We are IntechOpen,
the world's leading publisher of
Open Access books
Built by scientists, for scientists

5,800
Open access books available

144,000
International authors and editors

180M
Downloads

154
Countries delivered to

Our authors are among the

TOP 1%
most cited scientists

12.2%
Contributors from top 500 universities

BOOK CITATION INDEX
CLARIVATE ANALYTICS
INDEXED

WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?
Contact book.department@intechopen.com

Numbers displayed above are based on latest data collected.
For more information visit www.intechopen.com

Flight Control System Design Optimisation via Genetic Programming

Anna Bourmistrova and Sergey Khantsis
Royal Melbourne Institute of Technology
Australia

1. Introduction

This chapter presents a methodology which is developed to design a controller that satisfies the objectives of shipboard recovery of a fixed-wing UAV. The methodology itself is comprehensive and should be readily applicable for different types of UAVs and various task objectives. With appropriate modification of control law representation, the methodology can be applied to a broad range of control problems. Development of the recovery controller for the UAV Ariel is a design example to support the methodology.

This chapter focuses on adaptation of Evolutionary Algorithms for aircraft control problems. It starts from analysis of typical control laws and control design techniques. Then, the structure of the UAV controller and the representation of the control laws suitable for evolutionary design are developed. This is followed by the development of the general evolutionary design algorithm, which is then applied to the UAV recovery problem. Finally the presented results demonstrate robust properties of the developed control system.

2. Aircraft flight control

2.1 Overview of types of feedback control

Not unlike the generic control approach, aircraft flight control is built around a feedback concept. Its basic scheme is shown in Fig. 1. The controller is fed by the difference between the commanded reference signal r and the system output y. It generates the system control inputs u according to one or another algorithm.


Figure 1. Feedback system (one degree of freedom)

One of the main tasks of flight control as an engineering discipline is design of the controllers which enable a given aircraft to complete a defined mission in the most optimal manner, where optimality is based on mission objective. A number of techniques of producing a required controller have been developed over the last hundred years since the first altitude hold autopilot was introduced by Sperry Gyroscope Company in 1912. Most of the techniques make certain assumptions about the controlled system (i.e. the aircraft), most

notably linearity of its dynamics and rigidity of the airframe, to simplify synthesis of the controller. A brief overview of types of feedback control is presented below.

2.1.1 On-Off control

The simplest feedback control is the *on-off control*, also referred to among engineers as *bang-bang control*. This control law can be expressed as follows:

$$u = u_0 + \begin{cases} k, z < a \\ 0, z > a \end{cases}$$

(1)

where u_0 and k are arbitrary offset and gain suitable for the given system, and a is the threshold (set point). It can be extended to a multiple choice situation.

This law is not particularly suitable for direct flight control in normal situations. Indeed, with only two (or several) discrete control input levels, the system output will tend to oscillate about the set point, no matter how well damped the system is, because the control signal will not switch until the set point is already passed. Moreover, if the dynamics of the system is fast or a significant amount of noise in the system output is present, the controller will be switching also fast ('hunting'), possibly causing extensive wear to the control actuators.

To prevent such behaviour, a special 'dead zone' (or 'deadband') is established around the set point where no switching occurs. Obviously, this reduces the accuracy of control. However, the on-off control comes into play when event handling is required. A classic example of application of such control for aircraft is stall recovery. When angle of attack exceeds a specified value, a nose-down elevator command is issued. A similar logic can be implemented for overload protection or ground collision avoidance.

Another important area in the aerospace field where on-off rules can be successfully applied is internal switching between the controllers (or controller parameters). This approach is known as Gain scheduling (Rugh, 2000). The technique takes advantage of a set of relatively simple controllers optimised for different points of the flight envelope (or other conditions). However, it was found that rapid switching may cause stability and robustness problems (Shamma & Athans, 1992). One of the popular simple solutions is to 'blend' (interpolate) the output of two or more controllers, which effectively turns simple on-off switching into a complicated control method.

2.1.2 PID control

The proportional-integral-derivative (PID) control is probably the most widely used type of control, thanks to the simplicity of its formulation and in most cases, predictable characteristics. In a closed-loop system like that in Fig. 1, its control law is

$$u = K_P \varepsilon + K_I \int \varepsilon dt + K_D \dot{\varepsilon}$$

(2)

where the parameters K_P, K_I and K_D are coefficients for proportional, integral and derivative components of the input error signal respectively. By adjusting these three parameters, the desired closed-loop dynamics can be obtained.

Probably the most problematic issue with the PID control is due to the differential term. While being important for good response time and high-speed dynamics, the differential component keenly suffers from both instrumental noise and calculation errors. Indeed, even

a small amount of noise can greatly affect the slope of the input signal. At the same time, numerical calculation of the derivative (for a digital controller in particular) must be done with a fairly small time step to obtain correct slope at a given point. Although lowpass filtering applied to the input signal smoothens the signal, it severely compromises the usefulness of the derivative term because the low-pass filter and derivative control effectively cancel each other out.

In contrast, the integral term averages its input, which tends to eliminate noise. However, a common problem associated with the integral control owes exactly to its 'memory'. When a large input value persists over a significant amount of time, the integral term becomes also large and remains large even after the input diminishes. This causes a significant overshoot to the opposite values and the process continues.

In general, integral control has a negative impact on stability and care must be taken when adjusting the integral coefficient. Limiting the integrator state is a common aid for this problem. A more elegant approach involves 'washing out' the integrator state by incorporating a simple gain feedback, effectively implementing a low-pass filter to the input signal. PID control found wide application in the aerospace field, especially where near-linear behaviour takes place, for example, in various hold and tracking autopilots such as attitude hold and flight path following. The techniques of selecting the optimal PID coefficients are well established and widely accepted. Typically, they use the frequency domain as a design region and utilise phase margins and gain margins to illustrate robustness. However, design of a PID controller for nonlinear or multi-input multi-output (MIMO) systems where significant coupling between the different system inputs and outputs exists is complicated. As the aircraft control objectives evolved, new design and control techniques were developing. Although many of them essentially represent an elaborated version of PID control, they are outlined in the following sections.

2.1.3 Linear optimal control

The concept of optimality in mathematics refers to minimisation of a certain problem dependent cost functional:

$$J = \int_0^T g(t,x,u)dt + h(x(T)) \qquad (3)$$

where T is the final time, u is the control inputs and x is the state of the controlled system. The last term represents the final cost that depends on the state in which the system ends up. Optimal control is, therefore, finding the control law $u(t)$ that minimises J. In general, for an arbitrary system and cost functional, only numerical search can find the optimal (or near optimal) solution.

There is a number of types of linear optimal control, such as the *Linear Quadratic Regulator (LQR)*, the *Linear Quadratic Gaussian with Loop Transfer Recovery (LQG/LTR)* and the extended Kalman filter (EKF). The theory and application of these control techniques and Kalman filtering is detailed in many common control and signal processing textbooks, such as for example (Brown & Hwang, 1992; Bryson & Ho, 1975; Kalman, 1960; Maciejowski, 1989).

The Kalman filter relies on the model of the system (in its predicting part), and the guaranteed performance of the controller can easily be lost when unmodelled dynamics, disturbances and measurement noise are introduced. Indeed, robustness is a challenging issue of modern control designs. Also, the well known in classic linear design phase and

gain margin concepts cannot be easily applied to the multivariable systems that modern control is so suited for. These problems led to the introduction of robust modern control theory.

2.1.4 Robust modern control

Unlike traditional optimal control, robust optimal control minimises the influence of various types of uncertainties in addition to (or even instead of) performance and control energy optimisation. Generally, this implies design of a possibly low gain controller with reduced sensitivity to input changes. As a consequence, robust controllers often tend to be conservative and slow. On the other hand, they may be thought as the stabilising controllers for the whole set of plants (which include a range of uncertainties) and not only for the modelled system, which is a more demanding task.

The majority of modern robust control techniques have origins in the classical frequency domain methods. The key modification of the classic methods is shifting from eigenvalues to singular values (of the transfer function that describes the system), the singular value Bode plot being the major indicator of multivariable feedback system performance (Doyle & Stein, 1981).

Probably the most popular modern robust control design techniques (particularly in the aerospace field) are H_2 and H_∞ control, also known as the frequency-weighted LQG synthesis and the small gain problem respectively. These techniques and the underlying theories are thoroughly described in several works, notably (Doyle at al., 1989; Kwakernaak, 1993; Zhou & Doyle, 1998).

The H_∞ control design found extensive use for aircraft flight control. One of the first such applications was the development of controllers for the longitudinal control of a Harrier jump jet (Hyde, 1991). This work has been subsequently extended in (Postlethwaite & Bates, 1999) to fully integrated longitudinal, lateral and propulsive control. Other works include (Kaminer et al., 1990), where a lateral autopilot for a large civil aircraft is designed, and (Khammash & Zou, 1999), where a robust longitudinal controller subject to aircraft weight and c.g. uncertainty is demonstrated.

A mixed H_2 / H_∞ approach is applied in (Shue & Agarwal, 1999) to design an autoland controller for a large commercial aircraft. The method employed here utilises the H_2 controller for slow trajectory tracking and the H_∞ controller for fast dynamic robustness and disturbance rejection.

Several H_∞ controllers have been tried to accomplish the UAV shipboard launch task (Crump, 2002). It has been found that these controllers perform quite well in nominal situations. However, in the presence of large disturbances which place the aircraft well beyond its linear design operating point, the controllers performed poorly (sometimes extremely). At the same time, inability to include even simple static nonlinearities such as time delays and saturations made it difficult to synthesise a practical controller for this task within linear approach. Another deficiency found is common to all frequency domain techniques: the frequency domain performance specifications cannot be rigidly translated into time and spatial domain specifications. Meanwhile, time and especially spatial (trajectory) constrains are crucial for both launch and recovery tasks.

The time domain performance can be accounted for directly in the time domain l_1 design (Blanchini & Sznaier, 1994; Dahleh & Pearson, 1987).

It is often extended to include the frequency domain objectives, resulting in a mixed norm approach. However, l_1 design is plagued by the excessive order of the generated controllers. This is usually solved by reducing the problem to suboptimal control, imposing several restrictions on the system and performance specifications (Sznaier & Holmes, 1996). The application of l_1 approach to flight control is discussed in (Skogestad & Postlewaite, 1997), with the conclusion that controllers with excessive order will generally be produced when using practical constraints.

There have been attempts to solve the H_∞ optimal control problems for nonlinear systems. However, these methods usually rely on very limiting assumptions about the model, uncertainties and disturbance structure. The mathematical development of nonlinear H_∞ control can be found in (Dalsamo & Egeland, 1995). An example of nonlinear control of an agile missile is given in (Wise, 1996).

Despite a very complicated solution algorithm, this work is limited by a linear assumption on the vehicle aerodynamics, reducing the benefits gained from the use of nonlinear control.

2.1.5 Nonlinear control

Linear control design techniques have been used for flight control problems for many years. One of the reasons why aircraft can be controlled quite well by linear controllers is that they behave almost linearly through most of their flight envelope. However, when the aircraft is required to pass through a highly nonlinear dynamic region or when other complicated control objectives are set, it has been found by several researchers that it is difficult to obtain practical controllers based on linear design techniques. The UAV shipboard launch and recovery tasks are substantially nonlinear problems. The sources of nonlinearities are the aerodynamic forces generated at low airspeeds and high angles of attack (especially when wind disturbances are present); trajectory constraints imposed due to proximity of ground (water) and ship installations; kinematic nonlinearities when active manoeuvring is required; actuator saturations and some more.

