Optimal Persistent Surveillance using Coordinated Soaring

Dmitriy Makovkin[*]

Jack W. Langelaan[†]

The Pennsylvania State University, University Park, PA 16802, USA

Coordinated soaring by a flock of small unmanned aerial vehicles (sUAVs) provides a means of conserving fuel while performing aerial tasks. The ability to exploit thermal columns in the atmospheric boundary layer allows sUAVs to remain airborne without expending any onboard sources of energy, i.e., soaring flight. This paper presents an analysis of the cruising phase during coordinated soaring where a flock of sUAVs relies on thermal exploitation to maximize endurance for monitoring-type missions. To this end, a maneuver is investigated that involves each sUAV repeating a round-trip between a thermal and a monitoring/surveillance target so as to maintain continuous monitoring of the target. The focus is on minimizing the number of agents required to maintain continuous, persistent surveillance of the target for given atmospheric conditions (thermal strength and distance between the thermal and monitoring target) and on maximizing a free parameter (time or distance) when the number of agents is specified. It will be shown that the optimal cruising speed for maximizing the endurance of monitoring-type missions varies between the best L/D speed and the MacCready speed and depends on the "aggregate thermal strength" of a given cycle, or equivalently, the ratio of the time that one sUAV spends away from the target to the time that it spends at the target. An examination of multiple-thermal exploitation is then presented, followed by an evaluation of the flight simulations used to support the results.

Nomenclature

A, B	Auxiliary variables for the optimal airspeed equations
a, b, c	Coefficients for the quadratic fit of a sink polar
d	Distance between a thermal and monitoring target in single-thermal exploitation
d_c	Combined length of the first and final leg in multiple-thermal exploitation
d_{free}	Free exploration distance in multiple-agent scenarios
E	Exploration job
Δh	Global working altitude
L/D	Lift-to-drag ratio
N	Number of agents required for zero-altitude-loss continuous monitoring
$\bar{N}$	Corrected N accounting for transition $T \to M$
$\lceil N \rceil$	Integer number of agents required for zero-altitude-loss continuous monitoring
M	Monitoring job
s	Cruising sink rate
$s_{(L/D)}$	Sink rate from cruising at $v_{L/D}$
s_{min}	Minimum sink rate of an aircraft
s_s	Monitoring sink rate of an aircraft
T	Expected climb rate in a thermal
t_{away}	Time spent away from the monitoring target
t_c	Cruising duration

---

[*]Graduate Student, Department of Aerospace Engineering, Student Member AIAA.

[†]Associate Professor, Department of Aerospace Engineering, Associate Fellow AIAA.

American Institute of Aeronautics and Astronautics

t_{free}	Free loitering time in multiple-agent scenarios
t_m	Monitoring duration
t_t	Thermalling duration
t_{tm}	Transient maneuver duration
v	Cruising airspeed
v_c	Optimal cruising airspeed for single-agent scenarios
$v_{L/D}$	Best L/D airspeed
v_{Mc}	MacCready speed
v_{min}	Cruising airspeed for s_{min}
v_{opt}	Optimal cruising airspeed for multiple-agent scenarios
w	Speed of sinking airmass

I.   Introduction

SMALL unmanned aerial vehicles (sUAVs) excel at tasks unsuitable for their larger counterparts. During the last 15 years, there has been an increasing interest in thermal exploitation strategies and controllers as a means of mitigating the limited on-board fuel capacity of sUAVs.[1–6] One area yet to be investigated is the cruising stage during monitoring-type missions where a sUAV travels repeatedly between a thermal column and a monitoring target. Here, this topic is explored with a focus on maximizing airborne time, i.e., endurance, for missions requiring continuous uninterrupted monitoring of a target.

Thermal columns, or "thermals" for short, are ascending masses of air that act as the primary convective units of the atmospheric boundary layer by equilibrating the temperature and pressure gradients between the ground and atmosphere.[7] Some typical defining characteristics of thermals include the following: diameters from less than 100 m to nearly 1000 m, heights of nearly 1500 m, vertical air currents ranging from 1 to 4 m/s, and lifespans between 5 and 30 minutes.[8] These values, however, depend strongly on time of day and season; thermal activity peaks between the hours of noon and sunset and the months of spring and summer when the ground can become significantly warmer than the surrounding air.

Previous studies have established the effectiveness of thermal exploitation for both single aircraft and coordinated aircraft scenarios. Allen presents two important facts for a single soaring aircraft. First, during summer and winter, convective lift (thermals) can give a 12 and 6 hour increase in endurance, respectively, for an electric-powered sUAV with a nominal endurance of two hours.[5] Second, "performance increase has a low sensitivity to many key simulation parameters including, aircraft glide slope, number of updrafts, updraft lifetime, updraft velocity, and height-above-ground upper limit."[5] Depenbusch demonstrates that, on average, with the goal of remaining airborne for as long as possible, a single sUAV doubles its endurance when keeping a memory of discovered thermal locations, thermal strengths, and associated covariances.[6] He also shows that, compared to a single sUAV using thermal mapping, a 23% average increase in endurance, defined as the soaring time until one sUAV reaches the ground, is provided by two coordinated sUAVs that explore the environment and share measurements with each other. Furthermore, four coordinated sUAVs experience a 98% average endurance increase over the two coordinated sUAVs.

Persistent monitoring soaring scenarios confine the aircraft to a specific region of a few square kilometers around a monitoring target. Examples of monitoring targets are areas that require military surveillance, search-and-rescue missions, or meteorological studies such as pollution studies, wind-gust modeling, and weather monitoring. Bethke et al. have examined persistent surveillance by a team of quad-rotors, where periodic recharge at known "filling stations" is performed as part of the mission.[9–11] Cutler et al. examine surveillance with a single UAV that exploits ridge lift, and develop a "seeability" metric as a means of quantifying video image quality during search and rescue missions.[12] When a flock of sUAVs is used in conjunction with thermal lift, one aircraft is assigned to monitor the target while the other aircraft explore and exploit the surrounding thermals, creating a map of thermal lift while maintaining altitude. In this paper, the exploration stage is disregarded and it is assumed that the strengths and locations of thermals are known. This results in a scenario where the aircraft are engaged in a cycle of monitoring a target, cruising to a thermal, exploiting a thermal, and cruising back to the target in a way that keeps at least one agent at the target at all times, resulting in continuous monitoring of the target. This scenario is restricted further by requiring each cycle to begin at a fixed altitude, thus, if the conditions were theoretically constant, the mission endurance would be infinite. In this fashion, focus is placed on maximizing mission endurance by

optimizing the cruising phase of continuous monitoring missions. It is important to note that the altitude and continuous monitoring constraints allow one to fully define the scenario and formulate governing equations; it will become clear that the results in this paper are not restricted to scenarios with these constraints.

The various monitoring scenarios will be compared with each other by the quantification of the number of agents needed to achieve continuous or persistent monitoring in a given scenario. The optimal cruising airspeed, as well as the optimal airspeed held between thermals during multiple-thermal exploitation, can be found by differentiating this equation, where the resulting speed minimizes the number of agents required for a persistent monitoring cycle. Minimizing the number of agents is a way of minimizing the amount of resources required for a given cycle, thus maximizing the efficiency of the scenario. It will be shown that this is equivalent to minimizing the ratio of the amount of time an agent spends away from the monitoring target to the amount of time an agent spends at the monitoring target. Minimizing this ratio is of interest in many types of scenarios involving one or more agents cruising between a thermal and some location requiring an extended presence. Instead of cruising between one monitoring target and one or more thermals, it may be necessary to alternate between multiple monitoring targets, perimeter surveillance for instance, or between exploration targets. The optimal airspeeds in these scenarios will take the same form because the objective is ultimately the same in each case: to gain altitude as quickly as possible without compromising the time spent at a target. In spite of this, as the analysis proceeds, one will see why this approach does not apply to all scenarios and this will lead to the more general concept of the "aggregate thermal" of a given cycle, for which the MacCready speed will prove to be optimal.

The remainder of this paper is organized as follows: Section II establishes the scenario and the simplifying assumptions; Section III develops the theory for optimizing the cruising phase of monitoring missions involving one agent , Section IV expands the theory to persistent monitoring missions where multiple agents are involved; Section V presents simulation results from the Silent Wings Soaring Simulator; and Section VI contains concluding remarks.

II. Problem Definition

The basic scenario is illustrated in Figure 1. A flock of sUAVs, hereafter referred to as "agents," is engaged in a cycle of monitoring a target and exploiting a thermal. Once an agent loses enough altitude at the monitoring target, it cruises to a thermal to regain altitude before cruising back to the target to resume monitoring. Continuous or persistent monitoring becomes possible when a new high flying agent immediately replaces the previous low flying monitoring agent that has departed to regain altitude. An altitude constraint is imposed to fully define the problem: altitude must be conserved in each cycle, that is, each cycle must begin at the same altitude.

(a) Single-thermal exploitation

(b) Multiple-thermal explotation
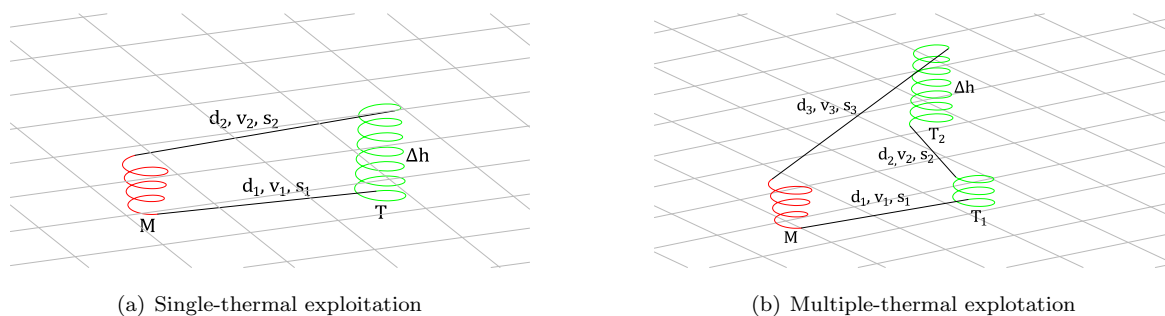
Figure 1. Persistent surveillance cycles. Single-thermal case shown on left; multiple-thermal case shown on right.

Given a thermal's strength and distance from a target, the problem is to compute the minimum number of agents required for continuous, persistent monitoring of the target. For this, a number of environmental assumptions are beneficial in simplifying the analysis:

1. Wind is absent in the environment.

2. Thermal exploitation is constant with altitude.

American Institute of Aeronautics and Astronautics

3. The altitude floor is constant.

4. The altitude ceiling is constant.

5. All agents have identical properties and capabilities.

6. Atmospheric conditions persist indefinitely.

7. The duration of transient maneuvers is negligible compared with other phases of flight.

8. Changes in altitude during transient maneuvers are negligible for sUAVs.

This leaves a scenario where thermalling and monitoring dynamics are absent and atmospheric conditions are constant. This allows one to focus entirely on the cruising phase during persistent monitoring missions without altering the significance of the analytical results. Each of these assumptions will be lifted after the introductory theory is established.

III. Minimizing the Number of Agents

A. Single-Thermal Exploitation

Intuitively, the number of agents required for persistent monitoring is equal to the time required for one agent to complete the circuit divided by the time that one agent spends monitoring the target:

$$N = \frac{t_c + t_t + t_{tm} + t_m}{t_m} \approx \frac{t_c + t_t}{t_m} + 1 \tag{1}$$

where t_c, t_t, t_m, and t_{tm} are the cruising, thermalling, monitoring, and transient maneuver durations, respectively. Note that t_{tm} represents the time required to transition between cruising and monitoring or between thermalling and cruising and vice versa. The duration of transient maneuvers will be neglected due to its insignificance when compared with t_c, t_t, and t_m.

Substituting the system parameters and applying the environmental assumptions gives

$$N = \frac{\frac{d}{v_1} + \frac{d}{v_2} + \frac{\Delta h}{T}}{s_s^{-1}\left(\Delta h - s_1\frac{d}{v_1} - s_2\frac{d}{v_2}\right)} + 1 \tag{2}$$

with distance d, expected thermal climb rate T, working altitude Δh, surveillance or monitoring sink rate s_s, the cruising airspeed and sink rate toward the thermal v_1 and s_1, and the cruising airspeed and sink rate toward the monitoring target v_2 and s_2. The working altitude is assumed to be constant and equal to the altitude ceiling minus the altitude floor for the agents. It is temporarily assumed that the two cruising paths are of the same distance d. The general case where the two cruising paths differ, for example, when an aircraft diverts to explore the surroundings or when multiple thermals are exploited, will be treated in section B. To clarify equation (2), the sum of the first two terms in the numerator represents the cruising duration of one agent for a given thermal with an expected climb rate T and a distance d away from the target. Cruise airspeeds v_1 and v_2 are not assumed to be equal since it may be optimal to reach a thermal by cruising at one speed, e.g., the MacCready speed v_{Mc}, whereas cruising back toward the monitoring target may be optimized by cruising at another speed, e.g., the best L/D speed $v_{L/D}$. Note that, in short, the MacCready speed is the cruising airspeed that minimizes the time that an aircraft takes to reach the top of a thermal given an estimated thermal climb rate. The MacCready speed is independent of distance and is used by glider pilots during cross-country soaring races to cruise between thermals.[13] The third term in the numerator of equation (2) represents the time spent exploiting the thermal, t_t. For the time being, it is assumed that minimizing N, i.e., maximizing efficiency, translates to exploiting the thermal over the entire working altitude Δh, i.e., not departing the thermal prior to reaching the altitude ceiling and not arriving at the thermal above the altitude floor; this will be proved in the following paragraph. Lastly, the denominator of equation (2) represents the time spent monitoring the target t_m at the aircraft's minimum sink speed.The simplicity of t_t and t_m amounts to an elevator model of thermalling and monitoring: once an agent reaches either location, he sinks or rises at a constant rate until departing.

