Neuro-Genetic Adaptive Attitude Control

by

Dimitris C. Dracopoulos and Antonia J. Jones

Department of Computing, Imperial College London

DEPARTMENT OF COMPUTING

IMPERIAL COLLEGE OF SCIENCE, TECHNOLOGY

AND MEDICINE

UNIVERSITY OF LONDON

180 Queen's Gate, London, SW7 2BZ

Telephone: 071-589 5111 Ext: 5017/5096

Telefax: 071-581 8024

E-Mail: ajj@doc.ic.ac.uk

Published: Neural Computing & Applications, **2**(4):183-204, 1994. ISSN 0-941-0643.

Abstract. It has previously been demonstrated that for smooth dynamic systems, using relatively few sample points from a single trajectory, a neural network can be trained to perform very accurate short term prediction over a large part of the phase space. In this paper we exploit the capability of a locally predictive network (LPN) to derive an adaptive control architecture for a satellite equipped with controllable, bidirectional thrusters on each of the three principal axes. It is assumed that a hardware implementation of the neural network is available. The inputs for the network are a small history of system states up to the present time and a set of current control inputs, the outputs are the next system state. Once the LPN has been successfully trained, at each time step a genetic algorithm searches the space of hypothetical control inputs. Given a set of control signals the LPN is used to predict the state of the system at the next sample point. This enables the 'fitness' of each set of hypothetical control torques to be evaluated very rapidly. In effect the genetic algorithm determines a satisfactory solution to the inverse kinematic problem in time to apply the solution (set of control torques) at the next control point. With the exception of the neuromodelling (which is repeated only when the system dynamics change) the whole process is then repeated. The results presented indicate that such an architecture is easily able to master the attitude control problem for arbitrary slew angles, with arbitrary a priori *unknown* dynamics and noise in the sensor system.

Keywords. Neural Nets, Genetic Algorithms, Attitude Control, Euler equations, Adaptive Control, Satellites, Modelling, Prediction.

CONTENTS

List of Figures

Neuro-Genetic Adaptive Attitude Control

Dimitris C. Dracopoulos and Antonia J. Jones

Introduction.

The aim of this work is to demonstrate applications of techniques from neuromodelling and genetic algorithms to specific problems of adaptive control. In previous work [Dracopoulos 1993] we gave various examples of neuromodels used to make short range predictions for dynamic systems of various complexities. In this paper we shall address the issue of control.

In choosing a specific dynamic system to be modelled a number of criteria must be considered. The work described in this paper studies the feasibility of applying neural architectures to the control of highly non-linear systems, and so is aimed at exposing the *principles* involved, rather than modelling specific real systems such as aircraft. The system to be modelled should be low dimensional, in order not to obscure the essential nature of the problem, but exhibit sufficiently complex non-linear behaviour so as to pose problems for conventional control system methodologies.

The system chosen is indeed low dimensional: a rigid body, equipped with thrusters which provide controllable torques about three orthogonal axis. It is assumed that the rigid body is not acted upon by any external forces. The dynamics of such a system are classical and described by the Euler equations. We briefly outline the basic equations in a later section. Numerical integration is then used to simulate the dynamics of the rigid body system. In some special cases precise analytic solutions can be found and these can be used to check the accuracy of the simulator. Using the simulator it is a simple matter to introduce further complications for the adaptive control system, such as varying the geometry of the system, creating malfunctions of the thrusters, introducing noise in the sensor system, or imposing externally applied torques.

Neurocontrol and Genetic algorithms.

Since 1988 many papers have been published on neurocontrol and in general there are four basic approaches: supervised control, direct inverse control, neural adaptive control and reinforcement learning.

In *supervised control* there is a need for a training set of $\boldsymbol{x}(t)$ and $\boldsymbol{u}^*(t)$, where $\boldsymbol{x}(t)$ is the system state and $\boldsymbol{u}^*(t)$ is a suitable control input. The control actions often come from recording the actions of human beings solving the control problem. One early example of a supervised control system was Widrow's broom balancer in the 1960's. Since then many other supervised control applications have appeared, for example see [Miller 1990].

Many neurocontrol examples use *direct inverse control* based on some form of supervised learning. The network outputs are the desired control signals and the inputs are the state of the system. The history of inputs and outputs comes from letting the system move around and recording actual states and control signals. However, in general the problem of finding an inverse model of a dynamic system is *ill posed* in the sense that the solution to the problem is not unique. This can pose major problems in some types of applications.

Classical *model reference adaptive control* (MRAC) techniques solve a problem very similar to that of direct inverse control [Miller 1990]. The user, or higher order controller, is asked to supply a reference model that outputs a desired system trajectory. The control system is asked to send control signals to the plant (system), so as to make the plant follow the desired trajectory. It is usually assumed that the plant and the controller are both linear, but that the parameters of the plant are not known ahead of time. Various techniques have shown how to adapt the parameters of the control system so as to solve this problem. Neural adaptive control has been defined as the use of artificial neural networks in place of more classical mappings, within the classical designs of adaptive control theory [Miller 1990]. This approach is still quite new and is mainly represented by the work of Narendra at Yale University.

Narendra has developed three designs so far, that extend MRAC principles to neural networks [Narendra 1992]. The first is essentially just direct inverse control, adapted very slightly to allow the desired trajectory to come from a reference model. The second, which he called his 'most powerful and general' in 1990, is very similar to a strategy also used by Jordan. In this approach, the problem of following a trajectory is converted into a problem in optimal control [Bryson 1975], simply by trying to minimise the gap between the desired trajectory and the actual trajectory. Narendra's third design is similar to the second design, except that radial basis functions are employed in place of multilayer perceptrons, and certain constraints are placed on the system.

Control based on *reinforcement learning* has its genesis in psychological learning theories, see [Mendel 1970], in which the response of an organism (or system) to a stimulus is strengthened by reward and weakened by punishment. The characteristic features of reinforcement learning algorithms differ in important ways from the corresponding features of supervised learning algorithms. Supervised learning algorithms are usually based on the existence of an error vector rather than an evaluation. Many researchers have used variations of reinforcement control techniques in a variety of applications [Werbos 1988], [Werbos 1989], [Miller 1990].

Genetic algorithms have been applied to a variety of control problems, see for example [Davis 1991], [Kristinsson 1992], [Odetayo 1989], [Grefenstette 1989].

Attitude control.

The orientation control of a rigid body has important applications from pointing and slewing of aircraft, helicopter, spacecraft and satellites, to the orientation control of a rigid object held by single or multiple robot arms. A *rigid body*, can be defined as a system of particles, in which the distances r_{ij} (the distance between the i*th* and the j*th* particle) are fixed and cannot vary with time. In what follows, sometimes the rigid body will be referred to as a satellite or spacecraft, since the need for attitude control systems arises frequently in aerospace technology, but the underlying definitions and methods will be generally applicable to rigid bodies.

The *attitude* of a satellite is its orientation in space. Its rigid body motion is specified by its position, velocity, attitude and attitude motion. The first two quantities describe the translational motion of the centre of mass of the spacecraft and are the subject of what is variously called *celestial mechanics*, *orbit determination*, or *space navigation*, depending on the aspect of the problem that is emphasized. The latter two quantities describe the rotational motion of the body about the centre of mass and are the variables considered in this paper.

Generally orbit and attitude are interdependent. For example, in a low altitude Earth orbit, the attitude will affect the atmospheric drag which will affect the orbit; the orbit determines the satellite position which determines both the atmospheric density and the magnetic field strength which will, in turn, affect the attitude. Although knowledge of the satellite orbit is frequently required to perform attitude determination and control functions, here this dynamical coupling will not be considered and it will be assumed that the time history of the rigid spacecraft position is known and has been supplied by some process external to the attitude determination and control system.

One distinction between orbit and attitude problems is related to their historical development. Predicting the orbital motion of celestial objects was the initial motivation of much of Newton's work. Thus, although the space age has brought with it vast new areas of orbit analysis, a large amount of theory directly related to celestial mechanics has existed for several centuries. In contrast, although some of the techniques are old, most of the advances in attitude determination and control have occurred since the launch of *Sputnik* in 1957. The result is that the language of attitude determination and control are still evolving [Wertz 1978].

Attitude analysis may be divided into determination, prediction and control. *Attitude determination* is the process of computing the orientation of a spacecraft relative to either an inertial reference or some object of interest, such as the Earth. This typically involves several types of sensors on each spacecraft and sophisticated data processing procedures. The accuracy limit is usually determined by a combination of processing procedures and system hardware. In the present paper it will be initially assumed that the rigid body under consideration has sufficient sensors which are normally very accurate. In the final example noise is added to the sensor data. *Attitude prediction* is the process of forecasting the future orientation of the satellite using dynamical models. The limiting features are the knowledge of the applied and environmental torques and the accuracy of the mathematical model of satellite dynamics.

Attitude control is the process of orienting the satellite in a specified, predetermined direction. It consists of two areas: attitude stabilization, which is the process of maintaining an existing orientation, and attitude manoeuvre control, which is the process of controlling the reorientation of the satellite from one attitude to another. However, the two areas are not totally distinct. For example, the stabilization of a satellite with one axis towards the Earth, implies a continuous change (manoeuvre) in its inertial orientation. The limiting factor for attitude control is typically the performance of the manoeuvre hardware and software, and the control electronics, although with autonomous control systems, it may be the accuracy of orbit or attitude information (factors which are not addressed in the present paper) [Wertz 1978].

Some form of attitude determination and control is required for nearly all spacecraft. For engineering or flight-related functions, attitude determination is required only to provide a reference for control. Attitude control is required to avoid solar or atmospheric damage to sensitive components, to control heat dissipation, to point directional antennas and solar panels (for power generation), and to orient rockets used for orbit manoeuvres.

One method for categorizing spacecraft is the procedure by which they are stabilized. The simplest procedure is to spin the spacecraft. The angular momentum of a *spin-stabilized* spacecraft will remain approximately fixed in inertial space for extended periods, because the external torques which affect it are extremely small in most cases. However, the rotational orientation of the spacecraft about the spin axis is not controlled in such a system. If the orientation of three mutually perpendicular spacecraft axes must be controlled, then the spacecraft is *three-axis stabilized*. In this case, some form of active control is usually required because environmental torques, although small, will normally cause the spacecraft orientation to drift slowly. Three-axis stabilized spacecraft may be either non-spinning (fixed in inertial space) or fixed relative to a possibly rotating reference frame, as occurs for an Earth satellite which maintains one face towards the Earth and therefore is spinning at one rotation per orbit. Many missions consist of some phases in which the spacecraft is spin-stabilized and some phases in which it is three-axis stabilized. Some spacecraft have multiple components, some of which are spinning and some of which are three-axis stabilized [Chobotov 1991].

In general, a spacecraft attitude control system consists of the following four major functional sections: sensing, logic, actuation, and vehicle dynamics. The logic programs the electronic signals in a correct sequence to the torque producing elements (actuation), which in turn rotate the spacecraft about its centre of mass. The resulting motion (dynamics) is then monitored by the vehicle sensors which thus close the loop of the spacecraft attitude control system.