In contrast to the linear systems, the characteristics of the nonlinear systems are not simply classified and there are no general methods comparable in power to those of linear analysis. Nonlinear techniques are quite often designed for individual cases, regularly with no mathematical justification and no clear idea of re-applicability of the methods.

Some of the more popular nonlinear control techniques are covered in textbooks (Atherton, 1975; Graham & McRuler, 1971).

The most basic nonlinear control laws and the On-off control and Gain scheduling, noting that these controllers often lack robustness when the controllers are scheduled rapidly.

Another modern control technique remotely related to the on-off control is *variable structure* (also known as *sliding mode*) control (DeCarlo et al., 1988; Utkin, 1978). In this approach, a hypersurface (in state space) called *sliding surface* or *switching surface* is selected so that the system trajectory exhibits desirable behaviour when confined to this hypersurface. Depending on whether the current state is above or below the sliding surface, a different control gain is applied. Unlike gain scheduling, the method involves high speed switching to keep the system on the sliding surface.

A completely different approach is to enable applicability of the well known linear control methods to control nonlinear systems. This can be achieved using *nonlinear dynamic inversion*. This process, also known as *feedback linearisation*, involves online approximate linearisation of a nonlinear plant via feedback. Dynamic inversion gained particular

attention in aviation industry in the late 1980s and 90s, aiming to control high performance fighters during high angle of attack manoeuvres (known as *supermanoeuvres*). One of the early applications is NASA High Angle of Attack Vehicle (HARV) (Bugajski & Enns, 1992). In this work, quite good simulation performance results are obtained; however, with the inversion based on the same simulation model, any possible discrepancies are transferred into the controller, leading to questionable results in physical implementation with respect to incorrectly modelled dynamics.

2.1.6 Intelligent control

Intelligent control is a general and somewhat bold term that describes a diverse collection of relatively novel and non-traditional control techniques based on the so called *soft computing* approach. These include neural networks, fuzzy logic, adaptive control, genetic algorithms and several others. Often they are combined with each other as well as with more traditional methods; for example, fuzzy logic controller parameters being optimised using genetic algorithms or a neural network driving a traditional linear controller.

Neural network (NN), very basically, is a network of simple nonlinear processing elements (neurons) which can exhibit complex global behaviour determined by element parameters and the connections between the processing elements. The use of artificial neural networks for control problems receives an increased attention over the last two decades. It has been shown that a certain class of NN can approximate any continuous nonlinear function with any desired accuracy (Spooner, 2002). This property allows to employ NN for *system identification* purposes, which can be performed both offline and online. This approach is used in (Coley, 1998; Kim & Calise, 1997) for flight control of a tilt-rotor aircraft and the F-18 fighter.

Another area of intensive application of NN is fault tolerant control, made possible due to online adaptation capability of NN. In the recent work (Pashilkar et al., 2006), an 'add-on' neural controller is developed to increase auto-landing capabilities of a damaged aircraft with one of two stuck control surfaces. A significant increase in successful landings rate is shown, especially when the control surfaces are stuck at large deflections.

Fuzzy logic control also gained some popularity among flight control engineers. This type of control relies on approximate reasoning based on a set of rules where intermediate positions between 'false' and 'truth' are possible. Fuzzy logic control may be especially useful when a decision must be made between several controversial conditions. For example, in (Fernandez-Montesinos, 1999) fuzzy logic is used for windshear recovery in order to decide whether energy should be given to altitude or airspeed, based upon the current situation.

Adaptive control term covers a set of various control techniques that are capable of online adaptation. A good survey of adaptive control methods is given in (Astrom & Wittenmark, 1995). The applications of adaptive control is generally biased towards control for large time scales so that the controller has sufficient time to learn how to behave. This makes the relatively short-time recovery process unsuitable for online adaptation.

Evolutionary and *genetic algorithms* (EAs, GAs) are global optimisation techniques applicable to a broad area of engineering problems. They can be used to optimise the parameters of various control systems, from simple PID controllers (Zein-Sabatto & Zheng, 1997) to fuzzy logic and neural network driven controllers (Bourmistrova, 2001; Kaise & Fujimoto, 1999). Another common design approach is evolutionary optimisation of trajectories, accompanied by a suitable tracking controller (e.g. (Wang & Zalzala, 1996)). An elaborated study of applications of EAs to control and system identification problems can be found in (Uzrem, 2003).

Unlike the majority of other techniques, Genetic Algorithms (in the form of *Genetic Programming*) are able to evolve not only the parameters, but also the *structure* of the controller.

In general, EAs require substantial computational power and thus are more suitable for offline optimisation. However, online evolutionary-based controllers have also been successfully designed and used. The *model predictive control* is typically employed for this purpose, where the controller constantly evolves (or refines) control laws using an integrated simulation model of the controlled system. A comprehensive description of this approach is given in (Onnen et al., 1997).

2.2 Flight control for the UAV recovery task

Aircraft control at recovery stage of flight can be conventionally separated into two closely related, but distinctive tasks: guidance and flight control. Guidance is the high-level ('outer loop') control intended to accomplish a defined mission. This may be path following, target tracking or various navigation tasks. Flight control is aimed at providing the most suitable conditions for guidance by maintaining a range of flight parameters at their optimal levels and delivering the best possible handling characteristics.

2.2.1 Traditional landing

In a typical landing procedure, the aircraft follows a defined glide path. The current position of the aircraft with respect to the glidepath is measured in a variety of ways, ranging from pilot's eyesight to automatic radio equipment such as the Instrument Landing System (ILS). Basically, the objective of the pilot (or autopilot) is to keep the aircraft on the glidepath, removing any positioning error caused by disturbances and aircraft's dynamics. This stage of landing is known as *approach*. Approach may be divided into *initial approach*, in which the aircraft makes contact with ('fixes') the approach navigation system (or just makes visual contact with the runway) and aligns with the runway; and *final approach*, when the aircraft descends along a (usually) straight line. In conventional landing, final approach is followed by a flare manoeuvre or nose-up input to soften the touchdown; however, flare is typically not performed for shipboard landing due to random ship motion and severe constraints on the landing space.

The guidance task during final approach involves trajectory tracking with both horizontal and vertical errors (and their rates) readily available. It is important to note that the errors being physically measured are *angular* deviations of the aircraft position (Fig.2) as seen from the designated touchdown point (or more precisely, from where the radars or antennas or other guidance systems are located). They can be converted to linear errors Δh if the *distance L* to the aircraft is known.


Figure 2. Angular and linear glidepath errors

However, in many applications precise distance measurement is unavailable. Nevertheless, successful landing can be carried out even without continuous distance information. This comes from the fact that exactly the angular errors are relevant to precise landing. Indeed, the goal is to bring the aircraft to a specified point on the runway (or on the deck) in a certain state. The choice of the landing trajectory only serves this purpose, considering also possible limitations and secondary objectives such as avoiding terrain obstacles, minimising noise level and fuel consumption and so on. The angular errors provide an adequate measurement of the current aircraft position with respect to the glidepath.

However, if the landing guidance system takes no account of distance and is built around the angular error only, it may cause stability problems at close distances, because the increasing sensitivity of the angular errors to any linear displacement effectively amplifies the system gain.

2.2.2 UAV control for shipboard recovery

As it was seen from the discussion in the previous section, landing of an aircraft is a well established procedure which involves following a predefined flight path. More often than not, this is a rectilinear trajectory on which the aircraft can be stabilised, and the control interventions are needed only to compensate disturbances and other sources of errors.

The position errors with respect to the ideal glidepath can be measured relatively easily. Shipboard landing on air carriers is principally similar; the main differences are much tighter error tolerances and absence of flare manoeuvres before touchdown. The periodic ship motion does have an effect on touchdown; however, it does not affect significantly the glidepath, which is projected assuming the average deck position. The choice of the landing deck size, glideslope, aircraft sink rate and other parameters is made to account for any actual deck position at the moment of touchdown. For example, a steeper glideslope (typically 4°) is used to provide a safe altitude clearance at the deck ramp for its worst possible position (i.e. ship pitched nose down and heaved up). This makes unnecessary to correct the ideal reference trajectory on the fly.

For the UAV shipboard recovery, ship oscillations in high sea cause periodic displacement of the recovery window (the area where capture can be done) several times greater than the size of the window itself. This fact (and also the assumption that the final recovery window position cannot be predicted for a sufficient time ahead) makes it impossible to project an optimal flight path when the final approach starts. Instead, the UAV must constantly track the actual position of the window and approach so that the final miss is minimised. Therefore, it turns out that the UAV recovery problem resembles that of homing guidance rather than typical landing. While stabilisation on a known steady flight path can be done relatively easy with a PID controller, homing guidance to a moving target often requires a more sophisticated control.

Not surprisingly, homing guidance found particularly wide application in ballistic missiles development, hence the accepted terminology owes to this engineering area. In the context of UAV recovery, the UAV moves almost straight towards the 'target' from the beginning (compared to typical missile intercept scenarios) and thus the velocity vector and the line of sight almost coincide.

However, there are two major difficulties that can compromise the effectiveness of Proportional Navigation (PN) for UAV recovery. First, PN laws are known as generating excessive acceleration demands near the target. For a UAV with limited manoeuvrability,

such demands may be prohibitive, especially at low approach airspeeds. On the other hand, the PN guidance strategy does not change during the flight. Several alternative guidance strategies with more favourable acceleration demands exist, e.g. augmented proportional navigation. However, it is unlikely they can sufficiently improve the guidance to an oscillating target such as ship's deck.

2.3 UAV controller structure

The objective is therefore to synthesise such guidance strategy that enables reliable UAV recovery, and to produce a controller that implements the target tracking guidance strategy. The evolutionary design (ED) method applied for this task allows to evolve automatically both the structure and the parameters of the control laws, thus potentially enabling to generate a 'full' controller, which links available measurements directly with the aircraft control inputs (throttle, ailerons, rudder and elevator) and implements both the guidance strategy and flight control (Fig. 3):


Figure 3. Full controller with embedded guidance strategy

However, this approach, even though appealing at first and requiring minimum initial knowledge, proves to be impractical as the computational demands of the evolutionary algorithms (EAs) soar exponentially with the dimensionality of the problem. It is therefore desirable to reduce complexity of the problem by reducing the number of inputs/outputs and limiting, if appropriate, possible structures of the controllers.

Another difficulty is the evaluation of the controller's performance. In ED, performance or fitness is multiobjective. Therefore, it is highly desirable to decompose this complex task into several simpler problems and to solve them separately.

A natural way of such decomposition is separating the trajectory control (guidance) and flight control. The guidance controller issues commands u_g to the flight controller, which executes these commands by manipulating the control surfaces of the UAV (Fig. 4). These two controllers can be synthesised separately using appropriate fitness evaluation for each case.


Figure 4. UAV recovery control diagram

2.3.1 Guidance controller

The internal structure of the controller is defined by the automatic evolutionary design based on predefined set of inputs and outputs. It is desirable to keep the number of inputs and outputs to minimum, but without considerably compromising potential performance.

An additional requirement to the outputs u_g is that these signals should be straightforwardly executable by the flight controller. This means that u_g should represent a group of measurable flight parameters such as body accelerations, velocities and Euler angles, which the flight controller can easily track.

The structure of the output part of the guidance controller is as shown in Fig. 5. With this scheme, the guidance laws produce general requests to change trajectory in horizontal and vertical planes. The kinematic converter then recalculates these requests to the form convenient for the flight controller. Both the bank angle and normal body load factor n_y can be relatively easily tracked, with the sensors providing direct measurements of their actual values. At the same time, this approach allows to evolve the horizontal and vertical guidance laws separately, which may be desirable due to different dynamics of the UAV's longitudinal and lateral motion and also due to computational limitations.

