

PARACHUTE MODEL VALIDATION USING IMAGE SEQUENCES

Shortis, M. R.¹, Robson, S.², Jones, T. W.³ and Lunsford, C. B.³

¹RMIT University, Australia

²University College London, U.K.

³NASA Langley Research Center, U.S.A.

mark.shortis@rmit.edu.au

ABSTRACT

Parachute systems play a critical role in many science and military missions. Currently, NASA and the U.S. Army air delivery systems programs are seeking measurement technologies to support experimental and qualification testing of new and modified parachute concepts. Experimental testing to measure the physical parameters involved during the inflation and decent phases of parachutes is a key requirement to obtain a quantitative assessment of the parachute performance. The paper will provide further detail on the rationale, the system design, the experiments and the data analysis.

1. INTRODUCTION

Scientists are developing new parachute concepts for replacement of the aging T10 canopy which dates back to the Korean War (see figure 1). The evolution of the modern soldier has seen the overall combined payload weight reach the limits of the older canopies. To validate air flow models for their new concepts, the Army has begun a series of both ground based and air-drop tests of scalable parachute models for database development and correlation. The Army has sought external expertise in various measurement technologies to support this effort.



Figure 1a. Standard US Army T10 parachute canopy.



Figure 1b. An example of new parachute design.



Figure 1c. Apollo 15 landing, similar in concept to the new NASA Crew Launch Vehicle

The technology development activity is seen to have generic benefit where the entry, descent and landing systems incorporate a parachute. Ground validation testing to measure the physical parameters involved during the inflation and decent phases of parachutes is a key requirement.

Historically the shape of the canopy during inflation and descent has been determined qualitatively, however modern video imaging technology has provided a tool capable of supporting quantitative measurement. LaRC has a long history of optical systems and photogrammetric measurement of aerospace models, dating back to the 1970s (see Shortis and Snow, 1997). Most recently this technology has been applied to the In-Space Propulsion programs 20m Solar Sail ground validation test program (Shortis *et al.*, 2002). The two programs share many similarities in terms of

metrology; their large scale and “membrane like” structure exhibit very non-linear aspects during deployment and inflation. A goal of this effort was the development of a scalable imaging system that supports both ground based platform drops and wind tunnel testing, and high altitude airdrops.

The objectives of these measurements are to monitor the aerodynamics and structural dynamics of parachutes and to provide a database for the verification and validation of the mathematical models developed to predict the performance of these parachute concepts. Scale model ground testing of parachute systems has been conducted in both wind tunnels and low altitude airdrops, however the efforts here are aimed at correlating shape information with the other parameters including acceleration, velocity, load and canopy configurations.

2. EXPERIMENT DESIGN AND INSTRUMENTATION

A series of tests with a drogue parachute (6 feet diameter) and a generic circular parachute scale model (7.5 feet diameter) provided by the Army have been conducted in the High Bay facility at LaRC and in the Plum Brook Space Power Facility (SPF) at NASA Glenn Research Center. The tests at LaRC were conducted as a trial to validate the approach used for image sequence capture, camera calibrations and the sequence processing. SPF has the capability to provide a controlled environment of pressure and temperature that may be used in the future to replicate high altitude conditions. The tests at SPF were conducted in ambient conditions, however the chamber was maintained at constant temperature and has the added benefit of minimisation of air currents to eliminate any impact on the parachute during test drops. SPF has the additional advantage of a longer drop time from the 40m ceiling, as compared to 25m at LaRC.




Figure 2a. Experimental set up at Plum Brook Space Power Facility.




Figure 2b. Parachute model in the pre-drop position.

The drogue parachute comprised a conventional nylon fabric canopy made up of 12 pseudo-triangular sections known as gores, connected to the payload by support lines. The scale model parachute is manufactured from continuous fabric, with 16 support lines. In order to increase visibility, the cord support lines were replaced by thin steel wires. The payload acted both as ballast and an instrumentation package. The parachute was dropped along a guide wire attached to the ceiling and tensioned using 100kg of weights suspended just above the floor. The guide wire passed unimpeded through the payload and the vent of the parachute. Figure 2 shows a schematic of the set up at SPF and the parachute model in the pre-drop position.

