

UAV Localisation & Control Through Computer Vision

William Bath, Jonathan Paxman

ARC Centre of Excellence in Autonomous Systems (CAS)

Faculty of Engineering,

University of Technology, Sydney,

NSW Australia

william.h.bath@student.uts.edu.au, jpaxman@cas.edu.au

Abstract

A real time monocular vision based localisation system developed for use on a four rotor indoor unmanned aerial vehicle (UAV) is presented. The system makes use of augmented/mixed reality techniques, a video camera and a priori knowledge of features (markers), but not their location. Kalman filtering is used to filter the observed vision data and provide an optimal estimate of the state of the UAV. The Localisation and Mapping Application (LAMA) developed offers high levels of connectivity permitting development of distributed systems and is flexible enough to be used for most robotic platforms requiring indoor localisation. Accurate localisation of the UAV in six degrees of freedom, mapping in terms of markers and reliable control of the yaw axis have been achieved.

1 Introduction

Vision is the primary means by which many organisms perceive the world and serves as a basis for their interactions therein. Research in vision for robotic systems aims to achieve the utility and flexibility of biological vision, to enable robotic systems to interact intelligently with complex and dynamic environments.

At present the application of computer vision in the field of localization and mapping within the robotic realm is still in its infancy with methods for dealing with a changing and imperfect environment still in the process of being solved. Conventional SLAM algorithms can require exhaustive computation with the computational expense increasing with the square of the number of features. By investigating a system that uses known and readily detectable features, this system permits localisation within the world reducing the complexity of the SLAM problem to such a degree that localisation, and to a degree mapping, can be performed in real time on a machine that is not highly specified.

At the core of the localisation system are a set of

algorithms that form part of mixed or augmented reality systems. Mixed reality is a system of aligning and overlaying computer generated 3D objects onto a video stream from the real world. The stream and rendered object are combined and fed back to the user, via virtual reality (VR) goggles or other device providing the illusion that the virtual objects actually exist in the real world.

Part of the process in creating this illusion involves obtaining the transformation of the marker with respect to the camera. This transformation can be inverted to find the camera position with respect to the marker; making it a natural candidate for use in a localisation application. We apply these techniques by mounting a small wireless camera on the UAV, and markers of known structure within the UAV operating environment.

Though this process is simple in principle, there are a number of obstacles involved that must be overcome for the system to function accurately. For example, the system must be carefully calibrated to account for the relative translation and rotation of the camera with respect to the centre of mass of the UAV. Using the above simplification the camera would have to be mounted perfectly at the centre of mass of the UAV and also be perfectly aligned along the Z-axis of the UAV. If these criteria were not fulfilled and the camera is not at the centre of rotation, cross coupling between rotations and translations would quickly become a problem leading to inaccurate measurements of the UAV position and hence propagating into poor control. Other issues such as the limitations of the camera and noisy data on marker positions need to be addressed as the control of an indoor aerial vehicle must be precise which can only be achieved through accurate and reliable estimates of the state of the UAV. Furthermore as the camera used has a finite resolution and suffers from distortions induced at manufacture, a means must be determined to improve the quality of the image data to improve the quality of the observations and hence the localising capability of the system.

Kalman filtering is intended to play a substantial part in the development of reliably estimating the state of the UAV to provide accurate information for use in projecting the UAVs' position within the world. At this time due to hardware issues a successful flight has not

been achieved to test the performance of the filter.

2 What is the MXRToolkit?

MXRToolkit is a set of C/C++ libraries created for research and development of Augment Reality (AR) applications. Augmented reality is the process of detecting known markers within a video frame, calculating the Euclidean transformation with respect to the camera, rendering a 3D model to that transformation and overlaying it on the video feed such that the 3D object appears attached to the marker. More accurately MXR allows for the registration of virtual objects within the real world facilitating creation of the illusion that virtual objects exist in the same space as real objects [Malbezin et al., 2002; MXR].