Input measurements to the guidance controller should be those relevant to trajectory. First of all, this is all available positioning information, pitch and yaw angles and airspeed. They do not account for steady wind, but still provide substantial information regarding the current 'shape' of trajectory.

The yaw angle fed into the controllers is corrected by the 'reference' yaw ψ_0, which is perpendicular to the arresting wire in the direction of anticipated approach. A zero yaw indicates that the UAV is pointed perpendicularly to the arresting wire (ignoring ship oscillations), which is the ideal condition in the absence of side wind and when the UAV moves along the ideal glidepath. This is similar to rotating the ground reference frame $O_g x_g y_g z_g$ by the correction yaw angle ψ_0. The rotated frame is referred as *approach ground reference frame*.


Figure 5. Guidance controller

The derived quantities from raw measurements are the vertical and lateral velocity components with respect to the approach ground reference frame.

Determination of the current UAV position and velocities with respect to the arresting wire is crucial for successful recovery. While previously discussed flight parameters may only help to improve the guidance quality, positioning carries direct responsibility for recovery.


Figure 6. Positioning scheme

The system is based on radio distance metering and provides ten independent raw measurements (Fig. 6): three distances d_1, d_2 and d_3 from the UAV to the radio transmitters located at both ends of the recovery boom which supports the arresting wire and at the base of recovery mast; three rates of change of these distances; distance differences ($d_1 - d_2$) and ($d_3 - d_2$); and rates of change of the differences.

The guidance laws evolution process is potentially capable to produce the laws directly from raw measurements, automatically finding necessary relationships between the provided data and the required output.

The target spot elevation is chosen to be a constant, and the value is determined in view of the expected cable sag obtained from simulation of the cable model. For normal approach speed, the cable sag varies between 3.5 and 4.5 m. Accordingly, the target spot elevation is chosen to be $h_T = 2$ m above the arresting wire.

The recovery procedure, if successful, lasts until the cable hook captures the arresting wire. This happens when the UAV have moved about the full length of the cable *past* the recovery boom. However, position measurements may be unavailable beyond the boom threshold, and even shortly before the crossing the readings may become unreliable. For these reasons, the terminal phase of approach, from the distance about 6–10 m until the capture (or detection of a miss), should be handled separately.

It is possible to disconnect the guidance controller several metres before the recovery boom without affecting the quality of guidance. The allowed error (approximately 2 m in all directions, determined by the lengths of the arresting wire and the cable) should absorb the absence of controlled guidance in the last 0.3 to 1 second (depending on the headwind) in most situations.

2.3.2 Flight controller

Flight controller receives two inputs from the guidance controller: bank angle demand γ^d and normal body load factor demand n_y^d. It should track these inputs as precisely as possible by manipulating four aircraft controls: throttle, ailerons, rudder and elevator.

The available measurements from the onboard sensors are body angular rates ω_x, ω_y, ω_z from rate gyros, Euler angles γ, ψ, θ from strapdown INS, body accelerations n_x, n_y, n_z from the respective accelerometers, airspeed V_a, aerial angles a and β, actual deflection of the control surfaces δ_a, δ_r, δ_e, and engine rotation speed N_{rpm}.

For simplicity of design, the controller is subdivided into independent longitudinal and lateral components. In longitudinal branch, elevator tracks the input signal n_y^d, while throttle is responsible for maintaining a required airspeed. In lateral control, naturally, ailerons track γ^d, while rudder minimises sideforce by keeping n_z near zero.

3. Evolutionary design

The Evolutionary Design (ED) presented in this section, generally, takes no assumptions regarding the system and thus can be used for wide variety of problems, including nonlinear systems with unknown structure. In many parts, this is a novel technique, and the application of ED to the UAV guidance and control problems demonstrates the potential of this design method.

The core of evolutionary design is a specially tailored evolutionary algorithm (EA) which evolves both the structure and parameters of the control laws.

Since the algorithm is used for *creative* work only at the *design* stage, its performance is rather of secondary importance as long as the calculations take a sensible amount of time. The major requirements to automatic design methods are *quality* of the result and *exploration* abilities.

Although the basic framework of an EA is quite simple, there are three key elements that must be prepared before the algorithm can work. They are:

- representation of phenotype (control laws in our case) suitable for genetic operations (genome encoding);
- simulation environment, which enables to implement the control laws within the closed loop system;
- fitness evaluation function, which assesses the performance of given control laws.

These elements, as well as the whole algorithm outline, are addressed below.

Parallel evolution of both the structure and the parameters of a controller can be implemented in a variety of ways. One of the few successfully employed variants is the *block structure controller evolution* (Koza et al., 2000).

In this work the ED algorithm enables to evolve suitable control laws within a reasonable time by utilising gradual evolution with the principle of strong casualty. This means that structure alterations are performed so that the information gained so far in the structure of the control law is preserved. Addition of a new block, though being random, does not cause disruption to the structure. Instead, it adds a new dimension and new potential which may evolve later during numerical optimisation. The principle of strong casualty is often regarded as an important property for the success of continuous evolution (Sendhoff et al., 1997).

The addition of new points or blocks is carried out as a separate dedicated operation (unlike sporadic structure alterations in the sub-tree crossover), and is termed *structure mutation*. Furthermore, in this work structure mutation is performed in a way known as *neutral structure mutation*. That's when the new block should be placed initially with zero coefficient. This will not produce any immediate improvement and may even deteriorate the result slightly because more blocks are used for the same approximation. However, further numerical optimisation should fairly quickly arrive at a better solution. The usefulness of neutral mutations has been demonstrated for the evolution of digital circuits (Van Laarhoven & Aarts, 1987) and aerodynamic shapes (Olhofer et al., 2001).

As a result, the ED algorithm basically represents a numerical EA with the inclusion of structure mutations mechanism.

3.1 Representation of the control laws

Control laws are represented as a combination of static functions and input signals, which are organised as a dynamic structure of *state equations* and *output equations* in form of continuous representation.

The controller being evolved has m inputs, r outputs and n states. The number of inputs and outputs is fixed. The algorithm allows varying number of states; however, in this work, the number of states is also fixed during the evolution. As a result, the controller comprises of n state equations and r output equations:

$$\begin{aligned}
\dot{x}_1 &= g_1(x,u) \\
\dot{x}_2 &= g_2(x,u) \\
&\vdots \\
\dot{x}_n &= g_n(x,u)
\end{aligned}
\quad \text{and} \quad
\begin{aligned}
y_1 &= f_1(x,u) \\
y_2 &= f_2(x,u) \\
&\vdots \\
y_n &= f_n(x,u)
\end{aligned}
\tag{4}$$

where u is size m vector of input signals, $x = [x_1, x_2, \ldots x_n]$ is size n vector of state variables, $y_{1\ldots r}$ are controller outputs. Initial value of all state variables is zero. All $n+r$ equations are built on the same principle and are evolved simultaneously. For structure mutations, a random equation is selected from this pool and mutated.

3.1.1 Representation of input signals

Input signals delivered to each particular controller are directly measured signals as well as the quantities derived from them.

Within each group, inputs are organised in the subgroups of 'compatible' parameters. Compatible parameters are those which have close relationship with each other, have the same dimensions and similarly scaled. The examples of compatible parameters are the pairs (n_x, n_{xg}), $(\omega_y, \dot{\psi})$, (V_a, V_{CL}). As a rule, only one of the compatible parameters is needed in a given control law. For this reason, the probability of selection of such parameters for the structure mutation is reduced by grouping them in the subgroups. Each subgroup receives equal chances to be selected. If the selected subgroup consists of more than one input, a single input is then selected with uniform probability.

Therefore, every controller input may be represented by a unique *code* consisting of three indices: the number of group a, the number of subgroup b and the number of item in the subgroup c. The code is designated as u(a,b,c).

3.1.2 Representation of control equations and the structure mutation

Each of the control equations (4) is encoded as described above. To this end, only one single output equation of the form $y = f(u)$ will be considered in this section. State variables x are considered as special inputs and have no effect on the encoding.

This is done to speed up the search, which could otherwise be hampered by the multitude of possible operations. It proved to be more effective to include all meaningful quantities derived from source measurements as independent inputs than to implement all the functions and operations which potentially allow to emerge all necessary quantities automatically in the course of evolution.

The encoding should allow a simple way to insert a new parameter in any place of the equation without disrupting its validity and in a way that this insertion initially does not affect the result, thus allowing neutral structure mutations.

Conceptually, the equation is a sum of input signals, in which:

- every input is multiplied by a numeric coefficient or another similarly constructed expression;
- the product of the input and its coefficient (whether numeric or expression) is raised to the power assigned to the input;
- a free (absolute) term is present.

The simplest possible expression is a constant:

$$y = k_0 \tag{5}$$

A linear combination of inputs plus a free term is also a valid expression:

$$y = k_2 u_2 + k_1 u_1 + k_0 \tag{6}$$

Any numeric constant can be replaced with another expression. An example of a full featured equation is

$$y = ((k_4 u_4 + k_3)u_3) - 0.5 + k_2 u_2 + (k_1 u_1)^2 + k_0 \tag{7}$$

This algorithm can be illustrated by encoding the example (7). The respective internal representation of this expression is:

- Equation: $y = ((k_4 u_4 + k_3)u_3) - 0.5 + k_2 u_2 + (k_1 u_1)^2 + k_0$
- Expression: { u(3,-0.5) u(4) 1 2 *u*(2) 3 u(1,2) 4 5 }
- Object parameters: [k_4 k_3 k_2 k_1 k_0]
- Strategy parameters: [s_4 s_3 s_2 s_1 s_0]

This syntax somewhat resembles Polish notation with implicit '+' and '*' operators before each variable. The representation ensures presence of a free term in any sub-expression, such as k_3 and k_0 in the example above.

The algorithm of structure mutation is presented below.

1. Select an input (or a state variable) at random: u(a,b,c).
2. Obtain the initial values of the numeric coefficient and the strategy parameter (initial step size). The initial coefficient *k* is selected as described above, it can be either 0 or 10^6.
3. Append the object parameters vector with the initial coefficient, and the strategy parameters vector with the initial step size.
4. Form a sub-expression consisting of the selected variable code and the obtained index: { u(a,b,c) *n* }.
5. With 40% probability, set the insertion point (locus) to 1; otherwise, select a numeric value in the expression at random with equal probability among all numeric values present and set the locus to the index of this value.
6. Insert the sub-expression into the original expression at the chosen locus (*before* the item pointed).

This procedure may produce redundant expressions when the selected variable already exists at the same level. Thus an algebraic simplification procedure is implemented. It parses given expression, recursively collects the factors of each variable encountered and then rebuilds the expression.

3.2 Simulation environment

Fitness evaluation of the controllers is based on the outcome of one or more simulation runs of the models involved with the controller being assessed in the loop.

Simulation environment for this study is constructed in the MATLAB/Simulink software. All the models are implemented as Simulink library blocks (Khantsis, 2006). The models participating in the evolution include:

- The *Ariel* UAV model;
- The atmospheric model;
- The ship model;
- The static cable model.

The sample rate should be kept at minimum, ensuring, however, numerical stability. The fastest and thus the most demanding dynamics in the model is contained in

- sensors (both the lags and noise);
- control actuators;
- turbulence model;
- and potentially, controllers.

These components should be especially carefully examined when adjusting the sample rate. The experimentally determined minimum sample rate for the model is 100 Hz. The model is integrated using the 4th order Runge-Kutta method.

