PRELIMINARY FLIGHT TESTING OF DRAGONFLY'S ELECTRO-OPTICAL TERRAIN SENSING FUNCTION

Ike R. Witte[1*], Steve N. Jenkins[1], Justin R. Thomas[1], Samuel E. Bibelhauser[1], Nishant L. Mehta[1]. [1]Johns Hopkins University Applied Physics Laboratory 11100 Johns Hopkins Road, Laurel, Maryland 20723.
*Isaac.Witte@jhuapl.edu

Abstract. *NASA's Dragonfly mission is designed to autonomously fly a rotorcraft over the surface of Titan, acquiring scientific data at interesting sites distributed across a wide area of Saturn's moon. The Electro-optical Terrain Sensing (ETS) function of Dragonfly's Mobility subsystem provides optical measurements to the onboard navigation filter, enabling in-flight terrain-relative navigation without any a priori reference maps. Members of the Dragonfly team recently conducted a series of flight tests at the Imperial Dunes in California to quantify the performance of an Earth-analogue of the Mobility subsystem's preliminary design and analyze areas of sensitivity. This paper summarizes the design of the ETS function, describes the Integrated Test Platform (ITP) that was flown during the test campaign, and describes the results and lessons learned from these flight tests.*

Introduction. NASA's Dragonfly mission to Titan is a mobile rotorcraft lander mission to characterize the surface composition of Saturn's moon to investigate prebiotic chemical processes in this carbon-rich extra-terrestrial environment [1]. To facilitate this objective, Dragonfly's Mobility subsystem is designed to autonomously guide, navigate, and control the rotorcraft during its flights across the surface of Titan. The Electro-Optical Terrain Sensing (ETS) function is a component of this subsystem that provides terrain-relative optical measurements to an onboard navigation filter[2]. Using these measurements, along with a suite of navigation sensors, Dragonfly is able to autonomously navigate relative to surface features on Titan, enabling precise trajectory tracking and landing.

A unique feature of Dragonfly's mission is autonomous flight over terrain that has not been previously imaged or scouted at relevant length scales. Radar imagery from Cassini's flybys of Titan are the best direct data available of the morphology surrounding Dragonfly's landing ellipse in the Shangri-La dune fields [3]. However, digital elevation models (DEMs) of the terrain that allow for absolute localization, or at the very least, simulated performance assessment, are not available. This constraint has driven two major aspects of the optical navigation algorithm for Dragonfly: First, the concept of a "breadcrumb" image was developed as a method to navigate relative to previously seen terrain on a return flight. Second, an image registration algorithm was selected for robustness to the unknown distribution of optically resolvable features.

Due to the image registration algorithm sensitivity to feature quality in the imagery, it is necessary to evaluate the performance of ETS in an environment as similar to the operating regime as possible. The team both rendered imagery with terrestrial digital terrain and realistic Titan lighting models developed using the data we do have, and tested ETS performance against real imagery captured over stressing terrain [4]. The Imperial Sand Dunes in southeastern California were selected as the best Earth analogue to Dragonfly's destination accessible given travel logistics. Thousands of images were collected during a series of flights of Dragonfly's Integrated Test Platform (ITP) over the Imperial Sand Dunes under a wide range of locations, illumination conditions, poses, and attitude rates. These data were used to thoroughly assess the preliminary design of the ETS algorithm and its interactions with the navigation function of Dragonfly's Mobility subsystem. The results are summarized in the subsequent sections.

Algorithm Description. While Voyager, Cassini, and Huygens have all helped to determine the environment on Titan the terrain is still relatively unknown at the length scales needed for low altitude surface flights[5][6][3][7]. Due the range of possible surface features and visual conditions, the optical navigation algorithm inside Dragonfly's mobility subsystem must be robust to feature-poor imagery while meeting the accuracy requirements to ensure successful navigation to locations of interest. The ETS function design, at its core, employs a phase-only normalized cross-correlation algorithm to produce both odometry and SLAM-inspired "breadcrumb" measurements. These measurements are provided to an extended Kalman filter, which along with additional measurements from a suite of sensors, estimates the full state of the rotorcraft in various Titan-fixed frames [8]. In order to enable real-time processing, the ETS algorithm is implemented across software and firmware, using an FPGA to accelerate the image processing on the Optical Vision Processing (OVP) system. The bulk of this design has been previously described [2], but for completeness and to illustrate updates and improvements the algorithm flow will be reviewed. A high level overview of the process is shown in figure 1.

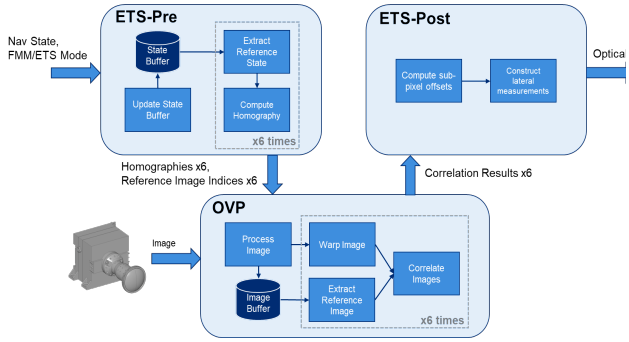

Figure 1. Overview of the ETS function, including acceleration by the OVP. In the figure, the term "Process Image" includes binning, calibration, normalization, and distortion correction.