The underside and leading edges, or skirts, of the canopies were covered with retro-reflective circular targets which were illuminated by ring lights or near-axis lights on the cameras (see figure 3). Early test configurations adopted at LaRC for the drogue parachute delineated the gores and were spatially dense to ensure accurate determination of the canopy shape. However it became apparent that target tracking would be problematic, especially in the critical period during inflation of the parachute. Later configurations on the scale model parachute used less targets on the canopy

or targets only on the skirt. Retro-reflective targets were also positioned near the connection point of the guide wire on the ceiling in each facility. The ceiling targets, including scale bars at SPF (see figure 4), were coordinated by a self-calibration network using a high resolution digital still camera and subsequently used as fixed references for the resection of the parachute images.




Figure 3a. Spatially dense target pattern. Early test configuration.

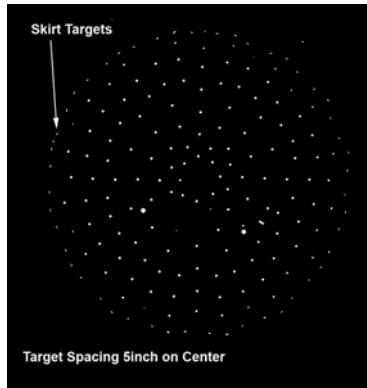


Figure 3b. Typical target pattern used for most drop tests.




Figure 3c. Canopy with skirt targets only.

The image capture system comprises six fixed Dalsa cameras on the floor surrounding the drop point and two “circuit board” CCD cameras fixed to the parachute payload. The Dalsa cameras have 1600 by 1200 resolution progressive scan sensors with 7.4 micrometre pixel spacing. The cameras have dual tap outputs which allow 30 frames per second read out rates via CameraLink interfaces. The cameras are synchronised based on a hardware trigger and the image sequences are captured using a multi-camera recording system from IO Industries P/L. The system provides a timestamp for synchronization checks in the corner of the recorded AVI image sequences. The synchronisation was validated at LaRC using a single retro-reflective target on a laser beam “chopper” rotating at 200rpm.

Based on simulations and the availability of 35mm and 50mm fixed lenses, four cameras were fitted with 35mm lenses and two cameras with 50mm lenses. One camera with a 50mm lens was placed very close to the drop point to ensure continuous visibility of all canopy targets. The four cameras with 35mm lenses were placed in a circular pattern approximately 3m and 5m from the drop point for LaRC and SPF respectively (see figure 4). These locations were adopted as a compromise between the geometric strength of the photogrammetric network and optimisation of the field of view in terms of the extent of the parachute drop and occlusions caused by the curvature of the canopy near the leading edge. The sixth camera was used at SPF but was unavailable for LaRC.



Figure 4a. Resection targets on the crane at the SPF ceiling.




Figure 4b. Overview of fixed cameras at the LaRC drop point.

A purpose-built, self-contained canister has been developed as the payload to provide ballast and capture digital images from the two circuit board cameras arranged in an upward-looking stereo

configuration. The cameras record monochrome 720 by 480 resolution, interlaced images at 30 frames per second. The images are captured via RS170 video interface to miniature digital video recorders (DVRs). The DVRs utilise solid state flash memory cards and a fixed 4 to 1 JPEG compression to reduce the storage requirements. Approximately 4 minutes of video can be stored for each camera on 1Gb flash cards. The image sequence is triggered by a radio frequency controller which transmits simultaneously to both DVRs, ensuring a synchronised start to the recording. The stereo pairs are subsequently synchronised based on an on-screen timer.

For the payload lighting system a 24 volt high-intensity multi-LED array was designed (see figure 5). The white light LEDs incorporate a non-standard lens to project a 52° divergence cone to encompass the canopy in its fully deployed shape. The LEDs are independently controlled by radio signal. In addition to economizing the power drain on the battery, this feature allows the lighting on/off operation to serve as a synchronization check for the recorded video from the on-board cameras and link to the image sequences from the ground based cameras. The canister also incorporates a number of other sensors, such as air velocity from a hot wire anemometer, a tri-axial accelerometer and four load cells on the suspension lines between the canopy and the payload.




Figure 5. Payload canister with ballast weights, circuit board cameras and LED lights.

