

# Can trained monkeys design flight controllers for hypersonic vehicles?

K. J. Austin\*      P. A. Jacobs†

(Received 5 October 2001; revised 3 October 2002)

## Abstract

The supersonic combustion ramjet is an as yet unproven propulsion system for hypersonic flight. Provided it can be developed into a practical vehicle, the ultimate success of sustained hypersonic flight will depend on configuring a robust and stable airframe-propulsion-control combination. To design the longitudinal flight controller for this inherently unstable vehicle we have applied a genetic algorithm, hence the trained monkeys metaphor in the title. Being a nondeterministic search method, there is no guarantee of generating a useful solution, yet given a little direction and enough time it is able to solve hard problems. The controller is built using fuzzy logic rules, directed at manipulating the vehicle's angle of attack through the actuation of symmetric

---

\*Department of Mechanical Engineering, The University of Queensland, Brisbane, AUSTRALIA. <mailto:austin@mech.uq.edu.au>

†as above, <mailto:peterj@mech.uq.edu.au>

<sup>0</sup>See <http://anziamj.austms.org.au/V44/CTAC2001/Aust> for this article,

© Austral. Mathematical Soc. 2003. Published 1 April 2003. ISSN 1446-8735

elevators. A preset structure for the rules is used whereby the design task is to configure the control surface through selection of the rule consequents. To direct the search for a controller design, the genetic algorithm uses simulated flight responses to a range of initial conditions, without linearization of the vehicle model and dynamics. Results for the genetic algorithm designed controller show longitudinal stability and disturbance rejection.

# Contents

<b>1</b>	<b>Introduction</b>	<b>C41</b>
<b>2</b>	<b>Hypersonic Vehicle Model</b>	<b>C44</b>
<b>3</b>	<b>Control Structure and Design</b>	<b>C48</b>
3.1	Fuzzy Logic controller . . . . .	C50
3.2	Genetic Algorithm . . . . .	C52
<b>4</b>	<b>Simulation Results</b>	<b>C54</b>
<b>5</b>	<b>Conclusions</b>	<b>C57</b>
	<b>References</b>	<b>C58</b>

# 1 Introduction

The supersonic combustion ramjet, or scramjet, is expected to extend the flight envelope of air-breathing engines and provide an alternative to rocket propulsion. Current high performance air-breathing engines such as turbojets and ramjets, are limited by the

thermal and structural loads that result from slowing the freestream air to subsonic conditions within the engine flowpath. Scramjets, by virtue of maintaining supersonic flow throughout, have the potential to provide broad access to the hypersonic regime [12]. Presently this hypersonic regime, where the Mach number is greater than 5, is only accessible by rockets [2]. Compared to rockets, scramjets benefit by capturing the oxygen from the atmosphere rather than having to carry it on board. The superior cycle efficiency that results means scramjets appear capable of a high specific impulse at high speeds and have a capacity to accelerate more mass than a rocket of the same size. The weight saving of not having to carry oxidant can be directed towards augmenting vehicle ruggedness and flexibility of use.

The beginning of scramjet research is marked by the 1950's experiments on combustion in supersonic airstreams [5]. Yet, despite the continuous research effort, it remains to be seen whether the scramjet can be developed into a practical propulsion system. Assuming practical scramjet operation is possible, the ultimate success of sustained hypersonic air-breathing flight will depend on configuring a robust and stable airframe-propulsion-control combination. Due to the extreme environmental conditions and the sensitivity of the engine to untrimmed flight, stringent attitude control and trajectory maintenance will be required. Maintenance of the vehicle attitude is therefore a requirement of vehicle stability and propulsive effectiveness.

