Self-Aware Helicopters:

Full-Scale Automated Landing and Obstacle Avoidance in Unmapped Environments

Lyle Chamberlain, Sebastian Scherer, Sanjiv Singh
Carnegie Mellon University, Pittsburgh PA 15213
(lyle@andrew.cmu.edu)

ABSTRACT

In this paper we present a perception and autonomy package that for the first time allows a full-scale unmanned helicopter (the Boeing Unmanned Little Bird) to automatically fly through unmapped, obstacle-laden terrain, find a landing zone, and perform a safe landing near a casualty, all with no human control or input. The system also demonstrates the ability to avoid obstacles while in low-altitude flight. The perception system consists of a 3D LADAR mapping unit with sufficient range, accuracy, and bandwidth to bring autonomous flight into the realm of full-scale aircraft. Efficient evaluation of this data and fast planning algorithms provide the aircraft with safe flight trajectories in real-time. We show the results of several fully autonomous landing and obstacle avoidance missions.

Figure 1: Gimbaled laser scanner mounted on the nose of the Boeing Unmanned Little Bird

NOTATION

AGL – Above Ground Level
ULB – Unmanned Little Bird
LADAR – Laser Detection and Ranging
UAS – Unmanned Aerial System
IMU – Inertial Measurement System
INS – Inertial Navigation System
GPS – Global Positioning System

INTRODUCTION

Unmanned aerial systems (UAS) have seen increased use in battlefield scenarios over recent years. High altitude surveillance, weapons designation, or communication relay missions are now commonplace. Generally, these missions are sutiable for fixed-wing aircraft because it is acceptable to fly at high altitudes and to always take-off/land-on prepared runways. UAS have not seen much use in missions that

Presented at the American Helicopter Society 67th Annual Forum, Virginia Beach, VA, May 3-5, 2011.

require operation at low altitude, unimproved landing sites, or less oversight from expert operators. Rotorcraft UAS platforms have been suggested as viable solutions for casualty extraction, cargo resupply, and low-altitude surveillance.

Before these low-altitude missions can be practical, there must be a reliable way of safely operating near power lines, buildings, or trees. Landing blind at unprepared sites can be catastrophic if small GPS errors cause the vehicle to land in an unintended spot (say, 10 feet to the left, on a small boulder). In general, operation in unprepared areas is prone to operator error because of the lack of established infrastructure and procedure for air operations in a particular location. If the operator gives an approach vector that has an unseen obstacle, or is in proximity to known obstacles, the UAS should have the "common sense" to do the right thing even if the operator acted in error.

Hence two key capabilities essential for low-altitude operations are landing zone selection, and obstacle avoidance. Automated landing has been implemented on several small UAS platform. Machine vision systems have been popular on small radio-controlled helicopters. (Bosch, Lacroix, & Caballero, 2006) used monocular images from a camera to identify planar areas for landing. (Johnson, Montgomery, & Matthies, 2005) used structure-from-motion to build 3D models of the terrain for guidance and landing of descending spacecraft. (Templeton, Shim, Geyer, & Sastry, 2007) also demonstrated a landing system based on similar structure-from-motion concepts.

While machine vision does work for many smaller aircraft that can hover very close to terrain for long periods of time, larger full-scale rotorcraft that are useful for casualty extraction or cargo resupply are not well suited to such flight profiles. In addition, the close standoff distances needed for high-resolution measurement can cause problems such as brownout, when the downwash from the rotor blades causes a vision-occluding dust storm.

LADAR systems have been used as high-resolution solutions to the autonomous landing problem. (Johnson, Klumpp, Collier, & Wolf, 2002) did early analysis and simulation that suggested that LADAR would be well suited for the similar problem of landing on Mars. (Whalley, Directorate, Schulein, & Theodore, 2008) implemented an autonomous landing system on a 4-meter RMax unmanned helicopter platform. The results of this work showed the enormous promise of compact 3D LADAR as a navigation aid for a UAS. In a similar vein, our group has demonstrated a perception system capable of identifying landing zones on a full-scale EC-135 helicopter (Scherer, Chamberlain, & Singh, AIAA Infotech@Aerospace 20-22 April 2010, Altanta, Georgia, 2010). This paper describes our efforts to implement a 3D LADAR-based perception and planning solution that accommodates the high speeds, shallow glide slopes on approach, and limited-hover requirements of a large cargo-carrying rotorcraft.

We also set out to extend our earlier work in obstacle avoidance, previously demonstrated on an RMax helicopter (Scherer, Singh, Chamberlain, & Elgersma, 2008), to a full-scale platform.

SENSING/PERCEPTION HARDWARE

The sensing system for an unmanned platform operating in an unmapped environment must provide two types of information: terrain information for analysis of potential landing zones, and proximity information about obstacles it may encounter en-route. Many sensing modalities are available for both tasks, with varying suitability for each task. The driving requirement in our case was high-resolution measurement of the terrain while flying at a reasonable speed and altitude for a full-scale helicopter. This terrain measurement requirement rules out radar and Machine Vision. A radar sensor with sufficient angular resolution would require a large antenna, and the depth and angular resolution of vision systems is inadequate except at close range. Recent advances in scanning ladar technology have improved the range and bandwidth significantly, opening this modality as a prime candidate for a viable perception system on a full-scale helicopter.

