

Development and Testing of the Miniature Aerial Delivery System Snowflake

Oleg A. Yakimenko*

Naval Postgraduate School, Monterey, CA 93943-5107

Nathan J. Slegers†

University of Alabama, Huntsville, AL 35899-0266

Robyn A. Tiaden‡

U.S. Army Yuma Proving Ground, Yuma, AZ 85365-9110

This paper discusses the current status of the development of the miniature aerial delivery system to be further employed in a variety of different research projects. It starts from the overall description of the system and proceeds with the discussion of the test results performed so far. Specifically, it addresses the effects of the changing ground winds on touchdown accuracy. The paper ends with conclusions and recommendations.

Abbreviations

ADS	=	Aerial Delivery System
AGL	=	Above Ground Level
CEP	=	Circular Error Probable
GNC	=	Guidance, Navigation and Control
MSL	=	Mean Sea Level
TNT	=	Tactical Network Topology
UAV	=	Unmanned Air Vehicle
YPG	=	U.S. Army Yuma Proving Ground

I. Introduction

THIS paper deals with the current state of development of a research testbed, a miniature ADS prototype known as Snowflake, capable of high-precision maneuvering and high touchdown accuracy for a variety of potential applications. Based on previous experience, several such systems have been fabricated and tested thus far. While the control algorithms implemented in this system are described in detail in Ref.1 this paper addresses the results of several test drops performed to date using general aviation aircraft and UH-1A helicopters.

The paper is organized as follows. Section II provides a brief description of the system and methods of its deployment. Section III analyses the results of the first drops of the system, performed in May of 2008 at Camp Roberts, CA. Section IV deals with the results of drops performed at the U.S Army Yuma Proving Ground, Yuma, AZ in October of 2008. The paper ends with conclusions and recommendations for the further development.

II. Description of the System and Methods of its Deployment

The fully deployed parafoil system “Snowflake” is shown in Fig.1. Figure 2 details a payload container, chosen to be a watertight, crushproof, and dustproof Pelican 1200 case (10.62”×9.68”×4.87” or 27cm×24.6cm×12.4cm), with the guidance, navigation and control (GNC) unit occupying a small fraction of its interior chamber (9.25”×7.12”×4.12” or 23.5cm×18.1cm×10.5cm) as shown in Fig.2b. The packed parafoil (about the size of a fist)

* Research Associate Professor, Department of Mechanical and Astronautical Engineering, Code MAE/Yk, oayakime@nps.edu, Associate Fellow AIAA.

† Assistant Professor, Department of Mechanical and Aerospace Engineering, slegers@mae.uah.edu, Member AIAA.

‡ Team Leader, Air Delivery Systems Branch, robyn.tiaden@us.army.mil, Member AIAA.

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE MAY 2009	2. REPORT TYPE	3. DATES COVERED 00-00-2009 to 00-00-2009			
4. TITLE AND SUBTITLE Development and Testing of the Miniature Aerial Delivery System Snowflake		5a. CONTRACT NUMBER			
		5b. GRANT NUMBER			
		5c. PROGRAM ELEMENT NUMBER			
6. AUTHOR(S)		5d. PROJECT NUMBER			
		5e. TASK NUMBER			
		5f. WORK UNIT NUMBER			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School, Department of Mechanical and Astronautical Engineering, Monterey, CA, 93943		8. PERFORMING ORGANIZATION REPORT NUMBER			
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)			
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)			
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified	Same as Report (SAR)	15	

resides in a side pocket as shown in Fig.2c. The canopy dimensions introduced in Fig.3 are outlined in Table 1. The dry weight of the Snowflake system is 4.3lbs (1.95kg).




Fig. 1 Fully deployed Snowflakes: ID #23 (a) and ID #47 (b).




Fig. 2 Payload container with embedded GNC unit.




Fig. 3 Canopy geometry.

Table 1. Snowflake canopy characteristics.

