Up, up, and away...

UAV's Endurance Gets a Lift by Latching onto Thermals

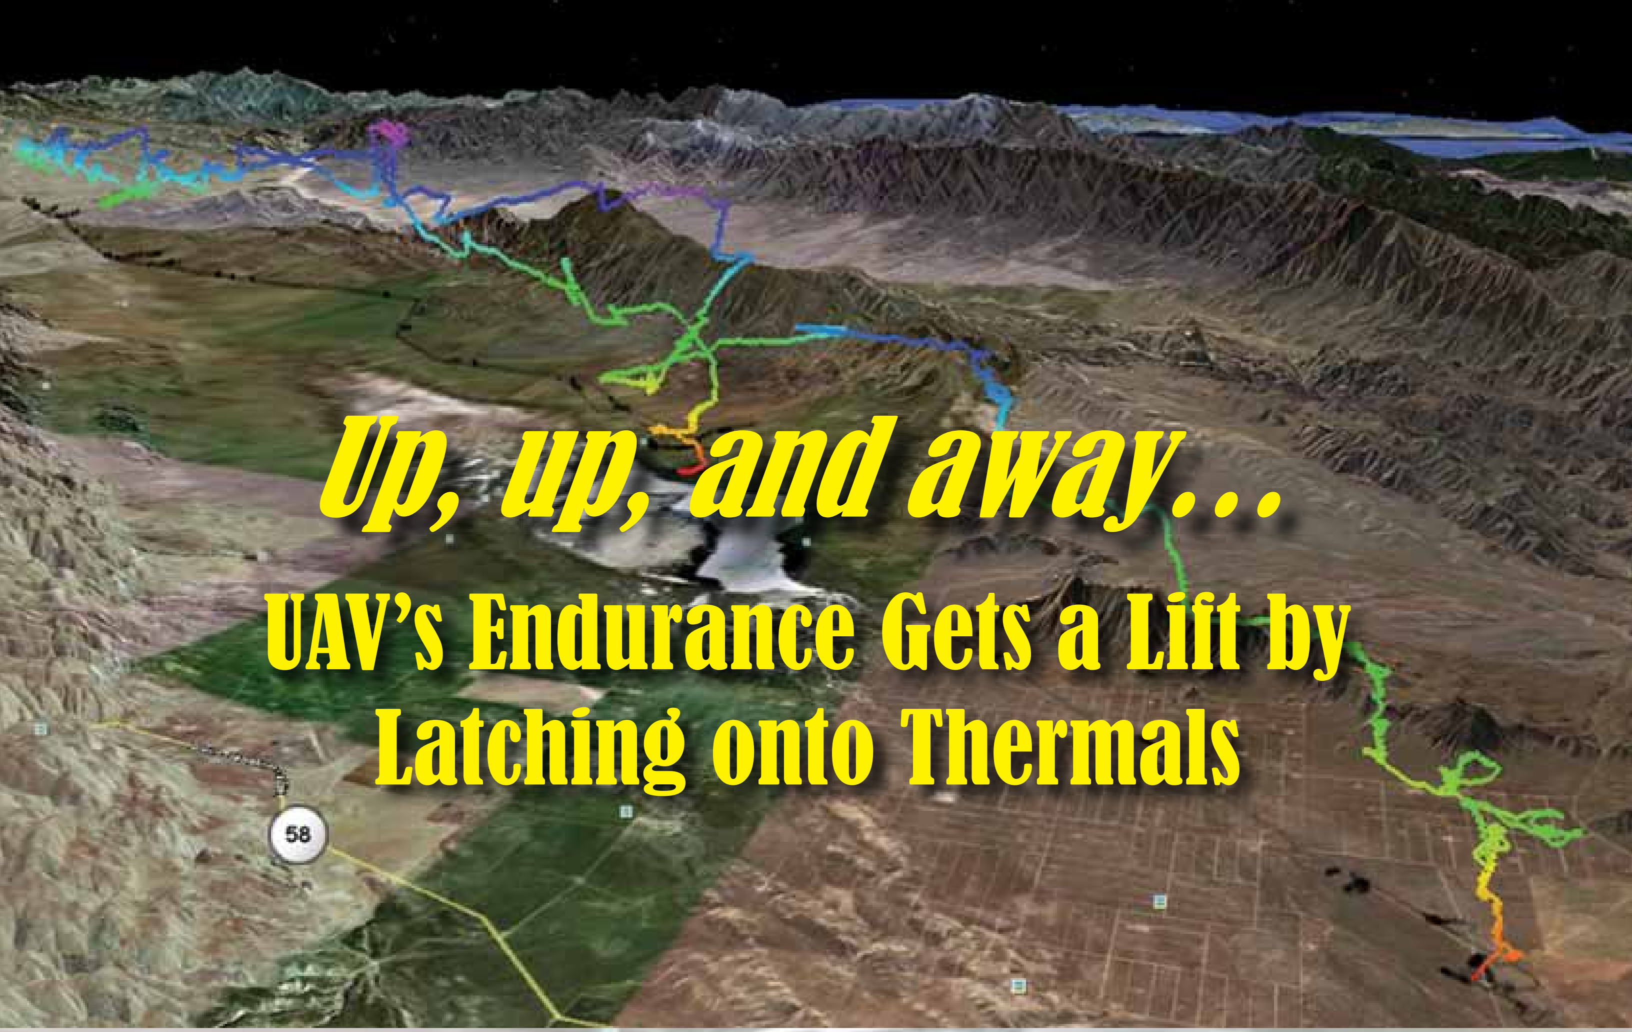

By using the naturally occurring air currents known as thermals, unmanned aerial vehicles (UAVs) can dramatically extend their operational endurance. NRL researchers demonstrated autonomous soaring using an unpowered, winch-launched unmanned sailplane, which soared within thermals to fly a notional road-monitoring mission for 5.3 hours beyond its nominal 3-minute capability. Since the autonomous soaring technology can be implemented as a software upgrade to existing UAV autopilots, NRL's Tactical Electronic Warfare Division expects that other UAVs such as the hydrogen-powered Ion Tiger will soar to new heights and longer flight times.

Caption: Using autonomous soaring algorithms, an unpowered sailplane flew a 113.5 km course in 4.5 hours.

Autonomous Soaring for Unmanned Aerial Vehicles

D.J. Edwards and M. Hazard
Tactical Electronic Warfare Division

NRL researchers have developed a way to extend the endurance of unmanned aerial vehicles (UAVs) by exploiting the naturally occurring rising air currents known as thermals. Unlike a typical UAV autopilot, which naively attempts to reject these atmospheric "disturbances," the new system adapts its flight path to maximize energy extraction. The system detects the presence of rising air using onboard sensors, finds the center of the thermal, and redirects the guidance system to circle in the thermal. These algorithms were incorporated into an unmanned sailplane, a winch-launched aircraft with a nominal endurance of only 3 min. With autonomous soaring engaged, the unpowered vehicle was able to fly for more than 5.3 h. Because implementing autonomous soaring might be as simple as a software upgrade — UAVs already carry the necessary sensors for thermal detection — this enabling technology