Prior to phase correlation, each incoming image undergoes several preprocessing steps. First, the image is 2x2 binned to 512x512 pixels, and the image is normalized to zero-mean. Following normalization the lens distortion is corrected using Brown's sixth-order polynomial barrel distortion model [9], characterized in equation 1 where x_0 and y_0 are the distorted pixel locations, r is the radial distance from the distorted pixel location to the camera principal point, k_1, k_2, and k_3 are the calibrated geometric distortion coefficients, and x_{corr} and y_{corr} are the corrected pixel locations.

$$x_{corr} = x_0(1 + k_1 r^2 + k_2 r^4 + k_3 r^6) \qquad (1)$$
$$y_{corr} = y_0(1 + k_1 r^2 + k_2 r^4 + k_3 r^6)$$

While these calibration and preprocessing steps are occurring on the OVP, ETS is simultaneously determining which reference image to correlate to. Due to the lack of an available digital elevation map of the Titan terrain, these reference images, along with their state estimates, are taken from a database of imagery captured earlier in flight or even on a separate flight. ETS uses this database to store images that are used in parallel for both localization (breadcrumbs) and velocimetry. An overlap estimation algorithm is employed to select reference images with a large expected overlap with the current image. Additional criteria for selection are used to provide the navigation filter with the most useful information possible. Velocimetry reference images are selected by choosing the oldest reference image that is still in the camera field of view. For breadcrumbs, the image with the maximum estimated overlap is chosen as the reference. Breadcrumbs from the current flight are designated "online" breadcrumbs, breadcrumbs from previous flights are designated "historic" breadcrumbs, and breadcrumbs with the desired landing site in the field of view are designated "terminal" breadcrumbs.

Once reference images are chosen, a set of homographies are computed in order to warp the current image into the reference image's frame, based on the a priori state estimates of the two images. The homography is computed using equation 2:

$$H_1^2 = AR_I^2(I_{3x3} - \frac{T_1^2}{d_2} n^T) R_I^{1^T} A^{-1} \qquad (2)$$

Where A is the intrinsic camera calibration matrix, R_I^X is the rotation matrix from the navigation frame to the camera frame of image X, I_{3x3} is a 3x3 identity matrix, T_1^2 is the 3D translation vector from camera pose 1 to camera pose 2, d_2 is the perpendicular distance from pose 2 to ground plane, which may be sloped, and n is the normal vector of the ground slope. The mean terrain slope can be estimated by the navigation filter by incorporating LiDAR range measurements. If the slope is flat and level (n = [0 0 1]), then d is simply the height of the rotorcraft above the ground level (AGL).

The hardware-accelerated image processing can perform several correlations can occur within the 1Hz time frame, which allows for multiple correlations focused on the breadcrumb measurement in order to maximize the likelihood of a successful correlation. Breadcrumb a priori state estimation errors are typically both larger and more uncorrelated than velocimetry measurements, resulting in more stressful correlations with lower amounts of overlap between the images. These extra correlations can be used to alter the state estimate, such as lateral position or altitude, or could be used to match portions of the same image to perform large image patch correlations between the two images. The current baseline, and the flight test results shown in this paper, uses the extra correlations to perturb the position estimate in down position, but future work will compare this approach to coarse patch correlations. At the time of testing, the ITP was unable to perform localization and velocimetry measurements in parallel, so for these flight tests the results shown produce one measurement or the other, but not both concurrently.

For each homography, the current image is warped into the reference image's frame using a come-from homographic warping. A come-from warping scheme iterates through the pixel locations of the destination image, uses the homography to determine where in the source image the pixel information will come from, and performs bilinear interpolation with the closest four pixels. A go-to warping scheme, on the other hand, iterates through the source image and directly maps the pixel information to the destination image. While more computationally complex, ETS employs come-from warping to avoid gaps in the warped image. Figures 2 and 3 shows images prior to and after warping into reference image frames from a recent flight test.

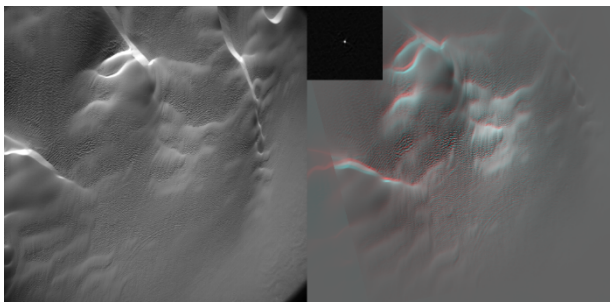

Figure 2. Demonstration of a velocimetry correction from flight tests over the Imperial Sand Dunes, May 2022. The current camera image (left) is warped into the reference image's frame (right) and the two are correlated (top). The right image shows a red-blue image comparison between the warped current image and the reference image, both post-windowing. Where the two align is grey, while where the two differ is either blue or red. Velocimetry measurements typically contain less input relative state knowledge error, so the images are well aligned and there is a clean correlation surface.

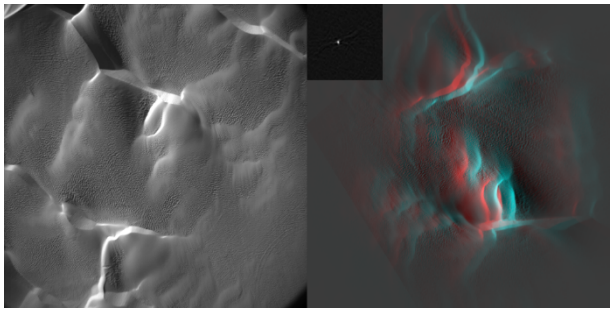

Figure 3. Demonstration of a breadcrumb correction from flight tests over the Imperial Sand Dunes, May 2022. Similar to figure 1, the current camera image (left) is warped into the reference image's frame (right) and the two are correlated (top). As the reference is a breadcrumb, the heading is close to the opposite direction of the reference image. In this case, there is a clear left-right shift between the warped current image and the reference image, which is recognized in the shift of the peak in the correlation surface. This is due to accumulated drift error between the time that the breadcrumb image was saved and the time it is used as a reference image. By successfully correlating to the breadcrumb, this accumulated drift error is reduced to the accuracy of the correlation, and subsequent breadcrumbs will show a closer alignment.

