

Design and Implementation of UAV Autoland Control Law Utilizing Barometric Pressure, GPS, and LIDAR- Lite for Altitude Determination

A project present to
The Faculty of the Department of Aerospace Engineering
San Jose State University

in partial fulfillment of the requirements for the degree
Master of Science in Aerospace Engineering

By

Frank C. Sanders

May 2015

approved by

Dr. Jeanine Hunter
Faculty Advisor



Design and Implementation of UAV Autoland Control Law Utilizing Barometric Pressure, GPS, and LIDAR-Lite for Altitude Determination

Frank C. Sanders

This paper presents the design and flight testing methodology to land a fixed wing Unmanned Aerial Vehicle (UAV) autonomously. Due to various hardware inaccuracies and limitations, Barometric Pressure, GPS, and LIDAR (Light Detection and Ranging) are used to determine the altitude for the increasing precision needs during landing. This autoland controller operates independent of ground support systems and depends solely on the on-aircraft systems.

I. Introduction

The subject of autonomously landing an aircraft has been widely studied as a means to prevent pilot error or aid a crew in the event of an emergency. Many approaches to the design problem have been purely theoretical and few have made it to full scale aircraft. Most of the proven implementations of autonomous landing control schemes have been applied to small unmanned aerial vehicles. The reasons are simple: several iterations of the control schemes can be implemented quickly without extensive robustness studies to ensure pilot safety, the relative cost of the hardware is kept low for both the initial investment and for damages/losses of aircraft, and the collateral damage risks are significantly lower due to lack of an on-board pilot and relatively low mass of the system.

Autoland an aircraft is an important challenge to overcome not only because it is one of the last remaining phases of flight to be automated, but also because landing is one of the most dangerous phases of flight. In favorable conditions, the landing phase of flight is trivial to any pilot. However, once atmospheric and physiological abnormalities begin to accumulate against the pilot, the landing procedure can easily become an incredibly difficult task. Due to the proximity to the ground during landing, a high level precision is required for the task, but is often difficult to attain during adverse conditions.

This autoland control law will alleviate the problems associated with adverse atmospheric conditions and, since there is no pilot input, will do away with human piloting error. The system is also entirely contained within the aircraft. Most other designs depend on a ground station sending glide-slope data to the aircraft; thus, making the aircraft dependent on only airports equipped to send that specific data. This aircraft will only require GPS coordinates of the landing zone, making the system independent of specific airports and useful for all airports.

II. Literature Review

A review of the literature has provided many interesting approaches to the design problem and many starting points for the different phases of the landing approach. The most promising designs have come from Sperry Flight Systems, German Aerospace Center (DLR-Oberpfaffenhofen), and Boeing Commercial Airplane Company. Each institution has developed a similar systems that are ground dependent, but these three are unique in that they have all own full scale articles. Although not all of their designs were successfully tuned and implemented, they all paved the path towards a system that operates entirely independent of the ground.

The Boeing Aircraft Company developed and flight tested a automatic landing are control laws for the B-737 and B-747 aircraft. Although their control laws did not handle the approach phase of landing, they

Graduate Student, SJSU, Control Science & Dynamical Systems (CSDy), frank.sanders@sjsu.edu.

did develop a system to control the most critical phase of landing. If the aircraft is not aligned straight on the runway, has excessive pitch or roll angles or rates, or has too much or too little sink rate then the entire portion of the landing will cause excessive control inputs and may lead to damage of the aircraft or ground items (runway lights, buildings, cars, etc) if the aircraft goes unstable.

The German Aerospace Center designed and flew an aircraft they designated the ATTAS (Advanced Technologies Testing Aircraft System). The project was government funded and was ultimately canceled before being totally problem free. However, their design procedure provides an excellent framework for future autoland development due to their pointed assumptions and associated pitfalls experienced when implementing the system on an aircraft.

Sperry Flight Systems was able to develop and perfect an autoland control law for NASA's Space Shuttle. Their architecture was unique in that they used a total energy management system due to the fact that the Shuttle is a glider. Their methodology utilized different assumptions than other systems, but their method will be quite useful in designing a ground independent system since the glide slope determination is well documented and defined.

III. Problem Definition

The autoland system must be able to guide the aircraft from a position in the sky, establish a glide slope, align itself with the runway until just before touchdown, and arrest the rate of descent to an acceptable level at touchdown. While the problem seems relatively simple, there is an overarching parameter that affects every aspect of the landing and poses the biggest difficulty in the design process: altitude. Altitude is the parameter that is present in every single phase of the landing process, and it is the most difficult parameter to estimate, especially with the changing precision requirements during landing.

Commercially available hardware each have their own inherent uncertainties. One cannot fully depend on one method of determining altitude, for the hardware uncertainties and operational limits each have their own pitfalls if depended on in the landing phase. For example, the GPS used in this system has an uncertainty of 10 feet. That large of an error is acceptable if one wants to estimate the altitude in the alignment phase of the landing. However, for the flare phase, if the altitude fed to the aircraft reads 10 feet above the actual altitude above the ground it will stall and crash. If the altitude fed to the aircraft is 10 feet lower than actual, the aircraft will plunge into the ground with an unacceptable sink rate. To remedy the hardware pitfalls, this control law depends on three different altitude determination systems: barometric pressure altitude, GPS altitude, and LIDAR. Their roles are discussed in detail in the Approach section.

IV. System Methodology

There are three phases in the landing sequence: Approach, Descent, and Flare. Each phase has its own specific altitude accuracy requirements, and each phase's systems and requirements will be discussed in detail.

