A 3D Color Terrain Modeling System for Small Autonomous Helicopters

by

J. Ryan Miller

The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

Spring 2002

A thesis submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Robotics

Abstract

This thesis develops a novel aerial terrain modeling system. The system is unique since it flies onboard a small autonomous helicopter and senses the structure and color of its surroundings to build accurate 3D terrain models. The system is capable of modeling terrain where current approaches are too expensive, too dangerous, or too difficult.

The prototype system is primarily composed of a mechanically aligned laser rangefinder and 1-pixel color camera, viewing the terrain through a common scan mechanism. The merit of this sensing approach is that range and color measurements are inherently collected from an identical terrain location.

This thesis presents a novel sensor, calibration methodology, and synchronization approach for a working terrain sensor prototype. The prototype's performance was verified by carrying out a number of real-world mapping missions. These missions range from geological feature modeling in the Arctic for NASA scientists, to mapping an urban building complex for DARPA researchers. The system has proven to be effective in over 50 modeling flights, which produced terrain models accurate to <20cm in 3D.

Acknowledgements

I would like to thank Omead Amidi for his continual support, and guidance during my graduate work at CMU. I genuinely appreciate the many hours he has spent working with me on this thesis, my research, and on the autonomous helicopter project. I have learned a great deal from our collaborations.

In addition to Omead, I sincerely appreciate the efforts of Chuck Thorpe and Takeo Kanade as my doctoral advisors. Chuck has always been a delight to work with, and is truly an inspirational basketball player. Takeo has been very supportive of my research and the autonomous project. I also thank Clint Kelly from SAIC for serving on my thesis committee and helping to support my research.

I would like to thank Mark DeLouis for being an integral part of the autonomous helicopter project, for his willingness to share his wealth of experience with me as we've worked together, and for teaching me how to fly model helicopters.

Thanks to Matt Arnold for generating many of the rendered images used in this thesis. I also sincerely appreciate Matt's friendship throughout my undergraduate and graduate years at CMU.

I would like to thank Todd Dudek, Marco La Civita, Greg Krauland, Bernard Mettler, and Randy Warner for their involvement in the project and my research.

I am fortunate to have the full support and love of my family. To my parents, Patti and Alec Miller, and sister Lara, I thank you for everything you've done for me, and for providing a truly wonderful family environment. I love you all.

I would like to thank Denise Coyle for here friendship, companionship and love. I appreciate her patience and support during my graduate studies, and look forward to our future togeather.

Last, but certainly not least, I must thank Shadow, Sunshine, Kafka, Tinker and Gypsy, my feline companions.

Table of Contents

Table of Figures

3 Terrain modeling system: Implementation 59

4 Terrain modeling system: Calibration and verification 121

Chapter 1

Introduction

In this thesis we describe a novel terrain sensing system, which can quickly and accurately capture the three-dimensional (3D) structure and appearance of large spans of terrain from desired vantage points. The system is based on an active laser rangefinding sensor for terrain localization, and a passive 1-pixel color camera for registered appearance measurements. The system operates onboard an autonomous helicopter that has localization and flight control capability.

A 3D digital model capturing structure and appearance of an area of interest is extremely valuable. Autonomous robots need such models of their surroundings for trajectory planning. Geographic information systems (GIS), the primary tool for construction planning, precision agriculture, and natural resource exploration, build accurate terrain models of large areas by merging individual local models. The military plans missions, guides missiles, and identifies targets based upon accurate terrain models. Simulation tools, virtual reality systems, and video games build realistic terrain models based on real-world data captured at sites of interest.

These models are useful only if they are of high quality (accuracy, density, and appearance). For example, archiving a historical structure such as the Statue of Liberty must be accurate enough to replicate the entire structure as actually built. Terrain

matching navigation systems such as cruise missiles and robotic ground vehicles cannot effectively navigate without accurate 3D structure and appearance models of their surroundings. An effective and useful modeler must maintain high quality while building models. Most useful are models of remote locations captured from desired vantage points, and over large areas, a task that has traditionally been difficult and extremely time consuming.

A helicopter is an ideal terrain mapping platform. The helicopter is highly maneuverable. It can take off and land without a runway, it can hover in place, and most importantly, it can serve as an eye-in-the-sky for observing remote or difficult to reach areas from desired vantage points. A small unmanned helicopter is especially desirable, as opposed to a full-sized helicopter, since the former can be easily transported by truck, uses a small amount of fuel, and is significantly less expensive to operate.

The maneuverability of the helicopter comes at a significant cost. Helicopters are highly unstable and dangerous to fly in close proximity to the terrain. An autonomous unmanned helicopter eliminates the danger to a human pilot, and can be smaller, safer, quieter and less expensive than a manned vehicle. New unmanned helicopter technology replaces the human pilot with an onboard computer control and sensing system which accurately positions the helicopter as desired and can virtually stabilize the terrain modeling sensors.

The primary terrain modeling sensors of the system are a laser rangefinder and a novel 1-pixel color sensor, both observing the terrain through a common scanning mechanism. This combination captures accurate color-enhanced 3D terrain measurement, and, when combined with the autonomous helicopter, enables an unsurpassed terrain modeling capability.

FIGURE 1-1. Terrain modeling system onboard CMU autonomous helicopter

This thesis describes how we developed and demonstrated such a terrain modeling system onboard a small (14 foot long, 165 pound) unmanned autonomous helicopter, shown in Figure 1-1. Both the system and the autonomous helicopter were developed at Carnegie Mellon University.

This terrain modeling system is capable of building accurate color-enhanced terrain models during flight. Figure 1-2 shows a typical colored point cloud representation of an outdoor scene. Each point corresponds to a single measurement of the terrain sensor, and consists of a 3D position with a red, green, blue color attribute. Figure 1-3 shows the same result from a different viewpoint. Figure 1-4 is a standard picture of the scene, taken from approximately the same viewpoint.

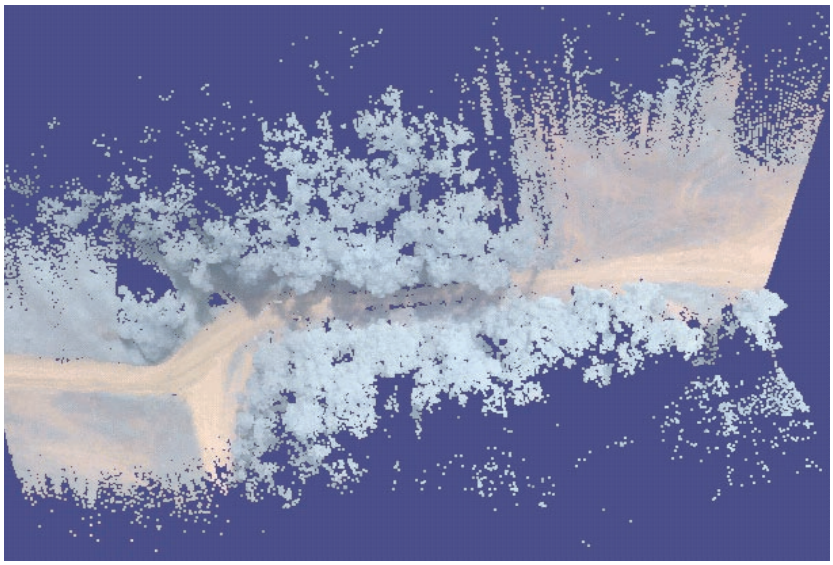

FIGURE 1-2. Typical resulting color-enhanced terrain model

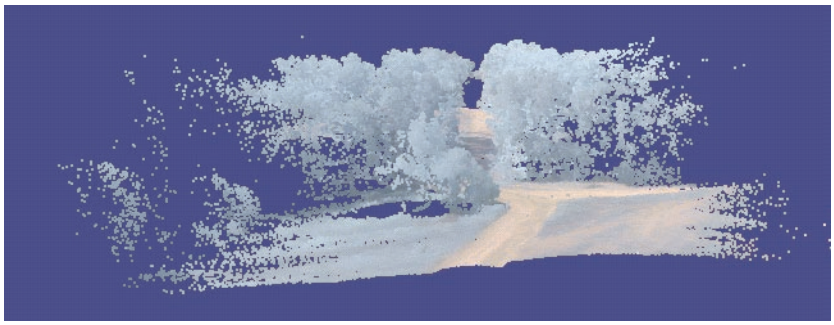

FIGURE 1-3. Alternative view of terrain model

FIGURE 1-4. Picture of scene

1.1 Modeling approach: A terrain sensor on-board a small autonomous helicopter

We claim that a small autonomous helicopter is an ideal platform for terrain modeling. In this section, we discuss the shortcomings of current practice for modeling remote areas. We motivate the need for autonomous helicopter-based modeling, and we present the challenges in building such a system.

1.1.1 Other terrain modeling approaches

Surveying equipment

Traditional surveying equipment and modern tools such as the Total Station [1] and portable GPS receivers [2] perform extremely accurate 3D position measurements (< 1cm) of a cooperative target (retroreflective prism or antenna), which is manually positioned at the desired measurement points. They are the best choice when measuring a small number of highly accurate points; however, they are intractably slow for comprehensive coverage of expansive terrain surfaces, due to the necessity for manual repositioning of the target or antenna for each measurement. Additionally, the points of interest must be accessible to a human surveyor holding the target, a significant problem in distant, dangerous, or hard-to-reach places. Finally, large distances may require multiple relocations of the instrument, requiring time for setup and referencing operations prior to taking measurements.

Precision ground scanners

Numerous systems exist for building models from a stationary platform using a scanning laser rangefinder [3][4][5]. These provide highly accurate models, but are limited in what they can practically scan. Since they measure from a single vantage point, only the immediate surroundings can be scanned, requiring many separate attempts to cover relatively small areas. These multiple scans must be registered to each other, requiring the placement of reference targets in the scene, or a difficult

automatic surface registration approach. Because the scanners are particularly vulnerable to any sensor movement during the scans, the scanner must be attached to solid ground. It is therefore difficult to get the scanner close enough to large objects to scan them, since scaffolding or scissor-lifts may not be stable enough.

Satellites and aerial photogrammetry

Remote observations over very large areas are possible using satellites [6] and high flying aircraft with photogrammetric equipment [7]. While the horizontal accuracy of these systems is amazing, the elevation measurement is relatively poor. Additionally, the elevation may be determined only at points where defining features in the terrain are found. Measurements can be performed only while there is no cloud cover. Finally, because only an overhead vantage point is possible, it is not possible to model terrain details on nearly vertical surfaces and under overhanging structures.

Precision airborne laser mapping

Presently, to build accurate topographical maps for applications such as power line inspection, tree monitoring, and flood plane determination, a number of commercial surveying companies use scanning laser rangefinders and GPS on-board manned helicopters [8]. The Institute for Photogrammetry [9] at the University of Stuttgart has an extensive research program investigating methods for improving sensor accuracy and automating the post-processing steps used to recover CAD models of the terrain. The size, weight and power requirements of these systems far exceed the available payload capability of the small helicopter. They are also limited to surveying from significant distances due to pilot safety requirements and potential interaction between the full-sized helicopter and the environment. The increased scan distance causes a magnification of angular errors, reducing accuracy. The increased scan distance also reduces measurement density. Finally, it is logistically difficult to mobilize full sized

helicopters to reach remote locations, since fuel and landing sites must be prepared in advance.

1.1.2 Our terrain modeling approach

A small, unmanned autonomous helicopter is an ideal platform for precision terrain modeling. The unconstrained helicopter is free to position its onboard terrain sensors as needed to best scan a site. However, the small helicopter's limited space and payload imposes significant restrictions on the on-board terrain sensor.

The small, unmanned autonomous helicopter is a particularly suitable platform for terrain modeling for several key reasons:

- The helicopter is small and light enough to be readily transported and deployed. Two people and a pickup truck can effectively transport and deploy such a system. The small size is also very advantageous when flying close to objects. Because, significantly less downwash and turbulence are generated by the helicopter, it is less disturbing to the environment.
- The small, unmanned autonomous helicopter is inexpensive to operate compared with full-sized aircraft. The full-sized aircraft requires extensive ground support (airport, fuel, hangars), hours of maintenance, and specialized equipment, all of which are very expensive. This makes the overall small helicopter mapping system an option for cost-sensitive applications, and provides an opportunity for large-scale modeling operations to afford multiple helicopters simultaneously.
- The unmanned helicopter is much safer than a manned aircraft, since it does not place a pilot at risk.
- The autonomous flight controls can accurately position and hold the helicopter at the optimal locations, and can precisely move it in an efficient path to cover the desired terrain from the desired vantage point. Additionally, the smooth and controlled flight provides consistent coverage of the terrain.
- Many sensors perform best when observing the terrain with a normal incidence

angle. By flying above the ground, the helicopter scanning system can achieve high quality measurements as opposed to those obtained by ground level scanners. Ground level scanners view the ground at oblique angles which returns less signal resulting in noisier measurements and a reduced range.

- Any terrain modeling system that is not fixed to a single location requires accurate knowledge of its actual position and attitude. An autonomous helicopter already has an onboard state estimation system as required to fly. This system can also be used for the terrain modeling functions.

The design and approach of the terrain modeling system can be demonstrated only with a real-world implementation. We follow a strapped-down implementation approach where the terrain sensors are not mechanically stabilized, but instead are rigidly affixed to the airframe. Strapped-down sensors are more compact and lighter than stabilized sensors; thus, they fulfill a major requirement for small autonomous helicopter flight. The disadvantage of this is that the sensors are not stationary as we scan the terrain; therefore, all readings are not observed from the same vantage point. For the helicopter to fly without any restrictions, the system must compensate for the actual sensor motions to build a coherent model of the terrain. We demonstrate that this approach can virtually stabilize the sensor suite and can practically map the terrain.

Figure 1-5 illustrates an example where existing terrain modeling systems fail to perform. A detailed 3D model of the entire dam surface requires mapping a large nearly vertical area (the walls of the dam). The scanner must look horizontally toward the wall from a low vantage point. From this vantage point, even areas under overhanging structures can be mapped effectively. Accurate and detailed coverage of the wall also requires that the vantage point be close to the wall. Here, a small, unmanned autonomous helicopter can safely fly in close proximity to the wall to position the onboard terrain sensors for the optimum vantage points. The helicopter can move to precisely cover the entire structure, while virtually stabilizing the terrain

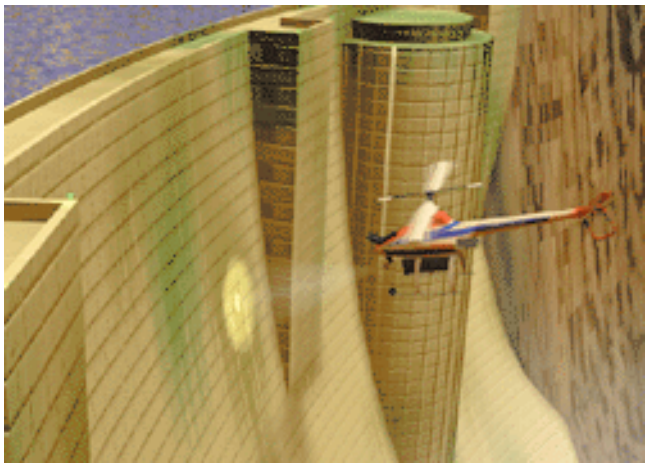

FIGURE 1-5. The autonomous helicopter provides a virtual scaffolding for terrain modeling in hard to reach areas.

sensing payloads from disturbances. When finished, it will return to its initial launch point, perhaps miles down river, to land.

1.1.3 Challenges of our terrain modeling approach

Terrain modeling from an autonomous small helicopter is challenging. The helicopter motions are highly dynamic and the onboard payload is limited.

The helicopter can move quickly and unexpectedly in all six degrees-of-freedom. Small helicopters can accelerate in the range of 0.5-0.7 g and can exhibit an angular velocity of greater than 100 degrees per second under normal operating conditions, and wind gusts can suddenly disturb the helicopter further. To keep up with the helicopter's high degree of maneuverability, a terrain modeling system must quickly compensate for its constantly changing viewpoint from the helicopter.

Additionally, vibrations due to the spinning rotor blades, gear train, and engine affect the terrain sensors. This agility requires accurate and timely state estimation and careful system-wide synchronization.

Helicopters have strictly limited payloads and available power. Typical small helicopters have a payload of only 20 to 75 pounds, and only a fraction of this can be allocated to the terrain modeling system. Similarly strict limitations must be enforced for electrical power consumption, equipment size, and placement for proper balancing. These limitations practically restrict the possible approaches in designing such a terrain-sensing system.

In this thesis, we address these challenges and describe how we demonstrated a capable terrain modeling system onboard a small autonomous helicopter.

1.2 Sensing approach: Terrain sensing by scanning a registered laser rangefinder and 1-pixel camera

We claim that a laser rangefinder and 1-pixel color camera, in combination with a simple scanning mechanism, produce a capable helicopter-based terrain sensor. In this section, we describe the need for registered appearance and geometry measurements, discuss the shortcomings of existing terrain sensors for helicopter-based modeling, present a novel terrain sensor design, and discuss some challenges in building such a sensor system.

Terrain models are composed of two primary classes of information about terrain: geometry and appearance. A 3D geometry sensor of the terrain measures the surface shape based on a collection of 3D coordinates sampled across the terrain surface. An appearance sensor of the terrain measures the terrain’s appearance from the point of view of the sensor by measuring color. We desire to build a single terrain model that captures both the geometry and appearance. Because, the two classes of measurements must be registered together, the appearance measurements are properly placed in 3D space, thereby allowing the unified color enhanced 3D model to be viewed from various viewpoints.

1.2.1 Other terrain sensing approaches

There are many possible sensor combinations that can measure terrain from an autonomous helicopter. In this subsection, we discuss several of the most likely candidates and evaluate their capability to produce high quality registered appearance (color) and 3D geometry terrain models.

Stereo vision

Stereo-based terrain modeling sensors [10][11][16] use two or more cameras to acquire sets of high resolution color images of the terrain, and compute the terrain geometry based upon triangulation of similar image areas. While the simultaneous capture of images allows a high quality appearance model to be developed, the quality of the geometry measurements is poor. Due to the stereo method's high depth sensitivity to image noise and resolution, accurate depth measurements are difficult to obtain. Some improvements in accuracy can be achieved by high quality cameras and optics, but these are large and heavy, and can not be carried onboard a small helicopter. Additionally, useful depth measurements are located only sparsely across the images, at uniquely identifiable corresponding features.

Laser scanner and video camera

The laser scanner and video camera combination [12] can directly measure the 3D location of a set of terrain points along its scan path. The laser rangefinder effectively measures the distance to the terrain along a narrow, pencil-like laser beam. When combined with the scanning mechanism and helicopter state, the rangefinder can build a high quality terrain geometry model. Meanwhile, a color video camera captures images of the terrain appearance, providing a high quality appearance model of the terrain. Unfortunately, it is quite difficult to register the geometry model with the appearance model. There are no good measures for evaluating the quality of registration between the camera image and the surface shape, a requirement for the

standard registration refinement techniques. An additional difficulty is that the fields of view of the two sensors are not matched, limiting the areas where both appearance and geometry data are potentially available.

Structure from motion

Structure-from-motion-based [10][15] terrain sensors use a single camera on-board the helicopter to capture a sequence of images as the helicopter flies. These techniques attempt to extract both the terrain geometry and appearance models from the images. As with any camera-based sensor, the appearance model is of high quality, but difficulties in accurately computing the terrain locations provide a poor quality geometric model. Additionally, most of these techniques require massive amounts of processing, which presently is difficult to realize in real time on-board a small helicopter.

SAR and line camera

A combination of a synthetic aperture radar (SAR) [13] and a line camera [14] allows direct measurement of the terrain geometry and appearance across a line on the ground. SAR is an advanced radar technique used to measure the terrain geometry along a strip of ground to one side of the helicopter as it flies forward. A color line camera (linear CCD) can measure the terrain appearance along a similar strip, building an appearance model.

The SAR measurements are advantageous because they can measure terrain from long distances, require no moving parts, and are capable of resolving terrain features significantly smaller than other radar-based sensors. However, ambiguities in the complicated SAR processing algorithms and measurement noise limit the accuracy of the resulting geometry model.

The appearance model is highly accurate and has an approximately matched field-of-view with the geometry model. However, registering the two models is still challenging. Determining when a color image is properly registered with a surface geometry is difficult for the same reason the laser scanner and video camera combination is difficult; there is no good cue for comparison. Additionally, the current SAR sensor hardware is too large and too heavy for use on a small helicopter, and requires enormous processing power.

1.2.2 Our terrain sensing approach

In this thesis, we describe how we developed a terrain sensing system based on an active laser rangefinding sensor for terrain localization, and a passive 1-pixel camera for obtaining registered appearance measurements (primarily of terrain color). This combination, when integrated with the autonomous helicopter's localization system, provides a comprehensive terrain modeling capability. The registered appearance and geometry data also serve as a framework for registering additional sensors such as high-resolution color cameras with the terrain model.

This system measures points on the terrain surface by using a scanning mechanism, a laser rangefinder and a 1-pixel camera. A scanning mechanism aims the laser beam in a known direction while the rangefinder directly measures distance. Simultaneously, a 1-pixel camera focuses on the same terrain point to measure its appearance. To ensure that the range measurement and appearance measurement are taken from the same point, the sensors are physically attached together, aligned, and share a common path through the scanning mechanism.

For simplicity, the 1-pixel camera measures only a single terrain point at a time, as opposed to imaging sensors such as a Charge Coupled Device (CCD) or focal plane array. This does, however, require a significantly faster measurement rate (>10,000 samples per second vs. 60), since a new reading is needed with each laser range

measurement. The primary advantage of this configuration is that each appearance measurement is taken of a known 3D location on the terrain, solving the registration problem. Also, since the laser and 1-pixel camera sample the same point, they have identical fields of view. Every geometry measurement therefore has a corresponding appearance measure, thereby avoiding the typical mismatch experienced in current systems.

The 1-pixel camera approach allows the measurement of a range of quantities. The camera contains multiple passive sensing channels, each measuring a different appearance attribute. Some channels can rely on sunlight to illuminate the target area and can measure characteristics such as a color band or polarization component of the light that reflects off of the surface and into the sensor. Other channels can measure target area emissions, such as infrared radiation. The number and choice of sensing channels can be tailored to match the desired overall capabilities. We primarily concentrate on sensing color with a three-channel sensor that measures the red, green and blue color components.

The terrain is observed by scanning the laser beam in a simple, fixed pattern across the terrain below the helicopter, as shown in Figure 1-6. Due to its small size and weight, a simple scanning mechanism is highly desirable for on-board the helicopter. However, the simplified scanning pattern prevents the helicopter from scanning large areas of terrain from a single location. Scanning large areas requires the helicopter's own movement to 'slide' the scan area across the terrain in a push-broom fashion. Since the autonomous helicopter can carry out this motion smoothly, we achieve even coverage of the areas.

In this thesis, we present the implementation, calibration, and methodology of a prototype terrain sensor, along with results obtained from its use.

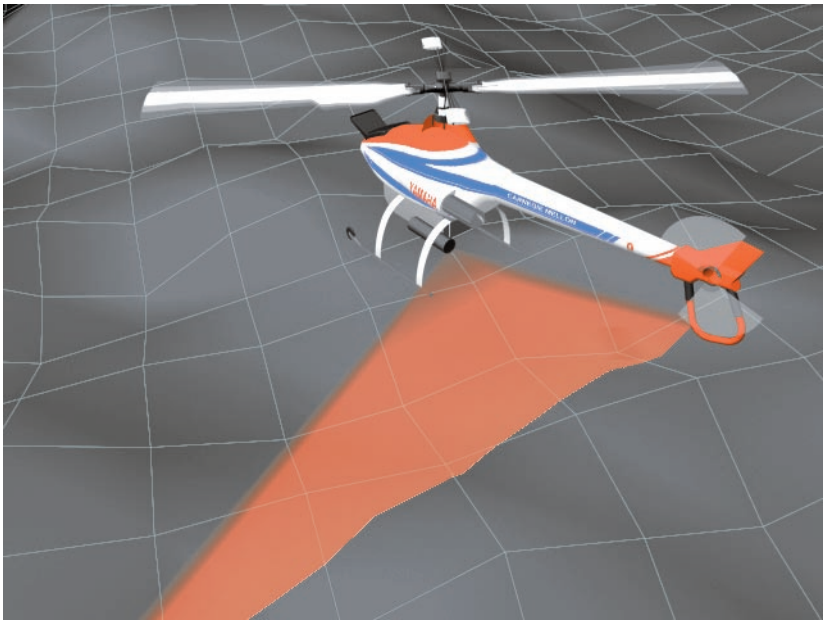

FIGURE 1-6. Pushbroom scanning configuration

1.2.3 Challenges of our terrain sensing approach

Multiple sensors are used on-board the helicopter for terrain modeling, state estimation, and additional payloads. Registration (spatial and temporal) of these various data sources is critical for accurate terrain modeling. Accurate registration is achieved by precisely adjusting each sensor's location and view angle from the helicopter, and recording the precise time each sensor measurement is taken. Accurate spatial registration is achieved by determining the sensor view through a set of calibration and adjustment procedures. Accurate temporal registration is achieved by a helicopter-wide hardware-based synchronization and timing system.

1.3 Thesis contributions

This thesis has three main contributions; in it, we describe how we:

1. Demonstrated the first integrated terrain modeling system onboard a small autonomous helicopter. The system accurately (< 10 cm) localizes the terrain using a

scanning laser rangefinder and measures a set of terrain appearance attributes using a 1-pixel camera. The system is tightly integrated with the autonomous helicopter to provide a coherent terrain modeling capability despite aggressive helicopter movement. The system weighs less than 10 pounds, making it usable onboard a small autonomous helicopter.

2. Designed and deployed a high-speed appearance attribute sensor. The multi-channel 1-pixel camera measures terrain appearance attributes (color) up to 12,000 times a second within a narrow (1^o) field-of-view. The 1-pixel camera is aligned, bore sited, and synchronized with the laser rangefinder to guarantee collocation of appearance measurements (color) and laser range measurements.
3. Developed effective calibration and synchronization methods for accurate terrain modeling. A system-wide synchronization system accurately integrates terrain sensor data with other autonomous helicopter state estimation sensors. Several novel calibration techniques were developed to align, focus, and aim the terrain sensor. Most significant is an in-flight calibration method that allows rapid reconfiguration of the scanning direction in the field.

1.4 Thesis overview

This thesis is divided into six chapters. Chapter 2 describes the architecture of the terrain modeling system as consisting of four primary components: the laser rangefinder, the 1-pixel camera color camera, a common mechanical scanning mechanism, and the integration of this terrain sensor with the onboard autonomous helicopter system. This chapter describes the requirements and design of each of these components.

In Chapter 3, we present our implementation of a terrain modeling system on-board a small autonomous helicopter. This chapter details the design and trade-offs employed to achieve a practical working system.

In Chapter 4, we present a series of calibration and verification procedures developed for the terrain modeling system, and detail each of the calibration procedures needed to set up the terrain modeling system to achieve its designed performance. We also describe the incremental verification process which ensured proper operation during the system development.

In Chapter 5, we present the results of deploying the helicopter-based terrain modeling system in a number of applications.

Finally, in Chapter 6, we present conclusions and future research directions of the work presented in this thesis. The future research directions presented include a modified 1-pixel camera for polarization measurements, and a CCD camera registration scheme for colored terrain models.

Chapter 2

Terrain modeling system: Architecture

We are interested in accurately modeling terrain from an autonomous helicopter. We consider all stationary objects near the helicopter to be part of the “terrain”. This includes the ground, trees, buildings, and vehicles. Since the autonomous helicopter is most useful when operating near the ground, it may frequently fly among the terrain. Therefore, we seek to build highly accurate (~20 cm) color-enhanced terrain models of nearby (< 200m) terrain.

In this chapter we present a general discussion of our terrain modeling system architecture. We first discuss design issues regarding the terrain sensor and its three primary components: a laser rangefinder, a 1-pixel camera, and an optical scanning mechanism. Then, we discuss the combination of an autonomous helicopter with the terrain sensor to complete the terrain modeling system.

2.1 Architecture overview

The class of terrain sensors for our terrain modeling system comprises three primary components: a laser rangefinder, the 1-pixel camera, and an optical scanning mechanism. Figure 2-1 shows the layout of these three components. The sensor is set up such that the principle viewing direction of the rangefinder and the 1-pixel camera are merged into a single optical path inside the sensor thereby ensuring that they

measure the identical terrain points. This single viewing ray can be mechanically steered by a moving mirror in the scanner mechanism toward different points on the terrain. Various terrain modeling sensors based on this basic layout can be developed with different choices for the laser rangefinder, 1-pixel camera, and scanning mechanism.

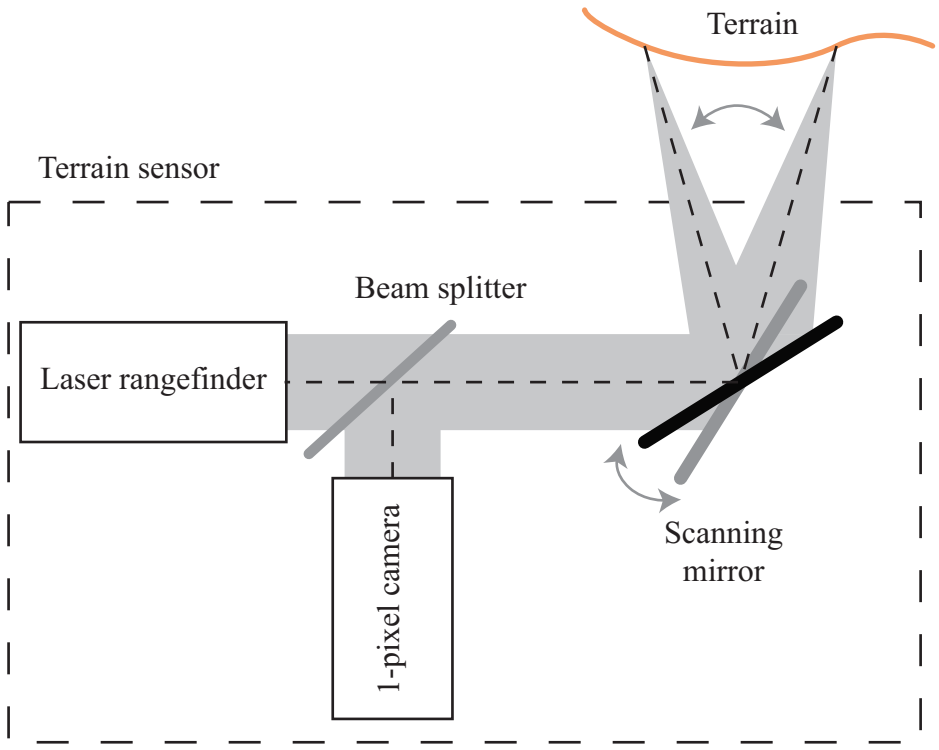

FIGURE 2-1. Principle components of terrain sensor

The terrain sensor is considered to be a single unit and is integrated with an autonomous helicopter to produce the complete terrain modeling system described in this thesis.

2.2 Laser rangefinder

The laser rangefinder measures the distance to the terrain along a narrow pencil-like laser beam. To measure distance, the laser rangefinder sends out a modulated laser beam, and analyzes the return light reflected off of the terrain. The range, combined with the scanning mechanism and helicopter state, accurately calculates the 3D location of the terrain. Most rangefinders repeatedly fire the laser so as to measure range at a desired frequency. Numerous laser rangefinders are commercially available. In this section, we discuss the important properties the laser rangefinder must possess, and the issues involved in selecting a suitable laser for our terrain mapping purpose.

2.2.1 Laser rangefinder objectives

We identify four primary objectives for the laser rangefinder for the terrain sensor. The laser rangefinder must:

1. *Measure distance from sensor to terrain along a single viewing direction to desired accuracy.*
2. *Possess sufficiently long range capability to measure distance from the helicopter to nearby terrain.*
3. *Acquire measurements at high enough rates to cover the terrain while the helicopter flies at a normal speed.*
4. *Be small, light, low powered, and robust enough for deployment on-board a small autonomous helicopter.*

2.2.2 Laser rangefinder model

The laser rangefinder measures the distance between the sensor and the terrain, along its principle axis. Understanding the characteristics and performance of the laser rangefinder is important to determining its effects on the overall terrain modeling system. The measurements from the laser rangefinder depend on the geometry of the

terrain and the laser rangefinder's exit aperture, divergence angle, min/max range, range accuracy, and sample rate. We examine these characteristics in more detail:

- The laser beam divergence angle (θ_{laser}) is the angle in which the laser beam spreads out after leaving the sensor. Specifically, the divergence angle is the angle between the 1/2 maximum intensity points on either side of the principle axis, commonly referred to as the full width at half maximum (FWHM).
- The laser exit aperture diameter (D_{laser}) is the size of the laser beam as it exits the sensor. Laser rangefinders typically have exit apertures in the range of 1-30mm. Theoretically, larger exit aperture lasers can be collimated better to achieve smaller divergence angles. However, in practice, some rangefinders intentionally have a large exit aperture and a larger divergence angle to ensure eye safety. This is effective because the beam is larger than the iris.
- The minimum and maximum range parameters ($RANGE_{min}$, $RANGE_{max}$) specify the valid distances over which measurements can be taken. The minimum range is often limited by the rangefinder's optical configuration, while the maximum range is typically found where the laser's reflection becomes too weak to detect. The surface reflectivity can have a significant effect on $RANGE_{max}$, since the amount of light reflected off a surface depends on surface properties of the object. A dark surface reflects significantly less light than a white surface does, resulting in a shorter $RANGE_{max}$.
- The range accuracy describes how close the measurement is expected to be to the actual distance. A statistical measure such as standard deviation is typically specified for the rangefinder when measuring a large, high reflectivity, planar target positioned normal to the laser beam. This is the most favorable condition, and consideration should be paid to factors that degrade this performance. The accuracy can be degraded when the target is observed from an off-normal incidence angle, since different parts of the laser beam are actually striking the target at different ranges. Some rangefinders are also affected by the strength of the reflected return.

This means that two objects at the same distance, but with different surface reflectivity, may be found at different ranges. Finally, some rangefinders produce extremely inaccurate results when one portion of the laser beam partially hits one target while the other portion hits another target that is located further away.

- The sampling rate ($RATE_{laser}$) specifies the number of measurements taken each second. Heat dissipation, data processing, and data transfer limitations inside the rangefinder typically constrain the sampling rate. A related consideration is the method the rangefinder uses to initiate a measurement. Some sensors allow an external trigger signal to start each measurement, while others continuously measure at a fixed sampling rate.

2.2.3 Laser rangefinder issues

Selecting an appropriate laser rangefinder requires consideration of a number of specific issues [19]. In this subsection, we discuss several of these issues in the context of aerial terrain mapping.

Rangefinding technique

Two dominant techniques for laser range measurement are used in commercial sensors: pulsed and Amplitude Modulated Continuous Wave (AMCW).

The pulsed laser rangefinder emits a short pulse of laser light and measures the time until the pulse reflects off the target and returns to the laser receiver. This time-of-flight is readily converted into range based on the speed of light. One advantage of the pulsed laser rangefinder for terrain mapping is that multiple reflections may be received if the laser spot hits multiple objects, such as a leaf on a tree and the ground behind it. In this case, a first reflection corresponding to the leaf distance will be seen at the time, followed by a second pulse from the ground.

Different variations of the pulsed laser exist based on how these multiple returns are handled. A first-pulse system detects the initial pulse, corresponding to the closest target (leaf). A last-pulse system [18] waits for all of the reflections to return, and measures the distance to the last one, corresponding to the farthest target (ground). The first-pulse system is desirable when scanning for purposes of obstacle avoidance, since the closest obstacle is detected. In fact, the last-pulse system could be disastrous for obstacle avoidance, as partial returns from close, narrow objects, such as power wires, will be intentionally ignored to show only the ground beneath them.

The last-pulse system is advantageous in adverse weather conditions where snow, rain, dust or other particulates in the air may reflect some of the laser beam before it reaches the terrain. The related last-pulse, trailing-edge rangefinder is able to measure the terrain surface even when the entire surface is enveloped in a cloud of dust or fog. The last-pulse laser is also easier to enclose in an environmentally protected enclosure, since minute reflections off of the enclosure window, a common problem with first-pulse systems, can be ignored.

The primary disadvantage of pulsed laser systems is the relatively slow sample rates. Fast, pulsed laser rangefinders typically have sample rates up to 30,000 hz, with the fastest rates being achieved by using multiple lasers fired in succession.

The second rangefinding method is Amplitude Modulated Continuous Wave (AMCW) [17]. The AMCW rangefinder continually modulates the laser beam with one sinusoid, or multiple sinusoids of different frequencies. Once received, the reflected light is compared with the outgoing light to determine the amount of phase shift at each of the modulation frequencies. The phase shift is directly related to range with a known modulation frequency. Using multiple frequencies lets the sensor establish the coarse distance, with the slower modulation, and achieve higher range resolution with the higher frequency modulations.

The AMCW laser's primarily advantage is its ability to sample at extremely high rates, approaching 1 million samples per second, and to measure range accurately (<5mm). This high rate provides excellent sample density, and allows the helicopter to fly faster, while still covering the terrain. However, unlike the pulsed laser, AMCW lasers are unable to handle multiple returns from different targets. The result of multiple returns is typically a position measurement which is completely invalid, possibly the range to somewhere between the targets. This result requires additional filtering to deal with spurious measurements. To minimize this occurrence, AMCW lasers typically have very small exit apertures and divergence angles, leading to a small spot size.

Laser spot size

The laser rangefinder determines distance by illuminating a small spot on the terrain with the laser, and then measuring reflected return. The size of this spot is dependent on range, laser divergence angle, and laser exit aperture size. The larger the spot, the bigger the terrain area considered for each measurement.

Successive laser samples across the terrain should be independent of each other; therefore, we desire to prevent their laser spots from overlapping. Similarly, too small a laser spot will leave very large unmeasured gaps which could potentially miss small or thin objects of interest, such as overhead electrical wires. Ideally, we desire a laser spot size similar to the spacing between successive measurements in space. Therefore, slow sample rate lasers would employ larger spot sizes to prevent large gaps, while fast sample rate lasers would employ small spot sizes, since the sample density is higher. Fortunately, this relationship exists in the commercial range finders. The very fast rangefinders typically have small spots, and slower rangefinders typically have larger spots.

Eye safety

Eye safety is an important issue when working with lasers. The eye can easily be damaged by absorbing excessive optical energy. Many of the lasers used in rangefinders are in the infrared region, and the eye's natural aversion responses (blinking, pain, and looking away) do not respond to light in this region. This presents a real hazard, as eye damage can occur before any sensation of pain or discomfort is felt; thus, it mandates that eye safety considerations be an integral part of the sensor design.

Eye safety considerations typically limit the maximum range of the rangefinder, since increased range requires additional power from the laser, posing a hazard to people. One technique that allows the laser range to be increased, while maintaining eye safety, is to use a large laser beam and spot size to spread the laser's energy over an area larger than the human pupil, thus reducing the maximum power that the eye may absorb.

Many times, a laser rangefinder that is not eye-safe when stationary can be made eye-safe when used with a scanner. By moving the beam around, the scanner prevents the eye from continuously absorbing the laser energy. In the worst case, the eye will have to absorb only periodic pulses of light from the laser as the scan pattern crosses the eye, a significant reduction.

Federal guidelines for the safe use of lasers are available from the United States Food and Drug Administration in regulations entitled "American National Standard for the Safe Use of Lasers" [20].

2.3 1-pixel camera

The 1-pixel camera must measure a set of appearance attributes (color of light, polarization, etc.) from a spot on the terrain as it is illuminated by sunlight. When aligned and synchronized with a laser rangefinder, registered measurements of the terrain range and appearance can be acquired simultaneously. We developed a custom 1-pixel camera, since no such sensor is commercially available. Custom development opens many possibilities which enable the optimization of the sensor for our aerial terrain modeling purposes.

2.3.1 1-pixel camera design objectives

We identify four primary objectives for the 1-pixel camera design. The design must:

1. *Passively measure the color appearance of the terrain at a single spot in typical daylight conditions.*
2. *Measure appearance at the same point on the terrain as that of the laser spot.*
3. *Measure appearance simultaneously with each laser reading.*
4. *Be small, light, low powered, and robust enough for use on a small helicopter.*

2.3.2 1-pixel camera design model

In this section we describe the process by which the 1-pixel camera observes the terrain. Many factors, including illumination, terrain surface properties, and sensor characteristics affect the measurement. This understanding of the measurement process is essential for making optimal trade-offs in the design of the sensor. We discuss two important aspects of the sensor design: the fundamental characteristics of the sensor and the expected performance of such a sensor in an actual environment. First, we present the fundamental 1-pixel camera characteristics which theoretically model the behavior of the sensor. These characteristics determine how rays of light from the environment enter the sensor, and what the sensor outputs in response to the

incoming light. Second, we examine how the terrain surface and external illumination affect the characteristics of the expected light from the scene.

1-pixel camera: Sensor model

The 1-pixel camera contains one or more sensing channels that are aligned to look along parallel principle axes. Each channel is composed of optics, a photodetector and associated electronics, and can be approximately modeled with a set of six parameters: effective aperture, field of view, depth-of-focus, spectral sensitivity function, minimum detectable signal, and bandwidth.

- The effective aperture (A_{sensor}) of a channel, the cross-sectional area at the front of the sensor, collects light. Only light rays passing through this collection area can be measured. In some channel designs, the effective aperture is based on the front diameter of the optics. In others, internal aperture stops determine the effective collection area.
- The field of view (FOV_{sensor}) of a channel is a measure of its spatial sensitivity. Typically, the channel is most sensitive along its principle axis, with reducing sensitivity as the off-axis angle increases. The field of view is specified as the angle within which ½ of the total channel sensitivity is contained (assuming that the sensor is focused at infinity). The field of view and effective aperture determine the size of the sensing spot on the terrain, as illustrated in Figure 2-2.
- The depth-of-focus of a channel specifies the minimum ($FOCUS_{min}$) and maximum ($FOCUS_{max}$) sensor-to-terrain distances for which the terrain is in focus. At any in-focus distance, the circle-of-confusion due to defocusing is less than a specified level. Typically, the allowable defocusing angle is related to the field of view, for example, 20% of FOV_{sensor}, since it diffuses the sensing spot over larger areas.
- The spectral sensitivity function, $S(\lambda)$, of a channel describes the expected output of the photodetector as a wavelength-dependent function of the incident light power. The spectral sensitivity function is affected by each of the optical elements

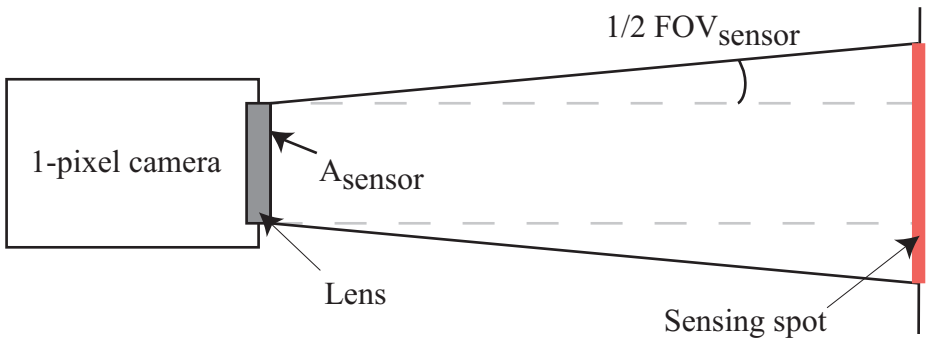

FIGURE 2-2. Effective aperture and FOV control sensing spot size

(lenses, filters, mirrors) in the sensing channel as well as the electro-optical characteristics of the photodetector and associated circuitry. The function specifies output photocurrent (in Amps) per unit of light power (in Watts) as a function of wavelength (in nm).

The incoming light is described by its Spectral Power Density (SPD) function. (This is a standard approach used for colorimetry calculations). The SPD specifies the light power (in Watts) per unit change in wavelength (nm) at various wavelengths.

The resulting photocurrent for a particular SPD of incident light is found by integrating the product of the SPD and spectral sensitivity function over all wavelengths.

- The minimum detectable signal (D_{min}) for a channel describes the channel's lowest photocurrent output (in Amps) that is measurable in the presence of noise. This parameter incorporates various sources of noise, quantization, jitter, biases, and integration time. It is a useful parameter as a reference for gauging the quantity of signal expected from the photodiode in particular situations. An imaging situation (sensor design and certain illumination environment) where the brightest surface produces a photocurrent similar to D_{min} is not desirable since the entire scene is "in the noise". Typically, a minimum photocurrent of $10 \cdot D_{min}$ is needed to acquire

usable measurements.

- The bandwidth (BW_{-3db}) of a channel describes the time taken for the sensor output to respond to changing light input. Sufficient bandwidth is necessary to achieve independent measurements from successive readings.

1-pixel camera: Terrain appearance model

In order to determine the optimal camera setup, it is useful to calculate the expected output from the camera, given knowledge of the lighting conditions and the nature of the objects to be imaged. Realistic physics-based appearance modeling is a very complex task, but the intent here is to identify a restricted model for guiding the sensor design. The optical principles used in this analysis are covered in many sources, such as [21] and [35]. The restricted situation we consider is illustrated in Figure 2-3.

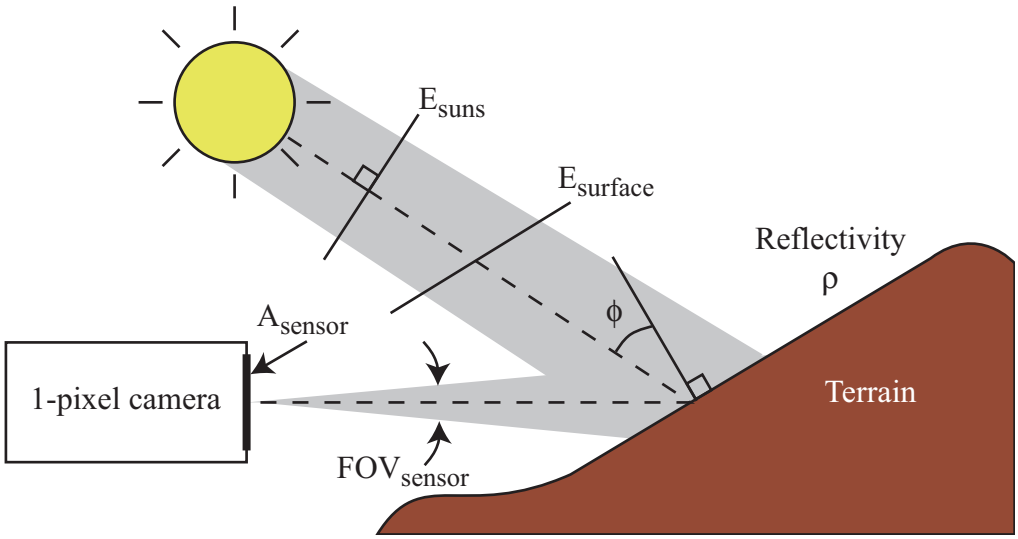

FIGURE 2-3. Simplified imaging model

The amount of light falling on the scene is known as the *irradiance* and is expressed as a power density (W/m^2). We assume that sunlight uniformly illuminates the terrain with an irradiance of $E_{sun,}$ and a spectral power distribution following that of a 5800K

black body radiator between 300nm and 900nm. Typical values of E_{sun} in various conditions are listed in Table 1 [22].

Condition	E_{sun}
Direct Sunlight	681.0 W/m^2
Full daylight	68.1 W/m^2
Overcast Day	6.81 W/m^2
Very dark day	0.681W/m^2

TABLE 1. Typical sun irradiance values

We consider the simple case where a single, planar surface patch is being observed by the sensor. The patch is considered to be an extended surface, since it is larger than the 1-pixel camera's sensing spot. If the sunlight is not arriving normal to the surface, a $1m^2$ area perpendicular to the sun beam projects over a larger area on the terrain surface. Therefore, the irradiance arriving at the surface is reduced to compensate for the angle of incidence, ϕ.

$$E_{surface} = E_{sun} \cdot \cos(\phi) \qquad (W)$$

The surface is assumed to be a diffuse Lambertian reflecting surface with a reflectance coefficient ρ. Colored surfaces could easily be considered by making ρ a function of wavelength, but a uniform reflectance (white) is acceptable for our evaluations. Further, it is common in conventional photography to assume an average scene reflectance of 18%, so we use ρ=0.18 as our typical scene value.