With actual flight of a scramjet powered vehicle yet to be achieved, the development of control approaches has been through ground based experimental facilities and computational simulation. The most advanced program is NASA's Hyper-X vehicle [6], a scaled down space-plane concept. A recent flight test of Hyper-X prematurely ended with the spectacular failure of the booster elevators [18]. If

it were not for the elevator failure, we would have seen the most advanced test of a scramjet vehicle, and in particular, the control approach. Hypersonic flight control for Hyper-X, or X-43 as it is also known, is realized by processing guidance commands and sensor feedback signals to produce aerodynamic surface commands. The control laws were designed using classical linear control design techniques, covering the 7 seconds of expected scramjet operation and the controlled descent. The longitudinal control law uses angle-of-attack error and pitch rate to derive a symmetric command for the all-moving wing. Feedback gains are scheduled with angle of attack and Mach number, and the surface command is further compensated for changes in dynamic pressure.

The same basic configuration of a feedback loop for longitudinal control is used here, but the controller transfer functions are provided by a set of fuzzy rule bases. Fuzzy logic control (FLC) was chosen for its promise of providing robust control of uncertain systems. Outer loop guidance rules provide a target attitude for trajectory maintenance, while inner loop rules based on angle-of-attack provide stability augmentation while controlling attitude. Whilst it is possible to encode a robust control law within the FLC structure, the realization of such a controller is dependent on the measure of control performance and the design procedure. In this case the objective function used to design the controller is a set of flight simulations. Given a test vector of controller parameters, the objective function returns a performance measure based on the nonlinear flight response to a set of initial conditions. Exposing the controller to a wide range of flight conditions and vehicle arrangements is necessary to achieve full coverage of the input space and should satisfy the disturbance rejection required of an uncertain system.

The question of the capability of *trained monkeys*, as posed by

the title of this paper, refers to the numerical optimization procedure used to configure the controller. Using a genetic algorithm as the optimization tool provides a black box-type design, analogous to the concept of trained monkeys working in parallel, and similar to the early human parallel computers [9]. In genetic algorithms the power of natural evolution is exploited. A population of individuals, which in some form are encoded with the control parameters, undergoes simulated evolution. The purpose of the evolution is the improvement of controller performance as defined by the objective function.

In the following sections the vehicle model is introduced and the control structure and design procedure is defined. Results show the evolution of the controller and simulated flight responses.

## 2 Hypersonic Vehicle Model

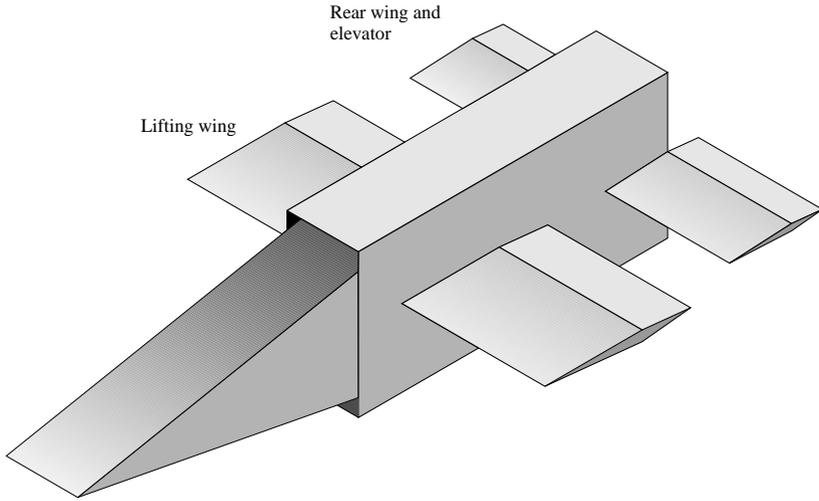
The vehicle concept used in this study is based on the axisymmetric configuration investigated at The University of Queensland [11]. It was originally proposed for a small payload vehicle concept, with the scramjet providing propulsion for the second stage of a three stage launcher [17]. Amongst the studies on flight characteristics, the axisymmetric vehicle is referred to as a winged cone model [15]. The principal design features are derived from decades of scramjet research at NASA. These include round combustors, swept compression surfaces, and a circumferential distribution of the engine modules around the vehicle axis of symmetry. In its generic form the axisymmetric vehicle maximizes the airflow capture area relative to the airframe area, promoting an adequate thrust margin while minimizing configuration drag at near-zero angle of attack.