Parameter	Value in United States' customary system of units	Value in International System of units
θ_1		80 °
θ_2		45 °
Airfoil thickness, h	0.35 ft	0.11 m
Wing span, b	4.5 ft	1.37 m
Chord, c	2.1 ft	0.64 m

The GNC unit consists of avionics and control actuators. The avionics include three accelerometers, three rate gyroscopes, a magnetometer, a global positioning system receiver and a barometric altimeter. All of these sensors, along with an autopilot processor, are integrated on a single circuit board (2"×1.37"×0.47" or 5.1cm×3.5cm×1.2cm) weighing only 0.6oz (17g). Among others, the developed GNC unit features precision data-logging, minimal power consumption, telemetry playback and hardware-in-the-loop simulator capability. Snowflake's autopilot uses standard rechargeable 7.4V 2100mAh batteries. It consumes 750mA of power in the full mode (controls and data downlink), 250mA in the autonomous mode (controls only), and drops down to just several mA in the sleep mode. The dimensions of the autopilot are about 3"×2"×1.5" (7.6cm×5.1cm×3.8cm), while its two servo motors (to move control lines) and batteries occupy roughly the same volume. Therefore, in total, the Snowflake GNC package occupies about 45in³ (737cm³), or less than 17% of the interior chamber of a Pelican 1200 case.

Mass-geometry parameters of the Snowflake system were used to develop an accurate six-degree-of-freedom model. The model was verified using actual drops and later used in computer simulations, providing necessary inputs to develop robust GNC algorithms as outlined in Ref.1.

As is, the system is capable of carrying an additional 4...5lbs (1.8...2.2kg) of payload and exhibits the following performance characteristics:

- Descent rate: 12 ft/s (3.66 m/s);
- Forward speed: 14 kts (7.2 m/s);
- Glide ratio: 2:1;
- Minimum turn radius 50 ft (15.2 m).

The system can be deployed by jumpers, and by piloted or unmanned air vehicles (UAV). To date the Snowflake system has been deployed from general aviation aircraft, Cessna-172 (Fig.4a), and UH-1A helicopter (Fig.5). Releasing a system consists of pushing the button (seen in Fig.2c), and throwing the system away from the aircraft (no static line). So far, the following release conditions were successfully tested:

- Release altitude above ground level (AGL): up to 4,000 ft (1,220 m);
- Release airspeed: up to 80 kts (41.2 m/s).

This accounts to a dynamic pressure of 0.15 psi (1KPa).

Upon release (from the side window of an aircraft as shown in Fig.4b, or from a side cargo ramp of a helicopter, Fig.5) the following two delays were implemented as a safety precaution:

- Canopy opening delay (to clear the aircraft tail or the helicopter downwash): 3 seconds;
- Controls activation delay (to allow for untwisting of the risers): 4 seconds.

Therefore, the canopy opens after losing 144.8ft (44.1m) of the deployment altitude, and the actual autonomous control of the system starts after losing an additional 257.5ft (78.5m) of altitude, (about 400ft or 120m loss in total).




Fig. 4 Cessna-172SP aircraft (a) at the ramp and a deployment window (b).



Fig. 5 An UH-1A helicopter with open side cargo ramps.

In the case of deploying from an UAV, a four-Snowflake holding case, to be attached to the belly of a vehicle, has also been developed (see Fig.6). In this case, depending on the scenario, each of the four Snowflakes can be released either by an UAV autopilot upon reaching the corresponding waypoint or from the ground. Further development also includes deploying the Snowflake system from higher altitudes (~10,000ft) using the Quasar RD-400 high altitude research launch vehicle designed to carry small-diameter payloads. For the Quasar rocket vehicle deployment, the entire Snowflake avionics is to be repackaged into a 4"-diameter by 0.8'-tall container (0.1m×0.25m). This rocket system, shown in Fig.7, stands over eight feet (2.4m) tall and weights about 7lbs (3.2kg).