The percentage, ρ, of the incident light that reflects off of the surface spreads out into a full hemispherical volume according to Lambert's Law. The light leaving the surface is known as *radiance* and is expressed as a power density per solid viewing angle[1] (W /

1. A solid viewing angle describes the acceptance cone of light reaching the sensor. The base unit, the steradians, is equal to the total area on a unit sphere, centered at the sensor, falling within the FOV of the sensor. A true omnidirectional sensor, which accepts light from all directions, includes the entire sphere in its FOV, and has a solid viewing angle of 4π steradians.

m^2 / steradian). The Lambertian surface property ensures that the radiance, L, is the same in all visible directions. This means that the surface brightness appears to be the same from all viewing angles -- typical of a matte surface. The radiance is calculated as:

$$L = (E_{surface} \cdot \rho) / \pi \qquad (W / m^2 / sr)$$

The 1-pixel camera is viewing the surface, and measures all of the light rays hitting its sensitive area, A_{sensor}, arriving within its field of view, FOV_{sensor}. The solid angular field of view in steradians is $\pi \cdot \sin^2(FOV_{sensor}/2)$, so the amount of optical power collected by the sensor is:

$$P_{sensor} = L \cdot A_{sensor} \cdot \pi \cdot \sin^2(FOV_{sensor}/2) \qquad (W)$$

If we substitute our intermediate values:

$$P_{sensor} = E_{sun} \cdot \cos(\phi) \cdot \rho \cdot A_{sensor} \cdot \sin^2(FOV_{sensor}/2)(W)$$

This final equation describes how the sensor's collected energy is dependent on the intensity of the sun's illumination, angle of the surface to the light beam, reflectivity of the surface, area of the sensor's collecting lens, and field of view of the sensor. Notice that it does not depend either on the distance between the sensor and terrain surface, or on the difference between the viewing angle and the surface normal. This independence is important for successively imaging the terrain from a moving platform. We can reasonably expect that, as we take measurements from various locations and directions, the appearance of the terrain will remain constant.

1-pixel camera: Simulating the models

The 1-pixel camera and terrain appearance models form the basis for simulating different 1-pixel camera designs. The primary consideration is that the 1-pixel camera

develops a sufficiently large output signal (photocurrent) for a typical terrain environment illuminated with typical day-time sunlight.

The simulation begins by using the terrain appearance model to calculate the total optical power reaching each of the sensing channels, P_{sensor}. As a general rule-of-thumb, we evaluate sensor response to full daylight irradiance arriving at 45º to the terrain surface with a terrain surface reflectance of 18%.

Assuming neutral colored terrain, the optical power reaching the sensing channels has the same spectral distribution as sunlight. Therefore, the spectral power distribution of the light entering each channel is the total power (P_{sensor}) times the normalized spectral weighting function for sunlight, $Q(\lambda)$.

The total photocurrent output from the channel can be determined by integrating the product of the spectral sensitivity function and the spectral power distribution.

$$I_{ph} = \int_{\lambda=300}^{900} P_{sensor} \cdot Q(\lambda) \cdot S(\lambda) d\lambda$$

The calculated photocurrents for each channel can then be compared with the minimum detectable signal (D_{min}) to evaluate how well the terrain will be observed. As mentioned before, a minimum of $10 \cdot D_{min}$ is needed to reasonably observe the signal.

2.3.3 1-pixel camera issues

Desired spectral sensitivity

What is the desired spectral sensitivity of the 1-pixel camera? The spectral sensitivity of each sensing channel determines the range of light wavelengths to which it is sensitive. In the case of measuring color, we want the three channels to measure the standard red, green, and blue color components or bands.

All color imaging sensors do not have the same spectral sensitivities. In principle, this means that each system measures a different color triple (Red, Green, Blue) when observing the same colored object. There does not seem to be a single, correct answer to the desired spectral sensitivity question.

One approach we can follow is to design the 1-pixel camera with a spectral sensitivity closely matched to another particular color sensor, such as a particular 3-CCD color video camera, which may be on-board the helicopter. This allows the 1-pixel camera and that video camera to measure similar color values for the same object. This property could be valuable for machine vision processing of images from the 3-CCD camera in combination with the 1-pixel camera's colored terrain data, where the color measurements are used for terrain matching purposes.

Figure 2-4 shows the relative spectral sensitivity function for each of the 3 CCD elements in a Pulnix 3-CCD camera [23]. A useful guideline for designing the 1-pixel camera's optical components is to match the shape of these curves. This can be accomplished by properly selecting the optical components and photosensor in each channel.

Adjustability

Our approach requires mechanical alignment of the 1-pixel camera channels and the rangefinder's laser beam. Once assembled, each of the 1-pixel camera sensing channels may need to be physically adjusted to align its principle-viewing axis with the laser beam of the laser rangefinder. The design should provide a mechanism for adjusting the viewing ray and locking the adjusted setup in place to prevent any changes despite the harsh vibration of the helicopter.

Because the focus of each sensing channel may also need to be adjustable, an ideal mechanism allows adjustment and locking of the focus.

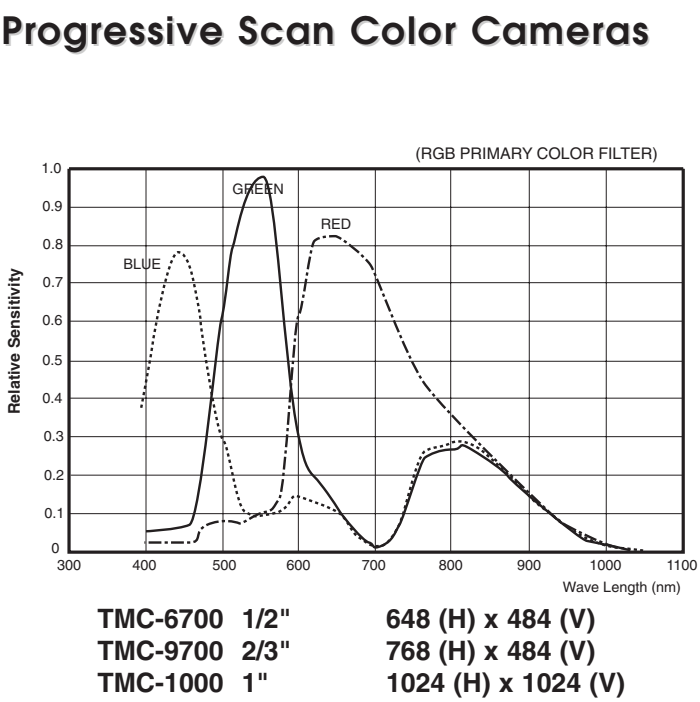

FIGURE 2-4. Spectral sensitivity curves for Pulnix 3-CCD video camera

Off-axis light rejection

The ideal 1-pixel camera is not affected by any light rays in the environment, except for the ones from the sensing spot on the terrain. Unfortunately, direct sunlight and off-axis light from other bright objects can get into the sensor and severely corrupt the measurement. The sensor design must be particularly capable of keeping these light rays out. Internal baffles, external shields, and other techniques must be deployed to retain the measurements' accuracy.

Capturing sufficient light

Capturing sufficient light for the measurements is challenging. The 1-pixel camera must take measurements at the same rate as the laser rangefinder, typically over 10,000

measurements per second. This severely limits the period of time during which light can be collected for each measurement, and therefore limits the total amount of optical energy reaching the photodetector. As the sensor model indicates, the signal can be increased by varying a number of parameters in the sensor design:

- The aperture, A_{sensor}, may be increased to collect more light. The collected light increases by the square of the aperture diameter, so significant improvements can be achieved with a larger aperture. At some point, however, physical limitations in the sensor will prevent further aperture increases. Another consideration is that as the aperture is increased, the depth-of-focus range is typically reduced. So, a large aperture may prevent the sensor from focusing at desirable close distances.
- The field of view, FOV_{sensor}, may be increased to collect light from a larger area of the terrain. As with the aperture, the collected light increases by the square of the FOV. The disadvantage of increasing the FOV is that each measurement is taken over a larger area, thereby reducing the resolving capability of the sensor.
- The efficiency of the sensor may be increased with better quality optics, more sensitive photodiodes, or low noise amplification circuitry. These approaches aim to boost the channel's spectral sensitivity function, $S(\lambda)$. Clearly, cost and complexity issues govern these improvements. Anti-reflection coated lenses, filters, and mirrors can significantly reduce the loss of optical power as light passes through each component. Electronic amplification clearly boosts the signal, but typically increases the noise level and may reduce the bandwidth. More noise increases the minimum detectable signal, D_{min}, possibly negating the signal gains, and lower bandwidths may prevent consecutive measurements from being independent. High sensitivity photodiodes, such as the avalanche photodiode (APD), can provide a significantly increased sensitivity, but are expensive and require extensive support circuitry.
- Finally, the minimum detectable signal, D_{min}, may be reduced by noise reduction techniques. This has an affect similar to increasing the signal amplitude. Typi-

cally, cost and complexity issues limit the available improvements.

We must focus on and examine the trade-offs and design decisions to realize a 1-pixel camera which is optimized for the helicopter terrain modeling application. This optimization is a major advantage of a custom-designed sensor.

2.4 Scanning mechanism

We desire to model several different types of terrain. The scanning mechanism, combined with the helicopter's movements, must view the terrain of interest. In some cases, we desire to quickly model large areas from above, for example mapping large fields or tree-top canopies. In this case, the scanner would continuously map the terrain beneath as the helicopter flies a prescribed trajectory. Other times, we wish to build a highly detailed model of a small site at a large distance from the takeoff point. The helicopter can quickly fly to the remote site, slowly maneuver as needed to allow the scanner to completely scan the area, and then fly back. Finally, we may model vertical and overhanging features (tree line, monument, or building), where the helicopter must carefully fly beside the terrain while the scanner models the surfaces either in front of or to the side of the helicopter.

The scanning mechanism takes a single light path from the laser rangefinder and 1-pixel camera and redirects it in different directions toward the terrain. A scan pattern describes the directions the light beam follows over time with respect to a stationary scanner. Numerous scan patterns are possible, and must be matched to desired tasks, taking into consideration the way the scanner is mounted to the helicopter. It is especially challenging to develop a scanning mechanism that is practical for flight on-board a small helicopter, due to size and weight limitations, harsh vibrations, and tight requirements for beam directional control and stability.

In this section, we discuss the design of potential scanning mechanisms for flight on-board an autonomous helicopter. These mechanisms are designed with extreme importance given to minimizing weight while ensuring adequate ruggedness.

2.4.1 Scanning mechanism design objectives

We identify four primary objectives for the optical scanning mechanism. The scanning mechanism design must:

1. Provide accurate measurements of beam direction synchronously with the laser rangefinder and 1-pixel camera samples.
2. Be small, light weight, low powered, and robust enough to operate on-board a small helicopter. In particular, the beam redirection must be accurately controlled despite constant vibrations from the helicopter.
3. Direct the sensor's optical path in a scan pattern that adequately covers the desired terrain areas when combined with the helicopter's own motion.
4. Ensure that the multiple sensing channels are uniformly redirected to maintain their alignment.

2.4.2 Scanner configuration issues

Adaptive vs. non-adaptive scanner

Payload restrictions dictate a configuration where the entire terrain sensor is rigidly attached to the helicopter airframe, as opposed to being mechanically stabilized. Therefore, any attitude change of the helicopter will also move the sensing spot on the terrain. There are two basic approaches to compensate for this effect: adaptive, and non-adaptive scanning.

An adaptive scanning approach attempts to control the terrain location being measured despite the helicopter's attitude changes. To do this, the scanning mechanism adjusts the scan pattern, in real time, in response to changes in the helicopter's attitude. Adaptive scanning allows a uniform sample distribution over the terrain, a property which may be valuable in some situations. Unfortunately, adaptive scanners are relatively large and complex mechanisms which allow the scan pattern to be controlled with high enough bandwidth to compensate for the helicopter's motions.

The non-adaptive scanning approach uses a scan pattern which is fixed and does not change with the helicopter's motion. Without feedback, the helicopter's attitude changes distort the distribution of samples on the ground, making them non-uniform.

Instead of mechanical adjustments, 3-D computations compensate for the helicopter's position and attitude to accurately localize the actual measurement point. The primary advantage of the non-adaptive scan approach is that the scanning mechanism is much smaller and lighter than its adaptive counterpart.

Conclusion: Due to its reduced complexity and lower weight, the non-adaptive scanning approach is preferred for use on-board a small autonomous helicopter. The greatest advantage of the adaptive scanning approach is that the sampling distribution can be done uniformly over the terrain surface. Our use of an autonomously flown helicopter allows the non-uniform scanner to still produce fairly uniform scan distributions. The autonomous flight system smoothly controls the helicopter in flight to minimize the drastic attitude changes common in human-piloted helicopters.

1-DOF vs. 2-DOF scanner

Many scanner designs use moving mirrors to reflect the light path toward the desired viewing direction. The viewing direction can be described by a pair of angles, θ and ϕ, measured from a fixed portion of the sensor. The path the viewing direction follows over time defines the scan pattern. Different mechanisms provide different scan patterns. One way to classify these scan patterns is based on the number of degrees-of-freedom (DOF) the mechanism provides. Since the scan space is two dimensional (θ and ϕ), the scanning mechanism would typically have either one or two degrees-of-freedom.

The 2-DOF scanner provides two independent controls (motors), which typically allow the scanner to completely cover a region in θ and ϕ. This is typical of fixed, ground-based scanners because the 2-DOF's allow an entire area to be scanned. Typical scan patterns might be a raster scan, where measurements are taken along a series of horizontal lines at successively different vertical angles. A 2-DOF scanner on-board the helicopter would be advantageous when having the helicopter either land

or hover in a single location and scan an area of the terrain. The disadvantage is that the 2-DOF scanning mechanisms are inherently more complicated and heavier than the 1-DOF scanner. Additionally, it is challenging to design a 2-DOF scanner which has the ability to scan over a large field of view. Many are restricted to relatively small (60° - 90°) windows, making it difficult to look both underneath and to the side of the helicopter.

The 1-DOF scanner provides only one control (motor) which controls the scan rate along a fixed scan pattern. The 1-DOF scanner allows only the locus of points along its scan pattern (typically simple patterns, such as a cone or plane) to be measured. on-board the helicopter, the 1-DOF scanner can be used to scan large terrain areas by using the helicopter itself as the additional degree-of-freedom that the scanner is lacking. The biggest advantages of the 1-DOF scanner are its simplicity and low weight, compared with those of the 2-DOF scanning mechanism.

Conclusion: The 1-DOF scanning mechanism is the preferred payload for a small autonomous helicopter. The 1-DOF scanning mechanism is capable of scanning virtually anything the 2-DOF scanner can, and is significantly simpler and lighter. The 1-DOF scanner does require the helicopter's motion while scanning, so the precision flight capabilities of the autonomous helicopter are critical for getting the scanner into all of the desired places.

2.4.3 Planar 1-DOF scanning configuration for autonomous helicopter mapping

One surprisingly capable scanning configuration for the autonomous helicopter is a planar, 1-DOF scanner which scans beneath the helicopter and to its sides [24]. This configuration is shown in Figure 2-5. The scan plane is normal to the direction of

forward flight, so successive cross sections of the terrain beneath and to the side of the helicopter can be collected.

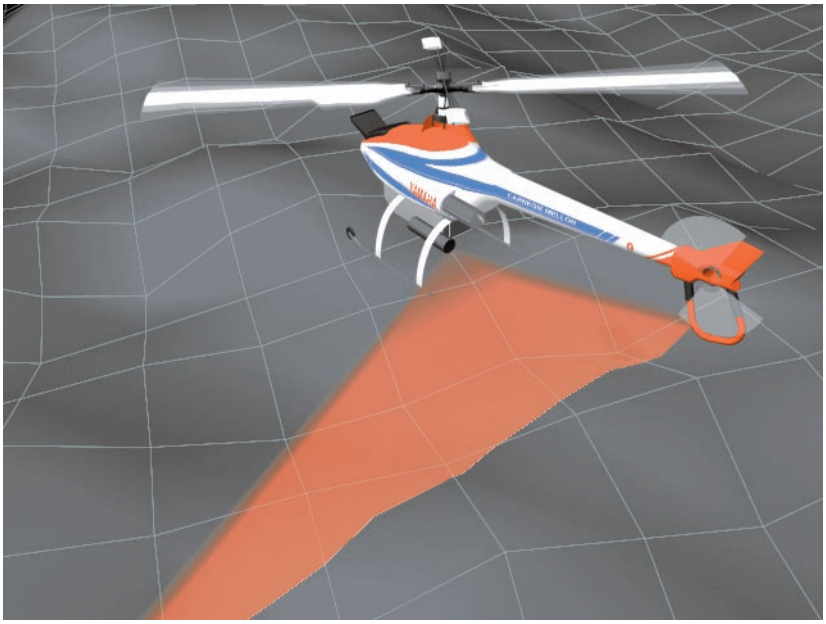

FIGURE 2-5. Planar, 1-DOF scanning configuration

The preferred implementation of the scanning mechanism (see Figure 2-6) consists of a spinning mirror which sweeps the terrain sensors across the scan plane. This configuration is mechanically simple, and can be designed to be lighter than other

scanners. It is also adequate for virtually all of the terrain modeling tasks we desire to accomplish.

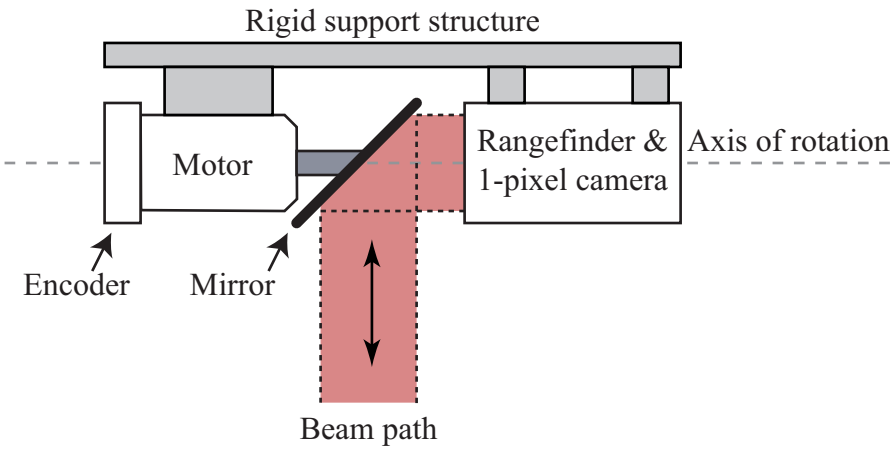

FIGURE 2-6. Planar, 1-DOF scanning mechanism

The speed of the mirror determines how often the plane is scanned. The laser rangefinder/1-pixel sampling rate and mirror speed together specify the sample density along the scan plane. The helicopter's forward velocity controls the distance between successive scans over the terrain.

Since each of these parameters (mirror speed, sampling rate, and helicopter velocity) can be controlled, the entire system can be readily reconfigured to handle different terrain modeling tasks. Slow mirror speed and slow forward motion generate highly dense models. Fast mirror speed and fast forward flight allow large areas to be mapped at a lower density. Also, since the scanner can look both beneath and to the side of the helicopter, vertical or even overhanging surfaces, such as buildings, can be mapped by flying alongside of them.

The 1-DOF planar scanning configuration also simplifies the measurement of beam direction. The angle of the spinning mirror, a directly measurable quantity, determines the beam's pointing direction within the scan plane.

One disadvantage of this configuration is that the entire plane is scanned, whether it is useful or not. A portion of the measurements taken each pass is wasted when the beam shoots into the sky or underside of the helicopter fuselage. Often, less than 50% of the scan plane is providing useful measurement. This inefficiency is acceptable, considering the advantages of a simple scanning mechanism. Also, the useful region of the scan plane changes depending upon the situation. Therefore, scanning the entire plane allows the region of interest to be easily changed in software.

Another obvious limitation of the downward-looking planar scanner is that it does not look ahead of the helicopter. Therefore, this design would not be useful for terrain avoidance tasks during forward flight.

2.5 Autonomous helicopter

The entire terrain sensor must be integrated on-board an autonomous helicopter. The helicopter provides support for the terrain sensor during terrain modeling. To globally localize each terrain measurement, the helicopter's autonomous flight systems provide critical position and attitude information.

Since we follow a rigidly mounted sensor design, localizing a point on the terrain requires the exact position and attitude of the terrain sensor as well as the terrain measurements taken from the sensor. All measurements from the terrain sensor (range, bearing, etc.) are relative, and therefore must be adjusted to compensate for the sensor's continuous motion. Figure 2-7 illustrates the kinematics of each terrain sample. The measurement process begins with helicopter localization in the world, thereby forming the coordinate frame of the terrain sensor. The 3D terrain location in the sensor's coordinate frame is found from this location by the terrain sensors measurements. The 3D terrain location can then be transformed to the world frame based on this coordinate frame..

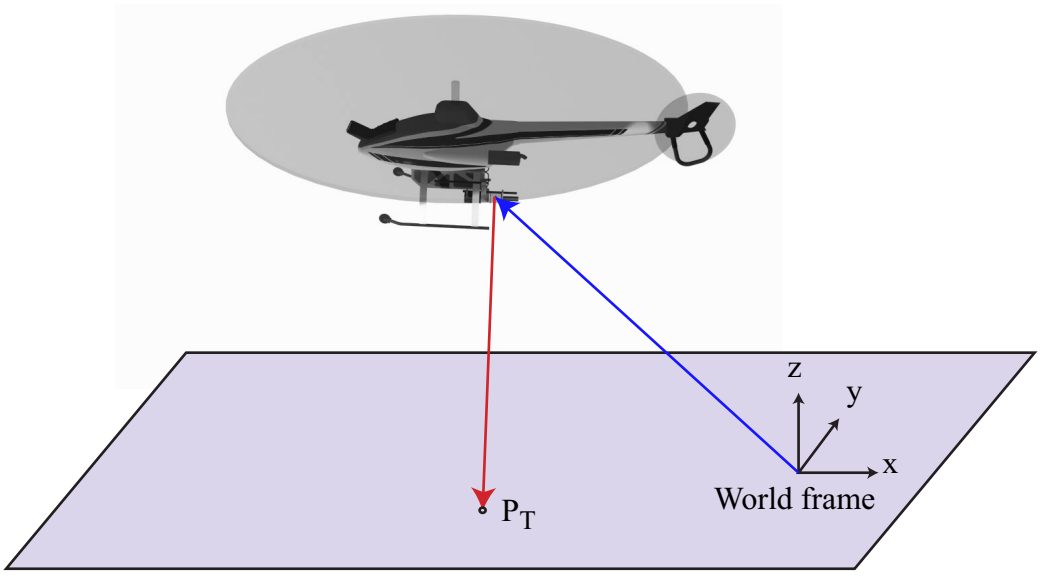

FIGURE 2-7. Localization of terrain from helicopter

The accuracy of each 3-D terrain measurement, and therefore the resulting terrain models, are dependent on both the accuracy of the terrain sensor (range, bearing, alignment, etc.), and also the accuracy of the helicopter's own measurement of position and attitude. Of particular concern is the attitude accuracy, since the effect of these errors grow as the target-sensor range increases.

2.5.1 Terrain sensor integration with autonomous helicopter system

Figure 2-8 is a block diagram showing the connection between the terrain sensor and autonomous helicopter for terrain modeling. While the terrain sensor measures the distance, bearing, and appearance of the terrain, the state estimator continually determines the exact 3-D position and attitude of the helicopter. The combination of the terrain and state data generates the accurate, world-referenced terrain measurements.

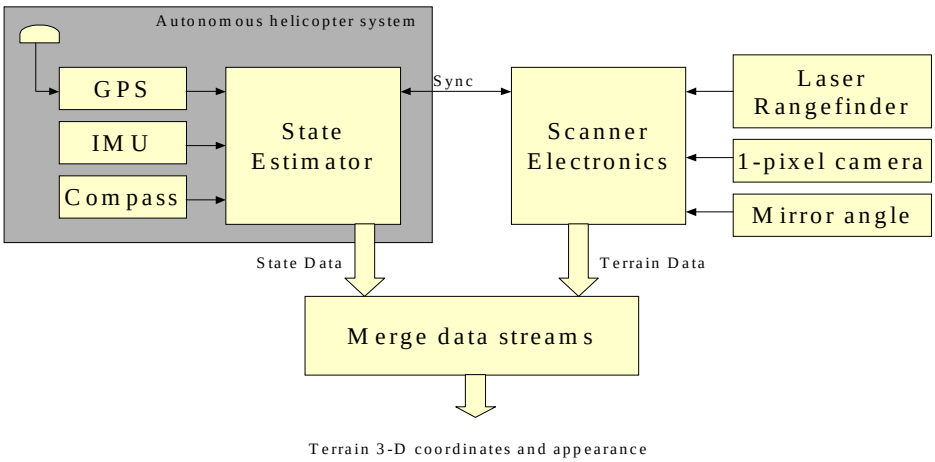

FIGURE 2-8. Block diagram of integrated terrain sensor and autonomous helicopter

The state estimator is a critical element of the terrain modeling system. It can use many different navigation sensors. The diagram shows a typical navigation sensor

suite, consisting of a differential global positioning system (GPS) receiver, inertial measurement unit (IMU), and magnetic compass.

The specific navigation sensors and techniques used are not of great importance to the terrain modeling system; rather, it is the achievable performance requirements which are critical. In this thesis, we rely on the existing autonomous helicopter system to provide the required state estimation capabilities. We do not concentrate on the details of state estimation. Instead, in the remainder of this section, we concentrate on the state estimation system characteristics and attributes which impact the terrain modeling process.

A secondary issue we consider when integrating the terrain sensor into an autonomous helicopter is that of providing a mechanism for the flight systems to use the terrain models onboard in real time. While this feedback of terrain models is not the primary focus of this thesis, it is strongly desirable to have these real time terrain modeling capabilities for obstacle detection, terrain referenced flight, and aiding various onboard perception systems. We must develop this real time terrain modeling capability to allow further research in these areas.

2.5.2 State estimator design objectives

The final accuracy of our terrain models is critically dependent upon the accuracy and capabilities of the onboard state estimation system. A number of important issues must be considered to determine whether a state estimation system is suitable. We identify four primary requirements for the state estimation system. The state estimator must:

1. *Measure the position and attitude of the helicopter with sufficient absolute accuracy to meet the terrain modeling accuracy requirements.*
2. *Provide a method for synchronizing the state estimates with external devices, i.e., the terrain sensor.*
3. *Estimate the uncertainty in the state estimator's position and attitude measure-*

ments.

4. *Provide the state estimates at sufficiently high rates to allow the helicopter motions to be tracked.*

2.5.3 Absolute accuracy requirements for state estimator

If we hope to build accurate terrain models, the state estimator needs to be accurate. The level of accuracy is dependent on the level of terrain modeling accuracy desired, and on the capabilities of the terrain sensor itself. In this sub-section, we discuss the need for absolute accuracy in state estimates, and look at the effect the state estimator accuracy has on the ultimate terrain model accuracy.

Absolute vs. relative accuracy

The type of terrain modeling we wish to perform requires the state estimator to be accurate in an absolute sense. Fundamentally, there are two measures of accuracy: absolute and relative. Absolute accuracy is based on the difference between the estimated state and the actual state. (Actual state is referenced to a ground-fixed world coordinate system in this case.) Relative accuracy is based only on the system's ability to estimate the change in state caused by the helicopter's motions from an undetermined starting point. Therefore, the estimates typically differ from the actual state. The relative accuracy of a system is always better than its absolute accuracy, since the absolute measurement includes all of the relative error sources in addition to error sources which constantly bias the system.

The importance of absolute accuracy for terrain modeling can be easily seen in a simple 2-D example. Figure 2-9 shows a situation where the helicopter measures a single point on the terrain at two different times. At time 1, the helicopter is heading North and is West of the point, At time 2, the helicopter is heading South and is East of the point. The ideal (absolutely accurate) state estimator will show the heading to be North at time 1, and South at time 2. Likewise, the position estimate will be at the

correct point. From this information, the system correctly calculates the location of each terrain sample being at the same place.

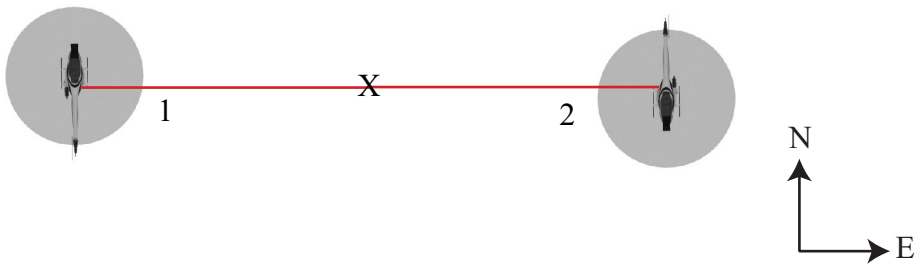

FIGURE 2-9. Actual situation where helicopter measures the same terrain spot two time

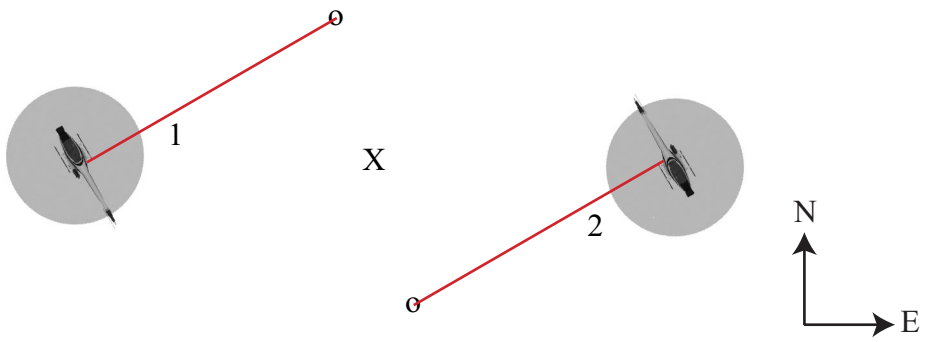

FIGURE 2-10. With heading offset, system thinks it measures two different spots

Next, consider the same situation, but assume that the heading estimate is no longer absolutely accurate. For illustration purposes, we offset the actual heading by 30º. Figure 2-10 illustrates how the terrain modeling system views the world. At time 1, the location of the helicopter is correct (position is still absolutely accurate), but the estimated heading is 30º West of North. The terrain sensor measured the ground at this time at location X, but the system uses the estimated position and offset heading to

calculate the location of the terrain point to be at the upper circle. At time 2, the helicopter is actually measuring the terrain at the point indicated by an X; but, since the heading has drifted, the computed location of the terrain point is wrong again, and is indicated by the lower circle.

It should be noted that, in the second situation, the relative accuracy of the state estimator was good. The heading did indeed change 180º between time 1 and time 2. However, the absolute accuracy of the heading was very poor, since it was constantly off by 30º. This is why the terrain system must rely on an absolutely accurate state estimator.

Target-less system

A technique for exploiting a relative estimator is to place artificial reference targets on the terrain at known locations. When the helicopter scans the area, each time one of the artificial targets is visible, its known location can be used to correct errors in the state estimator. The result can be a significant improvement in terrain modeling accuracy when compared with areas where no targets are used. This is particularly effective for improving the absolute accuracy of a state estimation system which exhibits either a constant bias or a slowly drifting bias.

Many photogrammetric and laser-based surveying systems rely on targets in the area being surveyed to provide a high level of accuracy and a mechanism for registration with existing maps.

In this thesis, we are concerned only with terrain modeling situations where reference targets have not been placed in the terrain. Since the autonomous helicopter is most useful for accessing remote and difficult to reach locations, it is undesirable to require preparation of the environment prior to scanning. This demands that the state estimator be capable of measuring sufficiently accurate absolute coordinates to enable mapping without the aid of ground targets.

How accurate will the system be?

Determining the exact accuracy of the terrain modeling system is a complicated procedure that is intimately tied to the specific implementation details, sensor choices, and unknown flight conditions. However, by considering a simplified model of the system, we can gain useful insights into the factors that limit the terrain model accuracy.

Figure 2-11 illustrates a 2-D version of the terrain localization situation that occurs for each terrain measurement. The location of the terrain sensor ($\mathbf{X}_{sensor}$, $\mathbf{Y}_{sensor}$) in this 2-D world is measured by the state estimator. This 2-D situation easily extends to the 3-D case. The current pitch angle of the helicopter (θ) is also determined by the state estimator. Since the terrain sensor is rigidly mounted onto the helicopter, the terrain sensor is also pitched forward by this amount. The scanning mechanism in the terrain sensor is aiming the laser beam at an angle (ϕ) from the top of the sensor, and the laser rangefinder measures the distance ($\mathbf{R}$) from the sensor to the terrain.

The 2-D coordinate of terrain spot is given by:

$$\mathbf{X}_{terrain} = \mathbf{X}_{sensor} + \mathbf{R} \cdot \cos(\theta + \phi)$$

$$\mathbf{Y}_{terrain} = \mathbf{Y}_{sensor} - \mathbf{R} \cdot \sin(\theta + \phi)$$

The accuracy of the terrain localization is affected by the inaccuracies of all of these measurements. The errors in the resulting terrain measurements ($\mathbf{X}_{terrain}$, $\mathbf{Y}_{terrain}$) can be generally attributed to:

- An error in the position measurement ($\mathbf{X}_{sensor}$, $\mathbf{Y}_{sensor}$) that directly corresponds to an error of equal magnitude in the terrain measurement. A 5 cm error in position results in a 5 cm error in the terrain location.
- An error in the range measurement ($\mathbf{R}$) also has a direct effect on the error of the terrain measurement. A 5 cm error in range maps into a 5 cm error in terrain loca-

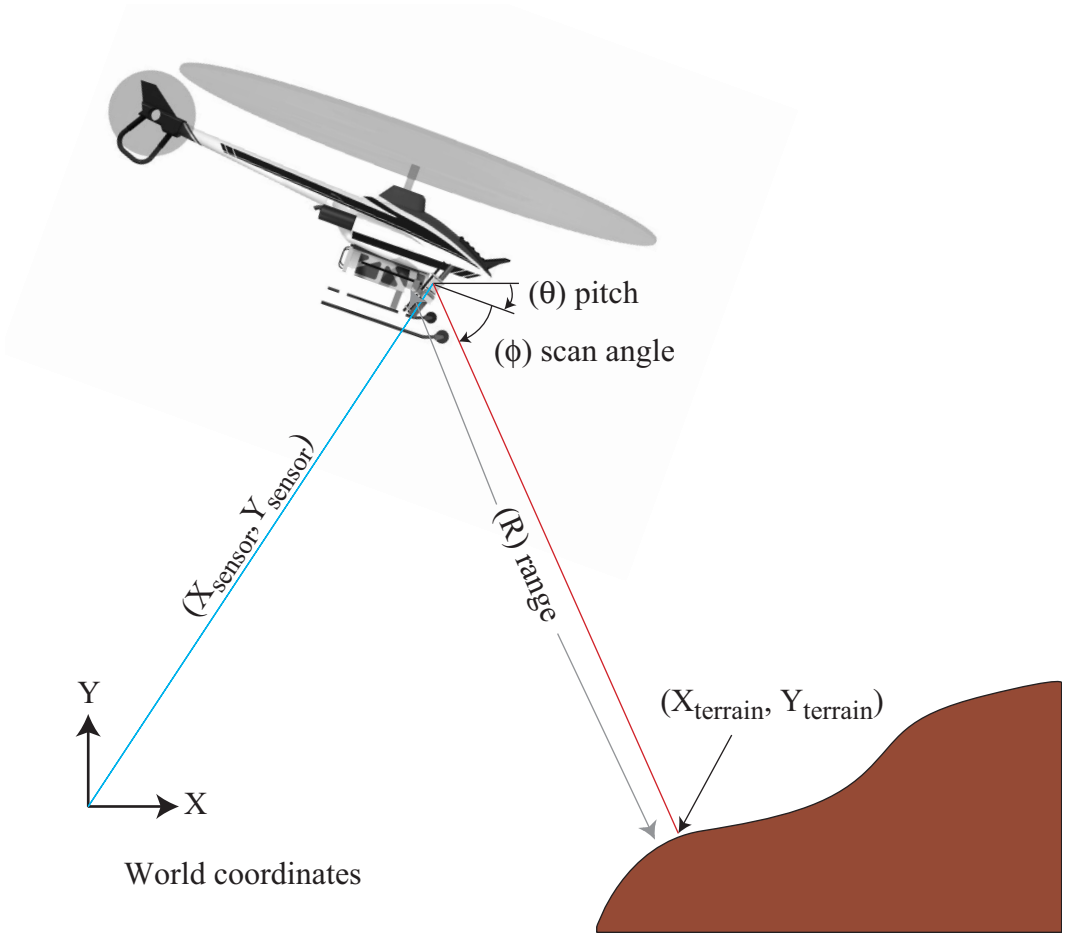

FIGURE 2-11. 2-D simplified model of terrain localization

tion, regardless of the direction of the laser beam.

- Errors in helicopter pitch (**θ**) or the scanner mirror angle (**ϕ**) are equivalent to each other, since they both affect the computed direction of the laser beam. An error in the direction of the computed laser beam is multiplied by the distance from the sensor to the terrain. Since the helicopter may often be located up to 100 meters from the terrain, the effect of these errors is significant. The error in these angles produces a corresponding error in terrain location which is equal to the range times the sine of the angular error. (**R**·sin(**θ**$_{\mathbf{err}}$)) For example, a 0.1° total error in these angles produces an 8.7 cm error in the terrain location when the terrain is 50

meters away.

The errors from each of these three sources combine to affect the terrain localization. Generally speaking, if one of the sources has significantly larger errors than the others, it will dominate, and limit, the achievable terrain modeling accuracy. Therefore, the system design must attempt to balance these errors for similar magnitude. Efforts for improvement should be directed to the dominant source of error. Clearly, it is not helpful to have highly accurate position measurements if the attitude measurements are poor.

2.5.4 Synchronization requirements for the state estimator

The previous discussion on localizing the terrain assumes that a terrain measurement (laser range, beam direction, and appearance) is taken with the helicopter at a known state (position and attitude). The helicopter is a highly maneuverable platform, and can move with high velocity and high angular rates. Therefore, the state is constantly changing, and we must register each terrain measurement with the actual helicopter state at the sample time. Failure to properly synchronize these measurements degrades the accuracy of the terrain model, particularly when the helicopter is in rapid motion.

An integrated synchronization system is required to ensure that these measurements can be properly aligned. One method for synchronizing the entire system is to use a central clock to timestamp each measurement, and later align the measurements based on this sample time. The state estimator can internally synchronizes its various sensors and provides an appropriate timestamp with each state estimation output. However, this is a non-trivial task, as the various sensors may run asynchronously and at various data rates. Interpolation and filtering in the state estimator can shift the effective time of a measurement and must be properly compensated for. For terrain modeling, our requirement is that, after all of its internal processing, the state estimator accurately reports the time at which each estimate is valid.

The terrain sensor itself must be designed to simultaneously perform all of its measurements (range, appearance, and beam direction), so that they are implicitly synchronized with each other. The terrain sensor hardware must record a timestamp for each measurement as well. With this information, the state estimator data can be accurately aligned temporally with the terrain measurements.

2.5.5 Uncertainty estimate requirements for state estimator

The terrain modeling system can benefit greatly from uncertainty measurements associated with each state estimate. While this is not a strict requirement for building a terrain model, the additional uncertainty information provided in real time can be used to refine, analyze, and merge the model.

The state estimator determines the state based on data from noisy sensors with imperfect characteristics, i.e., real-world devices. Therefore, the state estimates themselves may contain errors. We must describe the statistical nature of these errors in order to improve our terrain measurements. Traditionally, the basic measure of statistical error is the standard deviation.

Typically, a standard deviation estimate is associated with each parameter in the state estimate. Because we expect the state estimator to provide 3 components of position (x,y,z), and 3 components of attitude (roll, pitch, yaw), it should also provide the corresponding six standard deviations. The standard deviation and the corresponding covariance of the state are stored in the state covariance matrix.

Practically speaking, the standard deviations in the state estimate changes with various conditions over time. (GPS satellite location, previous motion path, time since visual landmark fix, etc.) Knowing the estimated standard deviation of each state variable, in addition to the state estimate, allows us to continually update the accuracy of each terrain measurement. In other words, we can establish a new covariance matrix for

each terrain measurement by combining the state covariance, measurement covariance, and the current configuration of the helicopter.

2.5.6 Update rate requirements for state estimator

The terrain modeling system needs to know the state of the helicopter each time a terrain measurement is taken. Typically, more than 10,000 terrain measurements are taken per second; therefore, the helicopter state must be determined at this rate as well.

Practically, the state estimator cannot measure state at these high rates, nor can it process and output estimates this fast. Instead, we expect the terrain sensor to provide high-rate terrain measurements, and the state estimator to provide less frequent updates of helicopter state. It is the responsibility of the terrain modeling system to interpolate appropriately between state estimates to determine the proper state for each terrain measurement.

The question is: How often should the state estimator provide a state estimate? In short, we need the state estimator to provide state often enough to fully describe the helicopter's motion and its primary vibration components. Due to the mass and moments of inertia of the helicopter, physics limits how quickly the helicopter's position and attitude can change. For the Yamaha R-50, the helicopter used in this research, an update rate of 20 Hz is sufficient to describe the primary motion of the fuselage.

Vibrations on-board the helicopter can be significant, causing the attitude of the terrain sensor to continually change. Figure 2-12 shows the roll of a Yamaha R-50 helicopter in flight. The plot shows the constant vibrations superimposed on the lower speed fuselage attitude changes. Our strapped-down approach need not attempt to completely isolate the vibrations from the terrain sensor, but instead can ensure that the state estimator and terrain sensors are rigidly attached to each other, and that both experience the same vibrations. Therefore, as long as the state estimator senses the

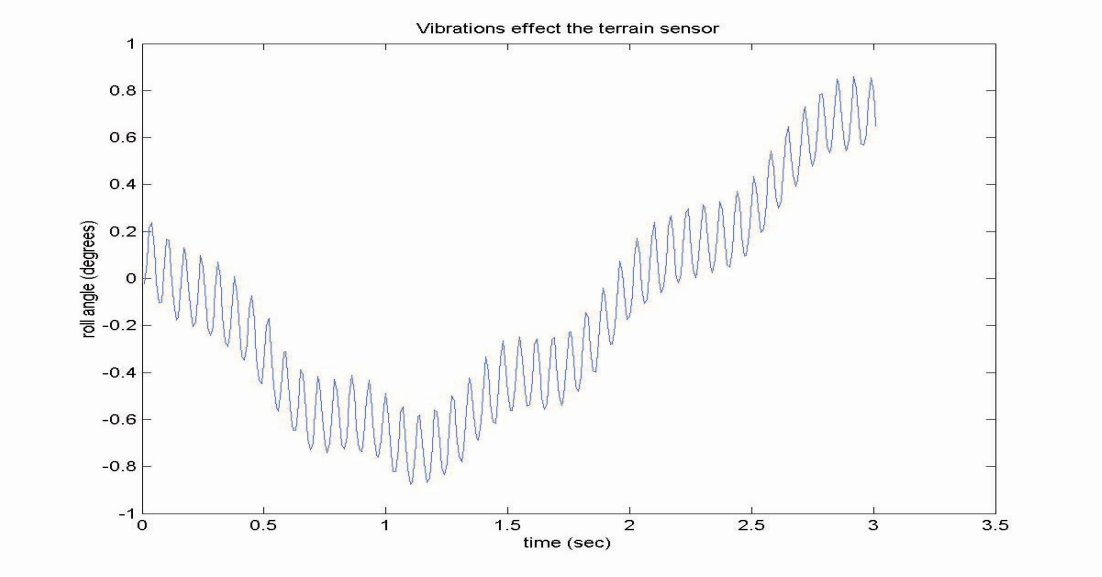

FIGURE 2-12. Fuselage attitude changes and vibrations affect the terrain sensor

movement caused by these vibrations, the terrain modeling system will compensate the terrain measurements appropriately.

The primary vibration source on the Yamaha R-50 helicopter is associated with the rotation of the main rotor blades, which rotate at 860 rpm (~15 rotations per second). Figure 2-13 shows the spectral content of R-50 attitude during flight. The 15hz vibration component is readily visible. At least six readings per rotation are needed to reasonably track these vibrations, requiring an update rate of at least 90 hz.

The terrain modeling system on a small helicopter does require a higher update rate than would be required on a full-sized helicopter. Full sized helicopters have more mass and inertia, and therefore react more slowly to disturbances. Also, full-sized helicopters operate with a slower main rotor speed, making it easier to track the vibrations.

Fortunately, as we will demonstrate in later chapters, current state estimator design can provide the necessary reaction time for accurate and effective terrain modeling.

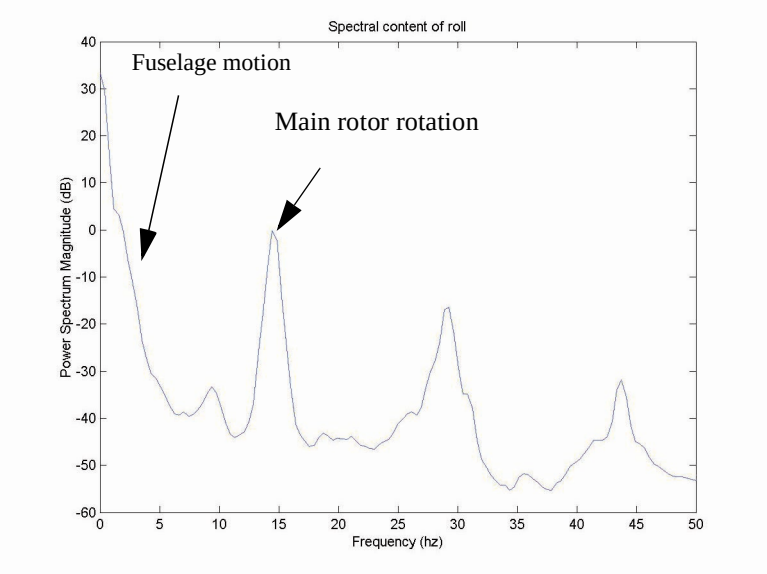

FIGURE 2-13. Spectral content of helicopter attitude shows principle vibrations

Chapter 3

Terrain modeling system: Implementation

This chapter presents an implementation of our terrain modeling system on-board an autonomous helicopter. The system follows the design described in Chapter 2.

We describe how we demonstrated our terrain modeling approach with a real-world working implementation. This implementation proves that the inherent challenges in building accurate terrain models by systems flying on-board a small helicopter can be overcome. The implementation also shows the value of measuring registered color with terrain geometry.

In this chapter, we first present an overview of the system implementation and then proceed to detail each of the system's five principle components: laser rangefinder, color 1-pixel camera, scanning mechanism, system-wide synchronization, and terrain measurement processing. Finally, we overview the CMU autonomous helicopter that supports this implementation.

3.1 Implementation overview

We built a prototype terrain modeling system based on the architecture presented in Chapter 2. The system can successfully build high accuracy (20 cm) color terrain models of the area below the host autonomous helicopter. The implementation focuses

on demonstrating high modeling accuracy, despite the inherent challenges of modeling from the helicopter.

This implementation demonstrates a finely tuned integration of main system components in a practical working system. A rigidly mounted 1-DOF planar scanner measures terrain near the helicopter during flight. The scanner traces the beam generated by a rugged commercial laser rangefinder, which rapidly (12,000hz) measures the distance to the terrain. Simultaneously, a color 1-pixel camera measures the appearance of the identical terrain point, as viewed from the helicopter. Since no suitable commercial sensor for the 1-pixel camera exists, we developed a custom sensor that can keep up with the laser to measure the color of the terrain within a narrow (<1º) field of view. To compensate for the helicopter's movements during flight, a precisely synchronized state estimator tracks the helicopter's motions and vibration.

Figure 3-1 shows the terrain sensor mounted onto the side of one of Carnegie Mellon University's autonomous helicopters. The terrain sensor scans a plane normal to the helicopter's forward direction, which spans a wide (>270º) unoccluded field of view (FOV). This configuration allows the scanning of terrain beneath, or to the sides of, the helicopter.