Unfortunately cameras like all sensors exhibit characteristics that can negatively impact the accuracy of the measurements if the data is used without consideration of these effects. These adverse effects or imperfections of the camera can be quantitatively described and accounted for using camera calibration models. In the application discussed, the Brown-Conrady model otherwise known as the “plumb bob model” is used.[Bouguet, 2004]

2.1 Camera Calibration

As suggested above the camera is subject to imperfections both in the manufacture of the components as well as in the assembly thereof. Specifically the dominating effects in the camera used for these experiments were:

- **Tangential Distortion:** A monochromatic distortion effect that is usually the result of the lens being incorrectly aligned at manufacture with respect to the CCD/CMOS element; such that the centre of the lens is not positioned above the centre of the element. This results in asymmetry of the field of view about the centre of the image plane specifically in this instance the field of view left of the centre of the image was 51.5 degrees whilst on the right it was only 32.5 degrees. This results in a greater portion of the more distorted components of the lens occupying the visible image and can cause vignetting.[Baggot]
- **Barrel Distortion:** A monochromatic effect where the image is progressively increasingly warped in on itself the further from the centre of the image it is. This effect is otherwise known as the fish-eye effect.[Baggot]

The afore mentioned aberrations and distortions give rise to large uncertainties and cause erroneous measurements on the data extracted from the image. Using the camera information as is and attempting to input it to the MXRToolkit would be met with limited success. Marker detection algorithms searching for quadrilaterals in the frame can have difficulty reconciling markers from the background; furthermore due to warping of the markers, information regarding the translation and rotation of the marker will be unreliable without accounting for the effects of distortion.

Calibration not only facilitates accurate extraction of information from within the image but also provides a mapping from the 2D image plane (units of pixels) into real world dimensions (i.e. mm) allowing for accurate measurements to be performed upon the images content in

terms of real world units. One calibration tool was provided with the MXR toolkit however it proved inadequate for the application so the “Camera Calibration Toolbox For Matlab” [Bouguet, 2004] was used. Specifically the parameters determined during calibration are focal length, principal point, skew, and distortion coefficients. The values determined for the camera used in the application are shown below for reference.

This model accounted for most but not all of the distortion effects. Specifically it was possible to calibrate the image such that all the effects were corrected for on the right side of the image plane but on the left side as you approached the edge of the image there was still some barrel distortion prevalent. The reason for this is due to the large tangential distortion which occurs due to misalignment of the lens to the CMOS element at manufacture.

The parameters for the camera used in this application are shown below for reference:

Focal Length: $fc = [430.70166 \ 443.99140] \pm [3.91245 \ 4.68504]$
Principal point: $cc = [335.27123 \ 245.19995] \pm [6.79368 \ 7.12709]$
Skew: $\alpha_c = [0.00000] \pm [0.00000] \Rightarrow$
angle of pixel axes = 90.00000 ± 0.00000 degrees
Distortion: $kc = [-0.28303 \ 0.02815 \ -0.00194 \ -0.00008 \ 0.00000] \pm [0.01423 \ 0.01245 \ 0.00287 \ 0.00218 \ 0.00000]$
Pixel error: $err = [2.89176 \ 2.15818]$

2.2 How Does MXR Work?

MXR Toolkit requires only image capture capability, known markers and a Win32 platform in order to function. The process is straightforward and is shown below; the mathematical details of each step are omitted here for brevity.

The process:

1. An image stream is read into the PC via a video capture device and camera.
2. The image is thresholded to simplify processing.
3. Quadrilaterals within the frame are detected.
4. The central content of the quadrilateral is compared to a known set of markers a confidence value is calculated, compared with a threshold for that known marker and a “detected” flag set for true or false if the marker is in frame or not.
5. The mxrFrame object that corresponds to that marker then has the Euclidean transformation (marker with respect to the camera) for the detected marker loaded into an array.
6. Using the transformation a 3D object is rendered and overlayed onto the video feed using the marker in the feed as a registration point.