To prove that the number of agents required for monitoring the target is minimized when Δh is maximized, start with the inequality $N_1 < N_2$, where N_1 represents a scenario with Δh_1 and N_2 represents an identical

scenario but with Δh_2 such that $\Delta h_1 > \Delta h_2$. The goal is to reduce this inequality to show that $N_1 < N_2$ if and only if $\Delta h_1 > \Delta h_2$:

$$\frac{\frac{d}{v_1} + \frac{d}{v_2} + \frac{\Delta h_1}{T}}{s_s^{-1}\left(\Delta h_1 - s_1\frac{d}{v_1} - s_2\frac{d}{v_2}\right)} + 1 < \frac{\frac{d}{v_1} + \frac{d}{v_2} + \frac{\Delta h_2}{T}}{s_s^{-1}\left(\Delta h_2 - s_1\frac{d}{v_1} - s_2\frac{d}{v_2}\right)} + 1 \tag{3}$$

Multiplying both sides by the product of the denominators, expanding, and canceling like terms results in

$$s_s\Delta h_2 d\left(\frac{1}{v_1} + \frac{1}{v_2}\right) - \frac{\Delta h_1 d}{T}\left(\frac{s_1}{v_1} + \frac{s_2}{v_2}\right) < s_s\Delta h_1 d\left(\frac{1}{v_1} + \frac{1}{v_2}\right) - \frac{\Delta h_2 d}{T}\left(\frac{s_1}{v_1} + \frac{s_2}{v_2}\right) \tag{4}$$

Collecting Δh_1-terms on the right-hand side and Δh_2-terms on the left-hand side gives

$$\Delta h_2 d\left[s_s\left(\frac{1}{v_1} + \frac{1}{v_2}\right) + T^{-1}\left(\frac{s_1}{v_1} + \frac{s_2}{v_2}\right)\right] < \Delta h_1 d\left[s_s\left(\frac{1}{v_1} + \frac{1}{v_2}\right) + T^{-1}\left(\frac{s_1}{v_1} + \frac{s_2}{v_2}\right)\right] \tag{5}$$

The terms in the square brackets are equal and the inequality reduces to the expected result

$$\Delta h_2 < \Delta h_1 \tag{6}$$

Thus, given the environmental assumptions, maximizing the working altitude Δh will always minimize the required number of agents. In the case of persistent surveillance using a single thermal optimal performance is obtained when agents exploit the full height of the thermal: a "partial recharge" will necessarily result in an increase in the number of agents. It will later be shown that this is not necessarily true in cases where multiple thermals are exploited.

The optimal v_1 and v_2 can be found by differentiating equation (2) with respect to v_1 and v_2. To do this, the aircraft's sink polar must be expressed in terms of the cruising airspeed. For optimization calculations and controller design, this can be expressed fairly well with a quadratic fit: $s = av^2 + bv + c$.[13] The constants a and c are positive, whereas the constant b is negative since the sign convention in this paper is to assign positive values to sink rates. Flight path angle is assumed to be small, so that airspeed is equal to the horizontal speed (for typical gliders the flight path angle is approximately two degrees). The optimal v_1 is found to equal

$$v_1 = \frac{-1 + \sqrt{1 + A_1 B_1}}{A_1} \tag{7}$$

where

$$A_1 = \frac{1}{d}\left(\frac{d}{v_2} + \frac{\Delta h}{T}\right) = \frac{1}{d}(t_{away} - t_{c1}) \tag{8}$$

$$B_1 = \frac{1}{ad}\left[\Delta h - d\left(\frac{s_2}{v_2} + b\right) + c(t_{away} - t_{c1})\right] \tag{9}$$

Similarly, the equation for the optimal v_2 has an identical form due to the structure of equation (2):

$$v_2 = \frac{-1 + \sqrt{1 + A_2 B_2}}{A_2} \tag{10}$$

where

$$A_2 = \frac{1}{d}\left(\frac{d}{v_1} + \frac{\Delta h}{T}\right) = \frac{1}{d}(t_{away} - t_{c2}) \tag{11}$$

$$B_2 = \frac{1}{ad}\left[\Delta h - d\left(\frac{s_1}{v_1} + b\right) + c(t_{away} - t_{c2})\right] \tag{12}$$

The variables t_{c1} and t_{c2} were introduced to represent the two segments comprising t_c. The time spent away from the target is represented by t_{away}. The interdependence and identical form of v_1 and v_2 implies that $v_1 = v_2$. Indeed, after constructing the constrained optimization problem for equation (2) using MatLab's

American Institute of Aeronautics and Astronautics

(a) $\Delta h = 700m$ (b) $\Delta h = 350m$
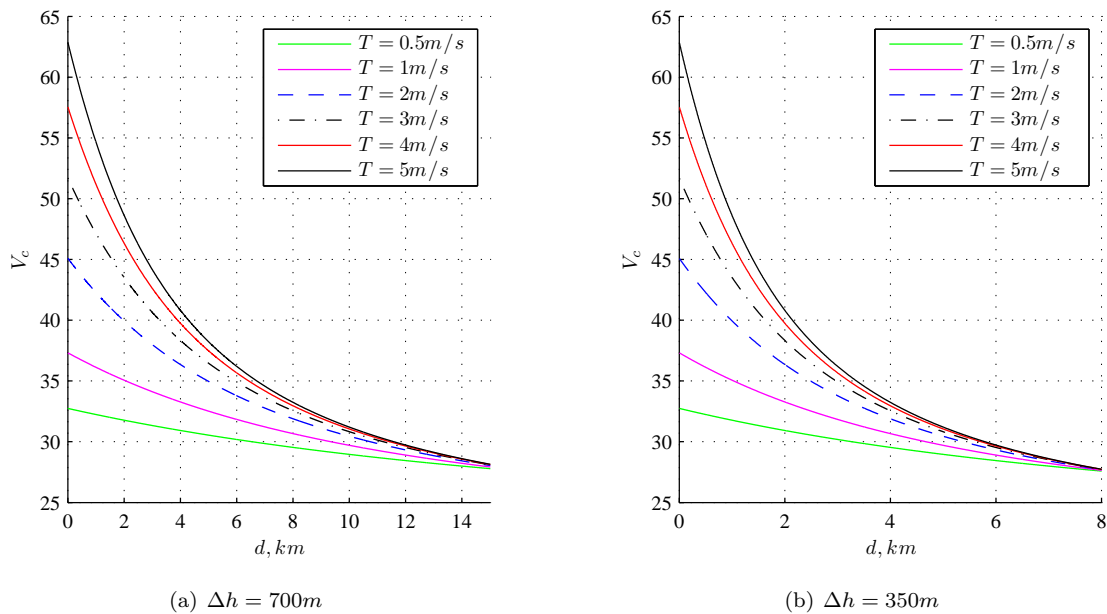
Figure 2. Optimal cruising airspeed for single-agent scenarios with the ASW-27B glider, where d represents the distance between T and M and $s_s = 0.52$ m/s.

`fmincon` function, it was observed that given any set of environmental conditions, the optimal v_1 and v_2 were both equal to

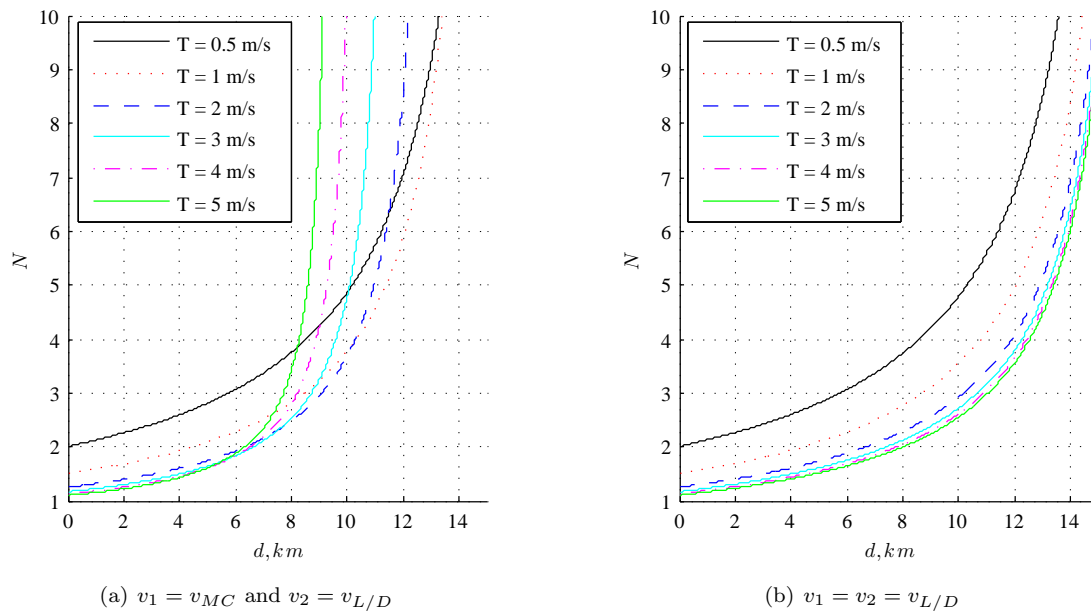$$v = \frac{-2 + \sqrt{4 + AB}}{A} \qquad (13)$$

where

$$A = \frac{\Delta h}{Td} = \frac{1}{d}(t_{away} - t_c) \qquad (14)$$

$$B = \frac{1}{ad}\left(\Delta h - 2bd + \frac{c\Delta h}{T}\right) = \frac{1}{ad}\left(\Delta h - 2bd\right) + A\frac{c}{a} \qquad (15)$$

At first glance, it appears that this equation is circular because of the dependence on t_c. This, however, is not the case because t_c is included in t_{away}; subtracting t_c eliminates the dependence on the two cruising segments. Symmetry within the circuit is expected because flying away from the target affects the circuit in the same way as flying toward the target. Both airspeeds affect the time spent at the target, where v_1 affects the departing altitude and v_2 affects the arriving altitude. Both airspeeds also affect the time spent away from the target with equal weight. Thus, the optimal v_1 that balances these two effects so as to minimize the N required for the system will equal the optimal v_2. It is worthwhile to mention that B may be rewritten using the relations $v_{L/D} = \sqrt{c/a}$, $v_{min} = -b/2a$, $s_{L/D} = 2c + bv_{L/D}$, and $s_{min} = \frac{b}{2}v_{min} + c$, where $v_{L/D}$ is the best L/D speed, v_{min} is the speed at the minimum sink rate of the aircraft, and $s_{L/D}$ and s_{min} are their corresponding sink rates.[13]

Figure 2 plots values of v (later renamed to v_c) as a function of T and d with a working altitude of 700 and 350 meters for a Schleicher ASW-27 glider. The two graphs in the figure converge to the same values as d approaches zero and as d approaches infinity. An alternative formulation in Section IV.B using an aggregate thermal will show that the limits of v are equal to the MacCready speeds as d approaches zero and the best L/D speed as d approaches infinity.

An important observation is that the optimal airspeed does not equal $v_{L/D}$, given by $\sqrt{c/a}$, nor does it equal v_{Mc}. This fact became clear after comparing N using v_{Mc} with N using $v_{L/D}$ during single-thermal exploitation. Equation (2) was solved for distances up to 15 km, a working altitude of 700 m, and thermal climb rates up to 5 m/s for a Schliecher ASW-27B glider (all results presented in this paper are based on the ASW-27B). The results are shown in Figure 3.

American Institute of Aeronautics and Astronautics

(a) $v_1 = v_{MC}$ and $v_2 = v_{L/D}$ (b) $v_1 = v_2 = v_{L/D}$

Figure 3. Number of agents required for single-thermal exploitation with a Schleicher ASW-27B glider with $\Delta h = 700m$ and $s_s = 0.52$. Note the instances where stronger thermals are outperformed by weaker thermals when v_{Mc} is used to cruise toward a thermal as opposed to $v_{L/D}$.