In our study the vehicle model is in-line with the flight simulation

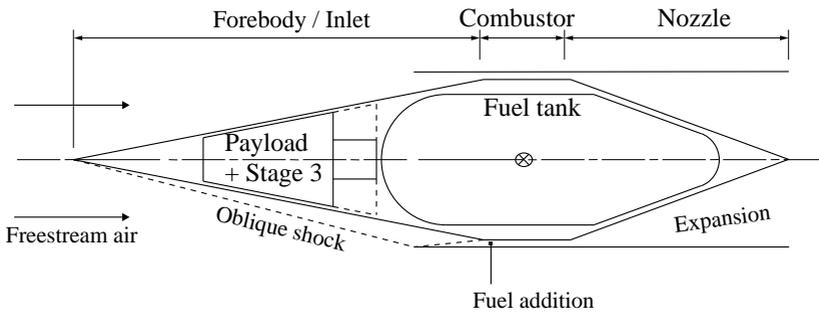
software, thereby avoiding the reduction of flight and control characteristics to a collection of linear models. For each integration time step, a complete aero-propulsive simulation is performed. Three dimensional simulation is computationally expensive so a simplification of the axisymmetric configuration, compatible with a longitudinal flight simulation, is used. Shown in Figure 1, the vehicle has a box section, thus providing two-dimensional flow paths through the propulsion system and over the airframe. The basic shape of the vehicle and engine are otherwise maintained and, importantly, the operational dependencies on attitude and flight conditions remain.

Since a fixed geometry inlet is used, a nominal operating condition at the high Mach number end of the trajectory sets the engine geometry and the overall airframe dimensions [19]. Accordingly, the engine design condition is set at a Mach number of 15, an altitude of 30 km, and zero angle of attack flight. The high dynamic pressure which this equates to means engine performance is given preference over the potential vehicle structural capabilities. Roughly 8.25 m in length, the vehicle has a dry mass of 2500 kg. For small vehicles such as this one, hydrocarbon fuels such as ethane are preferred over hydrogen, as they offer a higher energy density, simpler storage needs, and reasonable specific impulse values [16]. In the configuration shown in Figure 1, 2485 kg of fuel can be stored, offering approximately four minutes of scramjet operation.

The nominal flight trajectory is a constant dynamic pressure extrapolation of the design flight condition. This provides a mapping of velocity versus altitude, using a standard atmosphere model [1]. Because this trajectory equates to near-level flight, the vehicle utilizes a lifting wing with a diamond profile. The centreline of the wing is set at an angle of incidence of  $3^\circ$  with respect to the vehicle centreline. Open-loop stability of such a configuration requires prohibitively large aerodynamic surfaces, so closed-loop stability is



Overall airframe configuration.



Sectional view of the scramjet engine.

FIGURE 1: Baseline design for the scramjet powered stage of a small launch vehicle concept, featuring two-dimensional flow paths. The sectional view of the engine shows a possible arrangement of the payload and fuel, and the flow features for one of the engine flow paths.

provided by an all-moving rear wing arrangement. There is also a significantly smaller contribution from the relative distribution of the freestream air between the modules. Differential throttling with fixed fuel input, provides passive stability augmentation. Active throttling of the engines is also possible for control augmentation, but it is not used here because it lessens the acceleration capability of the vehicle. To counter the large moments generated by the intake surfaces, the rear wing provides relatively large actuation surfaces. For the results shown in this paper an actuation rate limit of  $2 \text{ rad/s}$  is used.

Numerical simulation of the aero-propulsive features includes those associated with internal engine flow and external flow over the wings and control surfaces. Referring to Figure 1, the inlet wedge splits the freestream flow, directing the air into two engine modules. This is done through the action of a pair of oblique shock waves. Following mixing and combustion of the fuel with the supersonic airstream, the products are expanded by the nozzle. Net thrust is available as the difference between the thrust generated by expansion of exhaust gases, and the total drag on the engine and airframe. Instantaneous forces and moments are computed using a combination of one and two-dimensional flow models. It is recognized that many of the real flow features crucial to engine operation are neglected when using these simplified models. However, the basic longitudinal features are dictated by the inlet processing, which despite ignoring viscous effects, is reasonably captured.

