllers during computer control engagement.




Figure 5-13. Safety system and control interfaces




Figure 5-14. High altitude flight using GPS

After ground testing trials, the helicopter was flown by the human pilot to compare vision and GPS positioning during actual flight. Lateral and longitudinal positioning were tested under different conditions and attitude variation. Stereo height measurement was compared with height sensed by the laser rangefinder and global altitude measured by the GPS. As will be shown in the next section, the data from vision, GPS, and laser rangefinder proved to be consistent enough to warrant computer control experiments.

Initial computer control trials were performed at high (~15 m) altitudes to allow the safety pilot ample time to manually override the computer in case of sudden loss of control. Figure 5-14 shows the helicopter during high altitude flight tests. Latitude and longitudinal control were tested first by mixing human control for height and heading with the computer commands. Helicopter control loops were conservatively tuned for low precision flight with GPS positioning as the back up in case of vision system failure. Slowly, heading and height control were enabled as the computer control proved effective in stabilizing the helicopter.




Figure 5-15. Low altitude precise flight using vision

Relying on the backup control system, more dangerous low (3-5 m) altitude tests were performed using vision-based feedback. Again, longitudinal and lateral control were tested before height control based on stereo vision was switched on. The laser rangefinder was actively used to check the consistency of stereo range from vision.

Figure 5-15 shows the vision-based helicopter control flight at **4-5** meters off the ground and two processed images a few seconds apart as transmitted from the helicopter. The images are blurry due to preprocessing with a Gaussian filter. The odometer successfully retained visual lock on poorly contrasting images taken from the grassy terrain under harsh vibration and varying heading angle during long hovering intervals. The processed images show the two tracked templates as the odometer locked onto a dry grass patch. The helicopter heading changed between the images but the grassy path remained trackable. The on-board PD controller precisely hovered the helicopter in one spot within 0.5 meter of the desired location during these intervals. Data from these experiments are presented in the next section.

5.4.3 Experimental Results

To test all of the on-board positioning systems, the helicopter was flown in an approximately circular pattern by the human pilot. The starting point of the pattern was at the summit of a small hill with significant (20-30 degrees) sloping terrain. For comparison, data were logged from vision, GPS, and the laser rangefinder in parallel.

5.4.3.1 Position Estimation

Figure 5-16 shows vision and GPS data collected while the helicopter was flow in the circular test path. The two dimensions are the X and Y axes of the local navigation frame, with Y pointing North and X East. Vision and GPS estimates matched accurately in the Y dimension but there was a consistent 20% difference in the X dimension. This discrepancy was attributed to the downhill grade of the terrain in this dimension. This significant grade violates the visual odometer's flat ground assumption and adds a systematic bias to position measurement. No significant drift was detected in the circular test run.




Figure 5-16. Test path for positioning evaluation




Figure 5-17. Lateral (X) positioning

To demonstrate the effects of the slope more precisely, Figure 5-17 and 5-18 show the lateral position (**X**) and longitudinal position (**Y**) with time. **GPS** and vision positioning match within 1.7 meters laterally and 0.7 meters longitudinally for the duration of the circular flight. In addition, it is important to observe the smoothness of the vision data, updated at 60 Hz, compared to the GPS data, updated at 5 Hz.




Figure 5-18. Longitudinal Positioning

Helicopter height measured by stereo vision, laser rangefinder, and GPS are plotted in Figure 5-19 below. Range measured by stereo and laser height matched within **30** cm. The GPS reports global altitude, which is biased to match stereo and laser at the beginning of each test flight. The GPS height estimate is not affected by ground slope and therefore may not always match stereo and laser range data at all times. For the flight test shown below, the GPS height is within **40** cm of stereo and laser rangefinder data.




Figure 5-19. Helicopter height

5.4.3.2 Velocity Estimation

Figures 5-20 and 5-21 show the helicopter lateral and longitudinal velocity measured by vision and GPS during the test flight. Velocity measured by image center pixel velocity is noisy, but matched the GPS velocity within with 50 cm/s. Vision-based noise is mainly caused by camera vibration and is limited to approximately 25 cm/s for both lateral and longitudinal directions. This small noise was not considered significant and was successfully employed by the on-board control system for helicopter stabilization.