3.3 Fitness evaluation

Fitness evaluation is based on objectives which are individual for each controller and therefore is calculated individually. As discussed above having several specific controllers makes it easier to define their particular objectives and reduces the number of these objectives as compared to one all-purpose autopilot.

Fitness evaluation of a controller can be divided into two main stages. First is the preparation of the sample task and simulation of the model with the controller in the loop. The second stage is analysis of the results obtained from the simulation and evaluation of the fitness as such. These steps are often repeated several times, making the controller perform varying tasks in order to obtain a balanced measure of its performance via multi-task fitness evaluation.

Fitness value may reflect different levels of performance achieved by the controller, including 'hard bounds.' For example, if the controller crashes the UAV, a high penalty score may be assigned.

Other parameters of flight taken into account is control usage (or control effort): it is desirable to keep control usage at minimum. Tracking abilities, if applicable, can be estimated by the weighted sum of the tracking error.

The total fitness value is calculated as the weighted sum of all estimates:

$$F = W_c C_c + W_f C_f + W_e C_e + \ldots \tag{8}$$

The exact value of the weighting coefficients W, as well as the number of estimates taken into account, is individual for each controller. As a rule, the weighting coefficients are chosen empirically. It should be noted that the signals are dimensionalised and thus the coefficients may vary significantly. For example, aerodynamic surfaces deflection is measured in degrees and may range from –16 to 16 (for rudder and ailerons); considering 10-fold saturation limits, the peak values may reach ±160. At the same time, throttle range is 0.01 to 1, thus the cost for its usage will be about 30 times as low (for similar control dynamics).

3.4 Evolutionary design algorithm outline

The Evolutionary Design algorithm is implemented as a function with two input arguments: population size and number of generations to run. Other, more specific parameters are initialised within the program. The algorithm framework is as follows.

1. Initialise all parameters. If the evolution is continued, go to step 3.
2. Create initial population.
3. Evaluate fitness of each individual.
4. Save the full state to a temporary file for emergency recovery.
5. If the required number of generations is reached, return.

6. Sort the population according to the fitness of each individual.
7. If elitism is enabled, copy the best scored individual into the new population.
8. Until the new population is filled up:
9. Select the next individual from the sorted list (starting from the elite member).
10. Every (ks)th generation, with probability Ps, perform structure mutation of the selected individual.
11. Reproduce n offspring individuals from the selected (and possibly mutated) individual.
12. Put these n new individuals to the new population.
13. Continue the loop from step 8.
14. Increase the generation counter.
15. Continue from step 3.

Several steps of this algorithm need further clarification.

Initial population is created according to the specified task. By default, it is initialised with the control laws of the form $y = const$ with randomly chosen constants. Most of the controllers, however, are initialised with more meaningful control laws. For example, tracking controllers may be initialised with the laws $y = k_1\varepsilon + k_0$, where ε is the respective error signal and the coefficients k are sampled at random (taking into account default step sizes for the respective signal).

It can be seen that *selection* is performed deterministically. This is the most commonly used way in ES. The populations used in this study were of moderate size, usually 24 to 49 members. Selection pressure is determined by the number of the offspring n of each individual. The smaller n, the lower the selection pressure. For nearly all runs in this work $n = 2$, which means that half of the population is selected. This is a rather mild level of selection pressure.

The parameters determining *structure mutation* occurrence, k_s and P_s, both change during the evolution. As discussed previously, the number of structure mutations should decrease as the complexity of the controllers grows and more time to optimise the coefficients is required. The probability of structure mutation P_s is normally high in the beginning (0.7 to 1.0) and then decrease exponentially to moderate levels (0.4 to 0.6) with the exponent 0.97 to the generation number. Default value of k_s is set according to overall complexity of the controller being evolved. For simpler single-output controllers, as a rule, $k_s = 10$; for more complex controllers $k_s = 20$. However, in the beginning of evolution, k_s is reduced by half until the generation 20 and 100 respectively.

Reproduction is performed simultaneously with mutation, as it is typically done in ES, with the exception that this operation is performed separately for each selected member.

4. Controller synthesis and testing

The UAV control system is synthesised in several steps. First, the flight controller is produced. This requires several stages, since the flight controller is designed separately for longitudinal and lateral channels. When the flight controller is obtained, the guidance control laws are evolved.

Application of the ED algorithm to control laws evolution is fairly straightforward: 1) preparation of the sample task for the controller, 2) execution of the simulation model for the given sample task and 3) analysis of the obtained performance and evaluation of the fitness value. When both the model and fitness evaluation are prepared, the final evolution may be started. Typically, the algorithm is run for 100–200 generations (depending on

complexity of the controller being evolved). The convergence and the resulting design is then analysed and the evolution, if necessary, is continued.

4.1 Step 1: PID autothrottle

Initially a simple PID variant of autothrottle is evolved - to ensure a more or less accurate airspeed hold. At the next stage, its evolution is continued in a full form together with the elevator control law. The PID structure of the controller may be ensured by appropriate initialisation of the initial population and by disabling structure mutations. Therefore, the algorithm works as a numerical optimisation procedure. The structure of the autothrottle control law is following:

$$\begin{aligned}
\dot{x}_1 &= k_1 x_1 + k_2 \Delta V_a \\
\delta_t &= k_3 x_1 + k_4 \Delta V_a + k_5 \dot{V}_a + k_6
\end{aligned}$$

(9)

where $\Delta V_a = V^d - V_a$ is the airspeed error signal, δ_t is the throttle position command and $k_{1\ldots6}$ are the coefficients to be optimised.

A 30-second flight is allocated for performance measurement. Such a long flight is needed because throttle response is quite slow. Since robustness of the PID controller to discrepancies in the UAV model is not topical at this stage (it will be addressed in further evolution), and because structure of the controller is fixed so that irrelevant measurements cannot be attracted, only a single simulation run for each fitness evaluation is performed.

Elevator inputs provide the main source of disturbances for training the autothrottle. A 2.5-degree nose-up step input is executed at time $t = 5$ s. It produces nearly steady climb without reaching stall angles of attack and saturation on the throttle (except for, possibly, dynamic transition moments).

At $t = 14$ s, a similar nose-down elevator step is commanded, entering the UAV into glide with small descent angle. In addition, a 3 m/s tailwind gust is imposed at $t = 23$ s. Generally, manoeuvres and wind disturbances are the most prominent sources of airspeed variations, therefore they are included for autothrottle assessment.

Initial airspeed is set to 23 m/s, i.e. 1 m/s above the required airspeed. This provides an additional small scale disturbance. Initial altitude is 100 m, providing enough elevation to avoid crash for any sensible throttle control.

Algorithm settings are as follows. Since the structure is fixed, large population size N is unnecessary. Population size of 25 members has been used in most runs. $N = 13$ showed similar results in terms of fitness evaluation demands (red line on the convergence graph in Fig. 7a; adjusted to the population size 25 for comparison). Fitness is evaluated deterministically because all random signals are repeatable from run to run. Elitism is enabled. Fitness is calculated with the following weighting coefficients:

$$F_t = 2000 C_e(V_e) + 1000 C_c(\delta_t) + 2000 C_f(\delta_t)$$

(10)

It should be noted that the *commanded* signal δ_t (which enters the actuator) is used for evaluation.

Convergence graphs for three independent runs of the ED algorithm are presented in Fig. 7a. They show the best fitness values for each generation. On average, 2000 to 2500

simulation runs is needed to achieve satisfactory performance. Because this is only an intermediate controller, full optimisation is not necessary at this stage.
The best controller obtained in the experiments is the following:

$$\dot{x}_1 = -2.9966x_1 + 85.57\Delta V_a$$
$$\delta_t = 0.012616x_1 + 0.089324\Delta V_a - 0.044004\dot{V}_a + 0.23827$$

(11)

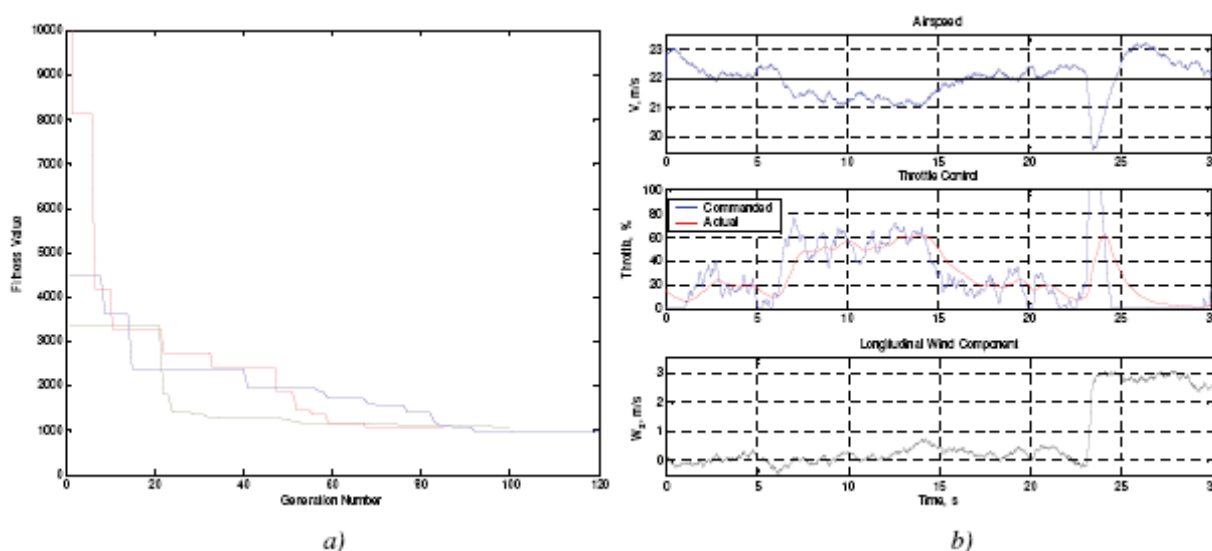Sample response of this controller is illustrated in Fig. 7b.

a)   b)

Figure 7. PID autothrottle current-best fitness (a) and sample response (b)

4.2 Step 2: Longitudinal control

With a simple autothrottle available, elevator control can be developed to provide tracking of the normal body load factor demand n_y^d .

It is desirable that the sample training task closely reflect the conditions the controller will be subject to in a real guidance task. Unfortunately, this is not possible at the early stages of multistage design process as guidance laws can be evolved only after a capable flight controller has been produced. In the meantime, a replacement pseudoguidance task must be used for the flight controller training.
As a suitable replacement of the guidance law, a simple static altitude hold is used:

$$n_y^d = 0.03\left(H^d - H\right) + \cos\theta$$

(12)

A low gain is chosen to ensure guaranteed stability. The $\cos\theta$ term takes into account gravity component, assuming that roll angles are small. However, real guidance in high sea will most likely require a more active control, especially closer to the recovery boom. For this reason, a synthetic 'chirp' signal $c(t)$ is added starting from 10th second of flight. This signal represents a sine wave which frequency increases linearly from $f_0 = 0.01$ Hz at $t_0 = 10$ s to $f_1 = 0.5$ Hz at $t_1 = 20$ s. For $t < 10$ s, $c(t) = 0$, providing 'training' for smooth control. Chirp signal is often used for response analysis of nonlinear systems. Altogether, the reference input signal is

$$n_y^d(t) = 0.03\big(H^d - H(t)\big) + \cos\theta(t) + 0.2c(t) \tag{13}$$

The preset altitude H^d is 5 m above the initial altitude H_0 = 100 m. This causes a small positive step input in the beginning of flight.