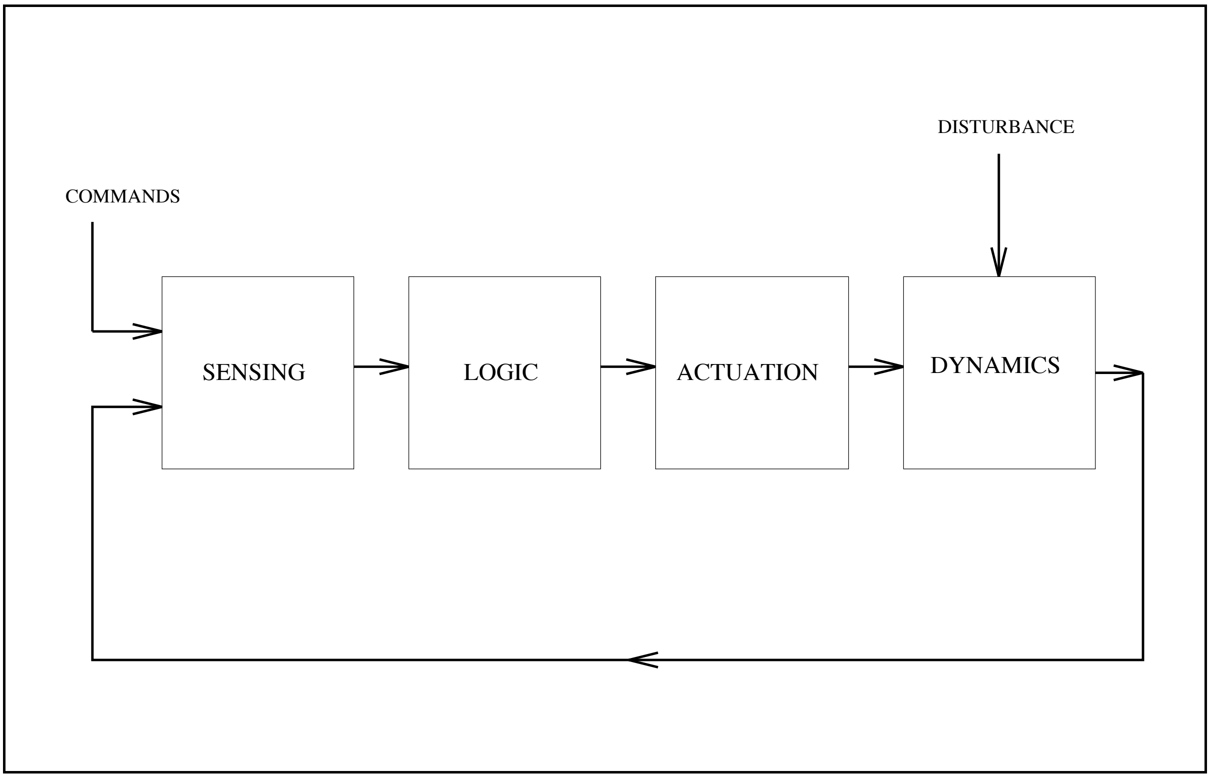

Figure 1 Schematic overview of a satellite attitude control system.

Rigid body rotation.

Angular velocities of the rigid body.

In this section, we consider the mathematical system which describes the rotational motion of a rigid body. This is determined by a system of first order differential equations (Euler equations).

$$\begin{aligned} I_1\dot{\omega}_1 &= (I_2 - I_3)\omega_2\omega_3 + G_1 \\ I_2\dot{\omega}_2 &= (I_3 - I_1)\omega_3\omega_1 + G_2 \\ I_3\dot{\omega}_3 &= (I_1 - I_2)\omega_1\omega_2 + G_3 \end{aligned} \tag{1}$$

where I_1, I_2, I_3 are the principal moments of inertia, ω_1, ω_2, ω_3 are angular velocities about the principal axes fixed in the rigid body, and G_1, G_2, G_3 are torques applied about these axes.

In general the Euler equations are more complicated than Lorentz's system and for certain choices of I_1, I_2, I_3 and G_1, G_2, G_3 exhibit both strange attractors and limit cycles [Leipnik 1981]. For example, if we choose $I_1 = 3$, $I_2 = 2$, $I_3 = 1$ and

$$\begin{pmatrix} G_1 \\ G_2 \\ G_3 \end{pmatrix} = \begin{pmatrix} -1.2 & 0 & \sqrt{6}/2 \\ 0 & 0.35 & 0 \\ -\sqrt{6} & 0 & -0.4 \end{pmatrix} \begin{pmatrix} \omega_1 \\ \omega_2 \\ \omega_3 \end{pmatrix} \tag{2}$$

then the Euler equations produce a binary system of strange attractors for which the attractor of an orbit is determined by the location of the initial point of the orbit. See, for example, Figure 18 of [Dracopoulos 1993]. The characteristic Lyapunov exponent for this system, calculated using the method in [Shimada 1979] is about 0.236.

Orientation of the rigid body.

Since the frame of reference adopted for the equations of motion is fixed to the rigid body, and moves with it, the position and orientation of the object cannot described relative to this frame. Following Meyer [Meyer 1966] the orientation may be described locally by three angles φ, θ, ψ which represent consecutive rotations about a set of orthonormal axes **i**, **j**, **k** fixed in the body, with origin at the centre of mass of the body. By performing this sequence of rotations we bring the (**i**, **j**, **k**) frame into alignment with an inertial frame (**I**, **J**, **K**).

We may then obtain the kinematic equation and derive the relationship between the inertial frame orientation angles and the body frame angular velocities. Writing the kinematic equation as

$$\begin{aligned} \begin{pmatrix} \omega_1 \\ \omega_2 \\ \omega_3 \end{pmatrix} &= \begin{pmatrix} \dot{\varphi} \\ 0 \\ 0 \end{pmatrix} + \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\varphi & \sin\varphi \\ 0 & -\sin\varphi & \cos\varphi \end{pmatrix} \begin{pmatrix} 0 \\ \dot{\theta} \\ 0 \end{pmatrix} \\ &+ \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\varphi & \sin\varphi \\ 0 & -\sin\varphi & \cos\varphi \end{pmatrix} \begin{pmatrix} \cos\theta & 0 & -\sin\theta \\ 0 & 1 & 0 \\ \sin\theta & 0 & \cos\theta \end{pmatrix} \begin{pmatrix} 0 \\ 0 \\ \dot{\psi} \end{pmatrix} \end{aligned} \tag{3}$$

we obtain

$$\begin{pmatrix} \omega_1 \\ \omega_2 \\ \omega_3 \end{pmatrix} = \begin{pmatrix} 1 & 0 & -\sin\theta \\ 0 & \cos\varphi & \sin\varphi\,\cos\theta \\ 0 & -\sin\varphi & \cos\varphi\,\cos\theta \end{pmatrix} \begin{pmatrix} \dot{\varphi} \\ \dot{\theta} \\ \dot{\psi} \end{pmatrix} \tag{4}$$

Inverting the matrix we obtain

$$\begin{pmatrix} \dot{\varphi} \\ \dot{\theta} \\ \dot{\psi} \end{pmatrix} = \begin{pmatrix} 1 & \sin\varphi\,\tan\theta & \cos\varphi\,\tan\theta \\ 0 & \cos\varphi & -\sin\varphi \\ 0 & \sin\varphi\,\sec\theta & \cos\varphi\,\sec\theta \end{pmatrix} \begin{pmatrix} \omega_1 \\ \omega_2 \\ \omega_3 \end{pmatrix} \tag{5}$$

These angles do not exactly coincide with the usual definition of Euler angles which perform successive rotations about the **k**, **j**´, **k**´´ axes, where **j** - > **j**´ under a rotation about **k**, etc., so as to bring the body axes into coincidence with the inertial frame (**I**, **J**, **K**). However, they serve the same purpose and give rise to a convenient set of equations. If we substitute ω_1, ω_2, ω_3, as given by (4), into (1) we obtain three differential equations for φ, θ, ψ. The solution of these equations in general form presents a formidable problem.

Our simulator for this system used the Adams-Moulton method for numerical integration [Conte 1980]. The simulator is the basic tool for the experiments we shall shortly describe. However, we first place the results in context by reviewing the existing solutions for the attitude control problem.

A literature review for the attitude control problem.

Most of the existing research for the *attitude control problem* has been done since 1957. Many authors have tried to derive new control rules for this problem. Although almost thirty years old, the work of Meyer at *NASA* [Meyer 1966], [Meyer 1971] is still considered one of the most seminal investigations in the area. Meyer, after expressing the attitude error in terms of an error matrix, synthesized a class of control laws for which the control inputs (torques) are functions of the real eigenvector of the error matrix [Wilkinson 1965], [Gollub 1989] and the angular velocity of the controlled body. However, the method is not directly applicable to on-off control, e.g. reaction jet control systems, since the results are valid mostly for reaction wheel control systems. In addition Meyer's techniques are not appropriate for the *adaptive* attitude control problem (when the system dynamics change in an unknown way), since the control laws derived assume that the analytic form of the dynamic system under consideration is known.

In [Jahangir 1993], Jahangir and Howe address the minimum-time attitude control problem for a symmetric rigid body missile. They describe a scheme of generating thruster firing times as functions of the initial and desired state of the missile. The scheme involves the transformation of the state variables and the integration of the transformed states and equations. Although they avoid the iterative part of solving a two boundary value problem, their method assumes memory data storage for table lookup, and numerical integration.

Singh et al. [Singh 1993] used linearization theory to represent the nonlinear dynamics of a space station and discussed the attitude control problem for space vehicles employing control moment gyros. A similar approach using also linearization theory is considered in [Hermes 1985], but this time the problem of a spinning satellite with two small jets is discussed. Attitude control using eigenvector analysis on the linearized attitude equations of motion, for a spinning symmetrical satellite in an elliptic orbit, is used in [Calico 1983]. Attitude control of an inverted pendulum using linearized dynamics of the system is discussed in [Furuta 1984]. [Kline-Schoder 1993] presents a linear regulator feedback law for controlling a tethered satellite. The dynamics and the design considerations for the attitude control of a two-dimensional tethered satellite can be found in [Humble 1990], [Lemke 1987]. Paielli and Bach [Paielli 1993] used linear error dynamics (by considering only the linear terms in the Euler parameters) to derive an attitude control law.

[Long 1992] introduces a coordinate frame for the rest-to-rest reorientation of a satellite, which transforms the original nonlinear problem into a linear one. The approach is especially attractive for optimal control attitude control, since it reduces the 7 X 7 matrix computations to 2 X 2 matrix computations. In [Nicosia 1992] a nonlinear *observer* is proposed, for reconstructing the state variables of a spacecraft. Then existing feedback control laws are used, giving a system which is asymptotically stable in a specific region. An observer based method, for a reaction wheel attitude controller, is proposed in [Tcuchiya 1982], while various control laws for a three reaction wheel, three magnetic torque configuration are found in [Junkins 1982].

Bloch et. al. [Bloch 1992] discusses methods of geometric mechanics for stability analysis of rigid body dynamics. Another feedback controller incorporating gain-scheduled adaptation of the attitude gains, is developed for a linearized model of a gravity gradient stabilized spacecraft in [Parlos 1992]. The simulations are done for the *Space Station Freedom*. Wie et al. evaluate the performance and stability of both classical and modern controllers for the *Space Station Freedom* [Wie 1990]. A simulator description of the *Space Station Freedom* with the McDonnell Douglas Space Systems Co./Honeywell Attitude Control System can be found in [Pope 1991]. More general methods for the design of spacecraft simulators can be found in [Thomas 1991]. A plume flow model to calculate the forces and the heat transfer caused by the firing of attitude control thrusters on satellites, is developed in [Dettleff 1986].

In [Piper 1992], it is shown, based on an analysis of a simple-axis problem, that if the momentum wheel assembly performance parameters are adequately matched to the spacecraft, it is possible to achieve a globally stable equilibrium. S. Vadali [Vadali 1992] derives various attitude control laws, for a space station, in the case of absence of disturbances using Lyapunov's second method. Satellite attitude control laws using Lyapunov's methods are also the subjects of [Chavy 1991], [Van Den Bosch 1986]. The application of a game theoretic control approach, combined with an internal feedback loop decomposition for uncertainties in the moments of inertia of a space station (which are considered constant in time) is described in [Rhee 1992].

[Singh 1989] considers a control law, for a class of uncertain nonlinear systems which can be decoupled by state variable feedback. The law is based on the technique of variable structure and is applied for the control of an orbiting spacecraft which uses reaction jets. Uncertainties in the spacecraft parameters are also considered in [Singh 1987], [Singh 1984]. In [De 1992] a method based on the algebraic theory of rational fractions, for the control of a spinning satellite using gyro-torquers, is discussed. Attitude control using gyro-torquers is also considered in [Iyer 1989]. The attitude control (pitch and roll) of automobiles is the topic in [ElMadany 1992].

Optimal control using nonlinear programming techniques with application to satellite attitude control is discussed in [Herman 1991]. Gregory and Peters propose a device [Gregory 1992] for the direct measurement of the attitude stability of a space station, while [Sheela 1991] present a technique for attitude determination. Three axis attitude control of a rigid body spacecraft using a sliding-mode control law is described in [Dodds 1991]. The approach is valid as long as sliding motion is maintained and the extreme values of the plant dynamic parameters are known.