Fig. 6 Schematics of a Snowflakes holding case to be used with an unmanned air vehicle (a) and a release mechanism (b).




Fig. 7 Schematics of a Quasar RD-400 research rocket vehicle with Snowflake payload residing close to the center of gravity.

III. Drops at Camp Roberts

This section presents the results of the first series of drops performed in May of 2008 during the Tactical Network Topology (TNT) experiments in Camp Roberts, CA. It starts from the overall description of the tests, followed by discussion of real and estimated wind profiles. It then proceeds with presenting the results of the drops and analysis of the factors that might affect the touchdown accuracy.

A. Experiment Set Up

As an example, Fig.8 presents a horizontal projection of one of the first trajectories, resulting in a miss distance of only 41.2m. Depicted by red crosses are several reference points established as the inertial trajectory frame to better guide the system. These points are the boundaries of the loiter area (four points), the final-turn-initiation point, and the target (see Ref.1 for more details).




Fig. 8 Horizontal projection of the descent trajectory for the Drop #23-1.

A two-hour window was dedicated to the experiments with Snowflake. To save time, several scenarios were calculated beforehand depending on the ground wind direction,. Five of these scenarios are shown in Fig.9. The

suggested target point determined for westerly or easterly winds, as shown in Fig.9a and 9b, respectively, is depicted with a circled cross. The suggested release points, A, B, C, and D, (shown in Fig.9 by the circles) were chosen upwind due to certain airspace limitations. The boundaries of the suggested loiter areas for each scenario are presented as a group of four crosses with the scenario ID number between them.




Fig. 9 Precomputed sets of reference points differing by direction of the ground winds.

By the time of experiment, the ground winds (to be discussed in detail next) were about 1.55m/s (3kts) from South-South-West (120° “from” or 300° “towards”), and therefore Scenario #2 with a release point B (shown in Fig.9a) was implemented. Two systems were released sequentially into the wind at the prescribed release altitude. The results of these drops are presented in Fig.10.




Fig. 10 Descent trajectories for Drops #23-2 (a) and #47-2 (b).

The miss distances for the two drops were 49.8m and 90.3m, respectively. It should be noted though that data recording for Drop #47-2 stopped prematurely, (when the system was 20.7m above the ground), so the actual touchdown point is not shown in Fig.10b. Since the wind profile affects the performance of any ADS the most, the next section discusses how good the onboard estimates of the winds were compared to pre and post-drop measurements obtained via weather balloon.

B. Real and Estimated Wind Profiles

Essentially, in its current configuration, the Snowflake ADS only needs to know the direction of the ground winds. The rest of the data which is needed for the GNC algorithms is being computed onboard while the system descends and progresses through all guidance phases.¹ However, for the sake of comparison of actual data with that estimated by the ADS, two weather balloon launches were conducted in association with ADS experiments, one before the first series of drops and another right after the second series of drops, at 9:00am and 12:00pm, respectively. Each time the parameters of the atmosphere, including magnitude and direction of wind, were measured twice – on the way up and on the way down (when reaching a certain altitude the Rawinsonde was released from the balloon). Figure 11 presents the vertical wind profiles gathered during these two balloon launches.




Fig. 11 Wind data for the balloon (going up) and Rawinsonde (going down) for 9:00am (a) and 12:00pm (b).

The first balloon went up to about 5,000ft MSL, but Fig.11a only presents altitudes relevant to the ADS drops. The Snowflake ADS was released at 2,800ft MSL and landed at 920ft MSL (the elevation of the McMillan airfield), so the vertical scale in Fig.11 goes from 0 to 1,800ft AGL. As seen from Fig.11a, at the higher altitudes (between 500ft and 1,400ft) both balloon and sonde measured about the same wind profile, while at the lower altitudes light winds exhibited quite big variations in wind direction. Because of a helium leak the second balloon did not go very high, only to about 1,500ft MSL (Fig.11b). As seen, in three hours the ground winds became stronger and there were less variations in their direction.

