CONTROLO 2008
8th Portuguese Conference on Automatic Control

EXPERIMENTS ON VISION BASED CONTROL OF AN INDOORS RF BLIMP

Marco Neves, João Marques, Alexandre Bernardino, José Santos-Victor

Institute for Systems and Robotics, Instituto Superior Técnico, Lisboa, Portugal.

Abstract: In this paper we describe a control methodology for the navigation of an autonomous lighter-than-air vehicle based on visual input. The control system consists in a Linear Quadratic Regulator (LQR) and a Linear Quadratic Gaussian observer (LQG). Due to non-linearities in the actuators and in the radio frequency control system, some non-linear control terms are added. The additional terms lead to significant improvements in the response to sudden orientation manoeuvres, either due to command following or to disturbance rejection. In this paper we explain our approach and illustrate experimentally the benefits due to the additional non-linear control terms. Both simulations and experiments with a real indoor blimp are shown, and validate our proposed methodologies.

Keywords: Vision-Based Control, Lighter-than-air Vehicles, LQR-LQG Control, Real-Time Systems, Trajectory Tracking.

1. INTRODUCTION

There is an increasing interest for lighter-than-air vehicles, to carry out tasks with little or no supervision in low speed and low altitude applications. Significant research efforts are being made in the utilization of such vehicles in dangerous and inaccessible environments, e.g. for environmental mapping [1] and planetary exploration [2]. Such systems have advantages over other airships (airplanes, helicopters) because they are inherently stable, but have the disadvantage of being slower and more subject to disturbances due to their light weight and payload. There is, thus, the need to design appropriate control methodologies to address their dominant dynamics properties. The goal of this work is to improve the control methodologies for the navigation of such vehicles by proposing effective solutions to address some common hard-to-model non-linearities in this class of vehicles.

The present research comes in the line of previous works done at the Computer Vision Lab at ISR-IST, having as main goal the implementation of vision based control and navigation of an autonomous indoor blimp. The first step towards the mentioned goal was made in the NARVAL project (ESPRIT-LTR Project 30185, NARVAL—Navigation of Autonomous Robots via Active Environmental Perception.), whose application was underwater environmental mapping. In that project, a blimp was used as a dry setup to develop and test both the visual servoing routines [3], mosaic building [4], and mosaic based navigation [5]. At this point, the developed control algorithms were based on manually tuned PID controllers. As to the vision algorithms, they were made assuming an approximately planar world. Recently the small blimp was replaced by a medium blimp. All the dynamical models were re-identified and a simulator was developed in order to test novel control algorithms, in particular a Linear Quadratic Regulator (LQR). In this paper we extended the previous work with the inclusion of an observer (LQG), non-linear control terms to improve manoeuvres and an overall better tuning of the controller design parameters.

The paper is organized as follows. Section 2 describes the blimp's system, including its dynamical model, actuation and sensing subsystems. Section 3 describes the application of the Linear Quadratic Regulator (LQR), the observer and additional non-linear control terms to the real system, discussing practical aspects of the implementation. In Section 4, both simulations and experimental results with the real blimp are performed, evaluated and discussed. Finally, in Section 5, we present conclusions and directions for future work.

2. BLIMP SETUP

The blimp used in this work is a medium size indoor vehicle with 2 meter long by 1 meter diameter, filled with Helium. A gondola is attached bellow the blimp, containing all the sensors and motors, except the tail motor. Figures 1 and 2 show our current setup. The blimp has four DC motors with voltage command (between -5V and 5V), one servo motor that allows to control the blimp's altitude, two frontal non differential motors for velocity control, and one tail motor for yaw control. The blimp has four tail stabilizers. The gondola carries one micro camera with an integrated TV transmission kit, one PWM receiver and two batteries, one to feed the motors and PWM receiver, and the other one to feed the camera and TV transmitter.

2.1. Decoupled Analysis

Figures 3 and 4 show the coordinate frames in the blimp's model and the placement of the thrusters. Since the motors are not aligned with the axis crossing the centre of mass of the blimp, and the centre of mass is not coincident with the centre of buoyancy, there is a significant coupling between the stern (tail motor) and frontal motors. The vehicle dynamics has a non-linear behaviour, which poses some challenges in terms of controller design. However, because the actuators have low-power, the blimp's dynamics is very slow, so we will consider this coupling negligible.