Figure 3 can be interpreted in two ways: one, the number of agents required in a cycle for a thermal of a given climb rate and distance, or two, the maximum distance away a thermal can be for a given number of available agents. v_{Mc} is immediately dismissed as the optimal cruising speed because there are cases where weaker and farther thermals require the same number of agents as stronger and closer thermals. This effect is dominant in thermals that are greater than 6 km away and is attributed to an inefficient exchange between an agent's altitude loss and cruising time for stronger thermals whose MacCready speeds are large. There is, however, a subtlety that is not shown clearly in Figure 3 which is that $v_{L/D}$ is not always more efficient than v_{Mc}, particularly for distances less than 2 km, and it was this realization that first led to the theory discussed in this section.

Table 1 lists the gains in the maximum allowable distance from using the optimal airspeed as opposed to $v_{L/D}$ given a working altitude of 700 m. The first entry is not applicable because thermals with climb rates of 0.5 m/s require three agents to complete a cycle regardless of their distance from the target. This is due to the minimum sink rate, which was assumed to be the sink rate during monitoring, being slightly

Table 1. The additional maximum allowable distance between a monitoring target and a thermal provided by v_{opt} as opposed to $v_{L/D}$ given a working altitude of 700 m and a surveillance sink rate or 0.52 m/s. The values in parenthesis represent the percentage gains in the maximum allowable distance with respect to the maximum allowable distance given by $v_{L/D}$.

Climb Rate	Number of Agents				
	2	3	4	5	10
0.5 m/s	na[*]	85 m (1.5%)	65 m (0.8%)	45 m (0.5%)	15 m (0.1%)
1 m/s	205 m (4.6%)	130 m (1.5%)	85 m (0.8%)	55 m (0.5%)	15 m (0.1%)
2 m/s	290 m (4.3%)	155 m (1.5%)	90 m (0.8%)	60 m (0.5%)	15 m (0.1%)
3 m/s	325 m (4.4%)	160 m (1.5%)	90 m (0.7%)	60 m (0.5%)	15 m (0.1%)
4 m/s	340 m (4.4%)	165 m (1.5%)	90 m (0.7%)	60 m (0.5%)	15 m (0.1%)
5 m/s	345 m (4.3%)	165 m (1.5%)	95 m (0.8%)	60 m (0.5%)	15 m (0.1%)

[*] not applicable, $N > 2$ for thermal climb rates of 0.5 m/s.

American Institute of Aeronautics and Astronautics

larger than 0.5 m/s. The calculation of d cannot be carried out directly by providing a value of N because the value of the optimal airspeed depends on d. Instead, the calculation of d for a certain number of agents was carried out by varying the distance from zero to sixteen kilometers in steps of five meters. The distance was recorded when the value of N reached an integer number.

Table 1 shows that if the optimal airspeed is ever to be used instead of the best L/D airspeed, it should be used in situations where $N \leq 2$, otherwise, the gains in distance do not justify the computational effort after considering the uncertainty of dynamic environments during real missions. It then becomes of interest to quantify the benefit of cruising at the optimal airspeed when $N \leq 2$. However, the previous method of calculating the maximum distance away a thermal can be does not apply for non-integer values of N, and calculating the difference in the number of agents required for a cycle is slightly non-intuitive. Instead, equation (2) is reformulated to account for a savings in time by considering exploration distance, d_{free}, and loitering time, t_{free} (discussed in Section IV).

B.   Multiple-Thermal Exploitation

Multiple-thermal exploitation is of interest because a flock may encounter many scenarios where exploiting a combination of thermals would reduce the number of agents required for a persistent monitoring cycle. Two interesting cases are the presence of one thermal on the path toward another and the utilization of a nearby thermal to reach a better thermal that would otherwise be too far given the working altitude.

The number of agents required for persistent monitoring becomes

$$N = \frac{\frac{d_1}{v_1} + \frac{s_2 d_2}{v_2 T_1} + \frac{d_2}{v_2} + \frac{\Delta h}{T_2} + \frac{d_3}{v_3}}{s_s^{-1}\left(\Delta h - s_1 \frac{d_1}{v_1} - s_2 \frac{d_3}{v_3}\right)} + 1 \tag{16}$$

As illustrated by Figure 1(b), d_1, v_1, and s_1 represent the first leg of the cycle between the target and the first thermal, d_3, v_3, and s_3 represent the final leg of the cycle between the second thermal and the target, and d_2, v_2, and s_2 represent the leg between the two thermals T_1 and T_2. The second term in the numerator of equation (16) represents a partial exploitation of thermal T_1: multiple-thermal exploitation is practical if and only if $T_1 < T_2$, otherwise T_2 would be avoided altogether. Thus, thermal T_1 must be treated strictly as a via point that enables an agent to leave the monitoring target at a later time with the intention of increasing monitoring time at the cost of increasing transit time. It follows that an agent arrives at thermal T_1 at the altitude floor and departs at the height needed for reaching thermal T_2 at the altitude floor. The stronger T_2 is then exploited over the full working altitude.

After differentiating equation (16), the equations for the optimal v_1 and v_3 exhibit the same properties as those in section A: both equations are identical in form and dependent on each other and, as before, MatLab's `fmincon` confirms that the optimal v_1 is equivalent to the optimal v_3. Their value is given by

$$v_c = \frac{-1 + \sqrt{1 + AB}}{A} \tag{17}$$

where

$$A = \frac{1}{d_c}(t_{away} - t_c) \tag{18}$$

$$B = \frac{1}{a d_c}(\Delta h - b d_c) + A\frac{c}{a} \tag{19}$$

where d_c represents the total cruising distance minus the distance traveled between thermals: $d_c = d_1 + d_3$. In fact, equation (17) is the general form of equation (13) when $d_1 \neq d_3$ and although the terms within t_{away} and t_c differ for different scenarios, this general form for the optimal airspeed remains the same. For instance, if the cycle in Figure 1(b) was changed to $M \rightarrow T_1 \rightarrow T_2 \rightarrow T_3 \rightarrow E$, where E is an exploration job, the term "$t_{away} - t_c$" would now equal the sum of the cruising times and exploitation times in $T_1 \rightarrow T_2 \rightarrow T_3$, d_c would now be the sum of the distances between $M \rightarrow T_1$ and $T_3 \rightarrow E$, and equation (17) determines the new optimal airspeed. Note the difference between d and d_c: d represents the distance between M and T for single-thermal exploitation, whereas d_c represents the sum of the lengths of the first and last legs in multiple-thermal exploitation. If single-thermal exploitation is represented in the multiple-thermal formulation, then $d_c/2$ would equal d since the first and last cruise legs are identical.

Equation (16) may also be differentiated with respect to v_2 to determine the optimal cruising speed between thermals. Surprisingly, it is found that the optimal cruising speed between $T_1 \rightarrow T_2$ is equal to the MacCready speed in reverse, i.e., the MacCready speed for T_1 rather than T_2: $v = \sqrt{(c + T_1)a^{-1}}$. Since $T_1 < T_2$ in scenarios involving multiple-thermal exploitation, the time spent exploiting T_1 should be minimized. The slower, reversed MacCready speed in $T_1 \rightarrow T_2$ will minimize the height lost during this phase, which translates to a lower required altitude for reaching T_2.

The reader is reminded that if the cycle in Figure 1(b) is beneficial, i.e., $N_{M \rightarrow T_1 \rightarrow M} > N_{M \rightarrow T_2 \rightarrow M} > N_{M \rightarrow T_1 \rightarrow T_2 \rightarrow M}$, then the cycle $N_{M \rightarrow T_1 \rightarrow T_2 \rightarrow T_1 \rightarrow M}$ should be taken instead of simply $N_{M \rightarrow T_1 \rightarrow T_2 \rightarrow M}$. That is, if T_1 decreases the number of agents required for a cycle by acting as a via point between M and T_2 that allows an agent to depart M at a lower altitude, then it must likewise be possible for T_1 to decrease the number of agents by acting as a via point between T_2 and M that allows an agent to arrive at M at a higher altitude. Indeed, the optimal cruising speed in $T_2 \rightarrow T_1$ is equal in value to the reversed MacCready speed that optimized $T_1 \rightarrow T_2$. However, the difference between the two MacCready speeds is that the segment $T_2 \rightarrow T_1$ is optimized by the forward form of the MacCready speed, $v_{Mc} = \sqrt{(c + T_1)a^{-1}}$, instead of the reversed form $v_{Mc} = \sqrt{(c + T_2)a^{-1}}$. To reiterate, in its normal form, the MacCready speed optimizes a cruise to a thermal based on the thermal that the aircraft is *traveling to*. In this paper, the MacCready speed is said to be reversed when the thermal that the aircraft is *starting from* is used to optimize an inter-thermal cruise in monitoring scenarios.

IV.   Maximizing a Free Parameter

Rather than minimizing the number of agents, this section seeks to maximize a free parameter for a given (but sufficient to perform continuous surveillance) number of agents. This free parameter (time or distance) can be used to perform other tasks such as exploration.

A.   Free time

Loitering time will be analyzed by considering the following equation:

$$N = \frac{\frac{2d}{v} + \frac{\Delta h}{T} + t_{free}}{s_s^{-1}\left(\Delta h - s\frac{2d}{v}\right)} + 1 \tag{20}$$

Equation (20) is a form of equation (2) with an additional segment of $T \rightarrow L$, where L represents a loitering period t_{free} at the thermal. In this paper, loitering is restricted to an agent remaining at the altitude ceiling within a thermal. In reality, atmospheric dynamics necessitate the prediction of dissipating thermals. By remaining within a thermal for an extended period, the loitering agent has a better sense of whether or not a thermal is dissipating, which becomes critical in the decision making process for the subsequent agents seeking a thermal. Alternatively, t_{free}, along with d_{free}, can be a way of measuring the margin of error within a circuit.

Equation (20) is utilized by solving for t_{free}:

$$t_{free} = (N - 1)\left(\Delta h - s\frac{2d}{v}\right)\frac{1}{s_s} - \left(\frac{2d}{v} + \frac{\Delta h}{T}\right) \tag{21}$$

The way to solve this equation is to first calculate the optimal number of agents without t_{free}, i.e., by solving equations (2) and (13), and rounding up to the nearest integer number of agents, $\lceil N \rceil$. Rounding N up to the nearest integer is necessary to account for the additional t_{free} since free time is available only if a circuit has more than enough agents. A difference in agents can now be interpreted as a difference in t_{free}; this can be seen clearly if equation (20) is rearranged in the following manner:

$$\lceil N \rceil - \left(\frac{\frac{2d}{v} + \frac{\Delta h}{T}}{s_s^{-1}\left(\Delta h - s\frac{2d}{v}\right)} + 1\right) = \frac{t_{free}}{s_s^{-1}\left(\Delta h - s\frac{2d}{v}\right)} \tag{22}$$

where the left-hand side equals the difference between the integer number of agents required in reality and the non-integer number of agents required in theory.

After comparing the values of t_{free} given by v_c to the values of t_{free} given by $v_{L/D}$, it was seen that v_c did not always provide a larger t_{free} than $v_{L/D}$ despite requiring a smaller N for a circuit. This result reinstated

(a) $\Delta h = 700m$ (b) $\Delta h = 350m$
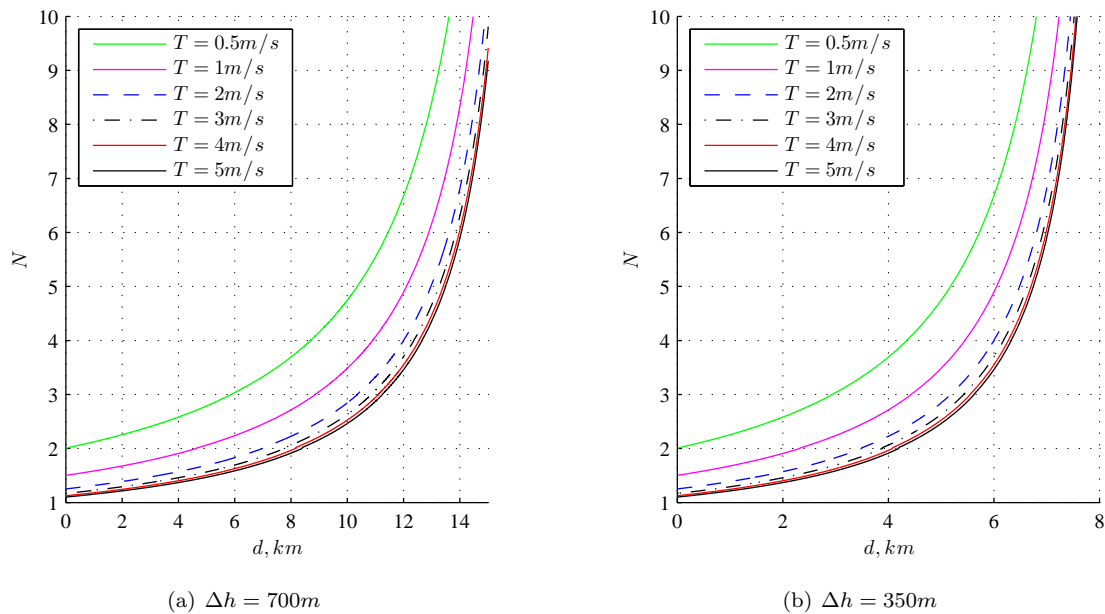
Figure 4. Number of agents required for single-thermal exploitation with v_{opt} for the Schleicher glider with a surveillance sink rate of 0.52 m/s.