The controlled flight study has been limited to the longitudinal characteristics. The restricted dynamics are formed by constraining the general six degree-of-freedom flight equations for a rigid body aircraft in flight relative to a spherical, rotating Earth [7]. More detail on the aero-propulsive modeling and simulation package is available in [3].

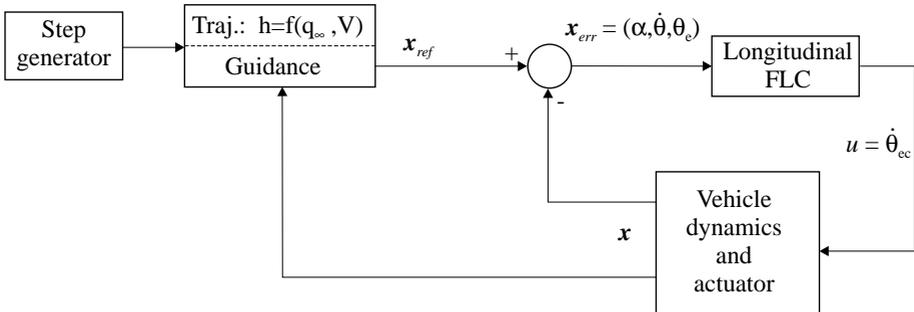
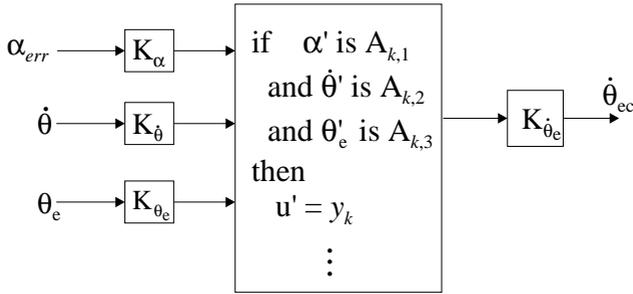


FIGURE 2: Closed-loop attitude and trajectory control model for longitudinal flight. The guidance trajectory block uses the vehicle flight speed  $V$  and a preset flight dynamic pressure  $q_\infty$ , to set a desired altitude, which is transformed into a reference vehicle state  $x_{ref}$ .

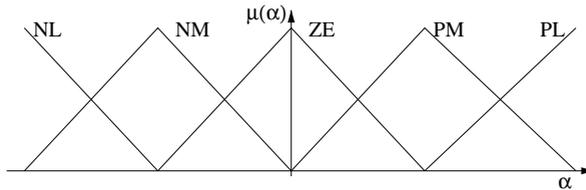
### 3 Control Structure and Design

The two primary tasks of the controller are to maintain vehicle stability and follow a constant dynamic pressure trajectory. While zero angle of attack is desired from the point of view of vehicle acceleration, non-zero angle of attack is needed to correct trajectory deviations or eliminate environmental disturbances. However, with large moments generated at relatively small angles of attack, it is not possible for the vehicle to make rapid or large angle of attack manoeuvres.

A conventional structure for longitudinal control is used, as shown in Figure 2. The flight control function is established by processing outer loop guidance commands to generate a target attitude for trajectory hold. An inner loop attitude control law based on angle-of-attack (AOA) maintains stability while tracking the attitude command. Results for the inner loop alone are presented in this paper.



(a) Angle of attack rule base operation.



(b) Fuzzy input variable description.

FIGURE 3: Longitudinal fuzzy logic controller structure: (a) showing the  $k$ th rule and (b) a typical interpretation of a fuzzy input variable.

The longitudinal control law expresses a functional relationship between a subset of the vehicle state  $(\alpha_{err}, \dot{\theta}, \theta_e)$ , representing angle-of-attack error, pitch rate, and elevator angle respectively, and the actuation command  $\dot{\theta}_{ec}$ , or elevator rate.