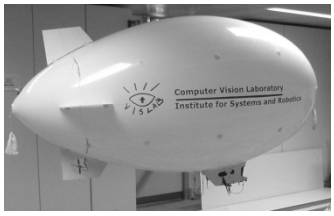

Figure 1 – The Blimp at the ISR Computer Vision Laboratory (VisLab).

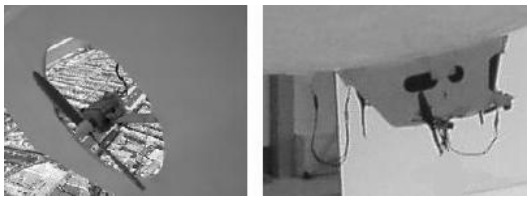

Figure 2- Detail of the stern thruster in the tail (left), and the gondola with the forward thrusters, pitch control servo and video camera (right).

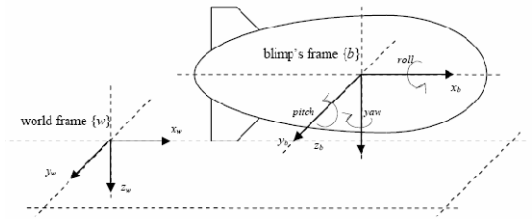

Figure 3 - Coordinate frames in the blimp's model.

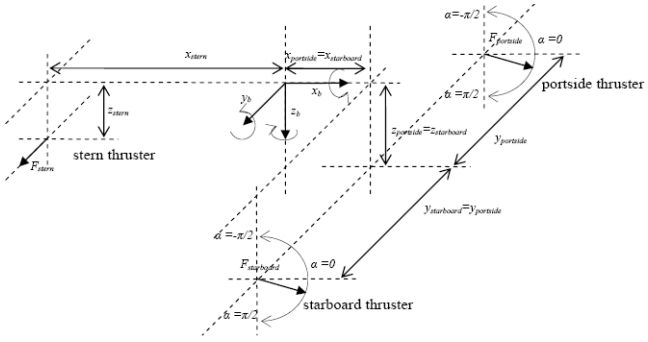

Figure 4 - Placement of the thrusters in the blimp's body.

In practice this is an acceptable approximation and simplifies significantly the controller's design. However, it should be noticed that this is only true as long as the controller is dimensioned to have slow dynamics, otherwise the mentioned decoupling would not be valid.

With the assumptions stated above, the system may be separated in two non correlated sub-systems, one that describes the blimp's behavior in the vertical plane (XZ system) and another one that modulates the vertical rotational behavior (heading system). The complete state-space model can thus be separated in two decoupled sub-systems, reducing the complexity of the control system, and mostly, allowing us to design distinct individual controllers for each sub-system.

The heading sub-system has an input that corresponds to the actuation in the stern thruster, and its state vector is, $x(t) = [v_y(t), w_x(t), w_z(t), \phi(t), \psi(t)]^T$, with $\{\phi, \psi\}$ = (roll, yaw).

The XZ sub-system has two inputs, the actuations in the servo and in the frontal thrusters. The state vector has the following variables:

$$x(t) = [v_x(t), v_z(t), w_y(t), z(t)]^T$$

2.2. Visual Sensing

The sole sensor employed in this work is a video camera installed in the gondola, looking downwards. We assume that the observed environment is approximately planar. This is a reasonable assumption if the travelling height is significantly larger that the ground relief variations. In the experiments, we use a floor-map image with 4.8m×3.6m, consisting of an aerial photo of the area surrounding the Instituto Superior Técnico – Alameda Campus.

In previous work we have developed real-time computer vision algorithms for blimp's position and velocity estimation with respect to the floor-map [5]. The image captured by the camera aboard the blimp is transmitted to a computer that estimates blimp's position and posture in the world reference frame {x,y,z,roll,pitch,yaw}, as well as the linear and angular velocities in the blimp's frame. With these variables available, the controller computes desired actuation to the motors. The command values are sent

back to the blimp, through a radio frequency system connected to the computer.

