The 2009 IEEE/RSJ International Conference on
Intelligent Robots and Systems
October 11-15, 2009 St. Louis, USA

A Stereo Vision System for UAV Guidance

Richard J. D. Moore, Saul Thurrowgood, Daniel Bland, Dean Soccol & Mandyam V. Srinivasan

***Abstract*—This study describes a novel, vision-based system for guidance of UAVs. The system uses two cameras, each associated with a specially-shaped reflective surface, to obtain stereo information on the height above ground and the distances to potential obstacles. The camera-mirror system has the advantage that it remaps the world onto a cylindrical co-ordinate system that simplifies and speeds up range computations, and defines a collision-free cylinder through which the aircraft can pass without encountering obstacles. The result is a computationally efficient approach to vision-based aircraft guidance that is particularly suited to terrain and gorge following, obstacle avoidance, and landing. The feasibility of the system is demonstrated in laboratory and field tests.**

I. INTRODUCTION

There is considerable interest in the design of guidance systems for UAVs that use passive sensing (such as vision), rather than active sensing which can be bulky, expensive and stealth-compromising. A recent trend in biologically inspired vision systems for aircraft guidance has been to exploit optic flow information for collision avoidance, terrain following, gorge following and landing (e.g. [1], [2], [3], [4], [5]). However, the acquisition of optic flow requires that the aircraft be in motion, which makes such systems unviable during periods of slow flight or hover.

Stereo vision, on the other hand, enables estimation of range regardless of the aircraft's speed, and functions even when the aircraft is stationary. Moreover, systems that rely on optic flow for extracting range information need to discount components of optic flow that are induced by rotations of the aircraft, and use only those components of optic flow that are generated by the translational component of motion. The reason is that it is only the translational components of optic flow that provide information on the range to objects in the environment. Vision systems that exploit stereo information do not require this computationally elaborate procedure.

Wide-angle stereo systems have previously been designed for aircraft (e.g. [6], [7], [8]), but they have rarely been tailored to the specific needs of aircraft guidance, such as terrain and gorge following, obstacle detection, and landing. Here we describe a stereo system that is specifically designed to serve these requirements.

A. System concept

The concept of the system is best described by considering an assembly in which a camera views a specially shaped reflective surface (mirror), as illustrated in Fig. 1. The profile of the mirror is designed to ensure that equally spaced points on the ground, on a line parallel to the camera's optical axis, are imaged to points that are equally spaced in the camera's image plane. The derivation of the profile of this mirror is given in [9], and is not repeated here. The aim of that study, however, was to simplify the computation of optic flow, while the goal of the present study is to simplify the computation of range from stereo.

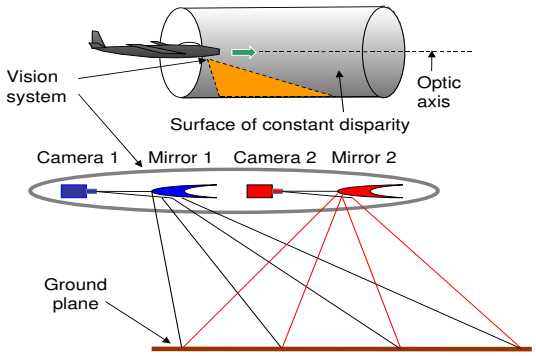

Fig. 1. Schematic illustration of the conceptual stereo visual system, surface of constant disparity and collision-free cylinder.

Consider a system in which two such assemblies are arranged coaxially, as shown in Fig. 1. Each camera views the environment through a mirror, as described above. It follows that the disparity in the positions of the images in the two cameras of any point on the ground will be inversely proportional to the radial distance of that point from the common optical axis of the two cameras. Therefore, surfaces of constant disparity will be cylinders, as shown in Fig. 1. Furthermore, the disparity in image positions will be one-dimensional only, simplifying the calculation of disparity and reducing it to a computation of one-dimensional optic flow. The maximum disparity, as measured by this system, defines the radius of a collision-free cylinder through which the aircraft can fly without encountering any obstacles. The system is therefore well suited to providing information for visual guidance in the context of tasks such as terrain and gorge following, obstacle detection, and landing.