that maximizing the endurance of a coordinated soaring mission is an integer optimization problem; the objective is to optimize the resources available, $\lceil N \rceil$, rather than minimizing the number of resources, N. Thus, the optimal cruising airspeed during real missions involving multiple agents must be derived by differentiating equation (21) with respect to v instead of differentiating equation (16). The new optimal cruising airspeed is given by equation (23).

$$v_{opt} = \sqrt{\frac{c + s_s(\lceil N \rceil - 1)^{-1}}{a}} \tag{23}$$

where the calculation of v_{opt} requires calculating v_c and $\lceil N \rceil$. The value of v_c actually converges to v_{opt} as N approaches an integer value, showing that the two optimization problems are related. In addition, v_{opt} converges to $v_{L/D}$, $\sqrt{c/a}$, as $\lceil N \rceil$ tends to infinite. A pleasant consequence of equation (23) is that v_{opt} is constant for a given $\lceil N \rceil$, which means that d can be divided into ranges of constant v_{opt} rather than calculating a different optimal airspeed for every value of d, as was done with v_c. The reader must keep in mind the difference between v_c and v_{opt}: v_c is the optimal cruising airspeed in scenarios involving one agent, whereas v_{opt} is the optimal cruising airspeed for multiple-agent scenarios where the monitoring target can

Table 2. The operating range for v_{opt} in terms of d given six thermal climb rates for the Schleicher ASW-27B, a working altitude of 700 m, and a surveillance sink rate of 0.52 m/s.

	Optimal Airspeed				
Climb Rate	32.8 m/s	30.2 m/s	29.3 m/s	28.9 m/s	28.6 m/s
	$\lceil N \rceil = 2$	$\lceil N \rceil = 3$	$\lceil N \rceil = 4$	$\lceil N \rceil = 5$	$\lceil N \rceil = 6$
0.5 m/s	na[*]	0 - 5.88 km	5.89 - 8.69 km	8.70 - 10.35 km	10.36 - 11.44 km
1 m/s	0 - 4.64 km	4.65 - 8.85 km	8.86 - 10.89 km	10.90 - 12.10 km	12.11 - 12.89 km
2 m/s	0 - 6.98 km	6.99 - 10.36 km	10.37 - 12.00 km	12.01 - 12.97 km	12.98 - 13.61 km
3 m/s	0 - 7.77 km	7.78 - 10.86 km	10.87 - 12.36 km	12.37 - 13.26 km	13.27 - 13.85 km
4 m/s	0 - 8.16 km	8.17 - 11.10 km	11.11 - 12.55 km	12.56 - 13.40 km	13.41 - 13.97 km
5 m/s	0 - 8.39 km	8.40 - 11.25 km	11.26 - 12.65 km	12.66 - 13.49 km	13.50 - 14.05 km

[*] not applicable, $N > 2$ for thermal climb rates of 0.5 m/s.

American Institute of Aeronautics and Astronautics

be monitored indefinitely in theory. The reference cruising airspeed for any coordinated soaring monitoring mission can equal up to about three or four different values; thermals requiring $\lceil N \rceil \geq 5$ are rarely the best options and v_{opt} changes very slightly beyond that. Table 2 lists ranges of d for scenarios involving up to 5 agents given six different thermal climb rates. Figure 4 plots N versus d when the cruising airspeed is v_c, also for six values of T.

B.   The "Aggregate Thermal"

It is interesting to note that the second term in the numerator of equation (23) represents a rate equal to the available altitude at the target divided by the time spent away from the target:

$$\frac{s_s}{\lceil N \rceil - 1} = \frac{\Delta h - s\frac{2d}{v}}{\frac{2d}{v} + \frac{\Delta h}{T} + t_{free}} \tag{24}$$

The right-hand side of equation (24) is a rate that can be thought of as an aggregate thermal and the left-hand side is its relationship with N. The effects of d and Δh are incorporated into the thermal climb rate for a more relevant metric in describing thermals given a particular environmental scenario. Thus, the aggregate thermal is a way of discounting a thermal's strength by the energy required to reach it. This explains the resemblance of equation (23) to the equation for the MacCready speed. Hence, if a circuit has an aggregate thermal strength of 0.5 m/s, the optimal cruising speed is equal to the MacCready speed for a thermal of strength 0.5 m/s: 32.8 m/s. The aggregate thermal concept exists in multiple-thermal scenarios as well and is a general way of combining multiple thermals, the distances between them, and the lengths of the first and last legs of the circuit into one rate.

The utility of this is that instead of thinking in terms of the number of agents required for a particular circuit, one can think in terms of the aggregate thermal for the circuit, avoiding confusion in scenarios where only one agent is available or when the zero-altitude-loss and continuous monitoring constraints do not apply. Equation (23) does not directly apply to scenarios with only one agent because maximizing t_{free} is no longer of interest since there is no free time available. Instead, $\lceil N \rceil$ is replaced by the non-integer N where the interest reverts to minimizing the ratio of the time spent away from the target to the time spent at the target, which is equivalent to the approach for finding v_c in Section III.A. This is precisely the reason why equation (17) gives the same values for the optimal airspeed as equation (23) without the ceiling operator:

$$v_c = \frac{-1 + \sqrt{1 + AB}}{A} = \sqrt{\frac{c + s_s(N-1)^{-1}}{a}} \tag{25}$$

This was a coincidental result that was realized by finding optimal speeds for different conditions, calculating N, plugging N into equation (25), and arriving at the same optimal speed. Equation (25) shows that v_c converges to $v_{L/D}$ and v_{Mc} as d approaches infinite and zero, respectively, pictured in Figure 2. The critical difference between the two equations is that the latter equation, for the continuous monitoring scenario, is circular, i.e., v_c depends on N, but N is a function of v_c, and cannot be used to find v_c for single-agent circuits.

In spite of this, there are other scenarios where the optimal cruising speed is known and the aggregate thermal strength can be determined. In any scenario, the formulation of the aggregate thermal involves defining a resource gained in units of distance, normally altitude, and the time it took to gain that resource. Consider a scenario where the goal is to create a thermal map of a region. The available agents explore the region, keeping a memory of thermal locations, strengths, uncertainty, etc., until needing to gain altitude to continue exploration. Each agent can then choose the thermal that benefits them the most by calculating aggregate thermal strengths. For this case, the aggregate thermal is equal to the estimated height gained from a thermal, relative to an aircraft's current altitude, divided by the estimated time to gain that altitude given the thermal climb rate, the estimated final altitude, the distance between an agent and a thermal, and the cruising speed. The optimal cruising speed, in this scenario, will reach the final altitude the fastest and this speed is known as the MacCready speed.[13] Note that if the agent's next exploration region is established before thermalling, then the MacCready speed will not be optimal, evident from Figure 3. Analogous to v_c maximizing the aggregate thermal, the MacCready speed optimizes the aggregate thermal strength by minimizing the time to reach the final altitude. Knowing that the optimal airspeed in this scenario is independent of N eliminates the circularity discussed in the previous paragraph and allows the aggregate thermal strength to be used as a metric for choosing the "strongest" thermal based on an agents relative position to a thermal.

American Institute of Aeronautics and Astronautics

C. Free distance

It has been mentioned that another way of quantifying the advantage of a circuit requiring less agents is through exploration distance, d_{free}. This can be thought of as an alternative to using t_{free} as loitering time. An agent now has the option of cruising slightly off course to explore nearby regions, the advantage of which is a higher likelihood of locating new thermals or identifying areas of sink to be avoided by subsequent agents. Consider equation (26):

$$ N = \frac{\frac{2d}{v} + \frac{\Delta h}{T} + \frac{d_{free}}{v_{free}}}{s_s^{-1}\left(\Delta h - s\frac{2d}{v} - \frac{d_{free}}{v_{free}}\right)} + 1 \tag{26} $$

Contrary to t_{free}, d_{free} affects the arrival height at the monitoring target. Nevertheless, solving for d_{free} and differentiating with respect to v leads to the same optimal airspeed that maximized t_{free}; this is expected since d_{free} can be thought of as an extension of the cruising legs.

Figure 5 and Figure 6 show the values of t_{free} and d_{free}, respectively. Solid lines represent the additional times and distances provided by v_{opt} compared to $v_{L/D}$, whereas the dashed lines represent the total times and distances provided by v_{opt}; the dashed lines are meant to put the solid lines into perspective. Breaks in the plots denote distances where the number of agents changes; the graphs show distances that require up to five agents. The dashed lines extend past the solid lines in all cases because the values of the additional times and distances were not recorded after the t_{free} from $v_{L/D}$ became negative. In other words, the dashed lines extend past the solid lines because the maximum d using v_{opt} is larger than the maximum d using $v_{L/D}$. The additional distance gained is constant given a certain number of agents since the free time rises at a constant rate given a fixed optimal airspeed. The loitering time gained is less than one minute for all three thermal strengths and the distance gained is less than 700 m. These gains become useful when N approaches an integer value, i.e., toward the right sides of the dashed lines, where resources begin to dwindle.

D. Multiple Thermals

In addition to the theory presented on multiple-thermal exploitation thus far, a better intuitive understanding is sought as to when multiple-thermal exploitation is preferred over single-thermal exploitation. It has been noted that, at the least, multiple-thermal exploitation involves exploiting a weaker, but closer, T_1 as a via point to a stronger, yet farther, T_2. However, the importance of the location of T_1 with respect to T_2 and M, the disparity between climb rate and distance, and the benefit of multiple-thermal exploitation remains to be explored. For simplicity, multiple-thermal exploitation cycles are limited to the partial cycle: $M \to T_1 \to T_2 \to M$, or equivalently, $M \to T_2 \to T_1 \to M$; and the full cycle $M \to T_1 \to T_2 \to T_1 \to M$. As mentioned previously, though the full cycle is preferred over the partial cycle, atmospheric dynamics will at times present the partial cycle as the best option due to thermal creation and dissipation.

Figure 7 shows polar maps of regions indicating whether single-thermal exploitation is preferred over multiple-thermal exploitation. The location of T_2 and M is constant in each map, where M is located at the origin and T_2 is located along the $90°$ line at radius equal to d_3 and is represented by a star; the value d_3 is indicated at the top of each map. Given T_1, T_2, d_3, these maps were created by choosing a location for T_1, thus setting d_1, d_2, and calculating N for the cycles $M \to T_1 \to M$, $M \to T_2 \to M$, and $M \to T_1 \to T_2 \to T_1 \to M$, represented by N_{T_1}, N_{T_2}, and $N_{T_1+T_2+T_1}$, respectively. If the location of T_1 resulted in $N_{T_2} > N_{T_1} < N_{T_1+T_2+T_1}$, where single-thermal exploitation of T_1 requires the least N, then the location was given a red marker. If the location of T_1 resulted in $N_{T_1} > N_{T_2} < N_{T_1+T_2+T_1}$, where single-thermal exploitation of T_2 requires the least N, then the location was given a blue marker. Similarly, if the location of T_1 resulted in $N_{T_1} > N_{T_1+T_2+T_1} < N_{T_2}$, where the full cycle, $M \to T_1 \to T_2 \to T_1 \to M$, requires the least N, then the location was given a black marker. Finally, regions of white within the radius d_3 represent locations where the zero-altitude-loss constraint could not be satisfied. In this manner, the maps were populated by choosing all possible locations of T_1 within d_3 and assigning the appropriate marker. Distance d_3 increases from left to right and top to bottom.