could be widely deployed at low cost. Our current research focuses on extending these benefits to powered UAVs such as the hydrogen fuel cell powered Ion Tiger. This article presents an overview of the autonomous soaring algorithms, flight test results, and potential directions for continued research.

INTRODUCTION

The increasing use of small unmanned aerial vehicles (UAVs) in military and civilian applications has been accompanied by a growing demand for improved endurance and range. These demands have been typically addressed by improvements in aerodynamic and structural efficiency, improved fuel-efficient propulsion systems, and the ongoing miniaturization of onboard computer and payload systems. Recently, more attention has been given to the extraction of energy from the atmosphere. Aircraft can make use of atmospheric updrafts, or thermals, to gain altitude without expenditure of onboard fuel stores. By intelligently tracking thermals, an unmanned aircraft can extend its range or endurance without carrying additional fuel or specialized sensors.

However, current UAV autopilots attempt to reject thermals and other atmospheric perturbations. Enabling autonomous soaring required the development of a thermal identification framework that could efficiently compare autopilot sensor data to a thermal model, allowing the autopilot to determine when it is flying in a thermal and choose the best flight path to maximize energy extraction. This basic soaring functionality was wrapped with a mission management algorithm that maintained the balance between soaring and timely arrival at the required destination. Finally, the thermal identification and mission management algorithms were implemented on a UAV sailplane for flight testing. This article gives an overview of thermal sensing and identification, soaring guidance, and flight testing results.