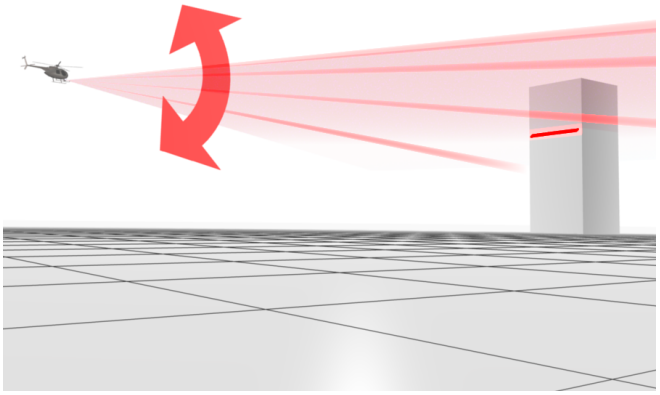

Figure 2: Obstacle-Detection Scan Configuration. Scanner faces forward and nods up and down to detect obstacles as the aircraft descends to lower altitudes.

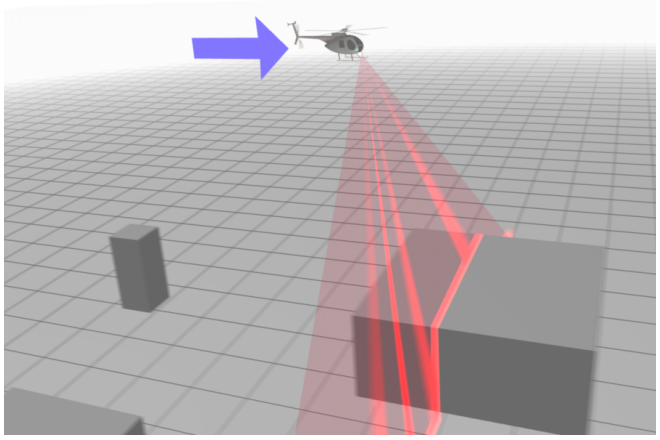

Figure 3: Landing Zone Survey Scan Configuration. Planar scanning sweeps terrain as the aircraft overflies potential landing zones from a safe altitude.

We designed and built a 3-D scanning ladar that operates in two modes: Forward-scanning for obstacle detection during low-altitude flight (Figure 2), and downward scanning for terrain mapping and landing zone search from a higher altitude (Figure 3). The scanner operates in two axes: a fast 100 Hz, 100-degree transverse scan and a slower "nodding" scan that cycles up and down. The fast-axis scan is provided off-the-shelf by the laser manufacturer, while the slow-axis is a single-axis gimbal that we fabricated in-house. The ladar is capable of pulsing at nearly 84 kHz, giving the point density for a high-resolution terrain map. The sensor has a range of 150 meters, centimeter accuracy, and can process multiple returns for obscurant and vegetation penetration. The Laser is Class 1 and eye-safe.

The motion control for actuation of both axes is precisely synchronized to a global timeframe so that for each measurement, the time is known and can be fused with INS data for point registration. Absolute encoders and a highly rigid frame allow the pointing angles to be known to <0.001 degree.

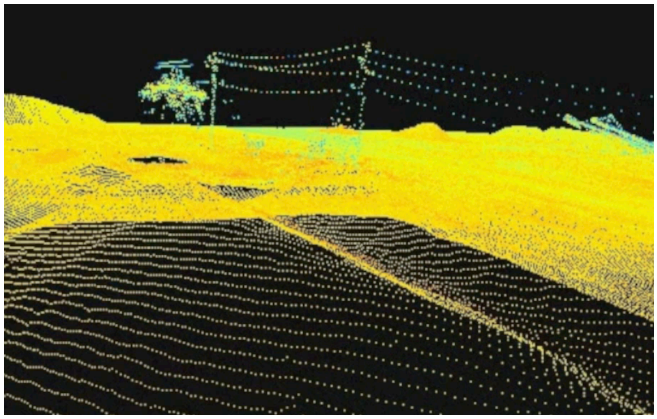

Figure 4: Example point cloud showing terrain and power lines.

The ladar is rigidly coupled to a high-accuracy GPS/INS system so that the attitude of the whole sensor head is precisely known at all times despite vibration. Otherwise, the map will become "blurred" due to errors in the measurement of where the instrument was pointed when each point was measured. Each measured point is registered into a global reference frame. The INS uses a ring-laser gyro with <10-deg/hour drift, differential GPS with 1cm accuracy, and a Kalman filter that timestamps the measurements into a timeframe that is synchronized with the rest of the perception and computing system.

Figure 5: Close-up of the sensor head mounted on the nose of the Unmanned Little Bird.

The system was shown to be capable of detecting chain link fences, wires (Figure 4), and 4-inch flat pallets from up to 100 meters away (Figure **6**) during flight on a helicopter. Maps built without differential GPS are still locally consistent after one map-building run, but possibly inconsistent over time. In our experiments, we operate with 1-centimeter differential GPS, but the maps would still have had sufficient local consistency to pick up small obstacles, even if their actual registered position was slightly off.