As expected, it becomes less beneficial to exploit T_2 as d_3 increases. The maps show that at close distances, $d_3 < 1350m$, exclusively exploiting T_2 is desired over multiple-thermal exploitation; note that there are no red or black points in the first map, dummy points were plotted solely to produce a complete legend. The next map to the right shows a developing circular red region around the origin. With the development of the red region, a developing black region begins slightly within the red circle and extends
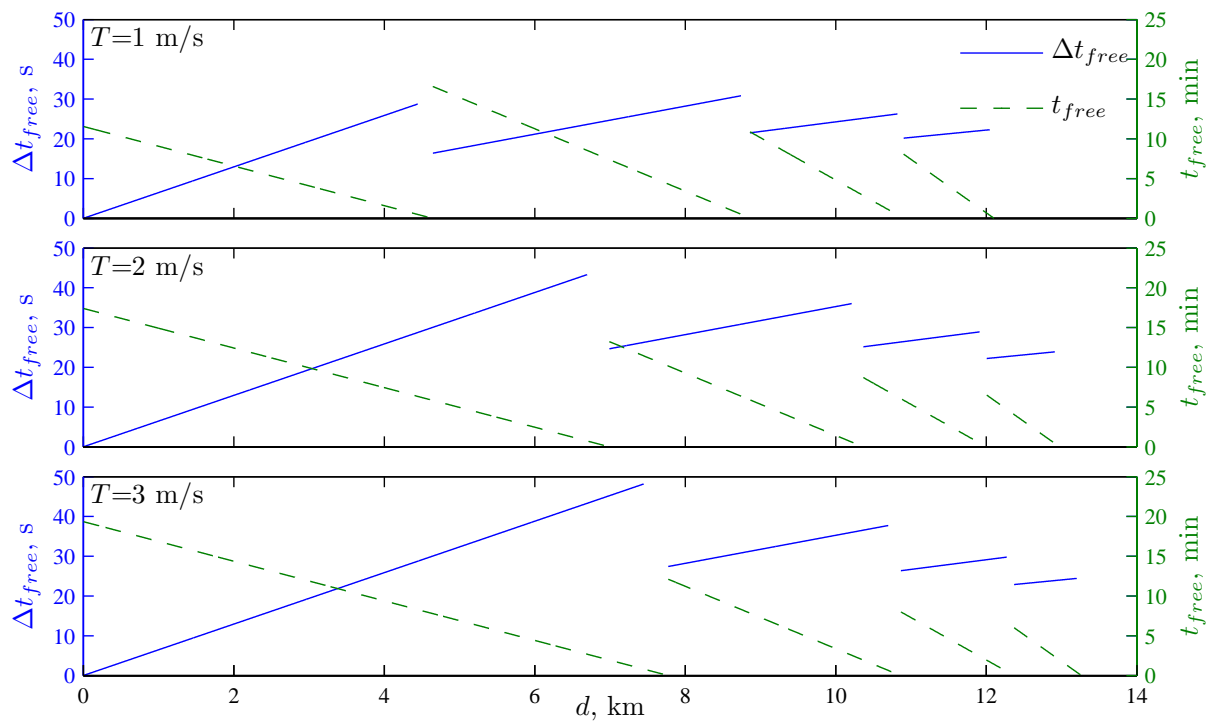
American Institute of Aeronautics and Astronautics

Figure 5. The t_{free} **gained from cruising at** v_{opt} **instead of** $v_{L/D}$. **The solid lines represent the additional** t_{free} **provided by** v_{opt} **and belong to the left axis, whereas the dashed lines represent the total** t_{free} **from** v_{opt} **and belong to the right axis. Top to bottom, the three thermal strengths are 1 m/s, 2 m/s, and 3 m/s. The working altitude is 700 m and the monitoring sink rate is 0.52 m/s.**
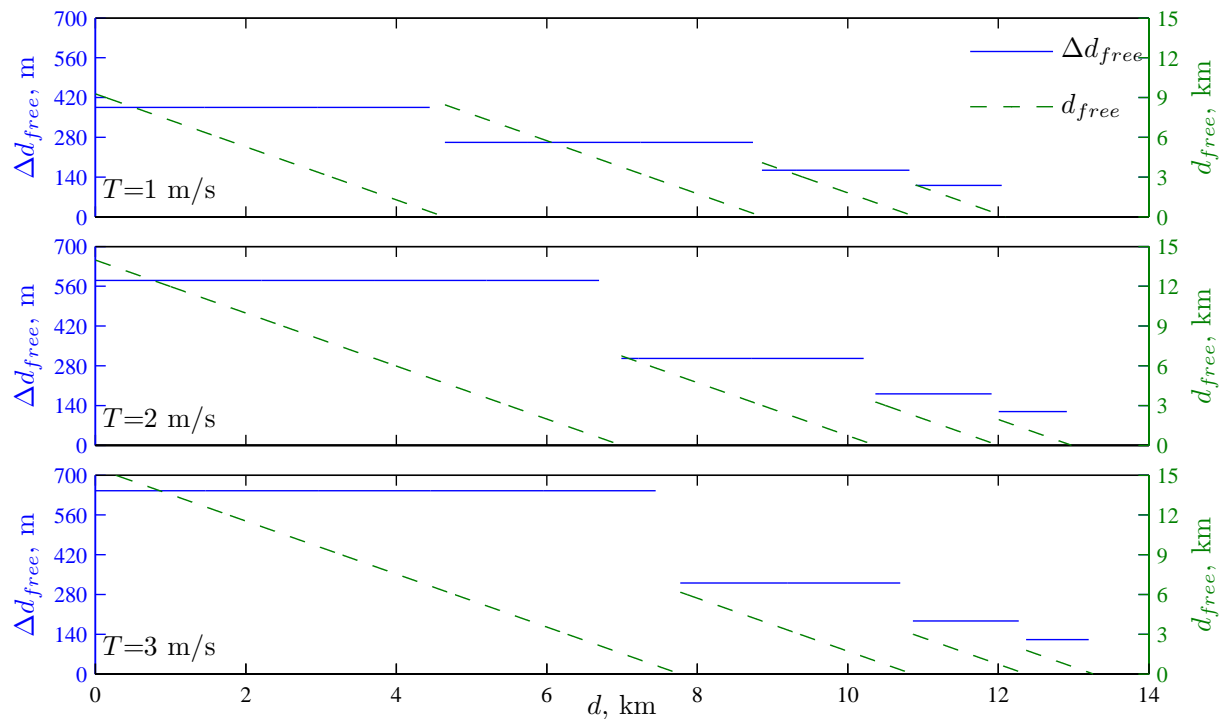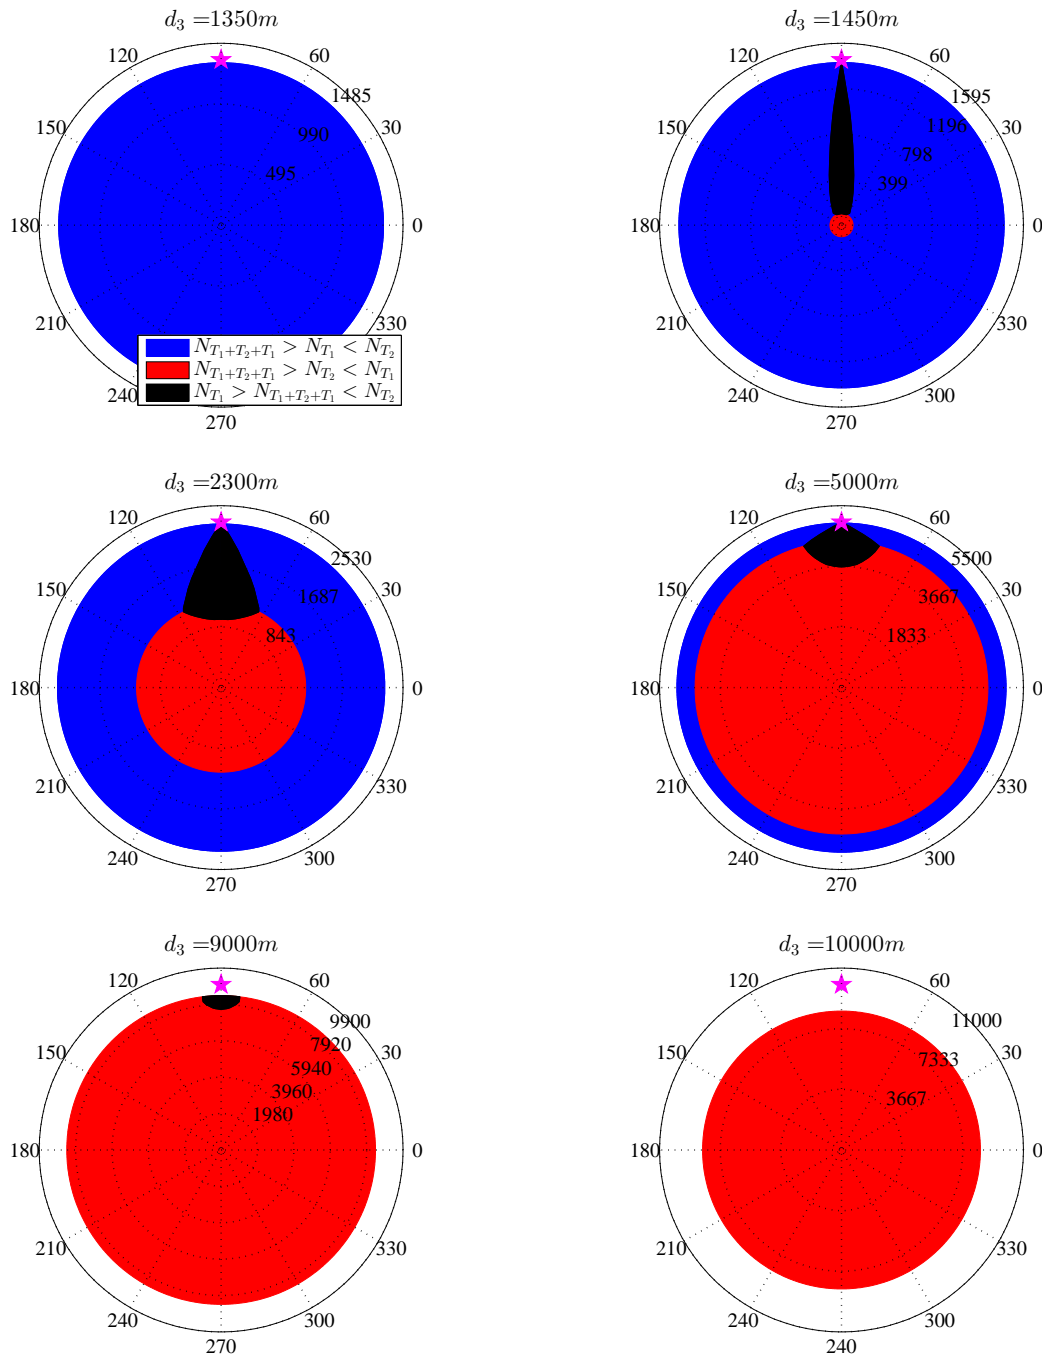

Figure 6. The d_{free} **gained from cruising at** v_{opt} **instead of** $v_{L/D}$. **The solid lines represent the additional** d_{free} **provided by** v_{opt} **and belong to the left axis, whereas the dashed lines represent the total** d_{free} **from** v_{opt} **and belong to the right axis.**

American Institute of Aeronautics and Astronautics

Figure 7. The polar maps of multiple-thermal exploitation for $\Delta h = 350m$, $s_s = 0.52m/s$, $T_1 = 1.5m/s$, and $T_2 = 3m/s$. The monitoring target is located at the origin and the location of T_2 is represented by a star along the $90°$ line at a radius equal to d_3. The distance d_3 varies for each map and is displayed at the top of each map. The maps were created by choosing locations for T_1 until populating the entire map and calculating the number of agents, N_{T_1}, N_{T_2}, and $N_{T_1+T_2+T_1}$ required for the cycles $M \to T_1 \to M$, $M \to T_2 \to M$, and $M \to T_1 \to T_2 \to T_1 \to M$, respectively. Red represents regions where exploitation of T_1 requires the least number of agents, blue represents regions where exploitation of T_2 requires the least number of agents, black represents regions where exploitation of $T_1 \to T_2 \to T_1$ requires the least number of agents, and the white within d_3 represents regions where a cycle that conserves altitude cannot exist.