THERMAL UPDRAFTS

Thermal updrafts are naturally occurring atmospheric convection currents that are created by uneven heating of the Earth's surface. Radiant solar energy is absorbed by terrain and transferred to the air. If a parcel of air becomes warmer than its surroundings, it will rise like a bubble in a boiling pot of water, creating a thermal. Each parcel of warm air in the thermal continues ascending (and expanding) until it cools to the temperature of the surrounding air and is no longer buoyant. Depending on atmospheric conditions, the top of the thermal may be hundreds or thousands of meters above the surface. Thermal formation and evolution are influenced by terrain texture and albedo, time of day, and ambient meteorological conditions. These factors make accurately modeling thermal activity a difficult task. Computational fluid dynamics (CFD) tools such as large-eddy simulation could be used to model thermals with high fidelity, but their computational cost makes them ill-suited for real-time simulation or control.

Instead, a simpler model that approximates the vertical wind profile of a single thermal is used. To reduce the complexity of the model, several assumptions are made: a thermal is radially symmetric about a center point, does not create horizontal winds, drifts with the ambient wind velocity, and is invariant with altitude and time. Figure 1 shows the modified three-dimensional Gaussian distribution chosen to represent vertical wind speeds in a thermal. Its main features are the characteristic "core" region of rising air and a ring of sinking air around the core. This reduced-order model is based on four parameters: the two-dimensional center location, peak updraft strength, and characteristic radius.

The simplified model is one component of the autonomous soaring framework. We also need some way to measure the instantaneous vertical wind speed, and a systematic way to update the estimated thermal location based on vertical wind speed measurements.

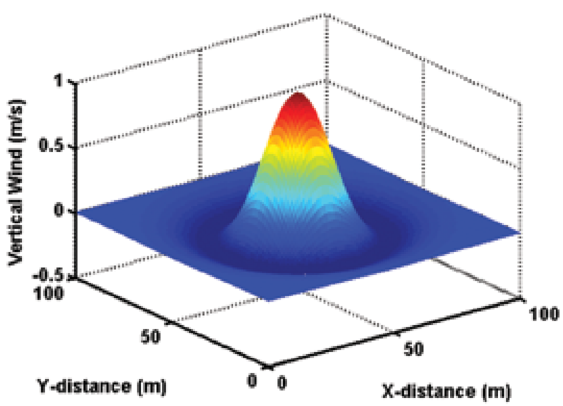

FIGURE 1
Modified Gaussian distribution updraft cross-section.

SENSING THERMALS

The updraft model represents a distribution of vertical wind motion, which is not easy to measure directly from a moving aircraft. Instead, we examined the problem from an energy/power perspective. In general, the total energy of the aircraft system is conserved aside from work done on the aircraft by external forces: thrust, drag, and any atmospheric contributions. Known contributors such as thrust and drag can be directly compensated based on wind tunnel and flight test data, leaving an "atmospheric power" signal. If thermals are the only atmospheric contribution, this power measurement is equivalent to vertical wind speed measurement and is compatible with our low-order model.

Accurate measurement of the aircraft's specific power state can be achieved using the same suite of sensors carried in a typical autopilot. Initial investigations computed total energy based on altitude and airspeed measurements, and then numerically differentiated this signal to estimate the instantaneous power. The differentiation step reduced the quality of the data, especially with the low-resolution sensors typical of small UAVs. In our improved system, the power state is computed directly using the onboard three-axis accelerometer, GPS velocity, and GPS position measurements. This alleviates the need to differentiate and significantly improves the data quality. The power signal is then compensated for drag effects using experimental sink polar data and for thrust effects using wind tunnel propeller efficiency data. Any remaining power contribution, whether positive or negative, is assumed to come from the atmosphere. Figure 2 shows the compensated specific power signal for a representative flight path that intersects a thermal.