Figure 3-2 presents a block diagram of the prototype terrain modeling system. The scanning mirror is controlled to sweep out the scan plane up to 20 times a second, and an optical encoder tracks the mirror angle as it spins. A terrain measurement is triggered when the mirror reaches specific angle increments (500 or 2000 measurements / rotation) to approximately match the 12 KHz measurement rate of the rangefinder.

When a terrain measurement is triggered, the laser rangefinder (Riegl LD-90) emits a short (9ns) burst of light (λ=908nm), and uses a time-of-flight ranging technique to

FIGURE 3-1. CMU's Yamaha R-50 based autonomous helicopter with terrain sensor

find the distance to the terrain. Simultaneously, a custom 1-pixel camera is boresited with the laser beam to passively measure the apparent color of the same terrain point.

Our sensor design relies upon precise mechanical alignment of 1-pixel camera and laser rangefinder during mapping flights. To achieve this, our implementation rigidly mounts the 1-pixel camera to the laser rangefinder head, and provides a mechanism which initially brings them into precise alignment and then locks them down so that the alignment remains unchanged during flight.

The high sampling rate of the laser rangefinder requires the 1-pixel camera to also perform measurements at 12 KHz. Measuring color with the 1-pixel camera at these high rates is a significant challenge, since high sampling rates and achieving a narrow FOV are conflicting requirements. Because the narrow FOV and the short measurement interval allow only a small amount of light to be collected, a high sensor sensitivity is required. Our implementation demonstrates a reasonable trade-off for achieving useful models in spite of these restrictions.

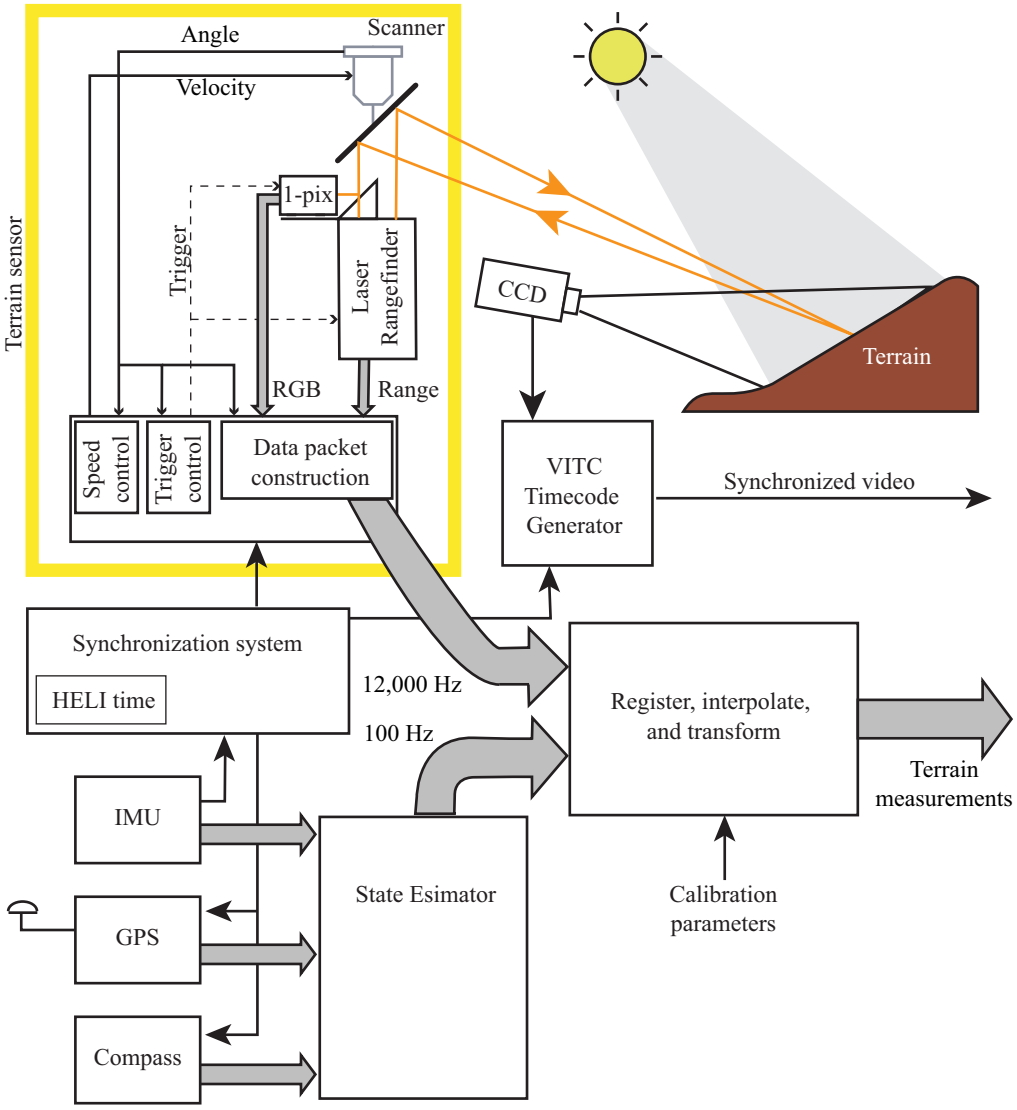

FIGURE 3-2. System overview diagram

The terrain sensor hardware collects the data relating mirror angle, range, laser reflectance intensity and color for each measurement, and sends this data to the onboard computing system for processing. This data allows the 3D location of the target to be determined relative to the terrain sensor. Localizing the target in a world-fixed reference frame requires the terrain sensor itself to be localized at the same time.

The autonomous helicopter's onboard state estimator tracks the helicopter's position and attitude using a strapped-down IMU (3-axis gyroscope and accelerometers), differential dual-frequency carrier-phase GPS receiver and a flux-gate compass. The state estimator fuses measurements from these sensors using a 12th order extended Kalman filter to provide estimated position and attitude measurements 100 times a second. The 100hz state estimate is sufficient to compensate for all body motions of the helicopter as well as for the most significant vibration components.

The precise time at which each of these sensors performs a measurement is regulated by a hardware synchronization system. The synchronization system contains a free-running clock which time-stamps each measurement for temporal alignment. To provide system wide synchronization, the time of each terrain measurement as well as that of other onboard sensor measurements, such as video imagery, are time-stamped.

The terrain processing system combines the terrain measurements with the state estimates to localize the target in the world. First, these two streams are temporally aligned and the state estimates are interpolated to determine a snapshot at the time each terrain measurement is taken. Geometric transformations then calculate the 3D terrain coordinates and appearance of the terrain. Further processing steps can use this data for specific applications.

The following sections further detail this terrain modeling system implementation.

3.2 Laser rangefinder

We chose the Riegl LD90 laser rangefinder [46] for the terrain sensor. The Riegl rangefinder is unique, since it separates the laser, receiver, and rangefinding electronics from the optical head. The electronics box is connected to the optical head by a pair of fiber optic cables which transfer the transmitted laser pulse and the received laser reflection from the target as shown in Figure 3-3. This is an extremely valuable feature. This allows us to mount only the smaller optical head in the scanning mechanism, thereby allowing the scanner to be smaller and providing more options for locations where it can be mounted on the helicopter. Also, since the electronics box is relatively heavy (5.25 lbs) and can be mounted separately, the box can be used to counterbalance the weight of the terrain scanner by mounting it on the opposite side of the helicopter. One disadvantage of this setup is the fragility of the fiber optic cables connecting the electronics box and optical head. Because these cables are so fragile, they require additional care during helicopter transport or maintenance.

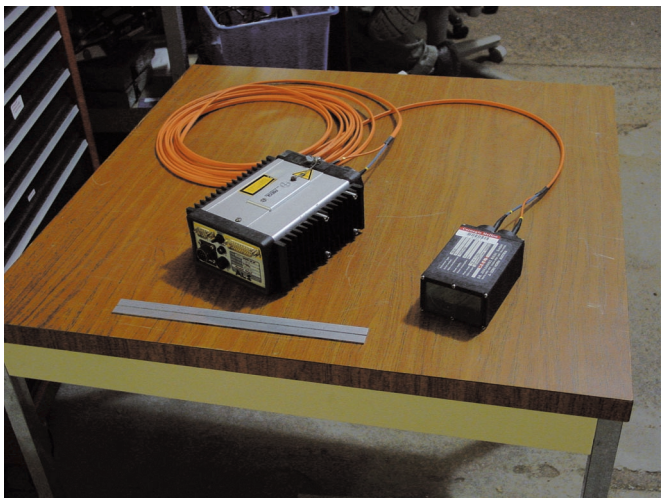

FIGURE 3-3. Picture of remote head Riegl LD90 rangefinder

The LD90 is a time-of-flight rangefinder that performs up to 12,000 range measurements per second ($RATE_{laser}$ = 12,000). This maximum rate is important, since it defines the highest rate of terrain measurements which can be captured. Therefore, the 1-pixel camera will need to be designed to also perform 12,000 samples per second. The LD90 can either be automatically triggered at 12 KHz, or externally triggered by a digital input. We generate an external trigger which dictates the precise time each measurement is taken.

The LD90 has a maximum range ($RANGE_{max}$) of approximately 100m on typical non-cooperative targets (natural terrain). This is a sufficient distance for scanning terrain near the helicopter without requiring dangerously close flight. Beyond this range, some strongly reflective objects are measurable, but we cannot expect to obtain a reliable measure of the terrain. The range measurement is a 16-bit digital output with a resolution of 0.4cm. The published accuracy of the range measurement is 2.5cm standard deviation.

The optical head provides a duplex rangefinding configuration, where the receive and transmit channels are located side-by-side. Figure 3-4 shows the optical head with the

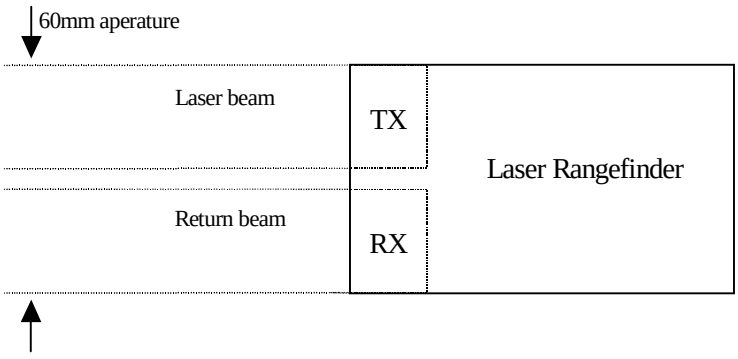

FIGURE 3-4. Duplex laser rangefinder configuration

transmit and receive beams.

The infrared (λ=908nm) laser beam has a 25mm diameter aperture at the optical head (D_{laser} = 25mm) and has a divergence angle of 3.2mrad (θ_{laser} = 0.18º). Compared with other laser rangefinders, the beam is significantly larger in diameter, resulting in a larger spot size on the terrain. The expected spot size over distance is shown in Figure 3-5. At the maximum 100m range, the spot size exceeds 30cm. This large size limits how closely spaced consecutive measurements can be independently taken. However, as will be discussed later in the chapter, the scanning mechanism typically moves the laser beam by a greater distance between measurements.

An important advantage of the large spot size is that the sensor is more robust to small amounts of dirt accumulating on the optics and to small particles in the air, such as light rain or snow. Since these constitute a small percentage of the area the laser is spread over, their minute reflections are not detected by the sensor as obstacles in the beam path. Smaller laser beams can be significantly blocked by even a single blemish on the optics. The other significant advantage of the large spot size is that it helps to ensure eye safety as was discussed in Chapter 2. The LD90 is not considered to be eye safe when stationary, but is eye safe when continually being redirected by the scanner.

In addition to measuring the range, the rangefinder outputs an 8-bit laser reflectance measurement. This provides a qualitative measure of signal strength, and can be used as a range reliability measure. High values prove that sufficient reflection was available to provide a reliable range measurement. Also, as we will discuss in the next chapter, the laser reflectance measurement is especially valuable for detecting the presence of a target with exceptional reflectivity (such as a retroreflective calibration target) within the laser’s spot.

The LD90 is a first-pulse rangefinding system. Therefore, it is unable to ignore significant reflections prior to hitting the terrain, and this requires that no paths exist for the transmitted beam to prematurely return to the receiver from the mechanism or protective enclosures.

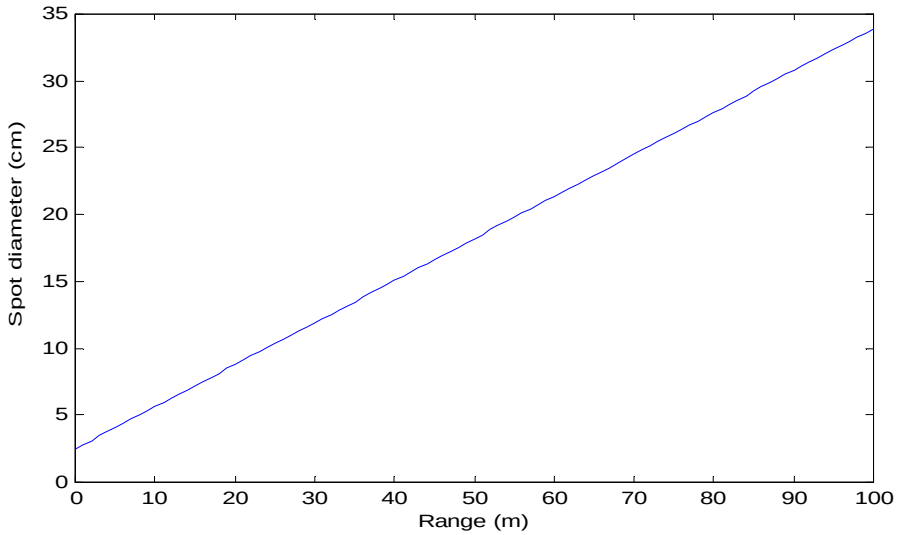

FIGURE 3-5. Laser spot diameter vs. range

3.3 1-pixel camera

In this section, we present details of a prototype 1-pixel camera developed for this thesis. Many different attributes could have been selected, but we chose to implement a color version of the sensor. Color is a good choice since a great deal of human understanding of the environment is based on color cues. Incorporating color appearance into the terrain model, therefore, significantly enhances the model's value.

To capture color, we build the 1-pixel camera with three sensing channels: one that measures the red component of light coming from the terrain, one that measures the green component, and one that measures the blue component. This is a common method for measuring color, since three properly chosen measurements (red, green, and blue) can almost completely represent all humanly identifiable colors.

This section begins with a description of the optical design developed for each of the sensing channels, based on the guidelines established in Chapter 2. Next, we discuss the actual mechanical implementation, including critical adjustments for focus and alignment. We then present the sensor's electronic data acquisition system. Finally, we examine the actual components used in the implementation to calculate a set of important characteristic parameters for this specific design.

3.3.1 Combined sensor configuration

In following our design approach, the 1-pixel camera is rigidly mounted and then boresited with the laser rangefinder. Since the optical paths of both the rangefinder and the 1-pixel camera must pass through the clear aperture of a scanning mechanism, the combined configuration of the two sensors significantly influences the size and weight of the scanning mechanism. The combined configuration must be considered during the design of the 1-pixel camera.

To prevent increasing the required clear aperture of the scanning mechanism further than is required for the rangefinder, our implementation optically overlaps the viewing paths of the 1-pixel camera with the laser rangefinder, to fit within a 60mm aperture, as shown in Figure 3-6.

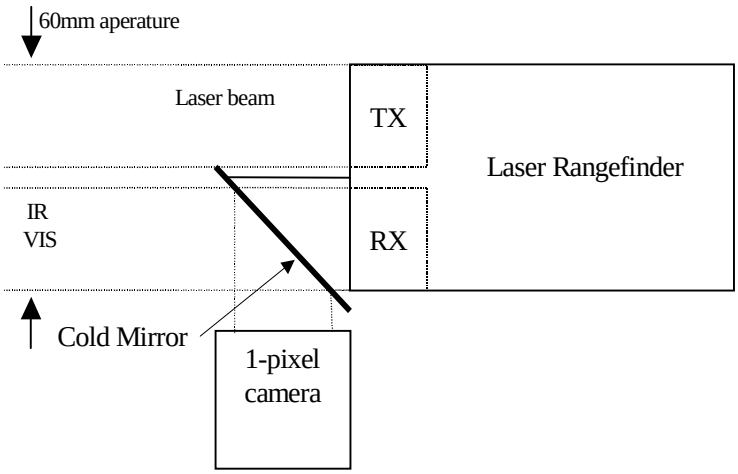

FIGURE 3-6. Configuration of rangefinder and 1-pixel camera

A wavelength selective mirror allows the 1-pixel camera and the laser rangefinder receive channel to share a common aperture. This “cold mirror” makes efficient use of the incoming light by reflecting the visible light rays toward the 1-pixel camera while allowing the infrared light to pass through it and into the laser rangefinder. Since the laser rangefinder operates strictly in the near-infrared region (908nm), the addition of the mirror does not significantly affect its operation.

In our implementation, we overlap only 1/2 of the scanner’s aperture, the half which contains the laser receive channel. This avoids any interference with the outgoing laser beam. We experimentally verified that passing the return beam through a cold mirror only negligibly reduces the maximum range of the rangefinder. The cold mirror is a good compromise, given the size and weight savings afforded by a smaller scanning mechanism. This approach, however, limits the available aperture for the 1-pixel camera to only 1/2 of the 60mm semi-circle.

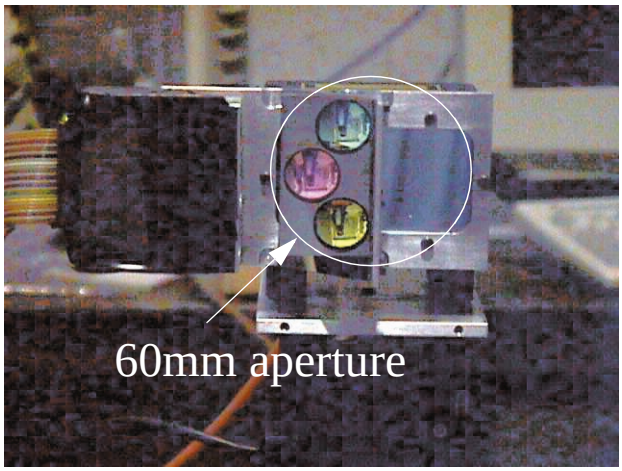

FIGURE 3-7. Front view of combined sensor

We must fit the three sensing channels into one half of the 60mm scanner aperture. Figure 3-7 shows a front view of the combined sensor with the scanner's clear aperture shown. Here we see that all of the needed optical paths have been arranged to fit within one-half of the clear aperture. The laser transmit lens is visible in the right half of the aperture. On the left half of the aperture, the three front lenses of the color-sensing channels appear. (The red channel is on top; the green channel is in the middle; and the blue channel is on the bottom). Notice that the lenses have been placed to maximize the lens diameter, while still having a clear view through the aperture. The color-sensing channels are located in the black box to the left of the aperture, and their reflection off the cold mirror is visible. Not visible in the left half of the aperture is the laser rangefinder receiving lens that is behind the cold mirror.

This configuration successfully combines the 1-pixel camera and laser rangefinder to share a 60mm scanner aperture. The decision to fit the three color sensing channels into one half of the 60mm aperture dictates the maximum channel diameter of 22mm for each channel, and establishes the relative positioning of the channels

Parameter summary	Channel diameter	22mm
	Scanner clear aperture	60mm

3.3.2 Optical configuration of a single channel

All of the color-sensing channels are essentially the same. Figure 3-8 shows the basic optical configuration of a single channel in the 1-pixel camera. As an aid to understanding the optical interaction of the 1-pixel camera with the Riegl laser rangefinder, both are shown, viewing the scene through a cold mirror. Each channel consists of a lens/pinhole pair which spatially controls the channel's FOV, an optical filter to select the desired colors of light, and a photodiode to convert optical energy into a photocurrent output. The photocurrent is ultimately measured by the 1-pixel camera electronics to determine light intensity.

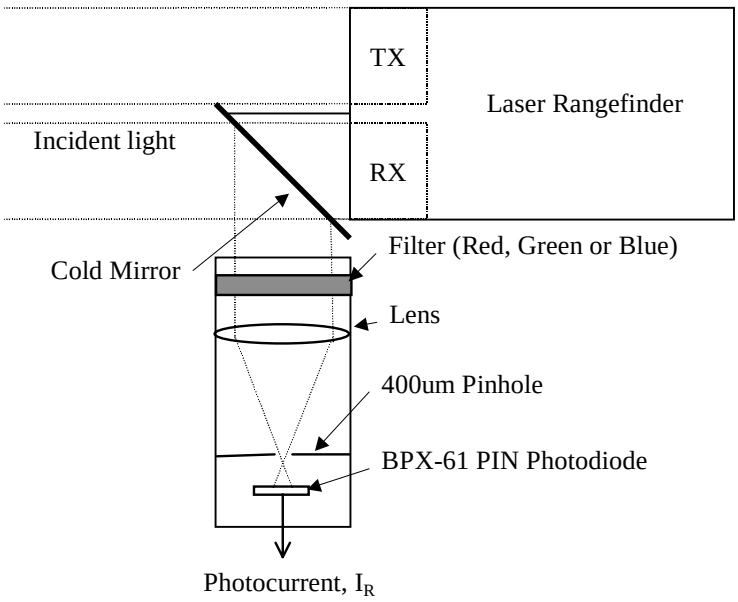

FIGURE 3-8. Optical configuration

Of particular interest is the path along which the visible light rays, arriving within the sensor's FOV, enter the sensor, and ultimately reach the photodiode. The goal of the optical channel is to funnel all properly colored incoming light along its viewing axis into a photodiode for detection, while rejecting all miscolored lights and off-axis rays. The optical configuration we use to achieve this is as follows:

- Incident light first hits the cold mirror. The cold mirror is mounted 45º to the light

path in front of the laser rangefinder. It allows infrared light to pass through (into the laser rangefinder), and reflects light in the visible range toward the color sensors.

- The visible light beam hits one of three color filters. The filters are composed of thin-film interference filter coatings on a glass substrate. They allow a certain band of light wavelengths to pass through while all other light reflects back out of the senor. Each sensing channel has a different filter. One passes red light, one blue, and one green.
- This visible light next encounters a lens, which focuses the desired collimated light beams toward a single point at the center of a pinhole.
- A thin metal pinhole with a small hole (400 um) is placed at the point where the light path converges. This is effectively an aperture stop that allows only light rays passing through that specific convergency point to cross the metal separator.

Any light passing through the pinhole hits the photodiode, which generates a photocurrent proportional to the light intensity.

3.3.3 Desired Field of view

We want each channel of the 1-pixel camera to measure light coming from a small spot on the terrain. To achieve this, the sensor needs to have a narrow FOV, as discussed in Chapter 2. Ideally, we want to match the FOV of the 1-pixel camera to the divergence angle of the laser rangefinder, so the color and range measurements are taken from the same sized spot on the terrain. Achieving this requires a 0.18º FOV. Unfortunately, due to restrictions imposed by the limited lens diameter and short integration period, not enough light can be collected from this narrow FOV to produce a good image in typical daylight. Therefore, to alleviate this problem, we chose to increase the sensor's FOV to approximately 1º. This provides a substantially increased signal (~25 times larger). Clearly, this compromise means that the 1-pixel camera is measuring the terrain appearance over a larger spot size than the laser rangefinder uses, but in

practice, successive samples are spaced far enough apart across the terrain to minimize overlapping measurements, thereby resulting in a representative color appearance measurement centered around each laser range measurement

Parameter summary	FOV_{sensor}	1^o

3.3.4 Lens selection

We must select an appropriate lens/pinhole combination which provides the desired FOV based on other sensor specifications; most important is the lens focal length. There are a number of factors which must be considered in selecting the lens focal length. We would like to keep the focal length as short as possible since the focal length dictates the physical length of the sensing channels. Long sensing channels would make the sensor larger and heavier.

On the other hand, using a shorter focal length presents several challenges which we must overcome:

- *Alignment is more difficult*: The direction of the principle viewing ray becomes more sensitive to any misalignments with decreasing distance from lens to pinhole. The principle viewing ray is centered in the lens and pointing in the direction in which the sensor is looking. In this case, it is a ray from the center of the pinhole through the center of the lens. The short focal length makes it more difficult to point the sensor in a desired direction and to keep it pointing there despite vibrations and other forces. We can overcome this problem by designing a mechanical adjustment that can compensate for these misalignments while being capable of rigidly locking down the adjusted optics. Later in this section, we will present this mechanical design.
- *Focusing is more difficult*: The distance from the lens to the pinhole determines where the channel is focused. Depending on where we want the sensor to be focused, this distance must be changed. Short focal length lenses require finer

adjustment ability to achieve proper focusing, and are more susceptible to the focus being affected by vibrations and other forces. We can overcome this problem by precisely adjusting the focus in the lab, and locking it in place. Later in this section, we present this approach.

- *Chromatic aberrations are more significant*: Chromatic aberrations are non-idealities which occur in real world lenses, where the focal length of the lens is not a constant distance, but rather varies with the wavelength (color) of the light passing through. With significant chromatic aberrations present, a sensor which is focused to properly view a red object (λ=700nm), will be out of focus when viewing a blue object (λ=450nm). As the focal length of a lens is reduced, the magnitude of the effect of chromatic aberrations increases.

  One approach for reducing chromatic aberrations is to use a doublet lens. The doublet consists of two lenses, each made of a different type of glass; these lenses are cemented together to act as a single lens. The doublet significant reduces chromatic aberrations for a given diameter and focal length lens. Most optical designs support the additional complication of the doublet to achieve a consistent focus independent of wavelength.

  However, in our color 1-pixel camera implementation, each sensing channel is attempting to image only one portion of the visible spectrum. So, the red channel needs to focus only on red colored light, thereby significantly reducing the effect of chromatic aberrations. Therefore, if we independently focus each channel for its sensing color, a simple singlet (standard one-piece lens) is sufficient.
- *Spherical aberrations are more significant*: Spherical aberrations are another non-ideality which occurs in many real-world lenses. The aberrations cause the focal length of the lens to vary depending on the light ray's radial distance from the lens center. This variance occurs because the standard lens-making approach can generate only spherically shaped surfaces, while the mathematically correct shape for the surface is slightly aspherical. As the focal length of the lens is reduced, the significance of the spherical aberrations increases, thereby reducing sensor focusabil-

ity.

Aspherical lenses which do not have spherical aberrations are available; however, they are typically of either low quality (molded plastic lenses) or are extremely expensive (diamond-turned glass lenses). Instead of turning to an aspherical lens, we use a plano-convex lens with the planar (flat) surface facing toward the pin-hole. This lens configuration provides the smallest amount of spherical aberrations available for a spherical surface lens for a given diameter and focal length.

We have previously determined that the maximum diameter for each channel is 22mm. While meeting the above criteria and minimizing focal length, we select a 22.5mm focal length plano-convex singlet lens with a diameter of 22mm. Because this is certainly a short focal length for such a large diameter lens, we must deal with the short focal length lens issues discussed above.

A final consideration in selecting the lens is whether an anti-reflective (AR) coated lens should be used. Typically, when a light beam traveling through the air enters glass (at the front surface of a lens), 4% is reflected away and lost. The same 4% reflection occurs when the light transitions from the glass back into the air at the back surface of the lens. This means that 8% of the light is lost each time light passes through an uncoated lens, ensuring that the lens is less than 92% efficient. AR coatings are designed to reduce this reflection, and can be applied to the front and back surfaces of the lens. AR coated lenses can approach 99% efficiency for light in their design wavelengths. AR coated optics do cost more than their uncoated equivalents, but the difference is marginal. The optics also must be handled carefully, since the thin AR

coating can be easily damaged. We chose to use an AR coated lens to help minimize the losses sustained when going through the lens.

Parameter summary	Lens diameter	22mm
	Lens focal length	22.5mm
	Lens type	Plano-convex
		Singlet
		AR-Coated

3.3.5 Pinhole selection

The FOV of the sensor is controlled by two elements: the focusing lens and the pinhole. These components must be appropriately selected and positioned in order to provide the desired sensing channel. Since we have already selected the focusing lens, we must find the proper pinhole size to produce our desired 1º FOV.

A raytrace diagram (Figure 3-9) helps in determining the pinhole size of the sensor. We assume that the sensor is focused at infinity, so the pinhole is spaced behind the lens at a distance equal to the lens focal length. This spacing focuses parallel rays of light (from distant objects) onto the pinhole plane. The sensor ultimately collects only the portion of the image that falls inside the pinhole. The field-of-view describes the full angular acceptance range of the sensor, and it is determined by finding the angle between incoming light rays which image to opposite edges of the pinhole. The FOV can be calculated as:

$$FOV_{sensor} = 2 \cdot \tan^{-1}(0.5 \cdot \text{Pinhole diameter} / \text{focal length})$$

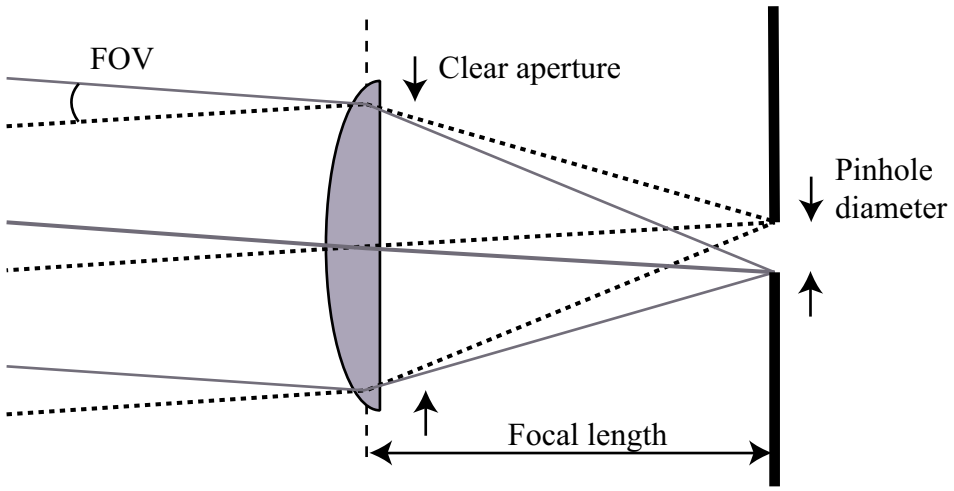

FIGURE 3-9. Raytrace of lens-pinhole configuration

Based on this equation, we need to use a 0.392mm diameter pinhole with our 22.5mm focal length lens to achieve the 1º FOV we selected for the sensor. We use a 400um pinhole, the closest commercially available pinhole diameter.

Parameter summary	Pinhole diameter	400um

3.3.6 Clear aperture

The clear aperture is the lens area that primarily collects light and images it onto the pinhole. In the configuration used for the 1-pixel camera, this is simply the area of the lens not covered by its mount. The total area of the clear aperture is $A_{sensor.}$

In the our implementation, the lens is 22mm in diameter. In the sensor mount, the clear aperture is restricted to 18mm, since the outer 2mm around the lens is used to hold it, and cannot gather light.

Parameter summary	Clear aperture diameter	18mm
	A_{sensor}	$254mm^2$

3.3.7 Optical filter selection and placement

In Chapter 2, we said that the ideal color 1-pixel camera would have a spectral sensitivity function which matches another color imaging sensor, such as a specific 3-CCD camera. The optical filter has a significant effect on the channel's spectral sensitivity, so achieving the proper overall spectral sensitivity is the primary objective when selecting the filter. The spectral sensitivity is not the only concern; we must also consider the type of filters and their effects on the overall design.

A suitable filter maximizes efficiency in the passband, to allow as much light as possible to reach the photodiode and maximizes attenuation in the stopband, to prevent all other wavelengths from passing through. For this reason, we select a set of three dichroic interference filters for the sensing channels. Dichroic interference filters use multiple layers of precisely deposited thin-films on a glass substrate to allow certain wavelengths of light to pass through with little attenuation, while other wavelengths are reflected from its surface. The high efficiency in the passband (> 80%) and low leakage in the stopband (~0%) make these filters particularly suitable for the sensor.

Figure 3-10 shows the percentage of transmission over wavelength for each of the three filters [29]. Later in this section, we will examine how these spectral transmittance curves interact with the other optical elements to achieve the desired color spectral sensitivity response.

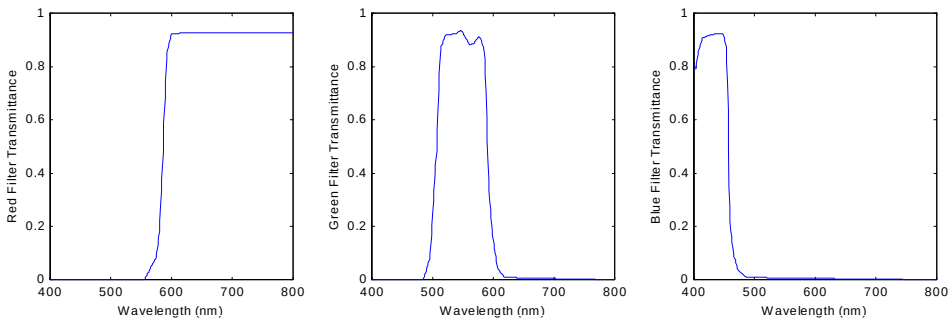

FIGURE 3-10. Transmittance curve T(λ) for the optical filters

One important factor to consider when using an interference filter is its dependence on incidence angle. The transmittance curve published with the filter (Figure 3-10) is valid only for light rays that hit the filter with an incidence angle of 0º (normal to the surface). Rays passing through the filter at other incidence angles are filtered by a different transmittance curve. Angles of incidence that exceed 20º have a dramatic effect on the transmittance curve, allowing the filter to pass light in totally different bands.

The convergence point is significant for our design, since, due to the very short focal length lens chosen, the light rays converge a short distance after they pass through the focusing lens. Placing the filter after the focusing lens requires the rays to pass through the filter at various angles of incidence. Rays passing through the center of the lens, the so called paraxial rays, will still have a 0º angle of incidence; however, rays passing through the outer edges of the lens, the so called marginal rays, will pass through the filter with a steep angle of incidence. In our design, the marginal rays are 9mm from the center of the lens (aperture is 18mm diameter), and the focal length of the lens is 22.5mm, resulting in an incidence angle of $\tan^{-1}$(9mm / 22.5mm) = 21.5º. Due to the varying angles of incidence, the result of placing the filter after the lens would be a sensing channel with a significantly distorted spectral sensitivity.

To ensure that this incidence angle effect does not compromise the color accuracy of the 1-pixel camera, we place the filter before the focusing lens in each channel, as shown in Figure 3-11. In this location, all of the light rays which will ultimately be measured by the photodiode are arriving nearly normal to the filter surface. In fact, since all light rays of interest are arriving within the FOV of the sensor, the largest incidence angle should be half of the FOV angle, or 0.5º. Within this range, the filters operate as expected.

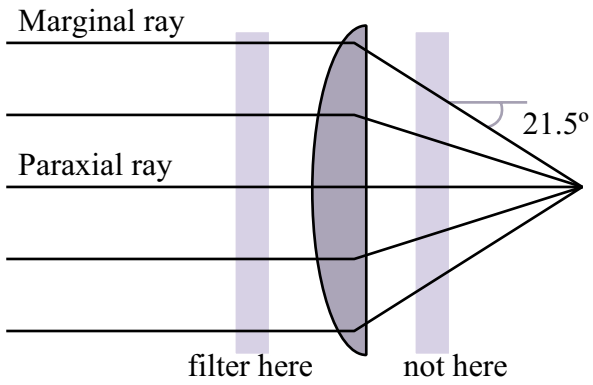

FIGURE 3-11. Interference filter should be placed in front of lens

3.3.8 Photodiode selection

The photodiode measures all of the light that successfully passes through the pinhole, and generates an electrical signal which is the sensing channel output. For the given amount of light passing through the pinhole, we desire to select a photodiode which generates a significant photocurrent signal and minimizes noise.

The photodiode generates a photocurrent output proportional to the optical power it receives. The diode has a wavelength-dependent sensitivity; Figure 3-12 shows the spectral sensitivity curve for the Osram BPX-61 photodiode [30] used in all three channels of the 1-pixel camera. This figure shows how the sensitivity of the photodiode changes over the visible band of wavelengths. The spectral sensitivity curve is of particular importance, since it indicates the amount of photocurrent output (Amps) generated for a certain optical input power (Watts) as a function of wavelength. The photodiode is more sensitive toward the red end of the spectrum than toward the blue end.

Most silicon-based photodiodes have a similarly shaped spectral response curve. The BPX-61 photodiode was selected because it has a slightly higher sensitivity in the blue region than others considered, and because it is a PIN photodiode. The PIN photodiode

has significantly less capacitance than the standard PN structure. Lower capacitance is desirable because it reduces the noise generated in the photodiode.

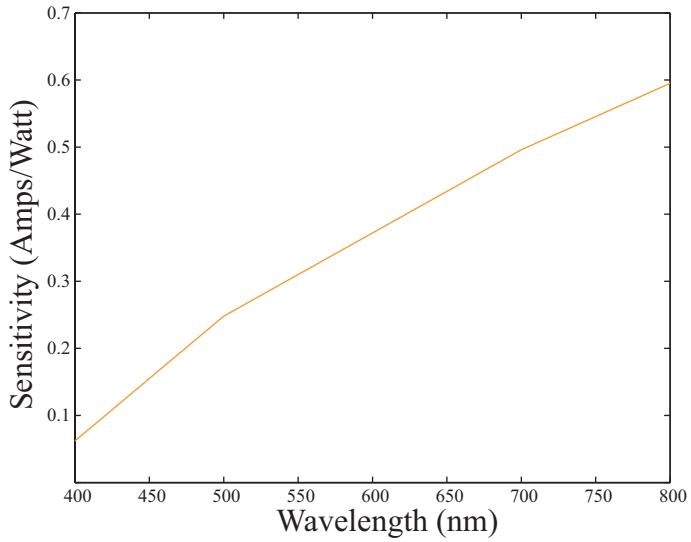

FIGURE 3-12. Relative sensitivity of photodiode

For simplicity, our design uses only a single focusing lens in each optical sensing channel, and that lens generates an image of the terrain spot onto a plane at the pinhole. The size of the pinhole diameter selects the sensor FOV. The photodiode is located behind the pinhole to collect all of the light passing through, as shown in Figure 3-13. As we mentioned in the filter discussion, the marginal light rays converge at an angle of 21.5º towards the pinhole. Therefore, the marginal light rays leave the pinhole and diverge with a half-angle of 21.5º.

Due to this rapid expansion of the light leaving the pinhole, we must place the photodiode physically close to the pinhole, and that ensure the sensing area of the photodiode is large enough to capture all of the light. Typically, the sensing element of a photodiode is encased in a metal or plastic package and is recessed approximately 2mm from the front surface of the package. Assuming that the front surface of the photodiode package is touching the pinhole, the light expands from the 0.4mm pinhole

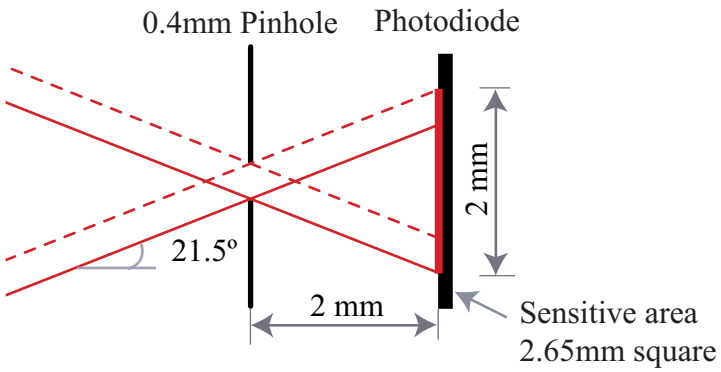

FIGURE 3-13. Illumination of photodiode

diameter to 2mm diameter. Therefore, we choose a photodiode with a sensing area larger than 2mm. The selected BPX-61 photodiode has a square sensing area with 2.65mm sides, thus ensuring that all light is collected.

Photodiodes do have a limited acceptance angle for incident light. The marginal light rays which reach the photodiode have an incidence angle of 21.5º. To ensure that the photodiode collects these rays, the photodiode must have an acceptance angle of at least 21.5º. Since the selected BPX-61 photodiode has an acceptance angle of +/- 55º, it easily handles this requirement.

3.3.9 Mechanical construction

We examine the mechanical construction of the 1-pixel camera supporting the selected components. The two primary goals in creating the mechanism are small size and weight, and ruggedness to handle the harsh helicopter environment.

Figure 3-14 shows a cross-section of one sensing channel. One of the most critical requirements for the mechanism is to hold the optical components exactly where they need to be, despite disturbances from vibrations and other forces. Due to our selection of a short focal length lens, the overall channel length is short and the entire sensor can

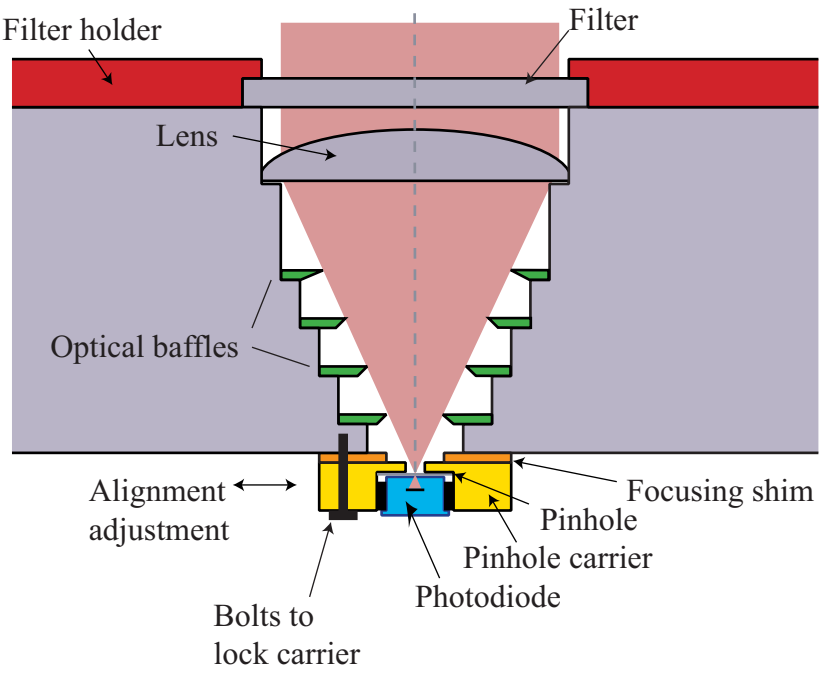

FIGURE 3-14. Optical channel construction

be small. This allows us to machine the three sensing tubes out of a solid piece of aluminum, thereby allowing the structure to be simple, rigid, and lightweight.

The 1-pixel camera contains four primary mechanical pieces: the primary optical block, a pinhole carrier, a focusing shim, and a replaceable filter holder. The 1-pixel camera is attached to the Riegl optical head by a common mounting structure.

Primary optical block

The primary optical block is a single piece of aluminum that contains the three sensing tubes, and serves as the primary mounting structure of all components. Lenses for the three sensing channels are rigidly mounted inside the top of each sensing tube, and are stationary with respect to one another. Each channel in the primary optical block includes a number of baffles to trap light rays which are not traveling directly from

lens to pinhole, and to prevent the rays from bouncing around inside the tube and interfering with the reading. The interior of the sensing tube is painted flat black to help minimize stray reflections as well.

Pinhole carrier

Unfortunately, the short focal length design and limitations of machining and component tolerances do not allow simply mounting of the pinhole rigidly as part of the optical block with the expectation that all three channels are exactly parallel to the primary viewing rays. Therefore, we mount each of the pinholes in a separate movable piece of aluminum, which attaches to the end of the optical channels. Prior to tightening the pinhole carrier's mounting bolts, the lateral position of the pinhole can be moved within approximately 1mm, thereby allowing the primary viewing ray to be adjusted by 2.5º. Once tightened, the pinhole is rigidly mounted with respect to the lens, and the primary viewing ray is fixed in the desired direction. We have devised an alignment procedure to properly adjust each channel. This procedure will be discussed in Chapter 4 as part of calibration.

A photodiode is the final destination for all light collected by the optical channel, and is mounted on the pinhole carrier just behind the pinhole. This ensures that all light rays passing through the pinhole fall onto the sensitive area of the photodiode. The photodiode moves with the pinhole during adjustment. The position of the light spot hitting the photodiode moves up to 0.05 mm due to adjustment, so there is enough peripheral sensitive area to accommodate this variation.

Focusing shim

To achieve the proper focus for each channel, the distance from the lens to the pinhole also needs to be adjustable. In our sensor, this distance can be adjusted by changing the thickness of the metal focusing shim, which spaces the pinhole holder away from the primary optical block. We employ an assortment of shims with different thicknesses,

thus allowing the proper focus to be achieved. One disadvantage of this design is that the pinhole carrier must be removed to change focus, making focus adjustments tedious. However, once locked in place, the distance from lens to pinhole will not change. The focusing procedure will be discussed in Chapter 4 in detail as part of calibration.

Replaceable filter holder

The optical filters are held in front of the primary optical block by a replaceable filter holder. The holder rigidly positions the three filters in front of each channel, and is itself rigidly bolted to the front of the optical block. The holder allows the filter set to be removed and changed without affecting the reset of the sensor.

Combined mount

Figure 3-15. shows how the 1-pixel camera's optical block, cold mirror and Riegl optical head are mounted directly to one another. thereby ensuring that they remain aligned during flight. This is critical since our approach requires the laser rangefinder and 1-pixel camera channels to be precisely mechanically aligned. Additionally, in this configuration the entire assembly can be mounted as a single unit in our system.

3.3.10 Measurement electronics

Up to this point, our discussion has concentrated on optically collecting the desired light rays, and then funneling them into a photodiode. In order to make usable measurements, we must convert the photocurrents generated by the photodiodes into digital numbers and pass these to the computer system. To match the sample rate of the laser rangefinder, successive, simultaneous measurements must occur on the three sensing channels at a rate of at least 12,000 times a second. These measurements must have sufficient sensitivity, dynamic range, and resolution to span the entire range of expected photocurrents (approximately 1nA to 1uA). The measurement bandwidth

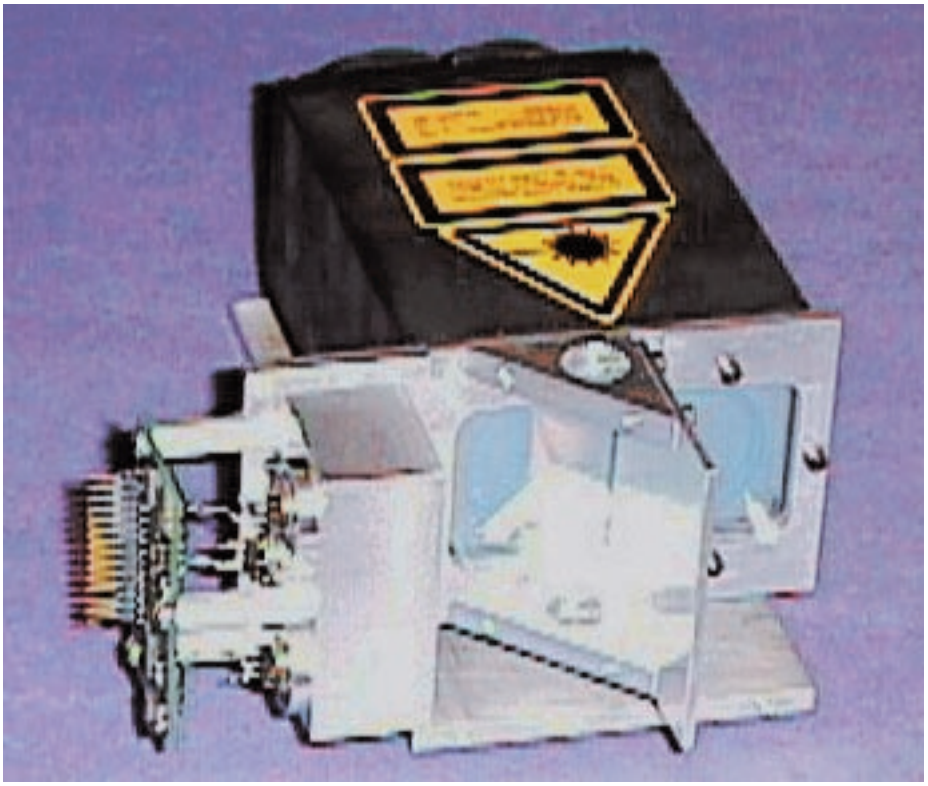

FIGURE 3-15. 1-pixel camera and cold mirror mounted to Riegl to ensure rigidity

must be sufficient to ensure that successive measurements are independent, and the noise level for the measurements must be kept small enough to avoid deterioration of the color measurement. A final point to consider is that these are precision DC current measurements, and any variation in measurement from channel to channel will adversely affect the measured color.