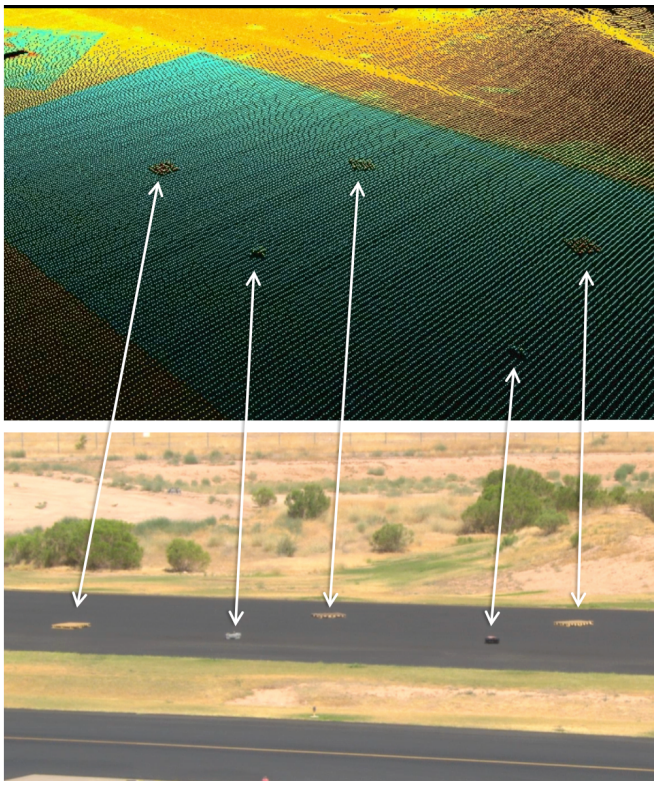

Figure 6: Obstacle registration near landing area. Detection of small objects (here, forklift pallets and plastic storage boxes) requires high resolution and precise coupling with the inertial measurement system in order to resolve small features from long distances. This map was made at 40 knots from 150 feet AGL.

Figure 7 shows the overall system design. Data from the ladar/INS feeds into the perception module which interprets the sensor data into meaningful perception primitives.. The planning system uses these primitives to decide how to achieve the mission objectives while keeping the aircraft safe and obeying the dynamic constraints of the aircraft. The result is a trajectory that is sent to the autopilot as often as necessary to respond to new information arriving from the perception system. The software is efficient, running in real-time on two off-the-shelf Core 2 Quattro PCs running at 2.2 GHz.

The Boeing Unmanned Little Bird (ULB) aircraft already has the capability of waypoint-mode autonomous flight and landing. Our system provides the autopilot on the ULB with obstacle-free trajectories to follow, and with verified landing zones.

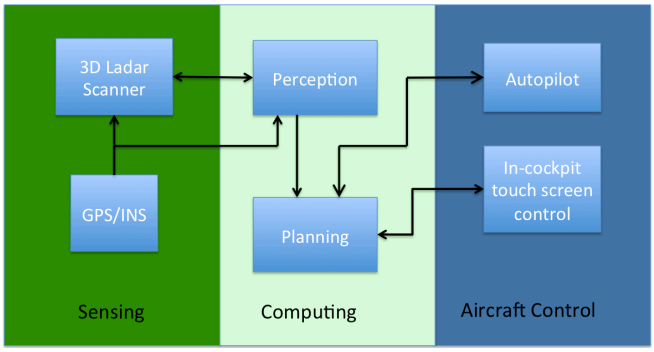

Figure 7: Hardware Functional Diagram. Data from the 3D scanning ladar is registered to a global reference frame using the GPS/INS. The planning module evaluates terrain and obstacles, and also plans safe trajectories to complete the mission. These trajectories are sent to the flight control system autopilot.