Figure 5-20. Lateral velocity




Figure 5-21. Longitudinal velocity

5.4.3.3 Computer Control Trials

For computer-controlled hovering, the helicopter was piloted off the ground by the safety pilot to an altitude of 4 meters and system control was switched to computer. Figures 5-22, 5-23, and 5-24 show the helicopter lateral, longitudinal, and height control accuracy using vision-based positioning feedback. The PD control system successfully hovered the helicopter within 0.5 meters of a desired location in the air.




Figure 5-22. Lateral position control accuracy




Figure 5-23. Longitudinal position control accuracy

All control axes exhibited slow oscillations with a 2-3 second period and the height controller consistently showed negative steady-state error due to the PD control system lack of integral control. The helicopter quickly drifted out of control once the computer control was switched off and the human pilot took over control for landing.




Figure 5-24. Height control accuracy

5.5 Summary

This chapter presented the experimental trials and system integration for outdoor autonomous helicopter flight. The on-board system integrates redundant position sensing capabilities for safe helicopter flight outdoors. A secondary positioning system is examined and integrated using carrier phase GPS. The GPS positioning is shown to be sufficiently accurate for low precision helicopter flight. For more redundancy, a laser rangefinder measures height in parallel with stereo vision and GPS to prevent loss of helicopter altitude due to system failures.

An elaborate safety system monitors system health and provides smooth computer control transitions from human-controlled flight. The safety system uses a heartbeat mechanism to detect failures in system components, including vision, GPS, control, and on-board power. In addition, the safety system can mix human and computer control for incremental tests with partial computer control. The on-board system is shown to stably hover the helicopter within 0.5 meters using vision as the primary source of position and velocity feedback.

Chapter 6.

Conclusions and Future Work

The work presented in this dissertation has demonstrated an airworthy autonomous helicopter with on-board vision for guidance and stability. This research shows that, when effectively integrated, vision-based object trackers and position estimators are capable of stabilizing highly responsive and difficult *to* control plants such as helicopters. In addition, this work has shown how close integration of powerful image processing elements with external sensors can be achieved through a new vision machine architecture for real-time and low latency image processing.

System evaluation plays a significant role in successful development of complex integrated systems such as autonomous helicopters. The research presented in this dissertation has demonstrated the advantages of an incremental design approach in which different system components, including position sensing, actuation and control, and human interfaces are independently evaluated by an array of innovative helicopter testbeds. The testbeds allow calibrated experiments by sensing helicopter ground-truth position, and provide safety by limiting helicopter speed and travel area.

This chapter summarizes the accomplishments and the future directions of this work in the areas of vision-based position estimation and low-latency vision machine architectures.

6.1 Accomplishments

The two accomplishments of the presented work include an autonomous vision-guided helicopter and a new vision machine architecture for real-time and low latency image processing. This section presents a summary of these accomplishments.

6.1.1 An Autonomous Vision-Guided Helicopter

This research has developed an autonomous helicopter guided and stabilized by a visual odometer. The odometer takes advantage of the abundant features in natural scenes to lock on to arbitrary ground targets for measuring helicopter displacement and altitude. The odometer maintains lock on two 40 by 40 pixel image segments or templates and actively tracks them at field rate (60 Hz) in parallel by high speed matching. When necessary, the odometer scales, rotates, and normalizes the templates in real-time for reliable tracking under abrupt attitude and height variations as well as the harsh vibration common to helicopters. The odometer also performs pixel velocity measurement and stereo image processing for helicopter velocity and height estimation. Indoor and outdoor flight tests have demonstrated vision-based position accuracies of 3- 10 cm during helicopter hovering.

The odometer is realized on-board an air worthy autonomous helicopter integrating custom-built vision processing, ground-pointing video cameras, a GPS receiver, a laser altimeter, a fluxgate compass, human interfaces, safety systems, telemetry, and PD-based control and actuation. (See Table 6-1 to 6-3 for helicopter specifications.)

The helicopter's first stable autonomous flight based on visual feedback was demonstrated on October 17, 1995. The helicopter positioning and control system successfully stabilized the helicopter within 0.5 meter of a desired location under different atmospheric and lighting conditions for approximately fifty test flights.