3. CONTROL METHODOLOGIES

In this section we present the controllers developed in this work for blimp's position and heading control. The design methodologies are based on linearized models of the decoupled heading and XZ subsystems, shown in Appendix A. Linear control (LQR) and observer (LQG) design methodologies are used to provide the system with good stability, disturbance rejection and noise immunity characteristics. Then, additional non-linear control terms are added to improve performance in command following and stern manoeuvres.

3.1. Linear Quadratic Regulator (LQR)

We consider each of the Blimp subsystem's (Heading and XZ) described by the Linear Dynamics equation:

$$\begin{cases} x(k+1) = Ax(k) + Bu(k) + w(k) \\ y(k) = Cx(k) + Dw(k) + v(k) \end{cases}$$

where w and v are Gaussian noise processes with covariance Q_N and R_N respectively. The (A, B, C, D) and the (Q_N, R_N) matrices are presented in appendices A and B respectively.

As the state variables are directly measurable with the image processing routines, we have first designed an LQR controller [6]. This type of controller computes a gain vector K that generates the control law $u(k)$ by:

$$u(k) = -Kx(k)$$

where $x(k)$ is the discrete state of the system. The regulator gain vector K is obtained by minimization of the quadratic cost functional:

$$J_c = \frac{1}{2}\sum_{k=0}^{\infty} [x^T(k)Q_1x(k) + u^TQ_2u(k)]$$

The manipulation of weight matrices, Q_1 and Q_2, allow us to quantify the relative importance each state variable has in the cost function and the amount of "energy" we are willing to spend in the control law. A proper definition of the weigh matrices allow to achieve a balance between a good system's dynamics, low power consumption, and low noise sensitivity. This method has the additional advantage of providing always stable control systems, thus we can concentrate the design effort in tuning the closed-loop system performance. In our case, these matrices tuned with the help some simulations and experimental work, until desired behavior was achieved.

3.2. The Observer (LQG)

To improve our system's noise immunity and robustness, we have included a state observer in the controller. Observers are typically used in two cases: (i) to estimate missing variables needed for state feedback control laws (e.g. LQR), or (ii) to filter process and sensor noises. In our case, we are mostly interested in the filtering properties of the state observer because our sensor (the video camera) already computes the whole state vector.

The Observer computes an estimate of System's state through the following equations:

$$\begin{cases} \hat{x}(k) = \bar{x}(k) + L\big(y(k) - C\bar{x}(k)\big) \\ \bar{x}(k+1) = A\bar{x}(k) + Bu(k) \end{cases}$$

The gain L is computed using optimal techniques (LQG or, equivalently, Steady State Kalman Filter) that minimize the cost function:

$$J_o = \frac{1}{2}\sum_{k=0}^{\infty} \Big[\big(y(k) - C\hat{x}(k)\big)^T\big(y(k) - C\hat{x}(k)\big)\Big]$$

With the proper adjustment of the process and sensor noise matrices (Q_N, R_N), we aim at reducing both sensor noise and errors arising from inaccurate dynamic modeling.

3.3. The non-linear control terms

During this project, we came across several problems due to blimp's actuator non-linear aspects: (i) motor static friction, (ii) exaggerated high torque in the yaw subsystem provided by the stern thruster, and (iii) low yaw damping at small velocities. To address these problems we propose the inclusion of the following non-linear terms.

Jump-Start: As it is known, a DC motor has different friction characteristics when stopped and when rotating. Static friction requires high control values to remove the motor from rest. This raises a problem in producing small corrections in the blimp's trajectory because the controller will not drive a sufficiently high value to start the motor. We have addressed this problem with an approach that we call "jump-start". This works as follows: we check for transitions from stopped to rotation or changes in movement direction and, when this is the case, a large control value to the motor is applied during five consecutive sample times (about 0.6 seconds). This way, the motor starts spinning before the low levels are requested, allowing the achievement of motor angular velocities that otherwise could not be used. Obviously this rule is motor dependent and should be tuned for different systems.

Yaw Control Saturation: In this blimp, the tail motor produces a very large yaw torque even with low actuations, because it is placed almost 1m away from the blimp's centre of rotation. To avoid reaching exaggerated angular velocities in the yaw subsystem that would be difficult to compensate, we saturate the available output values to half the full range of the motor actuation.