In Fig. 2, warmer colors show faster vertical air motion and cooler colors show little or no vertical air motion. Qualitatively, we can see a specific power distribution that matches key components of the updraft model, including a circular shape around a local peak. These measurements can be correlated with the low-order updraft model to estimate the thermal parameters that best fit the flight data.

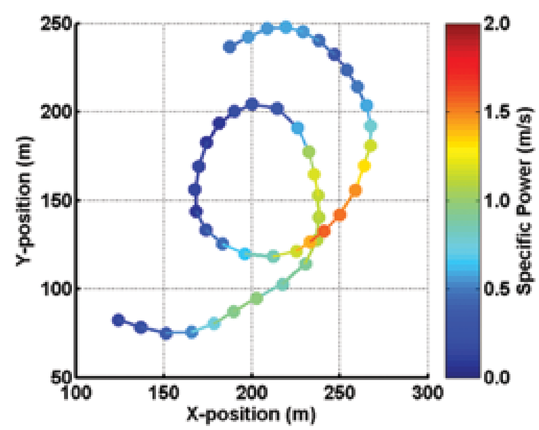

FIGURE 2
Measurement of specific power distribution of a thermal.

THERMAL PARAMETER IDENTIFICATION

Given vertical wind speed measurements, the next component is an algorithm to reduce the flight data into parameters that can be used to guide the aircraft. Our research investigated two options: batch methods, which accumulate sensor time history and iteratively solve for the model parameters; and recursive methods, which only use the most recent data to quickly update a previous parameter estimate. In general, recursive methods trade absolute accuracy for lower computational burden. Reducing computational burden is important to stay within the weight and power requirements of a small UAV.

The batch method is based on a similar approach developed by John Wharington that used an adaptive neural network.[1] To reduce computational workload for real-time use, the full method was reduced to an iterative nonlinear regression. In short, the method begins with a candidate center location node and computes the residual error associated with a two-dimensional fit in the remaining parameters (strength and radius) at that location. The residual is similarly evaluated at nearby locations and the candidate location with the smallest error is chosen as the winner. A new set of candidate locations is drawn in a smaller ring around the winner, and the process repeats until the algorithm converges on a location with stronger correlation to the model. While relatively computationally time consuming, this method of searching provides a robust means of multiparameter nonlinear curve fitting. In Fig. 3, several steps of the process are shown, starting from an initial location within the flight path and moving down the error surface toward the global minimum where the model best fits the measured data. The step sizes were purposely limited to serve as a low-pass filter on the center estimation.

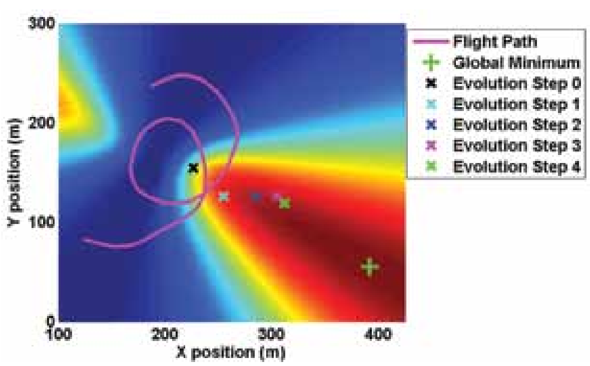

FIGURE 3
Evolutionary search method moving toward the global minimum.

A recursive unscented Kalman filter (UKF) method was also investigated to reduce the computational burden of the batch method. Kalman filters are a class of algorithms that track a set of parameters along with a covariance matrix that describes the filter's confidence in its parameter estimate. As each new measurement is received, the accuracy of the new measurement is weighed against the covariance matrix to compute the optimal balance between old data already fused into the parameter estimate and new measurements. In this work, we used the UKF, an extension of the classic linear Kalman filter, to allow using a fully nonlinear measurement model. The UKF takes a fixed time to complete at each time step, while the execution time of the batch method varies on the conditioning of the data set. The estimation accuracy of each algorithm is compared in Fig. 4, using the mean squared error in all four thermal model parameters as the metric. The batch method generally showed slightly lower error than the UKF, but both provide sufficient accuracy for autonomous soaring.[2] Simulations comparing the batch and recursive thermal identification algorithms showed the recursive method was between one and two orders of magnitude faster than the batch method.