Figure 12 combines data provided by two balloons together to show the spread of wind magnitude and direction within three hours of when Snowflake ADSs were deployed. In addition to that, Fig.13 shows the same data complemented with the wind estimates from two ADSs, #23-2 and #47-2. On the plots, solid circles represent the instants of control phase change.




Fig. 12 Spread of wind magnitude and direction versus altitude as recorded by the weather balloon at 9:00am and 12:00pm.




Fig. 13 Wind profiles measured by two weather balloons (the same as in Fig.11) and wind profiles as estimated by Snowflakes # 23-2 and 47-2.

According to Ref.1, Phase 0 corresponds to the time between pushing the activation button and when the ADS is fully deployed and ready to execute controls (no wind estimation), and Phase 1 corresponds to the initial autonomous descent towards the loiter area (with the initial wind guess uploaded to the system before deployment). Hence, the first circles correspond to the moments of switching between Phase 0 and Phase 1. The first time the wind is estimated is when the ADS gets to the loiter area (Phase 2, bounded by the second and third circles counted from the top of each subplot). As seen in Fig. 13, the magnitude of wind changes with altitude and on the average is slightly less than that of the measured 9:00am wind profile. Based on this estimate the beginning of the next phase is

calculated. During the turn toward the target, i.e. Phase 3, the wind estimation is halted. While on the downwind leg (Phase 4) the wind estimation is reactivated in order to compute the moment of the final turn maneuver. As seen during this phase the wind estimates for both drops happen to be within the two measured wind profiles, proving that onboard wind estimators do a fairly good job. During Phase 5 and 6, turning and final approach, respectively, the wind estimation is halted again.

The right plots in Figs. 11a, 11b, and 12 also show a vertical dashed line. This line represents a direction of the ground winds, 115° , that was assumed by the ADS. As seen, this direction is slightly off of what weather balloons provided, but the desire was to have ADS descending parallel to the runway during Phase 4, and that is why the Scenario #2 (Fig. 8) was picked up. As discussed in the next section, having light crosswinds during Phase 5 contributed to the touchdown error.

C. Drop Data Analysis

In general, by observing Figs. 8 and 10, it can be stated that all three systems released at Camp Roberts in May of 2008 behaved fairly well. No major malfunctions were encountered, all three systems were safely released from the piloted aircraft at a relatively high speed, all canopies had clean opening with no twisting, the switching between all phases went as planned, and the wind estimates were good enough to construct and track the entire maneuver down to and including Phase 4. Therefore, all objectives for the initial testing of the system were met.

In terms of touchdown accuracy, it should be noted that the test objective of less than 100m CEP was also met. Figure 14 represents the data collected for three Snowflake drops. As seen, this system demonstrated a superb performance. The average error for three drops (41.2m, 49.8m and 90.3m) was 60.5m and the median value turned out to be 50m.




Fig. 14 Performance for three drops of Snowflake ADS at TNT 08-03.

Phase 5 is significantly responsible for the touchdown accuracy. By looking at Figs. 8 and 10 it is clear that some improvement can be made to achieve better performance at this phase. Analysis of the data revealed three major sources of error and therefore three possible areas for improvement. (It should be noted that sometimes these errors cancel but sometimes they add up.)

A careful analysis revealed that the largest source of error for the initial drops happened to be a canopy bias. A closer look at Fig.10b reveals the most dramatic control bias. The parafoil is always to the right of the desired path during Phase 4 due to some bias (asymmetry) in controls. This is also observed during loiter. As the terminal phase is reached, instead of turning with of a diameter of 200ft (61m), the actual diameter happens to be 300ft. A simple calculation shows that the error in impact is approximately $(\pi + 1)e_{bias}$, where e_{bias} is the tracking error due to control bias. For Drops #23-2 and #47-2, the bias errors were 40ft (12.2m) and 100ft (30.5m), respectively, contributing 166ft (50.6m) and 414ft (126.2m) of error short.