For each drop test, the canopy and payload are hoisted up the guide wire from the pre-drop position (see figure 2). The fixed Dalsa cameras are triggered by the capture system and commence recording direct-to-disk. The payload cameras are then triggered by the radio control and commence recording to the DVRs. The payload LED lights are then switched on, also by radio control, with the secondary purpose of identifying a synchronisation point for the two independent systems. The canopy is then released. Other synchronisation points are provided by the release, the canopy reaching full inflation and, if needed, a manually fired external flash. Once the payload reaches the cushion at the end of the guide wire, all systems are commanded to stop recording and the LEDs are turned off to conserve power.

3. CALIBRATION

The two independent sub-systems of cameras are calibrated separately. Self-calibration and relative orientation of the fixed cameras is determined by lowering a calibration fixture down the guide wire used for the parachute drops (see figure 6). In an ideal case, the calibration fixture should travel the full extent of the drop, however in practice the number of locations can be limited by the logistics in the facility. Locations at the start, mid-point and bottom of the drop, with rotations of the fixture, are a practical solution that provides an optimum, 3D self-calibration network. The relative orientation, subsequently derived from the camera stations in the network, represents the full range of motion and thereby minimises any systematic errors that may ensue from a more limited sample of the drop extent.




Figure 6. Calibration fixture being lowered along the guide wire at LaRC.




Figure 7. Full frame image from a payload camera showing interline motion displacement.

This procedure was performed at LaRC, but could not be completed in full at SPF. To strengthen the self-calibration at SPF, a number of additional static images of the calibration fixture from all individual Dalsa cameras were incorporated into the network. This strategy improved only the camera calibration aspect of the self-calibration.

A self-calibration and relative orientation of the payload cameras was determined by using the stereo-cameras to image a static calibration fixture using a multi-station, convergent network. In this case the calibration fixture was positioned at the same range as the canopy and then moved around within the field of view to optimise the camera calibration and relative orientation.

A significant disadvantage of the interlaced images provided by the payload cameras is the motion displacement of target images, shown clearly in figure 7. Accordingly the payload cameras were calibrated in frame mode but were used for measurement and tracking in field mode. The relationship between calibrations in frame and field modes is described in Shortis and Snow (1995).

4. SEQUENCE PROCESSING

Target tracking commences with a simple setup procedure which is largely operator controlled. For convenience each set of acquired images is treated as a distinct epoch of data (see figure 8). To commence, either a previously established relative orientation of the cameras is used, or a minimum of four control targets on the ceiling are identified on all images in the first epoch. For the resection, a direct solution (Zeng and Wang, 1992) is then used to compute initial exterior orientation parameters for each camera. These initial values are then automatically refined by means of a robust least squares estimation resection. Following a successful resection, the computed exterior orientations of each camera are used to identify the imaged locations of any additional control targets. A further resection is automatically computed to ensure that all available information has been used to compute each camera location and orientation.

The next step is the identification and computation of the projected targets on the object surface. Two images from the first epoch are selected and all corresponding targets manually identified and measured to provide a spatially meaningful numbering sequence. Following an intersection to determine the three dimensional locations of these targets, their imaged locations in any remaining images in the set can be simply computed by the same back-driving procedure used to find the control targets. The back-driving procedure is robustified to account for occlusions by searching for each computed target according to its predicted image dimension and depth away from the camera. Any targets omitted on completion of this procedure, such as those inseparably close together or partially occluded by parachute support lines, are identified manually.