II. STEREO VISION SYSTEM

A. Design

In the present implementation of the system, two camera assemblies are rigidly mounted in a coaxial stereo configuration, as shown in Fig. 2a, to minimise vibration-induced measurement errors. We have not used physical mirrors, but instead used high resolution video cameras (PGR Grasshopper

This work was supported partly by US Army Research Office MURI ARMY-W911NF041076, Technical Monitor Dr Tom Doligalksi, US ONR Award N00014-04-1-0334, ARC Centre of Excellence Grant CE0561903, and a Queensland Smart State Premier's Fellowship.

Authors are associated with the Queensland Brain Institute & School of Information Technology and Electrical Engineering, University of Queensland, St Lucia, QLD 4072, Australia and ARC Centre of Excellence in Vision Science, Australia. Correspondence via r.moore1@uq.edu.au

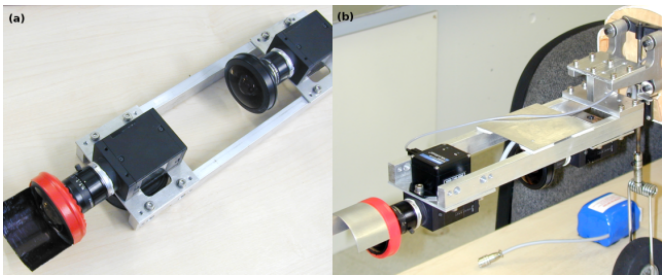

Fig. 2. (a) Implementation of the stereo vision system and (b) mounted on the aircraft.

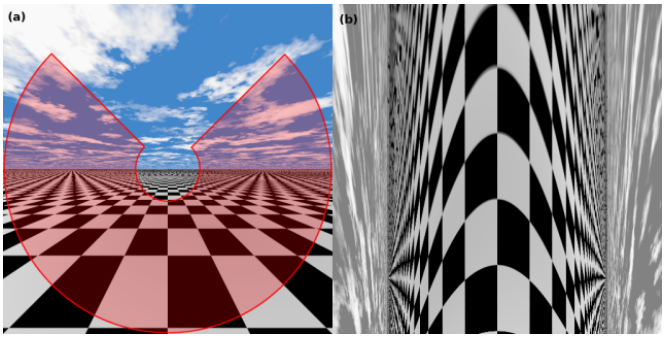

Fig. 3. (a) Raw image of a rendered scene and (b) the remapped image. The shaded area in (a) has been remapped to (b).

20S4M) equipped with wide-angle fish-eye lenses (Fujinon FE185C057HA-1) and simulated the imaging properties of the mirrors by means of software lookup tables. This method reduces the physical bulk of the system and enables us to implement and evaluate very quickly and cheaply a variety of imaging configurations involving lenses and/or mirrors. By making use of high resolution cameras and wide-angle lenses, resolution in the far-field and the field of view of the vision system are not compromised, whilst aberrations due to imperfections in the mirror surface are avoided.

The mapping produced by the present configuration is illustrated in Fig. 3. In this example a camera is positioned above a horizontal plane parallel to the camera's viewing axis. An image of the rendered scene is captured through a rectilinear lens with a 120° FoV. The shaded area of the raw image is unwrapped and transformed to produce the remapped image. Note that in Fig. 3b the perspective distortion, or foreshortening of the image along the optic axis, has been removed.

Range information is extracted from the remapped images by computing the image disparity between the stereo pairs. The algorithm used to compute the disparity is based on the sum of absolute differences (SAD) between images and is implemented using the Intel Integrated Performance Primitives library (see [10]). To remove low frequency image intensity gradients, which can confuse the SAD algorithm, the remapped images are convolved with a high-pass filter kernel (5px × 5px window) before the disparity is computed. Each correlation score is generated by computing the SAD between a window in the rear image with the corresponding window of the same size in the front image. The correlation scores for many windows can be quickly calculated by pre-computing a running integral of the absolute differences between the pixels in the two images. This process is then repeated as one image is increasingly offset with respect to the other. For each window, the disparity is then simply the image offset at which the computed SAD is a minimum.