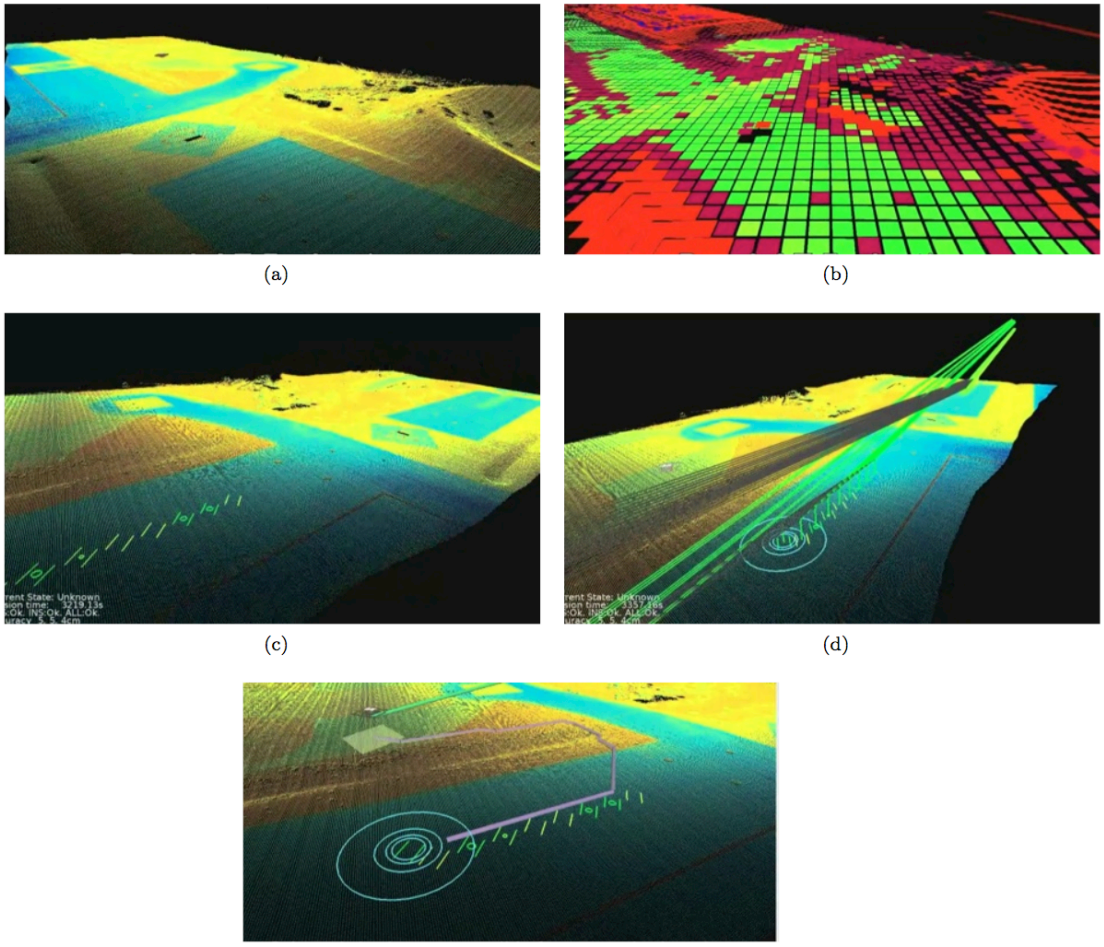

Figure 8: Landing Zone Evaluation is performed in 5 steps: (a) Overflight survey point cloud, (b) rough terrain evaluation, (c) fine terrain evaluation, (d) approach/abort path validation, and (e) ground path for dismount traversal to the simulated casualty

LANDING ZONE SELECTION AND LANDING

At the beginning of a mission, the aircraft is only given a GPS coordinate that marks the location of a casualty that needs to be picked up, and no-fly zones. No other information is given a-priori and hence the specific landing site for the helicopter is initially unknown.

Figure **8** shows the steps the system follows in order to find a landing site. First, the aircraft overflies the known position of the casualty in downward-scan survey configuration (see Figure 3) and builds a map of the surrounding terrain in the global coordinate frame. This map is shown in artificial color with reflectance values in Figure **8**a. Note that in this example there are several obstacles on what would normally be an ideal flat and level place to land on the tarmac. (See Figure **6** for a close up of the obstacles in a similar scan.)

A large detailed terrain map is hardly necessary to discount many sites as unsuitable. The first data-processing step, shown in Figure **8**b, does a *rough evaluation* to immediately remove rough or sloped areas from consideration. The green grid squares (or light, for those viewing in greyscale) are considered smooth enough for a closer look. Most of the terrain that is not on the tarmac is immediately discounted (red, or dark) from further consideration. Note that the clutter on the runway is also discounted in this step.

Now that the potential landing sites have been narrowed down to a few more feasible places, the system can perform a fine evaluation of each to decide which is best. The fine evaluation places a 3-dimensional virtual model of the helicopter on each grid cell and evaluates skid/ground contact, position of center-of-gravity within the support points (static tipover stability), wind direction, and clearance with adjacent terrain both beneath and adjacent to the rotor and tail. The software evaluates multiple headings at each potential landing spot. In this example, the Fine Evaluation step found good landing spots shown as parallel skid positions in Figure **8**c. The relative suitability of a given position is indicated by the size of the circle rendered between the landing skid markers. Note that these sites are located in the center of the tarmac. The *fine evaluation* step discarded the sites near the edge of the flight line or near the clutter obstacles because there was less clearance for the rotor/tail compared with the center locations.

Now that there is a list of acceptable touch-down points, the system evaluates the glide slope to each one from multiple directions, as well as the corresponding abort trajectory (Figure **8**d). These approach/abort paths are compared against the point cloud model and evaluated for proximity to obstacles such as buildings or power lines. Also, glide slopes with a headwind are preferred; crosswind is less acceptable, and downwind is rejected.

Simultaneous to the approach/abort step, the system also evaluates the path that a dismounted medical crew would have to follow to get to the landing zone from the reported casualty position. Figure **8**e shows a path from a selected site (marked by concentric rings) to the casualty (diamond). This ensures that the helicopter won't land on the wrong side of a ravine or fence.

The actual landing site is chosen by a multiple-objective optimization function that evaluates the results from the approach/abort and ground path metrics and comes up with a compromise between the two decision criteria. Once the final site is chosen, a path planner computes a landing pattern trajectory that will start its decent at the beginning of the approach path corresponding to the winning landing zone.

When the vehicle starts its decent at the beginning of the final glide to the landing zone, it switches the perception mode to forward facing for obstacle detection (Figure 2). This helps ensure that new or previously occluded objects in the glide path are detected and avoided by aborting the landing.