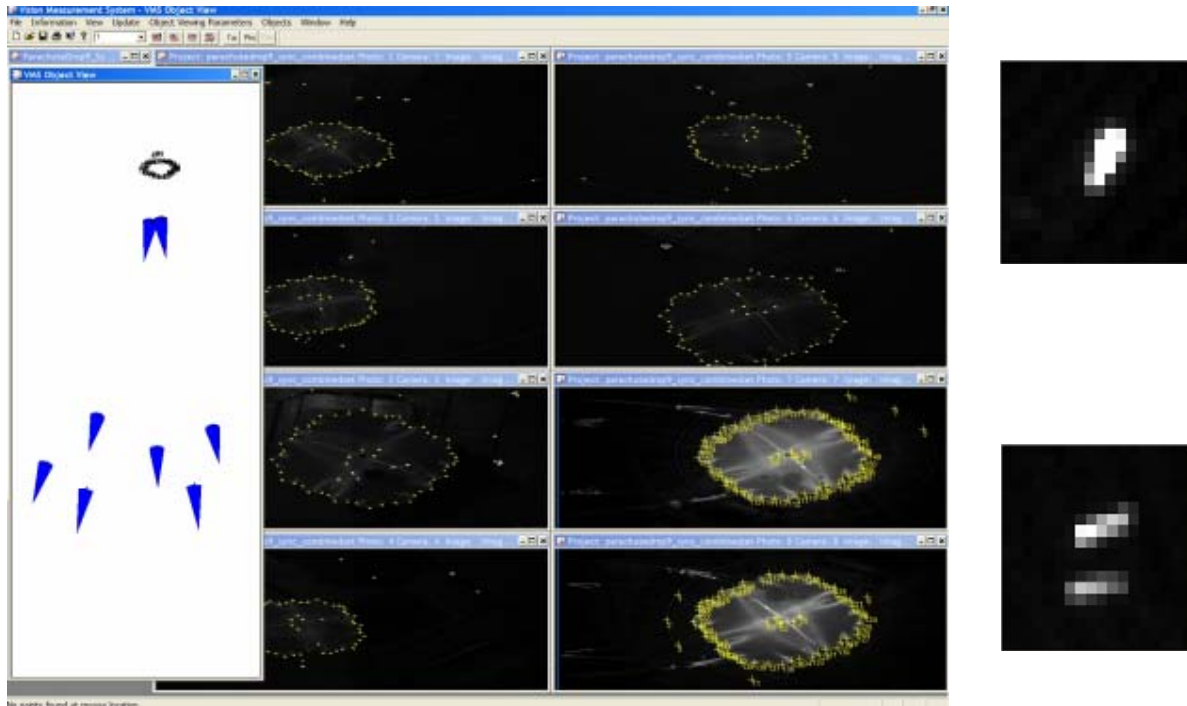


Figure 8. Object space schematic plus camera views for one epoch in an image sequence (note that this set has been acquired at a test electronic flash synchronisation point, giving a much more brightly illuminated canopy than is typical) and typical target images showing close proximity between targets during the collapsing phase of parachute deployment (bottom right).

The tracking method is based on knowledge of the prior locations of each target in up to three preceding epochs. On the basis of this information a set of motion equations (Shortis and Snow, 1997) can be used to predict the location of the target at the next epoch. Since initial values for the camera exterior orientation are known, the approximate imaged location of the target can be found.

For each new epoch the tracking process progresses according to five steps:

- a) The location of known control points by back-driving to their computed locations.
- b) Re-computation of the camera location by least squares resection.
- c) Back-driving to the predicted locations of the tracked targets for each image in turn.
- d) Computation, by intersection and then network solution, of the locations of all targets.
- e) Estimates of the next location of each tracked target are then computed.

A photogrammetric network adjustment for each epoch, incorporating all measurement data, is computed as part of step (d) once the resection/intersection procedure has been completed for each epoch. This was desirable in the case of these tests since control information was visible in all images and the ability to simultaneously determine optimal fit between longer range control points and closer target points on the parachute canopy takes all available information into account. This approach is in contrast to the use of the pre-determined relative orientation, which is based on a limited sample of control targets and calibration fixture locations on the guide wire. In effect, the multi-camera relative orientation risks extrapolation into an uncontrolled object volume.

The tracking process is repeated for each successive image set. In this way a set of object surface points can be rapidly acquired for each epoch. The process is also flexible in terms of the number of images used in each set and indeed the number of image sets. Each extra image set is simply treated as an additional epoch of data to be sequentially processed. In this particular example it is convenient to maintain the identification of each tracked target in all epochs, since this allows continuity of canopy motion to be assessed and modelled.

5. ANALYSIS OF RESULTS

Results from a network adjustment incorporating ground cameras, compared to both ground and payload cameras, are shown in table 1. It is evident that the addition of the two payload cameras into the network confers a significant improvement in the attainable precision. This is unsurprising given the imaging geometry seen in figure 8, where the upper two cones represent the payload cameras and demonstrate their much closer proximity to the parachute. The relatively large coordinate precisions from an externally constrained network which utilises control targets to define the datum result from the fact that the control targets are located on the ceiling several metres above the parachute canopy. In contrast, optimal precisions from an inner constraints network indicate precisions of targets on the parachute of the order of 1 to 6mm. In the case of the free network solution the payload cameras degrade the plan precisions of the target coordinates because of a reduced internal consistency, largely caused by the poor image quality of the payload camera images. Within an individual epoch, the level of precision attainable utilising inner constraints enables a much finer resolution for shape analysis.