In [Parkinson 1990] a Kalman filter is used to estimate local magnetic field and perform magnetic attitude control for the *GP-B* satellite, which is scheduled to be flown in 1996. Davis and Levinson [Davis 1990] introduced the use of a gravity-stabilized tether attached to a nonspinning part of a satellite, for enabling its attitude control by employing conventional control systems.

An enhancement in the solving techniques for the two-point boundary value optimal attitude control problem is presented in [Lin 1989]. Various problems, advantages and disadvantages in different choices, associated with the design of a space station control system are discussed in [Branets 1988]. [Singh 1987], [Singh 1988] presented an approach to the three-axis attitude control of a spacecraft-beam-tip body system, based on invertibility results. Momentum management systems designed to face the problem of momentum saturation of the control moment gyros, due to noncyclical external torques acting on the space station, are considered in [Woo 1988].

A near optimal orbit and attitude control system, for a plate-like rigid spacecraft in geostationary orbit, is presented in [Rahasingh 1987]. All the results and conclusions are based on simple linear models. A fuzzy logic controller for the control of a spacecraft is applied in [Daley 1987]. [Singh 1986] proposes a control law for single axis rotational manoeuvres of a spacecraft-beam-tip body (an antenna or reflector), in the presence of an unknown but bounded disturbance torque, acting on the spacecraft. In [Skaar 1986], impulse response functions are used for the selection of control switch times, in the bang-bang control of linear, elastic slewing satellites. A bang-bang control law is also presented in [Dodds 1981] and it is further extended in [Dodds 1984]. The nullification of the accumulated effect of the modelling errors, achieved by a correction in the control to induce the physical system to have a behaviour close to the reference model, is the subject in [Ceballos 1986]. A parameter optimization procedure is applied to find the gains of the described method.

Yuan and Stieber [Yuan 1986] employed a PI compensator augmented by a Kalman filter, to control the communications beams and the attitude angles of a flexible spacecraft. They explored two design methods: the first one based on eigenvalue analysis and the second based on singular value criteria. A review in attitude control systems and beam pointing accuracy can be found in [Muhlfelder 1986], while a general framework for the analysis of attitude tracking control problems can be found in [Wen 1991]. [Singh 1988] proposes a control law for asymptotic function reproducibility of a class of nonlinear systems, such that the output of the system tends asymptotically to a given function. Based on this control law, a nonlinear feedback control law is then derived for the attitude control of a satellite containing symmetric rotors.

The application of a controller consisted of a servo-compensator, a stabilizing feedback loop and a feedforward compensator, to the design of a vertical takeoff and landing aircraft flight control system is discussed in [De Melo 1984]. Hess [Hess 1984] studied aircraft attitude control systems, based on the optimal control model of the human pilot. The optimal control human pilot model has its genesis in the hypothesis that, with limitations and in specific well-defined control tasks, the human pilot can be described in terms of the operation of a linear optimal estimator and regulator. Geometric control theory for rigid body attitude control is considered in [Crouch 1984]. In addition an algorithm for stabilizing the system is outlined as proposed by [Hermes 1980]. A simple two-surface solar controller is described and applied for the attitude control of a spacecraft in [Pande 1982], [Venkatachalam 1982]. A proportional plus derivative control law for attitude control of non rigid body spacecraft is found in [Fenton 1981]. In [Shao-hua 1981] a decomposed controller which consists of two coupled electronic integrators is introduced, for satellite control.

The dynamics of *UOSAT* (low orbit satellite with a principal axis pointing towards the Earth centre and a minimum number of sensors and hardware) and its control using the on-board magnetorquer is given in [Hodgart 1982]. [Chretien 1982] describes the dynamic modelling and control of the *SPOT* French Earth observation satellites. In [Jahanshahi 1982] it is shown that the knowledge of the Voyager's limit cycle motion, as measured by the celestial and the inertial sensors, is adequate to estimate a selected number of errors, which adversely affect the spacecraft attitude knowledge.

Summarizing, so far none of the approaches to the attitude control problem have tried to use neural networks and genetic algorithms, techniques that have recently shown encouraging results in the control of simple systems [Miller 1990]. In addition, most of the attitude control techniques use linearized equations of rigid body motion, something which is not valid over a large part of the dynamic system phase space, and relative few relate to adaptive control.

Neuro-genetic control architectures.

Neurocontrol.

outlines a proposed architecture for control of a dynamic system using neural networks. This is just one possibility among many. For simplicity we shall sometimes use the analogy of an aircraft but the problem addressed is quite general.

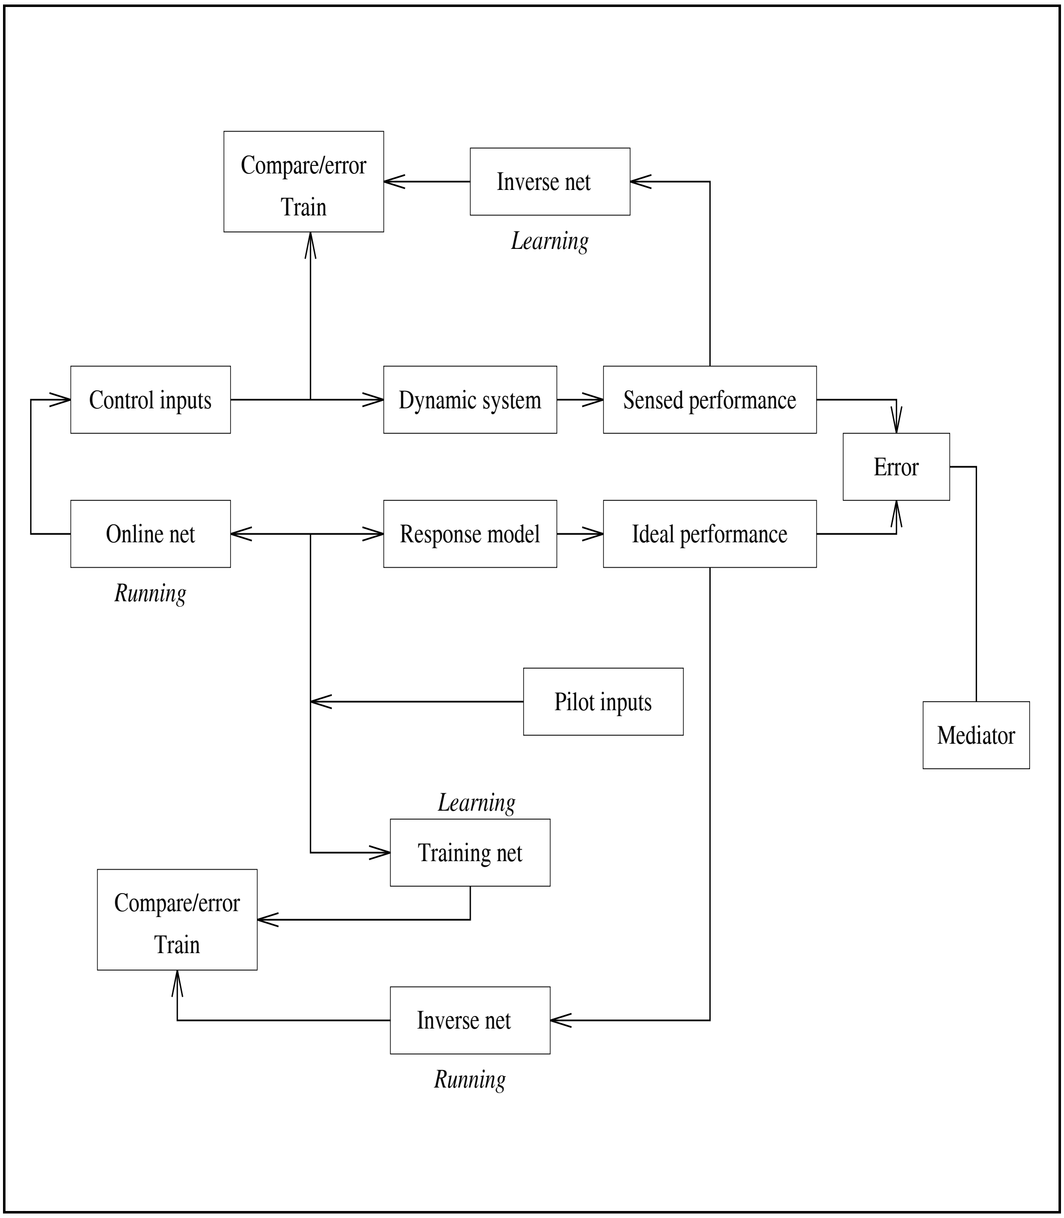

Figure 2 Schematic outline of a possible purely neural adaptive control architecture.

The rationale behind this diagram is as follows.

We assume a dynamic system with control inputs. We could, for example, think of the dynamic system as an aircraft with thrusters and control surfaces. In addition we suppose that a neural network, labelled the *Online network* acts as a mapping from a system we have labelled *Pilot* (although the *Pilot* module need not necessarily be human) to the actual control inputs of the dynamic system. The need for such an interface in modern high speed jets, for example, stems from the fact that stability is often sacrificed for performance. This leads to a highly non-linear response to the *Pilot* inputs. Usually such interfaces are relatively fixed (subject to the normal negative feedback loops of modern control systems). The interface can for example be trained to produce the correct control signals to the aircraft for (say) a shallow dive even though the actual control signals required are a highly non-linear function of speed and altitude and current fuel load.

However, suppose the dynamic system changes its expected characteristics. This could happen in a robotic manufacturing system, or an aircraft, through wear or damage. In this event the interface is no longer appropriate for the actual dynamic system. To some extent the interface could be made adaptive to this eventuality through AI technology, for example Expert Systems, but there are limitations to this approach. In order to have rules which cover all eventualities we must be able to foresee all possible combinations of damage modes. This is plainly unrealistic.

- The general idea behind neurocontrol proposals is that using neural networks the interface can be made adaptive to the currently prevailing dynamics of the actual system.

- In order to accomplish this it must be assumed that we can actually measure the response of the system to the control inputs.

This requires that the system be equipped with sensors to measure all principal system parameters.

- We must further assume that there is sufficient redundancy in the sensor system to provide reliable measurements even if the system as a whole is damaged.

The results of the sensor measurements is labelled *Sensed performance* in Figure 2.

We can further suppose that we know the desired response of the system to any particular *Pilot* signals. This we call the *Ideal Performance*. The *Response model* takes as inputs the *Pilot* control signals and computes the desired response (i.e. the *Ideal performance*). If the *Sensed performance* is incorrect an *Error* is generated. The existence of an error signal indicates a disparity between the actual response to *Pilot* signals and the desired response. However, it may be that with some different *Control inputs* the desired response can still be obtained, or at least a response that more closely approximates the desired response. This indicates an immediate need to modify the interface. Let us consider how this might be done.

In order to modify the interface we have to change the functional relationship between the *Pilot* signals and the *Sensed performance*. In Figure 2 this mapping is implemented by a (previously trained) neural net - the *Online net*. It follows that this net will have to be retrained and this in turn poses certain problems.

The inputs to the *Online net* are *Pilot signals* (which we know) and the outputs are appropriate *Control inputs* to the dynamic system (which we don't now know). In order to train a new *Online net* we need to know the inverse relationship between *Control inputs* and *Sensed performance*, i.e. we need a functional model from the *Ideal performance* to *Control inputs* which minimises the overall system *Error*.

This particular proposal is in essence: to learn the above relationship in a more or less continuous mode using a network called the *Inverse net*. The inverse net takes as input the *Sensed response* of the system and tries to output the *Control inputs* which would produce this response. As long as there is no significant difference between the actual *Control inputs* and the outputs produced by the *Inverse net* then we can consider the inverse net is functioning correctly (there is a caveat here that the whole of the space of control inputs and ideal performance needs to be measured).

If the system sustains damage then it is likely that both the *Inverse net* and the *Online net* will need to be modified. This modification is initiated by the *Mediator* in response to an unacceptable overall system *Error*. The approach taken by the *Mediator* may well be different from one system to another.