The transformation from the vehicle state to the control command is performed by fuzzy *if-then* rules as illustrated in Figure 3. Controller design is configured as a numerical optimization problem, to which a real-coded genetic algorithm is applied.

### 3.1 Fuzzy Logic controller

Fuzzy logic is the logic of human perception. Due to its power in interpreting human operation, fuzzy control has become a major application area, particularly for process control and tasks where traditional automatic control strategies are out-performed by human operators. In the hypersonic control problem we are exploiting other features of fuzzy control. Primarily these are a convenient structure for describing a complex and nonlinear control surface, and the promise of robustness against system uncertainty.

The basic operation of the controller can be drawn in three stages: input interpretation, parallel rule evaluation, and output generation. Input interpretation is simply the partitioning of the input space. Each input variable is partitioned across its domain space with fuzzy sets, as shown in Figure 3(b), where NL, NM, and ZE refer to negative large, negative medium, and zero respectively. The condition statement, if  $\alpha$  is NL, assesses to what degree  $\mu(x)$ , is  $\alpha$  satisfied by the description NL. Triangular fuzzy sets were used to partition the input variables, according to the following membership function.

$$\mu(x) = \begin{cases} 1 - \frac{|x-a|}{b}, & \text{if } |x-a| \leq b, \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

Symmetric partitions with 50% overlap forced nonlinear control elements to be formed by the consequent values rather than the predefined fuzzy sets. All fuzzy variables are defined within the range  $[-1, 1]$  forming a base set to which the input variables are scaled. The input gains  $K$  may then be interpreted as the inverse of the maximum expected or allowed value of the variable.

Each control rule relates a conditional statement, nominally involving all control inputs, and a consequent statement which sets

the control output. A simplified fuzzy inference method is used, where a scalar is used in place of a fuzzy set in the consequence of each rule. Application of the inference mechanism begins with the firing of each of the rules. In Figure 3,  $A_{k,i}$  defines the membership function of the  $i$ th input of the  $k$ th rule. The firing strength  $\beta_k$  of the  $k$ th rule, assuming all conditional statements are treated equally, is

$$\beta_k = \prod_{i=1}^n \mu_{A_{k,i}}(x'_i). \quad (2)$$

A weighted average, based on the firing strength, then provides the control output

$$u = \dot{\theta}_{ec} = K_{\dot{\theta}_e} \frac{\sum_{k=1}^r \beta_k \cdot y_k}{\sum_{k=1}^r \beta_k}. \quad (3)$$

The design approach for fuzzy control is generally dictated by the problem information available. For example, rules can be assembled using operator knowledge, physical constraints, desirable operating characteristics, or other heuristic information. Such information may establish the scaling of input and output parameters, the structure of the rules, or the consequent setting for each rule. An alternative approach is to use a design metric such as the system response to a given task, and through a numerical iterative approach, design some or all of the features of the controller. Here, the design of the controller has been configured as a numerical optimization problem. The rule base structure is preset, as is the partitioning of the input variables. Variable scaling is set using expected or allowed response amplitudes, leaving to the design task, the specification of rule consequents  $y_k$ .

## 3.2 Genetic Algorithm

As a design scheme, the genetic algorithm benefits from problem independence, allowing a black box design approach, and a search process which scales well with problem size. Both of these features are desirable for control design, especially where there is little understanding of the system behavior, a potentially large number of rules to configure, and a complex system model.

A genetic algorithm is a general search procedure, belonging to a larger set of evolutionary algorithms which simulate the process of natural evolution. The simple structure shown in Figure 4 has been used for this study. A randomly sourced population of individuals undergoes an evolutionary process of reproduction through selection for mating according to fitness, and recombination via crossover with mutation.

To use as a numerical optimization tool, it is necessary to have a method of encoding the problem as an artificial chromosome and a means of discriminating between good and bad solutions. Each individual in the population is a possible control configuration, encoded as an array of floating point values. We configure our genetic algorithm using real-coding rather than a binary alphabet, for the benefits it offers in reliability and search velocity on numerical optimization tasks [8, 13].