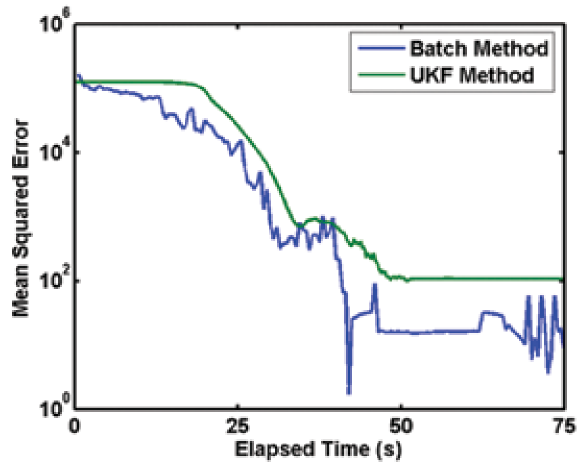

FIGURE 4
Comparison of mean squared error for batch and recursive methods.

SOARING GUIDANCE

Once a thermal is detected and located, the aircraft needs to be steered to fly inside the updraft core. Fortunately, many commercial autopilots include the ability to "orbit" around a GPS waypoint. The autonomous soaring algorithm uses this existing controller instead of interacting with the sensors and control surfaces directly. In other words, the autopilot (inner feedback loop) is nested inside the soaring controller (outer feedback loop), as shown in Fig. 5. This form of decoupled control is common to many aircraft control systems. Chief among the advantages of this architecture is the ability of the inner loop to run at high rates, typically

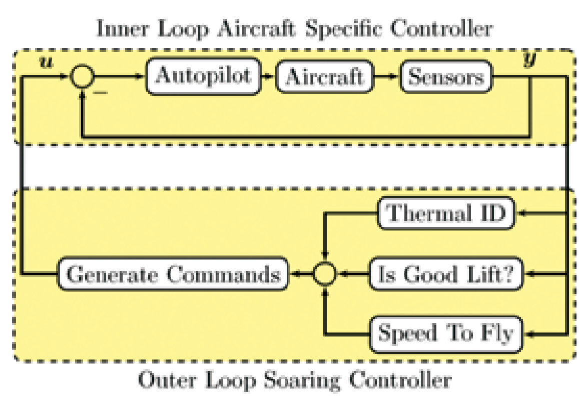

FIGURE 5
High-level autonomous soaring controller diagram.

50 Hz, while the outer loop may update more slowly — in the developed system, the orbit waypoint is updated at 2 Hz. The outer loop soaring controller also becomes autopilot and vehicle agnostic, making it easier and less expensive to adapt the soaring algorithms to a different unmanned vehicle.

MISSION MANAGEMENT

For an aircraft capable of autonomous soaring to be able to complete a useful mission, it must be able to balance progress along its commanded course with the possible endurance or range gains presented by thermals. At a very basic level, this requires flying efficiently along a course, stopping to investigate a thermal when it is encountered (latching), finding and soaring in the thermal, and returning to course (unlatching) if the energy gains are no longer worthwhile. An example of these behaviors in action is shown in Fig. 6. The system latches into soaring mode if the average vertical wind speed (thermal strength) exceeds a threshold, the thermal identification algorithm shows a good fit, and the vehicle is within mission limits and airspace boundaries. In soaring mode, the aircraft is commanded to fly an orbit around the thermal center as described in the previous section. The system unlatches if the average thermal strength drops below the threshold, with sufficient hysteresis to allow the system enough time to find and track the thermal after its initial engagement.