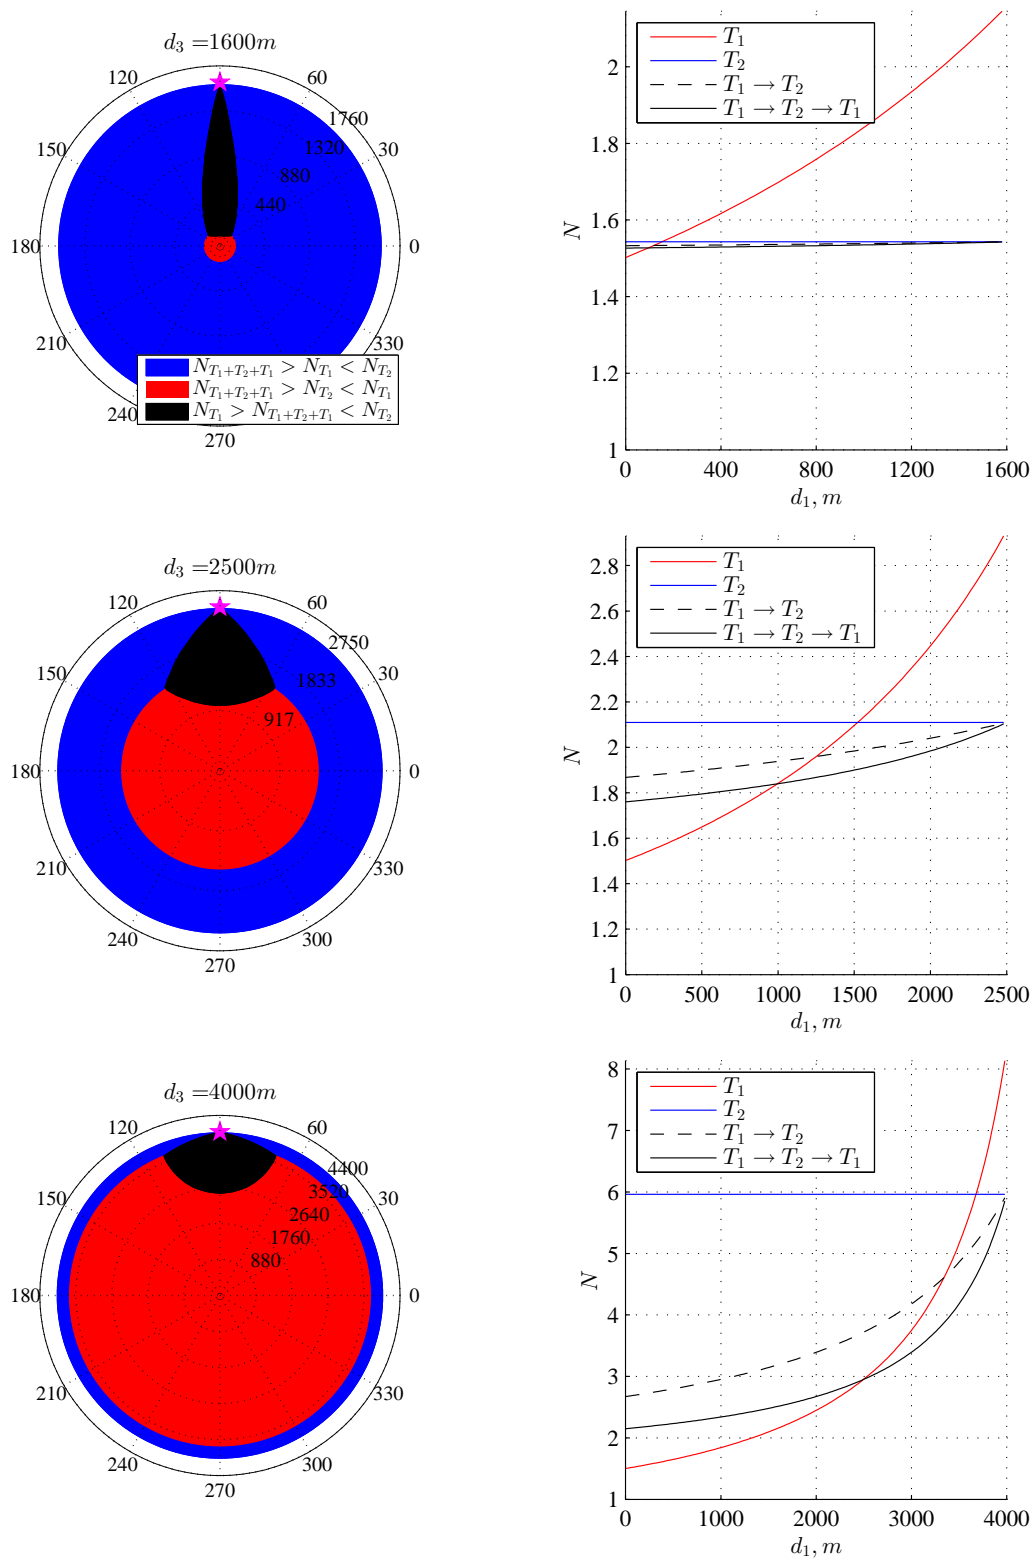American Institute of Aeronautics and Astronautics

Figure 8. The polar maps of multiple-thermal exploitation for $\Delta h = 200m$, $T_1 = 1m/s$, **and** $T_2 = 4.5m/s$ **and the corresponding** d_3 **annotated along the** $90°$ **line.**

American Institute of Aeronautics and Astronautics

to the location of T_2. As d_3 increases, the circular red region grows while the blue and black regions shrink, where the black region always "connects" the red region to T_2. Thus, exploiting T_1 as a via point is a maneuver that is present at larger distances of d_3 and d_1 and is beneficial even if T_2 can be "easily" reached. The third row of maps extend to distances where exploiting T_2 alone would result in a loss of altitude within a cycle. The map to the left shows a region where multiple-thermal exploitation would be beneficial if T_1 was sufficiently far away from the origin, whereas the map to the right no longer has this region because d_3 has become too large. These observations imply that if the zero-altitude-loss constraint cannot be satisfied for a thermal T_2, then multiple-thermal exploitation with T_1 is not worth considering because T_1 would also need to be very far away and the uncertainty in thermal dynamics during the cruising time would be too risky; it would be more efficient to explore the region for a better option. Another interesting observation is that the maximum angular range of the black region for all maps, not only the maps pictured in Figure 7, between $d_3 = 0m$ and $d_3 = 10,000m$ is approximately $75°$. The largest angular range for other conditions was roughly equal to $95°$, where the maximum range occurs very close to the origin in all maps. As a general rule, multiple-thermal exploitation is potentially beneficial if T_1 is located within a $75°$ region on the path to T_2, or $37.5°$ to both sides of T_2 (varying s_s does not affect the shapes in Figure 7 so long as N remains greater than 1).

Figure 8 shows the scenario described by $\Delta h = 200m$, $T_1 = 1m/s$, and $T_2 = 4.5m/s$ for three different values of d_3. The graphs on the right correspond to the polar plots on the left and plot N as a function of d_1, where d_1 is the distance between M and T_1 along the $90°$ line. The solid blue lines represent the exploitation of T_2 and are horizontal because d_3 is constant in each scenario. The solid red lines represent the exploitation of T_1, the dashed black lines represent the partial cycle, and the solid black lines represent the full cycle. The important message in this figure is that there will be scenarios where full-cycle exploitation will significantly change the number of agents required for a cycle: the scenario in the second row, starting from $d_1 \approx 1650m$ and ending at $d_1 \approx 2050m$, would require two agents ($\lceil N \rceil$) for a full-cycle multiple-thermal exploitation, whereas the other cycles would require three agents.

Replacing the full cycle in Figure 7 with a partial cycle affects the polar maps by changing only the black region. The effect is a flattening of the bottom rounded region, resulting in a relatively straight line between the two endpoints of the previously circular curve. As examples, the black region of the fourth polar map in Figure 7 would be roughly half as long due to the flattening of the circular curve and the black region of the fifth polar plot would disappear entirely. It is then concluded that a partial cycle decreases the area where multiple-thermal exploitation is optimal and enlarges the area where single-thermal exploitation of T_1 is optimal.

E.   Effect of Winds

This section is concluded by commenting on the assumptions made in Section III.A. Winds can be treated in the same way as in MacCready theory. Tailwind speeds are subtracted from the reference cruising airspeed, whereas headwind speeds are added to the reference cruising airspeed. The speed of sinking air masses are added to the c term in equations (23) and (25), where sinking speeds are positive and increase the reference cruising airspeed. In relation to equation (16), the speed of sinking air masses is also added to the c term in B:

$$B = \frac{1}{ad_c}\left(\Delta h - bd_c\right) + A\frac{c+w}{a} \tag{27}$$

where w is the speed of the sinking air mass. This was calculated by adding w to s_1 and s_3 in equation (16). Assumptions 2, 3, and 4 may be lifted as long as these quantities can be estimated; all that is required to calculate the optimal airspeed is to estimate the time of a circuit and the resulting altitude left over for monitoring, i.e., the ratio between the resource gained and the time it took to gain it. Assumption 5 may be lifted by keeping track of the a, b, c, and s_s for a nonhomogenous system of agents. Dealing with assumption 6, atmospheric dynamics, is an obstacle for any real-world scenario, irrespective of the cruising airspeed. Traveling at $v_{L/D}$ to a dissipating thermal will leave an agent just as stranded as if the agent had traveled at v_{opt}. The difference between the two cases is that traveling at $v_{L/D}$ would lessen altitude loss during cruise, hence, more time at the monitoring target. This, however, comes at the cost of flying slowly toward a thermal that may be dissipating, increasing the chances of arriving at a dead thermal. Should an aircraft arrive at a dead thermal or be stuck midway up the working altitude due to thermal dissipation, a new minimum allowable altitude should be established and the aggregate thermal metric should be used to choose the next thermal. The effects of atmospheric dynamics, the duration of transient maneuvers, and
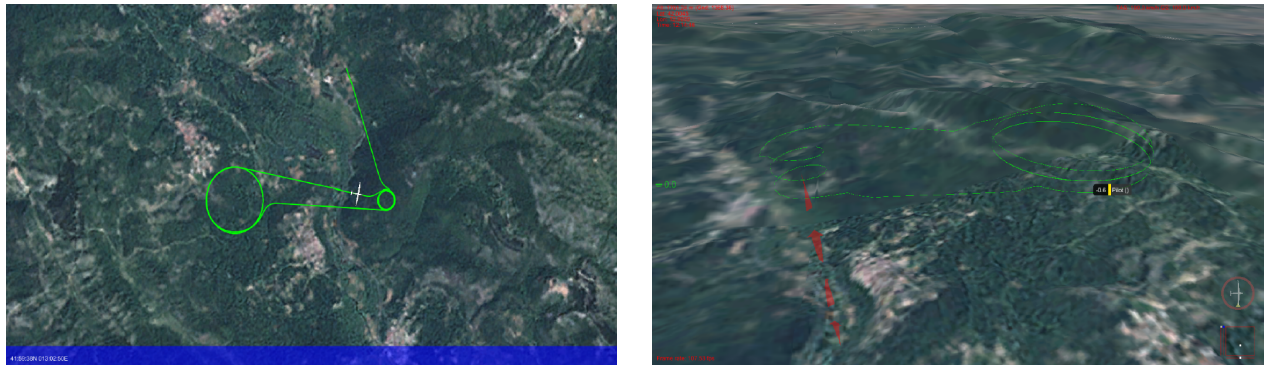
Figure 9. Silent Wings simulation flight paths during single-thermal exploitation with an ASW-27B glider.

the changes in altitude during transient maneuvers will be discussed further in Section V. As a supplement, Reichmann presents an important graphical analysis of the losses due to an incorrect expected climb rate in a thermal; he notes that the cruise speed will not be significantly affected until rather large errors are made and that the zero setting, for which the airspeed is $v_{L/D}$, should be avoided whenever possible.[13]

V.   Simulation Results

Simulations were carried out in the Silent Wings Soaring Simulator where the ASW-27B glider was used in each of the 24 autonomous real-time simulations. The aim of these simulations was to assess whether the discussed theory matched the behavior in reality, where unmodeled dynamics and simplifying assumptions are potential sources of error. Four single-thermal exploitation scenarios were chosen and each scenario was simulated six times: three times with the airspeed set at the optimal cruising airspeed, v_c, and three times with the airspeed set at the best L/D airspeed, $v_{L/D}$. The four chosen scenarios are summarized in table 3. Each simulation was initialized at a random altitude and location to cover a variety of thermal-departure/monitoring-departure headings. The interface with Silent Wings was established via MatLab's Simulink, where UDP was used to receive aircraft states and to autonomously send control commands to the glider. The Simulink model and the Silent Wings Simulator were run from separate computers to facilitate the addition of agents in future simulations involving multiple agents.

In each simulation, the environment was set to be windless except for one fixed thermal set at a known GPS position with a chosen climb rate and an inversion layer set at 2070 meters above sea level. The altitude ceiling for the glider was chosen to be 2000 meters above sea level, where the interaction with the inversion layer began at 1980 meters, i.e., 20 meters were allotted for "realizing" the dissipation before exiting the thermal.

The glider's initial behavior was set to cruise to the thermal and climb to the altitude ceiling, where the altitude ceiling was chosen as the starting location for the cycles in the results. Upon reaching the ceiling, the glider began to roll and change airspeed before cruising with the proper heading and airspeed toward the monitoring target, which was assigned a GPS position to match the distance specified by the scenario. The glider would then monitor the target until sinking to the departure altitude and cruising back to the thermal to repeat the cycle. Silent Wings screenshots show the typical flight path in Figure 9.