OBSTACLE AVOIDANCE

We extended our earlier work on a Yamaha RMax helicopter for obstacle avoidance in unmapped environments (Scherer, Singh, Chamberlain, & Elgersma, 2008). Obstacles are registered separately from the landing map in an evidence-grid representation of the world (Martin & Moravec, 1996). As the helicopter flies, its planned trajectory is compared to a C-Space expansion of the evidence grid (Scherer, Ferguson, & Singh, 2009). If the trajectory intersects an obstacle, then a new path is planned. In simulations, the path planner uses a hybrid trajectory optimization/path-planning algorithm utilizing a forward model of the helicopter dynamics to create a viable avoidance path. In actual flight tests, the limited waypoint-style interface to the flight controller necessitated a different planning method. We used the LIDT and D* lite algorithm to create a path around obstacles, and then forward-simulated the resulting waypoint paths to ensure that they did not intersect the obstacle.

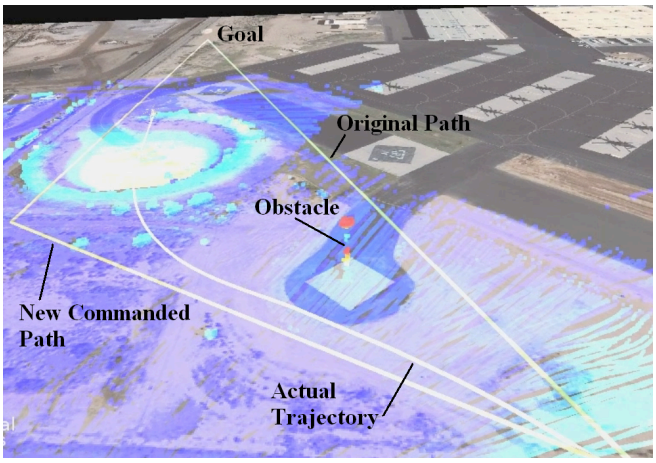

Figure 9: Data render during an autonomous obstacle avoidance maneuver at 21 knots ground speed. Points measured by the LADAR in-flight are overlaid against a satellite image for this visualization.

Figure 10: The system autonomously guides the Unmanned Little Bird around a previously unmapped obstacle (an extended manlift).

Figure **9** shows an example obstacle avoidance test and illustrates the resolution of commands we were able to send the Unmanned Little Bird controller. The original path was defined using a single waypoint, and the projected "original path" is the forward-simulated prediction of where the vehicle would fly. The new obstacle avoidance path is given using just three intermediate waypoints that are generated and verified using the D* and forward simulation approach outlined above. Even though the vehicle does not closely follow the connecting lines, the behavior is similar to what the path planner predicts using its forward simulation. It can be seen that operation in more cluttered airspace with slalom-like maneuvers will require a tighter integration between the flight controller and the path planner (as simulated as an ideal system design earlier in the project).

EXPERIMENTAL RESULTS

Tests were performed on the flight line of the Boeing test facility in Mesa Arizona. Autonomous landing and obstacle avoidance were tested in independent workups, close to the airfield.

Landing

After system integration and testing, a total of 8 landing missions were demonstrated with varying ground clutter and approach obstructions.

Figure 11 shows a typical autonomous landing mission. After autonomous takeoff and climb out, the aircraft approaches the flight line from the southwest. It performs the survey-mode flyover (typically at 40 knots, 150 ft) to search for acceptable landing sites. The resulting point cloud is shown overlaid on satellite imagery in Figure 11. After scanning the terrain and deciding on a landing zone, the system sets up an approach from the north and lands near the casualty.

Note, that the computation time for finding the landing zones is low enough that the best landing site has been selected within a few seconds after finishing the survey flyover and the system immediately begins to set up an approach.

A comparison of the landing approaches in Figure 11 and Figure 12 shows the effect that changing the priorities of the optimization process has on the decisions the system makes. In Figure 11, the optimization process was not heavily penalized for approaches that descend through unmapped terrain. This is acceptable since the system switches to forward-looking scanning (Figure 2) when flying at lower altitudes and could respond to unexpected obstacles. In this particular example, the system likely wanted to avoid the mining equipment in the quarry (upper right of scan) and the manlift we had placed near the flight line; consequently, it chose to land from the North.

In contrast, Figure 12 shows a mission where the optimization process was heavily penalized for planning any low-altitude flight where no map or survey existed. In this case the system chose to perform a button-hook turn at the end of the survey scan and land in the direction it came. It was forced to get closer to the mining equipment and manlift, but it only flew through areas that it already had mapped. Until we further test wave-off obstacle avoidance, this constraint is probably still the safer option.

Figure 13 shows a close-up of two missions flown with no penalty for descending outside the surveyed terrain. In Figure 13a, the system chooses to land in a large clear area on the flight line. It is a bit distant from the casualty (the red "+"), but the multi-objective optimization chooses to land there because it is a better landing site overall.

In Figure 13b, the best landing areas have been cluttered with low-profile obstacles. (The obstacle locations are annotated with the red X's because they are difficult to see in a small image. See Figure **6** for a close-up view of the type of clutter used and the resulting point cloud). In this case, the system decides to land just north of the casualty in a small space. The red circle around the casualty marker "+" denotes the safety radius around the casualty position. The vehicle has to use a north approach to land there due to the no-fly restrictions and the manlift obstacle to the northeast (see Figure 11).