Yaw Non-Linear Damping: Lighter-than-air vehicles have very little damping at low velocities. This is a problem, because it makes yaw very difficult to stabilize. After trying thoroughly to solve this problem, we found an effective non-linear control law able to address the problem. The law takes into account the z axis angular velocity (w_z) and the yaw (direction) error and its implementation is:

$u(k) = -2\text{sgn}(w_z)\Delta_{out}, \quad |w_z| > 0.11 \text{ and } |\psi_{err}| < 0.5$
$u(k) = u_{LQR}(k), \quad \text{otherwise}$

The non-linear control term is active if the blimp has a large yaw angular velocity and the yaw error (in absolute value) is lower than a desired value. When these conditions are met, the control will actuate in the opposite direction with an amplitude twice the maximum range (Δ_{out}) available to the LQR controller (because stern motor is artificially saturated, this actually corresponds to the maximum possible stern actuation).

With the introduction of the angular velocity (w_z) constraint, we guarantee that this actuation is only active while the blimp is spinning too quickly, enabling a fast reduction of the blimp's rotational momentum, and then allowing the LQR controller to do the remaining fine control adjustments.

4. RESULTS

In this section we present and discuss results obtained in simulations and experiments performed with the real setup. In previous work we have developed a simulator for the system, and use it as a first test bet for controller design. It provides a good approximation to reality except for the unmodeled nonlinearities described in the previous section.

Both the Controller and Observer matrices (K, L) were defined and tuned manually taking into account the knowledge about the state and noise units and amplitudes. Despite this may seem a tiresome process, it has to be done only once, and the results proved to be effective for all tested scenarios.

4.1. Path Following with LQG.

In the first tests, we will show simulations done with the LQR controller without the observer. These simulations will make clear the motives that led us to include an observer in the control system. In Figure 5 we can see the results of tracking a complex path with the LQR controller only. In some parts of the path, there are large deviations from expected mainly when high curvatures are demanded.

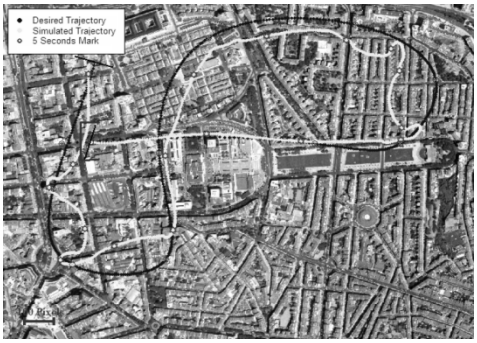

Figure 5 - Simulation: path following experiment using only the LQR controller.

Our next step was to include an observer (LQG) in the controller. This new feature will allow a better estimation of the state vector used by the LQR by filtering sensor noise and model errors, leading to greater accuracy, precision and robustness. The gain vectors for both LQR regulator and LQG observer are shown in appendix B. The results of the inclusion of the LQG can be observed in Figure 6.

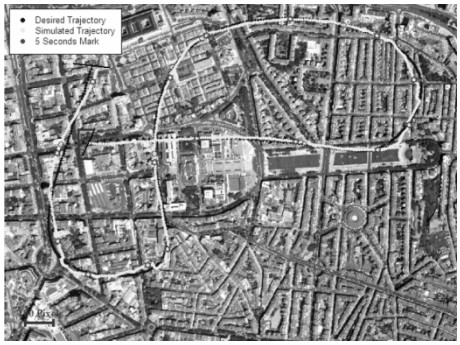

Figure 6- Simulation: path following using LQR and observer in the control.

To provide a time reference, we have marked the simulated trajectory with 5 seconds intervals (points over the trajectory). This way, one may have an idea of the time sequence along the path.

This controller brought significant improvements in terms of accuracy and precision, even in demanding trajectories. However, the new controller is about 1.5 times slower (the full path takes approximately 75 seconds to be travelled, in the new controller, against about 50 seconds without the LQG). Notwithstanding, we think the reduction in tracking speed is largely compensated by the drastic improvements in tracking precision.