The thermalling and monitoring controllers were simple radius hold controllers, for which the circling radius within the thermal and monitoring target was 150 meters and 500 meters, respectively; the respective commanded airspeeds were 31 m/s and 27.78 m/s, where 27.78 m/s equals the $v_{L/D}$ for the ASW-27B. The radii at these two locations were taken into account in the distance setting, thus, the 2 km setting corresponds to a separation of 2.65 km with the circling radii included. A switch case block was used for gain scheduling in the airspeed controller, where the gains were determined through trial and error. The polar fit coefficients required for calculating the optimal airspeeds were found by flying the glider in Silent Wings instead of consulting a documented source. The glider was flown with its flaps fully retracted and the coefficients were $a = 0.001559$, $b = -0.06475$, and $c = 1.174055$, with sink rate taken as positive with units of m/s.
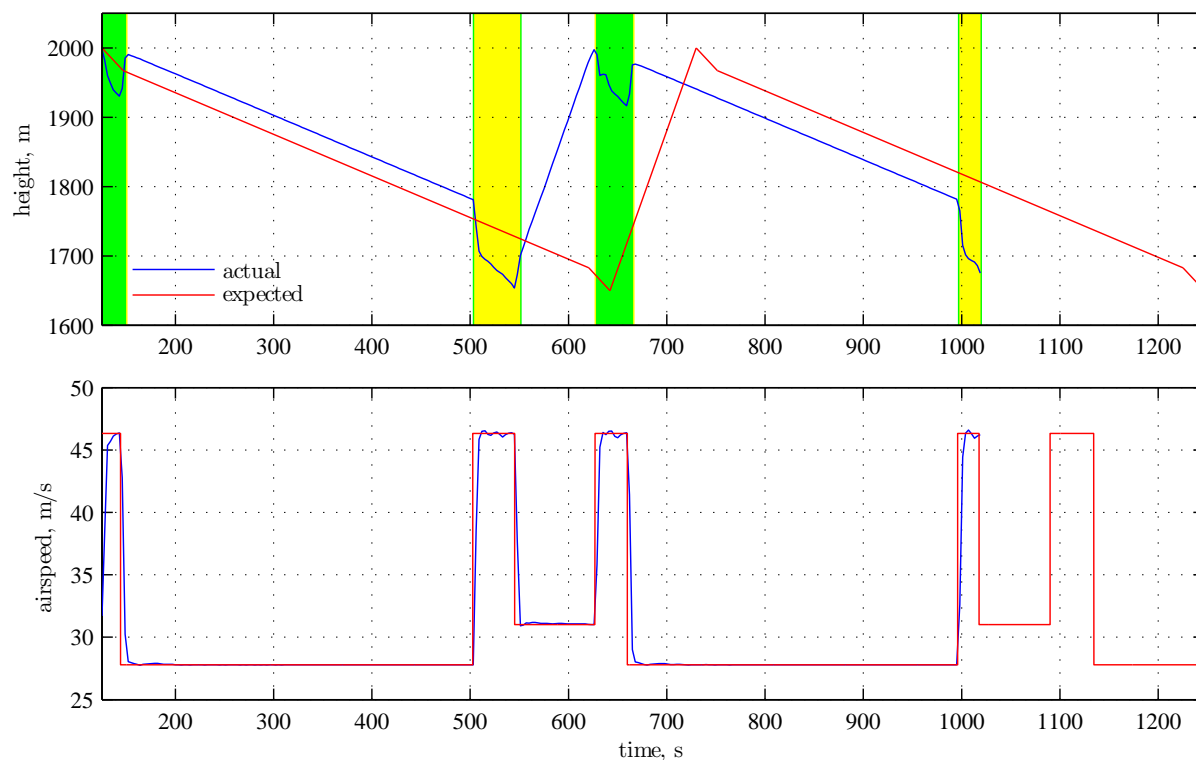
Figure 10. Simulation plot of the scenario $\Delta H = 350m$, $d = 1km$, and $T = 4m/s$, whose theoretical aggregate thermal strength is 2.14 m/s. The two graphs depict the gliders altitude and airspeed, where red lines represent expected altitude and commanded airspeed in the first and second graph, respectively, and blue lines represent the actual values during simulation. Shaded green regions represent $T \rightarrow M$ transitions and shaded yellow regions represent $M \rightarrow T$ transitions.

Figure 10 presents all of the important information obtained from one simulation. The top graph shows the actual height of the glider during simulation and the expected height from the theory discussed in section III.A, whereas the bottom graph shows the response of the airspeed controller. This figure represents the scenario $\Delta H = 350m$, $d = 1km$, and $T = 4m/s$, which had the strongest theoretical aggregate thermal of 2.14 m/s given the optimal airspeed obtained from equation (17), 46.35 m/s. These graphs were generated for all 24 simulations and were used to obtain the information presented in table 4 (discussed later in this section).

The stages illustrated in the height plot will now be explained from left to right. The shaded green region contains three stages. First, the glider starts at the top of the thermal at 2000 meters and enters a transition to begin cruising at v_c. This transition consists of rolling into the proper heading while simultaneously losing altitude to reach the desired airspeed, i.e., converting potential energy to kinetic energy. The glider then cruises at v_c until reaching the monitoring target, where it again enters a transition to reach the desired circling airspeed, this time trading kinetic energy for potential. Thus, the three stages in the green region are the two transitions with the one cruising stage in between them. Following the green region, the white region represents the monitoring stage where the glider holds a fixed radius and airspeed around the target. The glider eventually losses too much altitude and cruises back to the thermal, represented by the yellow region. Similar to the previous transition, the glider trades altitude for speed, cruises to the thermal, and trades speed for altitude before reaching the thermal. The cruising stage is seen clearly here as the straight line with a slope a bit less than -1 and lasting for approximately 45 seconds. Finally, the glider regains altitude at a relatively constant rate until repeating the cycle. The final, right-most point shown in the graph represents the bottom of the thermal of the second cycle, 1650 meters above sea level, where the actual arrival time is sooner than the expected arrival time. Note that the final yellow region does not include the second transition experienced just before starting the thermalling stage.

Table 3. The case descriptions of the four different scenarios tested in simulation. The values in parenthesis are the values associated with the scenarios for which the reference cruising airspeed was $v_{L/D}$.

			Scenario Description			
Case	Δh, m	d, km	T, m/s	v_c, m/s	N	$\bar{N}$
1	350	1	4	46.35	1.28 (1.31)	1.35 (1.30)[*]
2	350	2	4	39.76	1.47 (1.52)	1.53 (1.50)[*]
3	350	1	1	35.08	1.81 (1.82)	1.84 (1.80)[*]
4	350	2	1	33.28	2.08 (2.11)	2.12 (2.07)[*]

[*] $v_{L/D}$ outperforms v_c

Table 4. This table lists the number of agents required for all 24 simulations. Values in parenthesis correspond to simulations where $v_{L/D}$ was set as the reference cruising airspeed. In each case, the four values of N came from different starting points in the cycle calculations. Specifically, N_1 is for the cycle starting and ending at the altitude ceiling, N_2 is for the cycle starting and ending at the top of M, N_3 is for starting and ending at the bottom of M, and N_4 is for starting and ending at the altitude floor. Note that each N may be converted to an aggregate thermal strength via equation (24).

	Number of Agents from Simulations			
Case	N_1	N_2	N_3	N_4
1	1.36 (1.38)	1.40 (1.41)	1.41 (1.44)	1.43 (1.42)[*]
1	1.41 (1.37)[*]	1.42 (1.39)[*]	1.43 (1.43)	1.41 (1.43)
1	1.39 (1.37)[*]	1.39 (1.41)	1.38 (1.46)	1.39 (1.39)
2	1.61 (1.53)[*]	1.59 (1.51)[*]	1.61 (1.54)[*]	1.56 (1.52)[*]
2	1.63 (1.53)[*]	1.66 (1.52)[*]	1.72 (1.55)[*]	1.67 (1.53)[*]
2	1.57 (1.51)[*]	1.58 (1.51)[*]	1.59 (1.51)[*]	1.58 (1.52)[*]
3	1.90 (1.93)	1.94 (1.94)	2.01 (2.03)	2.00 (1.98)[*]
3	1.89 (1.93)	1.91 (1.95)	1.95 (2.04)	1.93 (1.98)
3	1.89 (1.92)	1.93 (1.96)	1.99 (1.94)[*]	1.97 (1.93)[*]
4	2.36 (2.11)[*]	2.31 (2.10)[*]	2.24 (2.10)[*]	2.23 (2.10)[*]
4	2.23 (2.12)[*]	2.21 (2.12)[*]	2.21 (2.12)[*]	2.22 (2.11)[*]
4	2.32 (2.11)[*]	2.29 (2.13)[*]	2.24 (2.12)[*]	2.23 (2.11)[*]

[*] $v_{L/D}$ outperforms v_c

American Institute of Aeronautics and Astronautics

There are a couple of subtle points that remain to be discussed. Starting from the first green region, one sees that the height lost during the initial transition is approximately entirely regained during the second transition. However, the same cannot be said of the two transitions in the first yellow region, where the height gained is about half of the height lost. This imbalance has to do with the characteristic region of sinking air surrounding all thermals. The reason why only half of the height is regained is because the glider enters a region of sinking air that was not encountered at the initial transition in the yellow region. In the green region, the reason why roughly all of the height is regained is because the glider spends extra time within the thermal during the initial transition, i.e., the glider engages in a banked turn while still within the region of lift. In addition, a small amount of the imbalance in the yellow region is attributed to the difference in airspeeds: the initial transition loses the altitude required to reach v_c from a speed of 27.78 m/s, whereas the final transition gains the altitude required to reach 31 m/s from v_c.

Figure 11. Simulation flight path for run one of case three for v_c. The thermalling stage is shown by the narrower spiral in black and the two monitoring stages are shown in blue and bold red, where the blue monitoring stage occurs before the thermalling stage and the red monitoring stage.

The final point to note is the difference in cruise durations: the first and final cruises required about half of the time required for the middle two cruises. This is most likely caused by a difference in transition durations. Take the flight path, shown in Figure 9, where the glider exits the thermal and monitoring target already in the direction of travel. Most of the simulations did not work out as nicely, i.e., the glider had to engage in a banked turn to reach the proper heading. The banked turns could take up to ten seconds and would manifest as circular curves stemming off from the reference radius at either the thermal or the monitoring target. This acted as an extension of the cruising segment and is the main cause for the difference in transit times. Figure 11 and Figure 12 illustrate these cruise extensions and are discussed after the next few paragraphs.

Table 3 describes the scenario settings for each of the four simulated cases and is used in conjunction to table 4. Table 4 lists the N for each of the 24 simulated cycles, half with v_c (values without parenthesis) and half with $v_{L/D}$ (values within parenthesis). The four values of N in table 4 correspond to four different starting points for a cycle calculation: N_1 is for the cycle starting and ending at the altitude ceiling, N_2 is for the cycle starting and ending at the top of M, N_3 is for starting and ending at the bottom of M, and N_4 is for starting and ending at the altitude floor. These N should be compared to the N predicted by theory, shown in Table 3, to get a sense of the accuracy of the theoretical model. Note that the expected values of N in table 3 were calculated with $s_s = 0.6$. Equation (24) may be used to convert these values to aggregate thermal strengths. Lastly, asterisks were placed wherever $v_{L/D}$ outperformed v_c. Note that the values in parenthesis have no relation to the values without parenthesis except for case number, due to the random starting altitude and location in each simulation run. This means that, for example, the first row of values within parenthesis (or without) may be switched with the second row, which may lead to the asterisks being omitted or located at different locations. Thus, the asterisks simply provide a "first-glance" comparison in performance between v_c and $v_{L/D}$. Instead, *all* of the values without parenthesis should be compared with *all* of the values with parenthesis within a given case.