Figure **14** shows a pair of tests when the optimization would be penalized for any flight through unmapped areas. Like the previous example, the system decides to land in the large open area shown in Figure **14**a when there is no clutter on the flight line. However, with clutter placed similar to the previous example, the system decides to land further down the flight line using the approach from the northeast (see Figure 12). In this case, the landing zone to the north of the casualty was discounted because the northeast approach would be blocked by the manlift obstacle.

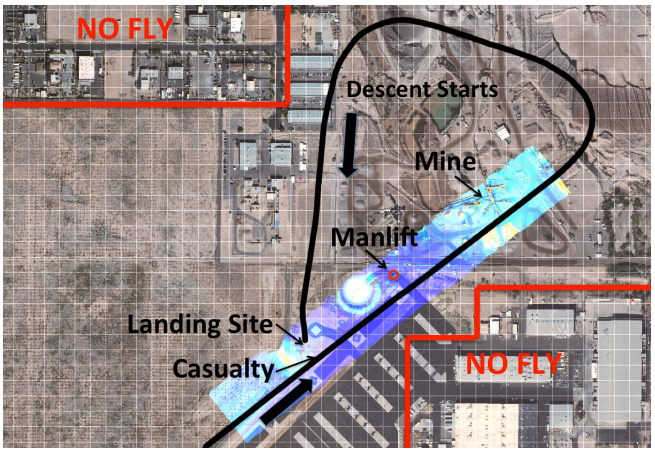

Figure 11: Typical landing mission including survey overflight, descent on final, and touchdown. Here, the system is not penalized for flying the final approach through unsurveyed terrain, and chooses to land from the north. The forward-facing scan would be responsible for detecting any obstacles in the final approach.

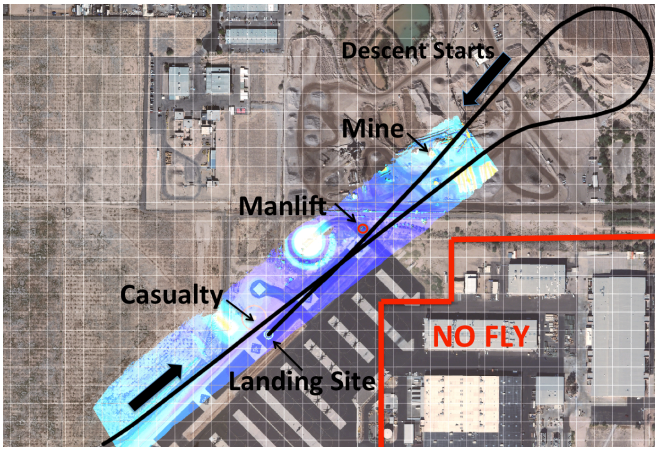

Figure 12: A landing mission where the system is penalized for descending on final through unsurveyed terrain. It chooses to perform a "button-hook" turn to final and descend through the area it just mapped.

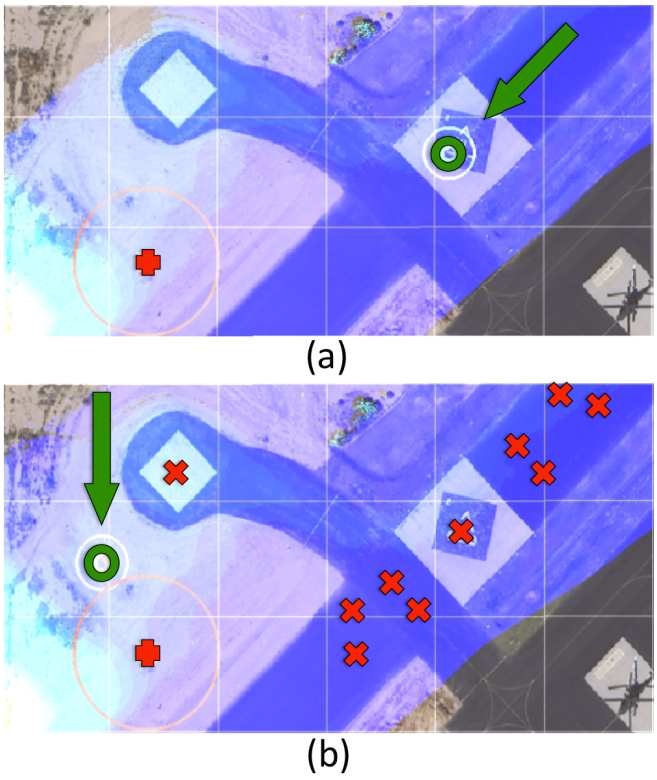

Figure 13: Example of chosen landing sites during a test series. (a) shows the landing decision and approach direction (green arrow and circle) for a casualty (red cross) located on sloped terrain but near large flat areas. (b) shows the result of having large areas covered with debris (denoted with red x's. See Figure 6 for close-up). The system lands from the north to avoid a structure that would occlude the glide slope from the North East.