TABLE I
SYSTEM PARAMETERS

	Calibration	Flight test
Stereo baseline ($d_{baseline}$)	200mm	
Remap image rows (h_{img}) / cols	384px / 128px	
Vertical FoV	20.6° to 70.8° from vertical	
Horizontal FoV	−100° to 100° from vertical	
Forward viewing factor (r)	2.5	
Detectable disparity (D_{pixel})	10px to 40px	0px to 15px
Operational altitude (d_{radial})	0.8m to 3.0m	2.0m ∼ 30m

To obtain sub-pixel disparity estimates, an equiangular fit (as described in [11]) is applied to the minimum and neighbouring SAD scores for each window. Disparities are computed to an accuracy of $\frac{1}{8}$px and represented as an unsigned byte, giving a maximum search range of 31.5px. Incorrect matches are rejected by re-computing the disparity for the reverse image order and discarding disparities which differ by more than two pixels from the original estimate. This bi-directional technique is effective at rejecting mismatches near stereo discontinuities but doubles the execution time of the algorithm. When tested on a 1.5GHz processor, the SAD algorithm (with bi-directional and sub-pixel search) runs at ∼ 48Hz, generating 384px×128px disparity images for a window size of 11px × 11px and a search range of 0px → 15px.

The pixel disparity, D_{pixel}, produced by a point at a radial distance, d_{radial}, from the optic axis of the system is given by

$$D_{pixel} = \frac{d_{baseline} \times h_{img}}{r} \times \frac{1}{d_{radial}}, \qquad (1)$$

where $d_{baseline}$ is the stereo baseline, h_{img} is the vertical resolution of the remapped images (i.e. the number of pixel rows) and r, the forward viewing factor, is the ratio of the total forward viewing distance to the height of the aircraft. The first term in (1) is simply a constant which depends on the system configuration. The system parameters and their values used in this study are listed in Table I. Increasing the magnitude of the pixel disparities, or increasing the stereo baseline, increases the flight ceiling of the system but also increases the disparity search range required to detect nearby objects. The current system design is a compromise between the desired flight envelope and the computational cost of the required disparity search range.

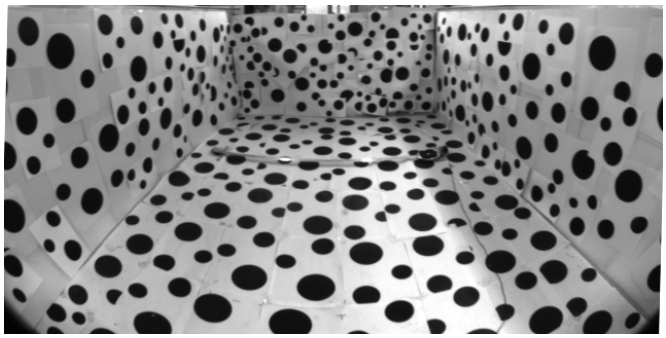

Fig. 4. Raw image of the indoor testing arena as seen by the front camera.

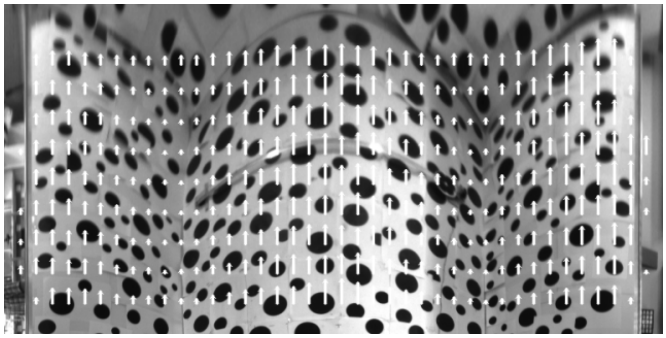

Fig. 5. Remapped image of the arena overlayed with the computed image disparities. The disparity vectors have been scaled to aid visualisation.

B. Calibration

To account for any idiosyncrasies, each camera assembly has been calibrated using the generic camera model described in [12]. The two assemblies have also been calibrated as a stereo pair so that any rotational misalignment can be compensated for during the remapping phase.

The performance and accuracy of the stereo system were evaluated in an artificially textured arena. A cropped image of the arena as viewed by the front camera is displayed in Fig. 4. The texture used to line the walls and floor of the arena is composed of black circles of varying diameter (65mm $\rightarrow$ 150mm) on a white background.

The image from the front camera has been remapped and is displayed in Fig. 5, overlayed with the image disparities computed with the corresponding remapped image from the rear camera. It can be seen that the disparity vectors in any given column have a constant magnitude. This verifies the expected result – that the image disparity depends only on the radial distance of the viewed point from the optic axis, or in other words, that surfaces of constant disparity are cylinders. Points on the rear wall of the arena display larger disparities since they are closer to the optic axis.