Flight time is limited to 20 s, which is close to normal final approach flight time (10–15 s), but allocates more time for better 'training' of slow throttle control. Overall fitness of the controller is calculated as an average of three fitness values obtained for each flight.

Elevator control law is initialised as follows:

$$\dot{x}_2 = 0$$
$$\delta_e = k_1 x_2 + k_2 \Delta n_y + k_3 \tag{14}$$

where $\Delta n_y = n_y^d - n_y$ is the load factor error and the coefficients $k_{1\ldots3}$ are sampled at random for the initial population. For simplicity, first-order dynamic control law is used. However, when combined with the throttle control, both control laws may include both state variables x_1 and x_2, forming a tightly coupled longitudinal control.


Figure 8. Longitudinal control sample response

Fitness estimation combines both throttle control fitness and elevator control fitness. They are constructed in a similar manner with the following weighting coefficients:

$$F_{te} = F_t + F_e = F_t + 500000C_e\big(n_y\big) + 50C_c\big(\delta_e\big) + 5000C_f\big(\delta_e\big) \tag{15}$$

The weights put slightly more emphasis on elevator control than on throttle control, because the former is more important for guidance. Autothrottle only supplements it to provide safe flight conditions.

Relatively large population size of 49 members has been used in all runs. Initial structure mutation probability P_s was set to 1.0; final value $P_s \rightarrow 0.6$. Structure mutations happened every (k_s = 5)th generation for the first 20 generations, then k_s = 10 until 100th generation, then k_s = 20. Satisfactory controllers were obtained after about 100 to 120 generations, which requires approximately 15000 simulation runs. The best performing controller has the following form:

$$\dot{x}_1 = \left(\left(-252.611\omega_z - 172.0765\right)\alpha - 4.411\right)x_1 + 75.0613\Delta V_a$$
$$\dot{x}_2 = -0.77101x_2 - 228.8894\Delta n_y$$
$$\delta_t = 0.0011588x_1 + 1.0565\Delta V_a - 0.6989\dot{V}_a + 0.24651 \tag{16}$$
$$\delta_e = 41.5594\theta + \left(-140.8115\alpha\right)^{0.5} + 0.89016x_2 - 26.1045\Delta n_y + 50.1835\omega_z + 3.871$$

Response of this controller to the sample input (in default configuration) is presented in Fig. 8.

4.3 Step 3: Lateral control

Lateral control consists of two channels: ailerons control and rudder control. As a rule, for an aerodynamically stable aircraft such as the *Ariel* UAV, lateral control is fairly simple and is not as vital for flight as longitudinal control. For this reason, both control laws, for ailerons and rudder, are evolved simultaneously in one step.

The set-up is largely similar to that used in the previous step. The just evolved longitudinal control laws (16) are connected to throttle and elevator. Both ailerons and rudder control laws are initialised in a similar manner to (14):

$$\dot{x}_3 = 0$$
$$\dot{x}_4 = 0$$
$$\delta_a = k_{11}x_3 + k_{12}\Delta\gamma + k_{13} \tag{17}$$
$$\delta_r = k_{21}x_4 + k_{22}n_z + k_{23}$$

Fitness is calculated as follows:

$$F_{ar} = F_a + F_r = 50000C_e(\gamma) + C_c(\delta_a) + 2000(\delta_a) + \dots 200000C_e(n_z) + 10C_c(\delta_r) + 2000C_f(\delta_r) \tag{18}$$

Algorithm settings are the same as for longitudinal control design. Fairly quick convergence (within 70–80 generations) was observed, with slow further progress. Considering that the initial state is zero, state variable x_3 is always zero. The controller is obtained as follows:

$$\dot{x}_4 = -39.2775x_4 - 44.57\omega_x + \left(\left(4.6845x_4\right)^{0.5} + 69.3478\right)n_z$$
$$\delta_a = -26.7812\Delta\gamma + 8.1188\omega_x + 0.57032 \tag{19}$$
$$\delta_r = -1.6916x_4 - 1.33n_z + 4.1764\omega_y + 21.2955\beta - 0.65184$$

Sample response of this controller is presented in Fig. 9.

Figure 9. Lateral control sample response

4.4 Step 4: Stall prevention mechanism

Active guidance control may produce high acceleration demands n_y^d, which may lead to exceeding the maximum allowed angle of attack $\alpha_{\max}$ and result in stall.

However, fitness evaluation of the control law with stall prevention would be more complex and time consuming. Separate estimation is required because the objectives are contradictory: adequate limitation of α invariably results in tracking degradation and vice versa.

Though, a separate control law which limits angle of attack is developed. The limiter should work only when dangerous angle of attack is reached, being deactivated in normal conditions.

The control law is evolved in a similar manner to other laws described above. The objective of the control law is to maintain $\alpha = \alpha_{\max}$ by correcting n_y^d as long as this demand is greater than required for safe flight. Also, free term is disregarded in both state and output equations.

Fitness is calculated similarly to all tracking controllers with respect to the command $\alpha = \alpha_{\max}$. Two exceptions are made, however. First, overshoot is penalised twice as heavily as undershoot, because exceeding the maximum angle of attack is considered to be hazardous. Second, since the limiter has no direct control on elevator, the amount of longitudinal oscillations is taken into account as the variance of pitch rate. Altogether, the fitness value is calculated as

$$F_a = 10^6 C_a + 0.1 C_c(\delta_e) + 1000 C_c(\omega_z) + 1000 C_f(\delta_e) \tag{20}$$
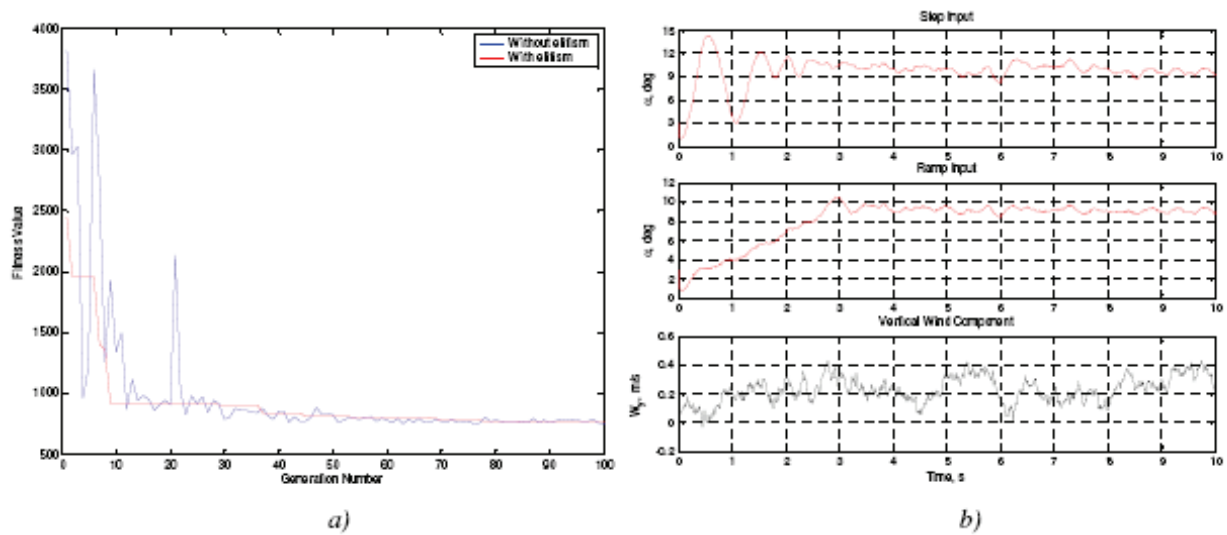

Figure 10. Angle of attack limiter current-best fitness (a) and sample responses (b)

Algorithm settings and initialisations are similar to the case of longitudinal control evolution, except that smaller population of 24 or 25 members is used because only one control law is evolved. Both elitist and non-elitist strategies have been tried with largely similar results. Quick convergence, within 40–50 generations, was observed in all cases, with slow further progress. Examples of current-best convergence are presented in Fig. 10a. A simpler solution obtained in elitist strategy has been selected:

$$\dot{x}_5 = (93.01\omega_z - 15.0798)(\alpha - \alpha_{\max})$$
$$n_y^{d.corr} = n_y^d - \Delta n_y^d = n_y^{d0} - 0.63208 x_5 - 25.246(\alpha - \alpha_{\max}) \tag{21}$$
$$K_{\max} = 60$$

Where $K_{\max}$ is a damping coefficient and n_y^{d0} is a frozen value for n_y^d. Sample responses of this controller are presented in Fig. 10.

4.5 Step 5: Guidance
At this point, flight controller synthesis is completed and guidance laws can be evolved. Guidance controller comprises two control laws, for the vertical and horizontal load factor demands, n_{yk}^d and n_{zk}^d respectively. This is equivalent to acceleration demands

$$a_{yk}^d = g n_{yk}^d \text{ and } a_{zk}^d = g n_{zk}^d \tag{22}$$

These demands are passed through the kinematic converter to form the flight controller inputs n_y^d and γ^d.

It is important to note that when the controller is evolved in a stochastic environment with many random factors, the evolution progress becomes highly stochastic as well. It is desirable to obtain the fitness value on the basis of multiple simulation runs, reducing

thereby sampling errors and making fitness value more deterministic and reliable. Excessive number of simulation runs, however, may slow down the evolution considerably, thus an optimal number should be found.

Second, for the same reason, stagnation of progress of the current best fitness is not a reliable indicator of stagnation of the evolution. Every generation may produce individuals with very good fitness which will not sustain further selections. In some cases, current best fitness may appear stagnated from the very first generation. Average population fitness may be considered as a better measure of evolution progress.

Third, elitism has little sense in a sufficiently stochastic environment and is usually not employed. If it is desirable, for some reason, to use elitist strategy, fitness of the elite solution must be recalculated every generation.

In the guidance task, the random factors include initial state of the UAV; Sea State, atmospheric parameters, and initial state of the ship.

Fitness is calculated as follows:

$$F_g = 40\Delta h_1^2 + 20\Delta z_1^2 + 50|\psi_1| + 25|\gamma_1| + 500C_c\left(n_{yk}^d\right) + 500C_c\left(n_{zk}^d\right) + 100C_f\left(n_{zk}^d\right) \tag{23}$$

where Δh_1 and Δz_1 are vertical and horizontal miss distances, and ψ_1 and γ_1 are final yaw and bank angles. Greater weight for vertical miss than that for horizontal miss is used because vertical miss allowance is smaller (approximately 3–4 m vs. 5 m) and also because vertical miss may result in a crash into the boom if the approach is too low.

Three simulation runs are performed for each fitness evaluation. To provide a more comprehensive estimation of guidance abilities, initial positions of the UAV in these N runs are maximally distributed in space, maintaining, at the same time, a sufficient random component.

Algorithm initialisation is the same as for longitudinal flight control laws evolution, except that elitism is not used and the population size is 48 members. The initial population is sampled with the control laws of the form

$$\dot{x}_6 = 0$$
$$\dot{x}_7 = 0$$
$$a_{yk}^d = k_{11}x_6 + k_{12} \tag{24}$$
$$a_{zk}^d = k_{21}x_7 + k_{22}$$

where all coefficients k are chosen at random. For convenience, the control laws are expressed in terms of accelerations, which are then converted to load factors according to (22). Since these control laws effectively represent 'empty' laws $y = const$ (plus a low-pass output filter with randomly chosen bandwidth), structure mutation is applied to each member of the initial population. Further structure mutations are performed as specified above. With such initialisation, no assumption about possible guidance law is taken.