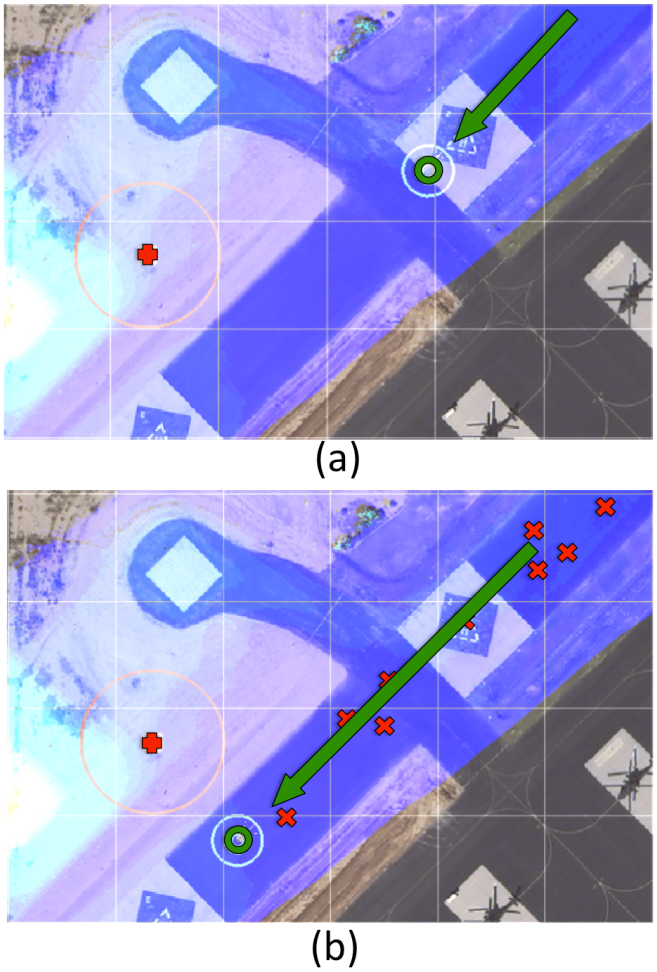

Figure 14: Example of chosen landing sites when the system is constrained to approach only through measured terrain. Similar to the tests in Figure 12, the system initially chooses the large open area near the helipad. However, when faced with clutter, it lands to the south of the casualty because this is the only clear area that has an obstacle-free glide slope available.

OBSTACLE AVOIDANCE

The obstacle avoidance tests demonstrated the ability of the system to detect and react to obstacles in real-time while flying at low altitude with the sensor in forward-scan mode. We did a total of five obstacle-avoidance runs against a real obstacle.

In these tests, the aircraft flew towards a GPS waypoint goal at approximately 20-meters altitude AGL. Unbeknownst to the system, we extended a 20-meter manlift crane directly into the flight path needed to achieve this waypoint. The system would need to detect the obstacle, recognize the threat, and change course to avoid a collision and still arrive at the goal point

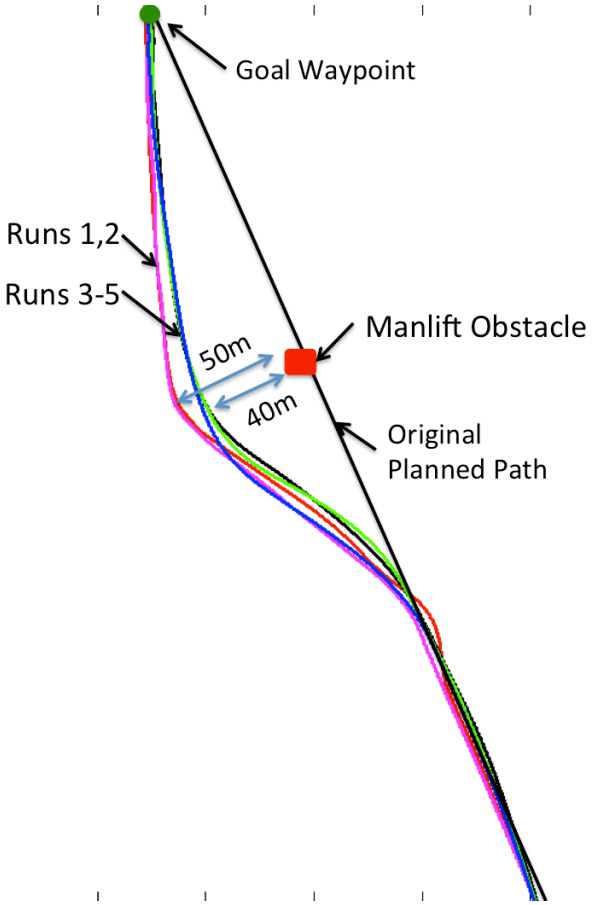

Figure **15**: Five obstacle-avoidance runs against an unmapped obstacle. Runs 1 and 2 were below 10 knots groundspeed, and runs 3-5 were at 21 knots. (Top-down view)

Figure **9** shows a snapshot of the data during one of these runs. Here, the manlift obstacle has been detected, a new route has been planned, and the helicopter has successfully avoided the obstacle and is arriving at its waypoint destination.

Safety was the primary concern during these tests, so we designed a test workup that allowed us to gradually verify the system without a large threat to the aircraft and crew. The first three runs were done with an altitude approximately 5 meters higher than the obstacle (25 meters AGL) – close enough that the system would still respond with a lateral maneuver, but high enough to miss the manlift if the vehicle did not respond in time.

The first two runs at 25 meters AGL were flown below 10 knots ground speed to verify that everything worked as planned. Figure **15** shows the flight traces from these runs. At low speed, the vehicle avoids the obstacle by about 50 meters (45 meters from edge of rotors). The third run was done at 21 knots ground speed at 25 meters AGL. The system responded in time and avoided the manlift by 40 meters (35 meters from edge of rotors).