The *Inverse net* (top in Figure 2) is trained on a real time basis. Thus after the system sustains damage the *Inverse net* will adapt to model the new inverse dynamics of the actual system. Let us assume that the *Inverse net* has successfully retrained itself.

The *Mediator* now uses the *Inverse net* (bottom in Figure 2) to train the *Training net* on a real time basis. For any given *Pilot* inputs the *Response model* computes the *Ideal performance* and taking this as input the Inverse net determines the required *Control inputs*. For the *Training net* the inputs are *Pilot* inputs and the outputs are the *Control inputs* from the *Inverse net*. If permitted the *Mediator* will train the *Training net* over the full space of possible *Pilot* inputs and Control inputs. When the discrepancy between the *Sensed performance* and the *Ideal performance* is lower than it was after the damage occurred the *Training net* and *Online net* are swapped. The new Online net now has a better performance than the previous Online net and the whole process can begin again with the training net starting from an initial state which was the outcome of the previous round of training.

Using a genetic algorithm for the inverse kinematics.

One major problem with the above idea (and many others for direct inverse control) is that the inverse kinematic problem is frequently ill posed, i.e. there are many control inputs which might produce a given *Sensed Response*. Worse, in general, if any plant state has an inverse image that is a non-convex region in control space, then there exist sets of samples in the region whose average is *outside* the region. Therefore any learning procedure for the *Inverse net* whose effect is to average over conflicting control signals (network outputs) for a given system state (network inputs) would be incorrect. Whilst such methods of directly learning inverse kinematics may be very effective in particular situations they cannot stand as a general methodology. Another potential pitfall is that, even for smooth dynamic systems, suitable solutions to the inverse kinematics are often *discontinuous* functions of time. This is illustrated by many of our simulation results for the attitude control problem. If the initial angular velocities are large the controller tends to produce very rapidly varying control torques, despite the fact that the progression towards the goal state is quite smooth. As the angular velocities become smaller the control torques tend to become much smoother. Whilst this does not entirely rule out the possibility that a neural network might learn such a functional transformation, it does make it very much harder to efficiently train such a network.

- For this reason we propose an alternative two part architecture, of which the first part is a LPN neuromodel that learns to predict a future system state given the few last states and control inputs.

Earlier results on neuromodelling [Dracopoulos 1993] indicate that this is a realistic prospect. The time scale for such a prediction might be as high as 0.5 - 1 sec. However, if the LPN were implemented in hardware then, once trained, it can output many thousands of evaluations (predictions) of the effect of a *hypothetical* control input. In practice, for the system under discussion around 5000 evaluations suffice to provide effective control.

- With this predictive capability we can now evaluate a large number of hypothetical control inputs and select the best. We do not do this directly but use a *genetic algorithm* to explore the space of hypothetical control inputs at any given moment and the LPN predictor to evaluate any member of the current population (of hypothetical control inputs).

Normally the evaluation phase (computing the fitness of a member of the population) is the most time consuming aspect of the genetic algorithm, but we propose to use a hardware neural network for this phase, thus allowing the evaluation of many thousands of sets of control inputs over the relevant prediction interval which may be a significant fraction of a second.

Of itself the GA is very simple. The control signals are represented by binary strings, with simple bit string manipulations for crossover and mutation. A variety of implementations could be used for the GA ranging from execution by a serial processor, to gate arrays with additional memory and processors (to provide a stack, for example). Given that the evaluation time per member of the population is very fast, the rest of the GA is quite simple. The particular implementation chosen would depend heavily on the system to be controlled and the speed of events in the real world.

We call this mixture of techniques a *neuro-genetic control architecture*. The learning phase of the neural network can be likened to basic mastery of motor control to the extent that one can predict the immediate consequences of any given set of actions. The genetic component of the system is analogous to a series of 'mind experiments', using many of these predictions to choose an actual set of control signals.

The sequence of events is thus (see Figure 3), t = k Δt (with Δt = 0.01 in the simulations):

At step *k-1*.	Controller applies $\boldsymbol{u}^*(k-1)$
	LPN predicts $\boldsymbol{x}^q(k)$
At times between *k - 1* and *k*.	Based on this prediction the GA tries to choose $\boldsymbol{u}^*(k)$ so as to optimise $\boldsymbol{x}(k+1)$
	The GA uses LPN to predict the result $\boldsymbol{x}^q(k+1)$ of hypothetical control inputs $\boldsymbol{u}^h(k)$ and hence evaluate the fitness of $\boldsymbol{u}^h(k)$.
At step *k*.	Controller applies $\boldsymbol{u}^*(k)$
	LPN predicts $\boldsymbol{x}^q(k+1)$
	...

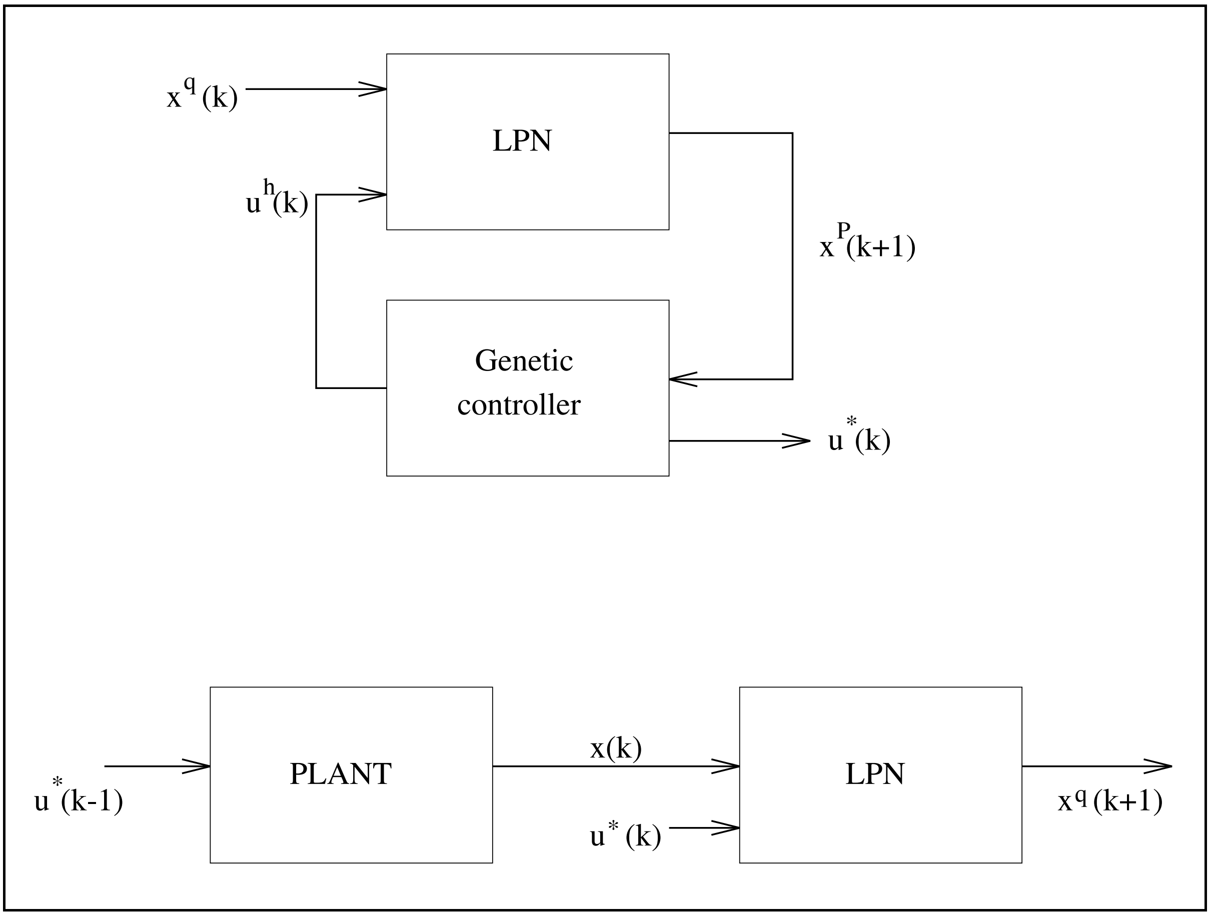

Figure 3 The Control Architecture. Top: Between k-1 and k the LPN uses $\mathbf{x}^q(k)$ and hypothetical control inputs $u^h(k)$. Bottom: When t = k Δt, $\mathbf{x}(k)$ is known and the LPN predicts $\mathbf{x}^q(k+1)$.

A locally predictive net for the attitude control problem.

As it was shown in [Dracopoulos 1993], using relatively few data points artificial neural networks have the capability to construct accurate local models of complex dynamic systems over a large region of the state space. This capability of local predictive networks (LPN) is essential for the Adaptive Control Architecture proposed in the previous section. If $\boldsymbol{x}(t)$ is the state vector of the system, such networks learn an input-output function of the form

$$((\boldsymbol{x}(t),\ \dot{\boldsymbol{x}}(t)),\ (\boldsymbol{x}(t-\Delta t),\ \dot{\boldsymbol{x}}(t-\Delta t)),\ \ldots,\ (\boldsymbol{x}(t-p\Delta t),\ \dot{\boldsymbol{x}}(t-p\Delta t))) \ \rightarrow\ (\boldsymbol{x}(t+\Delta t),\ \dot{\boldsymbol{x}}(t+\Delta t)) \qquad (6)$$

for some fixed $p \geq 1$.

In this section another example of local prediction using LPN networks is given. The Euler equations are again considered. This time the constructed LPN will have as additional inputs the control thrusts of the satellite. The control variables differ from the state vectors in that they are not normally dynamically deterministic (unless we impose a control law). The utility of the network is that, based on relatively few training samples, in any given system state and for *any* choice of control variables it can quite accurately predict the next system state. The next example is a simple illustration of the capability of such a network (but does not constitute an exhaustive test), see also [Dracopoulos 1993].

The values of the dynamic system were chosen to be $I_x = 3$, $I_y = 2$, $I_z = 1$. A LPN was trained to predict ahead for $\Delta t = 0.5$, using 501 data taken from the trajectory starting at $(\omega_1, \omega_2, \omega_3) = (3, 2, 1)$ with the thrusts given by the deterministic rule

$$\begin{pmatrix} G_1 \\ G_2 \\ G_3 \end{pmatrix} = \begin{pmatrix} -1.2 & 0 & \frac{\sqrt{6}}{2} \\ 0 & 0.35 & 0 \\ -\sqrt{6} & 0 & -0.4 \end{pmatrix} \begin{pmatrix} \omega_1 \\ \omega_2 \\ \omega_3 \end{pmatrix} \qquad (7)$$

Thus, the training data were chosen from the chaotic trajectory described in [Dracopoulos 1993]. The network architecture was 24-20-10-3, since a value of $p = 3$, in equation (6) was chosen, and the network had 12 inputs describing the current and (three) past states of the system and another 12 inputs describing the current and (three) past control inputs to the system.