2.3 MXR Performance

2.3.1 Overview

Testing of the raw data provided by the MXRToolkit was required to experimentally answer such questions as:

- What is the optimal marker size for the application?
- How far away from the camera a marker can reliably be detected?

- What is the mean value of the measured position compared with the actual position?
- How accurate and reliable is each element of the measurement and how it deteriorates with distance?
- What is the quality of the camera used, and how does it affect the quality of the data?

2.3.3 Experiment Results

The results of the experiment seem most promising with the observations exhibiting good stability. There is a very low standard deviation in the measurements on most axes which only increases slightly with distance from the marker. The Z displacement error exhibits approximately linear behaviour across all values and could be resolved simply by multiplying the Z measurement by a constant (in this case approximately 1.17). Roll and pitch exhibited issues regarding the sign of the measurement with it occasionally swapping for no apparent reason; however this phenomenon was only exhibited at the extremes of distance at which the marker can still be correctly identified.

The size of the marker and field of view of the camera does play a significant role in determining the effective range of the system. In these tests a 70mmx70mm marker was used which resulted in an effective range of approximately 545mm before marker identification rapidly deteriorated. Through induction the minimum size of a marker on a ceiling 2500mm above the floor that can be correctly identified would be approximately 321mmx321mm. In reality this marker should probably be closer to 500mmx500mm to ensure accurate detection of it from a line of sight distance of 3.8 meters. The LAMA application has scope for the use of markers of a variety of sizes. The strategic use of markers of different sizes could possibly be used to better effect by interspersing larger markers with smaller markers in between.

The testing did highlight some critical aspects of the system, one of which is camera quality and the ability of a given calibration method to compensate for camera error. Unfortunately the lens appears to have been poorly aligned with the CMOS element at the factory inducing a large amount of tangential distortion and ultimately resulting in the left side of the image being warped. Barrel distortion is also quite prevalent which is common in smaller inexpensive wireless cameras.

The tangential distortion is perhaps of greatest concern as it causes the field of view of the camera to be strongly skewed to one side of the lens. Also, due to the fact that the most distorted parts of the lens are in view, higher order distortional effects of the lens which are not accounted for within the calibration model are dominant towards the left side of the image and may in turn be creating the large discrepancies and non-linearity in position. Furthermore cross-axis coupling was noted between the translational and rotational axes due to this effect. As the marker is moved into the left side of the image plane where the distortion is located, the warping due to the distortion creates the illusion that the marker rolls to keep the facing the camera. The phenomenon obviously induces problems with control feedback as for this particular UAV, accelerations along x and y are a function of roll and pitch respectively.

The field of view as also strongly affected by the tangential distortions as the field of view of the camera

was not spherical but more egg like. The measured field of view was approximate 40 degrees up and 40 degrees down at the centre of the y-axis (image coordinates) and 51.5 degrees field of view to the left and 32.5 degrees to the right.

3 Camera Position Calibration

The MXRToolkit does not provide localisation, it simply provides a Euclidean transformation of a markers position relative to the cameras position; the frame of reference is the camera and as such it is a floating reference point.

A calibration procedure is required so that a series of marker-camera transformations can be combined to produce a UAV-world reference transformation. The procedure needs to account for any arbitrary pose and positioning of the camera on the UAV as well as any arbitrary positioning and pose of the marker with respect to the UAV takeoff position.

3.1 Calibration Procedure

First starting with the Euclidean transformation matrix given by the MXRToolkit, letting $C=Camera$ and $M=marker$:

cT_M

LAMA can then capture the transformation of the marker in the field of view with respect to the camera. Since establishing the position of the UAV takeoff point relative to the marker is of interest, the camera is positioned at the centre of the UAV takeoff point aligned along the Z axis of the marker as shown in Figure 2.




Figure 2: Calibration 1st step.

Capturing the Euclidean transformation of the marker with respect to the UAV pad, which also the initial position of the UAV, letting $O=Uav-Pad$, yields:

oT_M

Then mounting the camera on the UAV and replacing it upon the UAV pad, letting $MC=Mounted-Camera$ provides the transformation:

$${}^{MC}T_M$$




Figure 2: Calibration 2nd step.