When a measurement is triggered, the electronics simultaneously integrate the photocurrent from each photodiode through a short (64us) shutter period. The integrated charge is converted to a digital number, proportional to the charge, by analog-to-digital converters. After conversion, the three raw measurement numbers (red, green, and blue channels) are transferred to the computing system as part of the measurement digital data packet. A block diagram of the measurement electronics in the 1-pixel camera is shown in Figure 3-16.

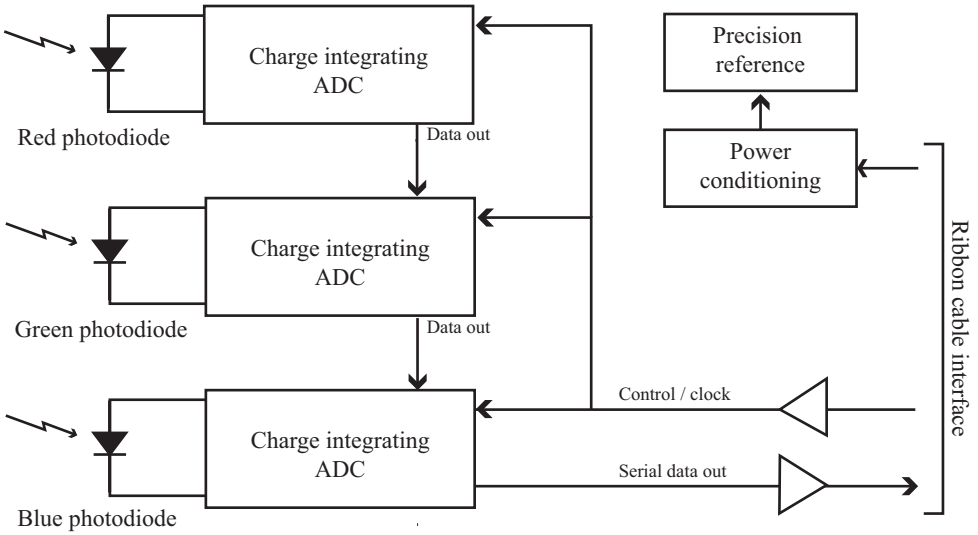

FIGURE 3-16. Block diagram of measurement electronics inside 1-pixel camera

Due to the small photocurrents being measured, it is imperative that low-noise design practices be utilized. To achieve this, we physically locate the sensitive measurement electronics inside the 1-pixel camera as shown in Figure 3-17, close to the photodiodes, and fully shielded from external noise sources. Because the interface from the 1-pixel camera to the main computing system is digital, it provides a lossless transmission.

Building sensitive, low-noise photodiode amplification and conversion circuitry is traditionally very challenging. Fortunately, much of the difficulty can be alleviated by using specialized integrated circuits that are designed to perform these functions. For the 1-pixel camera, we chose to use the DDC101 monolithic charge input ADC from Burr-Brown [26]. Due to its high level of integration the DDC101 greatly simplifies the measurement electronics.

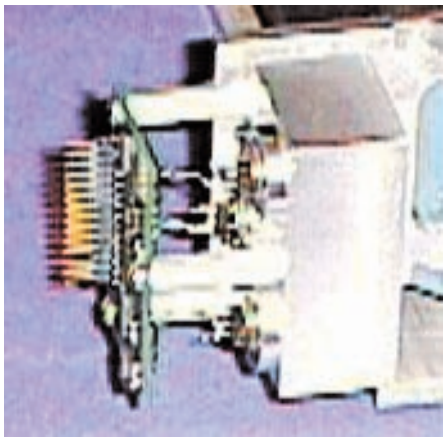

FIGURE 3-17. Sensitive measurement electronics located inside 1-pixel camera (outer shield removed for picture)

The DDC101 is designed to be directly interfaced to a photodiode, and provides all of the needed current integration, correlated double sampling, charge-to-digital conversion, filtering and serial digital transmission. When a measurement is requested, the current generated by the photodiode is integrated in a capacitor for 64us. This is the longest integration period the DDC101 will support while performing 12,000 samples per second (83us repetition period). Correlated double sampling circuitry removes constant biases within the DDC101, thereby reducing the DC errors. The amount of charge collected is linearly converted into a 20-bit digital value by an analog-to-digital converter.

With the 64us integration period, the DDC101 has a full scale input range of 0 to 7.8uA. With 20 bits of resolution, the smallest resolvable current is 7.4pA. Based on our previous modeling of the 1-pixel camera's behavior, we expect the maximum photocurrent measured from the terrain to be approximately 1uA. Therefore, the DDC101 has a wider dynamic range than we require. Typically, it is desirable to match the full scale input range with the expected input range to maximize the measurement resolution. However, due to the design of the DDC101, we are unable to make such an adjustment. As a result, the top 3 bits of the measurement will never be used, effectively limiting the resolution of the sensor to 17 bits.

We chose to operate the photodiodes in their photovoltaic mode. In photovoltaic mode, a 0 volt bias is maintained across the diode's leads by the DDC101's input circuitry. This mode is preferred for precision light measurement applications, since it provides significantly less noise than a reversed biased photodiode, as in the alternative photoconductive mode [27]. The typical disadvantage of operating in the photovoltaic mode is that the response time of the photodiode is slower. However, the photodiode response is still fast enough for the 1-pixel camera.

Three separate DDC101's are used to simultaneously sample the three channels. The serial data outputs from the three DDC101's are cascaded together to output a single 60-bit measurement at a clock rate of 2Mhz. Ribbon cable line drivers bring control and timing signals into the 1-pixel camera, and transmit the measurement data to the main computing system.

Interface electronics

A custom VME interface circuit board was developed to accept measurement data from the 1-pixel camera, laser rangefinder, and scanning mirror and to assemble them into a single data packet. Due to the tight timing requirements of this task, the processor can not perform that task.

The addition of this board to the onboard VME cage of our autonomous helicopter provides a direct interface to the terrain sensor for the main flight computers. The interface board shown in Figure 3-18 also provides hardware data buffering to reduce the real-time requirements on the processor, and hardware time synchronization for all onboard sensors as discussed in Section 3.5.

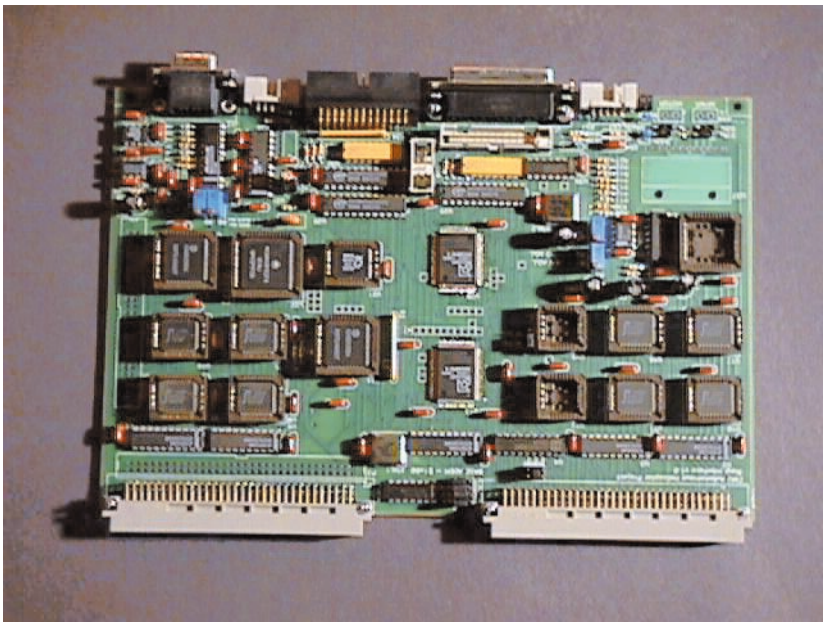

FIGURE 3-18. Custom VME interface board for terrain scanner

3.4 Scanning mechanism

We chose to implement a planar 1-dof scanner for the terrain sensor, as discussed in Chapter 2. This scanning configuration is capable of handling a wide range of mapping tasks, is simple to build, and weighs significantly less than more complex scanners.

A picture of the scanner is shown in Figure 3-19. The scanner consists of a mounting

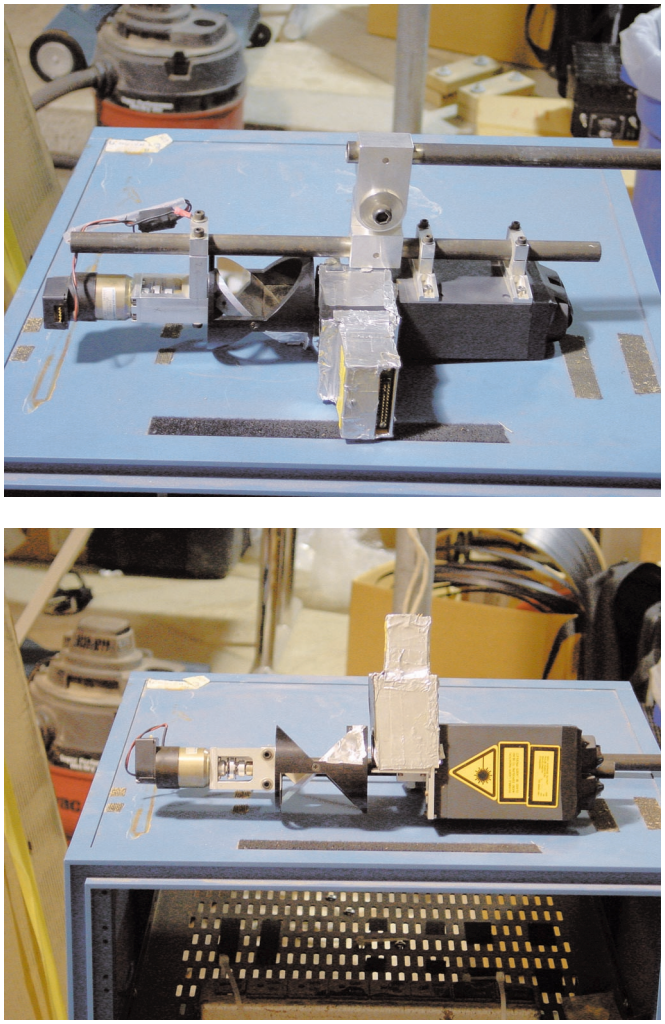

FIGURE 3-19. Scanning mechanism

structure which points the laser rangefinder / 1-pixel camera assembly toward a spinning mirror. The mirror scans the beam through a plane beneath the helicopter. This section concentrates on the spinning mirror assembly and on the overall scanner mount.

3.4.1 Mirror

The primary optical component of the scanning mechanism is its spinning mirror. The surface of the mirror is rigidly mounted at 45º to its rotational axis. This ensures that the optical path from the rangefinder / 1-pixel camera is reflected by 90º onto a plane normal to the rotation axis. During rotation, the optical path's direction is controlled by the angle of the mirror.

The spinning mirror, an Acuity, Inc. product [31], provides a 62mm clear aperture, has a front surface aluminum coated mirror mounted at 45º, and is dynamically balanced for vibration-free, high speed rotation. The mirror is neutral in color (meaning that it equally reflects light at all visible wavelengths), as opposed to the gold coated mirrors common in other laser scanning systems. Neutral color is significant in this design, since the 1-pixel camera is looking through the mirror, and a colored mirror would distort the color of the image.

As mentioned previously, the optical path from the rangefinder / 1-pixel camera requires a 60mm minimum diameter clear path through the scanner. To achieve this, the chosen mirror is elliptical in shape and has a minor diameter of 62mm, and a major diameter of 88mm (62mm / sin(45º)) to match our system's requirements.

The rotational axis of the mirror is attached to the scanner structure by a pair of preloaded bearings. These hold the mirror rigidly in place, while allowing it to spin freely.

3.4.2 Encoder

An optical encoder, which is mounted onto the rotational axis of the mirror to track its angle, is a 500 line quadrature encoder with an index pulse generated once per rotation. The 500 line encoder divides the full mirror rotation into 2000 discrete steps, each 0.18º wide. Since optical encoders measure only relative changes in angle, the

index pulse identifies one point in the rotation as a physical reference, thereby allowing the actual angle of the mirror to be tracked relative to this reference. The encoder is used for two purposes: to provide feedback for controlling the rotation of the mirror, and to trigger terrain measurements at precisely known mirror angles.

The scanner uses the encoder to initiate a measurement of the rangefinder and 1-pixel camera exactly 2000 times per rotation. This trigger occurs when the encoder transitions from one discrete step to the next, thereby ensuring that the mirror is at a precisely known angle. Notice that, since the encoder triggers the measurement at the transition from one angle step to another, the repeatability of the angle at which the measurement occurs is significantly better (0.01º) than the encoder resolution (0.18º). So, a measurement is triggered every 0.18º, but the direction of the beam is precisely known to 0.01º resolution. This is in contrast to many other scanners which fire the laser asynchronously at a fixed rate, and record the current angle of the mirror, based on the encoder feedback. In this case, the actual angle may be anywhere within the 0.18º step of the encoder, limiting the repeatability of such a scanner to worse than +/- 0.18º. In a second mode of operation, the scanner triggers a measurement every 4th encoder transition, resulting in 500 measurements per rotation. This mode allows the mirror speed to be increased.

3.4.3 Motor

A small DC brushed motor is attached to the mirror's rotation shaft via a flexible, zero-backlash rotational coupler. The motor continually spins the mirror at the desired rate during scanning, spining in one direction, as opposed to rotating back-and-forth, or to following other trajectories while scanning. This requires much less electrical energy, since the mirror's inertia sustains the rotation, and eliminates any accuracy differences between rotating in either direction.

To achieve a constant angular velocity, DSP-based motor control chip (HCTL-1100) drives the motor. The encoder provides feedback to the controller, thereby allowing an accurate closed loop velocity control. A constant, controllable mirror velocity is important for the system, since the measurements are triggered by the mirror angle. With too fast a velocity, the measurement requests may exceed the laser's limit of 12 Khz. With too slow a velocity, the measurements do not occur as fast as possible, therefore wasting some of the sensor's capability. Additionally, accurate motion control is important since the synchronization software assumes that the measurement rate is constant during the period between synchronization updates (20ms).

During normal scanning, the angular rate of the mirror is chosen to be either 6 Hz (360 rpm), when using a measurement rate of 2000 samples per rotation, or 24 Hz (1440 rpm) when taking 500 samples per rotation. Notice that both options perform approximately 12,000 samples per second. The choice of operating modes is based on the type of terrain mapping application.

3.4.4 Mount

The scanning mechanism is assembled as a single unit which rigidly holds the motor, encoder, and mirror unit in front of the laser rangefinder and the 1-pixel camera. The whole assembly is attached to the helicopter as one module. The mount is designed to maximize the amount of unobstructed view throughout a complete mirror rotation. This mount provides a full 270º of unobstructed mirror rotation, as shown in Figure 3-20a.

The next consideration is the scanner mounting location onboard the helicopter. Throughout this research, the scanner has been mounted in several ways. To support our primary goal of terrain modeling, the mount shown in Figure 3-21 has proven to be the most useful. Here, the scan plane is normal to the forward direction of flight, and is located just behind the landing skids to avoid being blocked by the skids. This allows a

(a)

(b)

FIGURE 3-20. Unobstructed scan window of scanner (a), and when mounted on helicopter (b)

230º unobstructed window where the helicopter body is not blocking the view, as shown in Figure 3-20b. Notice that this includes both the region under the helicopter, for scanning the terrain beneath, and also the region to the left of the helicopter, for scanning the terrain beside the helicopter.

Since this is a strapped down system, we rely upon the helicopter's attitude measurements to accurately track the attitude of the scanner. For this reason, the scanner must be mounted to the same physical structure as the helicopter's attitude

measuring sensors. IFigure 3-21 shows that the scanner and IMU are mounted near

FIGURE 3-21. Terrain scanner mount... scanning beneath the helicopter

each other on a single structure, thereby ensuring a rigid connection.

3.4.5 Weight

One of the principle driving factors in the design is to minimize the weight of the entire system. We have a limited payload on our helicopters, and a lower weight allows for a more versatile aerial platform. Table 1 outlines the weight of the primary components of the terrain sensor. With a total weight of under 10 lbs, the sensor is light enough to be used onboard CMU's autonomous helicopters.

Item	Weight
Laser rangefinder electronics box	5.25 lbs
Laser rangefinder optical head	1.25 lbs
1-pixel camera	0.75 lbs
Scanner	1.25 lbs
Total terrain sensor weight	9.50 lbs

TABLE 1. Weight of terrain sensor component

It should be noted that the 1-pixel camera comprises only a small fraction (< 10%) of the sensor's total weight, but adds a significant amount of additional terrain information, as compared with a range only terrain sensor.

3.4.6 Enclosure

The terrain sensor does not have a protective enclosure to isolate the scanner from the environment. An enclosure was not developed for two reasons. First, the enclosure would add a significant amount of weight to the terrain sensor (several pounds at least). Second, an enclosure was not required for developing, testing, and demonstrating the terrain modeling system. The individual components are sufficiently robust to handle the environmental conditions on the days we chose to fly (i.e., days with no significant precipitation). Fortunately, we have not experienced any failures due to the lack of enclosure.

However, an enclosure could be developed to completely enclose the scanning mechanism. The first primary issue is building a low-weight enclosure which will withstand the vibrations, strong downwash forces from the main rotor, as well as protect the sensor from the elements. Secondly, the enclosure must contain a clear window through which the laser rangefinder and 1-pixel camera can view the terrain. With the first pulse Riegl laser rangefinder and the scanning mechanism we chose to use, this is difficult, as the range finder may not see past the minute reflections off of the window's interior. One solution is to use a different type of laser rangefinder capable of selectively ignoring close reflections internally.

3.5 Synchronization system

In Chapter 2, we emphasized the importance of synchronizing all of the measurements on the helicopter. We developed a helicopter-wide synchronization system specifically for temporal alignment of measurements. This system incorporates a central hardware "clock," which determines the precise time all measurements within the helicopter system are captured. With this time-stamp, the system can temporally align all on-board sensors, thereby effectively synchronizing a "snapshot" of all measurements at an arbitrary instance in time.

The synchronization of the entire helicopter system is an important component of the terrain sensor. Lack of synchronization at any point in the system can limit the achievable terrain measurement accuracy. Therefore, we also present a method to synchronize the sensing elements of the helicopter state estimator with the terrain sensor. The system assumes that the state estimation system can properly maintain the timing relationships and can appropriately time-tag the resulting output.

While achieving system wide synchronization, we desire to minimize the impact on the various sensors, and the amount of required hardware elements. This requires development of a particular technique for synchronizing each sensor in the system.

The synchronization system generates a digital trigger signal for each of the sensors onboard the helicopter, keeps track of the precise time the trigger is asserted, and expects each sensor to constantly provide measurements with respect to the instance in time when the trigger is asserted.

Figure 3-22 outlines the trigger signal generation hardware. The Inertial Measurment Unit drives the system, by generating a digital pulse simultaneously with each of its measurements, precisely 400 times a second. Custom circuitry counts these pulses to keep track of the current "HELI time," and produces a trigger pulse for the Global

Positioning System, compass, terrain sensor, and video systems at a fixed multiple of the main trigger rate.

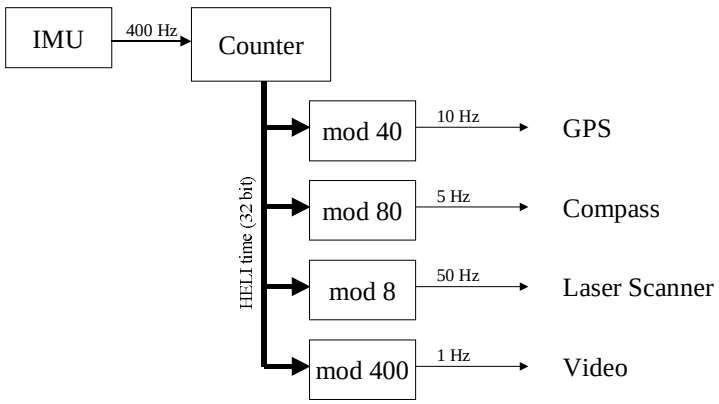

FIGURE 3-22. Trigger signal generation

The system starts following a synchronization reset event. This resets the HELI time to 0, and clears all of the dividers generating pulses at known multiples of HELI time. From this point on, the triggers are generated at precisely known times, based on the fixed divisors. The following sections discuss how each sensor uses the trigger pulse to generate measurements, and how the software accurately determines the HELI time at which each received measurement was captured.

In this triggering scheme, the hardware ensures that measurements occur at the proper times, and that the software can infer the time of each measurement from the constant timing relationships. More direct methods could explicitly measure the time of each measurement in hardware and send it to the computer, but this would require more hardware, processing, and complexity.

3.5.1 IMU synchronization

The IMU is the clock source for our synchronization system. It automatically triggers itself with a 400Hz clock, which synchronously triggers the rest of the on-board sensors. Following a synchronization reset, on-board software can easily track the HELI time of each IMU measurement by simply counting the number of incoming measurements.

3.5.2 GPS and compass synchronization

The GPS and compass sensors have trigger inputs. This means that they will produce a single measurement output at a precisely known time relative to an incoming trigger pulse. This allows direct hardware triggering of these sensors while the software keeps track of the HELI time of each measurement. Simply counting the number of measurements since the synchronization reset is sufficient to determine measurement timing. The GPS is triggered at a rate of 10 Hz. Therefore, the first three GPS measurements occur precisely at HELI times 40, 80 and 120. Similarly, the Compass is triggered at 5 Hz, so its first measurements occur at HELI times 80, 160, and 240.

A by-product of synchronizing with the GPS is that our arbitrary HELI time can be precisely related to the worldwide GPS system time. This allows the synchronization of the helicopter with many off-board devices, such as ground-based beacons, other autonomous helicopters, remote sensors, etc.

The accuracy of synchronization achieved with this method is directly related to the sensor timing accuracy. The GPS system provides remarkably accurate synchronization (< 1us uncertainty), since it is designed for precise synchronization tasks. The compass, on the other hand, is not as repeatable when triggered, so its synchronization accuracy is on the order of 100us.

3.5.3 Terrain sensor synchronization

The laser scanner measurements cannot be directly triggered by this system, due to its high sampling rate; therefore, a different approach is necessary for proper synchronization. As previously mentioned, an encoder on the scanning mirror triggers the laser rangefinder and 1-pixel camera to perform measurements at approximately 12 KHz. To determine the HELI time of each sample, the synchronization circuitry inserts a 1-bit flag into the rangefinder / 1-pixel camera's data stream to identify the terrain measurement which occurs immediately following each trigger pulse from the synchronization system. These marks occur at 50Hz.

Once in the computer, the terrain measurements are searched to locate tagged measurements. These measurements are assumed to have been collected at the time of the trigger pulse, so the first three flagged measurements occur at HELI times 8, 16 and 24. The HELI time for each un-flagged laser measurement is determined by linearly interpolating between adjacent flagged measurements, thereby providing precise times for each measurement.

This technique ensures that the HELI time of each flagged measurement is accurate to within one measurement time (1/12,000th of a second, or 84us), and that the unflagged measurements are close to this accuracy, since the angular velocity of the mirror is nearly constant over the 1/50th of a second period between triggers.

3.5.4 Video camera synchronization

The final sensor to be considered is an onboard NTSC video camera. The video camera acquires approximately 60 images per second by capturing light during a short shutter period, and outputs an analog video signal. We allow the camera to run asynchronously, while the synchronization system determines the precise HELI time that the shutter starts and ends for each image in the video.

Knowing the HELI time at which each image is taken is valuable. For example, this knowledge allows the helicopter's state to be estimated at the time of each image acquisition. It also allows vision based measurements to be used as measurements in the state estimator with accurate synchronization.

A custom Vertical Interlace TimeCode (VITC) generator assigns a unique number to each image coming from the camera. VITC is an established standard for assigning a unique time to each image of video, thereby establishing VIDEO time. This number is embedded into the National Television Standards Committee (NTSC) video signal on two lines in the vertical blanking interval. Since the timecode is a part of the video signal, it stays with the video, and can be digitized, transmitted to the ground, or recorded to video tape for later processing. At any time, the unique number for each image of the video can be determined by decoding the VITC lines in the vertical blanking interval.

While the video camera captures images, the synchronization system generates a trigger signal once per second for the camera. The trigger pulse is input to our custom timecode generator, which records the current image number as well as a time offset from the start of the current image. The current image starts at a specific point in the NTSC video stream, and the time offset is precisely measured with <1 us resolution. This information allows the precise HELI time at which the current image started to be determined, since the trigger pulses occur at HELI times 400, 800, 1200, etc. The video stream operates at a stable, known rate, so a single correspondence completely relates the video time with HELI time. We update the correspondence each time a trigger occurs (once per second) to prevent the timebases from slowly drifting apart.

To determine the time of image capture, we must determine the time at which the shutter opens and closes with respect to the start of the image. These times are dependent on the shutter duration selected, and are specific to the chosen camera.

Therefore, once the shutter times are determined for the deployed camera setup, the system can synchronize the images by determining the shutter start/stop HELI times.

The system can determine the HELI time at which the shutter starts and stops to better than 1us accuracy. However, since the image is captured over the entire shutter period (17ms to 100us), there is an inherent ambiguity in the precise time at which images are captured. Software can take advantage of the precise shutter times, but we do not address this aspect of the problem in this thesis.

3.6 Terrain measurement process

The data collected by the terrain sensors must be processed to build a terrain model. The dataflow and processing steps are illustrated in Figure 3-23. In this section, we discuss each of the involved processing steps.

3.6.1 Data streams

The terrain sensor generates a single measurement packet for each captured sample. This raw measurement contains the laser range, reflectance, mirror angle, RGB color appearance measurement, and the exact HELI time at which the sample occurred. These measurements occur up to 12,000 times per second.

The helicopter's state estimator generates a complete state estimate 100 times per second, based on various onboard state sensors. The state estimate contains the 3D position of the helicopter in the world, the attitude of the helicopter, the uncertainty estimates, and the exact HELI time of the estimate.

These two data streams must be combined to produce a consistent dataset for each terrain sample. This combination is achieved by aligning the two streams based on HELI time, and then resampling/interpolating the lower frequency state estimates to produce a match for each terrain measurement.

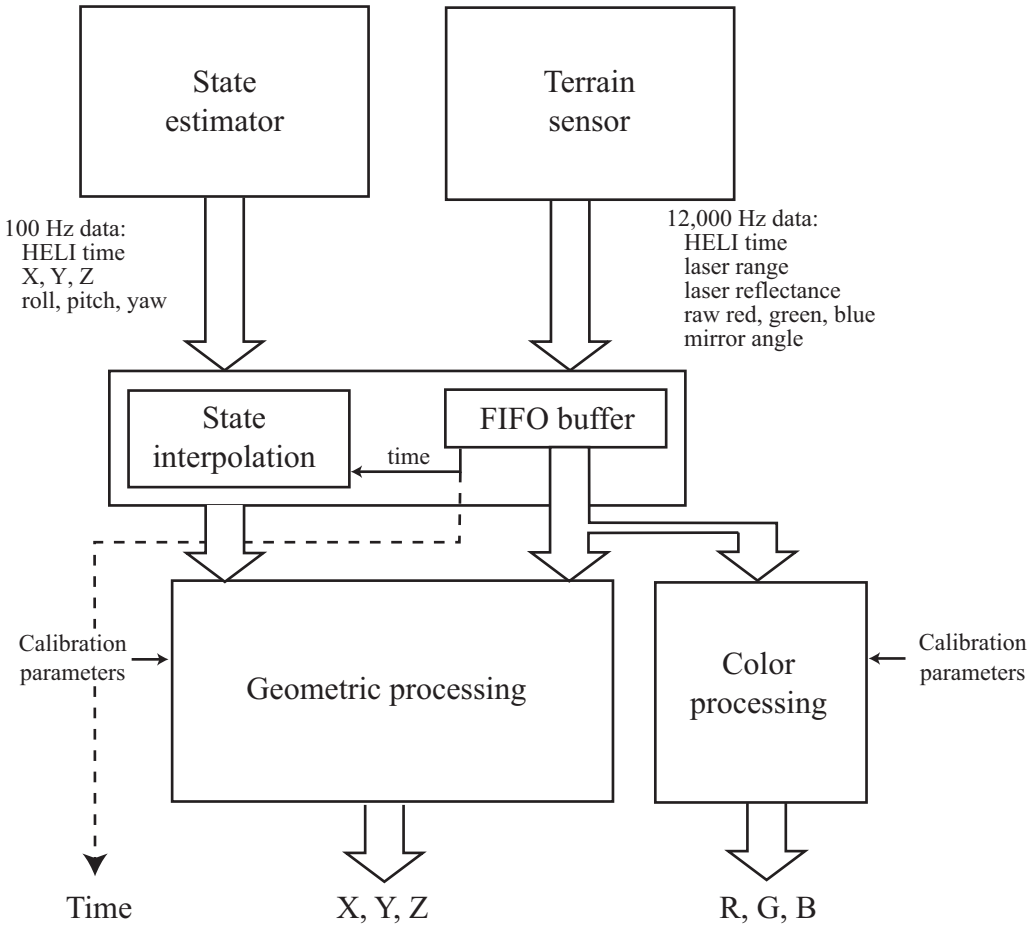

FIGURE 3-23. Dataflow and processing of terrain measurements

3.6.2 State Interpolation

The synchronization system ensures that we know the precise time at which measurements from each sensor are captured. However, these times do not necessarily coincide with the instant in time that a particular laser range measurement was made. We must interpolate between successive sensor readings to locate the sensor at the desired time instance.

3.6.3 Measurement rejection (filtering)

Not all terrain measurements are usable for terrain modeling. When a potentially inaccurate measurement is taken, it is preferable to eliminate the measurement.

Fortunately, the Riegl rangefinder is quite reliable, assuming that a sufficiently strong laser reflectance is received. A low reflectance intensity signals inaccuracies in the range measurement. To eliminate these points, a minimum threshold is applied to the laser intensity measurement. If the intensity does not exceed the threshold (typically 30 out of 255 maximum), the measurement is ignored.

Typically, the region-of-interest (ROI) is defined within the scan plane for the task at hand. Only measurements within this window are captured; all others are rejected to eliminate unnecessary processing.

Another case to eliminate includes ranges which are unusually short or long. Measurements with ranges less than a few meters or more than 100 meters are considered erroneous and are discarded. This is especially important when the portion of the scan plane in the ROI intersects the path of the main rotor blades. Some measurements inevitably hit the blades, and must be removed.

These simple rejection criteria provide effective filtering of the measurement data, and have been shown to produce relatively noise-free terrain models.

3.6.4 Measurement process

In order to build accurate 3-D maps of the terrain, the 3-D location of many target points must be determined. The target points are selected by successively firing a laser rangefinder toward the terrain in a uniform scan pattern. However, due to the motion of the helicopter and the unknown topography of the terrain, the target points are typically distributed in a non-uniform manner across the terrain surface. This requires

the measurement process to deal with each target separately, and imposes strict requirements on the synchronization of all measurements.

The measurement process of locating a single target is illustrated in Figure 3-24. In this section, we develop the equations and identify the parameters needed to do the terrain localization for a single measurement.

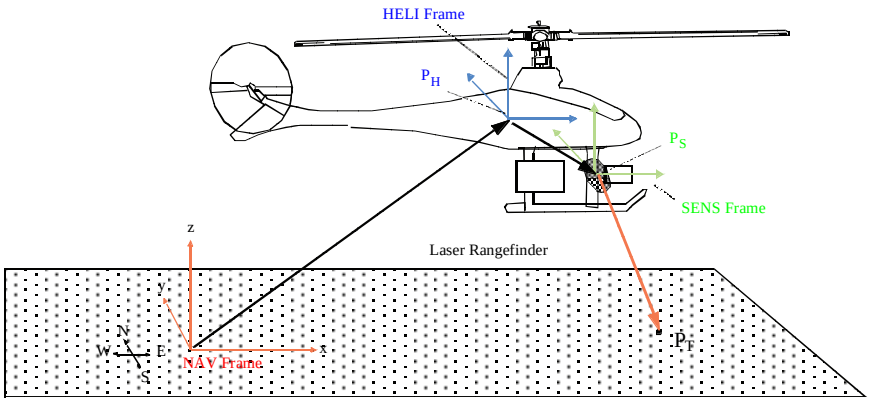

FIGURE 3-24. Measurement process

Coordinate systems

We deal with three primary coordinate systems.

- The navigation (NAV) frame is a stationary (earth-fixed) coordinate system that is the common reference for all state measurements. The location of the NAV frame's origin is arbitrary but once established, must remain stationary. The Z-axis points up (perpendicular to the local level), the X-axis points to the East, and the Y-axis points to the North. The coordinates in the NAV frame are commonly referred to as "world coordinates."
- The helicopter (HELI) frame is fixed to the helicopter fuselage, and is the common reference for all helicopter-based measurements. The HELI frame's origin is fixed at the nominal Center of Gravity (CG) position of the helicopter. The Z-axis points up (parallel to the main rotor shaft), the X-axis points forward (out the nose of the

helicopter), and the Y-axis points out of the left side of the helicopter.

- The sensor (SENS) frame is fixed to the terrain sensor. The SENS frame's origin is located along the spin axis of the scanner's mirror, at the point where it intersects the front (reflecting) surface of the mirror. This point is chosen since it can be considered to be the origin of the laser beam, regardless of the mirror's angle. The X-axis points along the spin axis of the mirror toward the motor; thus, it is always perpendicular to the reflected laser's beam. The Z-axis direction is defined to be perpendicular the X-axis and typically points "up." The Y-axis is therefore defined to be perpendicular to both the X- and Z-axes of the SENS frame. Formally, $Y_{SENS} = Z_{HELI} \times X_{SENS}$ and $Z_{SENS} = X_{SENS} \times Y_{SENS}$

Figure 3-25 shows the coordinate frames for two different scanner mounting configurations. The first configuration scans out a plane beneath the helicopter. The second scans a plane ahead of the helicopter.

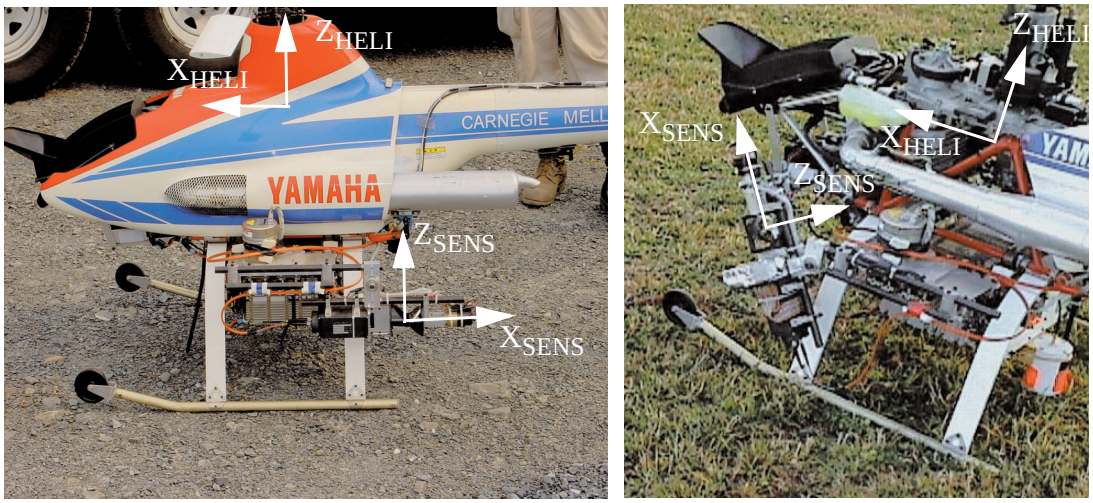

FIGURE 3-25. Coordinate frames for two different scanner mounts

Measurement variables

The terrain sensor and the helicopter's state estimator provide information on the target and the helicopter configuration for each measurement. These parameters are the real-time inputs to the measurement process.

- LASER_RANGE: The Riegl range finder measures the distance its laser beam must travel en route to hitting the target.
- MIRROR_ANGLE: An encoder on the scanning mirror's shaft determines the angle at which the laser beam exits the scanner. This angle is measured from an arbitrarily located index pulse within the encoder.
- ${}^{NAV}\boldsymbol{V}_{\text{heli}}$ and ${}^{NAV}_{HELI}\boldsymbol{R}$**:** The helicopter's state estimator provides the position and attitude of the helicopter at the instant each measurement is taken. Specifically, the estimator provides the full transformation from the NAV frame to the HELI frame. We consider this as consisting of two parts, a vector (${}^{NAV}\boldsymbol{V}_{\text{heli}}$), which is the location of the HELI frame origin in the NAV frame, and a 3x3 rotation matrix (${}^{NAV}_{HELI}\boldsymbol{R}$), which rotates vectors in the HELI frame into the NAV frame.

Calibration parameters

In addition to the measurement variables, a number of other parameters must be determined to complete the procedure. These are called the "calibration parameters," as they are fixed for all measurements, and must be determined only once. Specific methods for determining these parameters will be discussed in the following chapter on calibration.

- ${}^{HELI}\boldsymbol{V}_{\text{sensor}}$: The position of the SENS frame origin in the HELI frame. This vector is assumed to be fixed on the helicopter. Therefore, it must be re-determined each time the sensor mount is adjusted.

- ${}^{HELI}_{SENS}\boldsymbol{R}$: A 3x3 rotation matrix that rotates vectors in the SENS frame into the HELI frame. This vector is also assumed to be fixed, and must be re-determined each time the sensor mount is adjusted.
- LASER_RANGE_OFFSET: The distance the laser beam travels from the rangefinder to reach the spinning mirror. This offset, when subtracted from the LASER_RANGE measurement, provides the proper distance from the SENS frame origin to the target. This offset is assumed to be fixed, and must be re-determined only when the terrain sensor itself is adjusted.
- MIRROR_ANGLE_OFFSET: The angle from the Z-axis of the SENS frame to the encoder's index pulse. The sum of this offset with the MIRROR_ANGLE measurement provides the proper angle to the target in the SENS X-Z plane. This angle is assumed to be fixed, and must be re-determined each time the sensor mount is adjusted.

Terrain location calculations

Now that we have identified all of the measurement inputs and calibration parameters, we develop the measurement equation. The entire measurement process can be readily thought of as a kinematic chain which transforms the location of the target point from the SENS frame, to the HELI frame and finally to the NAV frame.

To compute the target location in the NAV frame, we first determine the location of the target in the SENS frame.

RANGE = LASER_RANGE - LASER_RANGE_OFFSET

ANGLE = MIRROR_ANGLE - MIRROR_ANGLE_OFFSET

$$ {}^{SENS}\boldsymbol{V}_{\text{target}} = \begin{bmatrix} 0 \\ -\text{RANGE} \cdot \sin(\text{ANGLE}) \\ \text{RANGE} \cdot \cos(\text{ANGLE}) \end{bmatrix} $$

Next, we transform this target vector into the HELI frame by applying the known rotation and translation determined during calibration.

$$^{HELI}\boldsymbol{V}_{\text{target}} = {}^{HELI}\boldsymbol{V}_{\text{sensor}} + {}^{HELI}_{SENS}\boldsymbol{R} \cdot {}^{SENS}\boldsymbol{V}_{\text{target}}$$

Finally, we transform the target vector into the NAV frame using the helicopter's state at the time of the laser measurement.

$$^{NAV}\boldsymbol{V}_{\text{target}} = {}^{NAV}\boldsymbol{V}_{\text{heli}} + {}^{NAV}_{HELI}\boldsymbol{R} \cdot {}^{HELI}\boldsymbol{V}_{\text{target}}$$

This calculation is performed for each terrain measurement, and yields the 3D position in the world of the terrain point. Since all of the measurements are carefully aligned in a single instance, the index of time is not shown explicitly. The above equations can also be rewritten in a single equation form. The following is the terrain measurement equation:

$$^{NAV}\boldsymbol{V}_{\text{target}} = {}^{NAV}\boldsymbol{V}_{\text{heli}} + {}^{NAV}_{HELI}\boldsymbol{R} \cdot \left({}^{HELI}\boldsymbol{V}_{\text{sensor}} + {}^{HELI}_{SENS}\boldsymbol{R} \cdot \left(\begin{bmatrix} 0 \\ -\text{RANGE} \cdot \sin(\text{ANGLE}) \\ \text{RANGE} \cdot \cos(\text{ANGLE}) \end{bmatrix} \right) \right)$$

3.6.5 Color processing

Separate from the 3D processing, the raw measurements of terrain appearance from the 1-pixel camera must be processed to generate an appropriate color measurement for the terrain point. The ultimate color is defined by ratios of red, green, and blue between 0 and 1. As we discussed before, it is desirable for the appearance results to match the colors from digitizing a standard 3CCD video camera. To achieve this, we must perform some operations to scale, transform, and offset the raw measurements.

The raw measurement from the 1-pixel camera is composed of three 20-bit values from the Analog to Digital Converter's (ADC's). Physically, these values represent the total optical energy collected by the optical channels during the shutter period. The

values are linearly proportional to the collected energy, as weighted by the spectral sensitivity of the optical channel. Therefore, the values can be thought of as measurements in the 1-pixel camera colorspace. The three basis vectors for the 1-pixel camera colorspace are determined by the spectral sensitivity of the 1-pixel camera (discussed in detail in Appendix A) which is shown in Figure 3-26.

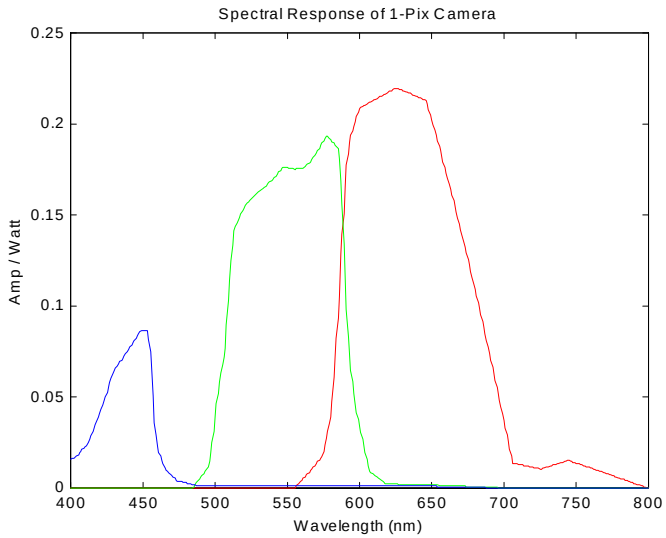

FIGURE 3-26. Basis vectors for 1-pixel camera colorspace

In order to convert these measurements to the desired colorspace, we must perform the transformation, as shown in Figure 3-27.

First, we subtract an offset from each raw measurement to compensate for any constant biases in the measurement electronics, and a fixed offset introduced by the ADC's.

Next, we multiply the three measurements by a 3x3 transformation matrix. This transformation converts the readings from the 1-pixel camera colorspace to the desired output colorspace. Generally, once a measurement is taken in one colorspace, it is not possible to exactly convert it to another colorspace. However, a best fit transformation

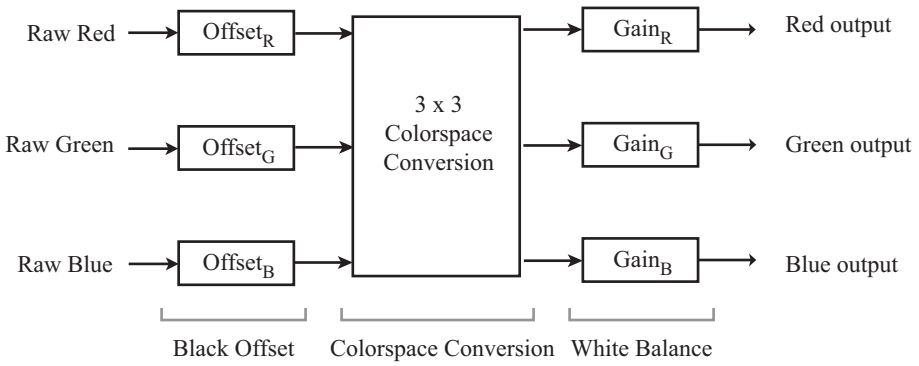

FIGURE 3-27. Color processing

(least squares) can be determined between two colorspaces with known spectral sensitivities, as described in [43].

Often, an identity matrix is used, resulting in the output colorspace being the same as the 1-pixel camera registers. At other times, if the output of the 1-pixel camera is to be compared with another color sensor, a non-identity transformation is employed. For work described in this thesis, we are not attempting to match a specific camera, so we typically use the 1-pixel camera's colorspace (identity transformation).

Each channel is scaled to adjust the white balance. White balancing ensures that the gains on the channels are matched, and that objects will appear to the human eye in the proper color. Typically, the white balance changes as the spectral content of the illumination (the sun) changes.

In our system, the parameters for the offset and white balance gain are determined experimentally during calibration (the procedure will be discussed in Chapter 4). The transformation matrix, on the other hand, is calculated based on our model of the 1-pixel camera's spectral sensitivity and the desired colorspace basis vectors.

Once this processing has been performed for each color measurement, the resulting Red, Green, Blue (RGB) values are associated with the computed 3D location in the terrain model.

Color image generation

It is worth noting that RGB values resulting from our camera measurements will not appear similar to an image taken by a digital or video camera, due to gamma correction [32]. Most commercial imaging devices intentionally apply some amount of gamma correction to the image prior to digital storage. This correction helps to improve a Cathode-Ray Tube (CRT) monitor's ability to display the image. Therefore, to properly display colors from the 1-pixel camera, we must apply our own gamma correction.

Gamma correction simply requires raising each Red, Green, Blue value to the gamma (γ) power. A typical value for gamma is 0.45. Image quality may be improved by slightly varying gamma to better match actual viewing conditions.

3.7 CMU autonomous helicopter overview

The terrain modeling system we describe in this thesis was built on-board one of the CMU autonomous helicopters. This small, autonomous helicopter has been under development at Carnegie Mellon University since 1991 [45]. In this section we briefly overview the particular autonomous helicopter used for work described in this thesis.

3.7.1 Yamaha R-50 helicopter

The helicopter is a modified Yamaha R-50 with on-board sensing, computing, and flight controls developed by the Autonomous Helicopter Project at CMU. The Yamaha R-50 is a commercially available industrial helicopter originally designed for remotely

operated crop-dusting. The general physical characteristics of the R-50 are provided in Table 2 and its dimensions are shown in Figure 3-28.