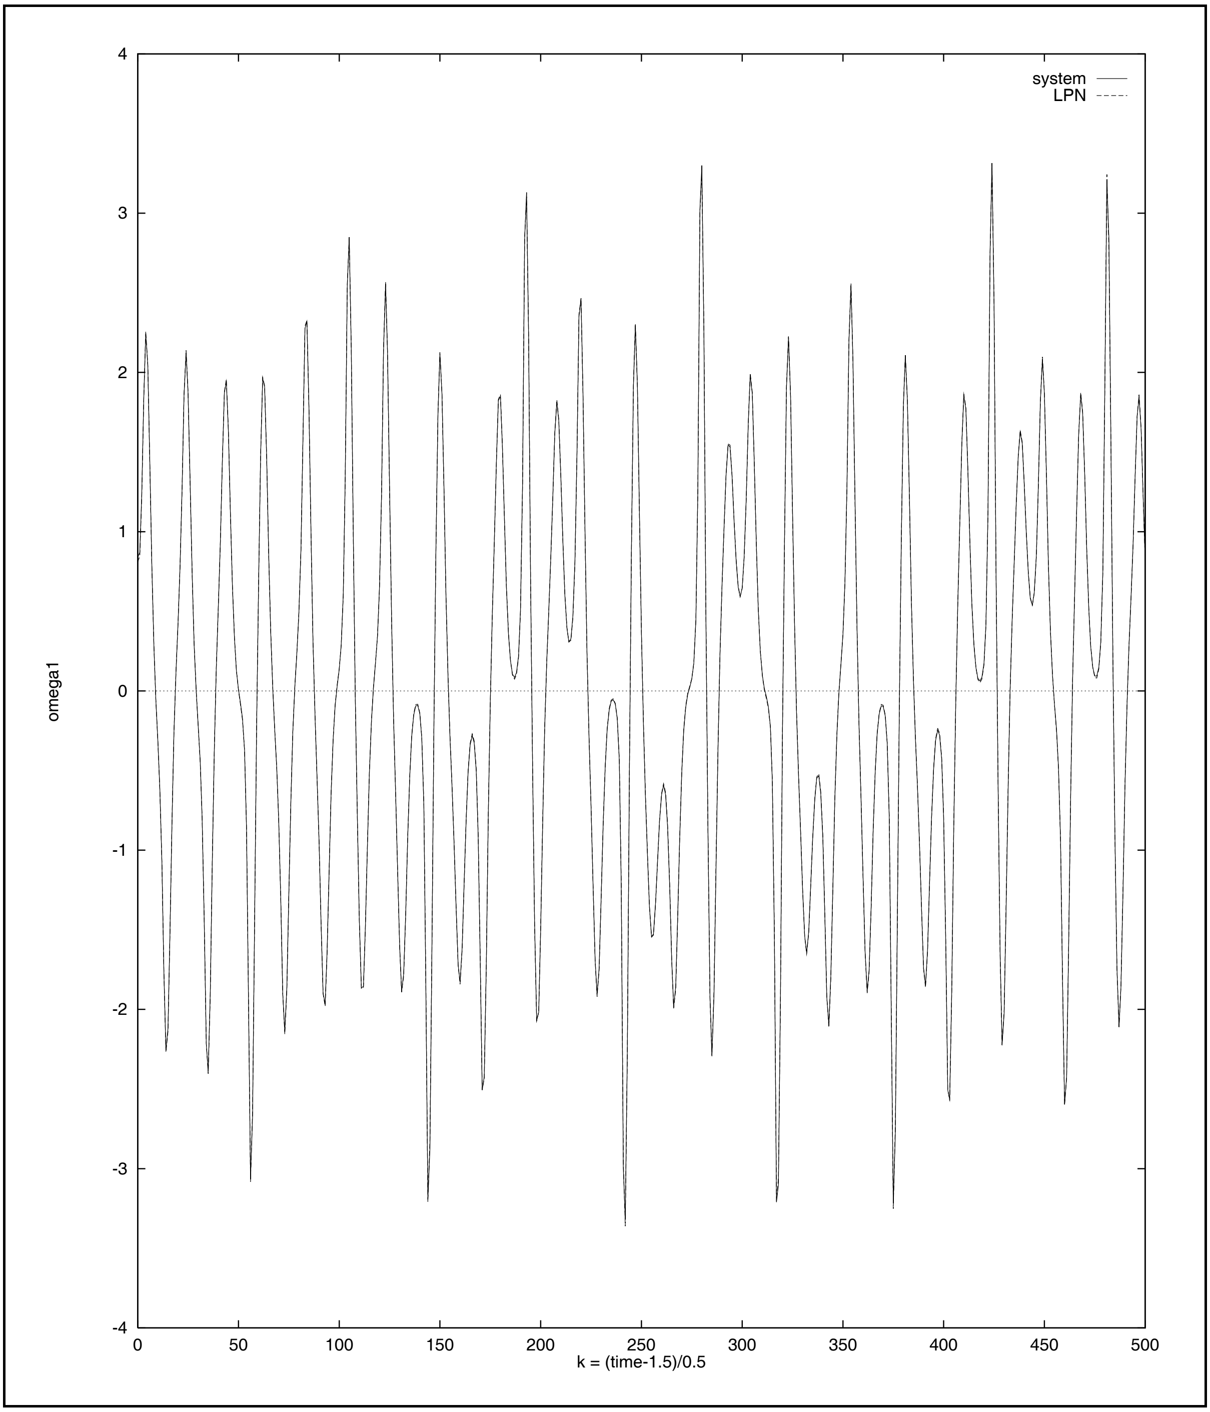

Figure 4 Euler system. How well the 24-20-10-3 LPN learnt the training data: x-angular velocity. Network predicting 0.5 sec ahead.

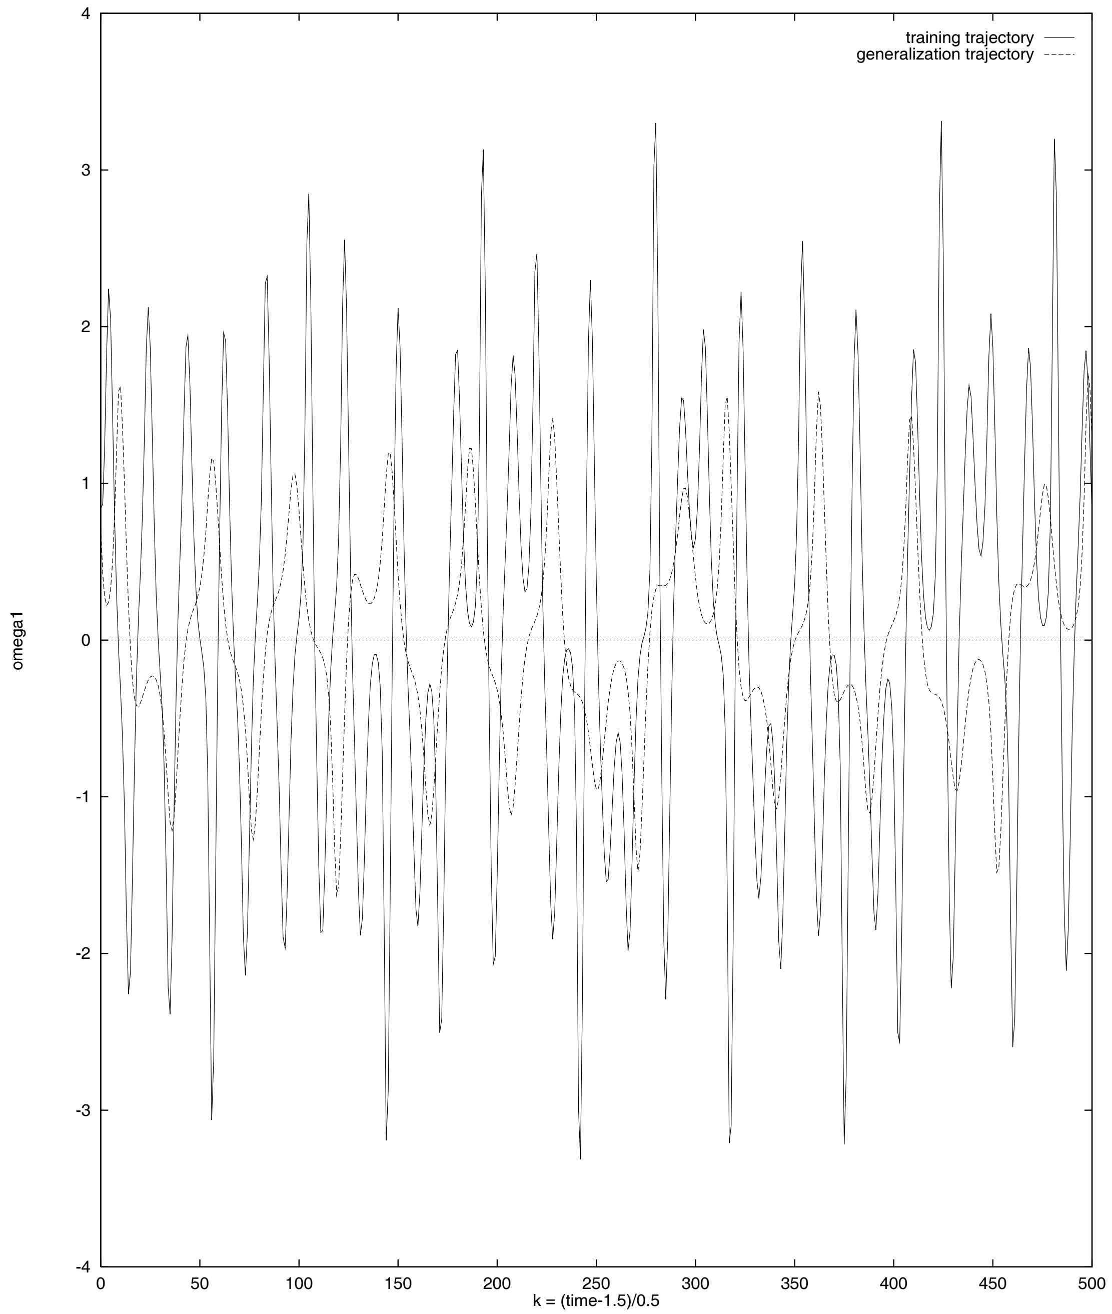

Figure 5 Euler system: x-angular velocity. The figure shows the both the trajectory used for training and the trajectory used for a test of the trained LPN for generalization.