In order to better compare the two trajectories, we have computed the statistics of tracking errors along the paths. We define the error as the distance between every point of the blimp's trajectory and the closest point in the desired path. The error statistics obtained are presented in Table 1. The introduction of the observer seems to be the key for this good performance. The LQG has proved very efficient in filtering the noise produced in the vision based estimation algorithms. However, this is still a simulation and does not reveal all aspects of the system behavior.

	LQR without observer	LQR with observer
Total Error	48 455 px	10 282 px
Mean Abs.	39.59 px	8.85 px
Std. Dev.	34.50 px	12.30 px

Table 1 - Error statistic of the simulated trajectories

4.2. Heading Control with yaw non-linear term

After promising tests in the simulated environment, we carried on to experiments with the real setup using the same controller. During these experiments we found that the orientation controller could not achieve a satisfactory performance. In Figure 7 we can observe the yaw system response to a change in

reference of about 90 deg.

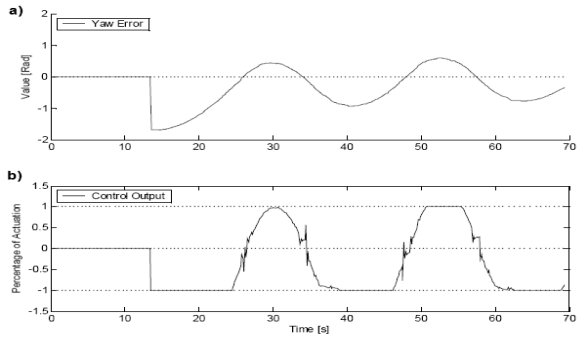

Figure 7- Step response to an approximately 90º yaw error without the non-linear control term. a) yaw error, b) yaw control output

We can notice that the control system has difficulty in stabilizing the yaw value close to the reference value, resulting in an oscillatory behavior. We have thus included the yaw control saturation and yaw non-linear damping techniques described previously. The yaw non-linear damping algorithm takes in account the z angular velocity and the yaw error, and after some tuning, appropriate threshold values for w_z and yaw error were chosen.

In Figure 8 we can observe the performance of the control systems with the added non-linear terms. In the figures, the gray zone indicates the non-linear control zone of actuation. As a result, we have a much quicker stabilization of the yaw value, without overshoot or oscillations. This way, we manage to properly align the blimp with the desired orientations, and avoid the oscillatory behavior around the desired yaw, like it was typical before.

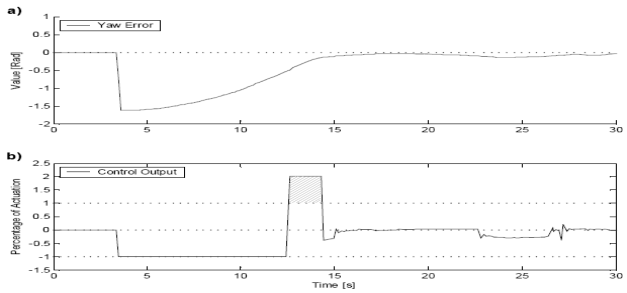

Figure 8 - Step response to an approximately 90º yaw error with the non-linear control term. a) yaw error, b) yaw control output

4.3. Positioning and Station Keeping

In this experiment, the objective is to maintain altitude and, starting from rest, stabilize around a given position. Results are shown in Figures 9 to 12. As one can see in Figure 9, the blimp converged to the desired point with a smooth trajectory. It starts from rest with an initial yaw error, rotates until it reaches the desired yaw, and then proceeds to the set point. This behavior is only possible because of the added yaw non-linear damping term that enables the controller to stop the blimp's rotational inertia and align it with the desired value without overshoot and oscillations. The effect is visible in Fig. 11, from the 17 to the 20 seconds, with an actuation above the maximum available to the LQR controller. In this plot we can also see the jump start effect, which enables the motor to achieve small angular velocities. As can be seen from the control outputs profiles, the desired actuation is quite smooth and is effective in its objective. The servo control is also very smooth and manages to keep the altitude error bellow 30 cm.