With the two images in the same reference frame, both images are windowed prior to correlation in order to avoid edge effects in the correlation surface. ETS and OVP support a flexible windowing scheme through an outer product of a 1D vector of coefficients. The baseline coefficients define a Hann window, but the flexibility allows ETS to use other windowing schemes such as Hamming, Tukey, or more custom windows such as a Hann window centered on a sub-image patch.

The OVP implements a phase-only cross correlation in frequency space to match the warped current image to the reference image and compute a delta between the two images in pixel space. The discrete 2D Fourier transforms are computed for both images, and the cross-power spectrum is taken by multiplying elementwise the Fourier transform of the first image by the complex conjugate of the Fourier transform of the second image. Rather than weight each spatial frequency in the correlation by its magnitude, phase-only correlation normalizes each "pixel" in the spectrum prior to applying the inverse Fourier transform. In addition, ETS and NCP support a frequency filter that is multiplied element-wise to the spectrum to selectively apply a weight to each frequency. The result is a correlation surface, C, shown in equation 3, defining the correlation strength between the two images under a lateral shift:

$$C = \mathcal{F}^{-1}\left\{\frac{\mathcal{F}\{Img1\} \circ \mathcal{F}\{Img2\}}{|\mathcal{F}\{Img1\} \circ \mathcal{F}\{Img2\}|}\right\} \qquad (3)$$

Where $\mathcal{F}$ denotes the Fourier transform, $Img1$ and $Img2$ denote the two images being correlated, and $\circ$ denotes an element-wise product. Following the computation of the correlation surface, both the peak and secondary peak are found, and the peak values, locations, and local surrounding neighborhoods are returned for sub-pixel fitting. An area around the peak location is ignored when searching for the secondary peak to ensure that both values do not come from the same correlation peak. The ratio of the value of the secondary peak to that of the primary peak is used to ensure a quality correlation has occurred. The peak ratio threshold for a valid correlation is set to a maximum of 0.6. If the peak ratio is >0.6, the identity of the correct solution is ambiguous and the correlation is marked as invalid. This can happen with multiple separate peaks or with a very broad single peak. Correlation surface information is found and passed out of the NCP for each homography in the set given as inputs. For breadcrumbs, where multiple correlations occur for the same image with difference scales, the solution resulting with the lowest peak ratio is chosen.

ETS produces measurements in the camera frame and in units of meters, so the 2D delta pixel result must be converted into meters. This conversion is done using equation 4:

$$\begin{bmatrix} d_x \\ d_y \\ d_z \end{bmatrix} = C_F^{2^T} \cdot A^{-1} \cdot \begin{bmatrix} d_i \\ d_j \\ 0 \end{bmatrix} \cdot range_2 \qquad (4)$$

Where d_x, d_y, and d_z are the measurement deltas in the desired frame F, C_F^2 is the transform from F to the reference image frame, A is the intrinsic camera calibration matrix, $range_2$ is the LiDAR range for the center of image 2 (essentially the image scale), and d_i, d_j, are the deltas in pixel space. To return the measurement in the frame of camera 2, C_F^2 is simply identity. The delta X and Y measurements are combined with the a priori pose estimate to compute the final measurement. These measurements are ingested into the navigation filter, along with information about the reference imagery to aid in the management of breadcrumb states.

Integrated Test Platform. As part of a process to validate the Mobility subsystem, the Dragonfly team built an Integrated Test Platform (ITP): a half-scale rotorcraft testbed designed to fly on Earth. The ITP allows for a sensor and processor box to be attached to its base that can control the rotors and autonomously fly the rotorcraft. The goal of the ITP is to serve as an early hardware-in-the-loop testbed to evaluate the Mobility subsystem in parallel with the truth model implemented as part of a closed-loop simulation (CLS) environment in MATLAB/Simulink. The same Mobility algorithm software that can be tested in the CLS, with the exception of an FPGA accelerated image processing component in place of a software emulation in the CLS, is compiled onto the sensor and processor box.

With this platform, a variety of flight modes and configurations are available, from remotely piloted with the Pixhawk autopilot, a commercial-off-the-shelf product, to controlled autonomously by the box throughout the entire flight, including autonomous takeoff and landing. While thus far, a select few components flown onboard the ITP are not flight-like hardware, e.g. a laser range finder in place of a LiDAR, the ETS function onboard the ITP operates much the same as the preliminary design for the Titan lander. The platform enables the Dragonfly team to gather performance data impossible to generate with a closed-loop simulation, and learn critical lessons in operating an autonomous rotorcraft.

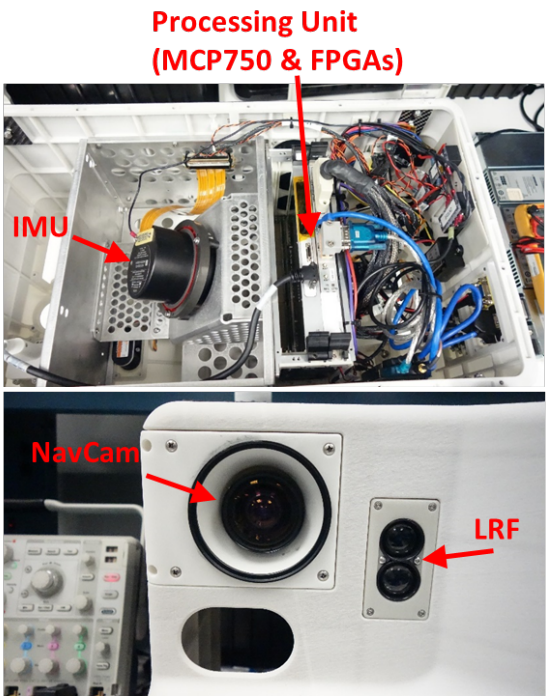

Figure 4. Labeled images of the ITP drone box with its various components, including processing units, navigation cameras, laser range finder (LRF), and IMU. The full drone box, with its sensor suite, weighs about 35lbs.