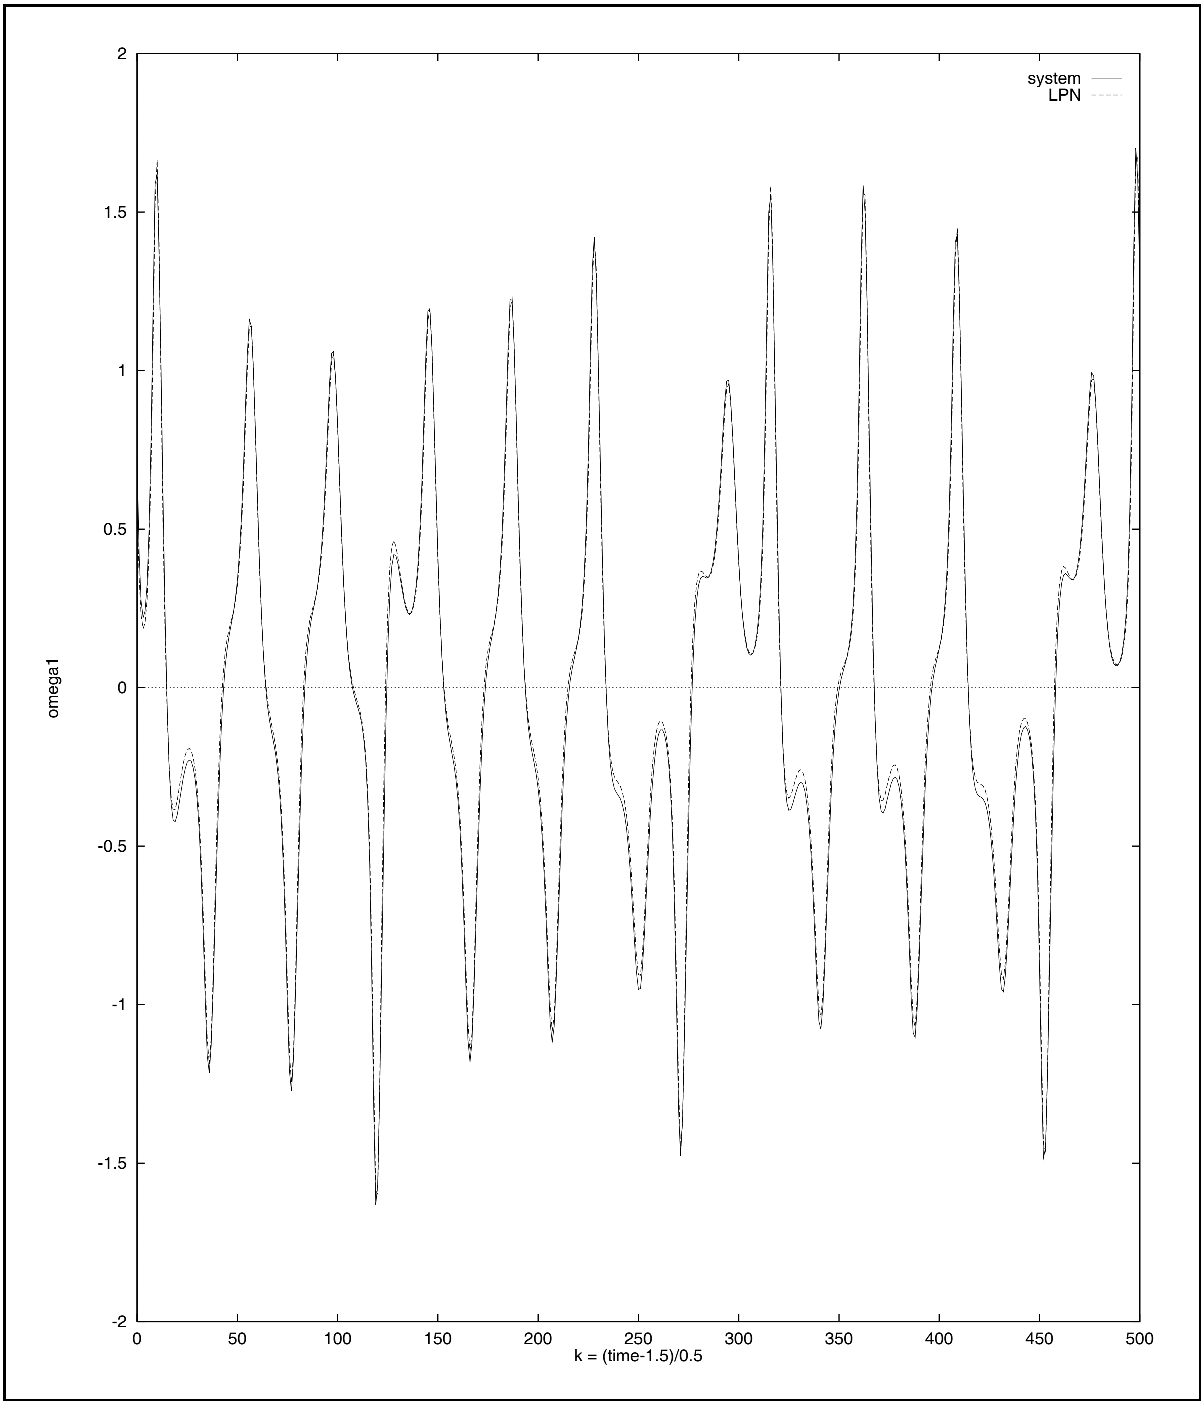

Figure 6 Euler system: x-angular velocity. Generalization of the 24-20-10-3 LPN on the trajectory starting at (3, 2, 1) with (G_1, G_2, G_3) determined by (8). Network predicting 0.5 sec ahead.

The training convergence of the above LPN (angular velocity ω_1) is shown in , while the generalization capability of the network was tested in a trajectory started at (ω_1, ω_2, ω_3) = (3, 2, 1) with the thrusts given by the deterministic rule

$$\begin{pmatrix} G_1 \\ G_2 \\ G_3 \end{pmatrix} = \begin{pmatrix} -1.5 & 0 & 0.5 \\ 0 & 0.2 & 0 \\ -2 & 0 & -0.3 \end{pmatrix} \begin{pmatrix} \omega_1 \\ \omega_2 \\ \omega_3 \end{pmatrix} \tag{8}$$

which is different from the rule (7) used in training. Thus the resulting trajectory is completely different from the training one (see Figure 5). The result of the generalization test is shown in Figure 6. These results extend and confirm those obtained earlier [Dracopoulos 1993].

It should be noted that for the Neuro-Genetic adaptive control architecture proposed in the previous section, and applied throughout the remainder of this paper, a prediction interval of $\Delta t = 0.5$ sec is sufficient to buy considerable computation time. Indeed since the time step over which the satellite operations occur (time at which the control thrusts of the satellite can change) is $\Delta t = 0.01$, prediction is actually required for only 0.01 sec ahead. Thus a prediction interval of $\Delta t = 0.5$ is 50 times larger than what is required, but still the accuracy is good as can be seen from the diagrams.

In the simulation results that follow, every time the system dynamics change in an unknown way, *it will be assumed* that a neuromodel describing the new dynamics of the plant has been trained according to this methodology. In practice, since vanilla backpropagation converges rather slowly and sometimes unreliably, this would require a hardware implementation of backpropagation, or some other method for rapidly learning a dynamic model. However, we put this issue on one side for the present and next investigate the viability of the genetic controller.

- In order to separate the two issues the experiments described for the genetic controller will use data (i.e. evaluations of hypothetical control signals) provided by the simulator rather than the neuromodel.

The genetic algorithm.

The function of the genetic algorithm is to optimise the choice of control thruster torques at any given moment. This amounts to a time constrained locally optimal solution to an inverse kinematic problem. A fast introduction to genetic algorithms and extensive references to the existing literature can be found in [Jones 1993] and we shall simply describe the necessary features of the present application.

Suppose, for example, that the target state for the system is $\varphi = 0$, $\dot{\varphi} = 0$, $\theta = 0$, $\dot{\theta} = 0$, and the attitude and angular rotation about the third axis are unspecified. With this target in mind, one possible choice of goal is to use the GA at any given moment to choose the control torques (G_1, G_2, G_3) so as to minimise the function

$$F(G_1,\ G_2,\ G_3) = |\varphi| + |\dot{\varphi} + \varphi| + |\theta| + |\dot{\theta} + \theta| \tag{9}$$

at the next sensor sample. Introduction of terms like $|\dot{\varphi} + \varphi|$ is motivated by two factors. First, if the angular velocities are initially high these terms will dominate, and the emphasis of the corresponding control actions will be mostly aimed at reducing them. Second, as the system approaches the target state these terms will act as damping so as to avoid overshooting the target state. Many other choices which accomplish the same goals are possible.

The GA will attempt to do this by maximizing the fitness function

$$v(G_1,\ G_2,\ G_3) = \frac{1}{1 + F} \tag{10}$$

Although different fitness functions based on an objective function have been used in the GA literature, all the simulations in this work showed that the choice of (10) always leads to good solutions (provided that the other GA parameters such as population size, mutation rates, etc. are chosen carefully). It is apparent in the simulations that this particular choice of fitness (which has non-zero derivative at F = 0) has the benefit of exaggerating the effect of small differences in the value of F, as F approaches zero (which happens in later generations). Thus, as the population of hypothetical control inputs improves, greater emphasis is placed on the small variations which make the difference between a good control individual and a very good one.

The control torques are each encoded as 16 bits. Allowing one bit for the sign this gives a range -32768 to +32768 for the magnitude of each torque, see Figure7. The string of 3X16 = 48 bits is taken as the genotype and a simple one point crossover or mutation is used on these bit strings.

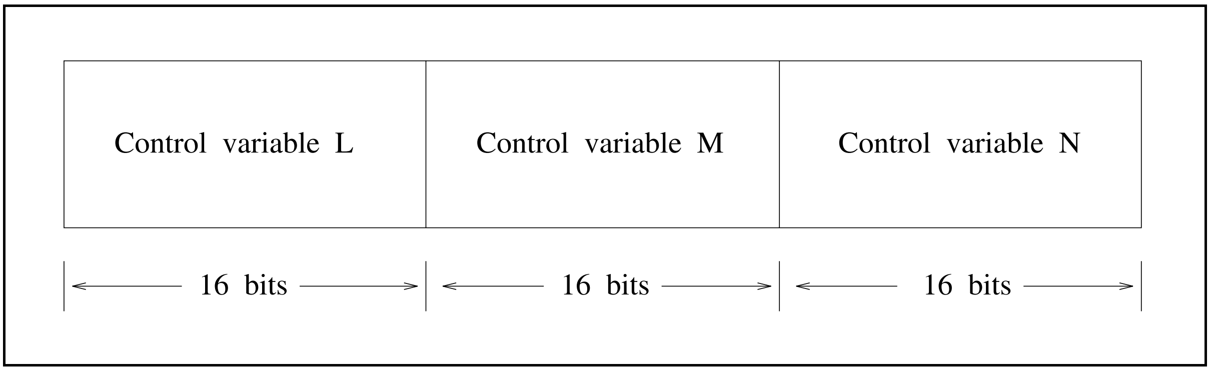

Figure 7 Encoding the control torques as a genotype.

Pseudocode for the GA used is given in Algorithm 1. A population size of 50 was used and the algorithm terminates after running for 100 generations (4802 individuals evaluated). The initial population is generated randomly and the main loop runs for 100 generations. During each iteration of this loop a new population is generated with 50 new individuals. These 50 consist of the two best from the previous generation, and 48 other individuals that are produced with the probabilistic application of the 1-point crossover and mutation. Each pair of individuals is selected for crossover based on a probability proportional to their fitness. The other parameters used by the GA were: probability of crossover 0.6 (otherwise simple replacement is used) and the resulting genotype is mutated with probability 0.0333 per bit [Grefenstette 1986].

```
population_size := 50
generation := 1
Initialize population with random binary strings
while generation ≤ 100 do
        Find the two best individuals of current population
        for i := to population_size - 2 step 2 do
                Select two individuals a, b based on fitness
                Probabilistically perform crossover
                if crossover_performed then
                        Copy the two offspring into new population
                else
                        Copy the two (mutated) individuals into new population
                endif
        endfor
        generation := generation + 1
endwhile

Procedure Select
        j := 1; sector := 0
        number := random()
        while sector ≤ number and j ≤ population_size do
                sector := sector + fitness_of_individual[j]
                j := j + 1
        return individual[j-1]
```

Algorithm 2 Pseudocode of the Genetic Algorithm used for the Attitude Control Problem.