The second contributor to the error is an estimate of the altitude above the ground and the descent rate, because these dictate the logic of the control algorithm. To this end, the analysis revealed that while for Drop #23-1 the estimate was almost perfect (the system hit the ground when it thought it was going to hit it, i.e. no estimate error), Snowflakes #23-2 and #47-2 experienced some errors. Specifically, for Drop #23-2 the error happened to be $-7m$ meaning that the system planned its final turn and descent based on the wrong altitude, 7m less than the actual one. (For Drop #47-1 the data recording stopped prematurely at 20.7m, so no data on altitude error is available.). With the glide ratio of 2:1, every meter in the vertical direction translates into two meters in the horizontal direction.

Let us do some math for Snowflake #23-2 (since we have all the data available). Although it landed 49.8m short, the controls asymmetry error contributed 50.6m. Hence, without this error we would expect a touchdown error to be $50.6\text{m}-49.8\text{m}=0.8\text{m}$. Now with a glide ratio of 2:1 the -7m error in altitude caused -14m horizontal error. So accounting for these two errors (that could have been avoided) the resulting accuracy would be $0.8\text{m}-14\text{m}=-13.2\text{m}$. Similarly, for Snowflake #47-2, fixing the controls asymmetry error would result in $126.2\text{m}-90.3\text{m}=35.9\text{m}$. Although no solid data is available on the altitude error, assuming the same -7m error (Snowflakes #23-2 and #47-2 were initialized before the drop to the same ground altitude and landed only two minutes apart) would result in $35.9\text{m}-14\text{m}=21.9\text{m}$ touchdown error. These values, -13.2m and 21.9m , match those found in simulations. The first error, i.e. controls asymmetry, can be easily eliminated by introducing an integral term in the Snowflake controller (which was accomplished for all further drops). The effect of the second error, altitude offset, could be mitigated if there was a capability to update the current barometric settings while the ADS was in the air.

However, the major contributors to touchdown errors are the unknown/changing ground winds. The latest wind estimate occurs at about 100m above the ground right before the final turn, and from this point down the entire trajectory is based on this estimate. Obviously, it is not what takes place in practice. The second series of experiments carried out at YPG addressed this source of error in greater detail.

IV. Drops at Yuma Proving Ground

A. Overall Tests Descriptions

Snowflake ADS airdrop tests were conducted at Sidewinder drop zone at YPG on October 20th – 21st of 2008. Figure 15 presents a snapshot of the test site featuring the wind tower that provided updates of the ground winds every minute. The UH-1A helicopter used for the drops was shown in Fig.5.



Fig. 15 Snowflake test site at YPG.

Figure 16a shows another view of the solar powered wind tower transmitting wind information to the YPG ground station residing on the truck shown on the left. About one minute prior to each double ADS drop, the Windpack system was released to measure actual winds. Four of these Windpack systems strapped in the helicopter's bay are shown in Fig.16b. Figure 17 shows the geodetic survey of an impact point.

In the course of two days, eight helicopter lifts were carried out and seven Windpacks were released from 2,500 to 3,000ft AGL. Each Windpack release was followed by deployment of two Snowflake ADSs from an altitude that was 500ft lower than that of the corresponding Windpack. Two targets were assigned in the vicinity of the wind tower and were used depending on the direction of the ground winds.

It should be mentioned that the first set of two ADS drops each day were planned (programmed on the ground before the first lift of the helicopter) with no knowledge about winds aloft, (the first Windpack was released with these two systems and its data was then used to plan the consequent drops). As a result, four Snowflake ADS (two pairs of the very first set of drops each day) were dropped at release points contrary to actual winds, so their GNC systems were never properly activated. Out of eleven remaining systems two got tangled and were eventually put into the manual control mode. The remaining nine systems performed well. The next section provides more specifics on these two-day drops and presents the overall picture of the achieved results.