Table 6-1. Helicopter Sensors

Sensor	Manufacturer	Specifications
Pair of Video Cameras	Sony XC75	NTSC, 1/1000 shutter, 752x582 CCD pixels, Externally synchronized non-interlaced single field output mode
Laser Rangefinder	Yamaha	0.1-20 m range, 1-2 % accuracy
GPS Receiver	NovAtel RT20	10-20cm accuracy at 5 Hz
Compass	KVH	0.5 degree resolution
Gyroscopes	Gyratron	0.2 degree resolution, 2-4 degrees per minute drift rate

Table 6-2. Helicopter On-Board Computing Hardware

Function	Manufacturer	Configuration
(7) DSP processors	TIM-40 Standard	50 MHz, zero-wait-state SRAM 1MB/ 1 MB global/local memory
Convolver (custom)	GEC-Plessey	(8x8) 10 MHz loadable mask, output multiplexor
A/D (custom)	BrookTree 252	4 multiplexed NTSC/PAL inputs, 256 LUT, programmable reference
D/A (custom)	BrookTree 473	pseudo-color, RGB output, sync on green
Real-Time Controller	Motorola MC68040 (MVME 162)	4MB DRAM, 25 MHz
Sensor Bridge (custom)	HP quad. decoder, AMD MACH 220 complex PLD state machines	10channel quad. encoder inputs, 4 independent A/D, 64 digital I/O lines, configurable data packets

Table 6-3. Autonomous Helicopter Specifications

Helicopter Model	Yamaha R50 -L12 (YACS)
Payload	20 Kg
Takeoff weight	67 Kg
Engine	2 cycle, 98 cc, 12 HP
Overall Length	3.5 meters
Main Rotor Diameter	3.1 meters
Body Length	2.7 meters
Autonomous Flight Duration	15 minutes
On-Board Power	7AH lead-acid battery

6.1.2 Real-time and Low Latency Vision

This dissertation has presented a new vision machine architecture for low-latency image processing. The architecture proved effective in efficiently integrating powerful image processors with external sensors to build a compact visual odometer machine flown on-board the prototype helicopter. Based on the philosophy that no single vision machine is suitable for all applications, the architecture provides a reconfigurable framework for designing vision systems tailored to specific applications. Processing capabilities are captured in modules which communicate via a uniform and high speed set of point-to-point links. Uniform communication means that all modules are electrically compatible and can be interconnected in different configurations for different tasks. Evidence of the architecture's configurability is that it has been used to develop a number of vision machines for medical image processing [51] in addition to commercially marketed vision systems [55] for robotic applications. In particular, the Kirin Brewery Company in Japan is supporting future research based on the architecture. The objective of the research is the development of vision-based factory inspection machines. As a prelude to this research, a prototype inspection machine was developed using the same basic modules as those of the visual odometer machine. Shown in Figure 6-4, the machine, developed in two months, is capable of detecting small imperfections in bottles and then rejecting them from the bottle conveyers at a maximum rate of 1200 bottles per minute.

Accomplishments



Figure 6-4. Bottle inspection machine

6.2 Future Work

The visual odometer's capabilities can be expanded to track particular objects of interest to perform autonomous search missions by the helicopter. Instead of locking on to an arbitrary object, the visual odometer can be guided by a secondary object detector to establish a lock onto a particular object and track it thereafter. In addition, the odometer could be configured recognize known environments such as a particular landing site by tracking known ground features to estimate helicopter position. In the following subsections, these two scenarios are investigated through preliminary experiments.

6.2.1 Object Tracking

Similar to the visual odometer, an object tracker can detect and track a particular object by template matching. A major difficulty of this approach is that object orientation is unknown and therefore all possible object orientations must be searched for in order to locate the object. Methods such as Karhunen-Loeve [49] expansion can be used to reduce computational complexity and storage of necessary templates. Another problem stems from varying helicopter altitude which will change the size of the object in the image. Close regulation and measurement of helicopter altitude is necessary to further reduce the complexity of this search.

Rotated template images differ slightly from each other and are highly correlated. Therefore, the image vector subspace required for their effective representation can be defined by a small number of eigenvectors or eigenimages. The eigenimages which best account for the distribution of image template vectors can be derived by Karhunen-Loeve expansion [50].