Payload	20 Kg
Takeoff weight	67 Kg
Engine	2 cycle, 98 cc, 12 HP, water cooled
Overall length	3.5 meters
Main rotor diameter	3.1 meters
Body length	2.7 meters
Rotor speed	870 rpm
Flight duration	1 hour
Top speed	85 mph

TABLE 2. Yamaha R-50 helicopter specifications

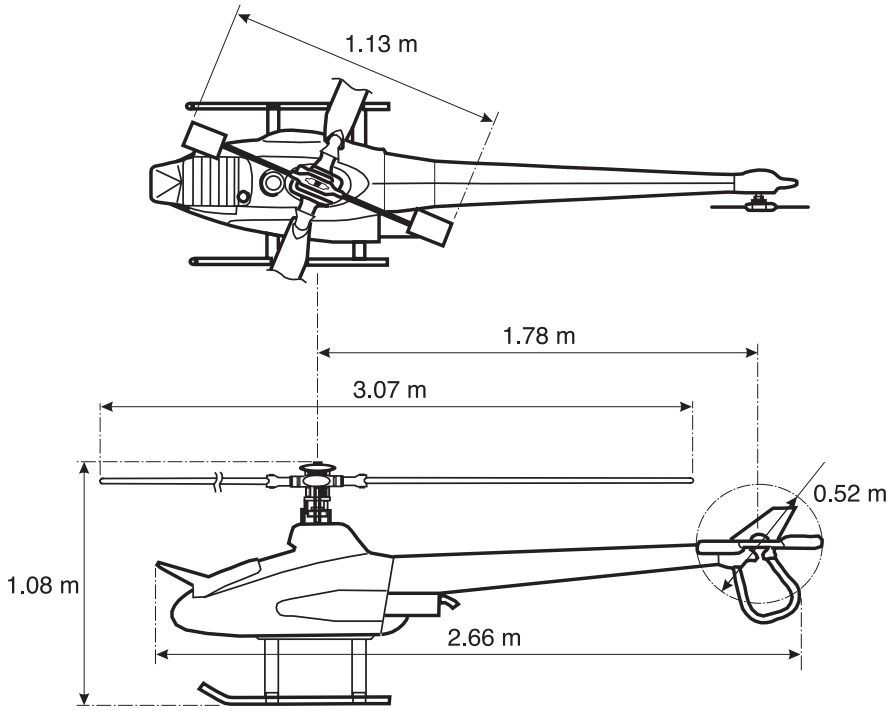

FIGURE 3-28. Dimensions of the Yamaha R-50 helicopter

3.7.2 State estimator

An onboard state estimation system fuses data from a differential GPS receiver, an inertial measurement unit (IMU), and a magnetic compass into a coherent measure of helicopter state. Specifications of these navigation sensors are listed in Table 3.

IMU	**3-axis raw angular rates ω_x, ω_y, ω_z and accelerations a_x, a_y, a_z** Litton LN-200: 3 silicon accelerometers, 3 fiber optic gyroscopes, < 1ms latency, 400 Hz data rate, 0.0027 deg angular resolution, 2mg acc resolution, SDLC interface
Differential GPS	**Geodetic latitude, longitude, altitude and UTC time** NovAtel Millennium RT2: 2 cm accuracy (<20 miles from differential base station), 10 Hz update rate
Compass	**Magnetic heading** KVH C100: digital flux-gate compass, 0.5 deg resolution, self calibrating, 5 Hz update rate

TABLE 3. Onboard navigation sensors

Measurements from these sensors are synchronized in hardware (as described in Section 3.5). A 12th order Extended Kalman filter combines the measurements to estimate the vehicle's state at 100 times per second. The state consists of the helicopter position, velocity, acceleration, altitude and attitude rates. The covariance matrix maintained in the filter provides continually updated uncertainty estimates of the state. Figure 3-29 shows an example of the accelerometer measurements from the INS extrapolating and smoothing the position estimates over the performance of GPS alone.

3.7.3 Autonomous flight control

The helicopter can autonomously take off, follow a desired smooth flight path, accurately hover at a point, and land. The on-board autonomous flight control uses the state estimates to update the helicopter control surfaces. The controller uses two nested PD control loops, the inner loop being the attitude loop that controls the velocity and the outer loop being the position loop [33].

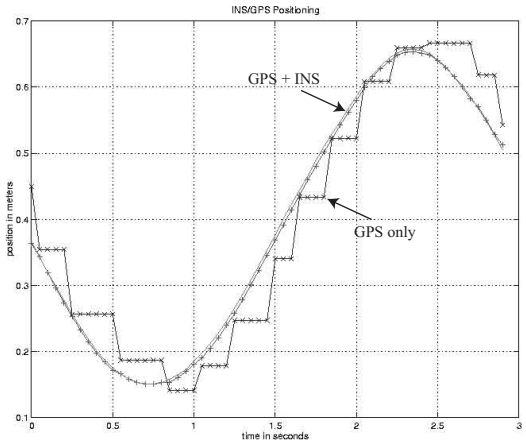

FIGURE 3-29. GPS + IMU performance

3.7.4 Onboard processing

A VME-based Motorola 68060 processor running VxWorks performs all onboard processing tasks, including the state estimator, flight controller, communications, and terrain data collection. A custom VME interface board provides connections to the onboard sensors, helicopter control actuators, and system status.

A wireless ethernet connection provides the primary telemetry link to the helicopter. Other data links include a radio modem for GPS differential corrections, and an analog video downlink. Table 4 provides the specifications of the links.

Wireless ethernet	**Primary telemetry link** Breezecom SA-10: 3 Mbps datarate, 2.4 GHz RF frequency
Wireless modem	**GPS differential corrections from base station** Pacific Crest RFM96: 2 watts output, 9600 baud datarate
Wireless video	**Video downlink** Dell-Star Technologies: 0.5m watt output, 2.4 GHz RF frequency

TABLE 4. Data links to helicopter

3.7.5 Terrain sensor support

The terrain sensor described in this thesis is rigidly mounted to the electronics chassis of the helicopter. The addition of the terrain sensor required some of the onboard equipment to be repositioned on the helicopter to allow a proper mount for the scanner, and to evenly distribute the mass on the helicopter.

A custom VME control board, shown in Figure 3-30,.was developed to interface the

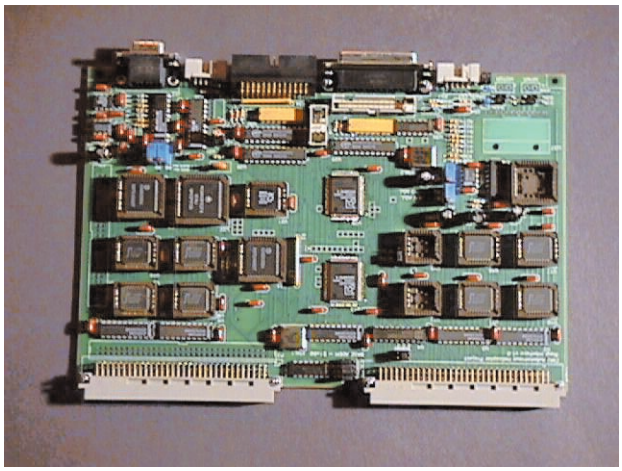

FIGURE 3-30. Terrain sensor interface board

terrain sensor with the onboard flight computer. This board manages the rangefinder, scanning mechanism and 1-pixel camera, and feeds the measurement data directly to to flight computer. Since the flight computer is busy with other tasks, terrain measurements are immediately relayed to a PC on the ground via the wireless ethernet. The ground PC logs the data to disk, and also performs real-time terrain modeling. This split processing approach was required since the onboard flight computer is not sufficiently powerful. Current processors, however, are sufficiently powerful enough to perform the tasks of the ground PC.

Chapter 4

Terrain modeling system: Calibration and verification

This chapter describes both the calibration techniques required for accurate terrain modeling, and the verification steps which demonstrate proper operation of different components.

In order to accurately measure target points, we needed to develop a set of calibration procedures for determining the unknown parameters in the measurement equations. Some of these quantities can be directly measured in the lab, while others require special techniques to achieve sufficiently accurate values.

The calibration and verification tasks are closely related and therefore are best presented together in this chapter.

4.1 Laser rangefinder calibration and verification

The Riegl laser rangefinder measures the distance to a target by transmitting a short laser pulse and then observing its time-of-flight as it reflects off the target. Due to various lengths of fiber optic cabling connecting the laser with its optical head and non-ideal characteristics of the rangefinder's electronics, the raw range data provided

by the rangefinder must be adjusted to provide an accurate distance measurement from the front of the lens to a target.

Prior to mounting the rangefinder in the scanning mechanism, the actual mapping from raw measurement to range is determined in the lab. We collect a series of range measurements to targets placed at a known distance from the face of the rangefinder. We verify that the range vs. raw reading mapping is linear, and can be described by a zero offset and scale factor. From the collected data, a best-fit line is constructed, and these parameters are extracted. In this case, the actual range (in meters) is:

$$\text{LASER\_RANGE} = 0.0040 \cdot (\text{data} - 3064)$$

The Riegl rangefinder has a measurement uncertainty of approximately 2 cm (std); therefore, we must measure the ground-truth range better than this level. Fortunately, 1 cm accuracy is reasonable to achieve with a quality measuring tape.

4.2 1-pixel camera calibration and verification

The 1-pixel camera requires a number of adjustments before it can operate properly. In this section, we present methods for adjusting the sensor's focus, line-of-site, and color calibration parameters, as well as verification of several important capabilities.

4.2.1 1-pixel camera bandwidth verification

The 1-pixel camera is required to perform up to 12,000 measurements per second. Because we expect each of these measurements to be independent, the photodiode and measurement electronics must have a sufficiently high bandwidth to pass the high frequency signals.

To verify that the bandwidth of the 1-pixel camera is sufficient, a sinusoidally modulated light source is sampled at 12 KHz by the 1-pixel camera. The modulation frequency is increased until the magnitude of the output sinusoid is reduced to 1/2 the magnitude it generates at low frequencies (10 hz).

This frequency is the -3db bandwidth of the sensor, and indicates the transient time the sensor needs to respond to a sudden change in input [25]. From the -3db bandwidth, the rise (or fall) time can be estimated as:

$$\tau_r = \frac{0.35}{f_{-3db}}.$$

Independent measurements require the rise time to be equal to the time between measurements, 1/12,000th of a second. This corresponds to a minimum bandwidth of 4.2 Khz.

We experimentally verify that the sensor exceeds the required 4.7 Khz bandwidth. In fact, the bandwidth of the detector/electronics components is actually greater than we can measure with this method, because of the limited sampling rate of 12 Khz. This ensures that the sensor can achieve independent readings from the 1-pixel camera.

4.2.2 1-pixel camera focus

We must clearly define focusing for the 1-pixel camera. Obviously, we have only one measurement (for each channel), and issues such as image blurring may seem irrelevant. However, if the sensor is not in focus, the sensing spot is blurred and light is captured over a larger area, thereby reducing the resolution of the sensor. The FOV of the sensor determines the size of the sensing spot at the plane of focus. The plane of focus is located at the theoretically ideal distance, where the target is in perfect focus. Everywhere else, there is some amount of blurring. In other words, targets at other distances have an additional circle of confusion, due to defocusing, combined with the FOV to provide a larger, less concentrated sensing spot. The amount of acceptable blurring determines the total range we consider to be in focus, the depth of field.

Effective focusing of the 1-pixel camera requires that we define the acceptable in-focus range, and adjust the sensor accordingly. In this section, we present our approach to focusing the 1-pixel camera.

Desired plane of focus

All targets at the plane of focus are perfectly in focus. Perfect focus means that a one-to-one relationship exists between points on the plane-of-focus and the sensor’s image plane as illustrated in Figure 4-1. The focus adjustment moves the plane-of-focus by varying the distance from the lens to the image plane.

For objects either closer or farther than the plane-of-focus, the one-to-one mapping fails as light rays from multiple points on the object converge to a single point on the imaging plane. Allowing an acceptable level of blur extends the range of in-focus distances beyond the plane of focus. To maximize performance, the focus must therefore be set appropriately to provide best performance over the full range of interest.

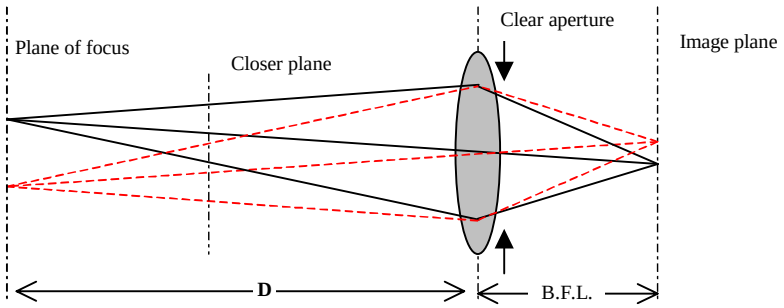

FIGURE 4-1. A simple imaging system

Every imaging system is tolerant to a certain amount of blur before noticeable performance degradation occurs. Typically, if the blur is less than 1 pixel (for a CCD camera) or less than the grain size (for film) the picture is still considered to be “in focus”.

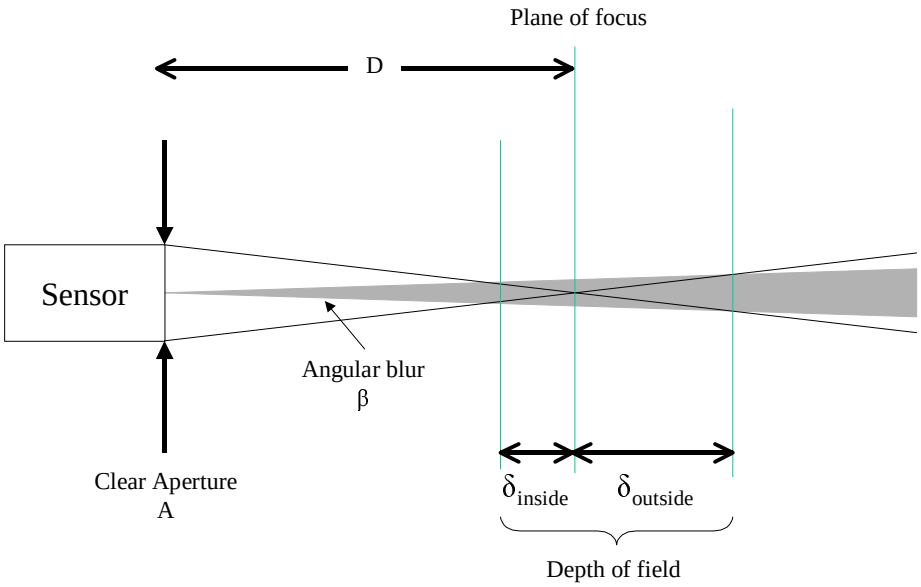

FIGURE 4-2. Depth of field geometry

The depth-of-field describes the range of distances on either side of the plane-of-focus that produce an acceptable blur level, as illustrated in Figure 4-2. The sensor is focused for distance, D, and has a clear aperture, A. The distances inside and outside the plane of focus with blur less than β (in radians) are:

$$\delta_{inside} = \frac{D^2\beta}{A + D\beta} \text{ and } \delta_{outside} = \frac{D^2\beta}{A - D\beta}$$

We want the 1-pixel camera to see distant objects, so one approach would push the plane of focus toward infinity. In this situation, parallel incident light rays are imaged to a single point on the image plane. The sensor is ideally focused for distant objects and blur increases for objects closer to the sensor. While functional, this approach unnecessarily reduces the usable depth-of-field, thereby limiting the sensor's ability to view close objects.

Alternatively, we can bring the plane-of-focus in closer to the sensor while allowing the depth-of-field to extend out to infinity. The maximum depth-of-field obtainable with infinity still in focus occurs when $\delta_{outside} = \infty$. This is achieved when the plane-of-focus is located at a specific distance called the hyperfocal distance [35]. The hyperfocal distance is calculated as:

$$D_h = \frac{A}{\beta}$$

where β is the maximum acceptable blur angle for the 1-pixel camera. Larger values of β increase the depth-of-field, and allow closer objects to be in focus. However, this also increases the spot size of the sensor by blurring the FOV. We choose to use $\beta = 0.2^o$ as a compromise. This limits the blurring of the FOV to 20%, and allows all objects greater than 2.6 meters away to be in focus, an acceptable distance from the helicopter. To achieve this result, we need to adjust the focus of sensor to place the plane-of-focus 5.15 meters away.

Focus adjustment

The focusing of each channel of the 1-pixel camera can be adjusted by changing the distance from the lens to the pinhole. As mentioned in Chapter 3, the design does not provide an easy method for changing this distance. Therefore, the focus procedure is tedious and requires inserting various metal shims to find the proper spacer thickness for each channel. This is a one-time operation for each channel.

Our focusing approach is based on measuring the sensing spot of the sensor when a particular focus shim is installed. Various focus shim thicknesses are tested to find the focus setting which achieves the most compact sensing spot. The sharpness of the sensing spot edge is a good indicator of focus. This adjustment is performed with the sensing spot measured at the hyperfocal distance (5.15 meters) to locate the plane-of-focus at the desired distance from the sensor. We also focus each of the three sensing channels independently, to compensate for potential differences among the channels, such as chromatic aberrations or manufacturing variations.

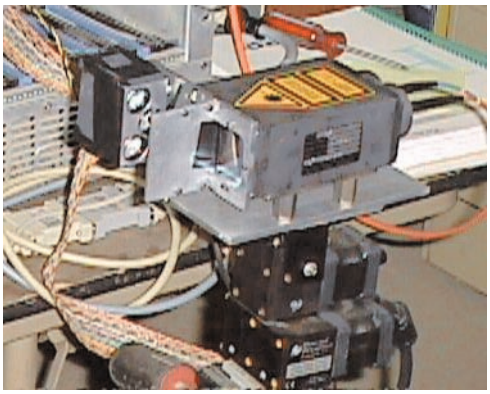

FIGURE 4-3. Verify sensing spot with pan-tilt unit (a) and point light source (b)

The sensing spot measurement is performed with a bright, single point light source located at the desired plane-of-focus in an otherwise dark environment. As shown in

Figure 4-3, the 1-pixel camera is mounted on a mechanical pan-tilt head that allows the sensor to be positioned precisely at various angles with respect to the incoming light. The sensor is slowly scanned in a raster format to cover a full range of pan-tilt values, with the 1-pixel camera taking a measurement at each step. This produces an image that can be interpreted as the angular sensitivity distribution of the sensor for each shim thickness.

Figure 4-4 illustrates the effect of the focus adjustment on the sensitivity distribution of the sensor. Figure 4-4(a) shows a properly focused sensor. The diameter of the spot is primarily determined by the sensor's FOV, and the sharp edges show its well-focused characteristic. Figures 4-4(b) and 4-4(c) are not as well focused. The gently sloping sides are characteristic of this unfocused condition.

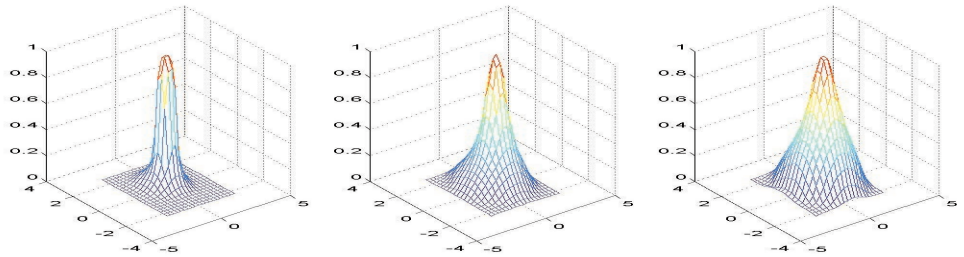

(a) (b) (c)

FIGURE 4-4. Effect of focus on sensing spot

The focusing approach requires systematic angular sensitivity measurements until all three channels exhibit well-focused spots. The sensor is then assembled with the proper sized focusing spacers permanently installed.

The focusing procedure also provides an opportunity to experimentally verify the sensor's FOV. In Chapter 3, we predicted FOV_{sensor}= 1.02º from the sensor design. Looking at the angular sensitivity after performing the focusing procedure, the actual ½ power spot diameter was found to be 1.1º.

4.2.3 Align 1-pixel camera channels with laser

The second sensor adjustment is to precisely align the 3 color sensing channels with the laser rangefinder to ensure they all image the same points on the terrain.

The central viewing ray of each sensing channel must pass through both the optical center of the lens and the center of the pinhole. In our sensor design, we adjust the viewing direction by translating the pinhole carrier laterally in two directions normal to the viewing ray. This adjustment is sufficient to establish the desired viewing direction. Once the pinhole carriers are properly positioned, three screws lock each of them in place to maintain the alignment.

Our alignment approach consists of two steps. First, an interactive alignment procedure guides the viewing rays into coincidence. Second, a verification procedure measures the actual misalignment of the channels to determine if the result is within the required tolerance. The procedures are repeated as needed to achieve proper alignment.

Interactive alignment

The interactive alignment procedure uses a remote CCD camera to “look back” at the sensor to visualize when the channels are aligned. This visualization is needed when moving the pinhole carrier to determine the proper alignment.

The interactive procedure has proved to be effective. A CCD camera is used to serve as a virtual terrain patch. The procedure attempts to align the laser and the 1-pixel camera so that they are observed to be in alignment from the CCD camera’s point of view. The procedure uses the 1-pixel camera as a light projector and a remote video camera as an observation point. The CCD camera is located approximately 50 meters away from the sensor and is zoomed in on the front of the sensor. At this distance, the lens of the

CCD camera is very small, providing a single point target, which provides good sensitivity. During the procedure, the camera's image is displayed on a nearby monitor.

First, we aim the laser rangefinder to hit the remote camera's lens. Then, we turn on the rangefinder and move the entire sensor assembly until the laser beam is centered on the remote camera's lens. The camera image provides effective feedback during the operation. The rangefinder's aperture glows brightly when aimed directly at the camera[1], as shown in Figure 4-5(b). With the laser aimed at the camera, the entire sensor is locked into place, thereby establishing the camera lens as our reference target.

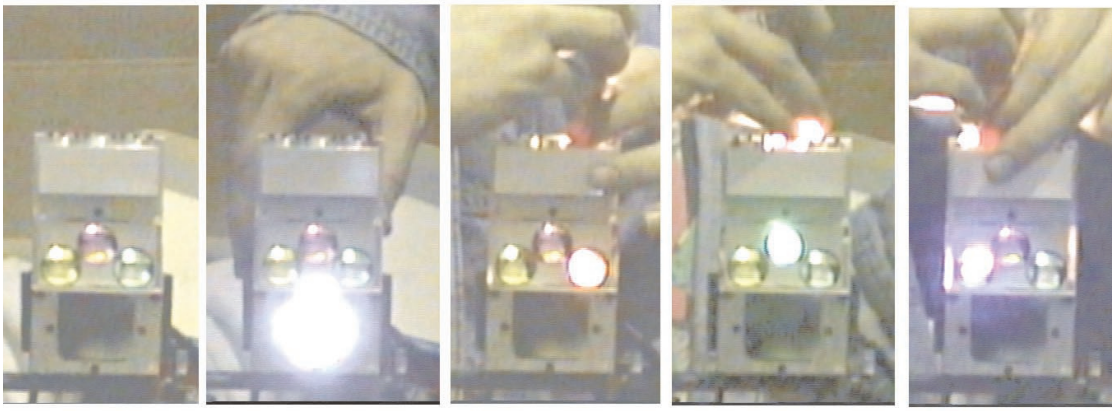

(a) (b) (c) (d) (e)

FIGURE 4-5. Remote camera view during alignment

Next, we adjust the viewing ray of each of the 1-pixel camera channels to also be centered on the camera lens. We achieve this by projecting light from the sensor and adjusting the pinhole positions to center the projection on the remote camera's lens. To project light through the 1-pixel camera, the photodiodes are removed from the assembly and a bright light is shined backwards through the pinhole (toward the CCD camera). As shown in Figure 4-5(c)-(e), the projected light makes the sensing tube

1. The camera does not have an IR blocking filter.

glow brightly in the camera's image when it is aimed directly at the camera lens, similar to the laser.

This technique relies on the human eye to center the spot for each channel. This task is fairly difficult for several reasons. The adjusting mechanism is sensitive; thus, small changes in alignment are very difficult. Also, the appearance of the spot varies subtilely as the alignment changes near the desired center. This variation makes discerning the precisely aligned center location quite challenging. In fact, significant observation and practice of the technique is necessary before one can reasonably align the sensor. This is the primary reason why we must verify the alignment by using an alternative method.

Alignment verification

We must verify the alignment to ensure that the interactive procedure successfully achieved the desired tolerance. We consider a 0.2º misalignment (20% of FOV) acceptable, since the majority of the sensing spots of the two channels overlap.

We verify the alignment by measuring the angular sensitivity of each channel and then measuring the angle between the centers of their respective sensitivity peaks. We also verify the angle between the center of the laser beam and the sensitivity peaks of the color channels.

The alignment verification procedure is very similar to the focus measurement, with two exceptions: the point light source is moved farther away (50 meters), and the target includes a retroreflective pattern surrounding the light source. The sensor is mounted on the pan-tilt unit and scanned over the area containing the target while simultaneously measuring range, reflectance, and red, green, blue color channels.

As shown in Figure 4-6, the center of the response for each channel is found in the resulting angular sensitivity images. These points correspond to the center of the target

light, as viewed by each channel. In a perfectly aligned sensor, we expect the points to be coincident for all channels.

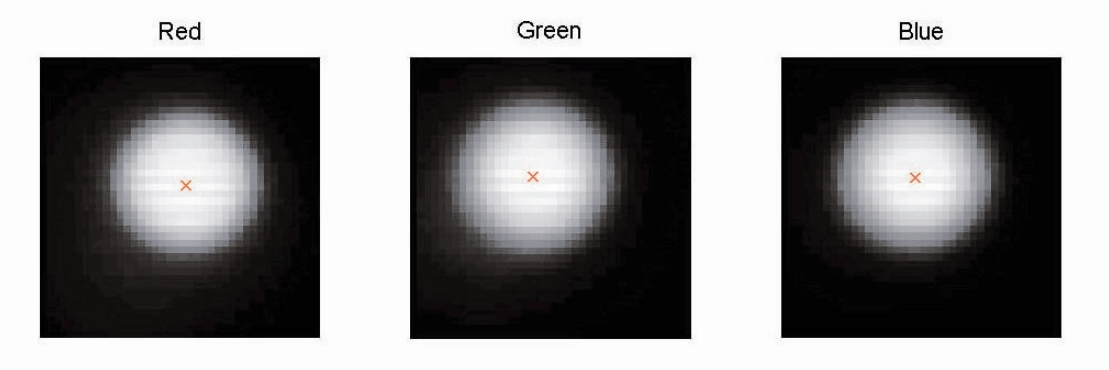

FIGURE 4-6. Angular sensitivity of three aligned sensing channels

The target center is next found in the laser rangefinder's image. Four pieces of retroreflective tape are radially centered on the light, as shown in Figure 4-7(a). This pattern is clearly visible in the reflectance image, and the center of this pattern can be directly compared with the points found for the color channels. Figure 4-7(b) shows the reflectance image and the target center locations from the red, green and blue sensors. In this manner, the laser-color misalignment can be measured.

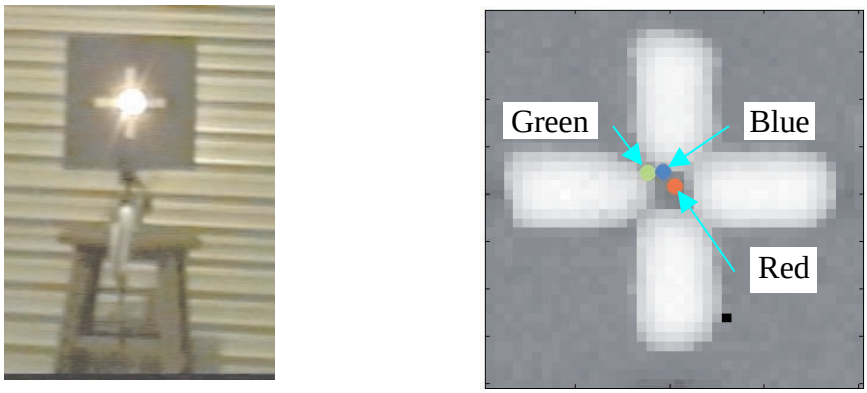

FIGURE 4-7. Alignment target (a), and reflectance image (b)

The resulting misalignment of the sensor is summarized in Table 1. These results

Red -Laser	0.08°
Green - Laser	0.18°
Blue - Laser	0.10°

TABLE 1. Misalignment summary

verify that each of the 1-pixel camera channels is aligned with the laser beam to within the required 0.2º. Therefore, the alignment is good.

4.2.4 Color calibration

The color data of the 1-pixel camera must be adjusted to match the red, green, blue values of a typical colorspace. The process of converting raw sensor measurements into RGB values (shown in Figure 4-8) requires calibration parameters for the black

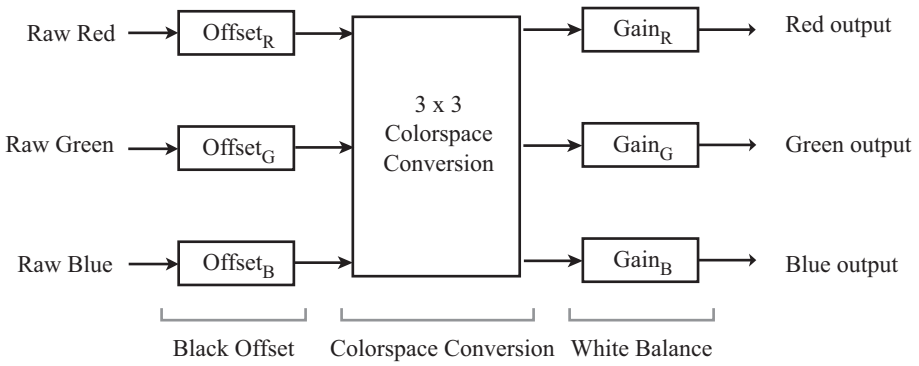

FIGURE 4-8. Black offset and white balance correction

offset, white balance, and colorspace conversion transform.

Black offset calibration establishes a reference for black surfaces, which are to be interpreted as an RGB value of 0,0,0. White balance calibration adjusts the gain of the red, green, and blue channels to provide accurate terrain surface color measurements.

This is not a constant gain adjustment. As discussed in Chapter 3, each of the channels has a different spectral sensitivity, and the spectral content of sunlight varies with different weather situations. Therefore, the proper weighting of the channels must be determined in the lighting environment expected during flight and may need to change according to different conditions.

White balance and black offset are adjusted by scanning a calibrated grayscale target with known surface properties. This is a field operation, as white balancing should be performed under the actual lighting conditions prior to flight.

The grayscale target is shown in Figure 4-9. The target consists of six 8.5" x 11" sheets of Munsell calibrated color material [34]. These sheets provide a lambertian surface with neutral color and a known reflectance. The reflectance of the six sheets (from black to white) is 3.1%, 9.0%, 19.8%, 36.2%, 59.1%, and 90.0%.

FIGURE 4-9. Grayscale target for white balancing

The target is positioned so that the sensor's scan plane intersects the center of all the sheets, as shown in Figure 4-10. In this figure, the scanner is not onboard the helicopter, but scanning can also be done with the scanner mounted on the helicopter. Either method allows each scan to measure the six targets in succession to determine the average raw RGB measurement from each sheet. The best offset and gain for the red, green and blue channels are found by performing a least squares fit to match the raw measurements with the ideal values.

Two issues worth noting are the changing illumination conditions and the 1-pixel camera focus during the calibration. First, it is likely that the illumination is

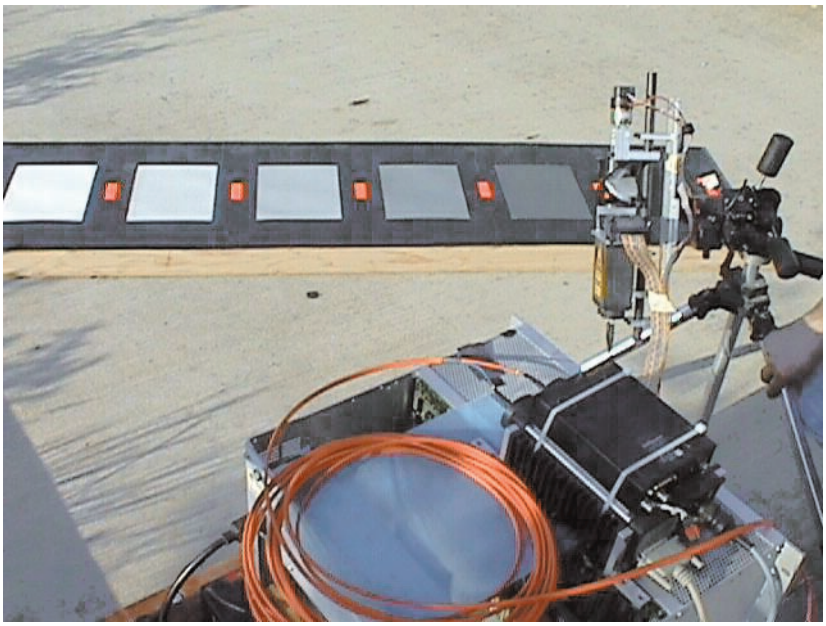

FIGURE 4-10. Color calibration setup

continually varying due to changing sky conditions. The color calibration requires a constant illumination of all six sheets. Therefore, we rapidly scan across the target (at 20 Hz), assume the illumination is constant during a single scan, and do the calibration independently for each scan. The second issue is that the 1-pixel camera may be out of focus at the close range (1-2 meters) where the target is located. This does not adversely affect the measurements, however, since the target is a uniform color, and is large enough to fully contain the out-of-focus sensitive spot.

A dataset taken in bright sunlight is provided as an example. Table 2 lists the average raw measurements for five of the sheets (the 3% sheet was not used for this test), along with the color output expected for each of the calibrated sheets. We use the 1-pixel camera input colorspace for the output colorspace; therefore, the transformation matrix is identity. Each channel is considered independently, and these data points are plotted as circles in Figure 4-11. The offset and gain for each channel is found by fitting a line to the five data points for each channel. The fit lines are shown as dashed lines. The calibration parameters determined for this situation are listed in Table 3.

Raw Measurements				Actual Color		
Red	Green	Blue		Red	Green	Blue
76794	34229	15638		0.900	0.900	0.900
53431	24488	11892		0.591	0.591	0.591
34617	16816	8922		0.362	0.362	0.362
22476	11963	7114		0.198	0.198	0.198
15495	9034	6007		0.090	0.090	0.090

TABLE 2. Example results from grayscale target

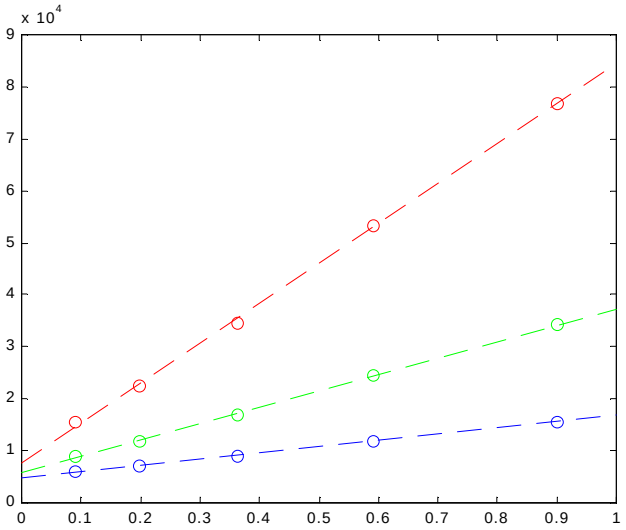

FIGURE 4-11. White balance

Channel	Offset	Gain
Red	7776	76568
Green	5862	31396
Blue	4772	12010

TABLE 3. Typical results from white balance calibration

As expected, each channel responds linearly to reflected light intensity, as shown by the close fit of the line. The parameters found for each channel are subsequently used to convert raw measurements into RGB values. These parameters remain valid while

the spectral content of the illuminating light (sunlight) is similar to the time of white balance calibration.

Noticeable variations in sunlight spectral content can be seen with different sky conditions. Clear days vs. overcast days represent the most significant change. Performing white balance calibration prior to flight compensates for day-to-day variations. Most problematic, however, are partly cloudy days when clouds obscure the sun periodically, thereby causing the spectral content to change often during a single flight.

The parameters, determined by white balancing, not only adjust the channel-to-channel gain to provide proper color readings, but also normalize the magnitude of the readings to provide RGB = (1.0, 1.0, 1.0) for an ideal 100% reflective, lambertian white surface in the illumination situation during calibration. This normalization is done for convenience, and we actually expect the illumination intensity to change over time. As such, unlike during calibration, the RGB readings may exceed 1.0 during periods of brighter sunlight.

4.2.5 Noise characterization

The 1-pixel camera measurements are affected by noise. Characterizing the noise is useful to ensure proper operation and understand the capability of extracting useful signals from the noise. We measured the signal variation during a short scan with the sensor covered to characterize the no-signal noise characteristics of the sensor.

Figure 4-12 shows a histogram of the raw readings acquired by a single channel during a period of 18 seconds. The mean value is 4092, with a standard deviation of ~3. The offset value is as expected, since the ADC used in the 1-pixel camera produces a raw output of 4096 when no input current is present. The shape of the signal distribution is consistent with a natural (normal) noise source, confirming proper operation.

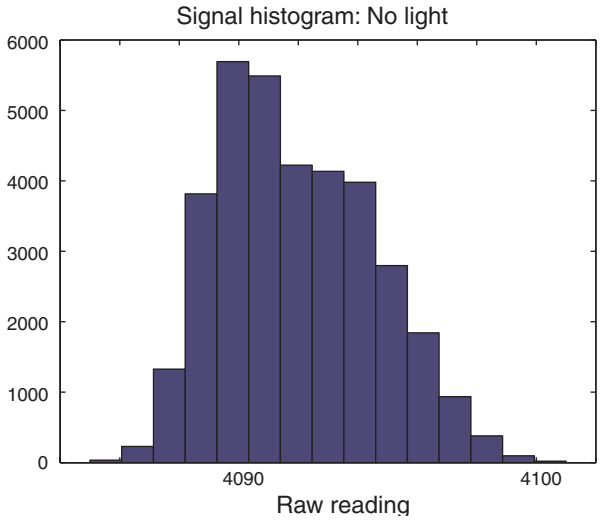

FIGURE 4-12. Histogram of no-signal noise

4.3 Combined sensor verification

We have developed two test configurations to verify proper operation of the combined 1-pixel camera and laser rangefinder system. One employs a pan-tilt head for scanning tests of the combined laser / 1-pixel camera sensor and the other uses a pan-tilt head to move the entire sensor, including the scanning mechanism. These tests provided an incremental testing approach. In this section, we present these verification steps.

4.3.1 Pan-tilt scanner (it works and aligned)

Figure 4-13 shows the 1-pixel camera and laser rangefinder mounted on a pan-tilt unit.

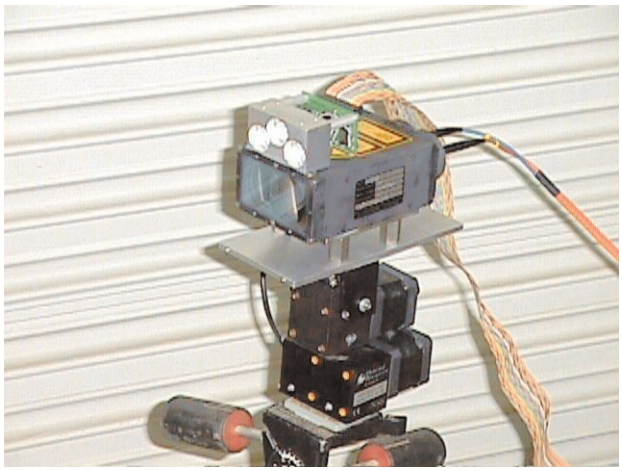

FIGURE 4-13. Laser / 1-pixel camera on pan-tilt test apparatus

They have been aligned for parallel viewing rays without the spinning mirror used on the helicopter. Slowly scanning the pan-tilt unit in a raster style over an area allows the construction of a 3D image of stationary objects surrounding the sensor.

Figure 4-14 shows an outdoor scene scanned to test the sensor. Figure 4-15 shows the resulting 3D colored terrain model of the scene from two viewpoints. The first is from the location of the scanner, and the second viewpoint shows the 3D nature of the model from another location. In this scan, some areas of the image were farther than the

FIGURE 4-14. Outdoor test scene

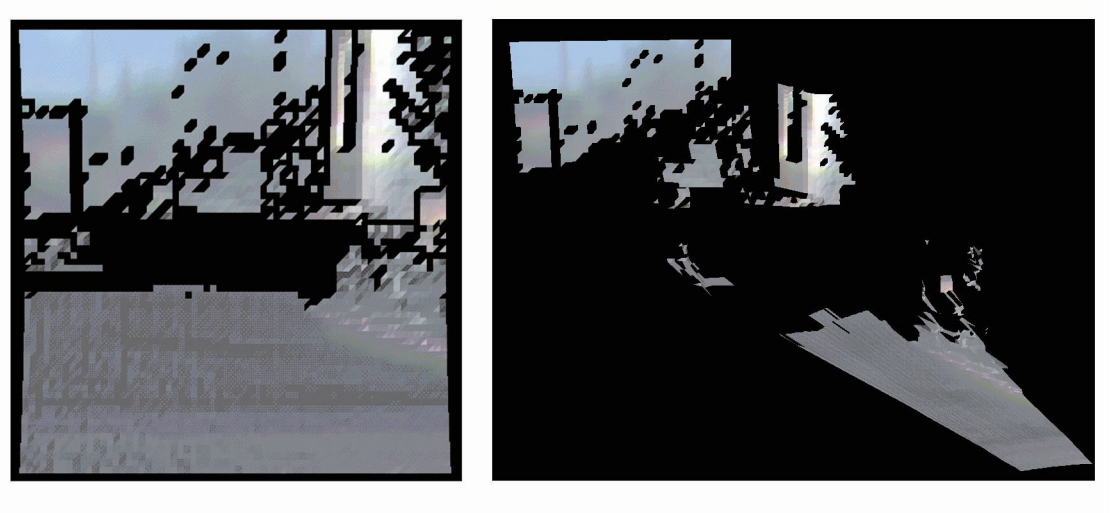

FIGURE 4-15. Two views of color/shape model

100m maximum range of the laser rangefinder, and no range could be determined. In these cases, a range of 100m was assumed, thereby allowing the background colors to be included in the model, although not at the proper location.

This test demonstrated that the 1-pixel camera measurement appears to match the environment. The laser and 1-pixel channels seem to be properly aligned, and the data acquisition circuitry is operational. The color gains were adjusted manually to provide a reasonable color balance. The primary problem encountered in this test was that the color was too dark. The magnitude of the measurements from the 1-pixel camera in areas directly illuminated by sunlight were just enough to be viewable. However, areas

not in direct sunlight were too dark for useful measurements. After this test, we modified the 1-pixel camera to capture more light, in order to obtain useful measurements. The primary modification changed the FOV of the 1-pixel camera to 1º.

4.3.2 Ground scanner (overlapped aperture, timing issues)

Figure 4-16 shows the ground scanner test configuration. This configuration tests all of

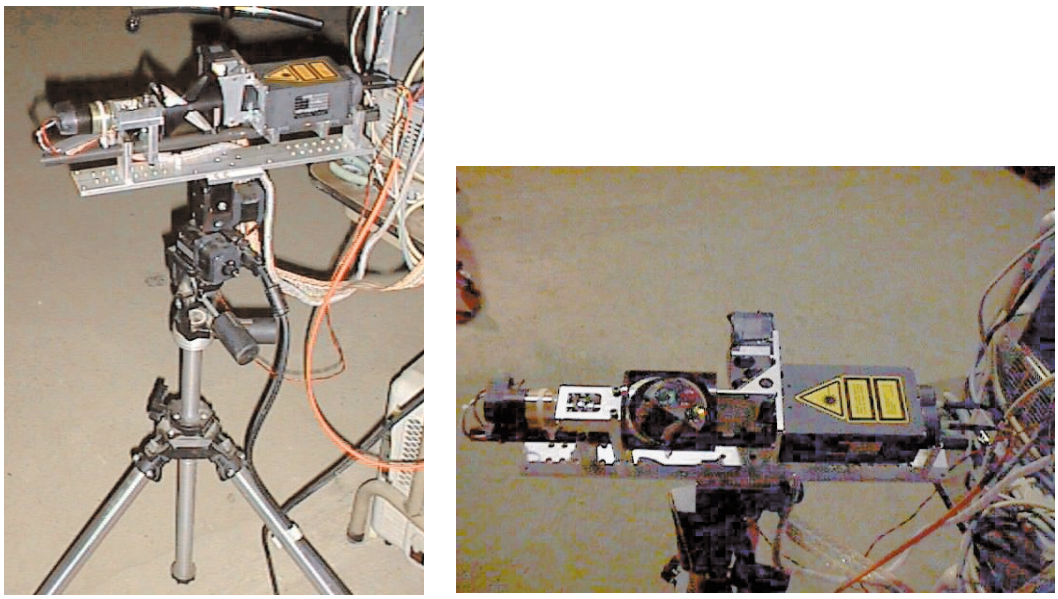

FIGURE 4-16. Ground scanner test setup

the primary scanner components before they are mounted on-board the helicopter. The 1-pixel camera is mounted to the laser rangefinder head with a cold mirror positioned to overlap their optical paths. The 1-pixel camera has been focused and aligned to ensure that all the sensors point to the same location. The spinning mirror is rigidly mounted in front of the sensors to scan its sensing spot over a vertical plane, centered at the mirror. This entire assembly is placed atop the pan-tilt unit, where only the pan axis is used to rotate it up to 360º.

This ground scanner configuration is similar to that of other ground scanners with 2 degrees of freedom. Onboard the helicopter, we will not have the additional degree-of-freedom provided by the pan axis of the pan-tilt unit, but for this test configuration, the 2-DOF setup allows an entire 3D volume to be scanned from a single location.

This test is intended to verify the high-speed timing of the rangefinder, 1-pixel camera signals, and the spinning mirror. Figure 4-17 shows a point cloud of the helicopter lab

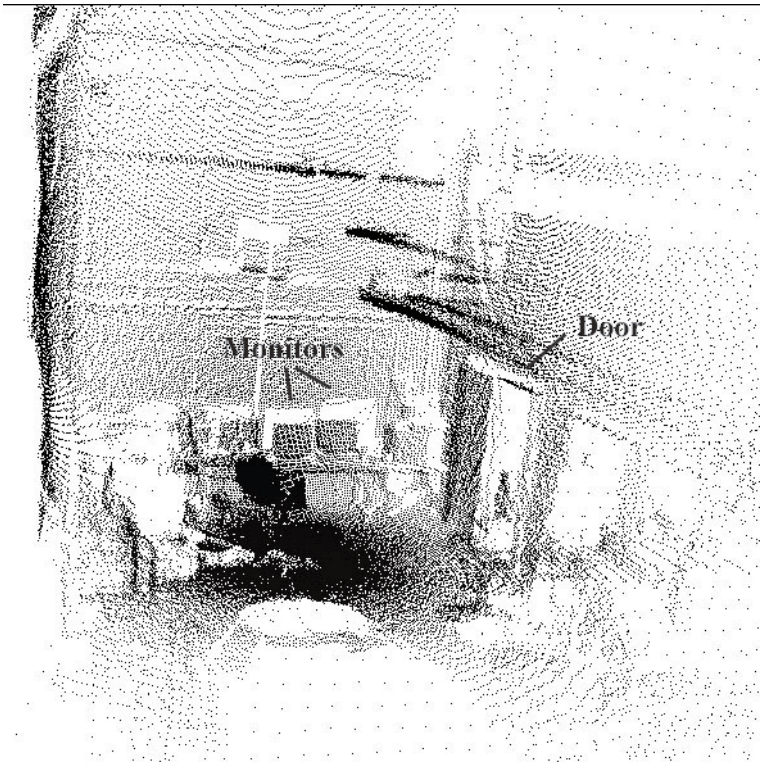

FIGURE 4-17. Scan inside room

at CMU. No color is included, since this scan was performed indoors, and there was not enough light for good color measurements.

4.4 Ground-based scanner mount calibration

For accurate terrain calculations, knowledge of the position and orientation of where the scanner is mounted to the helicopter is required. In the following subsections, we describe the calibration procedures to locate the scanner position and orientation with respect to the helicopter, and how we develop a ground-based method to determine scanner orientation, which is accurate, but time consuming. Later, in Section 4.6.2, we present an in-flight calibration method for fast recalibration in the field.

4.4.1 Mirror location

Once the scanner has been rigidly mounted on the helicopter, a number of parameters are needed to locate the scanning mechanism. First, the position of the spinning mirror center must be measured in the HELI frame, ${}^{HELI}\vec{V}_{\text{sensor}}$. This is accomplished effectively using basic measuring tools such as rulers and calipers. Two issues make this measurement problematic: first, the origin of the HELI frame is physically located inside the rotor drive train gearbox and cannot be directly accessed. Second, the axes of the HELI frame are not easy to determine precisely on the helicopter body, so the measured vector is likely to be skewed by a few degrees. Fortunately, the required accuracy of the scanner location on the helicopter is not exceptionally critical; a measurement within 1 cm of the actual position is sufficient. We found that careful measurements can readily provide this level of accuracy.