With plenty of margin, we lowered the altitude to just below the manlift height, approximately 18 meters AGL. Test runs 4 and 5 were run at 21 knots ground speed. Again the system responded appropriately and avoided with 35 meters of clearance to the rotor blades. (Figure **15**)

Our simulations show that the system would continue to safely avoid the obstacle at speeds up to 40 knots. The 40-knot limit may be acceptable for obstacle detection during a slow final approach, but sensors with longer range will be needed for low/fast flight.

CONCLUSIONS

We have shown a perception and planning system that is capable of guiding a full-scale helicopter in previously unmapped environments. The system successfully completed 8 landing missions with varied obstacle configurations. We also demonstrated 5 obstacle avoidance runs at speeds up to 21 knots against a tower-style obstacle. Simulations showed that 40 knots would be attainable with the current hardware. Longer-range ladar units will be needed for faster flight near the ground.

Our results show that high-accuracy inertial measurement and precise pointing registration allow a ladar system to create maps at the higher speeds and longer ranges needed by full-scale aircraft. While the full system was demonstrated successfully, the airspace constraints of the test environment precluded thorough methodical testing of all of the capabilities of the system. For example, aborted landings due to changing obstacle configurations were not tested. The effects of other terrain and casualty locations should also be investigated. Future work would involve testing in a less restrictive air space.

Obstacle avoidance was shown to work well in concept, but the limitations of the interface to the autopilot on the aircraft limited the types of avoidance maneuvers to simple swerving-style trajectories. A UAS design that will incorporate a perception system should have a flight guidance system that is designed to accept a continuously-updated trajectory from the planning/perception unit.

ACKNOWLEDGMENTS

Our team would like to thank TATRC (Telemedicine and Advanced Research Center) from the US Army Research and Materiel Command for funding this research. Our thanks go to our partners at Piasecki Aircraft Corp, whose expertise and financial commitment allowed us to test on a full-scale vehicle. The Boeing Company provided use of the Unmanned Little Bird and the test facilities in Mesa, AZ. We thank the ULB flight test team for making the integration and testing a safe and productive experience.

WORKS CITED

Whalley, M., Directorate, A., Schulein, G., & Theodore, C. (2008). Design and Flight Test Results for a Hemispherical LADAR Developed to Support Unmanned Rotorcraft Urban Operations Research. *American Helicopter*

Society 64th Annual Forum, Montreal, Canada, April 29 - May 1 .

Whalley, M., Freed, M., Harris, R., & Takahashi, M. (2005 йил 1-Jan). Design, Integration, and Flight Test Results for an Autonomous Surveillance Helicopter. *Proceedings of the AHS International Specialists' Meeting on Unmanned Rotorcraft* .

Barber, D., Griffiths, S., McLain, T., & Beard, R. (2005). Autonomous Landing of Miniature Aerial Vehicles. *Proceedings of the AIAA Infotech@Aerospace Conference* .

Bosch, S., Lacroix, S., & Caballero, F. (2006 йил 1-Oct). Autonomous Detection of Safe Landing Areas for an UAV from Monocular Images. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* , 5522 - 5527.

Johnson, A., Klumpp, A., Collier, J., & Wolf, A. (2002). Lidar-based hazard avoidance for safe landing on Mars. *AIAA Journal of Guidance Control and Dynamics , 25* (6), 1091-1099.

Johnson, A., Montgomery, J., & Matthies, L. (2005 йил 31-Mar). Vision Guided Landing of an Autonomous Helicopter in Hazardous Terrain. *Proceedings IEEE International Conference on Robotics and Automation* , 3966 - 3971.

Martin, M., & Moravec, H. (1996). *Robot Evidence Grids.*

Scherer, S., Chamberlain, L., & Singh, S. (2010 йил 10-Feb). Online Assessment of Landing Sites. *AIAA Infotech@Aerospace 20-22 April 2010, Altanta, Georgia* .

Scherer, S., Ferguson, D., & Singh, S. (2009). Efficient C-Space and Cost Function Updates in 3D for Unmanned Aerial Vehicles. *International Conference on Robotics and Automation.* IEEE.

Scherer, S., Singh, S., Chamberlain, L., & Elgersma, M. (2008). Flying Fast and Low Among Obstacles: Methodology and Experiments. *The International Journal of Robotics Research , 27* (5), 549-574.

Sprinkle, J., Eklund, J., & Sastry, S. (2005 йил 31-May). Deciding to land a UAV safely in real time. *Proceedings of the American Control Conference (ACC)* , 3506 - 3511 vol. 5.

Templeton, T., Shim, D., Geyer, C., & Sastry, S. (2007 йил 31-Mar). Autonomous Vision-based Landing and Terrain Mapping Using an MPC-controlled Unmanned Rotorcraft. *Proceedings IEEE International Conference on Robotics and Automation* , 1349 - 1356.

Tierney, S., & Langelaan, J. (2010 йил 15-Apr). Autorotation Path Planning Using Backwards Reachable Set and Optimal Control. *American Helicopter Society Forum* .

Tsenkov, P., Howlett, J., Whalley, M., Schulein, G., Takahasi, M., Rhinehart, M., et al. (2008 йил 4-Sep). A System for 3D Autonomous Rotorcraft Navigation in Urban Environments . *AIAA 2008* , 23.