The radial distance from the optic axis to each visible point can be quickly calculated from the computed disparity via (1). In order to quantify the accuracy of the stereo system, the relationship between estimated radial distance to the arena and viewing angle was plotted against the actual relationship (Fig. 6). The actual relationship was calculated from the known

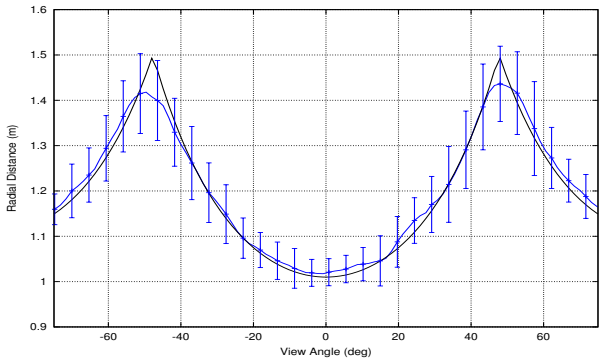

Fig. 6. Profile of the estimated radial distances to the arena wall/floor (blue line) overlayed with the actual radial distances at each viewing angle (black line). Error bars correspond to $\pm 2\sigma$ at each viewing angle.

geometry of the arena. Points corresponding to the rear wall of the arena have been omitted to simplify analysis. The error in the estimated radial distance at each viewing angle represents the variance in the estimates along the optic axis (ie. along a column of vectors in Fig. 5). It can be seen that the errors in the estimated radial distance are most significant for viewing directions corresponding to the corners of the arena, where the walls join the floor. This is a result of the non-zero size of the window used to compute disparity. A window size greater than one pixel would be expected to cause an underestimation of the radial distance to the corners of the arena, where surrounding pixels correspond to closer surfaces, and also a slight overestimation of the radial distance to the arena floor immediately beneath the cameras, where surrounding pixels correspond to surfaces which are further away, and indeed this is observed in Fig. 6.

If the direction of the ray originating at the nodal point of the rear camera and passing through the observed point is known, then the direct ray distance to the observed point (and hence its location in 3D space relative to the rear camera) can be calculated from the radial distance via simple trigonometry. The view rays for each pixel in the remapped rear image can be obtained from the calibration for the rear camera. Therefore, the three dimensional structure of the visible environment can be reconstructed from the computed disparities. As a further test of the system's performance, a three dimensional reconstruction of the arena was generated using this procedure. The reconstruction is shown in Fig. 7.

It can be seen from Figs. 6 & 7 that the stereo system is able to reconstruct a 3D environment (e.g. a valley) faithfully. The data is calculated from a single typical stereo pair and is unfiltered, however a small number of points were rejected during the disparity calculation step. Small errors in the reprojected viewing angles may arise from inaccurate calibration of the camera assemblies but are presumed to be negligible in this analysis. Therefore, the total error in the reconstruction can be specified as the error in the radial distance to the arena at each viewing angle. The standard deviation of this error, measured from approximately 2.5×10^4 reprojected points,

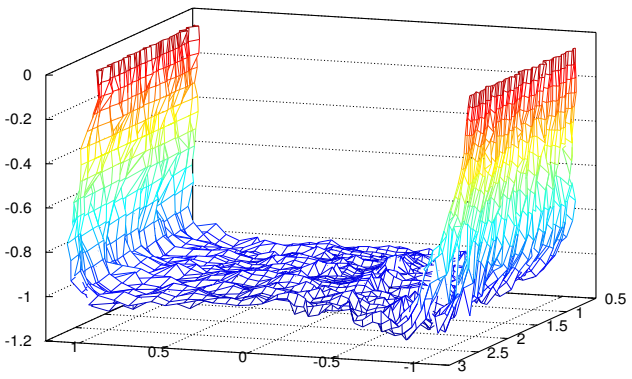

Fig. 7. 3D reconstruction of the arena used to calibrate the system. Measurements are in metres relative to the nodal point of the rear camera.