The problem definition provides an evaluation measure, referred to as the objective function, which in this case is extracted from a collection of flight simulations. When scaled using linear scaling with sigma truncation, the objective function becomes the fitness measure used to direct the search.

Stochastic remainder selection without replacement is used to select parents for mating, with complete replacement of the popula-

```
begin
   $t = 0$ 
  initialize population (random)
  evaluate population
  while ( $t < T$ ) do
     $t = t + 1$ 
    select parents for reproduction, biasing the fittest
    recombine individuals via crossover with mutation
    evaluate new population
  end
end
```

FIGURE 4: The basic genetic algorithm structure.

tion for each generation. This simply amounts to allowing multiple copies of parents with greater than average fitness in the mating array, and randomly selecting parents from that array for reproduction. New individuals are created by sharing parent information through crossover, and introducing new information through mutation. Arithmetic crossover [14] at a fixed probability  $p_c$  is used with the mixing parameter randomly generated each time recombination occurs. It is applied uniformly to the parent chromosomes. Each gene in the children's chromosome also undergoes mutation with probability  $p_m$ .

Arithmetic crossover generates offspring by blending the information from the parents. The non-uniform mutation operator proposed by Michalewicz and Janikow has been modified to avoid biasing the mutation to the centre of the search range. We refer to the modified operator as an adaptive range mutation. In this operation a nominally symmetric mutation range is established about the original value, with mutation yielding a value within that range.

The range is non-uniform across generations meaning the possibility of large mutations reduces with each generation. Details of the algorithm and its operators are given in [4].

## 4 Simulation Results

To provide stability augmentation for the accelerating scramjet vehicle, an angle of attack (AOA) control function is used. The control design task is the optimization of the rule based controller across the hypersonic trajectory, for possible AOA trim commands of  $\pm 3^\circ$ . A predefined rule base structure was used with control inputs  $\alpha$ ,  $\dot{\theta}$ , and  $\theta_e$ , and a control output command  $\dot{\theta}_e$ . Input and output variable scaling was preset to appropriate values, according to the maximum disturbance considered. In the example presented, the control function was represented by 27 fuzzy rules, and evolved over 500 generations using a population of 30 individuals. Genetic operators were applied with probabilities  $p_c = 0.6$  and  $p_m = 0.2$ .

Controller performance was assessed by applying a collection of closed-loop performance measures to a set of simulated flight responses. For the results presented here a set of 42 initial condition arrangements were generated, based about six nominal operating points from the proposed trajectory. The five performance measures that were used included: simulation time, average angle of attack error,  $(1/t_f) \int |\alpha_{err}| dt$ , the integral of absolute error,  $\int_0^{t_f} |\alpha_{err}| dt$ , and the maximum AOA error and pitch rate recorded in the last second of the simulation time. Each performance measure was scaled relative to a desired response tolerance, providing a maximum bound for each measure of 100. The overall objective function was a weighted sum of the individual performance measures. All but the simulation time component were applied non-uniformly across the design,

such that during the early design stages the primary selection pressure was on the development of control solutions which can at least prevent vehicle failure.

Figure 5 shows the evolution of the controller through the objective function history. The initial step-wise increase is primarily due to the introduction of stability robustness features amongst the population of control solutions. With the majority of test conditions satisfying the simulation time measure, the performance measures relating to the response characteristics lead to a rapid rise in the quality of the solution. The transition at generation 250 is due to the gradual reduction in the tolerance of the desired AOA response, placing greater demands on the steady state response once good control solutions have been established. Over the second half of the design, controller robustness has been established and the genetic algorithm is able to fine-tune a stable system. With only 27 rules, the generation of a stable system occurs rapidly with a few of the rules dominating the behaviour. For larger rule sets, formed through greater partitioning of the inputs, the rise in the objective function is more gradual, with more time required to establish robustness features.