The fully assembled drone box for the ITP is a two cubic foot, 35lb box that emulate sensors as shown in table 1.

Table 1: Comparison of ITP hardware to hardware planned for the Dragonfly Lander.

Component	Integrated Test Platform	Dragonfly Lander Baseline
Navigation Camera (NavCam)	Basler acA2040 CMOS FPA = 2048 x 2048 pixels Pixel size = 5.5 mm FOV = 90 deg	Malin LCAM CMOS FPA = 2048 x 2048 Pixel size = 4.8 mm FOV = 90 deg
NavCamera Lens	Kowa LM6HC f/# 1.8~16 (8 for ITP) f=6mm	Custom Design f/# 2.7 f=5.75mm
NavCam Pointing Orientation	Nadir Pointing	10° Forward Pitch
Single Board Computer	Motorola MCP750 233 MHz, 1 MB L2 cache optional† 128 MB RAM	BAE RAD750 V3 198 MHz, No L2 cache 1GB RAM
Navigation Coprocessor	COTS Virtex-5 FX130T, 70 MHz 512 MB DDR2 SDRAM 32 MB DDR SSRAM	Rad-tolerant Xilinx KU060, 80 MHz 2 GB DDR3 SDRAM 32 MB QDR SSRAM
LiDAR	Lightware sf11/c Laser Rangefinder Used for range to center of image	NASA GSFC Ocellus scanning LIDAR Used for range and 3D terrain mapping

Although effort has been taken to emulate a Dragonfly flight as closely as possible, there are several hardware and environmental differences that affect the ETS algorithm. Titan lighting conditions are much dimmer and more diffuse than typical terrestrial lighting conditions, which means that at certain times of day the ITP will see the drone shadow in its imagery. Self-shadow is a stressing image feature that is not anticipated in imagery during Titan flights. The lighting difference also causes the ITP integration times to typically be shorter. Testing was conducted during dawn and dusk to mitigate these affects, with the dusk flight most closely emulating the environment with dim imagery, no drone shadow or low-angle ripple-shadowing on the dunes, and integration times approaching expected Titan integration times.

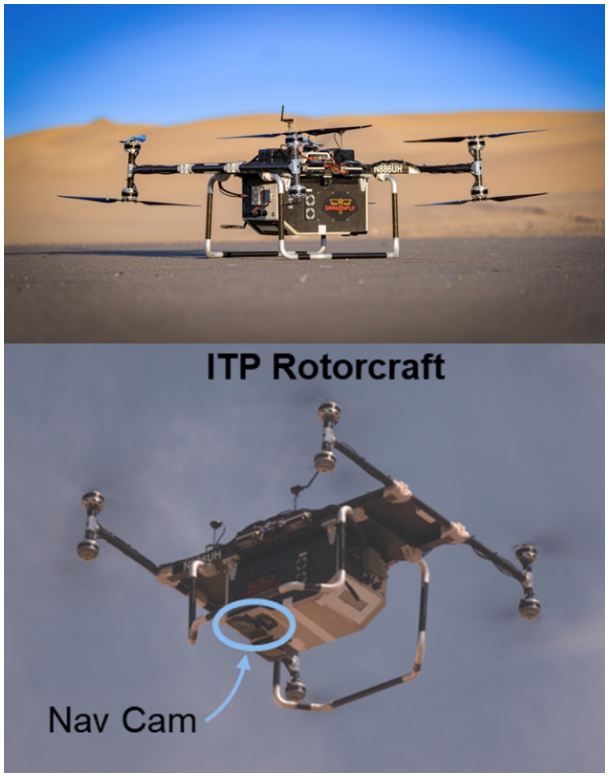

Figure 5. Images showing the full ITP, with drone box attached. The Imperial Dune fields are visible in the background in the top image, while the NavCams are visible in the bottom image.

Another difference between the ITP and the Dragonfly baseline is that the ITP employs a laser rangefinder rather than a LiDAR for image scale information. This has the effect of also precluding the use of terrain slope information in the ETS algorithm, as the range finder returns only the range to the center of the image rather than a slope estimate of the terrain within the FOV of the LiDAR. When slope information is not available, the ETS algorithm assumes that the underlying terrain is flat and level when computing the homography, leading to degraded performance over dune plinths and crests.

Additional flight differences between ITP and the anticipated Dragonfly lander baseline are detailed in table 2, including vehicle dynamics, flight envelope, and sensor alignment tolerances. In the ITP, the plastic and carbon fiber materials, in addition to the interchangeability of sensors and lenses, leads to greater body-to-sensor pointing knowledge errors, which can adversely affect the interpretation of the correlation results. Pointing knowledge for the actual lander is anticipated to be far more accurate.

Table 2: Differences between ITP flights and anticipated Dragonfly flights, including environmental differences.