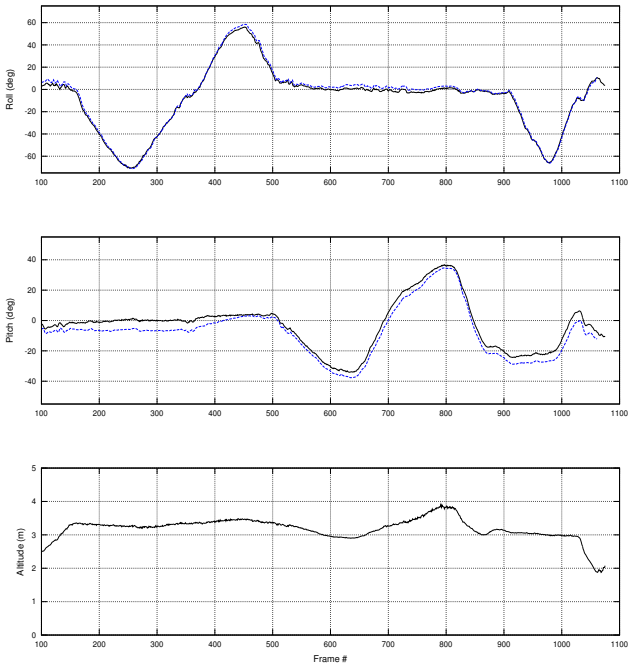

Fig. 8. The estimated roll (top), pitch (centre) and altitude (bottom) of the aircraft during ground-based testing. Also shown for comparison are the roll and pitch reported from the onboard inertials unit (dashed lines). Frames were captured at 8.3Hz.

was $\sigma = 3.5 \times 10^{-2}$m, with very little systematic bias. In other words, 95% (2σ) of the accepted disparity points were reprojected to within approximately 7.0×10^{-2}m of the true radial distance, although much of this error occurred at the corners of the gantry. When represented as a percentage of the estimated radial distance at each viewing angle, the absolute (unsigned) reprojection error was calculated as having a mean of 1.2% and a maximum of 5.6%. This error is a direct consequence of errors in the computed disparity.

C. Testing

Indoor testing has shown that the system is capable of accurately mapping a 3D environment. However, in order to provide guidance to an autonomous aircraft, the system must be capable of estimating the attidude and altitude of the aircraft. We assume that the ground directly beneath and in front of the aircraft (corresponding to the visible FoV, see Table I) can be modelled as a plane. The attitude and altitude of the aircraft can then be estimated relative to the ground plane. A theoretical analysis shows that the disparity measured from the remapped image of an infinite plane, as viewed from an aircraft with some attitude and altitude above the plane, can be expressed as

$$D_{pixel} = \frac{d_{baseline} \times h_{img}}{r} \times \frac{1}{d_{height}} \times [\cos(\theta_{pitch}) \cos(\theta_x - \theta_{roll}) - \tan(\theta_y) \sin(\theta_{pitch})], \quad (2)$$

where the first term is a system constant as described before, with the radial distance replaced by d_{height}, the height of the aircraft above the plane. The bracketed term describes the topology of the disparity surface, which depends on the roll, θ_{roll}, and pitch, θ_{pitch}, of the aircraft as well as two parameters, θ_y and θ_x, which define the vertical and the horizontal viewing direction respectively.

Using (2), the disparity surface for an infinite plane can be predicted given the roll, pitch, and height of the aircraft. Conversely, given the measured disparity surface, the roll, pitch, and height of the aircraft can be estimated by iteratively fitting the modelled surface to the measurements. This is a robust method of estimating aircraft attitude and altitude as the disparity data is used directly, hence there will be an even spread of data points over the fitted surface and the average error will be approximately equal for all data points. Currently, a non-linear derivative-free optimisation algorithm (implemented using the NLopt library [13] based on the Nelder-Mead simplex method [14]) is used to minimise the sum of squared errors between the modelled surface and the measured disparity.

As a next step, the visual system was taken outdoors and put through a series of motions to test the robustness of the SAD disparity algorithm to uncontrolled lighting conditions and textures, as well as to test the feasibility of estimating the aircraft attitude and altitude from the disparity surface fit. The accuracy of the results of fitting the disparity surface was judged by comparing the attitude estimated by the stereo system with the attitude reported by an onboard inertials unit (MicroStrain 3DM-GX2 IMU). The estimated and reported motions of the aircraft during the test are plotted in Fig. 8. It can be seen that fitting the disparity surface model described in (2) to the disparity data is a feasible method of extracting estimates of the aircraft's attitude and altitude. The error between the estimated and reported pitch angle visible around frame 250 is systematic and possibly due to a slight horizontal misalignment between the axes of the IMU and the cameras.