The thermal strength threshold for latching into soaring mode is determined by MacCready speed-ring theory. The speed ring is a manned sailplane instrument developed by the late Paul MacCready. It uses an

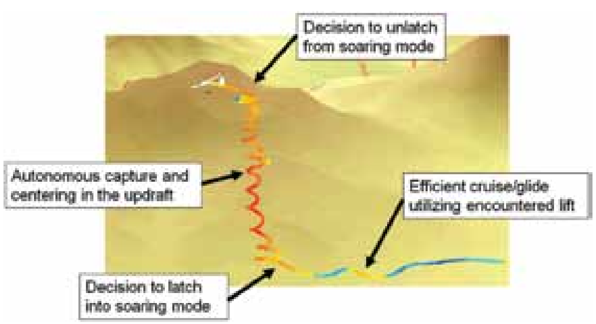

FIGURE 6
Soaring guidance modes (flight path color indicates specific power).

estimate of local average thermal strength to provide two pieces of information: the optimal cross-country speed to fly and the minimum thermal strength required to make stopping to soar worthwhile. In the test system, the speed ring setting is coupled with altitude. This ensures the aircraft flies conservatively at low altitudes, latching onto weaker thermals when gathering energy is crucial; as the vehicle gains altitude, thermals must get progressively stronger for latching to occur. Similarly, the aircraft flies faster at high altitudes to maximize its cross-country speed; the vehicle has sufficient altitude reserves, and the chance of intercepting a strong thermal improves with the distance traveled. At lower altitudes, the vehicle is commanded to fly slowly to conserve energy reserves (altitude) and maximize flight time to give a chance of finding even a weak thermal. Recently, this behavior has compared favorably with that of birds.[3]

IMPLEMENTATION

An off-the-shelf fiberglass composite sailplane served as a host vehicle for testing the soaring algorithms (Fig. 7). The Autonomous Locator of Updrafts (ALOFT) has a 4.2 m wingspan, a 5 kg takeoff weight, and a best glide slope of about about 24 m forward for every 1 m of descent. The aircraft is launched using a winch that hoists the vehicle like a kite to approximately 200 m altitude, giving a 3 min gliding descent when no thermals can be found.

In addition to the airframe and six electric control surface actuators, a few other items are needed to facilitate autonomous soaring. An off-the-shelf Piccolo II miniature autopilot from Cloud Cap Technologies provides the primary navigation solution and vehicle guidance. The autopilot communicates with a remote ground station via a 900 MHz spread-spectrum modem, giving operators the ability to monitor and retask the vehicle as necessary. A separate manual override circuit allows a safety pilot to completely circumvent the autopilot and fly the aircraft in remote control mode, providing a measure of redundancy in case of autopilot failure. An 87 Wh lithium-ion battery located in the aircraft nose provides power to the avionics and actuators for 8 h of continuous operation. ALOFT has no motor; therefore, any endurance beyond the nominal 3 min of flight time comes from the atmosphere.

FLIGHT TESTING

With the autonomous soaring algorithms installed on a UAV, the soaring algorithms were tested in a cross-country (XC) soaring scenario.[4] Cross-country soaring typically involves a "distance task" where the goal is to travel as far as possible on a single launch, or a "speed task" where the goal is to complete a predetermined course with the highest average speed. Satisfying these soaring tasks is a challenge for unmanned systems which, unlike a human pilot, do not have a priori knowledge of local thermal hotspots. Both tasks are similar to road-monitoring or communications relay missions.

FIGURE 7
ALOFT vehicle and internal hardware.

After completing a short in-air controls check following a manual launch, the autopilot followed a preprogrammed series of waypoints with the outer loop soaring controller managing the behavior of the vehicle. ALOFT soared in nearby thermals to reach a starting altitude of approximately 500 m, and then the mission task waypoints were activated to send the vehicle traveling on course. Along the way, the mission manager balanced its desire for altitude with the mission task of distance or speed. Once the task was complete or the vehicle breached a preset lower altitude limit, manual control was activated and the vehicle was landed to complete the test.

RESULTS

Over 170 flights were performed with the ALOFT glider, accounting for more than 70 h of flight time and 20 h of time spent actively soaring. Several milestones were reached over the course of the research program:

- September 30, 2007: First demonstration of autonomous cross-country soaring.
- October 5, 2008: Unofficially beat the world record for goal-and-return cross-country soaring (5 kg class), flying 48.6 km each way.
- May 23, 2009: Longest autonomous cross-country flight to date of 113.5 km over 4.5 h, for an average cross-country speed of 25.2 km/h.
- June 12, 2009: Highest endurance cross-country flight to date of 5.3 h.

Figure 8 shows the altitude time history from a distance task flight, showing how the algorithm moves in and out of soaring mode and operates in the altitude band of 500 m to 1500 m.