Fig. 16 Wind recording resources: wind tower (a) and Windpack parachutes (b).



Fig. 17 Geodetic survey of the touchdown point.

B. Drops Results

The trajectories of nine good Snowflake drops look similar to the ones shown in Figs. 8 and 10, two of which are presented and discussed in the next section. The impact points for these nine drops are shown in Figs. 18 and 19 with respect to the downwind leg of the trajectory and the target, as aligned with the intended final approach direction (into the wind). Figure 19 also shows the system ID number along with the direction / magnitude of the ground winds at impact.




Fig. 18 Consolidated drops reduced to the direction of an implied ground wind.




Fig. 19 Consolidated drops reduced to the direction of an implied ground wind (as in Fig.18) with the direction/magnitude of the actual wind at impact.

Impact data recorded by Snowflake onboard sensors (accurate within 3m) revealed the following miss distances (brought together as a vector to be further processed by the MATLAB script):

```
er=[42 90 16 63 10 45 81 40 40];
```

More accurate geodetic survey data provided by the YPG team (recorded with an accuracy of 0.04m) showed:

```
er=[31.61 107.88 16.03 63.83 11.67 46.61 87.69 34.66 34.26];
```

(YPG team data shown in Figs.18 and 19).

The following script, applied to these two vectors

```
er=sort(er)
mall=mean(er) % mean value for all nine drops, m
sall=std(er) % STD for all nine drops, m
m5=mean(er(1:5)) % mean value for the best five drops, m
s5=std(er(1:5)) % STD for the best five drops, m
CEP=median(er) % the median value (CEP), m
```

returned

```
er = 10 16 40 40 42 45 63 81 90
mall = 47.4444
sall = 26.7307
m5 = 29.6000
s5 = 15.3232
CEP= 42
```

and

```
er = 11.67 16.03 31.61 34.26 34.66 46.61 63.83 87.69 107.88
mall = 48.2489
sall = 32.3848
m5 = 25.6460
s5 = 10.9410
CEP= 34.6600
```

respectively. These results are consolidated in Table 2. (Note that as opposed to the common definition of CEP, where it is defined as a circle, centered about the mean, whose boundary is expected to include 50% of the population within it, we defined it as a circle, centered about the intended impact point, whose boundary includes 50% of the population within it.)

Table 2. The performance of Snowflake ADS.

	Snowflake (less accurate) data		YPG surveillance (very accurate) data	
	Nine drops	Best five drops	Nine drops	Best five drops
Average miss distance, m	48	27	48	26
Standard deviation, m	30	15	32	11
CEP, m		42		35

As compared to the drops at Camp Roberts, the CEP of the YPG drops was reduced by 30%, down to 35m. That was accomplished by eliminating just one of the three contributors to the touchdown error, canopy asymmetry. An integral controller introduced to the algorithm greatly reduced the canopy bias. Also, as seen from Figs. 18 and 19, due to the unaccounted for crosswind component, all impact points lie to the right of the intended landing direction. It should be noted that this was done intentionally (the intended landing direction was set slightly off the current ground wind direction to accommodate for possible change of the ground winds before the deployment of the system). Specifically, these offsets happen to be as presented in Table 3 resulting in the crosswind component at impact shown in Table 4. As seen, the spread in the crosswind direction is small, and in fact is of no concern compared to the spread in the upwind direction.

Table 3. Ground winds offsets as entered into the system before deployment.

Drop number	3	4	6	7	8	11	12	14	16
Ground wind offset, °	44	45	20	5	1	25	38	6	13

Table 4. Crosswind component at impact.