Such principal component analysis methods can dramatically reduce the number of templates required to locate a particular object template. In fact, an experimental system based on the work presented in [51] and [57] could detect a section of a toy Jeep which was placed under the indoor testbed helicopter within ± 80 degree orientation discrepancy using only 41 rotated (32x32) templates. The Jeep and the tracked template are shown in the two images of Figure 6-1.




Figure 6-1. Object Detection Experiment

The processing frequency for searching the entire image is 2.5 Hz, using three **C40** processors. Although relatively slow, this object search process can provide a resetting mechanism for the lower visual odometer which cycles at field rate. The detected object, the Jeep in this case, can serve as one of the tracked templates of the visual odometer, therefore positioning the helicopter relative to the object. This relative positioning can then be employed for aerial tracking while the helicopter is controlled in relation to the object.

6.2.2 Helicopter Positioning using Known Environments

In many common applications helicopters must fly close to known objects or land on predetermined landing pads. In most cases, it is reasonable to assume that special markings or features can be placed in view of the helicopter cameras to provide feedback for automatic close proximity hovering, landing, and take off. Aside from especially painted landing pads, it is desirable to use existing easy-to-detect features which may be dispersed irregularly, but at known positions, for relative helicopter position estimation. The traditional approach to this problem is to back project known **3-D** world features on the image plane and match them with the 2-D image features to estimate 3-D camera pose [52]. Other methods use projective transforms [53] and geometric invariants [54] for direct pose recovery from image features.

Following the back projection approach, experiments were conducted with easy-to-detect known ground features for helicopter positioning. A number of easily detectable ground features, white dots, are placed at known locations on a black background under the testbed helicopter as shown in Figure 6-2.




Figure 6-2. Experiment with known ground features

A feature detector located the white dots in the image by image thresholding. Figure 6-3 shows the detector's input and output images. The output image shows squares around detected features after lens calibration. The **squares** around features near the image boundaries do not line up perfectly with the raw features due to the larger lens distortion in these areas. The known ground features were then projected onto the image plane for matching. The proximity of image features to projected features was used as a matching criteria. In experimental trials, the closest 5-10 features were selected for position updates.




Figure 6-3. Feature detection

The transformation from the world to the image plane is nonlinear but continuous and well behaved, allowing linear functions to locally approximate it. Using the current helicopter location, a Jacobian was constructed to iteratively approximate the change in helicopter position and attitude given image feature displacement. Linear extrapolation by Newton-Raphson was employed to update helicopter position between successive images.

Experimental trials demonstrated that the hovering helicopter's movement in one field was small enough ($< 10\text{ pixels/feature}$) to yield satisfactory ($< 5\text{ cm}$ accuracy on average) position update with one or two iterations at **60 Hz** processing frequency. Positioning accuracy proved satisfactory for future work in helicopter landings and takeoffs,

6.3 Concluding Remarks

Robot helicopters are beginning to show their potential in an increasing number of applications. A small autonomous helicopter can perform aerial surveillance by transmitting images from on-board cameras taken at different altitudes and vantage points. This “eye-in-the-sky” capability, together with precise maneuvering of the helicopter, can provide a comprehensive picture of the environment central to scouting operations, site inspection, and movie production. Figure 6-5 shows a dangerous live power transmission line inspection and repair by an electrical worker sitting on a human piloted helicopter.’ **As** dangerous as they are, such applications are proving to be sufficiently cost effective to risk human lives, thus making a strong case for the necessity of autonomous robotic helicopters.




Figure 6-5. Electrical wire inspection

This dissertation has presented promising results in helicopter control using vision and it is my ultimate goal to apply these results to the development of robot helicopters for future real world applications.

1.Used with permission of Beyond 2000 television program.

Appendix A.

Six Degree-Of-Freedom Testbed

Helicopter control is difficult; careful experimentation is essential to build a working prototype robot helicopter. For calibrated experiments, research described in this dissertation led to the development of a six Degree-Of-Freedom (6-DOF) testbed for safe indoor helicopter flight. The testbed measures ground truth helicopter position and attitude as well as working as a safety device for preventing crashes and out-of-control flight.