From the final populations, the best solution is identified by calculating fitness of each member using $N = 25$ simulation runs, taking into account also success rate. The best performing controller is the following (unused state variables are removed) with 100% success rate in the 25 test runs with average fitness 69.52:

$$\dot{x}_6 = -3.548\omega_v$$
$$a_{yk}^d = 0.916x_6 - 54.3\omega_v - 0.1261 \tag{25}$$
$$a_{zk}^d = -107.68\omega_h + 0.0534V_{zT} + 0.4756$$

The other attempted approach is the pure proportional navigation (PPN) law for recovery which was compared with the obtained solutions (Duflos et al., 1999; Siouris & Leros, 1988). The best PN controller have been selected using the fitness values of each member of the final population averaged over 100 simulation runs. It is the following:

$$a_{yk}^d = -3.27V_{CLa}\omega_V$$
$$a_{zk}^d = -3.18V_{CLa}\omega_h \tag{26}$$

This controller received the fitness 88.85 and showed success rate 95%.

5. Controller testing

Testing is a crucial stage of any controller development process as the designer and end user must be satisfied that the controller meets its performance requirements. In addition, an allowable operational envelope must be obtained. Testing of a developed controller can be performed either within the physical system or in a comprehensive simulation environment. In physical implementation, the controller is working on a true model and hence there are no modelling errors to introduce uncertain effects into the system. In addition, physical testing accurately reflects the operational use of the controller. On the other hand, physical testing involves large time and cost demands for multiple tests. Moreover, the failure of the controller at any stage may lead to a dangerous situation and even to loss of the aircraft.

Testing within a simulation environment has the benefit that many different tests may be applied rapidly with no severe consequences in the event of controller failure. The whole range of operating conditions may be tested. On the downside is that potentially large modelling errors may be introduced, which may bias the results and cause an otherwise excellent controller to fail when implemented physically. In addition, testing situations which cannot occur in practice are possible in the simulation environment.

In this work, two main types of simulation tests are conducted – robustness and performance. Results are presented for robustness test, which is aimed at ensuring the controller has good robustness to modelling uncertainties and to test whether the controller is sensitive to specific perturbations.

5.1 Robustness tests

Testing the controller for robustness to model uncertainty involves perturbing the model and determining the effect upon controller performance. The perturbations can be performed in several ways. Physical quantities such as mass and wing area can be changed directly. Dynamics of the system can be varied by introducing additional dynamic elements and by changing internal variables such as aerodynamic coefficients. For a realistic test, all perturbations should be applied simultaneously to identify the worst case scenario. However, single perturbation tests (the *sensitivity analysis*) allow to analyse the degree of influence of each parameter and help to plan the robustness test more systematically. After

single perturbation tests multiple perturbation tests are carried out, where all parameters of the model are perturbed randomly within the identified limits.

5.2 Single perturbation tests

In this type of test, a single model variable is perturbed by a set amount and the effect upon the performance of the controller is determined. Performance evaluation can be measured in a manner similar to fitness evaluation of the guidance controller (equation (23)).

The additional parameters taken into account in performance measurement are impact speed V_{imp} and minimum altitude H_{min} attained during the approach.

Altogether, the performance cost (PC) is calculated as follows:

$$PC = 40\Delta h_1^2 + 20\Delta z_1^2 + 50|\psi_1| + 25|\gamma_1| + 50 f_V + 20 f_H + 10 C_c(\delta_a) + 10 C_c(\delta_r) + C_c(\delta_e) + \ldots \\ + 200 C_f(\delta_a) + 200 C_f(\delta_r) + 200 C_f(\delta_e) \tag{27}$$

where

$$f_V = \begin{cases} V_{imp} - 30, & V_{imp} > 30 \\ 0, & V_{imp} \le 30 \end{cases}, \quad f_H = \begin{cases} 10 - H_{min}, & H_{min} > 10 \\ 0, & H_{min} \le 10 \end{cases}. \tag{28}$$

Impact speed V_{imp} and minimum altitude H_{min} are measured in m/s and metres respectively. Other designations are as in (23). Unlike fitness evaluation in the flight controller evolution, the commanded control deflections δ_a, δ_r and δ_e are saturated as required for control actuators.

The absolute value of the PC obtained using (27) is not very illustrative for comparison between the results. Smaller values indicate better performance, but the 'ideal' zero value is unreachable because a minimum level of control activity is always present even in very calm environment. For this reason, a *Normalised Performance Cost* (NPC) will be used:

$$NPC = \frac{PC}{PC_{ref}} \tag{29}$$

where PC_{ref} is the reference Performance Cost obtained for the reference (unperturbed) model with the guidance controller being considered. NPC > 1 indicates deterioration of performance. However, the performance with the perturbed model may be better than the reference performance, thus NPC < 1 is also possible.

The environment delivers a great deal of uncertainty. To obtain a reliable estimate of performance, several tens (up to a hundred) of simulation runs in various conditions should be performed at every point. However, this would imply a prohibitively high computational cost. Meanwhile, there is no single 'typical' scenario that could encompass most of the real flight conditions. For these reasons, in this work few different scenarios were used for evaluation of a single Performance Cost to cover most of the possible real world conditions. The range of disturbances and the initial ship phase are chosen to provide a moderately conservative estimation. All random parameters (which include the turbulence time history

and the ship initial state) are reproducible between the PC estimations, therefore PC is calculated deterministically. Results presented in this work are for calm environment:

Calm environment. No wind, no turbulence; initial position of the UAV is at the ideal reference point: distance 300 m, elevation 14 m, zero sideways displacement. A small amount of ship motion corresponding to Sea State 2 (SWH)2 is, however, included. This scenario is useful to analyse the performance in benign environment and also to soften the conservative bias to the difficult conditions.

The first test determines the robustness to time delays in the measurement signals.

The tests are carried out for the two controllers selected in Section 4.5. The controller #1 is (26) (pure Proportional Navigation guidance) and the controller #2 is (25). Fig. 11a shows the NPC evaluated for these controllers for varying time delay.

Time delays occur in a physical system due to several factors, which include delays associated with the measurement devices, delays in encoding, decoding and transmitting the signals to the controllers, delays in controller computation and delays in the actuator systems.

The miss distance (averaged over the four scenarios) is shown in Fig. 11b. It can be seen that in terms of guidance accuracy, delays up to 0.07 s for the second controller and up to 0.09 s for the first controller cause little effect. However, considering the NPC value, which also takes into account control activity, sufficient degradation is observable from 0.035 second delay. This is expectable, because delays usually cause oscillations in the closed-loop system.
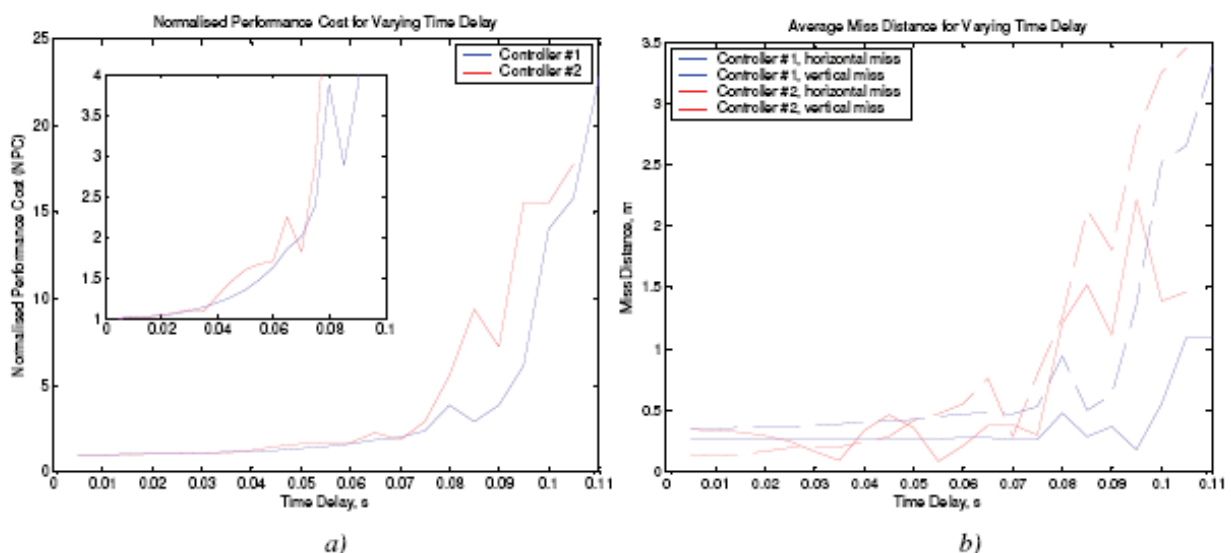

Figure 11.10 NPC (a) and miss distance (b) for controllers with time delays

For the selected NPC threshold, the maximum time delay can be defined as 0.065 s for the controller #1 and 0.06 s for the controller #2.

It should be noted, however, that perturbation of any one of the parameters does not provide a comprehensive picture of the effect on performance. For this reason, the effect of perturbation of empty mass m_{empt} will be closely examined as well. Unlike time delay, this perturbation should primarily affect the trajectory rather than the control activity.

m_{empt} is varied from 50% of the reference value (20.45 kg) to 200% linearly with 2% steps.

This is done separately for increasing and decreasing the variable until the threshold NPC = 1.8 is crossed, or a failure occurs, or the limit (200% or 50% respectively) is reached.

The parameters corresponding to aircraft geometry and configuration are tested in a similar manner. The range is increased to the scale factors between 0 and 10 (0 to 1000%) with step 0.05. NPC is linearly interpolated at the intermediate points. The allowable perturbations (as factors to the original values) are summarised in Table 1, where * denotes the extreme value tested.

Parameter	Lower limit factor		Upper limit factor	
	Controller 1	Controller 2	Controller 1	Controller 2
Empty mass (m_{empt})	0.50*	0.50*	1.54	1.52
Rolling moment of inertia (I_x)	0.20	0.20	10*	2.43
Yawing moment of inertia (I_y)	0*	0*	10*	10*
Pitching moment of inertia (I_z)	0*	0*	2.17	1.91
xy cross product of inertia (I_{xy})	0*	0*	10*	10*
Wing area (S)	0.74	0.87	1.55	1.60

Table 1. Allowable perturbations of UAV inertial properties and geometry

Fig. 12 demonstraites results for calm environment. Control signals for the worst two cases (with NPC > 3.8) are not shown, they exhibit extremely aggressive oscillations. Dashed cyan and green lines on the trajectory graphs represent the 'unrolled' along the flight path traces of the tips of the recovery boom (lateral position on the top view and vertical position on the side view). The bar on the right hand side illustrates the size of recovery window. The height and position of the window may be slightly different for different trajectories because it is determined by the shape of the arresting cable-hook. However, in most cases this is barely noticeable on the graphs since the difference in flight time and terminal part of trajectory is small. Note that the horizontal and vertical scale on the trajectory graphs is different.

The other parameters evaluated in this research were the perturbations of the power unit parameters, the aircraft aerodynamics parameters and sensor noise (Khantsis, 2006).

Overall, these tests have not exposed any significant robustness problems within the controllers. Large variations in single aircraft parameters caused very few control problems. However, such variations cannot realistically judge the performance of the controllers under simultaneous perturbations of multiple parameters. In order to compare the practical robustness of the aircraft, the following tests are performed.