These transformations can now be used to determine the transformation of the mounted camera with respect to the UAV like so:

$${}^{UAV}T_{MC} = {}^{UAV}T_M \times ({}^{MC}T_M)^{-1}$$

Note: ${}^B T_A = ({}^A T_B)^{-1}$ and ${}^A T_C = {}^A T_B \times {}^B T_C$

In future the UAVs position can now be calculated using the previously calculated transforms and the transformation of the marker with respect to the camera as so:

$${}^O T(t)_{UAV} = {}^O T_M \times ({}^{MC} T(t)_M)^{-1} \times {}^{MC} T_{UAV}$$

3.2 Marker Chaining

The use of a single marker for localisation is very limiting given that the field of view of the camera is limited. For this method of localisation to be useful, a means must be provided to localise new markers whose transformation is not known and hence localise the UAV when those markers are within the field of view; this school of thought lends itself to chaining markers together where two markers are in frame, one of which is known, the other of which is unknown.

Calculating the position of the UAV with respect to the UAV pad and then calculating the new markers position with respect to the UAV can provide the new markers position with respect to the UAV pad. This can mathematically be proven by first taking the Euclidean transformation of the marker with respect to the UAV pad, in the initial case this would be the first marker used for camera position calibration:

$${}^{UAV-Pad}T_{Marker-Known} = {}^{Mounted-Camera}T_{Marker-Known}$$

Then obtaining the Euclidean Transformation of the unknown marker:

$${}^{Mounted-Camera}T_{Marker-Unknown}$$

The new transformation of the unknown marker with respect to the UAV pad can now be simply calculated by:

$${}^{HPad}T_{MUnknown} = {}^{HPad}T_{MKnown} \times ({}^{MCamera}T(t)_{MKnown})^{-1} \times {}^{MCamera}T(t)_{MUnknown}$$

Unfortunately, erroneous detection of markers and noise in the measurements require that an average of the transformations be taken to establish a reliable estimate. This is not as straightforward as it seems. It is not mathematically valid to average Euclidean matrices on an element by element basis; other means are required.

Spherical Linear Interpolation (SLERP) is a mathematical means of interpolating rotations along a spherical surface. Here it can be used to average the rotations of the Euclidean transformations by first breaking them down into a rotation and a translation, then

converting the rotation matrices into quaternions. A weighted interpolation can then be performed on the old estimate quaternion and the new estimate quaternion. The translations can be averaged on a weighted element basis. The resulting quaternion can then be converted back into a rotation matrix and combined with the averaged translational element to form a new averaged Euclidean transformation.

4 The UAV Platform

The UAV used in this application is a commercially available product sold under the name “Draganflyer”. The Draganflyer is a quad rotored platform consisting of a central hub four carbon fibre arms upon at the end of which geared motors and fixed pitch rotor blades are attached. [DraganflyInnovations]

The Draganflyer platform has been successfully studied for automatic stabilisation [Castillo et al. 2004], but that work was in the context of a very precise, but tethered, high cost and low range position and orientation sensor (by Polhemus). This work forms part of an approach to automatic stabilisation using lower cost and more widely deployable sensors.

The four blades are divided into two sets of counter rotating blades which when all spin equally provide zero net torque on the airframe, yaw is induced by slowing down one pair and speeding up the other. Pitch and roll are also controlled by differential thrust in a similar manner by slowing down one motor whilst increasing the other for any given axis. Thrust controlling altitude along the z-axis is simply the sum of all thrusts generated by the rotors.

Due to the inherent instability of the platform the Draganflyer has low level differential feedback provided by means a 3-axis piezo gyroscope slowing down the dynamics sufficiently that a human pilot can fly it.

The radio communications protocol used to control the UAV is the standard Pulse Position Modulation (PPM). A device called the “PC-Buddy” converts RS-232C serial communications from a PC into PPM signals that can that was then plugged into the back of a Futaba Skysport 4.[RCElectronics] Thus the PC running the LAMA application can control the DraganFlyer directly.