FUTURE RESEARCH DIRECTIONS

The success of the autonomous soaring algorithms has opened up a number of different research avenues. An ongoing effort supported by an NRL Karle Fellowship seeks to implement soaring algorithms on other existing unmanned aircraft systems with complementary endurance technology. Ion Tiger, the NRL-devel-

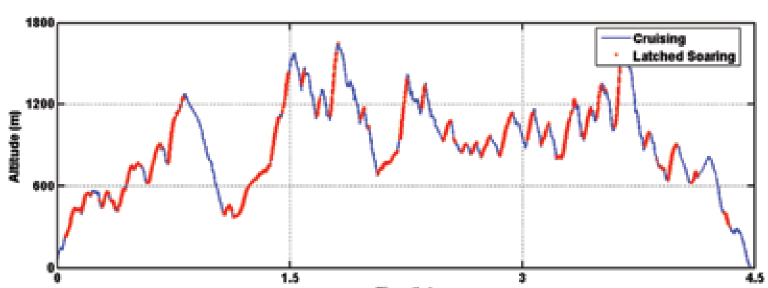

FIGURE 8
Altitude trace of a distance task flight.

oped ultralong-endurance hydrogen fuel cell powered aircraft, has demonstrated 26 h of endurance using gaseous hydrogen alone. Adding autonomous soaring capability to the system could entirely offset fuel consumption during times of convective activity, leading to further gains in endurance. Other research interests include improving the fidelity of thermal models and implementing the algorithms on micro-UAVs.

CONCLUSION

This project has culminated in the development of a system for autonomously locating and actively guiding a UAV to soar in thermals. In flight tests, an unpowered glider with a 3 min nominal endurance flew for more than 5 h in a realistic cross-country scenario, extending its flight time by autonomously extracting energy from ambient air motion. Any endurance improvement is of direct tactical benefit to many UAV missions. Because no additional hardware is required to enable autonomous soaring, the technique could be applied to many existing unmanned aircraft systems.

ACKNOWLEDGMENTS

The authors would like to thank C. Bovais, A. Propst, and L. Silverberg for their contributions.

[Sponsored by NRL]

References

[1] J. Wharington, "Autonomous Control of Soaring Aircraft by Reinforcement Learning," Ph.D. thesis, Royal Melbourne Institute of Technology, Melbourne, Australia (1998).

[2] M. Hazard, "Unscented Kalman Filtering for Real-Time Atmospheric Thermal Tracking," M.S. thesis, Mechanical and Aerospace Engineering Dept., North Carolina State University, Raleigh, NC (2010), http://www.lib.ncsu.edu/resolver/1840.16/6353.

[3] Z. Ákos, M. Nagy, and T. Vicsek, "Thermal Soaring Flight of Birds and Unmanned Aerial Vehicles," *IOP Journal of Bioinspiration and Biomimetics* **5**(4), 045003 (2010).

[4] D. Edwards and L. Silverberg, "Autonomous Soaring: The Montague Cross Country Challenge," *AIAA Journal of Aircraft* **47**(5), 1763–1769 (2010).

■

THE AUTHORS

DANIEL EDWARDS received his Ph.D. in aerospace engineering from North Carolina State University in 2010. He started at NRL as a SCEP student with the Spacecraft Engineering Department in 2003 and transferred to the Tactical Electronic Warfare Division in 2005. He was hired full-time after graduation in May 2010. His recent research and graduate work have concentrated on extending the endurance of unmanned aerial vehicles by using naturally occurring vertical wind motion. His ongoing work includes experiments combining atmospheric energy extraction techniques with advanced hydrogen fuel cell propulsion systems and developing LIDAR technology for remote sensing of vertical wind currents.

MATTHEW HAZARD received his M.S. in aerospace engineering from North Carolina State University in 2010, and has been working at NRL since 2008. His research interests include unmanned aircraft guidance, navigation, and controls, alternative methods for UAV launch and recovery, and cooperative UAV operations. He is the principal investigator for the NRLEAP 6.2 project, which is exploring autonomous wingtip-to-wingtip docking as a technique to improve the endurance of small UAVs.