Characteristic	ITP	Dragonfly Lander Baseline	ETS Implication
Camera Alignment	Alignment tolerances are lower	Alignment tolerances much more accurate (<= +/- 0.5° TBR)	Alignment knowledge errors manifest as ETS error
Lighting	Direct / Brighter Self Shadowing	Diffuse No Self Shadowing	Lower camera integration times on ITP, although some Yuma flights at the expected Titan values. Self-shadowing was problematic during ITP flights
NCP/OVP	Slower Image Processing	Faster Image Processing	ETS Timing for Lander is more favorable
Ranging	Laser Rangefinder	LIDAR	Slope knowledge more accurate for Lander, higher ETS performance expected
Flight Duration & Distance	~10 Minutes	30 Minutes	Breadcrumb Utilization (ITP flights don't use/need as many breadcrumbs)
Flight Altitude	100 meters	400 meters, up to 2km	ETS Performance (ITP ETS errors higher due to lower altitudes)
Attitude Dynamics	Faster, With More Disturbances	Slower, With Less Disturbances	Camera Motion Attitude/Rate Violations

In spite of these differences, the ITP has proven to be very effective in the development and testing of ETS. The existence of the ITP allows many of the Mobility subsystem flight algorithms to be flown and tested early and often. ETS can be tested closed-loop during substantial flights over relevant terrain with relevant hardware, and the successes and shortfalls during these flights are used in combination with simulation Monte Carlos to improve the design and implementation to create as robust a design as possible for the future lander. While not a Mobility requirements verification testbed, ITP serves as a technology validation platform that can highlight sensitivities to effects not yet fully modeled in simulation.

Flight Testing. In May of 2022, members of the Dragonfly team conducted a series of flight tests of the ITP at the Imperial Sand Dunes in southeastern California, gathering data on the performance of the subsystem's preliminary design. An image from ITP in flight is shown in Fig. 6. The Imperial Dunes were selected as a promising Earth-analogue to the expected morphology of Titan's Shangri-La dune fields. This section will describe the test campaign, show ETS performance over the flights, and discuss the lessons learned from the process.

Figure 6. Image of the Integrated Test Platform flying over the Imperial Sand Dunes.

Flight Tests. Flights were conducted across three days in a combination of three modes: manually piloted, Pixhawk autopilot controlled, and fully autonomous. In all modes, the sensor and processor box ran the full algorithm suite, even in the cases where the Mobility controller was disconnected. The flights occurred in all hours of sunlight, over a wide range of illumination conditions, culminating in four fully autonomous flights, roughly ten minutes each, flown in the early morning of the third day. Figure 7 shows a range of image conditions seen during the flight test campaign.

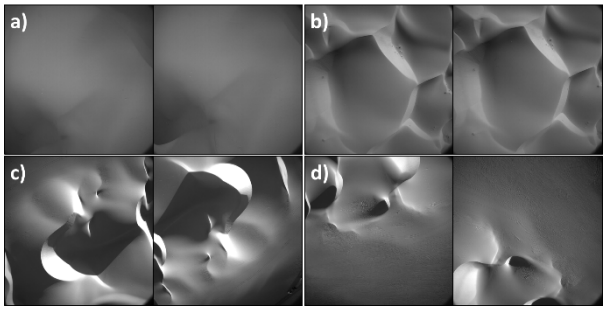

Figure 7. Four image pairs showing range of illumination conditions and feature content from across the flights on 5/12/2022. Image pair a) shows a velocimetry image pair during the dim lighting dusk flight, b) shows a velocimetry image pair of dune crests in a morning flight, and images c) and d) show breadcrumb image pairs with large heading changes during morning flights.

Flight tests during the first two days consisted of a series of checkout flights starting with Pixhawk and pilot control and progressing to fully autonomous flight. Checkout flights were performed to ensure nominal ITP performance after shipping the drone across the country, and to also check the stability of the optical measurements over the terrain. ETS was enabled during the checkout flights in the background to ensure performance was nominal prior to influencing the navigation and control of the rotorcraft. Following checkout flights, full autonomous flights were performed, but due to the time of day of the flights the drone shadow was present in the imagery. As will be discussed later in section C, self-shadowing is an optical artifact, unanticipated in the actual Dragonfly mission, that can cause biases and invalid measurements due to incorrect correlations. Plans were made to focus the third day flights at dawn and dusk to avoid self-shadowing and also to collect data during low light conditions, such as the imagery shown in figure 8.

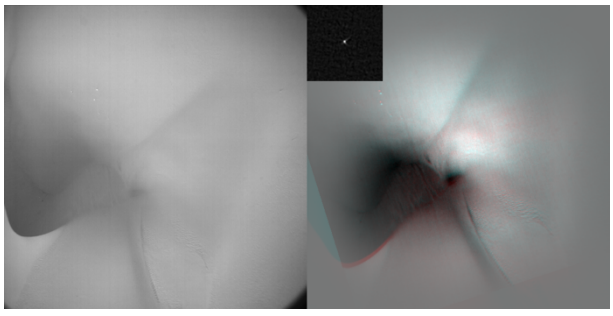

Figure 8. Example image from an ITP flight in the Imperial Dunes (left). Warped image, preprocessed for correlation and fused with reference image for display purposes (right). Correlation surface computed from the ETS algorithm showing a clear peak location (top).

During the test campaign, various flight profiles were chosen to reflect the types of flights anticipated on Titan: scout, leapfrog, and loop. As seen in figure 9, a scout flight takes off, travels a distance, turns around, and retraces its steps back to the takeoff location. ETS measurements during this type of flight will include velocimetry measurements and online breadcrumb measurements, as there are no historic breadcrumbs from a previous flight. Scouting flights scout a potential landing site for the next flight and return. The leapfrog flight builds on the scouting flight by traveling to the scouted location and then continuing on, scouting a new potential landing site, and then coming back to land at the site that the scouting flight found. Leapfrog flights employ velocimetry, online breadcrumb, and historic breadcrumb measurements. Finally, loops are a flight type that travels a distance without a significant period of path retracing, landing in a different spot than takeoff. Loops are not anticipated for use on Titan, but were useful here to produce a stressing, long-distance trajectory that overwhelmingly relies on velocimetry performance.