5.3 Simultaneous multiple perturbation tests
These tests involve simultaneous perturbation of all the aircraft variables by a random amount. This style of testing requires many simulations to ensure adequate coverage of the testing envelope. Normally distributed random scale factors are used for perturbations.
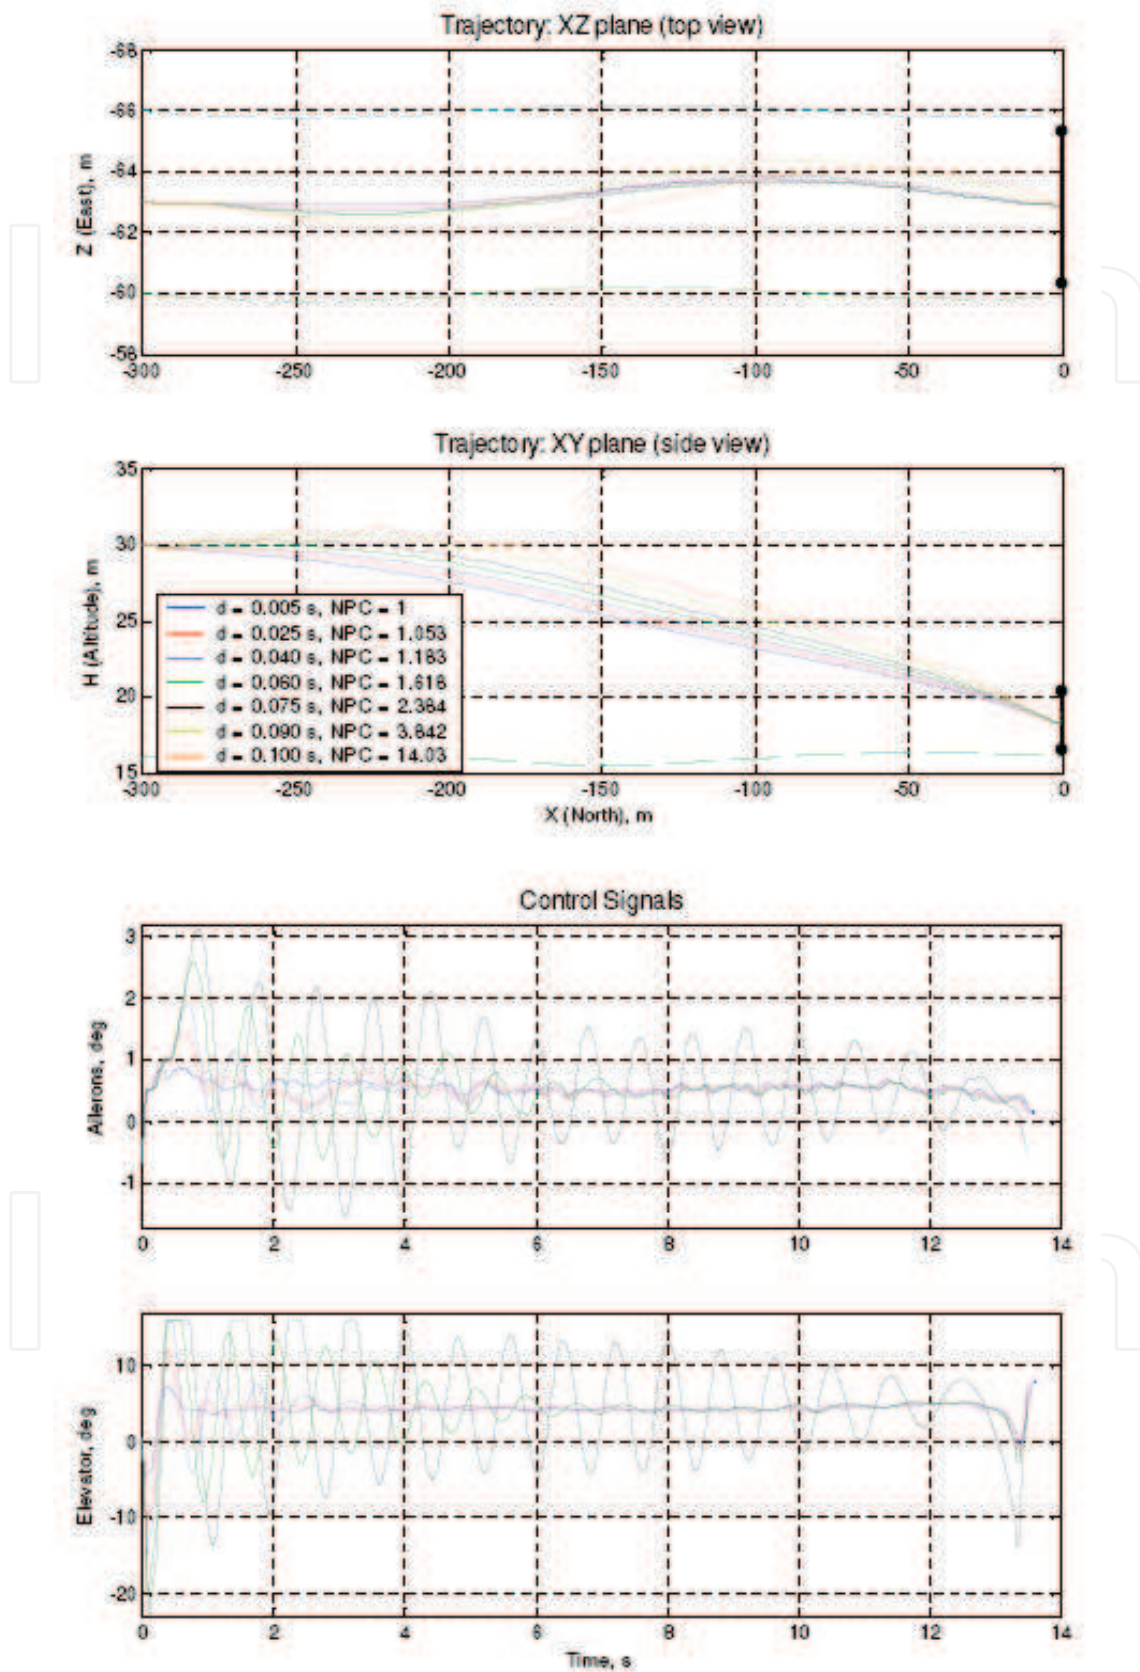
Figure 12. Flight path and control signals with time delays

The chosen standard deviations are approximately one tenth of the maximum allowed perturbations identified in the previous robustness tests.

For each of the two controllers, 1000 tests (4000 simulations considering four scenarios) have been carried out. Note that the perturbations were the same for both controllers in each case. This allows to compare the performance between the controllers in similar conditions. The second controller showed much poorer robustness. 90.4% of test cases showed NPC < 1.8. At the same time, only 69.8% of the tests reported successful recovery. Such a large proportion of cases which acceptable NPC value but without capture (about 20%) indicates that the controller delivers sufficiently inferior guidance strategy when subject to model uncertainty. This could be expected because this controller uses more positioning measurements. The majority of these cases involve a miss in the tailwind scenario, highlighting a potential guidance problem. The number of cases with a crash in one or more scenarios is 2.8%, which is a significant increase over the first controller as well.

From these tests, the Proportional Navigation controller is clearly preferable. However, it must also be tested over a wide range of scenarios.

6. Conclusions

In this chapter, an application of the Evolutionary Design (ED) is demonstrated. The aim of the design was to develop a controller which provides recovery of a fixed-wing UAV onto a ship under the full range of disturbances and uncertainties that are present in the real world environment.

The controller synthesis is a multistage process. However, the approach employed for synthesis of each block is very similar. Evolutionary algorithm is used as a tool to evolve and optimise the control laws. One of the greatest advantages of this methodology is that minimum or no a priori knowledge about the control methods is used, with the synthesis starting from the most basic proportional control or even from 'null' control laws. During the evolution, more complex and capable laws emerge automatically. As the resulting control laws demonstrate, evolution does not tend to produce parsimonious solutions.

The method demonstrating remarkable robustness in terms of convergence indicating that a near optimal solution can be found. In very limited cases, however, it may take too long time for the evolution to discover the core of a potentially optimal solution, and the process does not converge. More often than not, this hints at a poor choice of the algorithm parameters.

The most important and difficult problem in Evolutionary Design is preparation of the fitness evaluation procedure with predefined special intermediate problems. Computational considerations are also of the utmost importance. Robustness of EAs comes at the price of computational cost, with many thousands of fitness evaluations required.

The simulation testing covers the entire operational envelope and highlights several conditions under which recovery is risky. All environmental factors—sea wave, wind speed and turbulence—have been found to have a significant effect upon the probability of success. Combinations of several factors may result in very unfavourable conditions, even if each factor alone may not lead to a failure. For example, winds up to 12 m/s do not affect the recovery in a calm sea, and a severe ship motion corresponding to Sea State 5 also does not represent a serious threat in low winds. At the same time, strong winds in a high Sea State may be hazardous for the aircraft.

Probably the most important consideration in the context of this research is validity of the Evolutionary Design methodology. Whilst it is evident that ED did produce a capable controller that satisfies the original problem statement, several important observations can be made. First of all, in the absence of similar solutions to be compared with, it is unclear how close the result is to the optimum. Considering remarkable robustness of EAs for arriving at the global optimum, it may be believed that the controller is very close to the optimal solution. However, comparing performance of the two generated guidance controllers in different configurations, it may be concluded that there is still room for improvement. The main reason for that is believed to be the limited scope of the testing conditions used in the evolution process, which is mainly due to shortage of computational resources.

On the whole, Evolutionary Design is a useful and powerful tool for complex nonlinear control design. Unlike most other design methodologies, it tries to *solve* the problem at hand automatically, not merely to optimise a given structure. Although ED does not exclude necessity of a thorough testing, it can provide a near optimal solution if the whole range of conditions is taken into account in the fitness evaluation. In principle, no specific knowledge about the system is required, and the controllers can be considered as 'black boxes' whose internals are unimportant. Successful design of the controller for such a challenging task as shipboard recovery demonstrates great potential abilities of this novel technique.

6.1 Limitations of the Evolutionary Design

The first point to note is that ED, as well as any evolutionary method, does not produce inherently robust solutions in terms of both performance and unmodelled uncertainties. As noted above, behaviour of the solutions in the untested during the evolution conditions may be unpredictable. A sufficient coverage of all conditions included in the evolution process invariably requires a large number of simulations, which entails a high computational cost. When the amount of uncertainties in the system and the environment becomes too large, computational cost may become prohibitively high. At the same time, subdivision of the problem into smaller ones and reduction of the number of environmental factors in order to reduce the amount of uncertainties require good understanding of the system. This severely compromises one of the main advantages of ED, which states that little, if any, a priori knowledge of the system is needed.

Another limitation is that evolutionary techniques have little analytical background. Their application is based largely on empirical experience. In particular, most algorithm settings such as population size, selection pressure, probability of genetic operators etc. are adjusted using trial and error method. Whilst EAs are generally robust to these settings, it should be expected that the settings used for controller synthesis in this work may be far from optimal for another application, and their optimisation will take a significant time.

The main limitations of the controller produced in the ED procedure are due to limited validity of the system models available. For a complete control system design, higher fidelity models will be essential.

6.2 Research opportunities on the Evolutionary Design

The ED methodology proved to be easy to apply and extremely suitable for current application. However, the main effort has been directed towards the practical

implementation and application of this design method, with little investigation into its general functionality and efficiency.

An investigation into the most effective algorithm settings, such as the population sizes and probabilities of genetic operators, will also be beneficial. Extension and optimisation of the control laws representation may help to expand the area of possible applications as well as improve efficiency of the algorithm. A large room for improvement exists in the computer implementation of the algorithm. Whilst every effort has been made in this work to optimise the internal algorithmic efficiency, significant optimisation of the coding (programming) is possible.