Drop number	3	4	6	7	8	11	12	14	16
Crosswind component, m/s	0.778	0.728	4.354	4.051	3.243	3.640	3.124	3.123	2.759

C. Detailed Analysis of the Best and Worst Drops

As mentioned previously, nine systems behaved as they were supposed to. Upon release they transitioned to the loitering area, made several passes in estimating the current magnitude of the winds (along the predetermined final approach direction), and decided when to exit towards the target. Upon exiting they glided towards their target and then made a final turn, landing into the (presumed) wind.

With perfect knowledge of the winds, all Snowflakes would land almost exactly on the target. Unfortunately, the only winds available onboard are the wind estimates, provided by the GNC algorithm in real-time during two phases, loitering and downwind descent. Of course these estimates are not perfect, and the major problem is as follows. The last estimate on the wind magnitude comes at the end of the downwind leg, i.e. the construction of the final turn trajectory is based on this last update, obtained at an altitude of about 100m AGL. If the winds below 100m are the same as estimated at 100m then the system would land in close vicinity of the target, where the only source of error then would be an error in estimating winds at 100m. Of course, having the same winds throughout the last 100m of the drop is rarely the case (if ever at all). While the NPS-UAH team is currently working on the adaptive algorithms to accommodate these unknown variations in winds to be implemented during the next set of drops (in May of 2009), let us consider the achieved results from exactly this standpoint – how accurate the estimates of the winds were at 100 m and how much of a wind change the ADSs had to face on their final 100m to the ground. The data collected during the drops by Windpacks and wind tower were carefully analyzed and presented for two out of nine drops below. Analysis of these data allows one to really get into the cause of a downwind miss.

First, Figs. 20 and 21 show the complete sets of the recorded data for two of nine drops, the worst one (Fig.20) and the best one (Fig.21), respectively. These are Drops #4 and #6 as shown in Fig.19. The time difference between them is slightly over one hour.




Fig. 20 Snowflake performance for Drop #4.




Fig. 21 Snowflake performance for Drop #6.

Figure 22 presents the data collected by the wind tower. First, one can observe the drift in the barometric pressure that the estimates of the ADS altitude are based upon. As seen, the onboard barometric pressure sensor settings have to be changed constantly to accommodate these changes in the ground pressure. Failure to do that results in the ADS thinking that it is higher above the ground level than it actually is (because the barometric pressure drops down as the sun heats the surface). Secondly, during the four sets of drops the winds did change direction and speed. Starting with light winds in the morning (below 1 m/s) the ground winds got stronger up to about 4 m/s for the last sets of drops. Also, one can observe a drastic change of the wind direction before and during the first set of drops on each of these two days. The winds aloft changed their direction also and that is the reason the first set of drops failed to perform properly, i.e. the release point was computed based on the current ground winds only, and the morning desert winds demonstrated erratic behavior. After about 8:30-9:00 am the ground winds became more consistent (in the wind direction).

Next, Fig.23 presents the data collected by the Windpack a couple of minutes before the first (in a set of two) ADS release. In this figure they are presented as recorded by the Windpack, where altitude is provided with respect to the mean sea level. It is seen that during the last 100 m the winds do change and that is what causes touchdown errors. Figure 24 addresses this issue in more detail, providing the comparison of the Windpack data with that of the wind tower and Snowflake estimates. (In this figure, the Windpack altitude data is converted from MSL to AGL; and the Snowflake ADS altitude data is corrected to accommodate errors in altitude estimates, i.e. all data is altitude-synchronized.)

Figure 25 presents a more detailed analysis of the data shown in Fig.24. Specifically, the last phase of the guided descent, final turn to the target, is considered. The left portion of each plot shows the difference between the tailwinds predicted by the Snowflake ADS and actual winds measured by the Windpack, and the crosswind component that the ADS had to overcome on its way down. On the right portion of each plot the vertical speed as

recorded by the Snowflake ADS barometric altimeter is presented. Several calculated parameters shown on these plots represent a rough estimate of contribution that unaccounted for winds in the downwind direction might have on overshooting the target.