A. Approach

The first phase of the landing sequence is the approach. The aircraft needs to be in steady state flight and intersect the initial position (IP) where the next phase is initiated. When the controller is switched on, the aircraft will use GPS to fly through three waypoints before arriving at the descent IP in steady state flight. For altitude determination in this phase, the system will use a combination of the barometric pressure and GPS altitude. (REFERENCE) found that combining barometric altitude (accurate to within 7 feet) with GPS altitude (accurate to within 10 feet) reduces the altitude determination error to three feet. A three foot error at several hundred feet above the ground is acceptable to initiate the next phase. The flight path during this phase is presented in Figure 1.

B. Descent

When the aircraft reaches the IP, the data feedback to the controller will initiate the descent phase. During the descent phase, the aircraft will reduce power and begin a three degree descent towards another waypoint at the end of the runway. Control will be achieved with transfer functions relating the rate of decay of




Figure 1. Waypoint Tracking During Approach Phase

the combined GPS and Barometric variable. Once the aircraft's altitude is within range of the LIDAR, the LIDAR will solely provide altitude information to the controller for the rest of the descent until the are IP is reached.




Figure 2. Descent Phase

C. Flare

Once the aircraft reaches the are IP and is following the LIDAR altitude measurements, the aircraft begins the are routine. The details of the are routine have not been de ned yet, but during the are phase the power will be cut completely, sink rate arrested, and heading held until dynamic pressure reaches zero.

V. Work So Far

Much work went into determining how best to address the measurement error present in all distance measuring devices available on the market for aircraft of this size. The most promising sensors for carrying aloft on the aircraft were ultrasonic range finders and LIDAR.

The first method of determining altitude with a high degree of precision explored was the ultrasonic range finder. The ultrasonic range finder operates by emanating a 42kHz sound wave and measuring the time of flight of the return wave. This sensing method is accurate to one inch and is a perfect candidate for answering the precision problem during the flare phase. An ultrasonic range finder was mounted to the test aircraft and flown at various altitudes and speeds to determine the sensor's performance. Unfortunately, flight testing of the sensor readily exposed the sensor's flaws. Since the sensor depends on sound waves, the sensor is limited to the speed of sound. The collection rate of the sensor is 10Hz and not nearly up to the task of controlling the descent rate of an aircraft. In addition, the minimum sensing distance of the sensor was one foot (i.e. above one foot it would display the range, if below one foot it would read one foot). With the performance and ranging limitations in mind, the ultrasonic range finder was abandoned.




Figure 3. Ultrasonic Range Finder Attached to Flight Computer

After abandoning the ultrasonic range finder sensor, the search began for an alternative sensor that could operate at a much higher sampling rate. The intuitive answer was to look towards light based sensors, and LIDAR is the obvious light based ranging sensor that is readily available. However, the major obstacle with LIDAR is its size and cost. Until recently, LIDAR systems weighed a minimum of two pounds and cost several thousand dollars. Considering the highly limited budget of the project and the fact that the lightest LIDAR weighs the same as the test aircraft, alternatives to the commercially available sensors were thoroughly researched.

The search for an alternative LIDAR sensor yielded another sensor that is used in an unrelated, commercially available product: the Wii. The Nintendo Wii is a motion based video game console that allows users to wirelessly interact with the console via Bluetooth connection. However, in the front of the Wii remote is an infrared (IR) sensor that is capable of detecting up to four targets and outputs the x and y coordinates of each at 200Hz. Using the IR sensor, a LIDAR system (Figure 3) was designed and developed using a red laser to set at a known distance from the sensor to establish a sloped line. The system was designed so that as the aircraft descended, the y-coordinate would change and then be fed back and used to determine the height above the ground. Several months of development led to a successfully operating IR based LIDAR system. However, during the performance evaluations, the Wii LIDAR proved to be inadequate for the autoland task. While the sampling frequency was more than ample, the system was prone to locking on to IR sources that were unrelated to the task. In fact, while tracking an IR diode in the lab, the LIDAR was outputting mysterious coordinates that later turned out to be fluorescent ceiling lights. In addition, the effective range

of the LIDAR, with the laser that was intended to be mounted to the aircraft, was only about one and a half feet. The custom developed LIDAR was deemed unusable and was abandoned.




Figure 4. Custom Developed Wiimote LIDAR

Just shortly before the abandonment of the custom LIDAR, the company PulsedLight released their answer to the over-weight/over-priced problem: the LIDAR-Lite(Figure 5). This LIDAR system was developed for small robotics projects and weighs a mere 0.2 ounces at a cost of \$90.00. The LIDAR-Lite has a sampling frequency of 100Hz, maximum range of 130 feet, and an accuracy of one inch. This sensor t every requirement for the are phase and would be able to control the altitude all the way until the wheels touch the runway. Since the LIDAR-Lite was such a new product, only a few people had tested the best way to control it, since the sensor could be operated with pulse width modulation(PWM) or I²C. Since there was no clear optimal route to control the LIDAR, testing began. The LIDAR was mounted to a R/C helicopter (Figure 6) to evaluate actual ight collected data, as opposed to laboratory bench collected. Although the helicopter is not what the sensor will have ultimately been mounted to, the helicopter provided a vehicle with excessive, worst-case vibrational noise injected to the system and it allowed short up-and-down hops that cut down on excessive data collection. At the end of ight testing with the helicopter, it was decided that PWM was the best method to collect data. The LIDAR was then mounted to an R/C glider (Figure 7) and own in a manner similar to what is to be expected in an autonomous landing. The LIDAR performed as expected and is currently being mounted to the test aircraft.



Figure 5. The LIDAR-Lite



Figure 6. The LIDAR-Lite



Figure 7. The LIDAR-Lite