Next, we need to measure the distance from the front face of the laser rangefinder to the mirror center, LASER_RANGE_OFFSET. This measurement requires an accuracy of 1 cm, and is effectively determined with a ruler.

FIGURE 4-18. Forward mounted scanner configuration

For the mounting configuration shown in Figure 4-18, we found LASER_RANGE_OFFSET = 0.125 meters and

$$ {}^{HELI}V_{\text{sensor}} = \begin{bmatrix} 0.301 \\ 0.356 \\ -0.038 \end{bmatrix} \qquad \text{(meters)} $$

Any time the scanner mount is moved on the helicopter, the mirror position must be recalibrated. Typically, repositioning the scanner will not affect the distance from laser to mirror, so LASER_RANGE_OFFSET from previous calibrations may be reused.

4.4.2 Scanner Orientation

Determining the orientation of the scanner mount on the helicopter, ${}^{HELI}_{SENS}\boldsymbol{R}$, is a significantly more difficult challenge than finding its location. These angles cannot be directly measured to the required accuracy. The orientation must be known very accurately to eliminate systematic errors due to a misdirected beam. These errors

increase linearly with distance, so the orientation must be accurate to better than 0.02º to achieve the same 1 cm resultant error effect at 50 meters range.

We achieve the required accuracy by measuring the 3-D location which the emitted laser beam hits at long (> 50m) distances, instead of attempting to measure the angles of the scanner at the helicopter. This type of measurement requires two capabilities: the ability to see where the laser hits the ground, and the ability to accurately measure the 3-D coordinates of these points over a large distance. We are able to view the infrared laser spot using a commercial hand-held infrared (IR) scope (Find-R-Scope), and mark the spot's center. A precise differential GPS receiver can be used to measure the 3D coordinate of the spot. A drawback of this approach is that the GPS surveying equipment requires a clear view of the sky, and the IR viewer cannot see the laser spot in daylight.b Therefore, it must be performed outdoors and at night, a requirement that makes it logistically challenging.

Referring to Figure 4-19, we collect the data by performing the following procedure:

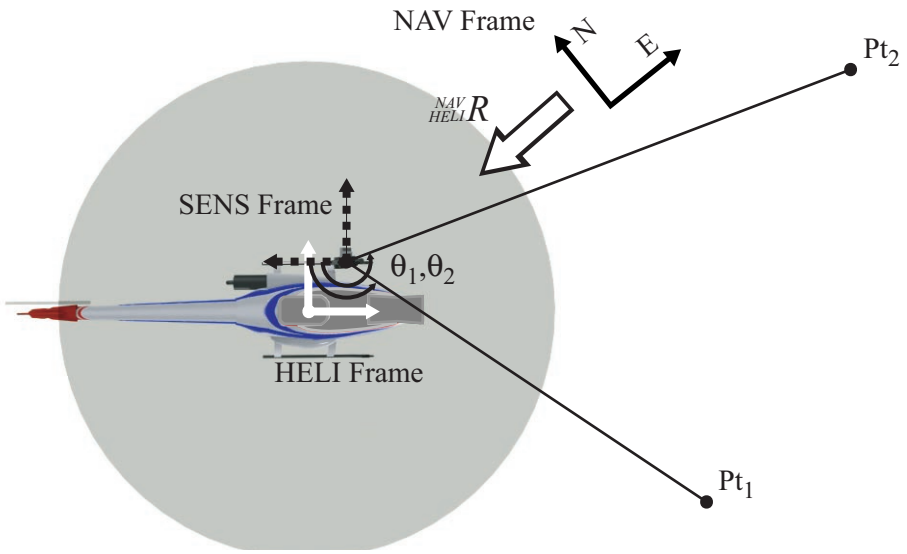

FIGURE 4-19. Scanner mount ground calibration

1. Position the helicopter such that the laser scan plane is roughly parallel to the

ground, and ensure that the helicopter is stationary.

2. Record the helicopter's attitude ($\overset{NAV}{HELI}\boldsymbol{R}$) as determined by the helicopter's state estimator.
3. Position the scanner mirror at a series of different angles and locate the laser spot at approximately 50 meters distance. Record the 3D coordinate of each laser spot (PT_k) and the corresponding mirror angle (θ_k).
4. Finally, measure and record the 3D coordinate of the mirror center itself, (O).

The scanner orientation on the helicopter, $\overset{HELI}{SENS}\boldsymbol{R}$, is determined as follows:

1. Calculate a unit vector in the NAV frame is calculated from the mirror center to each spot, based on the GPS measurements:

$$V_k = \frac{PT_k - O}{\left\| PT_k - O \right\|}$$

2. Find a unit vector in the SENS frame from the mirror center to each spot, based on mirror angle:

$$S_k = \begin{bmatrix} 0 \\ -\sin\theta_k \\ \cos\theta_k \end{bmatrix}$$

3. Align V_k and $\overset{NAV}{HELI}\boldsymbol{R} \cdot \overset{HELI}{SENS}\boldsymbol{R} \cdot S_k$ to ensure that when the correct scanner orientation is used to determine $\overset{HELI}{SENS}\boldsymbol{R}$ that the corresponding vectors are coincident in the NAV frame. This relationship provides an error metric for evaluating the fitness of an arbitrary scanner orientation, $\overset{HELI}{SENS}\boldsymbol{R}$. We define an error function evaluated over all of our measurements as:

$$\text{Error}(\overset{HELI}{SENS}\boldsymbol{R}) = \sum_k (\text{acos}(V_k \cdot (\overset{NAV}{HELI}\boldsymbol{R} \cdot \overset{HELI}{SENS}\boldsymbol{R} \cdot S_k)))^2$$

In other words, the error is the sum of the square of the angular misalignment

between each of the corresponding vector pairs.

4. Determine the best-fit scanner orientation by utilizing an iterative minimization technique (Nelder-Mead simplex) to minimize the error function. To perform the minimization, we must parameterize ${}^{HELI}_{SENS}\boldsymbol{R}$ with three independent variables, such as the $\text{Euler}_{Z'Y'X'}$ angles. The Nelder-Mead simplex algorithm minimizes the error function over these three angles. The best-fit parameters, therefore, are our best estimate of the calibrated scanner orientation.

To perform these measurements, we must make two assumptions: First, we assume that the front (reflecting) surface of the scanning mirror is flat and mounted at precisely 45º from the motor shaft. This is critical, as a non-45º mount will result in laser scanning in a cone, as opposed to the plane we assume. The second assumption is that atmospheric bending effects are negligible and the laser beam travels to the target in a straight line. Both of these assumptions are reasonable for our system, as we have carefully checked the mirror angle (measured to be 45.03º) and expect that over the relatively short ranges of our laser (100m), atmospheric bending effects are negligible.

4.5 Pre-flight system tests

We devised a test to build color 3D terrain models without flying the system. Test flights are demanding, and must be carried out only after careful verification. This test was designed to verify complete operation of the scanning system in a representative environment. All system components, including state estimation, scanning, synchronization, and data collection, were operating as if in flight during this test. Correct operation would yield a color 3D model of the surrounding area.

The helicopter is placed on a table with a rotary bearing (lazy Susan), which allows the helicopter to freely spin around the axis of its main rotor shaft. Since the scan plane extends out from either side of the helicopter, the area around the helicopter can be scanned by spinning the entire helicopter. Figure 4-20 shows the helicopter being manually yawed to scan the surroundings. The result of this mapping tests a significant portion of the system. However, effects of vibrations, translation, and roll/pitch rotation can not be verified with this test. We consider this an effective test of system performance, especially given the simple necessary setup

FIGURE 4-20. Ground scan setup

Figure 4-21 shows a picture of the outdoor test site chosen for the test. The site is a large open grass field with a concrete pad. To provide a variety of colors in the scene, we placed an array of nine 1 meter square color targets on the concrete pad. We also placed the grayscale color calibration target in the scene. To the right (not shown in the picture) was a blue van which we use as our base station and helicopter transportation vehicle.

To image the entire scene, the helicopter was manually yawed 360º over ~30 seconds, with the terrain sensor scanning a forward looking plane. The terrain sensor measurements and state estimates were combined to localize each 3D terrain point in the world, along with its associated color. These colored point clouds constitute a basic terrain model of the site.

The 3D color point cloud can be displayed from an arbitrary viewpoint. Figure 4-22 shows one viewpoint of the model. In the foreground, the grayscale color calibration target is visible. Prior to this scan, the color calibration procedure was performed with the scanner; the results were used in calculating the color of the model. Therefore, the calibration sheets appear to be neutral in color.

When building the model, only terrain points with valid range measurements were included. This explains the lack of samples from the trees in the background, since they are exceed the 100m range of the rangefinder. It also explains the two holes in the scene corresponding to black objects, which did not return enough of the laser light to provide a good measurement.

An interesting effect is visible on the pole located at the right side of the scene. In the picture, the pole is a single color (blue); however, in the terrain model, the color changes part way up the pole. This is a result of the large (1 degree) sensitive spot size of the 1-pixel camera; this spot exceeds the width of the pole. Therefore, the color of the pole includes its background color. In this case, the background of the pole (from

FIGURE 4-21. Picture of test scene

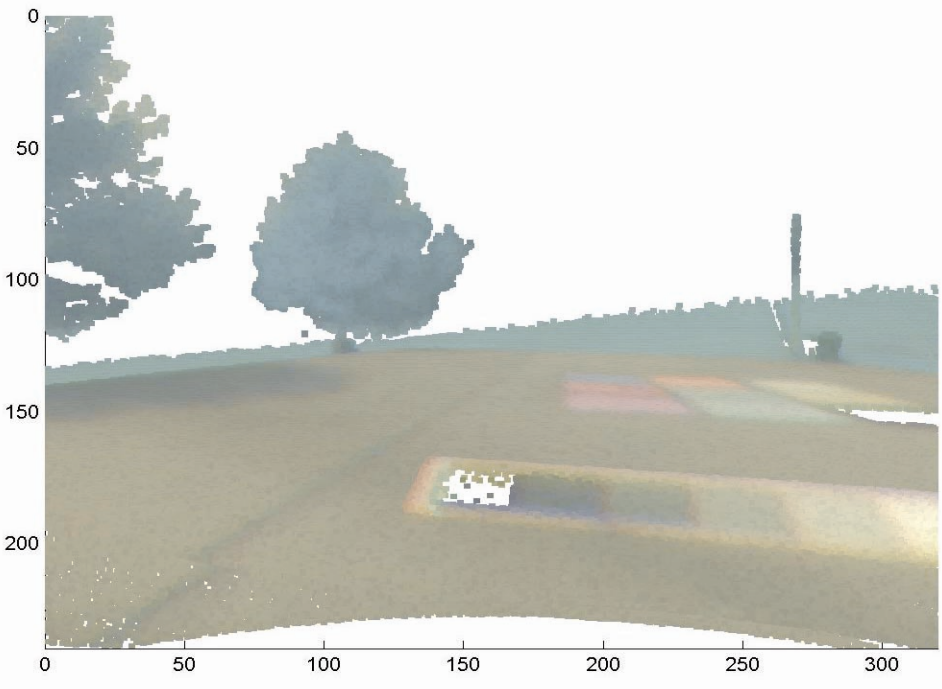

FIGURE 4-22. 3D color model of test scene

FIGURE 4-23. Picture of targets

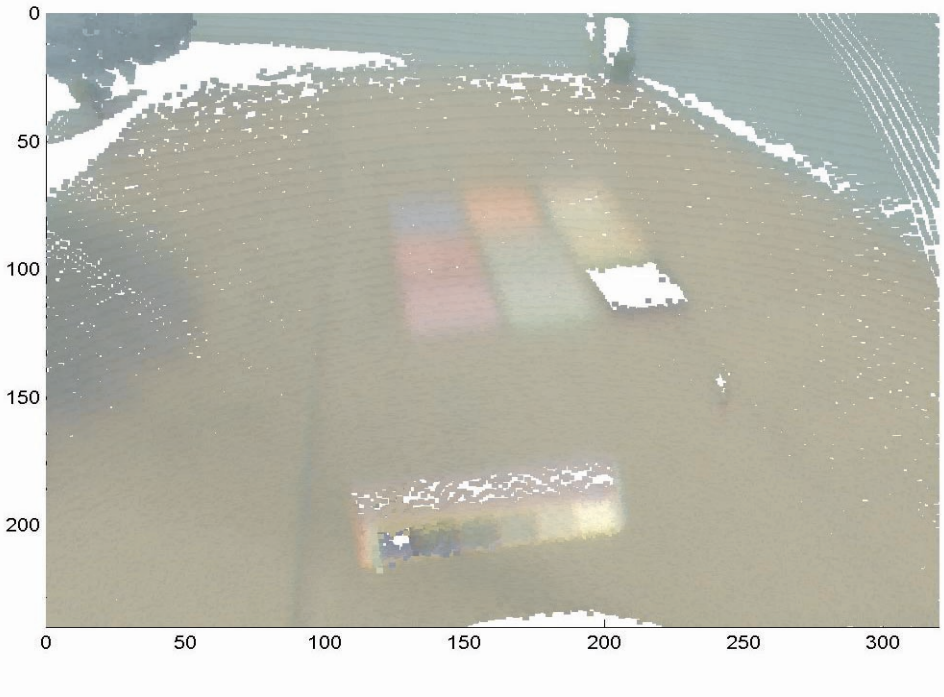

FIGURE 4-24. 3D color model of targets

FIGURE 4-25. Another view

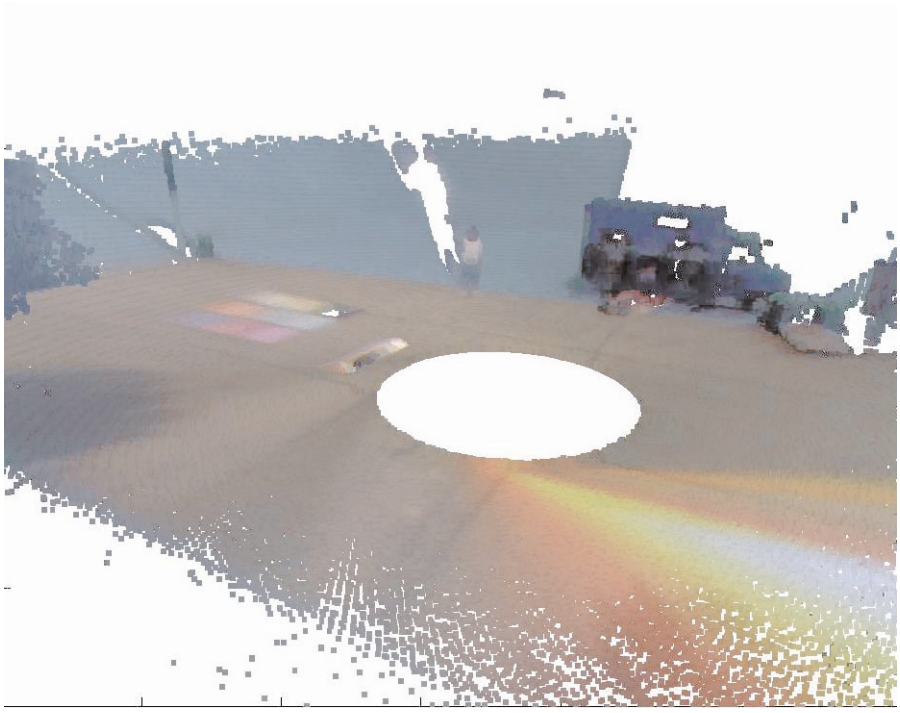

FIGURE 4-26. Remote view of terrain model

the scanner's point of view) changes from green grass to the dark under-tree area. This change corresponds to the change of the pole color from green to black in the model.

Figure 4-23 shows an elevated picture of the color targets placed on the concrete pad. Figure 4-24 shows these color targets in the terrain model. The color of each target is distinguishable, indicating that the sensor is providing good color measurements. Again, the black colored target absorbed too much of the laser light to allow measurements to be made. Therefore, we have a lack of data from this region.

Figure 4-25 shows another terrain model which includes our blue van and a person saying, "Hi". The lack of data points on the grass behind the person is caused by a range shadow, since these areas were not visible from the scanner's viewpoint. Also, a close look reveals that the person is wearing blue shorts and white socks. A final view of this terrain model, shown in Figure 4-26, was generated from a viewpoint a distance away from the scene. Here, most of the scene can be seen, including the color targets, the calibration target, and the blue van. The circle of missing data in the center of the pad is the helicopter's location during the scan. Since the scan plane is forward looking, it does not intersect the ground directly beneath the helicopter.

A troublesome effect is shown in the right side of the Figure 4-26. The bright, miscolored wedge in the image is a result of sunlight directly falling on the sensor. The 1-pixel camera is designed to reject off-axis light sources by the use of baffles and a pinhole. However, the energy in direct sunlight is large enough to achieve a significant interference signal, despite this attenuation. In later scans, additional light shielding was added around the sensor to minimize the opportunities for direct sunlight to reach the sensor.

4.6 Flight testing

4.6.1 Laser Reference Targets

For calibration and evaluation purposes, we use artificial targets that are easy to identify in the scanner's data. The targets provide a known relationship among all laser measurements and their corresponding 3D location in the world. This is extremely valuable in evaluating and calibrating the system.

We use small retroreflective targets and look for their signature in the laser reflectance data during each scan. Figure 4-27 shows the laser reflectance data across a single scanline in which a single sample hits a retroreflective target. The retroreflective target provides a much stronger spiked laser return than the diffuse surfaces in natural terrain, making it very easy to identify in the scan. Using this fact, computer programs can automatically search each scanline to identify these spikes, and locate the scanner mirror angle when the target was hit. In practice, we have found that the combination of an intensity threshold (> 130) and a minimum intensity jump (> 50) are a reliable detection criteria.

The retroreflective target is identified only when the laser spot overlaps the target. Therefore, for each target hit, the center of the laser beam is guaranteed to be within a small, known window around the center of the target. A 5cm diameter target and a 16cm diameter laser spot (50m range, 3.2mrad divergence) can locate the laser reflection point within 10cm of the target center. Most of the uncertainty in this measurement is attributed to the large laser spot size.

It is important to note that hitting a small target is a random event. It is quite possible to miss a target even if the system flies directly overhead. This error is primarily due to the non-uniform sample distribution that occurs during flight. Clearly, slower fly-overs and multiple passes improve the chances of hitting the target. Additionally, making the target itself larger also improves the odds at a cost, however, of reduced accuracy.

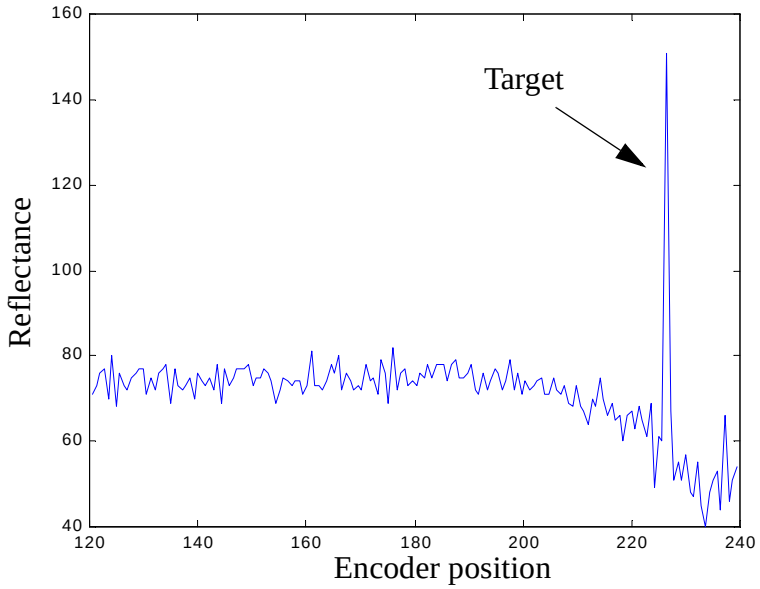

FIGURE 4-27. Retroreflective target signature in laser reflectance channel

It should also be noted that, while this method works well for the Riegl rangefinder, some other laser rangefinders (notably AMCW systems) have difficulty dealing with the especially strong returns from the retro-reflector, and would require an alternative method.

4.6.2 In-flight calibration

The ability to quickly recalibrate the sensor after reconfiguring the scanner mount is essential for most useful terrain modeling applications. The calibration procedures described earlier in this chapter are not suitable for rapid recalibration. They require extensive setup, and can only be performed in special environmental structures. It is desirable to develop alternative, possibly less accurate, calibration methods which are quicker and easier to perform, therefore enabling a whole class of mapping applications. Ideally, the scanner should be reconfigureable in the field, requiring new calibration before deployment.

The terrain sensor itself is a single unit which can be mounted onto the helicopter in different ways. Since the terrain sensor scans in a single plane, the position and orientation the sensor is mounted with respect to the helicopter determines where the terrain is scanned from the helicopter. When changing the sensor mount, the sensor unit itself is not effected, so many of the calibration parameters are unchanged. These include the laser range offset, mirror angle offset, 1-pixel camera calibration parameters, focusing, alignment, timing, etc. The two calibration parameters which do change are the position, ${}^{HELI}\boldsymbol{V}_{\text{sensor}}$, and orientation, ${}^{HELI}_{SENS}\boldsymbol{R}$, of the sensor on the helicopter.

Our approach for rapidly reconfiguring the scanner is to perform the calibration in flight. This approach, does not require the extensive ground procedures and estimates parameters from a small set of carefully collected terrain measurements during a calibration flight. A single fixed target on the ground is used as a reference point on the terrain. Multiple laser returns off the target are collected from various locations as the helicopter flies around the target, as shown in Figure 4-28. If the calibration parameters are correct, the system localizes all of these returns at the same 3D location. Any disagreement in the measurements indicates errors in the calibration parameters. Minimizing these disagreements corrects the calibration parameters.

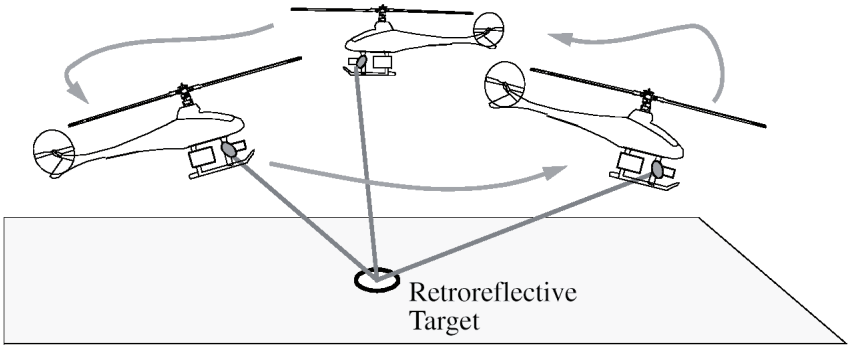

FIGURE 4-28. Multiple hits from single target for in-flight calibration

Potentially, both the position and orientation calibration parameters can be estimated from data collected in flight. However, we choose not to estimate the sensor position in this manner, since it can be measured with a ruler in the field prior to flight with sufficient accuracy. This reduces the in-flight calibration to 3 degrees-of-freedom, in the orientation rotation matrix. Therefore, when the scanner mount is reconfigured, we measure a new value for ${}^{HELI}\boldsymbol{V}_{\text{sensor}}$, and then use in-flight calibration to estimate ${}^{HELI}_{SENS}\boldsymbol{R}$.

In-flight calibration: Ideal case

Under ideal conditions, the sensor mount orientation can be determined from only two target measurements, each from a properly chosen sampling site. For each target measurement, the position and attitude of the helicopter, mirror angle and laser range are recorded.

The ideal situation assumes that the laser beam is narrow enough, and the target is small enough that every measurement which hits the target is actually measuring the same exact 3D point. Errors in other measurements, laser range, helicopter state, etc., are also ignored under ideal conditions. Both of these assumptions will be addressed in the next section.

First, to see that two measurements of the target can determine the scanner mount orientation, we apply the terrain measurement equation to each measurement.

$$ {}^{NAV}\boldsymbol{V}_{\text{target}(1)} = {}^{NAV}\boldsymbol{V}_{\text{heli}(1)} + {}^{NAV}_{HELI}\boldsymbol{R}_{(1)} \cdot \left({}^{HELI}\boldsymbol{V}_{\text{sensor}} + {}^{HELI}_{SENS}\boldsymbol{R} \cdot \left(\begin{bmatrix} 0 \\ -\text{RANGE}(1) \cdot \sin(\text{ANGLE}(1)) \\ \text{RANGE}(1) \cdot \cos(\text{ANGLE}(1)) \end{bmatrix} \right) \right) $$

$$ {}^{NAV}\boldsymbol{V}_{\text{target}(2)} = {}^{NAV}\boldsymbol{V}_{\text{heli}(2)} + {}^{NAV}_{HELI}\boldsymbol{R}_{(2)} \cdot \left({}^{HELI}\boldsymbol{V}_{\text{sensor}} + {}^{HELI}_{SENS}\boldsymbol{R} \cdot \left(\begin{bmatrix} 0 \\ -\text{RANGE}(2) \cdot \sin(\text{ANGLE}(2)) \\ \text{RANGE}(2) \cdot \cos(\text{ANGLE}(2)) \end{bmatrix} \right) \right) $$

Since both of these measurements are hitting the same 3D target point, the target position, itself, need not be known to solve for ${}^{HELI}_{SENS}\boldsymbol{R}$.

$$ {}^{NAV}\boldsymbol{V}_{\text{target}(1)} = {}^{NAV}\boldsymbol{V}_{\text{target}(2)} $$

Setting the left side of these two equations equal to each other, provides a single equation we must solve. The value of every term in the equation is known except for the 3x3 rotation matrix ${}^{HELI}_{SENS}\boldsymbol{R}$, which is the calibration parameter we seek. Even though the unknown rotation matrix contains 9 values, there are actually only 3 independent values to determine. The equation provides 3 scalar equations, one corresponding to each row. Therefore, in the absence of degenerate measurements, the calibration parameter can be calculated from these two measurements.

An important question is: From which sites should the two measurements be taken from to get the best estimate? Obviously, two identical measurements do not provide sufficient information for the calibration. There are other situations in which two measurements do not provide sufficient information, and we must avoid these conditions when calibrating.

In order to determine the degenerate situations, we must analyze the equations discussed for the two points above. In particular, we identify situations where the rotation matrix is not uniquely determined. From this analysis, we determined the following four degenerate conditions:

1. ANGLE(1) = ANGLE(2)
2. ${}^{NAV}_{HELI}\boldsymbol{R}_{(1)} = {}^{NAV}_{HELI}\boldsymbol{R}_{(2)}$
3. ${}^{NAV}\boldsymbol{V}_{\text{heli}(1)} + {}^{NAV}_{HELI}\boldsymbol{R}_{(1)} \cdot {}^{HELI}\boldsymbol{V}_{\text{sensor}} = {}^{NAV}\boldsymbol{V}_{\text{heli}(2)} + {}^{NAV}_{HELI}\boldsymbol{R}_{(2)} \cdot {}^{HELI}\boldsymbol{V}_{\text{sensor}}$
4. RANGE(1) = 0, or RANGE(2) = 0

The first criteria requires that the scanner's mirror angle must not be the same between the two target measurements. Figure 4-29a illustrates an example situation where the target is measured with the same mirror angle from two different sites. Figure 4-29b shows this situation as viewed from the HELI frame. Since both measurements occur along an identical axis relative to the helicopter, any scanner orientation which causes the scan plane to pass through this axis is consistent with the data and cannot be uniquely identified.

The second criteria requires that the helicopter's attitude must not be the same between the two target measurements. Figure 4-30 illustrates an example situation where the target is measured from two different sites with the same helicopter attitude. This condition poses a problem since the 3D position of the target is not known and cannot be determined from these measurements. The measurements are consistent with the target lying anywhere on a circle in space around an axis connecting the two scanner locations. Since a different scanner orientation can be found for each potential target location, the orientation cannot be uniquely identified.

The third condition requires that the position of the scanner must not be the same between the two measurements. Figure 4-31 illustrates an example situation where the

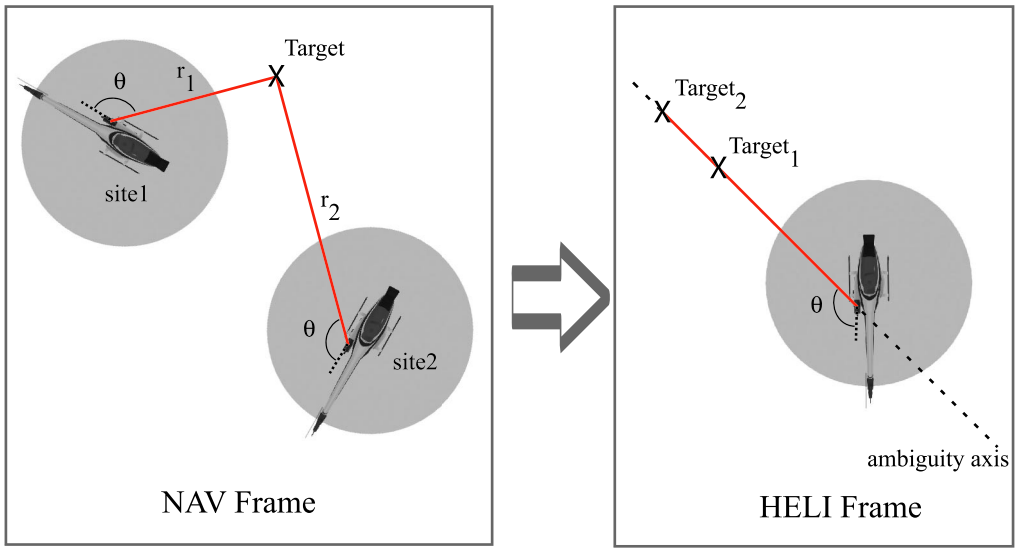

FIGURE 4-29. Degenerate condition 1: same mirror angle

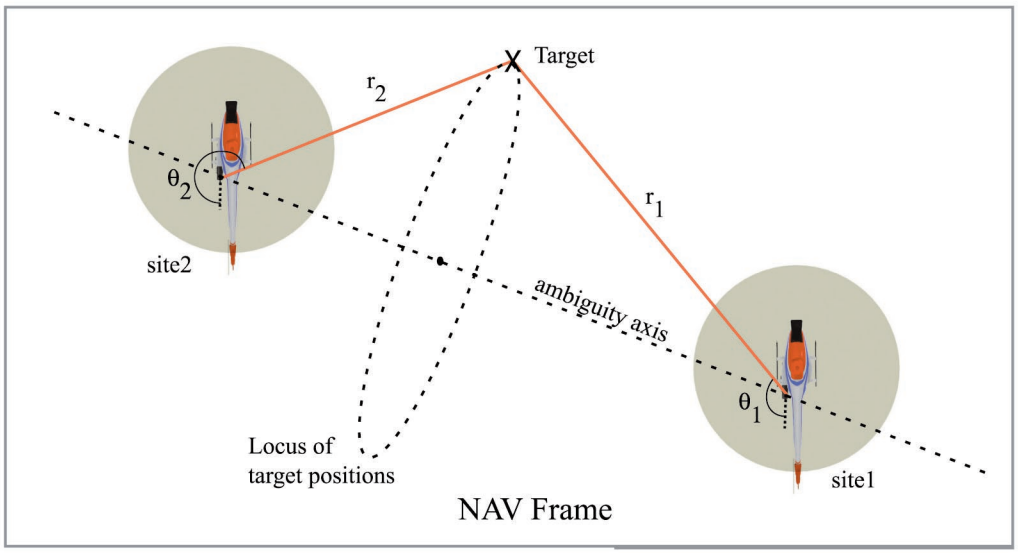

FIGURE 4-30. Degenerate condition 2: same helicopter attitude

target is measured with the scanner at the same location, but with a different attitude. In this case, the first measurement constrains the target point to lying anywhere on a sphere, centered around the scanner. The second measurement further restricts the target point to lie on a circle, but cannot uniquely determine a single position. The circle is centered around the scanner, and is in a plane normal to the axis around which the helicopter attitude changes. Since a different scanner orientation can be found for each potential target location on the circle, the orientation cannot be uniquely identified.

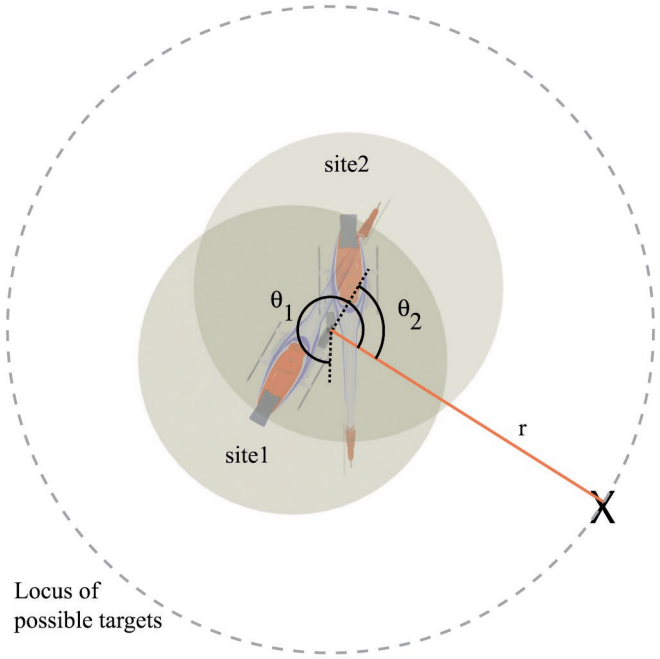

FIGURE 4-31. Degenerate condition 3: same scanner position

The last two degenerate cases are a result of not being able to determine the location of the calibration target from the measurements. If the target position were measured prior to flight, these problems could be avoided. However, it is not easy to accurately measure the target position. Therefore we keep the calibration procedure quick and simple by assume the target position is unknown.

In-flight calibration: Real-world case

In practice, two measurements are not sufficient to reliably calibrate the scanner mount. First, there is noise inherent in each measurement, preventing an exact solution of the calibration parameters. Second, the laser spot size and target shape are not infinitesimally small, so they cannot pinpoint a 3D location in space without significant uncertainty. These effects must be taken into account by a real-world calibration procedure.

To reduce noise, additional target measurements should be used in the calibration. A best-fit calibration from a set of measurements is better than using only two measurements. The challenge is to collect a useful set of measurements which avoids degenerate cases, reduces the noise, and does not bias the solution. An effective approach would select both multiple sites for measurements and acquire multiple measurements at each site to compensate for noise and avoid the degenerate conditions.

In the idealized discussion, we assumed all of the measurements are from the same exact 3D point, where the laser beam always hits the center of the target. However, this is not the case in practice. The laser spot size has an appreciable diameter, depending on the distance and the target has significant dimensions. The laser reflection from the target can actually be from any location within the ambiguous volume created by the overlap of the laser spot size and the target area. Figure 4-32 illustrates several potential measurements which “look” as if they are at the same place to the laser rangefinder, since they hit the calibration target, but are actually separated significantly on the ground. We must consider this effect in the calibration process to maximize the calibration accuracy. Namely, we must collect a sufficiently large set of hits surrounding the target from each site to statistically localize the center. Also, we must employ error models which consider the ambiguity of measurements due to finite spot size and target area. This type of uncertainty analysis and error modeling was

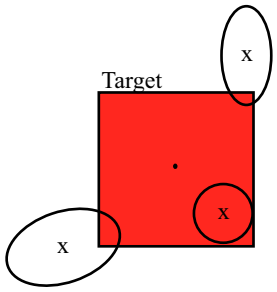

FIGURE 4-32. Center of laser spot not necessary at center of target

considered but not pursued as a research topic in this thesis. We employed large sample averaging as the first attempt in calibration data collection with good results.

In-flight calibration: Example

Here we present an example of this in-flight calibration from a mapping mission (See Section 5.4). The scanner was configured in a down looking scan plane. The target was a 12cm x 12cm retroreflector, and is shown in Figure 4-33, centered on a white ‘X’.

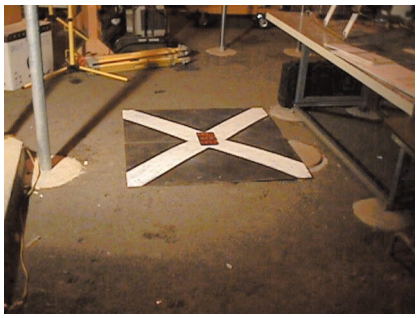

FIGURE 4-33. Picture of retroreflective target

The helicopter flew to acquire measurements from four different sites. These sites were chosen to provide a suitable set of measurements consistent with the criteria presented above. At each site, the helicopter hovered for a period of time to acquire a number of target measurements. Figure 4-34 shows the helicopter location and attitude

chosen for each site, and Table 4 lists the average measurements achieved from each site. Note that sites 1 and 2 have similar attitudes; sites 1 and 4 have the scanner at the same rough location; and sites 3 and 4 have almost the same mirror angles. These measurements do not lead to a degenerate case, however, since the overall data spans a significant range of each important parameter (position, attitude, mirror angle). With more than two sites, they do not all have to be different.

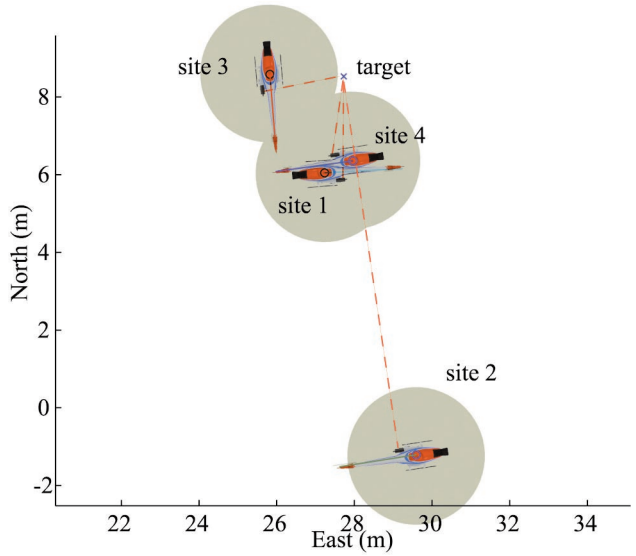

FIGURE 4-34. Four sites for calibration

The 89 target measurements acquired from these sites were used to determine the best-

	Mirror angle	Range (m)	Position (m)	Yaw	Number of samples
Site 1	168.0º	10.2 m	28.0, 6.4, 11.2	7.9º	8
Site 2	135.5º	13.9 m	29.6, -1.2, 11.2	8.9º	8
Site 3	190.6º	10.0 m	25.8, 8.6, 11.1	94.6º	63
Site 4	195.0º	10.3 m	27.2, 6.0, 11.2	-177º	10

TABLE 4. Average measurements from calibration sites

fit scanner orientation. Similar to the minimization procedure discussed in

Section 4.4.2, we parameterize the scanner orientation as a set of three Euler angles, and perform a Nelder-Mead minimization of an objective error function.

The error function compares where each measurement localizes the target, with the actual target location. The actual target location was taken to be at the mean of all 3D target coordinates, derived from each measurement.

The comparison was based on the 3D distance error between the estimated and actual target positions. This error value was weighted in favor of measurements taken at shorter range, since the target position is expected to be more uncertain as the distance increases. Attitude uncertainties and laser beam divergence has a greater effect from larger distances.

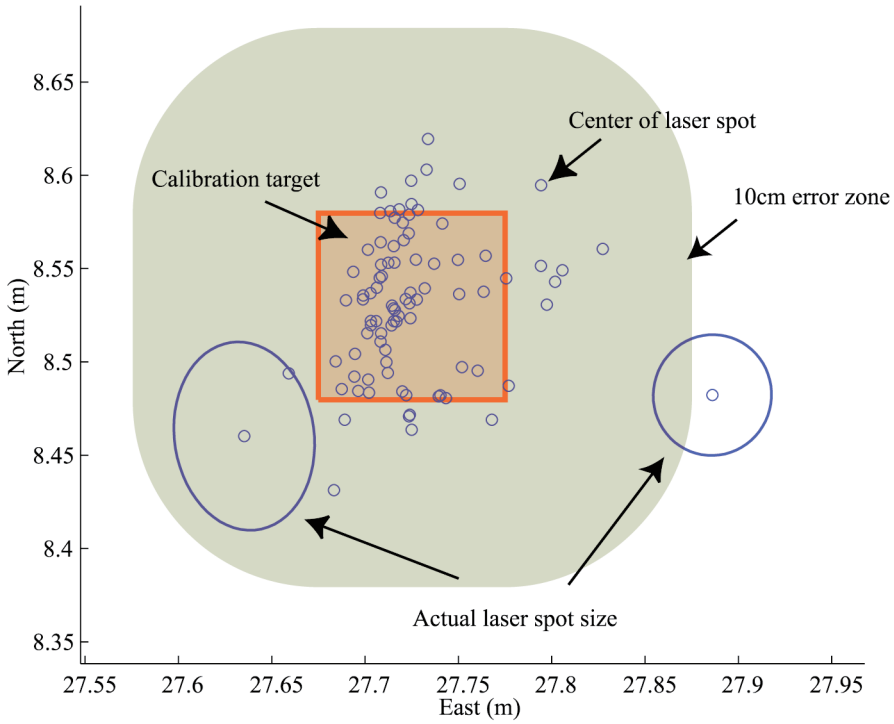

FIGURE 4-35. Resulting position of each target due to calibration

Figure 4-35 shows a 2D distribution of the localized center points for each sample using the optimized calibration parameters. The red square is 10cm x 10cm, and shows

the size of the target. The outline of the laser spot has only been drawn for two of the samples for clarity. A shaded region surrounding the target is also included to show the expected target region when allowing for 10cm of lateral error.

The objective of the optimization process is to determine the scanner mount orientation which bring all of the measurements together such that their laser beam spot overlaps the target. In the figure, a majority of the measurements are centered within the target. These calculations are therefore consistent with the target measurement invariant, that the target and spot must overlap. There are some other measurements, too far from the target for overlap, which are not consistent. The difference can be attributed to either measurement error of incorrect calibration parameters. Since the optimization procedure found the best calibration based on this data, the remaining difference is likely due to measurement error. These points are all consistent when 10cm of lateral error (a reasonable level) is allowed for, shown in the figure where the laser spot overlaps the 10cm error target.

The furthest sample is approximately 10 cm too far away, suggesting errors of at least 10 cm. This suggests a good calibration result, since the system is expected to have errors of 10-20cm.

4.6.3 Terrain modeling accuracy verification

To verify that the terrain modeling system was operating properly, we tested the accuracy of the terrain model. We placed a number of artificial targets across our test site. Each target was accurately surveyed to establish the ground truth measurement, and then flown over to map its 3D location. We could then evaluate the mapping system's accuracy by comparing the 3D position of each target in the model with the ground truth values.

Retroreflective targets are easy to identify in the terrain measurements, and can achieve a small enough target to enable us to verify accuracy. During this test, the terrain scanner was mounted in a forward looking scan plane. Due to the limited acceptance angle of the retroreflective material (+/-60º from normal), it was not possible to use a flat, horizontal target on the ground. To allow returns from all directions, we built cylindrical retroreflective targets on a pipe vertically protruding from the ground, as shown in Figure 4-36. A 1.5" wide retroreflective tape fully

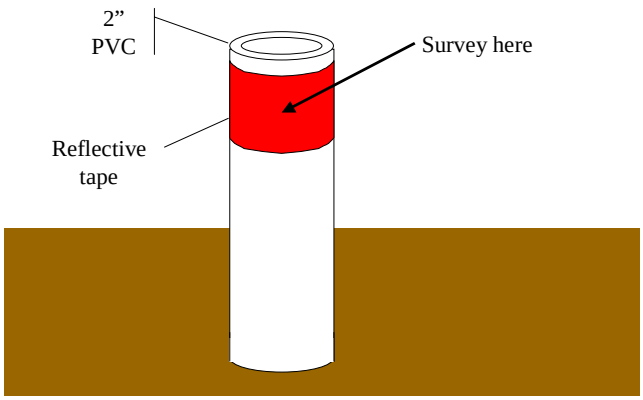

FIGURE 4-36. Cylindrical retroreflective target

encircled the outside of a 2” diameter PVC pipe target. This provided a maximum target size of 2” x 1.5”.

Overall, 31 targets were placed across the test site, covering an area approximately 70m x 70m.

To ensure accurate ground truth, the actual location of the center of the retroreflective cylinder for each PVC pipe was measured using two instruments. First, an accurate GPS receiver was used to measure the 3D position. Next, a precision surveying instrument, Leica total station, surveyed all of the targets as well. The GPS-based ground truth is desirable, since it measures position in the world geodetic frame, our desired mapping coordinates. A comparison of these measurements provide verification of the absolute accuracy of the terrain measurements. The total station measurements, on the other hand, are more accurate than the GPS (several millimeters, vs. several centimeters), so they provide a very accurate ground truth for the relative positions of the targets.

The helicopter autonomously flew over the field in a grid pattern, at approximately 15m altitude. The entire run consisted of 5 passes over the field with the helicopter flying very slowly (0.5 m/s) to maximize the chance of hitting each of the targets.

Figure 4-37 shows an overhead view of the results obtained from this flight. The dashed green line shows the actual flight path of the helicopter. The helicopter flew a regular grid pattern, starting at the top of the image, and flew five passes over the area with 15 m spacing between passes. The circles on the plot are the actual locations of the 31 retroreflective targets, based on the ground truth measurements.

As the helicopter flew, the terrain sensor scanned a plane in front of the helicopter. Each time a terrain measurement hit one of the targets, the location of the target was calculated by the terrain modeling system, and its location was marked with an ‘x’.

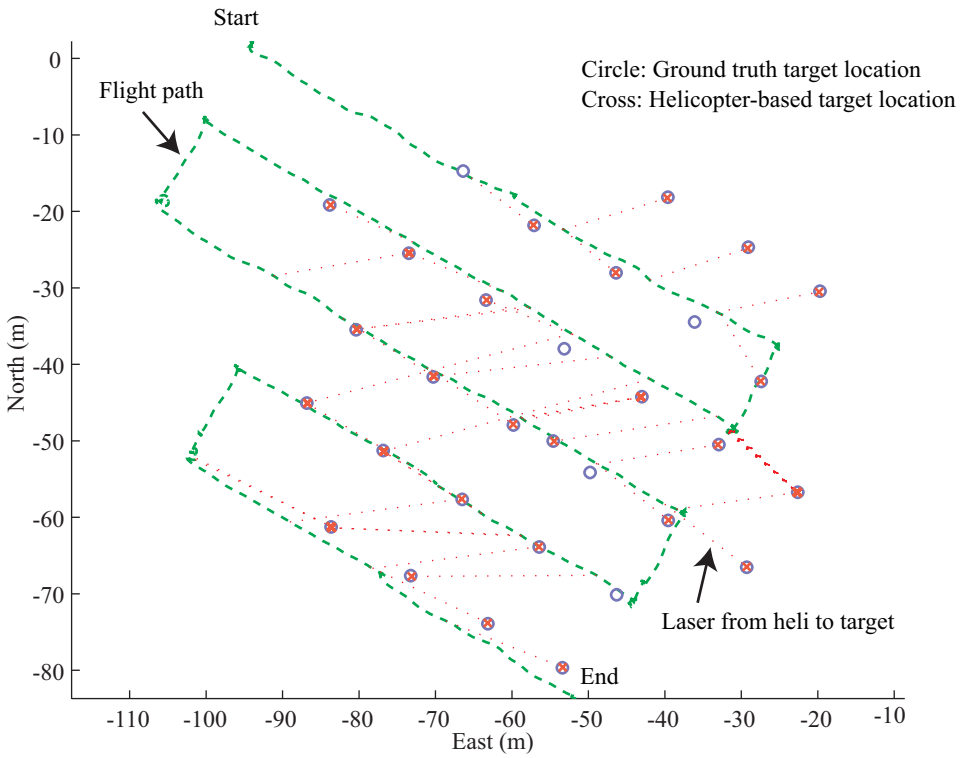

FIGURE 4-37. Accuracy verification flight

In this flight, 66 individual measurements hit targets. Some of the targets were not hit at all, while others had multiple hits because of the semi-random nature of hitting a target.

A dotted red line is also shown in the figure to indicate each measurement's optical path. One end starts at the helicopter, and the other goes to the computed position of the target. From these viewing rays, we can see that some of the targets were viewed multiple times from different locations.