Fig. 22 Wind tower data with vertical lines indicating Snowflake releases.




Fig. 23 Windpack data valid for Drop #4 (a) and Drop #6 (b).




Fig. 24 Windpack 2 data versus Snowflake estimate for Drop #4 (a) and Windpack 3 data versus Snowflake estimate for Drop #6 (b).




Fig. 25 Downwind and cross wind components of unaccounted winds for Drop #4 (a) and Drop #6 (b).

Consider Fig.24a. As seen from Fig.20, the last estimate of the along-the-course wind component by the Snowflake autopilot occurred at about 93.6m corrected altitude (109.1m – 15.5m) AGL, and provided the value of 5.3m/s. However, the real winds suddenly die and amount to only about 1 m/s all way down to the ground. Hence, there is a 4 m/s difference between the last Snowflake estimate and the actual winds (Fig.25a) all way down. These unaccounted for tail winds that acted during the last 27 s of the guided descent result in a large system overshoot of about 116 m, (which in this case would be even bigger with no error in altitude estimate).

Now, consider Fig.24b (Snowflake drop #6). Compared to the previous case, the wind estimate at the corrected 79m (86.9m – 8.23m) AGL (Fig.21) is almost perfect (5.2m/s) and the winds do not change much during the final turn (Fig.25b). This results in landing only 5 m short of the intended point of impact. The data for these two Snowflake drops is consolidated in Table 5. This table shows upwind and crosswind components of the miss distance and then analyses the possible contribution of two sources of error, which are an altitude estimation error and unaccounted for winds during the last phase of the flight, the final turn to the target. The last column shows the values of a miss distance that could be achieved if both errors were eliminated.

Table 5. Estimates of possible causes for over- / under-shooting the target.

Drop ID	Miss distance components, m		Altitude error, m	Correction due to the altitude error, m	Correction due to unknown winds, m	Estimated upwind miss distance, m
	Upwind	Crosswind				
4	102.13	34.77	15.50	31.00	-116.00	17.13
6	-5.11	15.19	8.23	16.46	-1.44	9.91

Another important finding is that the ground winds recorded by the wind tower at the moment of impact match those recorded by the Windpack fairly well. Further development of the Snowflake GNC algorithm could include a constant update of the ground winds to produce a more reliable wind profile during the final turn, along with the updates of the current barometric pressure to eliminate errors in estimates of the altitude above the ground.

V. Conclusions and Further Development

The Camp Roberts and YPG drops of the Snowflake ADS proved its overall concept and reliably. Of about 20 Snowflake drops, there were no major failures that would result in system loss. Specifically, the following performance was last demonstrated at YPG for the nine fully autonomous operable systems released properly:

- mean miss distance was 48m;
- best 5 out of 9 drops landed within 26m from the target;
- achieved CEP was 35m.

Further development should take into account uncertainties of the winds during the last ~100m (~25s) of system descent, and errors caused by drifting barometric pressure at the ground.

Acknowledgments

The authors would like to thank the USSOCOM-NPS Tactical Network Topology Field Experimentation Cooperative, and personally Prof. David Netzer, for funding this project and providing capability to test our system at the California Army National Guard training base. The authors would also like to extend their gratitude to the YPG testing team, and personally Mr. Robert Berlind of the Air Delivery Systems Branch (who passed away untimely on January 29, 2009), for their generous support in providing testing facilities during the drops at YPG. Finally, the authors would like to thank, Marina municipal airport manager, Mr. William Keeler, and the Center for Interdisciplinary Remotely-Piloted Aircraft Studies (CIRPAS) for allowing using the airport ramp for initial tests in early May of 2008.

References

¹ Slegers, N.J., and Yakimenko, O.A., “Optimal Control for Terminal Guidance of Autonomous Parafoils,” *Proceedings of the 20th AIAA Aerodynamic Decelerator Systems Technology Conference*, Seattle, WA, May 4-7, 2009.