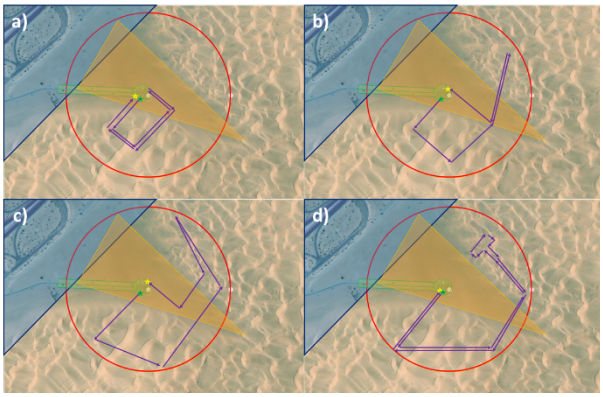

Figure 9. Flight profiles (purple) overlaid onto a map of the Imperial Sand Dunes, showing the maximum allowed range from takeoff (red) and areas with dune buggy tracks (yellow triangle) to avoid. Profiles are a) medium scout, b) medium leapfrog, c) long leapfrog, and d) long scout.

Results. For the flights on May 12, the third day of testing, ETS flights of interest were performed at 06:05, 06:38, 7:04, 7:44, and 19:30 local time at the Imperial Dunes in California. As table 3 shows, the morning flights consisted of two scouts and two leapfrogs, while the final flight was a long loop. The final flight was flown using the Pixhawk controller rather than the drone box, and collected valuable data of very low light conditions over bland terrain.

Table 3: Overall navigation performance for the May 12 ETS flights.

Local Time	Flight Profile	Control	Flight Time (sec)	Horizontal Flight Distance (m)	Landing Accuracy (m)
06:05	Medium Scout	Drone Box	460	1223	2.1
06:38	Medium Leapfrog	Drone Box	439.7	1428	2.2
07:04	Long Leapfrog	Drone Box	507.2	1705	11.8
07:44	Long Scout	Drone Box	672.2	2454	6.0
19:30	Long Loop	Pixhawk	624.7	2437	0.8

As ETS produces a measurement into the larger navigation filter, the focus of this paper is the performance of ETS rather than as the system as a whole. However, landing accuracies for the flights are shown for completeness, and will be expanded upon in a future paper detailing the overall navigation performance. Overall navigation accuracies were encouraging, with the ITP flying over a 1.4km leapfrog and landing within 2.2m of the desired landing site, and flying a 2.4km scout 6m landing accuracy.

To quantify ETS performance, measurement errors were analyzed in pixel space and split between velocimetry measurements and breadcrumb (localization) measurements. There were a total of 1995 valid velocimetry measurements and 389 online breadcrumb measurements across the flights analyzed from 5/12/2022. Upon returning from the field, it was determined that the camera intrinsic calibration matrix used during the flight tests was inaccurate, which affects measurement accuracy. Internal analyses have shown that principal distance calibration error leads to a velocimetry bias and an increased standard deviation of the breadcrumb error measurements, while a principal point calibration error leads to breadcrumb measurement biases. ETS performance using the as-flown and re-calibrated parameters is compared using a suite to replay the image data, with the re-calibrated parameters reducing biases and the spread of breadcrumb measurements.

While accurate intrinsic calibration reduced the biases, there was still an evident cross-track bias in breadcrumb measurements that was not present in velocimetry measurements. A likely source of this error is the body-to-camera pointing knowledge error. Velocimetry measurements are generally facing the same heading, so an angular error would be correlated and not result in as large of a bias as breadcrumb measurements, which in these flights typically had a heading difference close to 180 degrees. With the NavCam's geometry, even a 2 degree fixed angular knowledge error can lead to measurement errors of over 50 pixels (unbinned) due to the 180 degree heading difference. A 1 degree pointing knowledge error is feasible for the ITP due to the flexible materials of the drone box and the ability to swap sensor

in and out with hand tightening of fasteners. Using this knowledge, it was estimated that the camera pointing knowledge error was 1.4 degrees in body roll direction. Correcting for this value, the breadcrumb cross track biases largely disappear, and ETS performance with the new calibration parameters and updated sensor pointing angles is compared to the as-flown and new-calibration results in figure 10 for velocimetry measurements and 11 for online breadcrumb measurements, with the performances summarized in table 4.

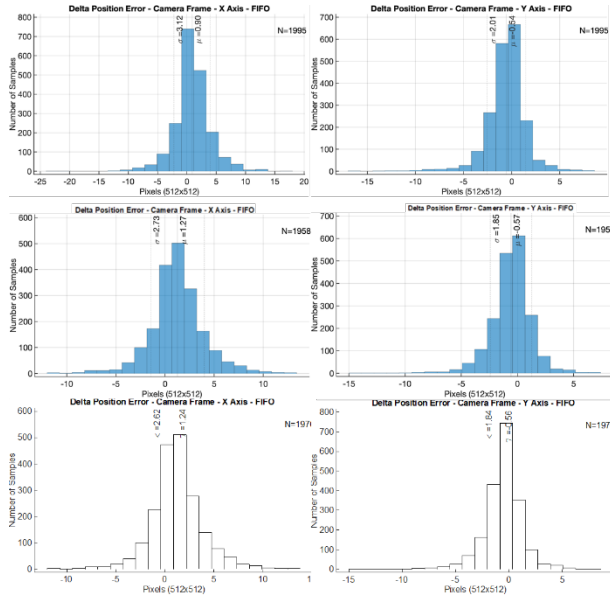

Figure 10. ETS measurement errors from across the 5/12/22 flights for velocimetry measurements as flown (top), with updated camera calibration (middle), and updated camera calibration with estimated camera pointing angle (bottom). Errors are shown in pixel space.