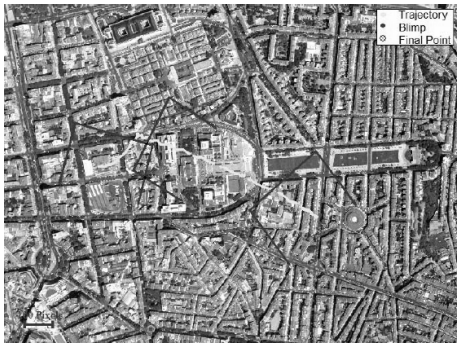

Figure 9 - Experimental: Step Response.

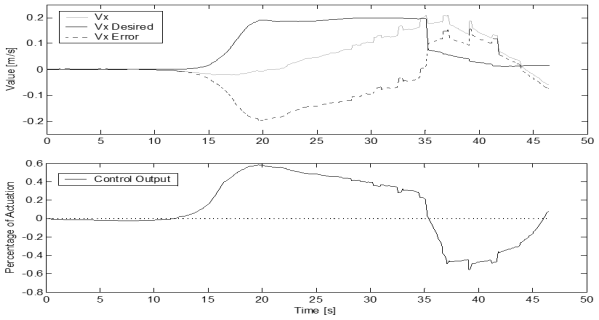

Figure 10 - Step Response. Top: Vx control variables; Bottom: Vx control output.

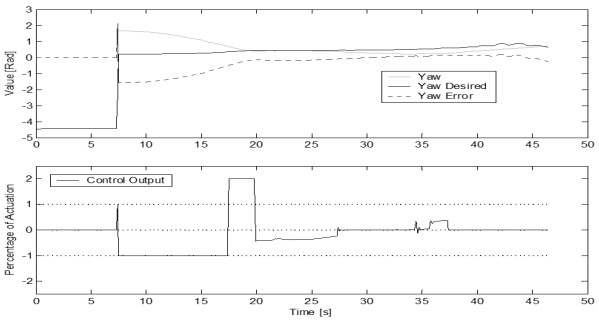

Figure 11 - Step Response. Top: yaw control variables; Bottom: yaw control output.

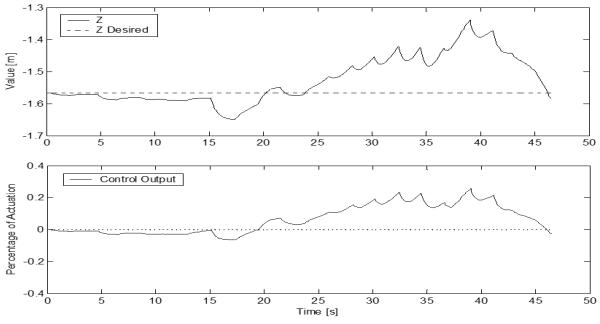

Figure 12 - Step Response. Top: z control variables; Bottom: z control output.

5. CONCLUSIONS AND FUTURE WORK

This work's main goal was to improve the blimp's performance in path following and station-keeping tasks. We have introduced real improvements in the controller's performance with the introduction of a LQG observer and additional non-linear terms. This was first noticed in the simulations performed, in which our controller showed to be very robust to noise, which led to precise and accurate path following behavior. When we used this same controller in the real setup, there were expected differences to the simulations. However, with the introduction of non-linear correcting terms the control law, we have achieved good performances in both path following and station keeping.

In future work, it would be important to characterize the noise introduced by the vision algorithm, leading to a better filtering action from the observer and consequently to a better overall performance. Also, a good alignment of the buoyancy center with the impulsion center, if achievable, will allow a better path following and station keeping.

ACKNOWLEDGMENTS

We would like to thank to all people from VisLAb involved in this project and to all former researchers that, through time and in different fields, have been essential to the continuous developments is this project and similar ones.

This work was supported by Fundação para a Ciência e a Tecnologia (ISR/IST plurianual funding) through the POS_Conhecimento Program that includes FEDER funds.

REFERENCES

[1] E. Hygounenc, I. Jung, P. Souères and S. Lacroix. The Autonomous Blimp Project of LAAS-CNRS: Achievements in Flight Control and Terrain Mapping, The International Journal of Robotics Research, Vol. 23, No. 4–5, April–May 2004, pp. 473-511.