As shown in Figure A- 1, the testbed supports an electrical model helicopter which can fly freely in a cone-shaped volume six feet wide and five feet tall. The helicopter is fastened to six poles by rods which are free to move through two-degree-of-freedom (2-DOF) joints. The joint angles are measured by shaft encoders and are used by the computer to calculate the helicopter's ground truth position and attitude during flight tests.

An important issue concerns the effects of the testbed components on the helicopter dynamics in free flight. To minimize inertial variations, the testbed is built from light-weight metal and composite materials custom-designed and fabricated to minimize weight and friction. Minimizing friction is especially critical; friction has a tendency to significantly dampen helicopter movement. This gives a false sense of control system stability on the testbed compared to untethered flight which is undesirable.




Figure A-1.6-DOF testbed

A.1 Testbed Helicopter

The testbed helicopter is an electric model, Kalt Whisper, as shown in Figure A-2. The helicopter is modified in several respects for testbed operation. Its power plant, a small DC motor, is built for high power, using a custom wound armature to operate more efficiently at higher voltages. The motor dissipates 0.5 HP and provides enough power to lift the helicopter, the rods, and the on-board sensors. Because of its small size and high power dissipation rate, the motor is actively cooled by forced air to prevent premature failure. In addition, the motor voltage is regulated using a servo loop to keep rotor revolutions as constant as possible with varying helicopter loads.




Figure A-2. 6-DOF testbed helicopter

A stereo pair of light-weight camera heads is mounted on the helicopter for vision along with vertical and directional gyroscopes for attitude sensing. The on-board sensors are shielded from helicopter vibration by a small suspension system built into the helicopter body. The suspension also dampens helicopter oscillations from the energy stored in the support rods just before takeoff.

A.1 Testbed Support Structure

The testbed helicopter is fastened to a planar light-weight structure made of three equal length aluminum tubes meeting at one point, designated as the testbed origin, and spread equally to span the area of an equilateral triangle as shown in Figure A-3.




Figure A-3. 6-DOF testbed helicopter mounting structure

Mounting sites at the edges of the aluminum structure connect to support rods which travel through bearings at each 2-DOF joint, as illustrated in Figure A-4. Two rods are connected to each mounting site and three of their four joint angles, (α, β, θ), are measured by shaft encoders to determine the 3D position of each mounting site. The helicopter position and orientation are then computed from the measured mounting site 3D locations.




Figure A-4. Testbed geometry

The support rods, 2-DOF joints, and helicopter catching mechanism are shown in Figure A-5. The support rods are made from graphite arrows generally used for archery. The rods move through frictionless air bearings at each joint and are terminated by a spring-loaded stopper. The stopper cushions collisions as the helicopter reaches the rod's extreme. A catcher mechanism connects each rod to the mounting site using support pins. The catcher mechanism stops the helicopter as it falls and the replaceable support pins absorb the impact energy by bending to prevent the rods from fracturing.





Figure A-5. Testbed joint and catching mechanism

Appendix B.

Lessons Learned

I have learned several valuable lessons while developing an autonomous vision-guided helicopter at Carnegie Mellon University. Others who strive to build complex robotic systems may benefit from these lessons:

1. *Follow an incremental and systematic design and evaluation approach:* I started my graduate work by diving in and performing a series of outdoor experiments with model helicopters. I equipped a model helicopter with an array of sensors and performed a number of outdoor experiments which produced little concrete results. I was plagued with malfunctions, did not have a clear idea of system deficiencies, and most importantly had no way of quantifying system performance. Learning the hard way, I began to follow an incremental and systematic approach by building a number of indoor calibrated testbeds to design and evaluate each system component independently. The testbeds significantly reduced the chance of failure by verifying each system component before it was integrated into the final system. In spite of the seemingly extra amount of effort, this incremental approach to building complex integrated systems appears to be the most favorable in the long run. This is especially true for helicopter control experiments in which one malfunction can cause a crash with major loss of time, resources, and safety.