To evaluate the accuracy of these results, we look at the difference between the ground-truth position of each target and the position determined by the terrain modeling system. In the figure, the computed locations of the targets (marked with 'x') coincides with the actual locations (marked with 'o'), within an average distance of 13 cm.

Figure 4-38 is a histogram showing the error distribution from this test. The error

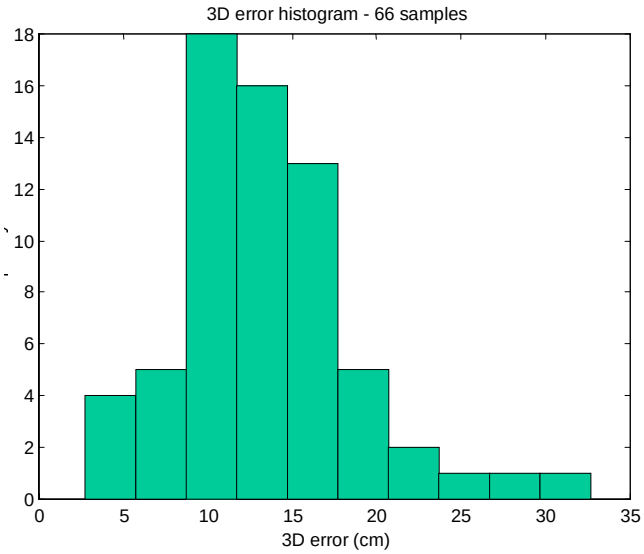

FIGURE 4-38. Histogram of 3D error, mean error = 13cm

metric is the actual 3D distance between computed and ground truth. The mean error is 13cm.

The ability of this method to determine the accuracy of the system is limited by the size of the laser beam spot as well as by the size and shape of the retroreflective target. Due to the strong return provided by the retroreflector, any portion of the laser spot hitting the retroreflective surface of the target reflects enough light to register as a target hit. In this test, the nominal spot size was 7 cm in diameter (range 15 m), and the target was 5 cm x 3.8 cm. This means that there is inherently up to 7 cm of uncertainty in the error measurements.

4.6.4 In flight color target test

We tested the ability of the 1-pixel camera to measure a variety of color targets during flight. The array of nine 1 meter square targets provide a set of dissimilar colors. Figure 4-39 shows the 3x3 color target placed on our test site.

We slowly flew the helicopter over the targets to scanned the area. Figure 4-40 shows the resulting color model of this area. In the model, each of the nine color targets are visible, and appear to have colors matching those from the image.

The black target is missing from the color model, since the black surface did not return enough laser light for the rangefinder to measure distance. Without the distance, the color cannot be properly positioned in the model.

FIGURE 4-39. Outdoor color target

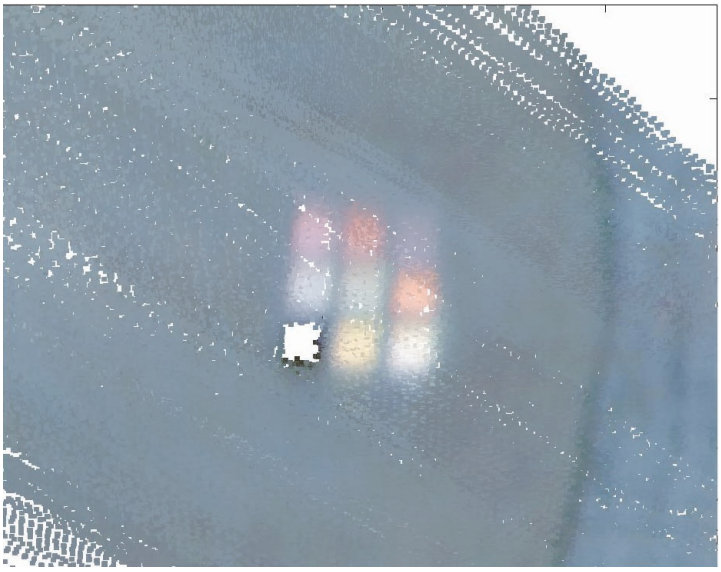

FIGURE 4-40. Color results off targets

Chapter 5

Terrain modeling system: Experimental Results

In this chapter, we present the real-world missions the terrain modeling system has carried out during the work on this thesis. Each mission was an opportunity to demonstrate the merits and to identify the shortcomings of the terrain modeling system at the time. Such missions validated the architecture, design, and implementation of the prototype terrain modeling system and explicitly show the contributions of this thesis.

The missions proved that a useful terrain modeling system must perform well in a variety of conditions. The system must handle changes in lighting, variations in terrain structure (buildings vs. landscape), and weather challenges.

We sought field test opportunities to exercise the system under different conditions while carrying out useful missions. We present these missions in this chapter. They constitute a substantial portion of the system's 50+ flights. These missions ranged from mapping the Haughton Crater in the Canadian Arctic, to mapping of the crash site of United Airline's Flight 93 following the September 11, 2001 terror attacks. The missions were performed throughout the period of this research, beginning in October

1997 through September 2001, and helped guide continual improvements to the system.

The following is a list of the missions:

Mission	In support of	Location	Date	Flight length
VSAM mapping	VSAM Project - DARPA	Bushy Run, PA site	October, 1997	2.2 km
Haughton Mars Project	NASA	Devon Island, NWT Canada	July, 1998	~5 km
Air support for ground testing	Perceptor - DARPA	Harmony, PA	May, 2001	~ 1 km
Tracked vehicle maneuvers area	Demo III - ARL	Ft. Indiantown Gap, PA	August, 2001	10 km
Mapped Surface coal mine	Perceptor - DARPA	Berlin, PA	September, 2001	1.4 km
Flight 93 crash site	NTSB, FBI	Shanksville, PA	September, 2001	2 km

5.1 VSAM mapping

Mission goals

In October 1997, we scanned a large site to construct a high resolution digital elevation map (DEM) for the DARPA Video Surveillance and Monitoring (VSAM) project at CMU [38]. The mission was the first field test of the terrain modeling system. The DEM, necessary for accurately modeling the variations in ground elevation throughout the site, or as an integral part of VSAM’s human/vehicle tracking algorithm. The VSAM system relied on an accurate elevation map; therefore, the terrain sensor was configured to best generate the DEM of the area.

Mission characteristics

The Bushy Run site, which was unoccupied during this mission, was previously a remote research facility of CMU. The site is approximately 300m x 300m, and has an asphalt road surrounding an open field that contains two large buildings; a number of trees surround the area. Our test pilot manually flew the helicopter approximately 10 meters above the ground as the system scanned the surrounding environment. All of the open areas were mapped by systematically flying the helicopter overhead.

Terrain modeling system status

The terrain modeling system used to map Bushy Run was an early version of the system built onboard an earlier and less sophisticated version of CMU's autonomous helicopter [39].

The autonomous helicopter's state estimation system at that time was less capable than that of our current system. The system used a single frequency GPS receiver with a nominal accuracy of 20cm (as opposed to 2cm), and a set of directional gyroscopes for attitude measurement (as opposed to the more accurate / less drift-prone Litton IMU). Additionally, the state estimator was limited to producing only 20 estimates per second (as opposed to the current 100 Hz version).

The terrain sensor measured only geometry (no color information), but used the same Riegl rangefinder and planar spinning mirror as the current system. The scanner plane was configured to scan the area beneath the helicopter. On this helicopter, the scan plane intersected the two landing skids of the helicopter, resulting in an incomplete scanline across the ground. Two shadows where the skids blocked the laser were produced.

Results

The first result presented is from a mapping flight over the road. The helicopter was manually flown completely around the asphalt road twice: first in the clockwise direction, followed by a flight in the reverse clockwise direction to fill the shadow areas. The flight was approximately 5 minutes in duration, during which over 2.5 million measurements were recorded. Figure 5-1a shows a point cloud visualization of the data from this test flight. For display purposes, each point is shaded, based on a combination of elevation and the return intensity of the laser. As such, the road is readily visible, since it absorbs the laser more readily than does the surrounding grass. It should also be noted that the helicopter was flown directly above the road where the highest density of the data was collected. The measurements decreased in density toward the center of the field, as we had expected. We performed other passes to fill in this area for the final DEM. Figure 5-1b shows a smaller portion of the dataset as viewed from above.

(a)

Figure 5-2a shows an aerial photograph of the Bushy Run site, which was acquired from the USGS. The image is geo-registered, such that the latitude and longitude of the four corners are known, and each pixel is exactly 0.5m wide. This provides an independently established ground-truth with which we can evaluate the results of our scanning system. Since the terrain modeling system uses GPS-based state estimation, the scan data is registered with the geo-registered image.

(b)

A single DEM of the site was generated by combining data from all of the scans. In the DEM, the size of each cell was 0.5 m x 0.5 m, covering a total area of 256 m x 256 m. Each 3-D terrain point measured by the laser scanner was registered to form the DEM by simply allowing it to fall into the closest cell. For each cell of the DEM, three statistics were computed: the mean elevation of all points landing in that cell, the variance in the elevation of those points, and the total number of points. These three matrices constituted the resulting DEM used by VSAM.

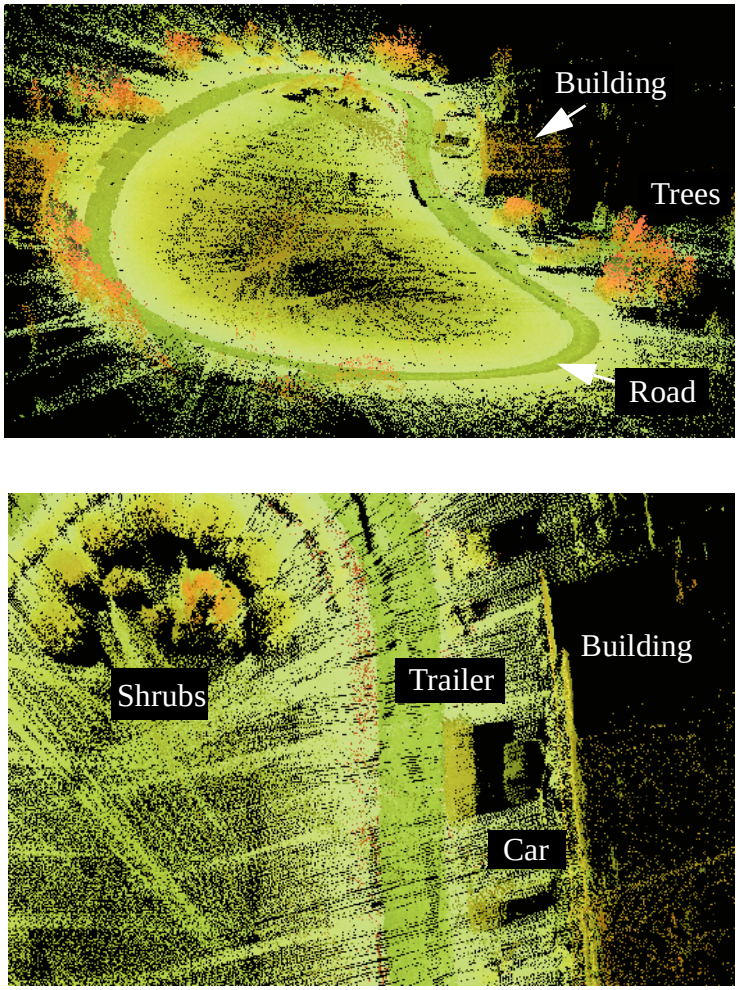

FIGURE 5-1. Perspective view of 3-D point cloud and overhead view of smaller area

Figure 5-2b shows the average elevation DEM, where the brighter pixels represent points with higher elevation. The DEM was intentionally generated such that the locations of the four corners are the same as in the aerial photo, with each pixel 0.5m wide. The intensity of each pixel indicates the average elevation measured for that location. A pixelwise comparison of these two images shows that the DEM closely matches the aerial photograph. This is an indication that the mapping system is producing accurate results, and that the global accuracy of the terrain model is at least 0.5 meters.

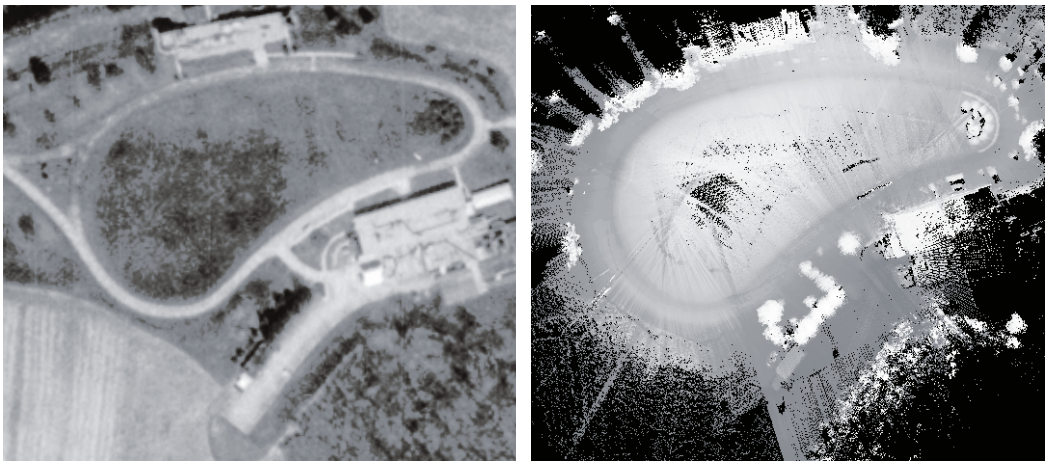

FIGURE 5-2. Aerial photograph of test site and corresponding DEM

This DEM was successfully incorporated into the VSAM system, and provided a more accurate and higher resolution terrain model (USGS elevation map with 30m spacing.) than was previously available.

5.2 Haughton Mars Project

Mission Goals

The primary purpose of our involvement in the Haughton Mars Project (HMP) was to assist NASA scientists in their geological study of Devon Island's Haughton Impact Crater, a Mars-analog environment [36]. The study of such environments is important to scientists in preparation for future manned missions to Mars. Of particular interest was an accurate 3D map of the crater because available maps did not represent the desired level of detail for NASA geologists. We were assigned a mission to investigate the feasibility of applying our autonomous helicopter technology to building accurate aerial maps of the crater.

This mission provided a unique opportunity to study a potential application of our unmanned aircraft technology; it was unique for several reasons. First, our aircraft could fly over uninhabited areas without concern for endangering human lives and property. Second, our system could perform a real-world mission to build highly accurate topological maps. Third, our helicopter could fly in a very different environment compared with its limited test flight area near Carnegie Mellon. This was especially valuable because the Arctic region's unique navigational challenges could test the robustness of our flight controls. Finally, our system's ease of transport and deployment could demonstrate the advantages and applicability of a small unmanned aerial aircraft for mapping tasks in inaccessible locations.

Terrain modeling system status

The HMP mission was the first to use the current autonomous helicopter system and state estimator, as described in Chapter 3.

This mission performed only geometry measurements with the laser scanner (no color data collection). The scan plane was positioned to scan across-track, beneath the helicopter, as was done in the Bushy Run mission. We removed part of the helicopter landing skids to prevent their occlusion of the laser scan.

For aerial imagery, we added an onboard CCD color video camera. The video was timecoded on-board, and relayed to a VCR on the ground via a wireless video link. The timecode inserted in the video allowed the helicopter state to be determined at the precise time each video image was captured.

The terrain modeling data was relayed to a ground PC and stored for later processing.

Mission characteristics

Logistically, flying over an island in the Arctic is challenging. The area is remote with no infrastructure. Our helicopters, personnel, fuel, and support equipment were successfully transported by a small Twin Otter airplane to the site. Figure 5-3 shows

FIGURE 5-3. Transporting the terrain modeling system

the terrain modeling helicopter (and a backup) onboard the plane. This experience reinforces how versatile the small helicopter can be, as opposed to a larger system. Once on the island, we camped with the helicopter equipment for two weeks during the mission. The second logistical challenge was to perform the actual mapping. As the helicopter flew, terrain data was downloaded to a computer on the ground. Since our mapping flights took us to different areas across the island, the entire system needed to be easily portable. We prepared the ground base station, a laptop computer and wireless communications on an all-terrain-vehicle (ATV). This allowed the ground equipment to be easily repositioned.

Two characteristics of the island made accurate state estimation a challenge. First, Devon Island is very close (300 miles) to the magnetic North pole, thus the onboard compass is rendered useless for heading measurements. To help deal with this problem, we utilized a special dual-antenna GPS receiver which determined the

heading of the baseline between the two antennas [37]. The GPS environment on the island itself is problematic. The U.S. GPS satellites do not orbit over the Earth's poles, so the satellites never reach a high elevation angle in the sky. This satellite geometry limits the GPS system accuracy, and at times makes it difficult for the GPS antenna to acquire the satellites. An additional problem is the surrounding hills which sometimes create significant multi-paths for the GPS satellite signals. We carefully monitored the GPS measurement quality and avoided using the data when a problem was expected to occur. This might lead to holes in the terrain model, but good measurements were acquired for the bulk of the mapping runs.

Results

Despite the environmental challenges, the helicopter's laser mapping system proved operational [40]. The system built accurate 3-D area maps by collecting time synchronized (within 1 μsec) laser, video, and vehicle state information at each site and merging this information after each flight. Following the data collection flights, the raw data logs were processed to generate 3-D models of the geographic features of the scanned crater portions.

Figure 5-4 illustrates the steps taken to process the raw data collected by the mapping system to generate a 3-D model of a set of cliffs in the crater. Figure 5-4a is a photograph of the cliffs. We show mapping of an area approximately 300 meters across and 50 meters high. Figure 5-4b is a point cloud visualization of the data collected during this flight, consisting of over one million 3-D data points sampled from the terrain in a single pass. Figure 5-4c is a triangular surface mesh generated from the point cloud by gridding and tessellating a manually selected viewpoint. The number of points was reduced (~40,000) to generate the surface connectivity. The missing data along the base of the cliff is the Haughton River. The data is missing because of the water's specular properties which do not return a measurable amount of the scanner's laser beam. The surface mesh is the type of 3-D model needed for numerous terrain calculations, and also for generating realistic terrain features in

virtual environments. Figure 5-4d is an example of a virtual world created by manually combining the surface mesh with various textures and rendering effects to aid in visualizing the details of the captured surface geometry.

During the flights at the Haughton Crater, we also collected color video from a fixed onboard camera (an image from the video is shown in Figure 5-5). The video signal was relayed to a base station on the ground via a wireless video link and recorded to videotape. Each image (field) of the video was tagged with a unique image number. The state-synchronized onboard VITC timecode generator digitally inserted the timecode into the actual video stream; that timecode was recorded on the video tape along with the images. The significance of that data collection is that the actual position and attitude of the helicopter was captured with every single field, thereby

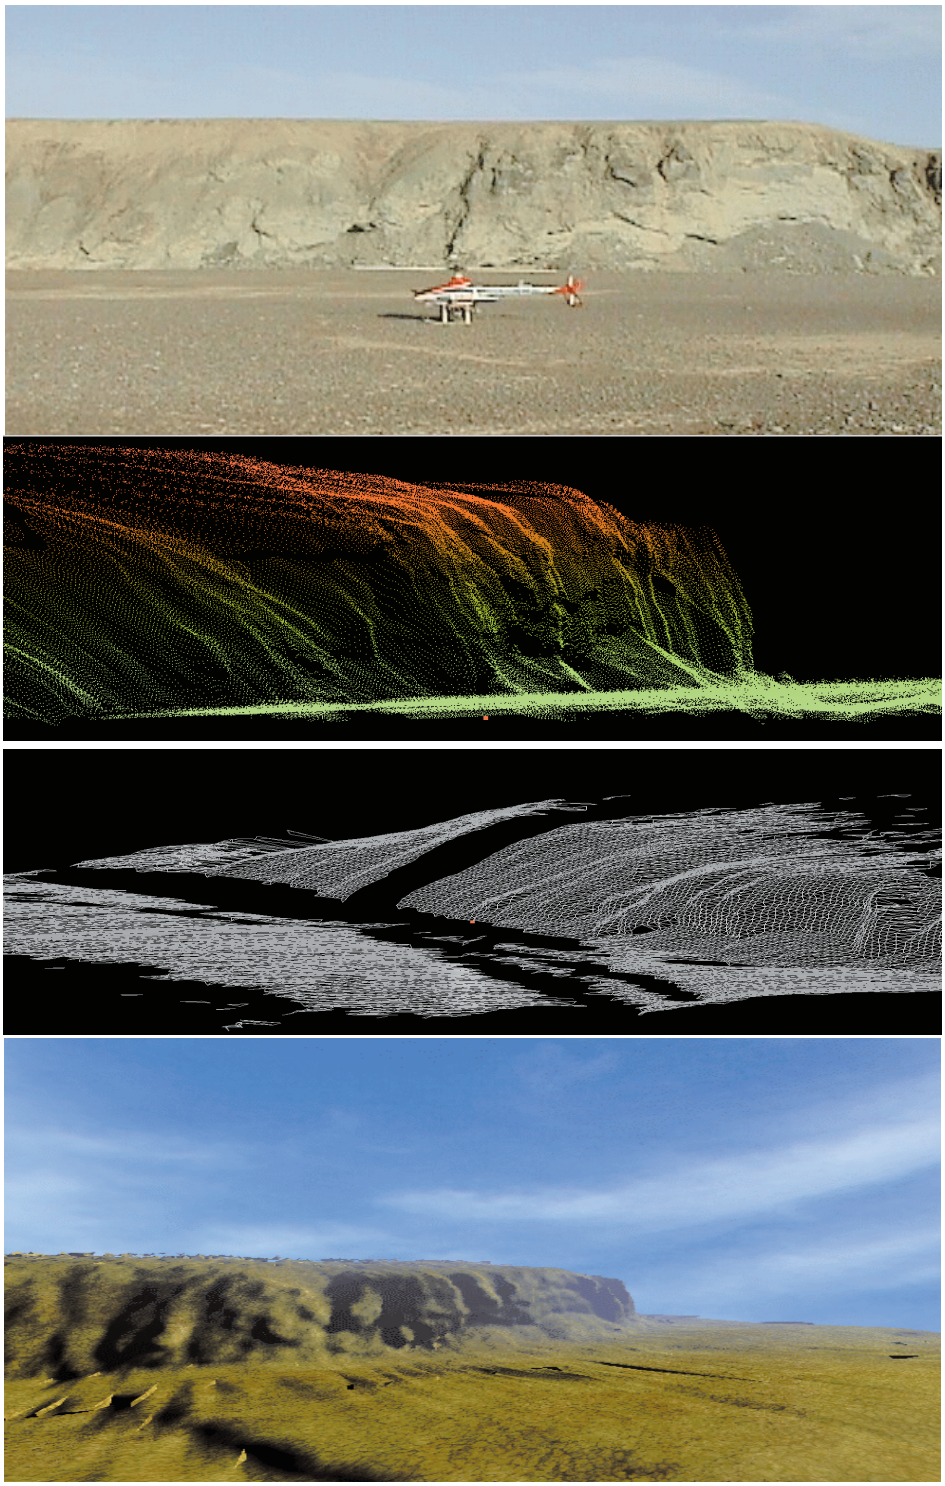

FIGURE 5-4. Mapping results of cliff

providing an accurate estimate of camera viewpoint for each image:

FIGURE 5-5. Video image captured during helicopter flight

Certain sensor and logistical problems prevented us from our goal of modeling the entire crater. However, the experience gained proved that the system can model large areas, and that, with further refinement and ruggedizing, our system could achieve this goal.

5.3 Map building for ground robot localization

In May 2001, our helicopter system was flown at our test field in Harmony, PA to provide aerial terrain mapping support for the ground-based Perceptor at CMU.

Mission Goals

The purpose of the experiment was to demonstrate 3D map matching between high resolution ground data and low resolution aerial data. This capability is useful for robot localization within the aerial map. Our helicopter-based terrain modeling system

was to first model the entire area. Then, a ground-based laser scanner was to scan a small region with high accuracy. The experiment involved matching the 3D surface geometry measured by the ground scanner with the helicopter-based terrain model. Since the aerial model was measured in the world coordinate frame, the map matching system could potentially determine the world coordinate of the ground-based scanner, a critical localization capability for ground-based robotics systems.

The primary goal of the terrain modeling system in this mission was to accurately model the site. The ground systems relied upon a complete terrain model; therefore, the terrain modeling system flew in patterns which would eliminate model holes caused by range shadows.

Terrain modeling system status

The complete terrain modeling system presented in this thesis was deployed during this mission. The laser rangefinder and 1-pixel camera were both used, although the supported research used only the geometry information. The scan plane was mounted in a forward-looking configuration, with the plane oriented approximately 30º below horizontal. At this time, the forward-looking configuration was being explored to scan upcoming terrain during flight. For this mission, however, the downward looking scanner was a better choice, as it provided a better angle of incidence for the laser.

Mission characteristics

The area of interest was a gently sloping grass field that spanned an area of approximately 200 m x 100 m. A gravel road passed through the center of the field, and a number of 80 cm high gravel piles (shown in Figure 5-6) were placed at various locations to provide a controlled terrain shape for map matching.

The helicopter was flown at an altitude of approximately 30m and performed five parallel passes over the site; each pass was spaced 12 m apart. The scanner measured a

FIGURE 5-6. Gravel piles on grass field

swath 75m wide across-track, which provided substantial overlap between passes. This was intentionally done to ensure that range shadows from one pass would be filled in the subsequent pass. The helicopter's flight control system was configured to fly (0.5 m/s) to acquire a high density of measurements.

Mapping results

We now examine the built terrain model of the test site. Figure 5-7 shows a portion of the collected colored point cloud. The point cloud is the result of three parallel passes over the area. A number of interesting features can be seen in this figure. The dirt road crossing the field can be readily seen, and grass can be seen growing between the two tire tracks in the center of the road. The five piles of gravel appearing in Figure 5-6 can also be seen. The other feature intentionally prepared on the site is a path of tall, uncut grass. The appearance of the grass in this region is clearly different from other areas of grass; however, the difference is subtle.

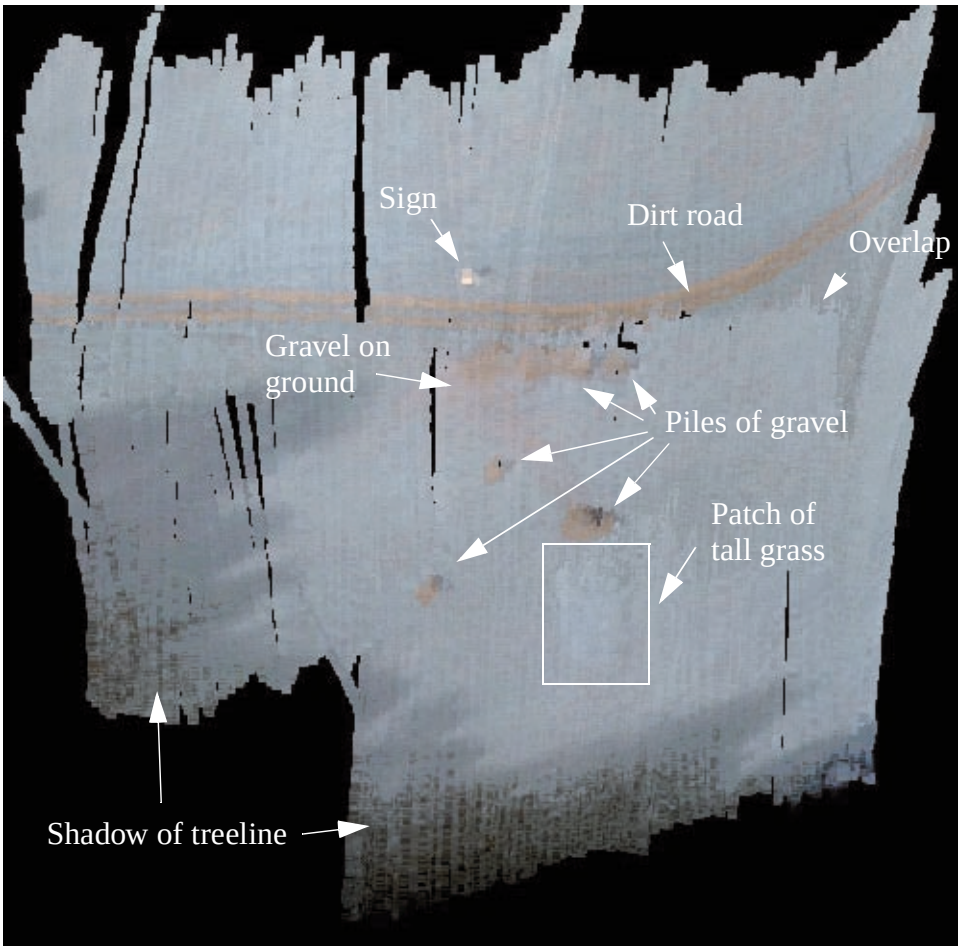

FIGURE 5-7. Color point cloud of gravel piles

Two interesting effects are worth noting. First, the darker areas of grass on the bottom and left of the image are caused by shadows cast by the treeline surrounding the field. The areas of the ground in the shadow are illuminated less strongly than other areas, so they appear darker. The other effect has to do with the area near the two piles of gravel along the road, which appears to have similar appearance to the piles. This area is where the gravel was located prior to moving it into the separate piles. As a result, there is residual gravel spread over the grass in this area. This generates a color that combines the gravel with grass.

Finally, an artifact from the overlap of two passes can be seen where indicated in the figure. The second pass (the lower one) shows the grass appearing to be brighter than that of the first pass. This difference in appearance was caused by a change in the overall intensity of sunlight during the time between passes over the area.

Experimental results

After the mapping flight, we delivered the colored point cloud terrain model to the Perceptor group. After discarding the color information, they built a digital elevation model of the area from our data. Figure 5-8 shows the resulting DEM created from the

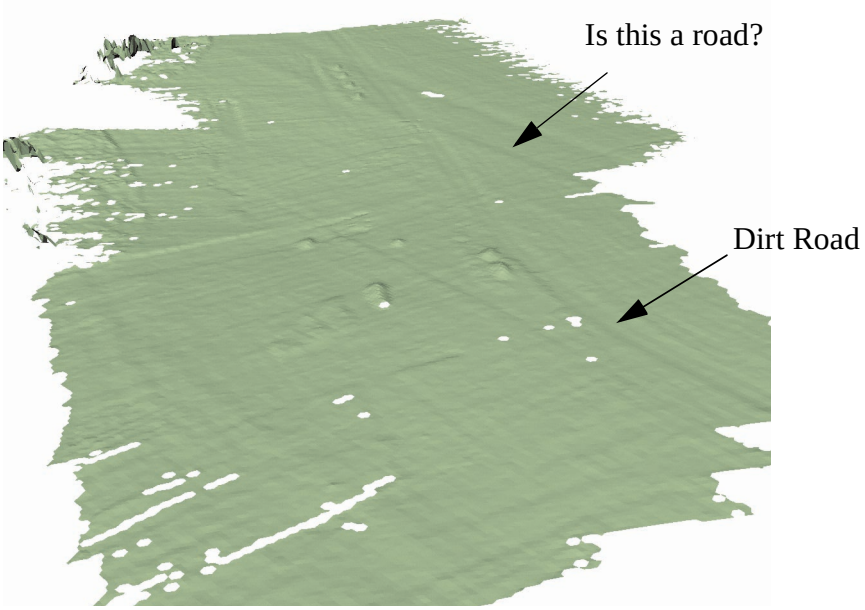

FIGURE 5-8. DEM of site

helicopter terrain measurements.

It is worth noting that, without the color attributes in the model, it is quite difficult to identify some of the features at the site. For example, the dirt road is actually visible in the DEM, since there is some depression in the ground from vehicles, and in the grass

growing alongside the road. However, it would be difficult to recognize this area as a road from the geometry alone. A diagonal depression similar to a road runs across the field near the marked area on the top of the image. This is not a road. Instead, it is a surface depression in the actual field completely covered by grass. It would be difficult to differentiate between the two from this data, but the addition of color substantially simplifies the road and non-road identification.

Figure 5-9 shows a result of the automatic matching of the helicopter built DEM (red)

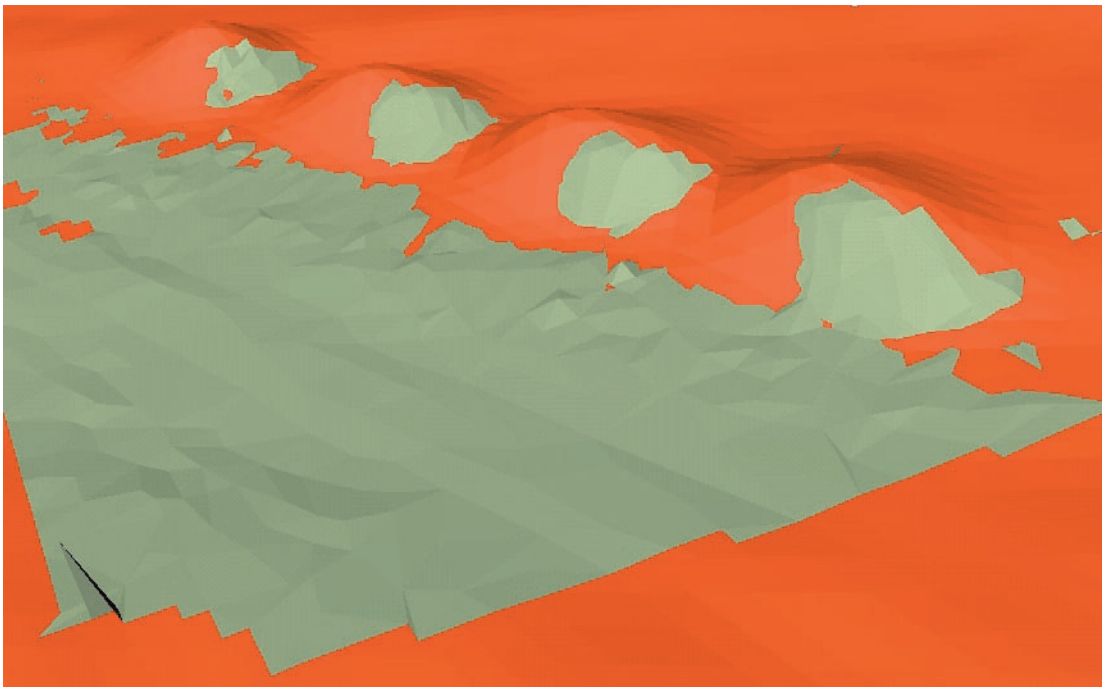

FIGURE 5-9. 3D matching results between aerial DEM and ground scanner

with a scan from a ground-based laser scanner (green). The shapes of these 4 hills match closely, showing virtually the same terrain model geometry as the one derived independently from the ground-based sensor data. The registration could globally localize the ground laser scanner to better than 50cm accuracy.

5.4 Tracked Vehicle Maneuvers Area

Mission Goals

The goal of this mission was to build terrain models along portions of the Tracked Vehicle Maneuvers Area (TVMA) testing site at Ft. Indiantown Gap, a National Guard military base in Pennsylvania. The mission supported autonomous ground robot research. The purpose of modeling these areas was to demonstrate the ability of the terrain modeling system to quickly model an area to aid in planning and navigation of a group of ground mobile robots entering the same area.

Terrain modeling system status

The complete terrain modeling system presented in this thesis was deployed during the TVMA mission. The laser rangefinder and 1-pixel camera were both used for this mission. The scanner was mounted in a downward-looking configuration to measure the terrain across-track, beneath the helicopter. The terrain and state measurements were transmitted to a computer on the ground via a wireless link. During flight, the terrain model was assembled and displayed in real-time to provide immediate feedback of progress.

Prior to arriving at the TVMA site, the scanner mount was not calibrated. To determine the exact scanner mount, we used the in-flight calibration technique, described in Chapter 4, during our first flight.

Mission characteristics

The TVMA site is an ideal location for autonomous helicopter terrain modeling experiments. It is a safe area for flying, since it is a large, open area which is isolated and uninhabited. It also provides a rich terrain environment for modeling. Dirt roads,

small vegetation, larger shrubs, and large trees are found throughout the site. Pronounced treelines, where the woods stop and open fields begin, are prevalent and provide a good example of situations which require modeling of vertical and overhanging surfaces in the terrain. Figure 5-10 is a typical picture of the TVMA range.

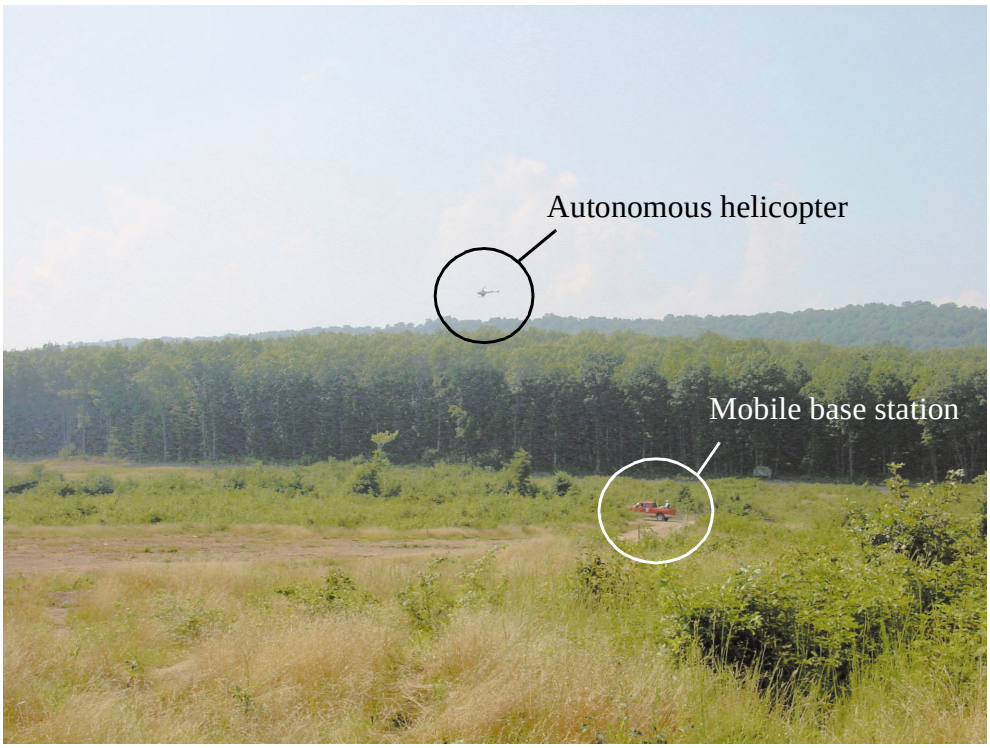

FIGURE 5-10. Picture of TVMA area

During the TVMA mission, the system had to cope with strong wind conditions. The autonomous flight capabilities of the system successfully compensated for the wind and flew the helicopter in uniform patterns over the areas of interest.

Because we were mapping over a large distance, we needed to continually move the base station to keep it close to the helicopter (<300m) since the wireless ethernet radio used to downlink the data has a relatively short operating range. To allow continuous

modeling, we mounted all of the base station equipment in a chase vehicle which followed the helicopter to stay close enough during flight.

To achieve consistent terrain coverage while mapping, an augmented flight mode was developed. In this mode, the control system flew the helicopter forward at a constant speed and altitude. During flight, an operator in the chase vehicle could steer the helicopter's direction with a joystick. The operator was therefore able to monitor the real-time terrain model, and guide the mapping flight path to achieve the desired coverage, while the control system ensured a smooth, consistent flight.

During mapping, the helicopter moved forward at a constant speed of 3 m/s, and the altitude was high enough to ensure that the helicopter was above the trees. Typically, this placed the helicopter approximately 50 meters above the ground.

Results

Figure 5-11 is a screen dump of our real-time display. As the helicopter was mapping the terrain, a color enhanced DEM was built and displayed in the large box on the lower right portion of the screen. The figure shows a composite DEM which combines nine separate mapping runs performed over the 3 day mission. The displayed map area is 2 km x 2 km. Since the primary objective of the mission was to map the areas along the dirt roads in the TVMA range, the composite DEM contains only data near the roads, as opposed to a more complete coverage of the entire area. The measured terrain color is displayed. From the color data, the differences between the brown open dirt areas and the green vegetation areas are readily apparent. Also apparent is some variation in intensity of color at different places. This variation was caused by the sunlight changing over the times the terrain was modeled, resulting in dark and light bands across the area.

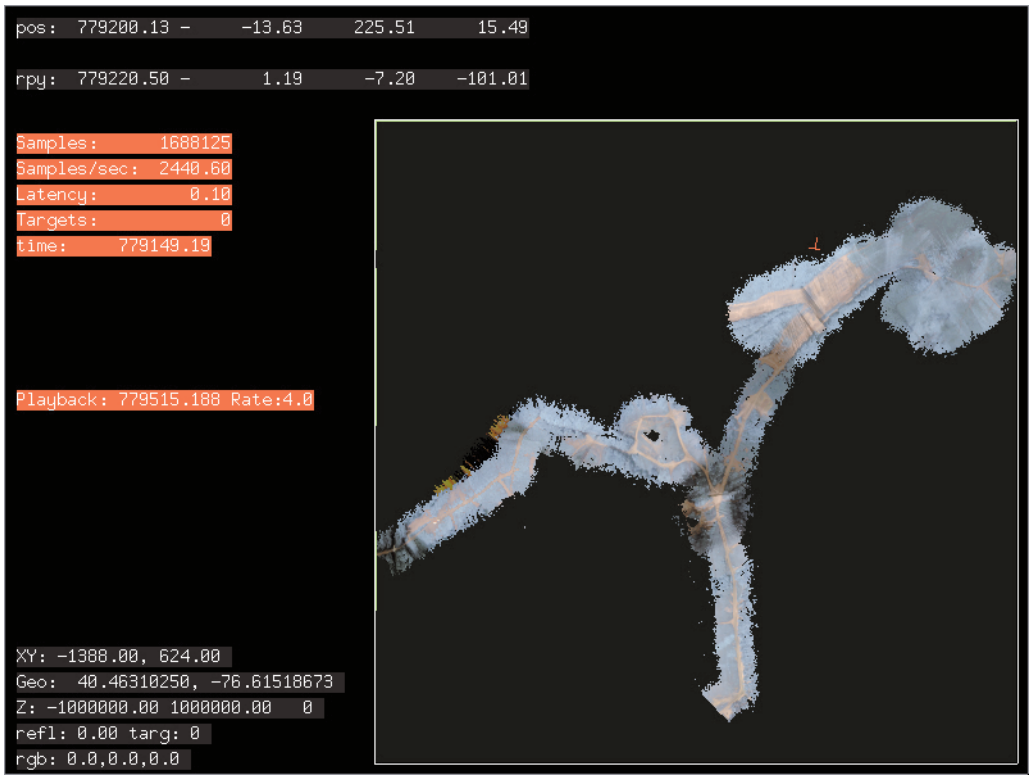

FIGURE 5-11. Composite overview DEM - color display, 4m resolution

Figure 5-12 shows the measured elevation of the same composite DEM. Dark areas are lower, and brighter areas are higher. In this view, the treelines are clearly visible, as the trees are tall (20-30 meters).

Showing the entire composite map requires zooming out to make each pixel 4m across, which does not show most of the terrain measurements collected. Figure 5-13 and Figure 5-14 show a zoomed in view of the DEM and terrain color in the area marked “A” in Figure 5-12.

In this area, the dirt road enters a patch of woods where the tops of the trees on either side of the road are almost enclosing the road. The helicopter flew directly over the road at approximately 50 meters altitude, clearing the tree tops by at least 20 meters.

FIGURE 5-12. Composite overview DEM - height display

Figure 5-15 shows another segment of the terrain model. In this segment, the helicopter flew lower than the treeline along the upper edge of the scan. Therefore, terrain measurements were acquired only on the nearly vertical surface of the treeline facing the helicopter. as opposed to the measurements from the top of the tree canopy as in Figure 5-13.

The real-time DEM is a 2.5D representation of the terrain. Only a single elevation and color measurement can be stored for each DEM cell. Since each measurement taken by the terrain sensor identifies a unique 3D point on the terrain, the DEM does not fully illustrate the acquired terrain measurements. To see all of the measurements, we must view a 3D point cloud. We can view a 3D colored model by displaying the point cloud with the measured 1-pixel color. To achieve this, we generate VRML colored point cloud files.

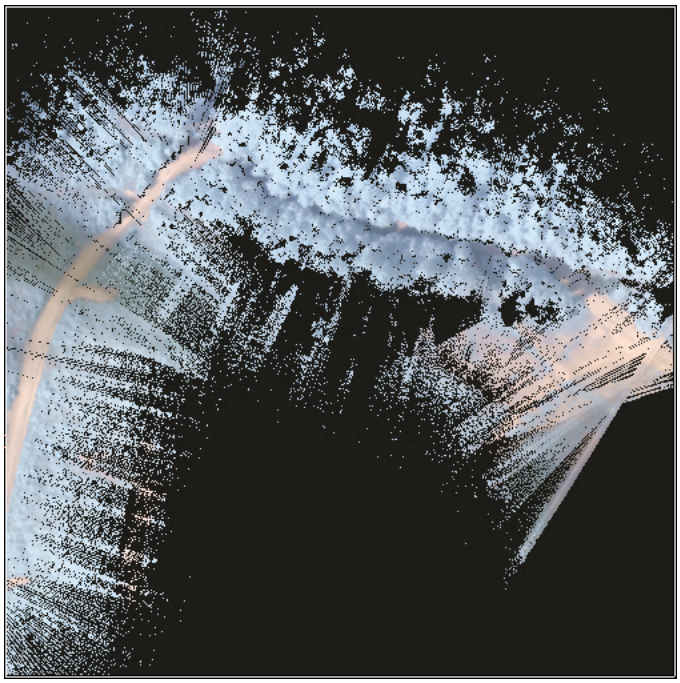

FIGURE 5-13. Color view of tall canopy over dirt road. 25cm resolution.

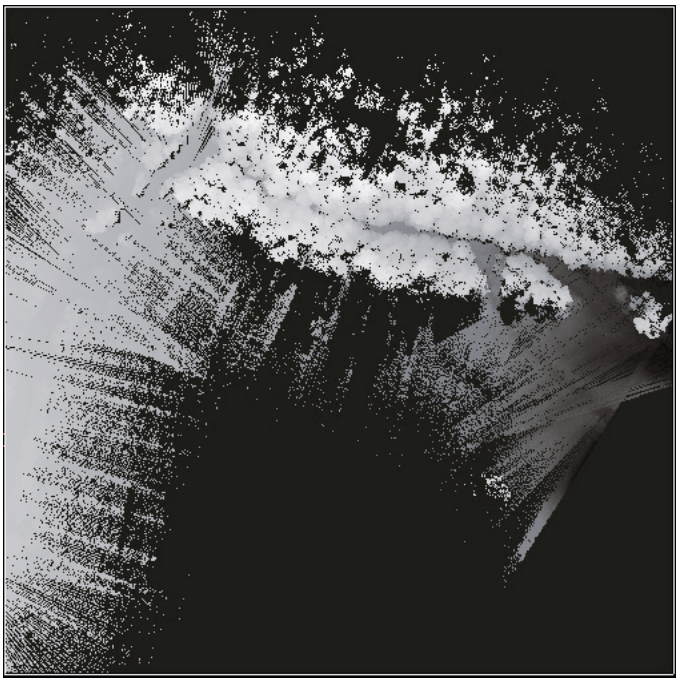

FIGURE 5-14. Height view of tall canopy over dirt road.