4 Control

The UAV is inherently unstable about the pitch and roll axes which make those the most difficult to control of all axes. Yaw however is the slowest in response and provided an excellent first step towards achieving control as it is, for the most part when simplified, decoupled from all other axes.

A simple PID controller was developed and implemented that worked on the data stream direct from the MXRTToolkit without the use of Kalman filtering or other optimal estimation methods. The noise was reduced by averaging the data read over 12 samples (effectively half a second) which was acceptable due to the slow time constant of the yaw axis.

Control was remarkably easy to establish with some “rule of thumb” estimates applied as to coefficients for proportional, integral and differential. Effectively very little differential control was required (possibly due to the

fact that differential feedback already existed on the lower level controller), with proportional and integral being dominant terms.

Control over X & Y position however was not possible using the raw inputs. Due to the combination of the fast dynamics of the pitch and roll axes, the uncertainty in the pitch and roll from the camera observations and the relatively slow sampling frequency (24fps) compared to the time constant, the data was not of sufficient quality to permit robust control.

At this point control was abandoned until an optimal estimation method could be implemented to improve the quality of the state estimate to a level that reliable control could be implemented.

5 Kalman Filtering

The MXRToolkit only provides information giving displacements. Due to the noise in the measurements the data cannot be used in backward differencing, attempting to do so results in the output of phenomenal accelerations and velocities that are in no way commensurate with the real world behaviour of the UAV. Averaging smoothes the data but at an unacceptable expense of lag with the measurements as due to the short mechanical time constant relative to the sampling frequency (frame rate) of the observation data, averaging provides an unacceptable level of lag.

Kalman filtering is a means of creating an optimal estimation of the state of the UAV whereby the outputs of a theoretical mathematical model describing the dynamics of the UAV are combined with observational data from the MXRToolkit. The two are dynamically blended together based upon the estimate error covariance to produce the best (optimal) estimate based upon the data at hand. In this way the noise level of the measurements is reduced without compromising the responsiveness of the measurements. [Grewal & Andrews, 1993]

5.1 Model Derivation

The UAV was analysed using first principles techniques in order to create a mathematical model of its response to control inputs.

Planar models analysing only a few axis at a time were created to facilitate analysis of the UAVs characteristics. Specifically the following models were developed:

1. A planar model involving yaw, x and y.
2. A planar model involving pitch, y and z.
3. A planar model involving roll, x and z.

Analysis of the equations and critical thinking about the workings of the system revealed that an optimal way to disseminate the system to create a reliable mode was through the creation of two models; one a planar yaw with world-centric θ (yaw), the other a combined ψ (pitch) and Φ (roll) with UAV-centric axes of x, y and z. This approach allows for velocities and accelerations upon all axes to be modelled in UAV-centric coordinates then passed through a filter combining them with world yaw coordinates to generate world-centric models of the UAV in terms of translational/rotational accelerations and velocities. These models are simplified down to first order effects and assume that the UAV will be travelling at low

speeds such that high speed effects like translational lift are negligible

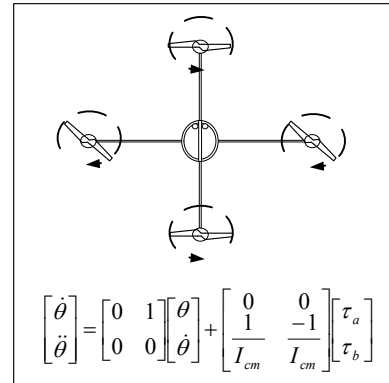
The combined model for pitch, roll, x, y and z is valid if some logical constraints are placed upon the UAVs dynamics. The x and y axes are effectively decoupled from each other so long as:

1. $\psi, \phi \approx 0$
2. $\dot{\theta} \approx 0$

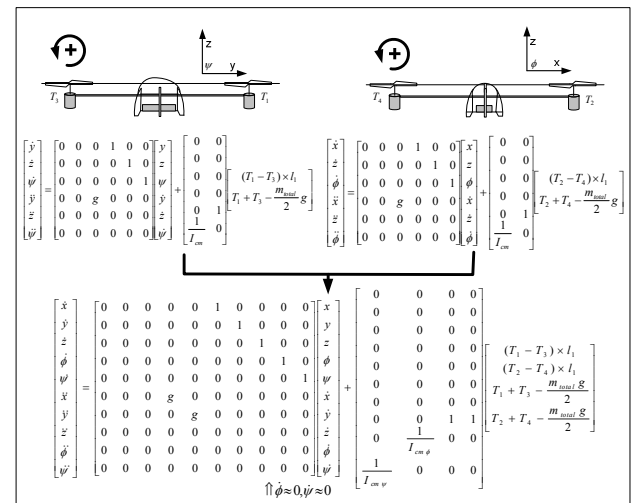
If these conditions remain true the Ψ -y and Φ -x, θ and z dynamics will remain decoupled and the model will work well. Unfortunately world centric yaw does not remain at zero and hence a translator is needed to permit the model to interact with the real world. The translator is based on simple trigonometry accounting for angular displacement and velocity on yaw axis.

The model for determining yaw has its own constraints specifically that there is no significant actuation on pitch or roll whilst yaw is being actuated. This can shift the centre of rotation away from the centre of mass of the UAV during substantial actuation of the pitch or roll axes.

Thus we have two sets of equations governing the systems dynamics. A translator is then implemented to compensate for any yaw by converting world-centric coordinates in which the UAV is at some arbitrary angle of yaw into UAV-centric co-ordinates. Thus the current real world estimate of the UAV can be converted into a form that is manageable by the model created It can be processed, an updated estimate performed and the data retranslated back from UAV-centric coordinates into world-centric coordinates.



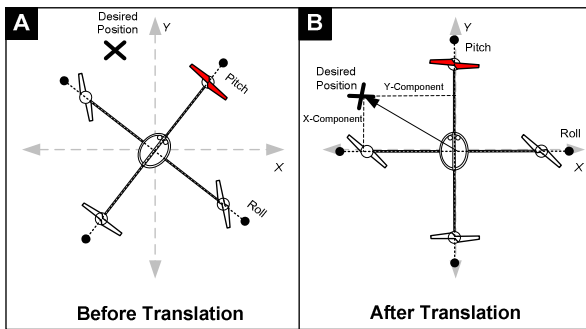
$$\begin{bmatrix} \dot{\theta} \\ \ddot{\theta} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \theta \\ \dot{\theta} \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ \frac{1}{I_{cm}} & \frac{-1}{I_{cm}} \end{bmatrix} \begin{bmatrix} \tau_a \\ \tau_b \end{bmatrix}$$



$$\begin{bmatrix} \dot{y} \\ \dot{z} \\ \dot{\psi} \\ \dot{\phi} \\ \dot{\chi} \\ \dot{x} \\ \dot{y} \\ \dot{z} \\ \dot{\psi} \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & g & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} y \\ z \\ \psi \\ \phi \\ \chi \\ x \\ y \\ z \\ \psi \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 1 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} (T_1 - T_2) \times I_1 \\ T_1 + T_2 - \frac{m_{total} g}{2} \end{bmatrix}$$

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{z} \\ \dot{\phi} \\ \dot{\psi} \\ \dot{\chi} \\ \dot{x} \\ \dot{y} \\ \dot{z} \\ \dot{\psi} \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & g & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ \phi \\ \psi \\ \chi \\ x \\ y \\ z \\ \psi \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ \frac{1}{I_{cm \psi}} & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} (T_1 - T_2) \times I_1 \\ (T_2 - T_1) \times I_1 \\ T_1 + T_2 - \frac{m_{total} g}{2} \\ T_2 + T_4 - \frac{m_{total} g}{2} \end{bmatrix}$$

$\uparrow \dot{\phi} \approx 0, \dot{\psi} \approx 0$



6 Conclusion

The work has produced a number of successful outcomes, principally:

- The application provides the means to simultaneously map the environment in terms of markers and localizing the UAV within that environment. This data is provided both quantitatively and visually in the form of a graphical representation of the known world.
- The system achieved successful automatic stabilisation of the yaw axis dynamics in the test flights.
- Monocular localisation is possible and yields good results for use in the indoor arena (optical flow techniques and feature extraction are on the agenda for the MXR projects future releases, this will permit greater expansion of the system and reduce a priori preparation of the environment).
- A means of camera position calibration that allows for an arbitrary mounting position and attitude of the camera has been created which is user friendly and performs well.
- Kalman filtering has been implemented, though it has not yet been fully tested in flight.

The undertaking of the project revealed several critical areas that require attention for possible future implementations of the system, these are outlined below.

- The limitations in the camera, specifically strong tangential distortions and barrel distortions (fish eye effect) limited the ability to acquire solid measurements of state. The results of these distortions were errors in the x, y, z, pitch and roll measurements of the UAV. Furthermore due to the distortions some measurements exhibited some non-linearity particularly regarding x and y. The purchase of a camera with improved manufacturing quality should correct these problems and represent substantially improved accuracy in the measurements as the issue is caused by the specific hardware implementation, not the LAMA system itself.
- Adjustable and intelligent thresholding of the images will be required to permit flight outside the laboratory environment. Logically this will need to be implemented within the MXR Toolkit. The authors of the toolkit have agreed to include intelligent thresholding in new releases of the toolkit. The auto white balance of the camera also permits some correction of the light; however the auto white balance of this particular

camera can create a large amount of noise in over exposed areas within the frame. An improved camera should correct this problem as well as correct the other image issues especially those pertaining to lens alignment.

The study has revealed areas that require improvement that are of a critical nature to the application and has provided potentially useful data for others engaged in research in similar areas. The results achieved are very positive and represent a large step forward in the case for monocular vision based localisation and control.

Acknowledgements

This work is supported by the ARC Centre of Excellence programme, funded by the Australian Research Council (ARC) and the New South Wales State Government.

References

- [Baggot] Baggot, *Aberrations*, <<http://cvu.strath.ac.uk/courseware/msc/jbaggot/aberration/aberration.html>>.
- [Bouguet, 2004] J. Y. Bouguet, *Camera Calibration Toolbox for Matlab*, 2004 <http://www.vision.caltech.edu/bouguetj/calib_doc/htmls/parameters.html>.
- [Castillo et al., 2004] P. Castillo, A. Dzul, and R. Lozano, *Real-Time Stabilization and Tracking of a Four-Rotor Mini Rotorcraft*, IEEE Transactions on Control Systems Technology, 12(4), July 2004.
- [DraganflyInnovations] DraganflyInnovations, *RCToys.com//Draganflyer V Ti Thermal Intelligence Gyro Stabilized Helicopter*, viewed 16 May 2005 <<http://www.rctoys.com/draganflyer5.php>>.
- [Grewal and Andrews, 1993] M. S. Grewal and P. A. Andrews, *Kalman Filtering: Theory and Practice*, 1993 Prentice Hall.
- [Malbezin et al., 2002] P. Malbezin, W. Piekarski, and Thomas, B, *Measuring ARToolkit accuracy in long distance tracking experiments*, 2002, Tinmith Augmented Reality Research, <<http://www.tinmith.net/papers/malbezin-art-2002.pdf>>.
- [McKerrow 2004] P. McKerrow, *Modelling the Draganflyer four-rotor helicopter*, in Proceedings of the IEEE Conference on Robotics and Automation, April 2004, (4) 3596- 3601.
- [MXR] Mixed Reality Lab Singapore, *MXR Toolkit*, <<http://www.mxrtoolkit.sourceforge.net/>>.
- [RCElectronics] RCElectronics ., *PC Serial Port to Buddy box Interface*, <<http://www.rcelectronics.co.uk/pcbuddy.htm>>.