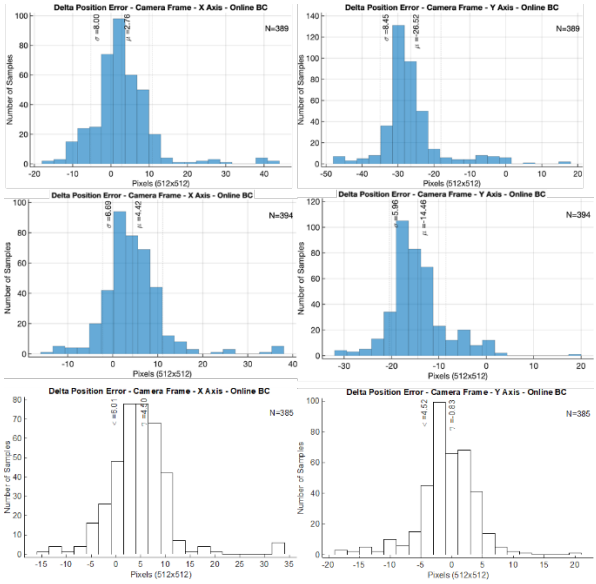

Figure 11. ETS measurement errors from across the 5/12/22 flights for online breadcrumb reference image comparisons as flown (top), with updated camera calibration (middle), and updated camera calibration with estimated camera pointing angle (bottom). Errors are shown in pixel space.

Table 4: Comparison of velocimetry (Vel.) and breadcrumb localization (OBC) measurements as flown, with updated intrinsic calibration parameters (Cal), and with updated calibration parameters along with estimated camera pointing angles (Cal + Angles).

Error Statistics	Vel. As Flown	Vel.: Cal	Vel.: Cal + Angles	OBC: As Flown	OBC: Cal	OBC: Cal + Angles
CamX Mean (m)	0.90	1.27	1.24	2.76	4.42	4.50
CamX STD (m)	3.12	2.73	2.62	8.00	6.69	6.01
CamY Mean (m)	0.54	0.57	0.56	26.52	14.46	0.83
CamY STD (m)	2.01	1.85	1.84	8.45	5.96	4.52

Replaying the trajectory using the updated calibration parameters nearly halved the online breadcrumb bias from 27 pixels to 14 pixels, which was further reduced by estimating the camera pointing knowledge error. Taking advantage of the updated camera calibration matrix also decreased the standard deviations of the ETS measurement errors, across the board.

Lessons Learned. The ability to fly the ETS algorithms in the field early in the mission development is a unique opportunity for the Dragonfly team. Sand collection on the camera window, shadowing conditions, auto-exposure performance, jitter, camera fixed pattern noise, and camera pointing knowledge error are all examples of

real impacts to image correlation performance that can be modeled in simulation, but present themselves in uncertain ways in a physical system. Data gathered from these flight tests have been critical in identifying areas of improvements in our CLS models, as well as ensuring the algorithms designed for Dragonfly are robust to the uncertainty in Titan's environment.

At its core, the ETS function creates optical measurements by attempting to match the pattern of the signal from the terrain across two images. If there is a competing pattern, then there will be secondary peaks that cause the peak ratio threshold to be exceeded and therefore an invalid measurement. Both self-shadowing and fixed pattern noise can be a source of this competing pattern, especially over feature-poor imagery with either washed-out or dim lighting conditions. During flights near midday, the shadow of the drone was observed in ETS imagery. When flying in a straight line, from one image to the next the signal from the underlying terrain shifted in the image but the shadow remained in the same location. After the warping step of the ETS function, while the terrain signals generally matched, there was a noticeable shift in the shadow location, as shown in figure 12. This offset led to invalid solutions, leading to the decision not to perform flight tests near midday. Analyses in simulation show that due to Titan's dim, diffuse lighting conditions, the issue of self-shadowing will not occur during the Dragonfly baseline mission, but this sensitivity to fixed optical artifacts is useful knowledge and current design work includes more robust outlier detection.

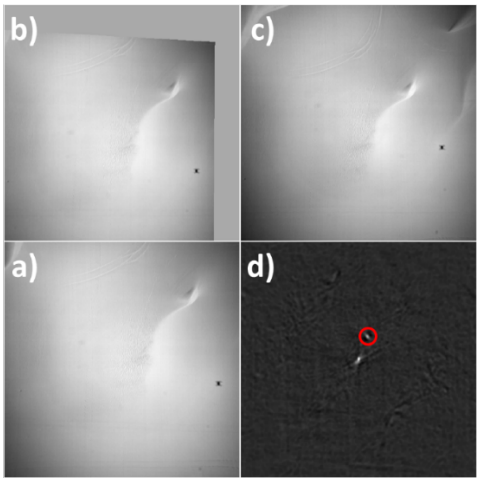

Figure 12. Demonstration of the effect of self-shadowing on the ETS function. When the current image (a) is warped (b) into the same frame as the reference image (c), the features of the terrain match up but the shadow (right side of imagery) is in a different location in (b) than in (c). This leads to a second peak (circled in red) in the correlation surface.