[2] A. Elfes, J. Montgomery, J. Hall, S. Joshi, J. Payne and C. Bergh. Autonomous Flight Control for a Titan Exploration Aerobot.

[3] S. Zwaan, A. Bernardino and J. Santos-Victor. Visual station keeping for floating robots in unstructured environments, Robotics and Autonomous Systems (Elsevier), 39(3-4), June 2002.

[4] N. Gracias and J. Santos-Victor . Underwater Mosaicing and Trajectory Reconstruction using Global Alignment, IEEE OCEANS 2001, Honolulu, November 2001.

[5] N. Gracias, S. Zwaan, A. Bernardino and J. Santos-Victor, Mosaic Based Navigation for Autonomous Underwater Vehicles, IEEE Journal of Oceanic Engineering, October 2003.

[6] G. Franklin, J. Powell and M. Workman, "Digital Control of Dynamic Systems", Prentice Hall, Third Edition, 1997.

APPENDIX

A. State-Space model of the decoupled sub-systems

We present in the following, the equations for the discrete state-space representation of the heading and XZ sub-systems.

A.1. Decoupled Heading Sub-System

$$A = \begin{bmatrix} 0.9608 & -0.0150 & -0.0004 & -0.2041 & 0 \\ -0.0598 & 0.9002 & 0.0006 & -1.3595 & 0 \\ -0.0033 & -0.0032 & 0.9836 & -0.0740 & 0 \\ -0.0021 & 0.0670 & -0.000 & 0.9515 & 0 \\ -0.0001 & -0.0001 & 0.0694 & -0.0026 & 1.0000 \end{bmatrix} \qquad B = \begin{bmatrix} 0.0186 \\ -0.1181 \\ -0.1384 \\ -0.0042 \\ -0.0049 \end{bmatrix}$$

$$C = [0 \quad 0 \quad 0 \quad 0 \quad 1] \qquad D = [0]$$

A.2. Decoupled XZ Sub-System

$$A = \begin{bmatrix} 0.9889 & -0.0006 & 0.0121 & 0 & 0 \\ -0.0001 & 0.9698 & 0.0011 & 0 & 0 \\ 0.0077 & 0.0029 & 0.9381 & 0 & 0 \\ 0.0069 & 0.0686 & 0.0001 & 1.0 & 0 \\ 0.0700 & 0 & 0 & 0 & 1.0 \end{bmatrix} \qquad B = \begin{bmatrix} 0.0346 & -0.0103 \\ -0.0015 & 0.0377 \\ 0.0815 & 0.0531 \\ 0.0001 & 0.0013 \\ 0 & 0 \end{bmatrix}$$

$$C = \begin{bmatrix} 1.0000 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1.0000 & 0 \\ 0.0700 & 0 & 0 & 0 & 1.0000 \end{bmatrix} \qquad D = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix}$$

B. Controller Gains - LQR and Observer

The gain matrices of the LQR and observer that led to the best performance controllers are presented below.

B.1. Heading Sub-system

$$Q_1 = diag(1,1,3,1,18), \quad Q_2 = 1$$

$$K = [0.053544 \quad 0.170343 \quad -1.735743 \quad 0.937068 \quad -3.232513]$$

$$Q_N = diag(10,10,50,10,150), \quad R_N = 0.01$$

$$L = [-0.005398 \quad 0.561435 \quad 0.112484 \quad -0.001 \quad 1.014475]^T$$

B.2. XZ Sub-system

$$Q_1 = diag(4,0.01,0.05,2,0.1), \quad Q_2 = diag(1,1)$$

$$K = \begin{bmatrix} 1.900133 & 0.021457 & 1.94884 & 0.177211 & 0.291133 \\ -0.052148 & 0.15544 & 0.031419 & 1.327408 & -0.039502 \end{bmatrix}$$

$$Q_N = diag(3,1,1,2,2), \quad R_N = diag(0.01,0.01,0.01)$$

$$L = \begin{bmatrix} -0.978375 & -0.000092 & 0.000581 \\ -0.000824 & 0.352473 & 0.0 \\ 0.039997 & 0.013009 & 0.000024 \\ 0.011739 & 1.038486 & 0.0 \\ 0.000592 & 0.0 & 0.995050 \end{bmatrix}$$