2. *Carefully verify that the traditional approach to the problem is not sufficient before developing a new theory or approach:* I learned this lesson while developing a helicopter control system. The highly unstable nature of helicopters led me to propose research in new adaptive control methods based on fuzzy logic and unsupervised learning techniques for my thesis work. To investigate the training capabilities of these methods, I designed a simple PD-based controller to stabilize and collect performance data from a model helicopter. To my surprise, the classical PD-controller worked quite well for hovering and low speed maneuvers which were the main modes of operation I was envisioning. The favorable performance was largely due to the high rate of positioning feedback and eliminated the need for rigorous system modeling. Therefore, instead of controls and system modeling, I focused on high-speed helicopter state estimation.
3. *Reach a well-dejined ultimate goal through a series of manageable sub-goals:* Throughout my graduate work, I maintained a single goal of stabilizing helicopters with on-board vision. I discovered that I reached this goal by reaching a number of short, manageable, and well-defined sub-goals. The sub-goals maintained my high level of energy and reduced the large burden of what at times felt like an impossible task. Without clear sub-goals, it is easy to lose focus and attack interesting research problems for the sake of “research” alone.
4. *Build what you need if you can't find it or can't afford it:* I wasted valuable time and resources by trying to use what was available to me instead of what I needed. I learned the hard way that just because a certain component is available it should not necessarily be part of the system. Instead of spending time on adapting available technology, I discovered that I could build certain components, such as image processing hardware, customized to my specifications and at a lower cost.
5. *Give equal importance to every system component:* The components of an integrated system can be thought of links in a chain. Every component is assembled for a purpose and must be treated with equal respect for successful operation of the entire system. At times, I have paid a high price for impatient implementation and improper use of seemingly unimportant system components. Successful systems are built with patience and a consistent level of craftsmanship for every single component.

Bibliography

- [1] R. W. Prouty. *Helicopter Performance, Stability, and Control*. Robert E. Kreiger Publishing Co., 1990.
- [2] W. Johnson. *Helicopter Theory*. Princeton University Press, 1980.
- [3] J. Kaletka and W. von Grunhagen. Identification of mathematical derivative models for the design of a model following control system. *Vertica*, 13(3):213–228, 1989.
- [4] D. R. Downing and W. H. Bryant. Flight test of a digital controller used in a helicopter auto-land system. *Automatica*, 23(3):295–300, 1987.
- [5] A. Yue, I. Postlethwaite, and G. Padfield. H Infinity design and the improvement of helicopter handling qualities. *Vertica*, 13(2): 119–132, 1989.
- [6] U. Christen, M. F. Weilenman, and H. P. Geering. Design of H_2 and H infinity controllers with two degrees of freedom. *Proceedings of the 1994 American Control Conference (Cat. No. 94CH3390-2)*, 1994.
- [7] M. F. Weilenman, U. Christen, and H. P. Geering. Robust helicopter position control at hover. *Proceedings of the 1994 American Control Conference (Cat. No. 94CH3390-2)*, 1994.
- [8] R. A. Hess and Y. C. Jung. An application of generalized predictive control to rotorcraft terrain-following flight. *IEEE Transactions on Systems, Man, and Cybernetics*, 19(5), 1989.
- [9] B. K. Townsend. The application of quadratic optimal cooperative control synthesis to a CH-47 helicopter. *Journal of the American Helicopter Society*, pages 33–44, January 1987.
- [10] R. A. Hess and K. K. Chane. Preview control pilot model for near-earth maneuvering helicopter flight. *Journal of Guidance*, 11(2):146–152, April 1988.
- [11] D. F. Enns. Multivariable flight control for an attack helicopter. *IEEE Control Systems Magazine*, April 1987.
- [12] M. Sugeno, T. Murofushi, J. Nishino, and H. Miwa. Helicopter flight control based on fuzzy logic. *Fuzzy Engineering toward Human Friendly Systems*, 1991.
- [13] H. Ekerol and D. Hodgson. A machine vision system for high speed object tracking using a moments algorithm. *Mechatronics*, 2(6):555–565, Dec. 1992.
- [14] H. Inoue, T. Tachikawa, and M. Inaba. Robot vision system with a correlation chip for real-time tracking, optical flow and depth map generation. *Proc. 1992 IEEE Int. Conf. on Robotics and Automation*, 2:1621–1626, May 1992.