A sample flight response is shown in Figure 6 for the best control solution generated after 500 generations. To show the flight response across the trajectory, every 2 seconds the reference AOA is randomly perturbed and the flight condition moved along the nominal trajectory. The controller rapidly trims to the reference condition and performs robustly across the flight envelope. To initiate a change in vehicle trim the elevator must first move in a direction away from the final trim position. This can lead to vehicle failure for large initial trim errors, placing limits on the allowed step commands in angle of attack and requiring representation in the objective function through a large test set. Larger rule bases can

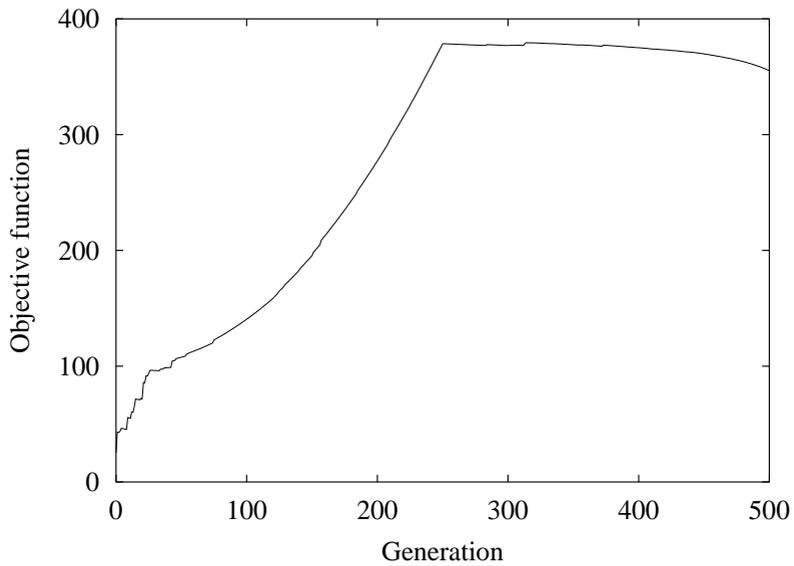


FIGURE 5: Evolution of the angle of attack controller depicted by the objective function history.

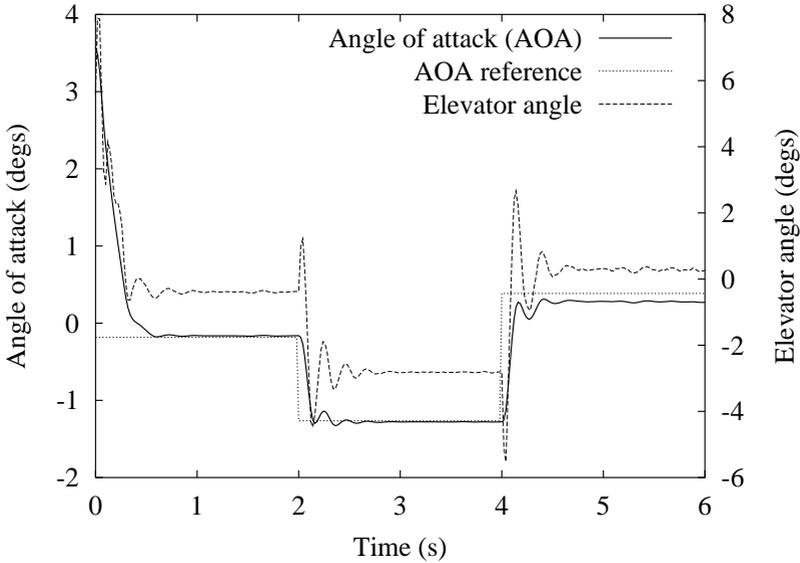


FIGURE 6: Closed-loop response of the scramjet vehicle to angle of attack commands.

provide greater separation between large error responses and fine-tuning the response, and therefore provide more desirable response characteristics.

## 5 Conclusions

The control design procedure presented in this paper utilizes a genetic algorithm as a function optimization tool. It is a computationally demanding task despite the relatively small population and generations used. While genetic algorithms justifiably lay claim to finding solutions to difficult problems, as with any approach, they are subject to the *no free lunch* theory [10]. Along with the time

required to evolve the control parameters, there are inherent difficulties with the black box approach. It requires careful consideration of the representation of the objective function, including their relative magnitudes. Undesirable features not readily apparent in the initial specification of the objective may well be exposed by the genetic search. It may also be difficult to determine the source of problems when the design fails to produce a useful controller. Possibilities include algorithm settings, algorithm structure, the system model, exceeding the capabilities of the vehicle through inappropriate initial conditions or disturbances, or an erroneous objective function. The authors have experienced all of these.