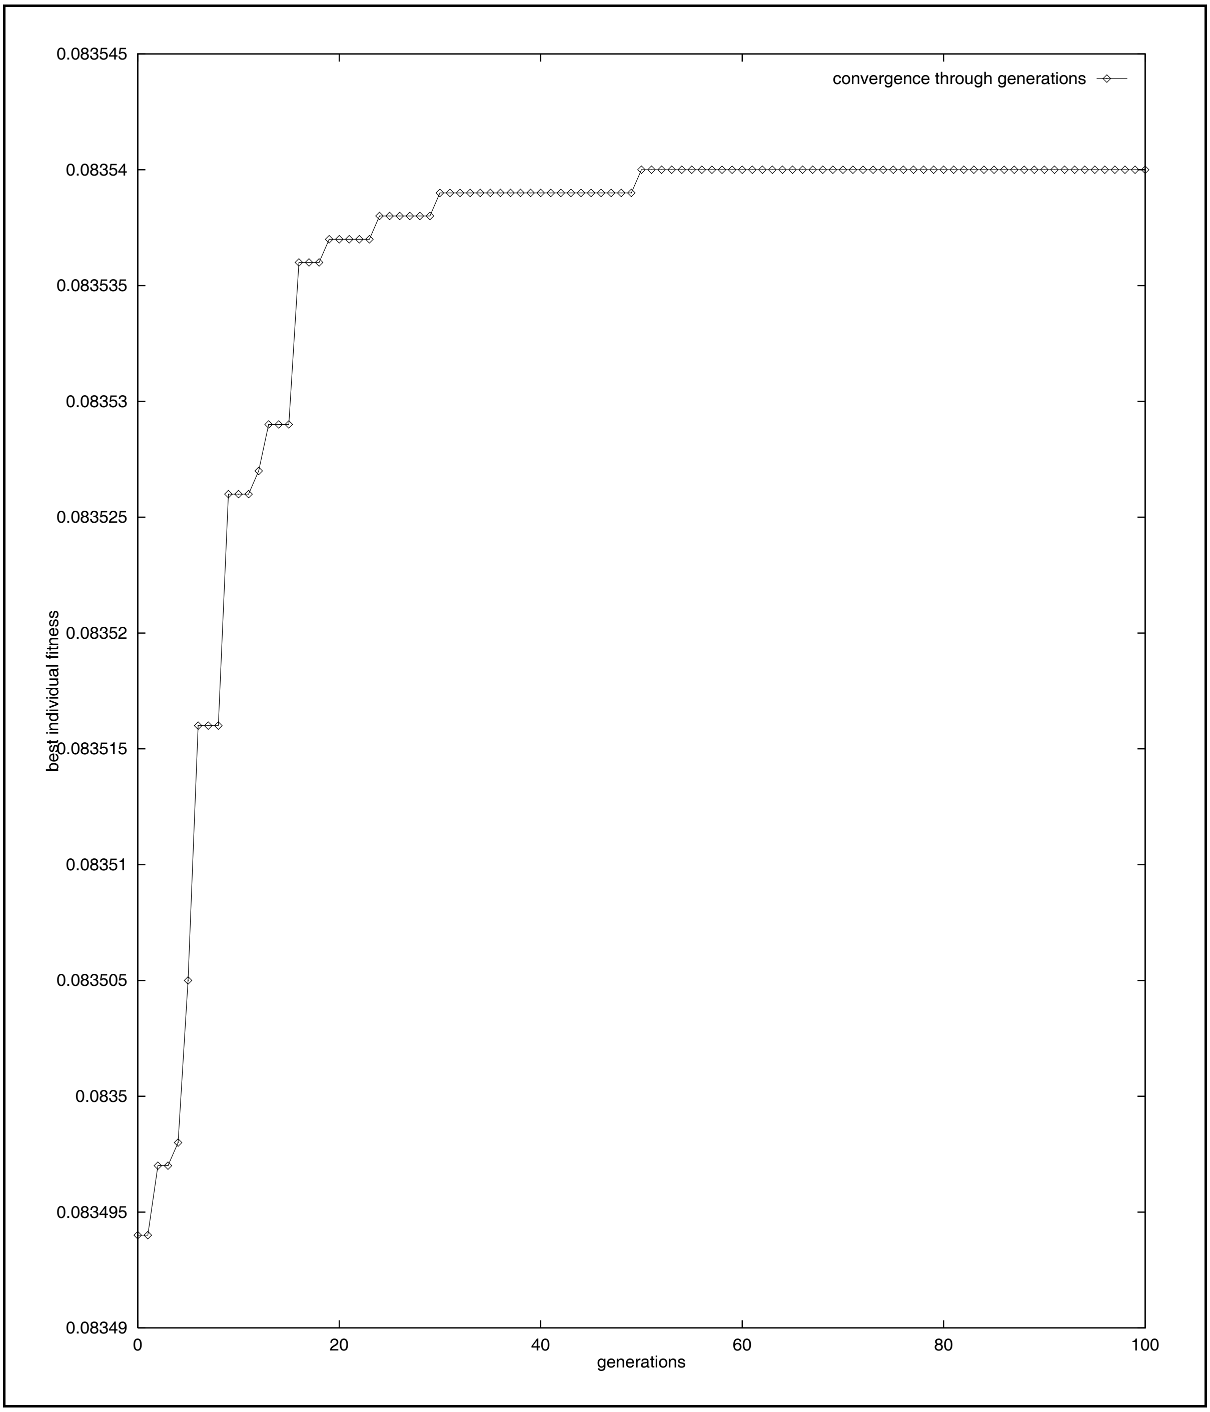

Figure 8 The evolution of the solution to the sample problem.

In Figure 8 the evolution towards a solution to the problem over 100 generations is shown. In this case the initial conditions at time $t = 0$ are $\omega_1 = 1.75$, $\omega_2 = 2.9$, $\omega_3 = 1.33$ and $\varphi = 2$, $\theta = 1$, $\psi = 3$, whilst $I_x = 1160$, $I_y = 23300$, $I_z = 2400$. Although in this particular case the final result was found in only 50 generations, in the general case experiments showed that a maximum of 100 generations was desirable.

The number of possible individuals in the global search space is $2^{16} \times 2^{16} \times 2^{16} = 2.81 \times 10^{14}$. The genetic algorithm processes only $48 \times 100 + 2 = 4802$ individuals between each application of a set of control torques, i.e. about 1.7×10^{-9} % of the possible control combinations. However, in all the simulation results run the genetic algorithm was able to find a 'good' solution of the problem whilst processing this small proportion of the possible gene strings. This point is very important for the real-time application of the proposed neuro-genetic architecture.

Directly estimating the effectiveness of this genetic algorithm is rather difficult because the true optimum is unknown. In principle it is possible to obtain an upper bound for the achievable fitness from energy considerations: given the fact that available thruster torques are bounded (and no external torques are acting upon the body) only system states in a certain region about the initial state are in fact reachable. However, the true test of the genetic algorithm is whether it produces sufficiently good solutions to achieve successful control and, as we shall see, by this test the genetic algorithm performs very well.

Spin stabilization of a satellite about a stable axis.

Assume, that for some unknown reason (damage), a satellite with specified dynamics, changes its characteristics so that the moments of inertia become $I_x = 1160$, $I_y = 23300$ and $I_z = 24000$.

During the period where the system dynamics change, unknown forces lead it to the state $(\omega_1, \omega_2, \omega_3) = (3, 2, 1)$, $(\varphi, \theta, \psi) = (2, 1, 3)$. The goal is to detumble the satellite about the *x, y* body axes, spin it about the *z* body axis, and reorient it, so that $\varphi = 0$, $\theta = 0$.

The application of the genetic adaptive control architecture, described by Figure 3, with the objective function given by (9), leads to the situation described by Figure 9 - Figure 14.Figure 9 shows the evolution in time of the angular velocities ω_1, ω_2 about the *x, y* body axes respectively. The genetic controller soon leads both to the prespecified value of zero. The third angular velocity is arbitrary, since it was unspecified in the control objectives.

Figure 10 shows the reorientation of the satellite for the angles φ, θ during the application of the genetic controller. While these are becoming zero, the satellite is rotating about its *z* axis, Figure 11. It should be noted that the controller not only leads the system to a desired state, but it maintains this state afterwards.

The applied thrusts during the genetic control of the satellite are shown in Figure 12, Figure 13, and Figure 14. We observe that during the time that the angular velocities are large the thrusts vary very rapidly, so that in some situations they can be seen as applying a kind of *bang-bang control*. As soon as the angular velocities obtain small values, the required and applied thrusts G_1, G_2, become 'smooth'. The third thrust G_3 is not subject to evolutionary pressure near the target state and consequently has no incentive to become small. This is a result of the particular choice of objective function in this case - the target state is not fully specified.

We speculate that to some extent, at least for the attitude control problem, the discontinuous nature of the control solutions found by the genetic algorithm may be an artifact of the particular choice of objective function. The description of orientation in terms of φ, θ, ψ is modulo 2π and the objective function takes no account of this fact. Discontinuities might therefore be expected, particularly for large angular velocities. We regard it as a matter of considerable interest to investigate constraints on the genetic algorithm which may lead, where possible, to smoother solutions. Particularly because intuition suggests that highly discontinuous control solutions, with rapid sign changes of the torques, are liable to be energy inefficient.

Any method of direct neural inverse control is inevitably going to encounter this problem in some cases. Training a backpropagation network requires training data, and if this data is elicited from a discontinuous control strategy one can expect severe problems in successfully training the network. In contrast, our experiments suggest that an (unconstrained) genetic algorithm will find very good (in the sense that the target is acquired rapidly) solutions to the inverse kinematic problem, even if these form a discontinuous function over time. Any quest for a smooth, energy minimal, adaptive control strategy will, in any case, inevitably involve a tradeoff of energy consumption against time-to-target. We plan to address these issues in later work.

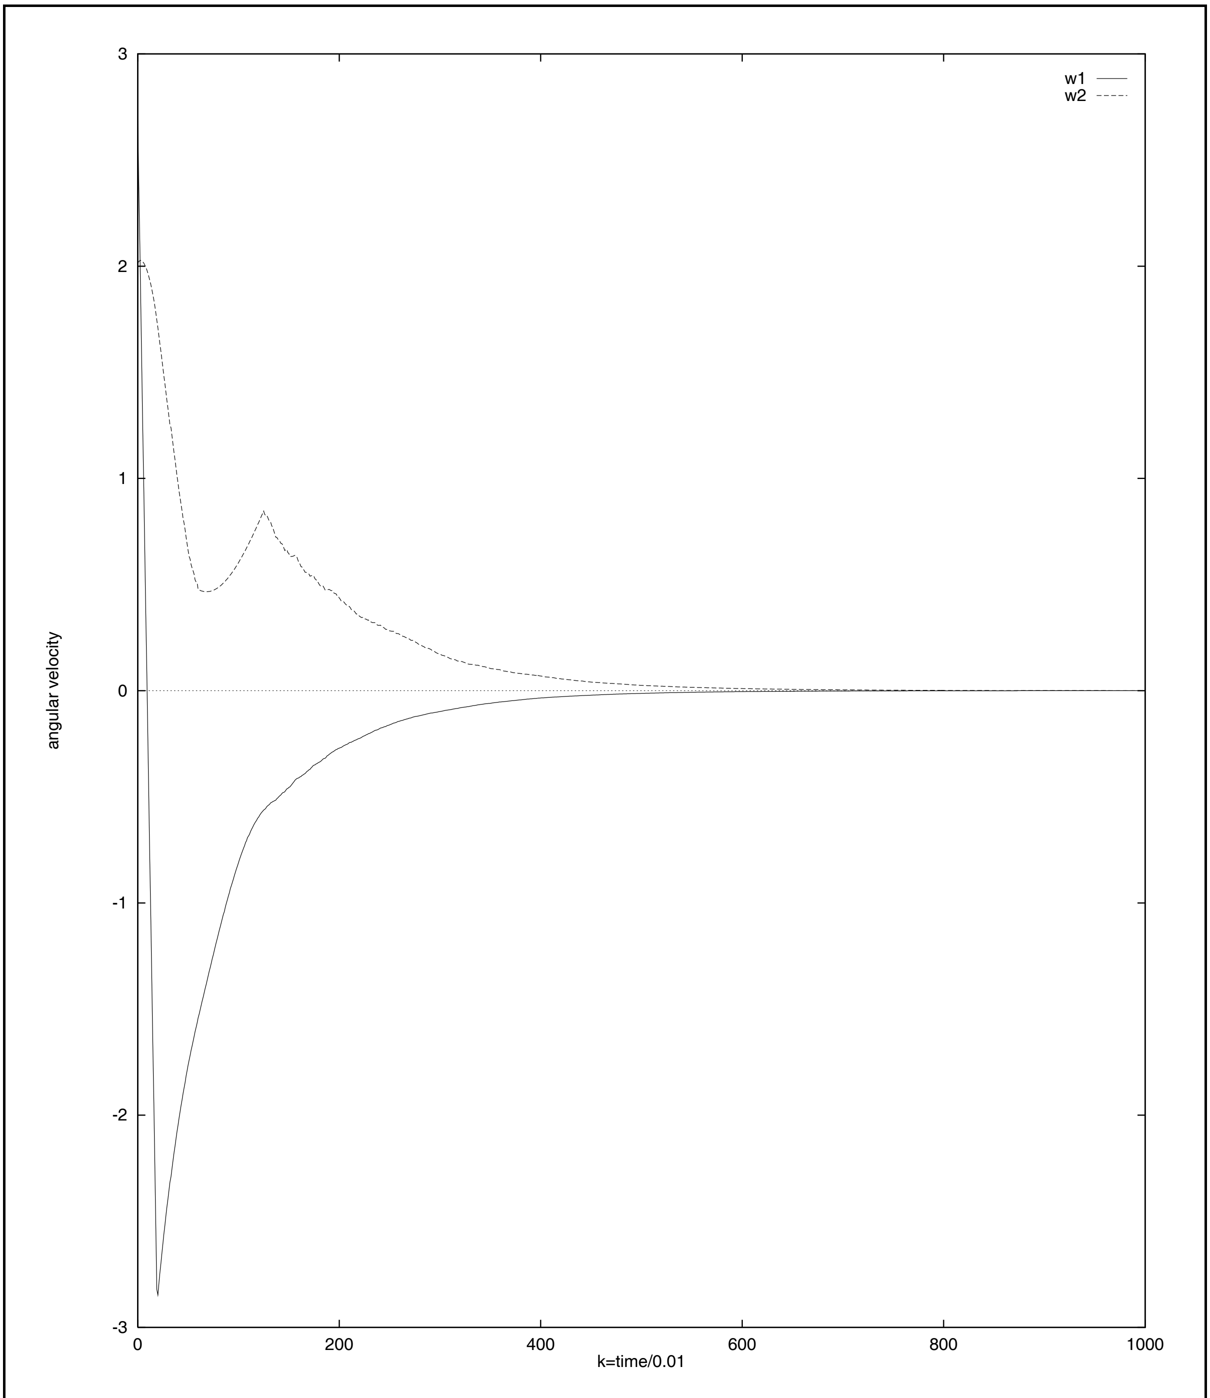

Figure 9 Angular velocities ω_1, ω_2 for the satellite during the application of the genetic controller. Initial conditions $(\omega_1, \omega_2, \omega_3) = (3, 2, 1)$ and $(\varphi, \theta, \psi) = (2, 1, 3)$.