- [15] N. Papanikolopoulos, P. Khosla, and T. Kanade. Visual tracking of a moving target by a camera mounted on a robot : a combination of control and vision. *IEEE Trans. Robot. Autom.*, 9(1):14–35, Feb. 1993.
- [16] A. Rizzi, L. Whitcomb, and D. Koditschek. Distributed real-time control of a spatial robot juggler. *Computer*, 25(5): 12–24, May 1992.
- [17] R. Luo and R. Mullen Jr. Combined vision / ultrasonics for multi-dimensional robotic tracking. *Sensor Fusion: Spatial Reasoning and Scene Interpretation; Proceedings of the SPIE*, 1003:113–122, 1989.
- [18] C. Harris and C. Stennett. Rapid - a video rate object tracker. *BMVC90 Proceedings of the British Machine Vision Conference*, pages 73–77, Sept. 1990.
- [19] C. E. Thorpe. *Vision and Navigation: the The Camegie Mellon Navlab*. Kluwer Academic Publishers, 1990.
- [20] D. Pomerleau. ALVIN¹: An autonomous land vehicle in a neural network. In D. Touretzky, editor, *Advances in Neural Information Processing Systems I*. Morgan Kaufmann, 1989.
- [21] D. Pomerleau and T. Jochem. Rapidly adapting machine vision for automated vehicle steering. *IEEE Expert*, 11(2), 1996.
- [22] M. Hebert and T. Kanade. 3-D vision for outdoor navigation by an autonomous vehicle. In *DARPA Workshop on Image Understanding*, 1988.
- [23] M. Hebert, T. Kanade, and I. Kweon. 3-D vision techniques for autonomous mobile robots. Technical Report CMU-RI-TR-88-12, Carnegie Mellon University, The Robotics Institute, August 1988.
- [24] E. Dickmanns and B. Mysliwetz. Recursive 3-D road and relative ego-state recognition. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 14(2): 199–213, Feb. 1992.
- [25] E. Dickmanns. A general dynamic vision architecture for UGV and UAV. *Journal of Applied Intelligence*, 2:251–270, 1992.
- [26] R. Suorsa and B. Sridhar. A parallel implementation of a multisensor feature-based range-estimation method. *Proc. 1993 IEEE Conference on Computer Vision and Pattern Recognition*, pages 379–385, 1993.
- [27] B. Sridhar, R. Suorsa, and B. Hussien. Vision-based obstacle detection for rotorcraft flight. *Journal of Robotics Systems*, 9(6):709–27, 1992.
- [28] B. Sridhar, R. Suorsa, and B. Hussien. Vision based techniques for rotorcraft low altitude flight. *Intelligent Robotics. Proceedings of the International Symposium*, 1571:27–37, 1991.
- [29] R. Michelson. Aerial robotics competition rules. Technical report, Georgia Tech Research Institute, Smyrna, Georgia, 1995.