D. Flight testing

The stereo system was tested *in situ* by mounting it to the nose (see Fig. 2b) of a medium-sized fixed-wing aircraft

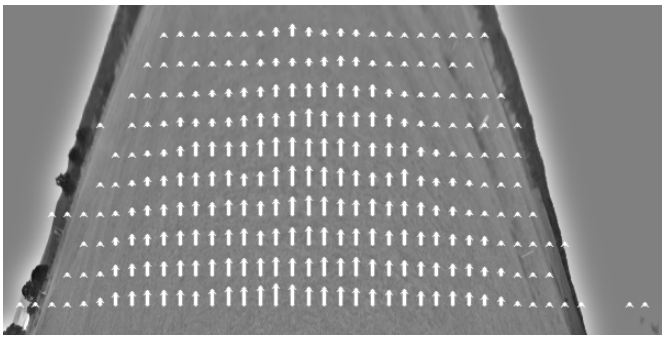

Fig. 9. Remapped image taken from the front camera during flight, overlayed with the computed disparity. The disparity vectors have been scaled to aid visualisation.

(Super Frontier Senior-46, wingspan 2040mm – modified so that the engine and propeller assembly is mounted above the wing). The system parameters used during the flight test are shown in Table I. It would seem that the upper limit to the operational altitude should be infinite as the SAD disparity algorithm is capable of detecting disparities down to zero pixels. However, it was assumed that disparities <1px would fall below the system noise threshold, giving a flight ceiling of $\sim$ 30m for obtaining measureable estimates of ground distance. The flight was conducted 'open-loop', so that manual control was retained throughout the test. During the flight, the stereo image data recorded by the system was compressed and saved to an onboard computer (Digital-Logic MSM945, which incorporates an Intel Core2 Duo 1.5GHz processor) for off-line analysis.

A remapped image taken from the front camera during the flight test is shown in Fig. 9, overlayed with the computed image disparity vectors. It can be seen that the disparity decreases in magnitude towards the top of the image, indicating that the optic axis, and hence the aircraft, is pitched upwards relative to the ground plane. This can also be deduced from the remapped profile of the horizon. For straight and level flight over a horizontal plane, the horizon lines should be vertical and equidistant from the centre of the image, as in Fig. 3b.

A segment of flight during which the aircraft decends and lands has been analysed. The attitude and altitude of the aircraft during the manoeuvre, as estimated by the stereo system, is displayed in Fig. 10. Also shown is the attitude of the aircraft as reported by the onboard IMU. It can be seen that the estimated motion of the aircraft is closely correlated with the motions reported by the IMU. However, there is a significant offset between the estimated and reported attitude of the aircraft at the beginning of the flight segment. There can be several reasons for this offset. Firstly, the intended flight envelope of the stereo system extends from 2m $\sim$ 30m above ground level. It can be seen from Fig. 10 that initially the aircraft is close to the flight ceiling. This causes noisy estimates of attitude and altitude because the disparity generated by the ground is very low and close to the noise threshold. Similarly, when the aircraft decends below 2m, as in frames 205 $\rightarrow$ 220 and again after frame 285, the disparity