External constraints network	6 ground and 2 payload cameras			6 ground cameras		
RMS image residual (micrometres)	4.42			3.93		
Number of redundancies in the network	816			588		
Relative precision for the network	1 : 318			1 : 215		
	X	Y	Z	X	Y	Z
Mean precision of canopy targets (mm)	155	123	79	216	177	134
RMS corrections on control (mm)	28	31	123	24	28	71
Inner constraints network						
Mean precision of canopy targets (mm)	3.0	6.3	3.2	0.9	1.5	5.1

Table 1. Comparisons of results for the various network solutions.

The particular parachutes analysed in this series of tests demonstrated a pulsation or “jelly fish” like motion as they descended, creating additional difficulties for target tracking. In particular, targets at the periphery of the parachute edge moved closer together and then further apart as the canopy descended (see figure 9). A solution to the case prior to the inevitable occurrence of occlusion of one target by another (see figure 8) is simply to reduce the dimensions of the search window used to find the targets in each image. In the case of occlusion, both targets were dropped from the tracking solution and then interactively re-acquired in a following epoch. This approach necessitated a tracking solution that required operator intervention in order to preserve consistent target numbering. Analysis of the observed motion typical of the parachutes under investigation is being undertaken in order to provide a better predictive model. An alternative, but currently less practical solution, is to increase the rate at which image sequences are acquired so that the relative image motion between data epochs is minimised.

As noted in the previous section, a network solution was used in preference to the pre-determined relative orientation at each epoch. This option was validated by the testing with the relative orientation solution for the SPF sequences, which exhibited RMS image errors of the order of 10 to 15 micrometres, significantly larger than the 3 to 5 micrometer errors produced by the bundle adjustments. Whilst the use of the multi-camera relative orientations would simplify the sequence processing, the use of an externally constrained network solution is necessary in this case to provide a link between the ground and payload cameras, and provide a consistent datum for all epochs.

A number of visualisations of the tracked targets have been generated to understand the changes in shape of the canopy, and in particular the variations of the leading edge or skirt of the parachute. The results presented here are extracted from the tracking carried out at SPF using the scale model parachute. Figure 9 shows the skirt with respect to a fixed coordinate datum as the parachute falls along the guide wire. The pulsations of the parachute soon after full inflation, easily identified in the image sequences, are evident. From this visualisation the magnitude of the pulsations can be estimated. Figure 10 shows the skirt with respect to the payload cameras. In this visualisation the

lighter colours indicate progression in time. The leading edge of the canopy demonstrates a sinusoidal twisting motion relative to the payload as the parachute expands and collapses during the pulsations, and this effect was not readily evident in the image sequences. Further analysis of the shape changes and spatial behaviour of the canopy, along with the information from the other sensors on the payload, is required to fully understand the aerodynamics of the parachute.




Figure 9. Side view of the skirt and canopy centre of a parachute with respect to a fixed datum.



Figure 10. Top view of the skirt and canopy centre of a parachute with respect to the payload cameras.

6. CONCLUSIONS

The approach to the characterisation of the parachute models described in this paper is successful in extracting sequences of target coordinates, but needs improvement in either the frame rate or the tracking solution utilised. Development work is ongoing to determine a more automated solution, based on investigations of a variety of surface model and photogrammetric geometry solutions.

REFERENCES

- Shortis M. R., Robson, S., Pappa, R. S., Jones, T. W. and Goad, W. K., 2002. Characterisation and tracking of membrane surfaces at NASA Langley Research Center. *International Archives of Photogrammetry and Remote Sensing*, 34(5): 90-94. ISSN 1682-1777.
- Shortis, M. R. and Snow, W. L., 1995. Calibration of CCD cameras for field and frame capture modes. Proceedings, *Conference on Digital Photogrammetry and Remote Sensing '95*, SPIE Vol. 2646, pp. 2-14.
- Shortis, M. R. and Snow, W. L., 1997. Videometric tracking of wind tunnel aerospace models at NASA Langley Research Centre. *The Photogrammetric Record*, 15(85): 673-689.
- Zeng, Z. and Wang, X., 1992. A general solution of a closed-form space resection. *Photogrammetric Engineering and Remote Sensing*, 58(3): 327-338.