For a single rule base to be applied to the entire trajectory, the design procedure must be exposed to a large array of initial conditions, inevitably compromising the local stability for broad range stability. Improved performance is likely to be available by introducing additional control inputs which allowed scheduling across the flight trajectory. However, there is no avoidance of the considerable computational time necessary to design such a controller.

In the context of the metaphor in the paper title, trained monkeys are capable of designing a hypersonic flight controller, but they need to be treated with care and patience.

## References

- [1] —. *U.S. Standard Atmosphere*. U.S. Government Printing Office, Washington, D.C., 1976. C45
- [2] J. D. Anderson. *Hypersonic and High Temperature Gas Dynamics*. McGraw-Hill, 1989. C42

- [3] K. J. Austin. *Evolutionary Design of Robust Flight Control for a Hypersonic Vehicle*. PhD thesis, The University of Queensland, 3 December 2002. C47
- [4] K. J. Austin and P. A. Jacobs. An adaptive range mutation operator for real-coded genetic algorithms. *Evolutionary Computation*, 2001. Submitted Sept. 2001. C54
- [5] F. S. Billig. Research on supersonic combustion. *Journal of Propulsion and Power*, 9(4):499–514, July-Aug 1993. C42
- [6] J. Davidson, F. Lallman, J. D. McMinn, J. Martin, J. Pahle, M. Stephenson, J. Selmon, and D. Bose. Flight control laws for NASA’s Hyper-X research vehicle. AIAA Paper 99-4124, 1999. C42
- [7] Bernard Etkin. *Dynamics of Atmospheric Flight*. Wiley, 1972. C47
- [8] David E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley Publishing Company Inc., 1989. C52
- [9] D. A. Grier. The math tables project of the work projects administration: The reluctant start of the computing era. *IEEE Annals of the History of Computing*, 20(3):33–50, 1998. C44
- [10] Y. C. Ho. The no free lunch theorem and the human-machine interface. *IEEE Control Systems*, 19(3):8–10, June 1999. C57
- [11] P. A. Jacobs, M. N. Macrossan, and R. J. Stalker. ARC collaborative research grant, final report - effects of scale on scramjet vehicles. Report 13/96, Department of Mechanical Engineering, The University of Queensland, September 1996. C44

- [12] R. A. Jones and P. W. Huber. Toward scramjet aircraft. *Astronautics and Aeronautics*, pages 38–48, February 1978. C42
- [13] Z. Michalewicz. *Genetic Algorithms + Data Structures = Evolution Programs*. Springer-Verlag, 3rd edition, 1996. C52
- [14] Z. Michalewicz and C. Z. Janikow. Genetic algorithms for numerical optimization. *Statistics and Computing*, 1:75–91, 1991. C53
- [15] E. Mooij. Numerical investigation of model reference adaptive control for hypersonic aircraft. *Journal of Guidance, Control, and Dynamics*, 24(2):315–323, 2001. C44
- [16] R. G. Morgan and G. A. Wallace. Ethane combustion in a hypersonic duct. In *Shock Tunnel Studies of Scramjet Phenomena*, volume 1 of *NASA Grant NAGW 674*. 1988. C45
- [17] E. C. Semple. Trajectory and performance of an expendable scramjet powered launcher. In *Seventh National Space Engineering Symposium*, pages 37–41. The Institution of Engineers, Australia, 1992. C44
- [18] B. A. Smith. Elevon failure precedes loss of first X-43A. *Aviation Week & Space Technology*, 154(24):50–51, June 2001. C42
- [19] R. J. Stalker. Thermodynamics and wave processes in high Mach number propulsive ducts. AIAA Paper 89-0261, January 1989. C45