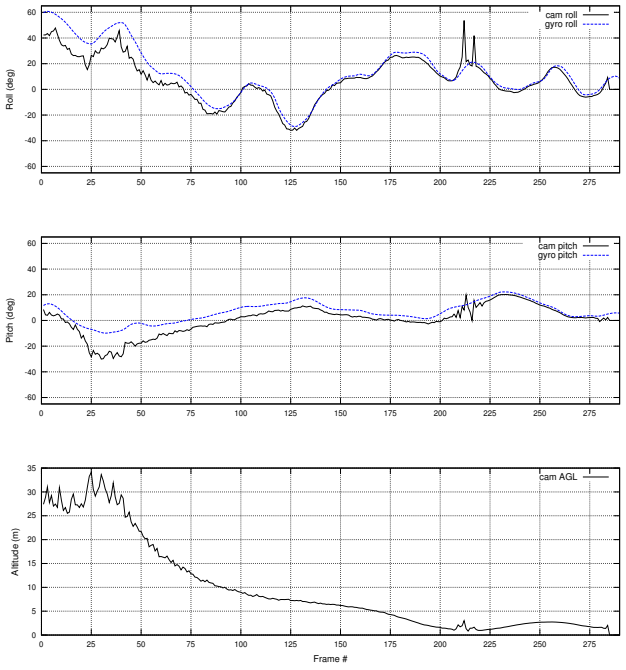

Fig. 10. The estimated roll (top), pitch (centre) and altitude (bottom) of the aircraft during a segment of flight, throughout which the aircraft was under manual control. For comparison, the roll and pitch reported by the inertials unit onboard the aircraft are also shown. Frames were captured at 25Hz.

generated by the ground is greater than the maximum searched disparity, which leads to erroneous estimates or an absence of data if the aircraft descends too low.

Secondly, the attitude and altitude estimates computed by the stereo system are relative to the local terrain. On the other hand, the attitude reported by the IMU is with respect to a more extended world frame. At the location of the flight test, the terrain in the vicinity of the aircraft at the beginning of the flight segment sloped upwards by several degrees to the front and left of the aircraft. This would manifest as a negatively greater estimate of the pitch and roll with respect to the corresponding values reported by the IMU. This is, in fact, what is observed in Fig. 10 and may contribute towards the discrepancy.

Additionally, the attitude reported by the IMU tends to drift over time, particularly following periods of high acceleration. The flight segment analysed in Fig. 10 is several minutes into the flight test and so the attitude reported by the IMU cannot be relied on completely. Therefore, given the accuracy of the system during testing and the close correlation between the estimated and reported motions of the aircraft during flight testing, the data presented in Fig. 10 validates the feasibility of the proposed method for providing guidance information for the stabilisation of aircraft attitude and altitude and for terrain following. A visualisation of the aircraft's motions during the flight test has been constructed from the attitude and altitude estimated by the stereo system, and is shown alongside the raw footage from the front camera in the attached video.

Generally, except for flight above a level plane, estimates of

the local and global attitude will differ. Which estimate is more useful depends on the application but, for terrain following tasks, knowing the attitude of the aircraft with respect to the local terrain can be of more importance than knowing the attitude with respect to a global world frame. Consider a situation in which an aircraft is flying above level ground. If the aircraft's trajectory brings it towards an obstacle, the assumption that the ground can be accurately represented by a level plane will be false. However, the fitted disparity surface will pass through both the ground and the obstacle, causing the aircraft to deviate away from the obstacle when correcting for the estimated error in attitude. The practicality of the stereo system may be extended further by incorporating other guidance strategies such as horizon detection [15]. Indeed, the system has been specifically designed to allow such extensions. At its most fundamental level the mapping produced by the vision system defines a collision-free cylinder (as previously described) which should facilitate navigation through gorges and urban streets, and under overhead structures such as bridges.

Although the results presented here were obtained from off-line analysis of recorded flight data, testing has shown that the onboard computer and vision system are capable of capturing stereo images, performing the remapping and pre-processing and computing stereo disparities at 30Hz. Stereo images are captured over an IEEE 1394b interface at a resolution of $1040\text{px} \times 1040\text{px}$ and synchronised to within $125\mu\text{s}$. Other parameters used are as listed in Table I for the flight test. Currently, the disparity surface fitting is performed off-line. For closed-loop operation the fitting time must be reduced, possibly by subsampling the disparity data before iteratively fitting the model. Alternatively, the attitude and altitude of the aircraft can also be estimated with respect to an assumed ground plane by reprojecting the disparity points into 3D coordinates, as for the reconstruction of the testing arena shown in Fig. 7. This approach is much faster than iteratively fitting a model to the disparity surface, as fitting a plane in 3D coordinates can be completed in a single, non-iterative step. However, the plane fitting is less robust to noise as neither the reprojected points nor the average reprojection error will be evenly distributed over the surface that is sampled. Even so, such an approach may prove necessary for a real-time implementation of the system.