- [30] A. H. Fagg, M. A. Lewis, J. F. Montgomery, and G. A. Bekey. The USC autonomous flying vehicle: An experiment in real-time behavior-based control. In *Proceedings of 1993 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS '93)*, Volume 2, pages 1 173–80, 1993.
- [31] M. A. Lewis, A. H. Fagg, and G. A. Bekey. USC autonomous flying vehicle: An experiment in real-time behavior-based control. In *Proceedings of 1993 IEEE International Conference on Robotics and Automation; Atlanta, GA*, Volume 2, pages 422–9, 1993.
- [32] N. C. Baker, D. C. Mackenzie, and S. Ingallis. Development of an autonomous aerial vehicle: a case study. In *Applied Intelligence: The International Journal of Artificial Intelligence, Neural Networks, and Complex Problem-Solving Technologies*, Volume 2, pages 27 1–97, 1992.
- [33] K. M. Black, J. O. Smith, and R. A. Roberts. The uta autonomous aerial vehicle-automatic control. In *Proceedings of the IEEE 1992 National Aerospace and Electronics Conference] NAECON 1992 (Cat. No.92CH3158-3)*; Dayton, OH, Volume xvii+1295, pages 489–96, 1992.
- [34] R. P. Paul. *Robot Manipulators: Mathematics, Programming, and Control*. MIT Press, Cambridge, 1980.
- [35] R. G. Willson. Modeling and calibration of automated zoom lenses. In *Videometrics III; Boston, MA, USA; 2-4 Nov. 1994*, Volume 2350, pages 170–86, 1994.
- [36] A. Rosenfeld and A. C. Kak. *Digital Picture Processing*. Academic Press: New York, 1976.
- [37] P. Anadan. A computational framework and an algorithm for the measurement of visual motion. In *International Journal of Computer Vision*, Volume 2, pages 283–3 10, 1989.
- [38] R. S. Cok and J. S. Gerstenberger. A t9000-based parallel image processor. In *Transputer Research and Applications, NATUG-6. Proceedings of the Sixth Conference of the North American Transputer Users Group; Vancouver; BC, Canada*, Volume viii+355, pages 142–52, 1993.
- [39] D. May and P. Shepherd, and R. Thompson. The T9000 transputer. In *IEEE 1992 International Conference on Computer Design: VLSI in Computers and Processors. ICCD '92 (Cat. No.92CH3189-8)*; Cambridge, MA, Volume viii+605, pages 209–12, 1992.
- [40] M. Atkins. Performance and the i860 microprocessor. *IEEE Micro*, 11(5):24–7, 72–78, 1991.
- [41] N. Margulis. i860 microprocessor internal architecture. *Microprocessors and Microsystems*, 14(2):89–96, 1990.
- [42] S. Parry. Parallel DSP looks to fill transputer gap. *New Electronics*, 24(10):43–4, 1991.
- [43] Texas Instruments. *Texas Instruments TMS320C4x Users Guide*, 1993.

- [44] D. Hartley and D. M. Harvey. Analysis of the TMS320C40 communication channels using timed petri nets. In *Applications and Theory of Petri Nets 1993. 14th International Conference Proceedings; Chicago, IL*, Volume ix591, pages **562–71**, **1993**.
- [45] E. D. Kaplan. *Understanding GPS Principles and Applications*. Artech House Publishing, Boston, **1995**.
- [46] Federal Radionavigation Plan, Dept. of Defense and Dept. of Transportation. *Global Positioning System Standard Positioning Service*, **1995**.
- [47] T. J. Ford and J. Neumann. NovAtel's RT20 - a real-time floating ambiguity positioning system. In *ION GPS-94, Salt Lake City*, **1994**.
- [48] K. Dowling, R. Guzikowski, J. Ladd, H. Pangels, S. Singh, and W. Whittaker. *Navlab: An Autonomous Navigation Testbed*, Chapter **12**. Kluwer Academic Publishers, **1990**.
- [49] S. Yoshimura and T. Kanade. Fast template matching based on the normalized correlation by using multiresolution eigenimages. In *IEEE/RSJ International Conference on Intelligent Robotics and Systems*, Munchen, Germany, August **1994**.
- [50] K. Fukunaga. *Introduction to Statistical Pattern Recognition*. Academic Press, Boston, second edition, **1990**.
- [51] M. Uenohara and T. Kanade. Vision-based object registration for real-time image overlay. In *Computer Vision, Virtual Reality and Robotics in Medicine. First International Conference, CVRMed '95; Nice, France*, Nice, France, April **1995**.
- [52] D. Lowe. Robust model-based motion tracking through the integration of search and estimation. In *International Journal of Computer Vision*, Volume **8**, pages **113–122**, **1992**.
- [53] O. D. Faugeras. *Three-dimensional Computer Vision: a geometric viewpoint*. MIT Press, Boston, **1993**.
- [54] J. L. Mundy. Image understanding research at GE. In *Proceedings of 23rd Image Understanding Workshop*, Volume 1, pages **143–7**, **1994**.
- [55] K²T Inc. *ZPI-40*, **1994**.
- [56] L. Matthies, R. Szeliski, and T. Kanade. Kalman filter-based algorithms for estimating depth from image sequences. In *International Journal of Computer Vision*, Volume **3**, pages **209–38**, **1989**.
- [57] O. Amidi, Y. Mesaki, and T. Kanade. Research on An Autonomous Vision-Guided Helicopter. In *Proceedings of AIAA/NASA Conference on Intelligent Robots in Field, Factory, Service, and Space*, League City, Texas, March **1994**.