Similar to self-shadowing, fixed pattern noise remains stationary in the image plane as the signal from the terrain underneath shifts, leading to multiple peaks in the correlation surface. Effects of fixed pattern noise were negligible on the correlation scores until the final flight of the flight test campaign in a dim lighting at dusk. With bland terrain and feature-poor imagery, the signal from the fixed pattern noise had a noticeable impact on the correlation scores, leading several measurements to exceed the peak ration threshold. While both dark current and fixed pattern noises should be removed via a subtractive and multiplicative calibration, respectively, in the field the effect of the fixed pattern noise was illustrated by summing the frames of the flight to determine the average value of each pixel, and playing the flight back, subtracting the pattern from each image as it became the current image. As is evident in figure 13, the simplistic removal of fixed pattern noise got rid of the secondary peak artifacts in the correlation surface, resulting in valid measurements that had previously been marked as invalid. Due to the observed sensitivity to fixed pattern noise, simulations for ETS will prioritize the incorporation of dark current and fixed pattern noises into the camera truth model, and the baseline ETS function will include dark current and fixed pattern noise removal.

Figure 13. Comparison of the peak ratios for the final, dim-lighting test flight of the campaign. The fixed pattern leads to measurements that exceed the peak

ratio test (top), but when the fixed pattern is removed those measurements are accepted (bottom). The iterations marked "invalid_solution" are due to either altitude or attitude thresholds exceeded, due to the higher rates the ITP experiences.

Finally, as illustrated in the results section, ETS shows sensitivity to miscalibration: both in the intrinsic camera matrix as well as in the body-to-camera transform knowledge. Future flight tests will continue to periodically calibrate the camera and use the updated calibration parameters, and a more rigorous calibration process will be performed on the NavCam flight models when they are delivered to ensure accurate knowledge. For camera pointing knowledge, however, the problem is more difficult to solve for the ITP. The Dragonfly baseline lander will undergo pointing alignment calibration prior to launch and will have a tight error budget, but for ITP the "plug-and-play", fast-paced nature of the drone box leads to alignments that are difficult to reproduce on the sub-degree accuracy needed for low measurement errors. The team is working to determine effective pre-flight calibration procedures, and it remains an active area of research.

Conclusion. The Dragonfly mobility team has performed a series of flight tests over relevant, desert dune terrain in order to test the mobility guidance, navigation and control application, including the optical navigation ETS function. The Dragonfly ITP setup successfully flew autonomously during flights of up to 2.4 km long and returned to land accurately. The test was beneficial to demonstrate the system as a whole, and within ETS to characterize both performances and sensitivities. ETS demonstrated the ability to create both velocimetry and localization measurements, allowing the rotorcraft to retrace its steps back to the landing site. ETS also showed sensitivities to calibration and optical artifacts, which will inform ongoing design efforts to make the algorithm robust and also will inform future testing procedures. Through the flight test campaign, a wealth of data over a wide range of conditions will be a valuable dataset for ongoing and future analyses, which the team can use to test new algorithm components through playback and image pair analysis.

Acknowledgements. The authors would like to thank the Dragonfly Mobility team for their persistence and consistently excellent work to develop the ITP platform, MGNC application, and support the long hours of the flight test campaign. We would also like to thank everyone who is working to make this incredible mission a reality as part of the Dragonfly team. Finally, we would like to thank JHUAPL's Jim Harris, who endured the long days of testing with us and captured some amazing video and imagery using the chase drone.

Figure 14. Some of the Dragonfly flight test team at the Imperial Sand Dunes, May 2022.

References.

[1] R. D. Lorenz *et al.*, "Dragonfly: A Rotorcraft Lander Concept for Scientific Exploration at Titan," *Johns Hopkins APL Technical Digest (Applied Physics Laboratory)*, vol. 34, no. 3, pp. 374–387, 2018.

[2] I. R. Witte *et al.*, "No GPS? No Problem! Exploring the Dunes of Titan with Dragonfly Using Visual Odometry," Jan. 2019. doi: 10.2514/6.2019-1177.

[3] R. D. Lorenz *et al.*, "The sand seas of Titan: Cassini RADAR observations of longitudinal dunes.," *Science*, vol. 312, no. 5774, pp. 724–7, May 2006, doi: 10.1126/science.1123257.

[4] C. A. Sawyer and N. L. Mehta, "Rendering the Titan Environment for Dragonfly," Aug. 2019.

[5] G. F. Lindal, G. E. Wood, H. B. Hotz, D. N. Sweetnam, V. R. Eshleman, and G. L. Tyler, "The atmosphere of Titan: An analysis of the Voyager 1 radio occultation measurements," *Icarus*, vol. 53, no. 2, pp. 348–363, Feb. 1983, doi: 10.1016/0019-1035(83)90155-0.

[6] D. F. Strobel and D. E. Shemansky, "EUV emission from Titan's upper atmosphere: Voyager 1 encounter," *J Geophys Res Space Phys*, vol. 87, no. A3, pp. 1361–1368, Mar. 1982, doi: 10.1029/JA087iA03p01361.

[7] J.-P. Lebreton *et al.*, "An overview of the descent and landing of the Huygens probe on Titan," *Nature*, vol. 438, no. 7069, pp. 758–764, Dec. 2005, doi: 10.1038/nature04347.

[8] B. Schilling, B. Villac, and D. Adams, "Preliminary Surface Navigation Analysis for the Dragonfly Mission," *Astrodynamics Specialist Conference, AAS*, pp. 19–772, Aug. 2019.

[9] D. C. Brown, "Close-Range Camera Calibration," *Photogrammetric Engineering*, vol. 37, no. 8, pp. 855–866, 1971.