III. CONCLUSIONS & FURTHER WORK

This study has described the design and implementation of a vision system which simplifies the computation of range from stereo in the context of aircraft guidance. Two video cameras are used in conjunction with wide-angle lenses to capture stereo images of the environment, and a special geometric remapping is employed to simplify the computation of range. The maximum disparity, as measured by this system, defines a collision-free space through which the aircraft can fly unobstructed. This system is therefore especially suited to providing information for visual guidance in the context of tasks such as terrain and gorge following, obstacle detection and avoidance, and take-off and landing.

Tests have shown that this system is able to reproduce accurately, and in real time, the three dimensional structure of simple environments, both indoors and outdoors. It has also been shown that the system is able to extract sufficient information to allow the autonomous guidance of an aircraft. Currently, data is recorded to an onboard computer and processed off-line. However, future work will involve 'closing the loop' and conducting all processing onboard the aircraft. As a next step, it is planned to adapt this system to provide autonomous visual guidance of an aircraft performing tasks related to terrain following, object avoidance, and landing.

ACKNOWLEDGMENT

Sincere thanks to Mr. David Brennan, who owns and maintains the airstrip at which the flight testing was done.

REFERENCES

[1] J. S. Chahl and M. V. Srinivasan, "Panoramic vision system for imaging, ranging and navigation in three dimensions," in *Proc. Field and Service Robotics Conference*, Pittsburgh, Aug. 1999, pp. 127–132.

[2] G. L. Barrows, J. S. Chahl, and M. V. Srinivasan, "Biologically inspired visual sensing and flight control," *The Aeronautical Journal*, vol. 107, no. 1069, pp. 159–168, 2003.

[3] M. V. Srinivasan, S. W. Zhang, J. S. Chahl, G. Stange, and M. Garratt, "An overview of insect inspired guidance for application in ground and airborne platforms," *Proc. Inst. Mech. Engnrs. Part G*, vol. 218, pp. 375–388, 2004.

[4] F. Ruffier and N. Franceschini, "Optic flow regulation: the key to aircraft automatic guidance," *Robotics and Autonomous Systems*, vol. 50, pp. 177–194, 2005.

[5] J. C. Zufferey and D. Floreano, "Fly-inspired visual steering of an ultralight indoor aircraft," *IEEE Trans. Robot.*, vol. 22, pp. 137–146, 2006.

[6] C. L. Tisse, O. Frank, and H. Durrant-Whyte, "Hemispherical depth perception for slow-flyers using coaxially aligned fisheye cameras," in *Proc. International Symposium on Flying Insects and Robots*, Ascona, Switzerland, Aug. 2007, p. 123.

[7] S. Thurrowgood, W. Stuerzl, D. Soccol, and M. V. Srinivasan, "A panoramic stereo imaging system for aircraft guidance," in *Proc. Ninth Australasian Conference on Robotics and Automation (ACRA'07)*, Brisbane, Australia, Dec. 2007.

[8] J. M. Roberts, P. I. Corke, and G. Buskey, "Low-cost flight control system for a small autonomous helicopter," in *Proc. International Conference on Robotics and Automation*, Taipei, Taiwan, Sep. 2003.

[9] M. V. Srinivasan, S. Thurrowgood, and D. Soccol, "An optical system for guidance of terrain following in UAV's," in *Proc. IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS'06)*, Sydney, Australia, 2006, pp. 51–56.

[10] (2009, Jul.) Intel. [Online]. Available: http://software.intel.com/sites/products/collateral/hpc/ipp/ipp_indepth.pdf

[11] M. Shimizu and M. Okutomi, "Significance and attributes of subpixel estimation on area-based matching," *Systems and Computers in Japan*, vol. 34, no. 12, 2003.

[12] J. Kannala and S. S. Brandt, "A generic camera model and calibration method for conventional, wide-angle, and fish-eye lenses," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 8, pp. 1335–1340, 2006.

[13] S. G. Johnson. (2009, Jul.) The NLopt nonlinear-optimization package. [Online]. Available: http://ab-initio.mit.edu/nlopt

[14] J. A. Nelder and R. Mead, "A simplex method for function minimization," *The Computer Journal*, vol. 7, pp. 308–313, 1965.

[15] S. Thurrowgood, D. Soccol, R. J. D. Moore, D. Bland, and M. Srinivasan, "A vision based system for attitude estimation of UAVs," in *Proc. International Conference on Intelligent Robots and Systems*, St Louis, MO, Oct. 2009, accepted for publication.