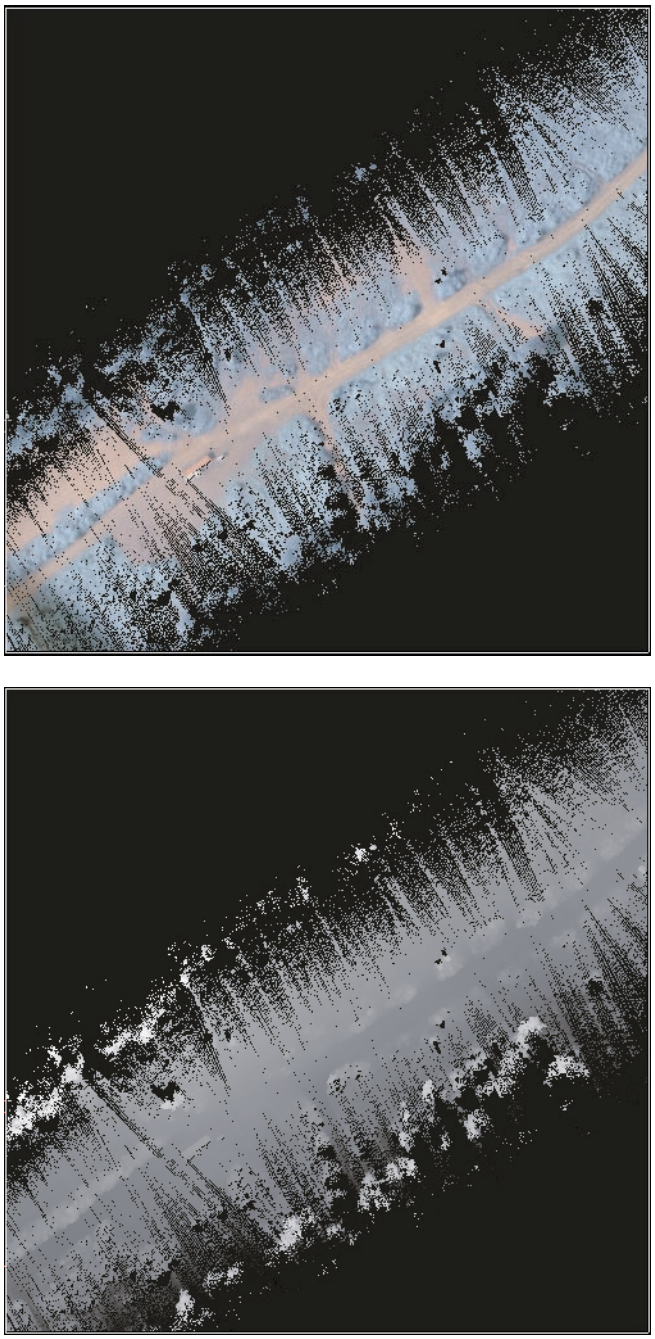

FIGURE 5-15. DEM of another road section

Figure 5-16 is a top view of a colored point cloud from another segment of the map. Two interesting points to note are: first, shadows can be seen on the road where it passes through the trees. The sunlight must be coming from the left side of this image

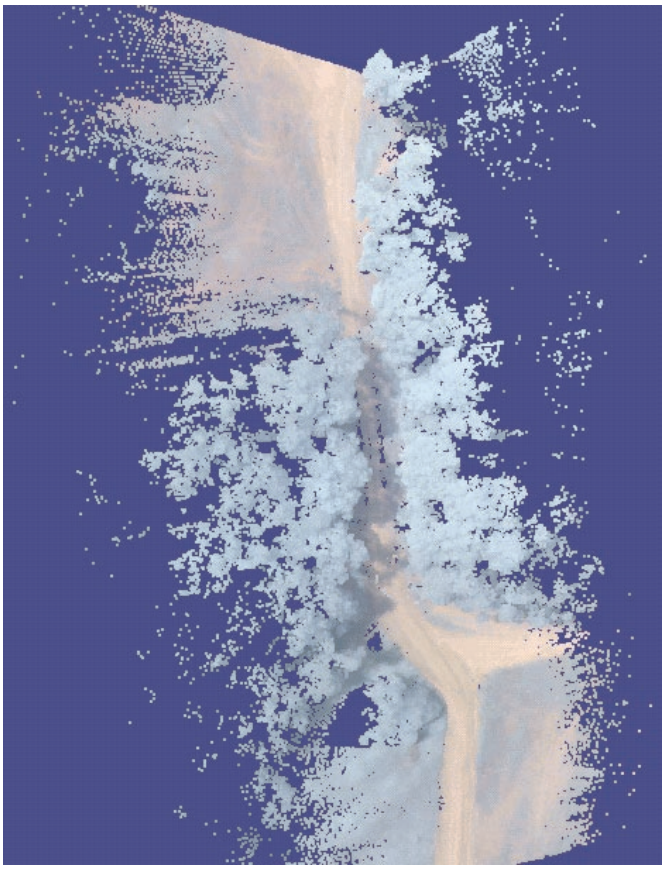

FIGURE 5-16. Top view colored point cloud

to cast these shadows from the trees; second, vehicle tracks in the ground are readily visible in the upper left portion of the image, illustrating some of the small features observable by the terrain modeling system.

Figure 5-17 is the same colored point cloud, viewed from the side. Here, the 3D shape of the trees are readily apparent. Also, we can look down the road through an opening in the trees.

Figure 5-18 shows a side view of another segment of the terrain map. Here, shrubs on either side of the road are clearly visible.

FIGURE 5-17. Side view of colored point cloud

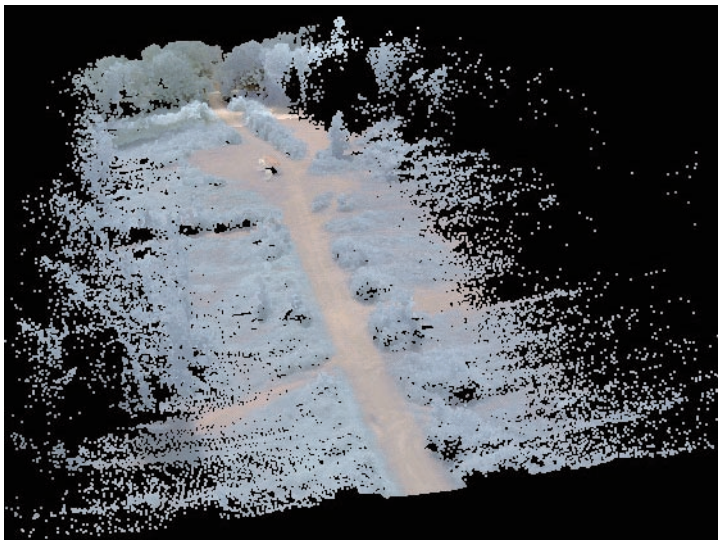

FIGURE 5-18. A second colored point cloud

In total, over 10km of mapping flight were performed during the TVMA mission, mapping over 500,000 square meters of terrain.

5.5 Surface coal mine

Mission Goals

The goal of this mission was to build a terrain model of a surface coal mine in Berlin, PA to support research by the Perceptor group at CMU. This mission was similar to a previous mission supporting their effort (described in Section 5.3), but was performed at a site with larger scale terrain features.

Terrain modeling system status

The complete terrain modeling system presented in this thesis was deployed during this mission. The laser rangefinder and 1-pixel camera were both used. To measure the terrain across-track, beneath the helicopter, the scan plane was mounted in a downward-looking configuration. Since the area being modeled was small enough, a fixed base station was setup.

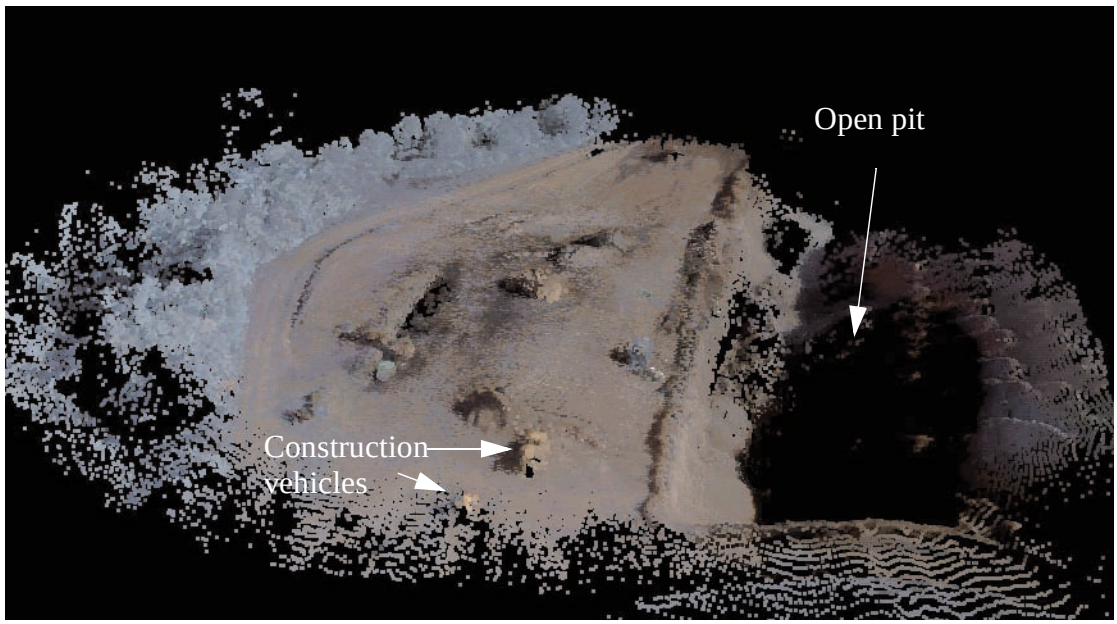

FIGURE 5-19. Surface coal mine colored point cloud

Mission characteristics

The surface coal mine site presented several unique challenges for our terrain modeling system. First, exposed coal was present throughout the site. The low reflectivity of the dark colored coal prevented return sufficient to exceed our minimum laser return threshold. This resulted in a lack of readings from some surfaces.

The second problem with the site was finding a takeoff/landing site for the helicopter. The mine is surrounded by trees, and virtually all of the exposed open areas on the site have a dirt ground. Takeoff and landing of our helicopter from a dirt ground is a problem, since a large amount of dirt is blown around by the helicopter downwash. We dealt with this by setting up a 12ft x 12ft piece of carpet for a takeoff / landing pad to minimize the dust.

Results

Similar analysis of the terrain models were performed using this data as to the analysis described in Section 5.3; comparing the helicopter scans to those of the Perceptor ground-based scans. A colored point cloud of the mine acquired by the helicopter is shown in Figure 5-19.

Chapter 6

Conclusions and future work

In this thesis, we have demonstrated that an accurate and capable terrain modeling system can be developed by carefully integrating a small autonomous helicopter with an onboard terrain modeling sensor. Numerous field tests have shown the value and maturity of the system for quick modeling of large areas in a way that is unobtainable with other approaches. This terrain modeling capability has significantly extended the usefulness of the CMU autonomous helicopters.

Additionally, in this thesis, we have introduced the concept of the 1-pixel camera and have shown the merit of combining active laser range measurement with the passive 1-pixel camera for building color 3D terrain models. The 1-pixel camera concept ensures properly registered color and geometry measurements in a simple, light weight design.

In this chapter, we summarize the accomplishments and present future directions for this research.

6.1 Accomplishments

The two primary accomplishments of this work include an autonomous helicopter based terrain modeling system and a novel color 3D terrain sensor. In this section, we summarize these accomplishments.

6.1.1 Autonomous helicopter based terrain modeling system

This research has contributed to the development a terrain modeling system based on a small autonomous helicopter system. The autonomous helicopter has proven to be an ideal platform for controlled and systematic terrain modeling. The system is accurate, localizing the 3D location of points with an average error of 15 cm.

Autonomous flight provides distinct advantages for terrain modeling. Preplanned flight patterns can be executed accurately, thereby ensuring complete coverage. Even though the scanner is strapped down to the helicopter, the precise autonomous flight of the system can provide uniform coverage of the terrain.

The system is readily portable, since the helicopter itself weighs approximately 160 lbs, and can be easily transported in a pickup truck. Fuel consumption on the order of 1 gallon of gasoline per hour, and relatively simple system maintenance make the system especially economical and practical.

The terrain modeling system has proven to be reliable in the field. Over 50 mapping flights have been flown with only minimal required maintenance. Aside from periodic cleaning of dust from the unenclosed scanner optics, no problems with the hardware have been encountered. Even in harsh environments such as Devon Island and Ft. Indiantown Gap, the system has performed well.

6.1.2 Color 3D terrain sensor

This research has produced a novel terrain modeling sensor design which combines a commercial laser rangefinder with a custom 1-pixel color camera to simultaneously measure the color and 3D position of points on the terrain. The ability to build accurate color terrain models has been demonstrated through a series of incremental verification experiments, and actual modeling missions.

During this research, we have become convinced that color measurement, in addition to geometry, is extremely valuable for understanding the environment. Even though the color measurements are sparse when compared with color images from a standard camera, they provide substantial cues when viewing the model over viewing the geometry alone.

The 1-pixel camera provides an ideal approach for providing this color capability to the terrain sensor. The camera can be added to the laser rangefinder either externally, as we did, or internally to the rangefinder unit. Our prototype is small, weighs less than 1 pound, and is inexpensive (<$500 in components). For these reasons, we believe that the 1-pixel camera concept should be incorporated into all commercial rangefinders, since the increased value of the data far exceeds the additional costs. In fact, a similar approach has recently been taken by Riegl in their LMS-Z210 3D imaging sensor with their True Color Channel [41]. Although this sensor is far too large and heavy for use onboard our helicopter, it is promising that, in the future, such systems may be commercially available.

6.2 Lessons learned

Through the work described in this thesis, we have learned a number of useful lessons. A few of them are:

6.2.1 Custom sensor design is time consuming

The 1-pixel camera initially appeared to be quite a simple device, one which could be built quickly. As development proceeded, however, numerous details emerged to complicate building the seemingly simple sensor. Systematic development, testing, and verification helped to ensure that the sensor was successful; however, this did require a substantial investment of time. To achieve the desired terrain modeling system, the custom design was worthwhile but, in general, developing custom sensors should be approached with caution.

6.2.2 Field recalibration is important

We found the need to quickly recalibrate the terrain sensor in the field to be more important than initially had been thought. Our initial approach to calibration was to rigidly attach the various components to the helicopter, calibrate it once in the lab, and then leave it alone. In reality, we found it necessary to remove or reposition the terrain sensor on the helicopter from time to time. In some cases, maintenance on other parts of the helicopter required their temporary removal from the aircraft. At other times, the mapping mission would require the scan plane to be repositioned. Also, the entire sensor was removed periodically when other experiments required using the helicopter. Each time, the sensor was remounted, the scanner mount had to be recalibrated.

Because our initial calibration approach was prohibitively time consuming, it required the development of rapid calibration techniques, such as the in-flight calibration discussed in Section 4.6.2. During the most recent modeling missions, we have been able to simply mount the sensor before going to the field, confident that it would be calibrated as part of the first flight of the day.

6.2.3 Sunlight-changes make consistent color challenging

The 1-pixel camera is a passive sensor, so it relies on sunlight to illuminate the terrain. Unfortunately, the illumination conditions created by sunlight are continually changing. Since the color balancing is generally calibrated prior to flight, the primary difficulty comes from changes in illumination during flight.

Days which are “partly cloudy” pose a significant problem for the constant illumination assumptions. As the clouds move across the sky, the lighting on the terrain is constantly changing.

6.2.4 Calibration is as important as design

One of our beliefs which has been reinforced by this work is that accurate calibration techniques are as important as good system design to system performance. Although some results from a system can typically be generated with poor or arbitrarily selected calibration parameters, carefully designed calibration techniques are essential to enable the system to reach its ultimate potential. For this reason, calibration techniques constitute a sizeable portion of this thesis.

6.3 Future work

Based on the developments described in this thesis, there are several directions for future work. In this section, we present several of the most directly related possibilities, which are either natural extensions of the thesis work, or would address specific problems encountered in our work.

6.3.1 Answer the question: “Why are the trees blue?”

One issue observed during this thesis was that the colors produced by the sensor were not as photo-realistic as might be desired in some cases. Direct comparisons between the color terrain model and a color photograph or digital image of a scene show that the colors from the terrain modeling system appear to be less saturated and with a different color hue. For example, most of the color modeling results presented in this thesis have included trees and other vegetation in the scene. These portions of the color models have consistently appeared “too blue” to humans looking at the scene. The goal of this future work is to determine how to make the 1-pixel camera produce more photo-realistic color measurements.

Photo-realistic color is a property of an imaging sensor where the color results appear to a human observer as if they are directly viewing the imaged scene. The raw output of many imaging sensors is not photo-realistic, and special color processing is required

to achieve it. The non photo-realistic color measurements currently generated by the 1-pixel sensor are useful for many of the terrain modeling applications we require. Terrain segmentation based on color differences can be reliably performed, even in a non-photo-realistic color space. Similarly, terrain classification based on color histograms and rudimentary color terrain visualizations do not require photo-realistic color. On the other hand, achieving photo-realistic color measurements is necessary for other potential applications, such as virtual reality modeling and color comparison with other color imagers. Therefore, understanding how to generate photo-realistic color measurements with the 1-pixel camera is an important piece of our future work.

In general, if the spectral sensitivity and color processing of the 1-pixel camera were to match that of a photo-realistic imaging systems (quality color camera), it would produce photo-realistic colors. Since the current sensor does not fully achieve this level, the first step is to identify where our current sensor deviates from this imaging system. We believe that one (or more) of the following statements may be true, accounting for this difference:

1. The color calibration procedure is not completely calibrating all of the essential color parameters, and the uncalibrated parameters are essential.
2. The actual spectral sensitivity of the 1-pixel camera is too far away from the human eye's spectral response, implying that no photo-realistic mapping from raw measurements to the displayed red, green, blue values exist.
3. The color processing algorithm used for each measurement is not sufficient to provide a photo-realistic mapping from raw measurements to displayed red, green, blue values.

Our future efforts to improve the photo-realism of the color will initially concentrate on these three issues. A brief explanation of the issues and our approach to investigate each of them follows.

Incomplete calibration

We feel that the most likely cause for the non-photo-realistic color is that the color calibration procedure presented in Section 4.2.4 does not determine all of the available calibration parameters. Specifically, the 3x3 colorspace conversion matrix described in the color processing path, shown in Figure 4-8, is assumed to be identity. This assumption is based on the premise that the color sensing channels are decoupled from each other. In other words, the blue output value is completely determined by the blue channel measurement, and similarly for the other channels.

The colorspace conversion matrix is provided in the color processing algorithm to allow a mixing to be performed between the various channels. When a non-identity matrix is used, the blue output value can depend on the red, green, and blue channel measurements. This color mixing is often performed in digital camera color processing algorithms [47] to compensate for spectral sensitivity mismatches between the sensor and the destination colorspace.

We had initially hoped that the spectral sensitivity achieved for our sensor would be close enough to the display colorspace that the identity colorspace conversion matrix we used would produce reasonably photo-realistic color results. The calibration procedure follows this approach by using a grayscale calibration target, shown in Figure 4-9. The grayscale target provides adequate information to properly balance the three channels relative to each other, ensuring that neutrally colored objects (equal red, green, and blue components) appear properly. However, no calibration information is available to determine if a proper mixture between channels is needed. Based on our decoupled channel assumption, non-neutrally colored objects should appear properly as well, since no color mixing is needed.

It may be the case that the spectral response of the sensor is not close enough to the display colorspace for this assumption to be valid. Therefore, we will need to enhance

the color calibration procedure in order to fully calibrate the colorspace conversion matrix, and achieve photo-realistic colors.

One possible method for doing this is to expand the calibration target to contain additional colors which span the entire color space, and provide information to determine the needed color mixture. The MacBeth ColorChecker [34] is a calibrated target routinely used for color calibration of video cameras. It consists of 24 2-inch squares of calibrated color material which represents typical colors observed in natural scenes. Images of this target can be used to calibrate the entire colorspace conversion matrix by using the camera's raw measurements and the calibrated target specifications. Unfortunately, the ColorChecker target is too small to be used with the 1-pixel camera. However, the calibrated material used in the ColorChecker is available in 8.5" x 11" sheets which are large enough to image with the 1-pixel camera. Therefore, a large target with these 24 color sheets could be constructed and used in a calibration procedure similar to the one we've presented. This would allow the complete set of color processing parameters to be calibrated.

Incompatible spectral sensitivity

Another possible cause of the non-photo-realistic color could be that the actual spectral sensitivity of the sensor is drastically different from the display colorspace, and there exists no mapping between the two which provides photo-realistic colors from the 1-pixel camera.

At present, this does not seem to be the case. First, the sensor is providing color measurements which are unique for the different colored terrain areas. This suggests that the sensor is detecting the different colors properly, even if they may not be photo-realistically mapping into the display colorspace. Second, the estimated spectral sensitivity of the 1-pixel camera appears to be reasonable. This supports the argument that a reasonable mapping should exist.

In Appendix A, the spectral sensitivity of the 1-pixel camera is calculated based on the specifications for the individual components used in each channel. When compared with the published spectral sensitivities for a number of commercial video cameras, a wide variation among the video cameras are apparent, and the 1-pixel camera does not appear to have a significantly different sensitivity compared with the video cameras. Therefore, we would not expect the 1-pixel camera to be any less capable of producing photo-realistic colors than the video cameras.

The results in Appendix A are calculated from datasheets, and it is possible that the actual sensor's spectral sensitivity is significantly different for some unknown reason. One method for investigating this possibility is to directly measure the actual spectral sensitivity of the 1-pixel camera. Unfortunately, this is a difficult test to perform without an elaborate and expensive test setup. Less involved methods for estimating the spectral sensitivity of video cameras have been suggested by numerous researchers [48][49]. We may be able to adapt one of these methods to the 1-pixel camera, and experimentally verify the spectral sensitivity.

Improper color processing algorithm

Finally, we should consider the possibility that the color processing algorithm we are using is not complete for photo-realistic color. The processing which we currently perform (offset, colorspace conversion, gain, and gamma correction) is derived from the operations performed in video cameras. This approach is justified, since the raw measurements from the 1-pixel camera are a linear measure of the optical energy captured, and is comparable to the raw output from each pixel of a CCD or CMOS imager. Therefore, it is reasonable to assume that the same processing steps should apply.

However, the specific proprietary operations used inside of many commercial video cameras are not available to us. The manufacturers may use additional tricks or

processing stages internally to achieve photo-realistic color images. If this is the case, we may need to perform additional research to understand these techniques and apply them to the 1-pixel camera color processing.

6.3.2 Alternative sensing -- polarization

In this thesis, we have concentrated on measuring terrain color using a 1-pixel camera. One of the strengths of the 1-pixel camera concept is that it can be readily adapted to use alternative sensing techniques. This could include increasing the number of sensing channels to measure additional terrain attributes (multi-spectral), changing the spectral sensitivity of the individual channels to measure specific bands of interest (near IR), or using specialized optical components to measure alternative properties of the incoming light (polarization).

To demonstrate this capability, we modified our 1-pixel camera to measure the polarization of light reflected from the terrain. Polarization is a property of all electromagnetic waves; it describes the orientation of the electric field over time. Linearly polarized light constrains the electric field to a fixed plane. Unpolarized light, on the other hand, allows the electric field to vary. It is well known that certain objects, such as metal, polarize reflected light. Others reflect unpolarized light, such as typical diffuse surfaces. Measuring the ratio of linear polarization received from the terrain can therefore provide a useful criteria for identifying the terrain composition [42].

To modify the sensor, we built a second filter holder for the 1-pixel camera; to hold three polarization filters instead of the color filters. The two filter holders are shown in Figure 6-1. A polarization filter allows only the portion of light polarized along a specific axis to pass. The sensitive axes for the three channels are 45º apart. This allows the measurement of the ratio of energy received along the different axes.

With the polarization filters in place, an outdoor scene was scanned from the ground, similar to the color scans in Section 4.5. Based on the three polarization channels, an

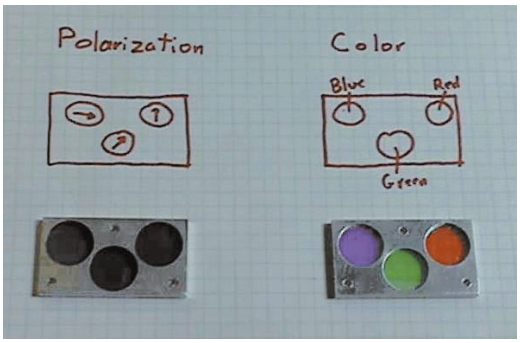

(a) (b)

FIGURE 6-1. (a) Polarization filters, (b) Color filters

intensity and polarization ratio image were created. These are shown in Figure 6-2.

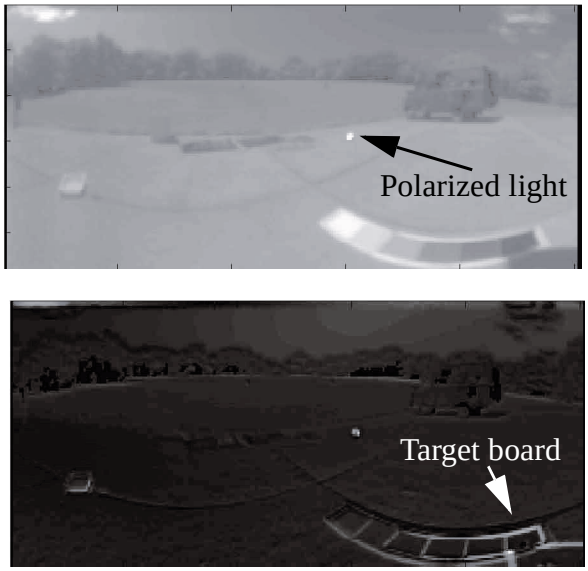

Intensity image

Polarization ratio image

FIGURE 6-2. Results of polarization scan

The intensity image is similar to a monochrome camera image, and the scene appears as expected. Notice the 6 gray scale targets. The polarization ratio image is bright where the optical energy from one area of the scene is polarized. In the scene, a small light (flashlight) with a front polarization filter produces a highly polarized light source in the scene. This spot can be clearly seen in the polarization image, since it has

a high polarization ratio. The remainder of the scene is primarily unpolarized. In particular, notice that the six calibration targets have similarly unpolarized reflections.

One unexpected result is the high polarization ratio from the board holding the six targets. The board itself has glossy lacquer finish, while the targets are matte. Therefore, a polarized specular reflection off of the board is possible. Alternatively, this could be an artifact resulting from the 1-pixel camera being out of focus at the short range (1 meter) to this spot. Also, A container of water and several metal objects were placed in the scene, but do not appear prominently in the images. Further experimentation with the polarized 1-pixel camera should resolve these issues and produce better results.

6.3.3 Register color image with terrain model

One of the exciting areas of future work involves automated registration of high-resolution color imagery from a CCD camera or other imaging sensor with the color terrain models generated by our system. This is useful for many purposes, such as camera localization, texture mapped model generation, enhanced computer vision based navigation, etc.

To facilitate the registration, the 1-pixel camera provides a key additional element. The color measurement associated with each geometric point provides a useful measure which can be compared with potential correspondences in a color image of a 2D imaging camera. The goal of such a registration procedure is to determine the actual pose of the camera when the image was taken. Figure 6-3 illustrates an iterative approach for this pose estimation method. Camera pose estimates can be iteratively evaluated by rendering the color terrain model from an estimated viewpoint and then performing a pixel-by-pixel comparison with the color image. The camera pose can be estimated and refined as the matching of these images becomes closer.

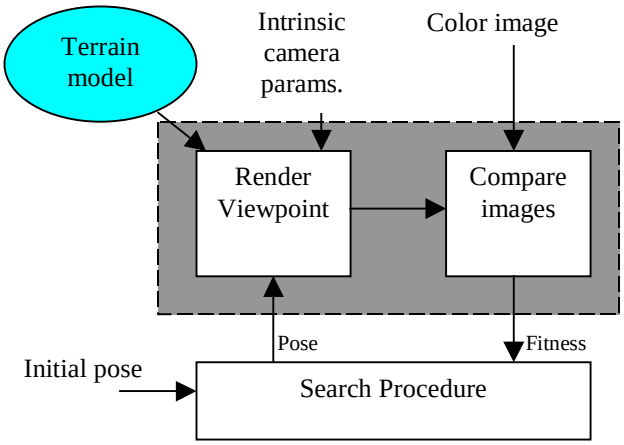

FIGURE 6-3. Automatic image - model registration

Performing the same registration without the intrinsic color measurements provided by the 1-pixel camera is virtually impossible. Therefore, we feel that this is a natural direction for future work.

6.3.4 Automated calibration flight

In this thesis, we have described how we have developed an in-flight calibration procedure which requires a series of measurements, off of a single retroreflective target, to be collected from various locations in space. To acquire these measurements, we flew the helicopter around the target area either under remote control or in an autonomous flight mode, where a remote operator would manually position the helicopter such that the scan plane approximately intersected the target, and the helicopter was perturbed slightly in an effort to hit the target. Once the target had been hit, real-time feedback notifies the operator, who had the option of collecting additional measurements from the same basic location, or else repositioning the helicopter for other viewpoints.

In practice, we found this manual approach for collecting the calibration measurements to be problematic. It is difficult for an operator on the ground to estimate where the scan plane is intersecting the terrain, so knowing where to move the

helicopter to actually hit the target is challenging. Similarly, it is difficult for the human operator to position the helicopter to collect calibration measurements from the optimal locations. Typically, we are happy with a set of measurements which are sufficiently distributed to allow the calibration to converge.

A significant improvement could be made by future work on an automated flight planner for in-flight calibration. This could provide a completely hands-off calibration procedure, and would reduce the flight time needed to perform the calibration data collection. The helicopter system could easily calculate the proper position which locates the scan plane close to the target. The calibration flight would start flying close to the target to get an initial, rough calibration which would be accurate enough to allow measurements from farther away. As the accuracy of the calibration parameters improves, long distance measurements could be acquired, and the optimal combinations of helicopter position, attitude, range and scan angle could be automatically carried out to guide the collection. These measurements are ultimately desirable, as they provide the most accurate calibrations.

Appendix A

Spectral sensitivity calculations

This appendix presents the expected spectral sensitivity function based on calculations for the 1-pixel camera. We model the theoretical spectral response of the sensor by examining the behavior of the sensor's collected light as it follows its path through the 1-pixel camera. Each channel has a corresponding spectral sensitivity function that describes the expected output of the photodiode for a certain incoming light power at particular wavelengths.

Incoming light is defined by its Spectral Power Density (SPD) function. (This is a standard approach used for colorimetry calculations [43]). The SPD specifies the light power (in Watts) per unit change in wavelength (nm) at various wavelengths. For example, the SPD of sunlight is approximately modeled by a 5800º K blackbody radiator, and is shown in Figure A-1.

Typically, the light entering the sensor has an unknown SPD. As this light passes through each of the components in its path, its SPD is altered before reaching the photodiode. The photodiode converts the total received light energy into its output photocurrent reading.

First, the cold mirror reflects the light being measured towards the sensing channels. The cold mirror is a partially reflective front-surface mirror where the amount of

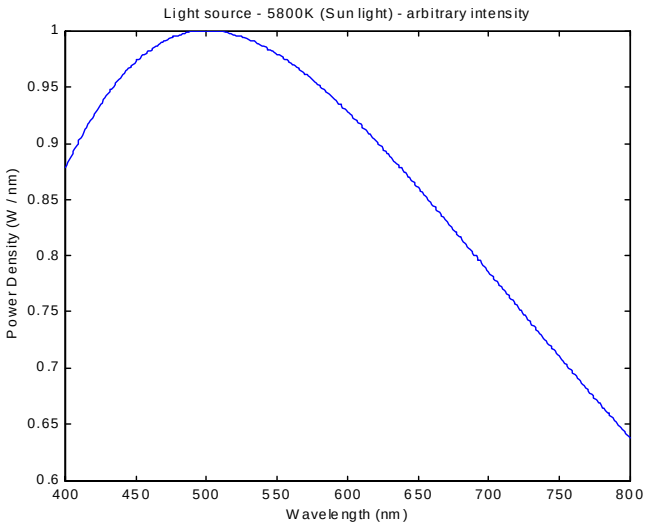

FIGURE A-1. Spectral power distribution of typical sunlight

reflectance is a function of wavelength. Figure A-1 shows the reflectance curve, $R_m(\lambda)$, for our mirror as provided by its manufacturer [29]. The SPD of the light leaving the mirror is the SPD of the incoming light attenuated by the mirror reflectance curve, shown below.

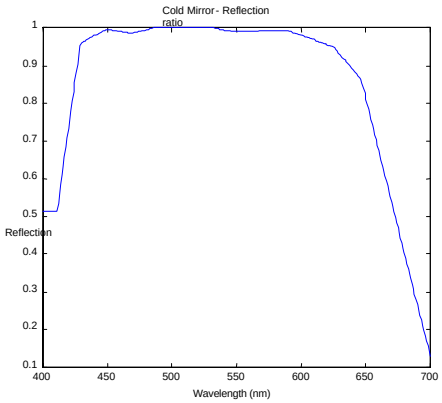

FIGURE A-1. Spectral reflectance curve for cold mirror - $R_m(\lambda)$

After reflecting off the cold mirror, the light passes through one of the three color filters. The filter attenuates the passing light as a function of wavelength, defined by

the filter transmittance. The transmittance curves provided by the manufacturer of the three filters are shown in Figure A-2.

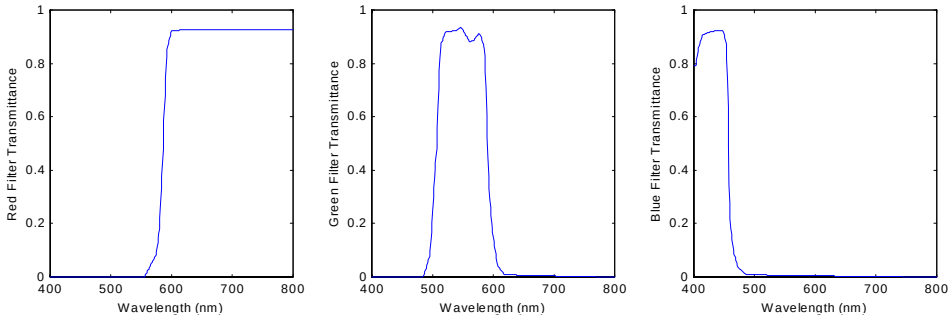

FIGURE A-2. Transmittance curves for filters

The light next encounters the focusing lens. The lens redirects the light, but does not significantly affect the shape of its SPD as the lens is neutral in color. It does, however, reduce the total power in the light due to energy losses at the air-glass and glass-air transitions. To account for this insertion loss, the SPD of the incoming light must be scaled by the lenses efficiency. The lenses selected are anti-reflection coated and provide an efficiency of approximately 98%. This is better than the typical 92% efficiency of uncoated lenses.

Only the rays of light leaving the lens and passing through the pinhole reach the photodiode. The pinhole itself has no significant effect on the light's SPD. The sensitive area of the photodiode is placed in close proximity to the pinhole to ensure that all direct rays coming through the pinhole are captured.

The photodiode generates a photocurrent output proportional to the optical power it receives. The diode has its own wavelength dependent sensitivity curve, and Figure A-3 shows the relative sensitivity for the Osram BPX-61 photodiode used in all three channels of the 1-pixel camera. Additionally, the manufacturer specifications define the peak spectral response as 0.62 Amps / Watt at 850 nm frequency. Thus, the total expected output current can be determined by taking the incident light SPD,

multiplying it by the photodiode's relative sensitivity function, and finally integrating it over the wavelength range to obtain the measured light power (watts). This wattage, multiplied by the peak spectral response value yields the output photocurrent (amps).

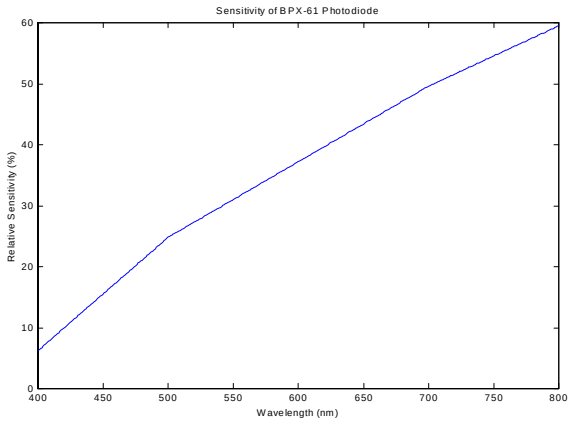

FIGURE A-3. Relative sensitivity of photodiode

These steps transform the SPD of incoming light into the 3 channel photocurrent values, we consolidate all attenuation steps into a single function directly mapping the incoming light SPD to photocurrent output values. Figure A-4 shows these curves for the three channels.

The calculations in this appendix can be summarized by a single gain constant for each channel. For a certain known spectral distribution of incomming light, the resulting photocurrent (in Amps) for a given input power (in Watts) can be calculated as we have described for each channel. For example, if we assume illumination by our sunlight model (Figure A-1), the following optical-to-electrical gains are computed:l

Optical-to-electrical gain		
Red channel	0.0615	A/W
Green channel	0.0497	A/W
Blue channel	0.0121	A/W

The spectral response shows that the 3 color channels are not evenly weighted. The

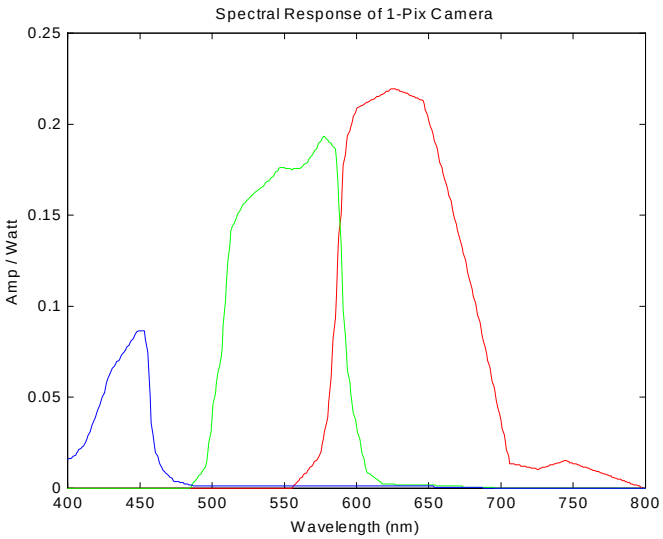

FIGURE A-4. Red, green, and blue spectral responses of 1-pixel camera

blue channel is significantly less sensitive than the others. In fact, the ratio of photocurrents when viewing a white object in typical sunlight (Figure A-1) would be 5.1 (red) to 4.1 (green) to 1.0 (blue). It is desirable to have balanced responses from the channels, as this maximizes the usable illumination range. The imbalanced response is not desirable as it requires a higher minimum illumination to generate a significant signal on all of the channels, and can tolerate a lower illumination before the higher sensitivity channels saturate. The BPX-61 photodiode was chosen for the prototype because it had the highest sensitivity in the blue region of all the diodes considered. One approach to improve the ratio, not implemented on the prototype, is to use a larger diameter lens for the blue channel to collect more light.

Appendix B

Terrain sensor specifications

B.1 Riegl laser rangefinder

Wavelength:	908nm
Exit aperture diameter:	25mm
Divergence:	3.2mrad
Accuracy:	+/-2.5 cm std.
Resolution:	0.4 cm
Measurement rate:	12,000 hz
Maximum range: (ρ=80%)	150m
Maximum range: (p=10%)	50m
Minimum range:	1m
Intensity resolution:	8 bit

B.2 1-pix camera

Number of channels:	3
Effective aperture:	18mm
FOV: (1/2 power diameter)	1.1 degrees
Depth of focus:	2m - infinity
Spectral response:	400 – 475nm (Blue Channel)
	500 – 600nm (Green Channel)
	590 – 690nm (Red Channel)
Relative sensitivity (5800K):	5 : 4 : 1 (red:green:blue)
Channel-to-channel alignment:	< 0.2 degrees
Channel-to-laser alignment:	< 0.2 degrees
Bandwidth: (-3dB)	> 5 KHz

Maximum sample rate:	12,000 Hz
Integration period:	64us
Resolution:	20 bits (17 bits effective)
Min illumination (ρ=18%, 5800K):	4.5 W/m^2
Max illumination (ρ=18%, 5800K):	30,000W/m^2
Noise (no illumination)	3 Lsb
Dimensions	5.5cm x 6.5cm x 13.0cm
Weight	340 g

References

[1] Leica Geosystems A.G., *TC1100 Total station applications manual,* www.leica-geosystems.com.

[2] Large, Goddard, and Landau, *eRTK: A New Generation of Solutions for Centimeter-Accurate Wide-Area Real-Time Positioning,* White paper from Trimble Navitagion Ltd. Dayton, Ohio.

[3] Langer, et al., *Imaging Laser for 3-D Surveying and CAD Modeling of Real World Environments,* International Journal of Robotics Research, Vol. 19, No. 11, November 2000, pp 1075-1088.

[4] Jacobs, G.L., *New large-scale 3D laser scanning, modeling and visualization technology provides advanced capabilities for scene reconstruction and interpretation*, Proceedings of the SPIE - The International Society for Optical Engineering, Vol.3905, pp 92-101.

[5] N. Pfeifer, G. Rottensteiner, *The Riegl Laser Scanner for the Survey of the Interiors of Schonbrunn Palace,* Proceedings of 5th Conference on Optical 3-D Measurement Techniques, Vienna, Austria, 2001.

[6] Baltsavias, Pateraki, and Zhang, *Radiometric and Geometric Evaluation of IKONOS Geo Images and their use for 3D Building Modelling*, Joint ISPRS Workshop “High Resolution Mapping from Space 2001”, Hannover, Germany, 19-21 September.

[7] E.M. Mikhail, J.S. Bethel & J.C. McGlone, *Introduction to Modern Photogrammetry,* Wiley, 2001.

[8] TopEye AB, *Laser Scanning: Innovative new technology presents new possibilities for spatial applications,* White paper from Top Eye AB. www.topeye.com

[9] N. Haala, D. Stallmann, and M. Cramer, *Geometric Processing of High Resolution Airborne Scanner Imagery using GPS-INS and Ground Control Points,* Third International Airborne Remote Sensing Conference and Exhibition, Copenhagen, Denmark, July 1997.

[10] T. Kanade, et al., *A Video-Rate Stereo Machine and Its New Applications*, 27th International Symposium on Industrial Robots, 1996, pp. 671-676.

[11] T. Williamson, *A High-Performance Stereo Vision System for Obstacle Detection*, doctoral dissertation, tech. report CMU-RI-TR-98-24, Robotics Institute, Carnegie Mellon University, September, 1998.

[12] McIntosh and Krupnik, *Integration of laser-derived DSMs and matched image edges for generating an accurate surface model*, ISPRS Journal of Photogrammetry & Remote sensing 2002, To appear.

[13] M.W. Barclay, N. K. Williams, *High resolution SAR/ISAR image from a helicopter platform*, Proceedings of Radar Systems (Conf. Publ. No 449), Edinburgh, UC, October 1997.

[14] N. Haala, D. Fritsch, D. Stallmann, M. Cramer, *On The Performance of Digital Airborne Pushbroom Cameras for Photogrammetric Data Processing -- A Case Study*, IAPRS Vol. 33, Part B4/1. pp. 324-331.

[15] Y. G. Leclerc, *Continuous Terrain Modeling from Image Sequences with Applications to Change Detection*, in Proceedings of the DARPA Image Understanding Workshop, (New Orleans, LA), May 1997.

[16] R. Mandelbaum, L. McDowell, L. Bogoni, B. Beich, M. Hansen, *Real-time stereo processing, obstacle detection, and terrain estimation from vehicle-mounted stereo cameras*, Proceedings of the Fourth IEEE Workshop on Applications of Computer Vision, pp 288-289, Princeton NJ, Oct. 19-21, 1998.

[17] D. Langer, M. Mettenleiter, F. Hartl, C. Frohlich, *Imaging Ladar for 3-D Surveying and CAD Modeling of Real World Environments*, The International Journal of Robotics Research, Vol. 19, No. 11, November 2000, pp. 1075-1088.

[18] *Last Pulse Measuring Technique*, Application note GI009, Riegl Laser Measurement Systems GmbH, Austria.

[19] M. Hebert and E. Krotkov, *3D Measurements from Imaging Laser Rarars: How Good Are They?*, Image and Vision Computing, vol. 19, No. 3, April, 1992.

[20] American National Standards Institute, Inc., *American National Standard for the Safe Use of Lasers, Z136.1.1993*. (ANSI Z136.1-1993).

[21] A. Ryer, *Light Measurement Handbook*, International Light, Inc. Newburyport, MA, 1997.

[22] RCA, *Electro-Optics Handbook*, RCA Solid State Division, Lancaster, PA, 1974.

[23] *Pulnix Imaging Products: Spectral Response Curves*, Pulnix America, Inc. Sunnyvale, CA.

[24] A. Samberg, *Mathematical model for time-varying system employing across-track scanning,* Proceedings of the SPIE, Vol 3065, 1997, pages 214-23.

[25] F. R. Horowitz and F. R. Hill, *The Art of Electronics, 2nd ed.* Cambridge, U.K.: Cambridge Univ. Press, 1993.

[26] Burr Brown, *DDC101 20-bit Analog-to-Digital Converter Datasheet*, BurrBrown Inc, 1993.

[27] J. Graeme, *Photodiode Amplifiers: Op Amp Solutions,* McGraw Hill, 1995.

[28] K. Booth and S. Hill, *The Essence of Optoelectronics,* Prentice Hall, 1998.

[29] *Product Specifications Data*, Edmund Industrial Optics, Barrington, NJ. www.edmundoptics.com

[30] B*PX-61 Silicon PIN Photodiode Datasheet,* Infineon Technologies AG. www.infineon.com

[31] *AccuRange Line Scanner Datasheet,* Acuity Research Inc. Menlo Park, CA. www.acuityresearch.com

[32] C. Poynton, *Frequently Asked Questions about Gamma*, <http:// www.inforamp.net/~poynton/>.

[33] O. Amidi, *An Autonomous Vision-Guided Helicopter,* doctoral dissertation, Robotics Institute, Carnegie Mellon University, 1996.

[34] Munsell Color Laboratory, GretagMacbeth, New Windsor, NY.

[35] R. Guenther, *Modern Optics,* Wiley Press, 1990.

[36] P. Lee, *The Haughton-Mars Project: Study of the Haughton Crater – Devon Island, Canadian High Arctic, Viewed as a Mars Analog,* NASA Tech Report.

[37] NovAtel, Inc., *BeeLine GPSCard Users Manual*, 1998.

[38] R.Collins, Y.Tsin, J.R.Miller and A.Lipton, *Using a DEM to Determine Geospatial Object Trajectories,* DARPA Image Understanding Workshop, Monterey, CA, November 1998, pp. 115-122. Also CMU Robotics Institute technical report CMU-RI-TR-98-19, July 1998.

[39] R. Miller and O. Amidi, *3-D Site Mapping with the CMU Autonomous Helicopter,* Proc. of 5th International Conference on Intelligent Autonomous Systems. (IAS-5), June 1998.

[40] R. Miller, O. Amidi, and M. DeLouis, *Arctic Test Flights of the CMU Autonomous Helicopter,* Proceedings of the Association for Unmanned Vehicle Systems International 1999. 26th Annual Symposium, Baltimore, MD.

[41] R. Reichert, J. Reigl, A. Ullrich, *Time-of-Flight Based 3-D Imaging Sensor with True-Colour Channel for Automated Texturing*, Proceedings of 5th Conference on Optical 3-D Measurement Techniques, Vienna, Austria, 2001.

[42] L. Wolff, *Applications of Polarization Camera Technology*, IEEE Expert, Vol. 10, Issue 5, October 1995.

[43] W. Driscoll, *Handbook of Optics*, Optical Society of America. McGraw-Hill, 1978.

[44] M. Kayton, *Avionics Navigation Systems, 2nd Edition*, Wiley-Interscience publication, New York, N.Y., 1996, pages 384-386.

[45] O. Amidi, T. Kanade and R. Miller, *Vision-based Autonomous Helicopter Research at Carnegie Mellon Robotics Institute,* Proceedings of Heli Japan ‘98. Gifu, Japan, April 1998.

[46] *Riegl LD90-3GF User’s Manual, Riegl Laser Measurement Systems*, Riegl Laser Measurement Systems GmbH. www.riegl-usa.com.

[47] *PB-R006 Photobit TrueColor Processor for CMOS Image Sensors -- Product Specification,* Photobit, Inc. January 2000.

[48] Barnard and Funt, *Camera characterization for color research*, Color Research and Application, 2001. John Wily & Sons, Inc.

[49] Vora, Farrell, Tietz, and Brainard, *Digital color cameras. 1. Response models*, Hewlett-Packard Laboratory, Technical Report HPL-97-53.


THE END