The methodology presented in this thesis can be used to develop a highly optimal controller capable of autonomous recovery of the UAV in adverse conditions. The Evolutionary Design method, which is the core of this methodology, potentially enables to produce controllers for a wide range of control problems. It allows to solve the problems automatically even when little or no knowledge about the controlled system available, which makes it a valuable tool for solving difficult control tasks.

7. References

Astrom, K. J. & Wittenmark B. (1995). A survey of adaptive control applications, *Proceedings of IEEE Conference on Decision and Control* , pp. 649-654, ISBN: 0-7803-2685-7, New Orleans, LA, USA, December 1995

Atherton, D. (1975). *Nonlinear Control Engineering*, Van Nostrand Reinhold Company, ISBN: 0442300174

Blanchini, F. & Sznaier, M. (1994). Rational L1 suboptimal compensators for continuous time systems,. *IEEE Transactions on Automatic Control*, Vol. 39, No 7, page numbers (1487-1492), (July 1994), ISSN: 0018-9286

Bourmistrova, A. (2001). Knowledge based control system design for autonomous flight vehicle, *PhD thesis*, RMIT University: Melbourne, September 2001

Brown, R. G. & Hwang, P. Y. C. (1992). *Introduction to Random Signals and Applied Kalman Filtering*. 2nd ed. John Wiley & Sons, ISBN-10: 0471128392, ISBN-13: 978-0471128397

Bryson, A. J. & Ho Y.-C. (1975). *Applied Optimal Control*. Hemisphere, Washington, DC, 1975

Bu, T., Sznaier, M. & Holmes, M. (1996). A linear matrix inequality approach to synthesizing low order L1 controllers, *Proceedings of IEEE Conference on Decision and Control* proceedings, pp. 1875-1880 vol.2, ISBN: 0-7803-3590-2, Kobe, Japan, December 1996

Bugajski, D. J. & Enns, D. F. (1992). Nonlinear control law with application to high angle-ofattack flight, *Journal of Guidance, Control and Dynamics*, Vol. 15, No 3, (May/June 1992), page numbers (761-767), ISSN 0731-5090

Coley, D. A. (1998). *An Introduction to Genetic Algorithms for Scientists and Engineers*. University of Exeter, ISBN 9810236026, 9789810236021, Published by World Scientific, 1999

Crump, M. R. (2002). The dynamics and control of catapult launching Unmanned Air Vehicles from moving platforms, *PhD thesis*, RMIT University: Melbourne, 1992

Dahleh, M. A. & Pearson, J. B. (1987). L1 optimal feedback compensators for continuous-time systems, *Proceedings of IEEE Transactions Automatic Control,* pp. 889-895, Vol AC-32, ISSN: 0018-9286, October 1987

Dalsamo, M. & Egeland, O. (1995). H∞ control of nonlinear passive systems by output feedback, *Proceedings of IEEE Conference on Decision and Control* proceedings, pp. 351 - 352 vol. 1, ISBN: 0-7803-2685-7, New Orleans, LA, December 1995

DeCarlo, R., Zak, S. & Mathews, G. (1988). Variable structure control of nonlinear multivariable systems: a tutorial. *IEEE*, Vol. 76, No 3, (March 1988), page numbers (212-232), ISSN: 0018-9219

Doyle, J. C., Glover, K., Khargonekar P. P. & Francis B. A. (1989). State-space solutions to standard H_2 and H∞ control problems, *IEEE Transactions on Automatic Control*, Vol. 34, No. 8, (August 1989), page numbers (831-847), ISSN: 0018-9286

Doyle, J. C. & Stein, G. (1981). Multivariable feedback design: concepts for a classical/modern synthesis,. *IEEE Transactions on Automatic Control*, Vol. AC-26, No 1, (February 1981), page numbers (4-16), ISSN: 0018-9286

Duflos, E., Penel, P. & Vanheeghe, P. (1999). 3D guidance law modelling, *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 35, No 1, page numbers (72-83), (January 1999), ISSN: 0018-9251

Graham, D. & McRuler, D. (1971). *Analysis of Nonlinear Control Systems,* Dover, ISBN-10: 0486610144, ISBN-13: 9780486610146

Hyde, R.(1991). The Aplication of Robust Control to VSTOL Aircraft, *PhD thesis*, Cambridge

Kaise, N. & Fujimoto, Y. (1999). Applying the evolutionary neural networks with genetic algorithms to control a rolling inverted pendulum, *Lecture Notes in Computer Science : Simulated Evolution and Learning*, Volume 1585/1999, (January 1999), page numbers (223-230), Springer Berlin / Heidelberg, ISBN 978-3-540-65907-5

Kalman, R. E. (1960). A new approach to linear filtering and prediction problems, *Transaction of the ASME – Journal of Basic Engineering*: p. 35-45, March 1960

Kaminer, I., Khargonekar, P. P. & Robel, G. (1990). Design of a localizer capture and track modes for a lateral autopilot using H∞ synthesis, *IEEE Control Systems Magazine*, Vol. 10, No 4, (June 1990), page numbers (13-21), ISSN: 0272-1708

Khammash, M. & Zou, L., Almquist, J.A.; Van Der Linden, C (1999). Robust aircraft pitch-axis control under weight and center of gravity uncertainty, *Proceedings of the 38th IEEE Conference on Decision and Control*, pp. 1970 - 1975 vol.2, ISBN: 0-7803-5250-5, Phoenix, AZ, December 1999

Khantsis, S., (2006). Control System Design Using Evolutionary Algorithms for Autonomous Shipboard Recovery of Unmanned Aerial Vehicles, *PhD thesis*, RMIT University: Melbourne, 1992

Kim, B. S. & Calise, A. J. (1997). Nonlinear flight control using neural networks,. *Journal of Guidance, Control and Dynamics*, Vol. 20, No 1, page numbers (26-33), (January 1997), ISSN 0731-5090

Koza, J. R., Keane, M. A., Yu, J., Bennett III, F. H. & Mydlowec, W. (2000). Automatic creation of human-competitive programs and controllers by means of genetic programming, *Genetic Programming and Evolvable Machines*, Vol. 1, No 1, page numbers (121-164), (January 2000), ISSN: 1389-2576

Kwakernaak, H. (1993). Robust control and H∞ optimization - tutorial paper, *Automatica*, Vol. 29, No 2, page numbers (255-273), (February 1993), ISSN: 0005-1098

Maciejowski, J. (1989). *Multivariable Feedback Design.* Addison Wesley, ISBN-10: 0201182432, ISBN-13: 978-0201182439, 1989

MathWorks Inc. *Robust Control Toolbox User's Guide.* 2001. Available online at http://www.mathworks.com/access/helpdesk/help/pdf_doc/robust/robust.pdf

Olhofer, M., Jin, Y. & Sendhoff, B. (2001). Adaptive encoding for aerodynamic shape optimization using Evolution Strategies, *Proceedings of Congress on Evolutionary Computation*, pp. 576-583 vol. 1, Seoul, South Korea, 2001, ISBN:1-59593-010-8

Onnen, C., Babuska, R., Kaymak, U., Sousa, J. M., Verbruggen, H. B. & Isermann, R. (1997). Genetic algorithms for optimization in predictive control, *Control Engineering Practice*, Vol. 5, No 10, page numbers (1363-1372), (October 1997), ISSN: 0967-0661

Pashilkar, A. A., Sundararajan, N. & Saratchandran, P. (1999). A fault-tolerant neural aided controller for aircraft auto-landing, *Aerospace Science and Technology*, Vol. 10, No 1, page numbers (49- 61), (January 2006), ISSN: 1270-9638

Postlethwaite, I. & Bates, D. (1999). Robust integrated flight and propulsion controller for the Harrier aircraft, *Journal of Guidance, Control and Dynamics*, Vol. 22, No 2, page numbers (286-290), (February 1999), ISSN 0731-5090

Rugh, W. J. & Shamma, J. S. (2000). Research on gain scheduling,. *Automatica*, Vol. 36, No 10, page numbers (1401- 1425), (October 2000), ISSN: 0005-1098

Sendhoff, B., Kreuz, M. & von Seelen, W. (1997). A condition for the genotype-phenotype mapping: Casualty, *Proceedings of 7th International Conference on Genetic Algorithms*, pp. 73-80, Michigan USA, July 1997, Morgan Kaufmann

Shamma, J. S. & Athans, M. (1992). Gain scheduling: potential hazards and possible remedies, *IEEE Control Systems Magazine*, Vol. 12, No 3, (June 1992), page numbers (101-107), ISSN: 0272-1708

Shue, S. & Agarwal, R. (1999). Design of automatic landing system using mixed H2 / H□□control, *Journal of Guidance, Control and Dynamics*, Vol. 22, No 1, page numbers (103-114), (January 1999), ISSN 0731-5090

Siouris, G. M., Leros, P. (1988). Minimum-time intercept guidance for tactical missiles, *Control Theory and Advanced Technology*, Vol. 4, No 2, page numbers (251-263), (February, 1988), ISSN 0911-0704

Skogestad, S. & Postlethwaite I. (1997). *Multivariable Feedback Control.* John Wiley and Sons, ISBN-10: 0471943304, ISBN-13: 978-0471943303

Spooner, J. T., Maggiore, M., Ordonez ,R. & Passino, K. M. (2002). *Stable Adaptive Control and Estimation for Nonlinear Systems: Neural and Fuzzy Approximator Techniques.* John Wiley and Sons, ISBN-10: 0471415464, ISBN-13: 978-0471415466

Ursem, R. K. (2003). Models for evolutionary algorithms and their applications in system identification and control optimization, *PhD dissertation*, Department of Computer Science, University of Aarhus, Denmark, 2003

van Laarhoven, P. J. M. & Aarts, E. H. L. (1987). *Simulated annealing: theory and applications.* Dordrecht, The Netherlands: D. Reidel, ISBN-10: 9027725136

Wang, Q. & Zalzala, A. M. S. (1996). Genetic control of near time-optimal motion for an industrial robot arm, *Proceedings of IEEE International Conference on Robotics and Automation*, pp. 2592-2597 vol. 3, ISBN: 0-7803-2988-0, Minneapolis, MN, April 1996

Wise, K. & *Nonlinear*, S. J. (1996). H∞ optimal control for agile missiles, *Journal of Guidance, Control and Dynamics*, Vol. 19, No 1, page numbers (157-165), (January 1996), ISSN 0731-5090

Zhou, K. & Doyle, J. C. (1998). *Essentials of Robust Control.* Prentice Hall, ISBN-10: 0135258332, ISBN-13: 978-0135258330

Aerial Vehicles

Edited by Thanh Mung Lam

This book contains 35 chapters written by experts in developing techniques for making aerial vehicles more intelligent, more reliable, more flexible in use, and safer in operation.It will also serve as an inspiration for further improvement of the design and application of aeral vehicles. The advanced techniques and research described here may also be applicable to other high-tech areas such as robotics, avionics, vetronics, and space.

How to reference

In order to correctly reference this scholarly work, feel free to copy and paste the following:

Anna Bourmistrova and Sergey Khantsis (2009). Flight Control System Design Optimisation via Genetic Programming, Aerial Vehicles, Thanh Mung Lam (Ed.), ISBN: 978-953-7619-41-1, InTech, Available from: http://www.intechopen.com/books/aerial_vehicles/flight_control_system_design_optimisation_via_genetic_prog ramming

INTECH
open science | open minds