A new variable, $\bar{N}$, was introduced to quantify the effect of an asymmetric transition in the cycles involving v_c, where the cruising airspeed is larger than the reference airspeed at either T or M. The altitude lost during the transition when entering M is balanced by the transition when leaving M, i.e., the gain in potential energy upon entering M is lost upon exiting M. Unfortunately, the same cannot be said for the transitions around T. Entering T results in an altitude gain that shortens the exploitation time of the thermal. This effect is not canceled because the transitional loss of potential energy at the altitude ceiling does not necessarily mean that the exploitation time is extended. For the exploitation time to increase, the glider would need to continue thermalling until the transitional loss has been regained, i.e., the transition would need to occur within the thermal. Nevertheless, this may have been inadvertently accounted for
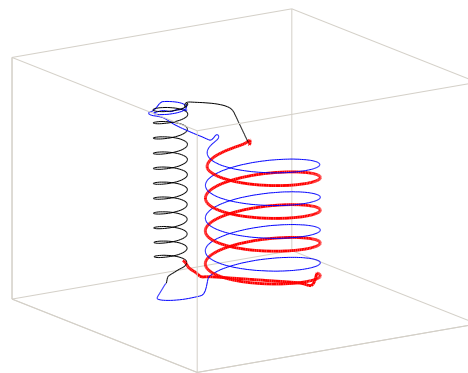
American Institute of Aeronautics and Astronautics
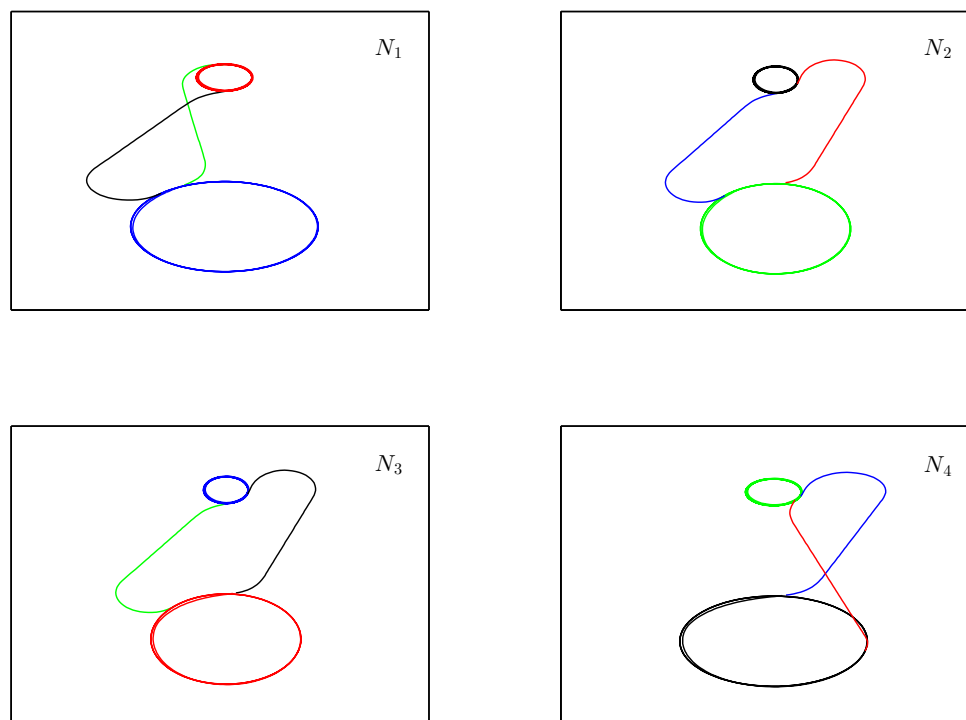
Figure 12. Simulation flight path for run one of case three for v_c divided into the individual N_1 through N_4. The colors of the segments represent chronological order: green, blue, black, and red represent the first, second, third, and fourth segments to occur. These departure trajectories are suboptimal and deviate from the estimated N.

since the glider does, at times, remain within the thermal during its transition trajectory. In any case, $\bar{N}$ includes a transitional correction factor by subtracting the transitional altitude loss from exiting T. The implicit assumption here is that the transitional altitude loss is not at all regained at the top of T. This correction factor appears to decently match the simulation results, presumably because of the surrounding area of sinking air around the thermals adding to the negative effect of shortening the monitoring time. The inaccuracy, however, necessitates a change to the calculation of N, at least for gliders as large as the ASW-27B.

A new calculation of N would account for the asymmetric transitional loss. This would add a dependency on the airspeeds at M and T, in addition to increasing the dynamics of the problem. The end result would be a change in the optimal airspeeds, which appears to be necessary judging by the performance of $v_{L/D}$ in simulation. Faster airspeeds would be preferred when cruising toward S because the transitional altitude gain would increase monitoring time, provided that reaching the faster airspeed does not have an overall negative effect on the system. Conversely, slower airspeeds would be preferred when cruising toward T because less altitude would be lost during the $M \rightarrow T$ transition and more altitude would be gained within T. One consequence of this may be to rely on the thermalling controller to reach the cruising airspeed upon reaching the top of the thermal. In this way, the transitional altitude loss from exiting the thermal would be entirely regained and its effect on the monitoring time would diminish.

Aside from a modification to the thermalling controller, Figure 11 and Figure 12 point toward an important change to the departure controller. Both of these figures show run one of case three cruising at v_c. Figure 11 shows the flight path for the entire simulation, whereas Figure 12 divides the simulation run into the four segments tabulated in table 4. In Figure 11, the thermalling stage is shown in black with the smaller radius and the monitoring stages are shown in blue and a thick red. The monitoring stage in blue preceded the monitoring stage in red, meaning that the monitoring in blue came before thermalling and the monitoring in red came after the thermalling (recall that the simulations started at the inversion layer). The different departure trajectories depicted in Figure 12 account for a large portion of the error and variation seen in table 4.

American Institute of Aeronautics and Astronautics

The colors represent the task order within the calculations of N_1 through N_4. Green, blue, black, and red correspond to orders one, two, three and four. In the plot showing N_1, the agent began by departing the thermal at the inversion layer and cruising to the monitoring target. From there on, the agent monitored the target, departed the target and arrived at the thermal, and finished at the top of the thermal. In the plot showing N_2, the agent began by monitoring the target and finished by cruising back to the top of the monitoring target. The important distinction between N_1 and N_2 in this run is the departure angle from the thermal. In N_1, the agent departs in a shorter trajectory while being near the outer region of lift. In N_2, the agent departs the thermal in a much longer trajectory in addition to being away from the region of lift. This is presumed to be exactly why N_1 is less than N_2 in run one of case three. In N_3, the agent now must suffer the

Figure 13. Simulation flight path for run three of case two for $v_{L/D}$ with the superposition of N_1 through N_4. These departure trajectories are near perfect and agree with the estimated N.

consequences introduced in N_2 by monitoring from a lower starting altitude, hence the lower red segment in Figure 11. This compounding of error is presumably why N_3 and N_4 perform even worse than N_2. The modified departure controller would either delay or prematurely end the thermalling or monitoring stages to minimize the length of the departing trajectory. Such a controller would cause the trajectories to be similar to the ones seen in run three of case two flown at $v_{L/D}$, shown in Figure 13. These trajectories are near perfect and have nearly identical values for N_1 through N_4 that agree with the predicted N for $v_{L/D}$.

Smaller, lower wing-loading aircraft more representative of small UAVs (e.g. the RnR Products SB-XC) are likely to be less influenced by the effects of transitions: flight speeds are lower and the difference in speed between thermalling/surveillance and the cruise flight conditions are smaller than for full-sized sailplanes. As a result the altitude changes at transitions (as kinetic and potential energy is exchanged) are smaller.

The flight paths between thermalling and surveillance suggest that some improvement in overall performance (i.e. closing the gap between analytical results and higher-order simulation results) can also be achieved if optimal flight paths are computed. For example, a transition flight path that is tangent to both the thermal and the surveillance paths will be significantly shorter in length than the current paths, resulting in both reduced flight-time during the transition and reduced altitude loss during transition.

VI.   Conclusion

This paper has presented a means to optimize persistent, continuous surveillance by a flock of soaring-capable autonomous aircraft. The minimum number of agents required to maintain persistent, continuous surveillance is computed by optimizing cruise speeds between the monitoring target and the energy source (in this case a thermal) and between the energy source and the monitoring target. These optimal speeds depend on thermal strength and distance to the target, and interestingly are generally not equal to either the speed for best L/D (which would minimize altitude loss during the cruise segment) or the MacCready speed (which minimizes the time required to fly some distance to a thermal and then climb back up to the starting altitude).

In addition to the case of single-thermal exploitation, the problem of exploiting multiple thermals (e.g. using a "mid-stage" thermal to enable exploitation of a stronger, but more distant, thermal) is addressed. Both optimal cruise speeds and conditions under which a multiple thermal cycle should be taken is addressed. The concept of an aggregate thermal is introduced as a unified metric for energy exploitation during surveillance, and the use of this aggregate thermal for other tasks (such as exploration) is discussed.

Finally, a method to maximize a free parameter (such as time or distance) is discussed. This free parameter can be used to complete a secondary task (such as exploration) during the surveillance cycle.

Results of simulations using both MatLab and a high fidelity soaring simulator (Silent Wings) illustrate the results.

Generally the gain obtained by flying at the optimal airspeed is quite small when compared with flying at
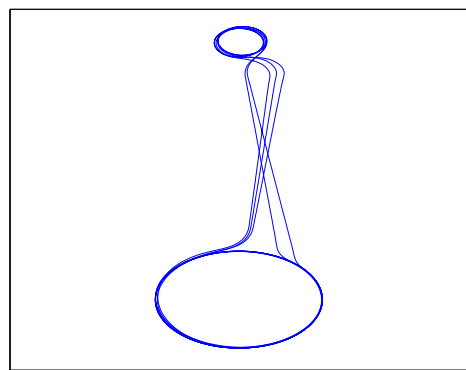
American Institute of Aeronautics and Astronautics

the speed for best L/D (i.e. best range). However, this gain can become important in resource-constrained situations (for example, when thermal conditions are such that an additional aircraft is required if optimal speeds are not flown).

In practice, secondary effects (such as optimal flight paths between thermals and the surveillance target) are very important, and naive approaches to the cruise portion of the mission can have a significant adverse impact on mission performance.

Acknowledgments

References

[1]Andersson, K., Kaminer, I., Jones, K. D., Dobrokhodov, V., and Lee, D.-J., "Cooperating UAVs Using Thermal Lift to Extend Endurance," *AIAA Unmanned Unlimited Conference*, Seattle, WA, 6-9 April 2009.

[2]Akos, Z., Nagy, M., Leven, S., and Vicsek, T., "Thermal Soaring Flight of Birds and UAVs," *Bioinspiration and Biomimetics*, Vol. 5, No. 4, Nov. 2010.

[3]Wharington, J., *Autonomous Control of Soaring Aircraft by Reinforcement Learning*, Ph.D. thesis, Royal Melbourne Institute of Technology, Melbourne, Australia, Nov. 1998.

[4]Edwards, D. J., "Implementation Details and Flight Test Results of an Autonomous Soaring Controller," *AIAA Guidance, Navigation, and Control Conference*, Reston, VA, Aug. 2008.

[5]Allen, M. J., "Autonomous Soaring for Improved Endurance of a Small Uninhabited Air Vehicle," *AIAA Aerospace Sciences Meeting and Exhibit*, Reno, NV, 10-13 Jan. 2005.

[6]Depenbusch, N. and Langelaan, J. W., "Coordinated Mapping and Exploration for Autonomous Soaring," *Infotech@Aerospace Conference*, AIAA Paper 2011-1436, St. Louis, MO, March 2011.

[7]Oke, T. R., *Boundary Layer Climates*, Methuen, London, 2nd ed., 1987, pp. 72–73.

[8]Kerlinger, P., *Boundary Layer Climates*, University of Chicago Press, Chicago, 1989, pp. 78–100.

[9]Bethke, B., How, J., and Vian, J., "Group Health Management of UAV Teams With Applications to Persistent Surveillance," *American Control Conference*, 2008.

[10]Bethke, B., How, J., and Vian, J., "Multi-UAV Persistent Surveillance With Communication Constraints and Health Management," *AIAA Guidance, Navigation and Control Conference*, Chicago, Illinois, 2009.

[11]Bethke, B., Redding, J., How, J., Vavrina, M., and Vian, J., "Agent Capability in Persistent Mission Planning using Approximate dynamic programming," *American Control Conference*, 2010.

[12]Cutler, M. J., McLain, T. W., Beard, R. W., and Capozzi, B., "Energy Harvesting and Mission Effectiveness for Small Unmanned Aircraft," *AIAA Guidance, Navigation, and Control Conference*, AIAA Paper 2010-8037, American Institute of Aeronautics and Astronautics, Toronto, Ontario, August 2-5 2010.

[13]Reichmann, H., *Cross-Country Soaring*, Soaring Society of America, Chicago, 1993, pp. 114–124.

[14]ASW 27B, *http://www.easternsailplane.com/eastern/eastern.php/Schleicher/ASW27B*, accessed June 4, 2013.

American Institute of Aeronautics and Astronautics

Appendix

Table 5. Properties of ASW-27B.[14]

Parameter	Value
Span including winglets	15 m (49.22 ft)
Wing area	9 m^2 (96.88 ft^2)
Aspect ratio	25
Fuselage length	6.55 m (21.49 ft)
Cockpit height	0.80 m (2.62 ft)
Cockpit width	0.64 m (2.10 ft)
Height at tail unit	1.30 m (4.27 ft)
Wing airfoils	DU 89-134/14 and DU 92-131/14MOD
Empty mass	235 kg (518 lb)
Flight mass max.	500 kg (1102 lb)
Mass of one wing	58 kg (128 lb)
Wing loading max.	55.56 kg/m^2 (11.38 lb/ft^2)
Wing loading min.	32.80 kg/m^2 (6.7 lb/ft^2)
Water ballast max.	165 L (43.59 US Gal)
Useful load max.	130 kg (287 lb)
Useful load in pilot seat	115 kg (254 lb)
For m=320 kg:	
Max. speed	285 km/h (154 kts)
Max. maneuver speed	215 km/h (116 kts)
Min. speed	70 km/h (37.8 kts)
Min. sink	0.52 m/s (102 ft/min)
Best L/D	28 at 100 km/h (54 kts)

American Institute of Aeronautics and Astronautics