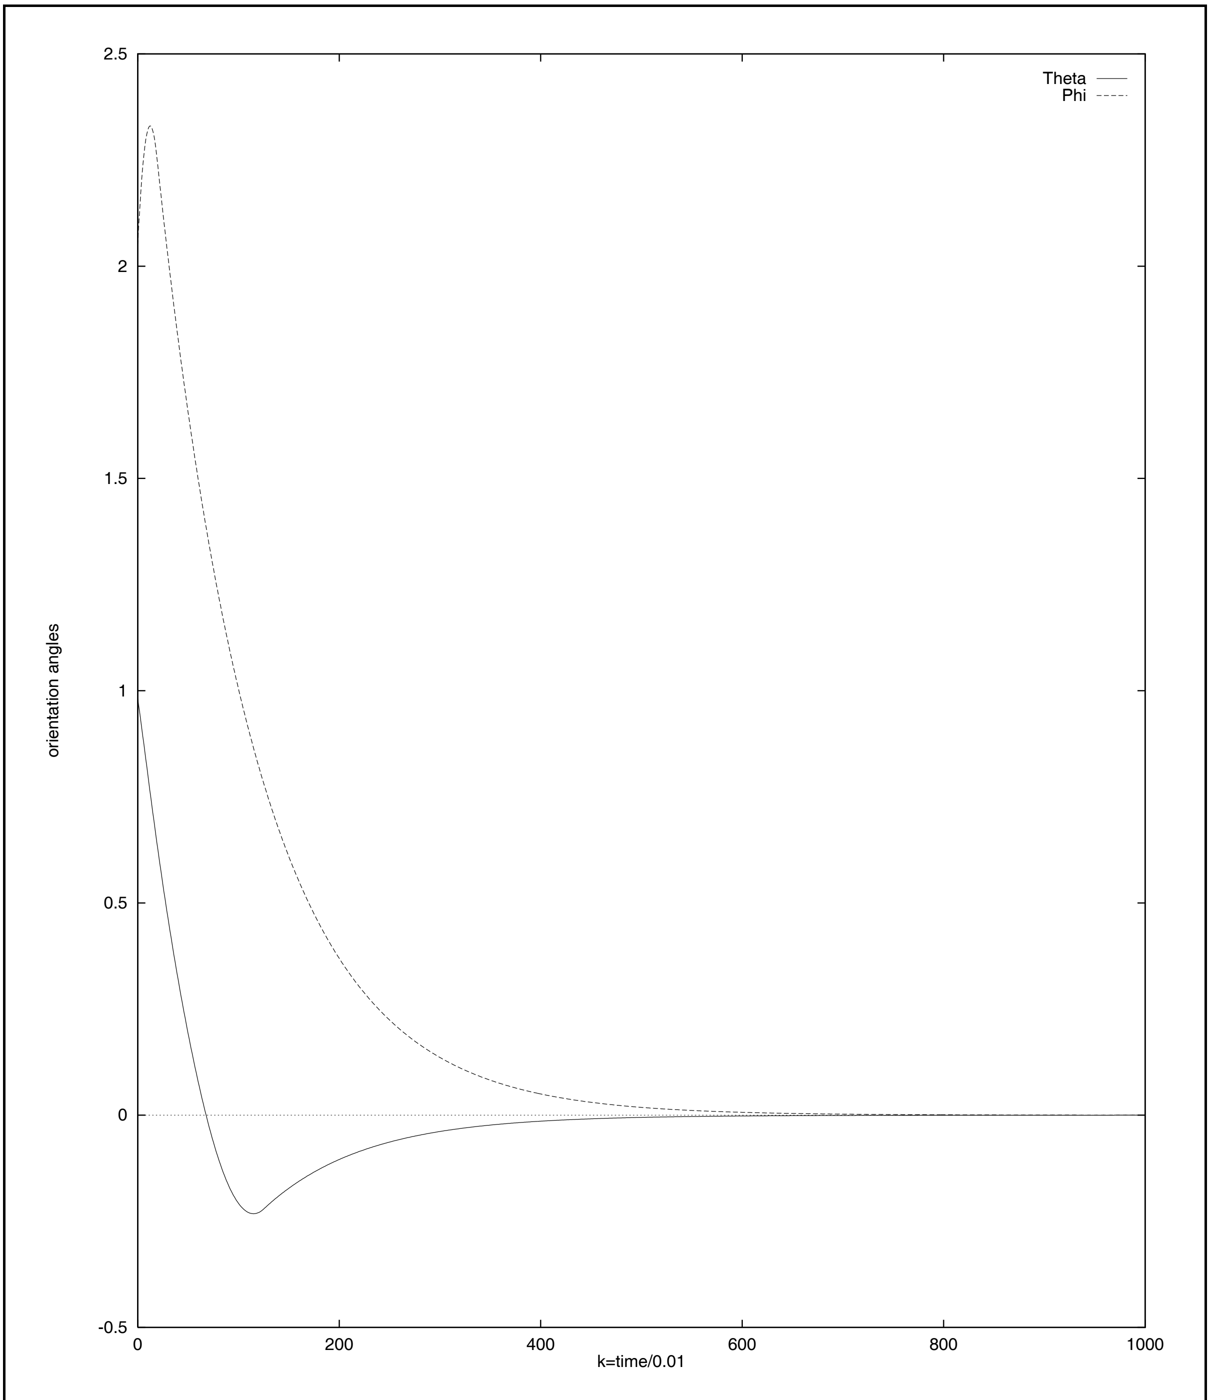

Figure 10 Orientation angles φ, θ for the satellite during the application of the genetic controller. Initial conditions (ω_1, ω_2, ω_3) = (3, 2, 1) and (φ, θ, ψ) = (2, 1, 3). Stable target.

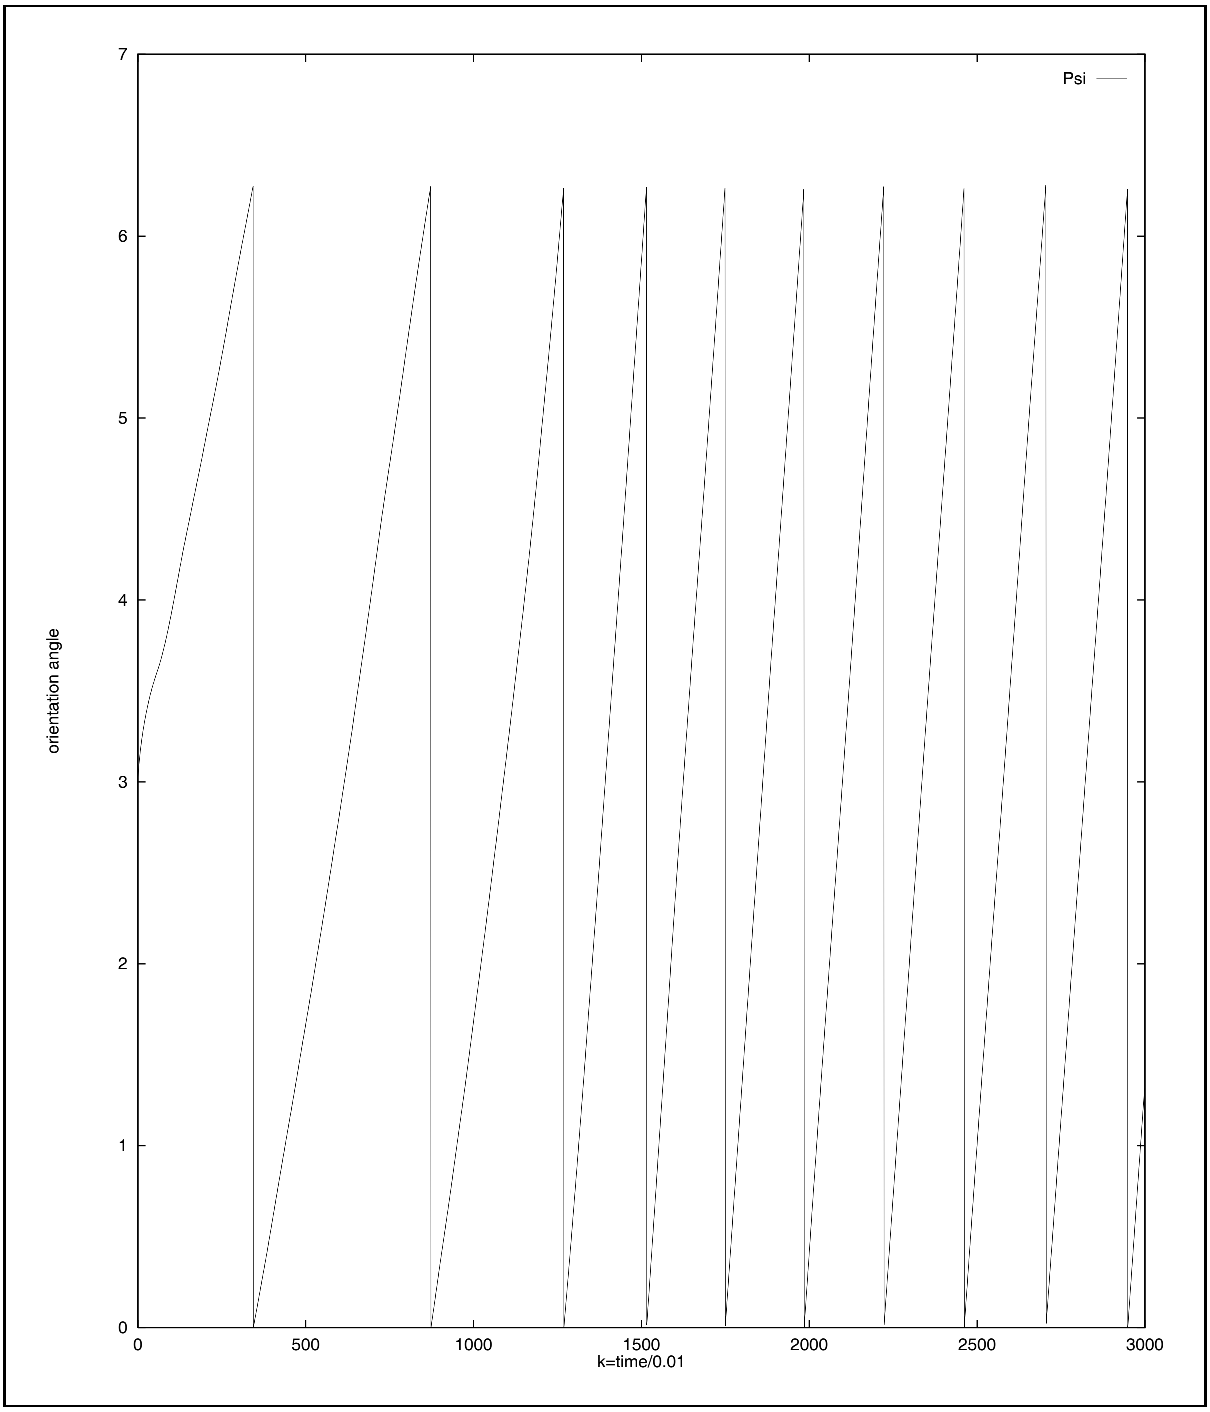

Figure 11 Orientation angle ψ during the application of the genetic controller. Initial conditions $(\omega_1, \omega_2, \omega_3) = (3, 2, 1)$ and $(\varphi, \theta, \psi) = (2, 1, 3)$. Stable target.

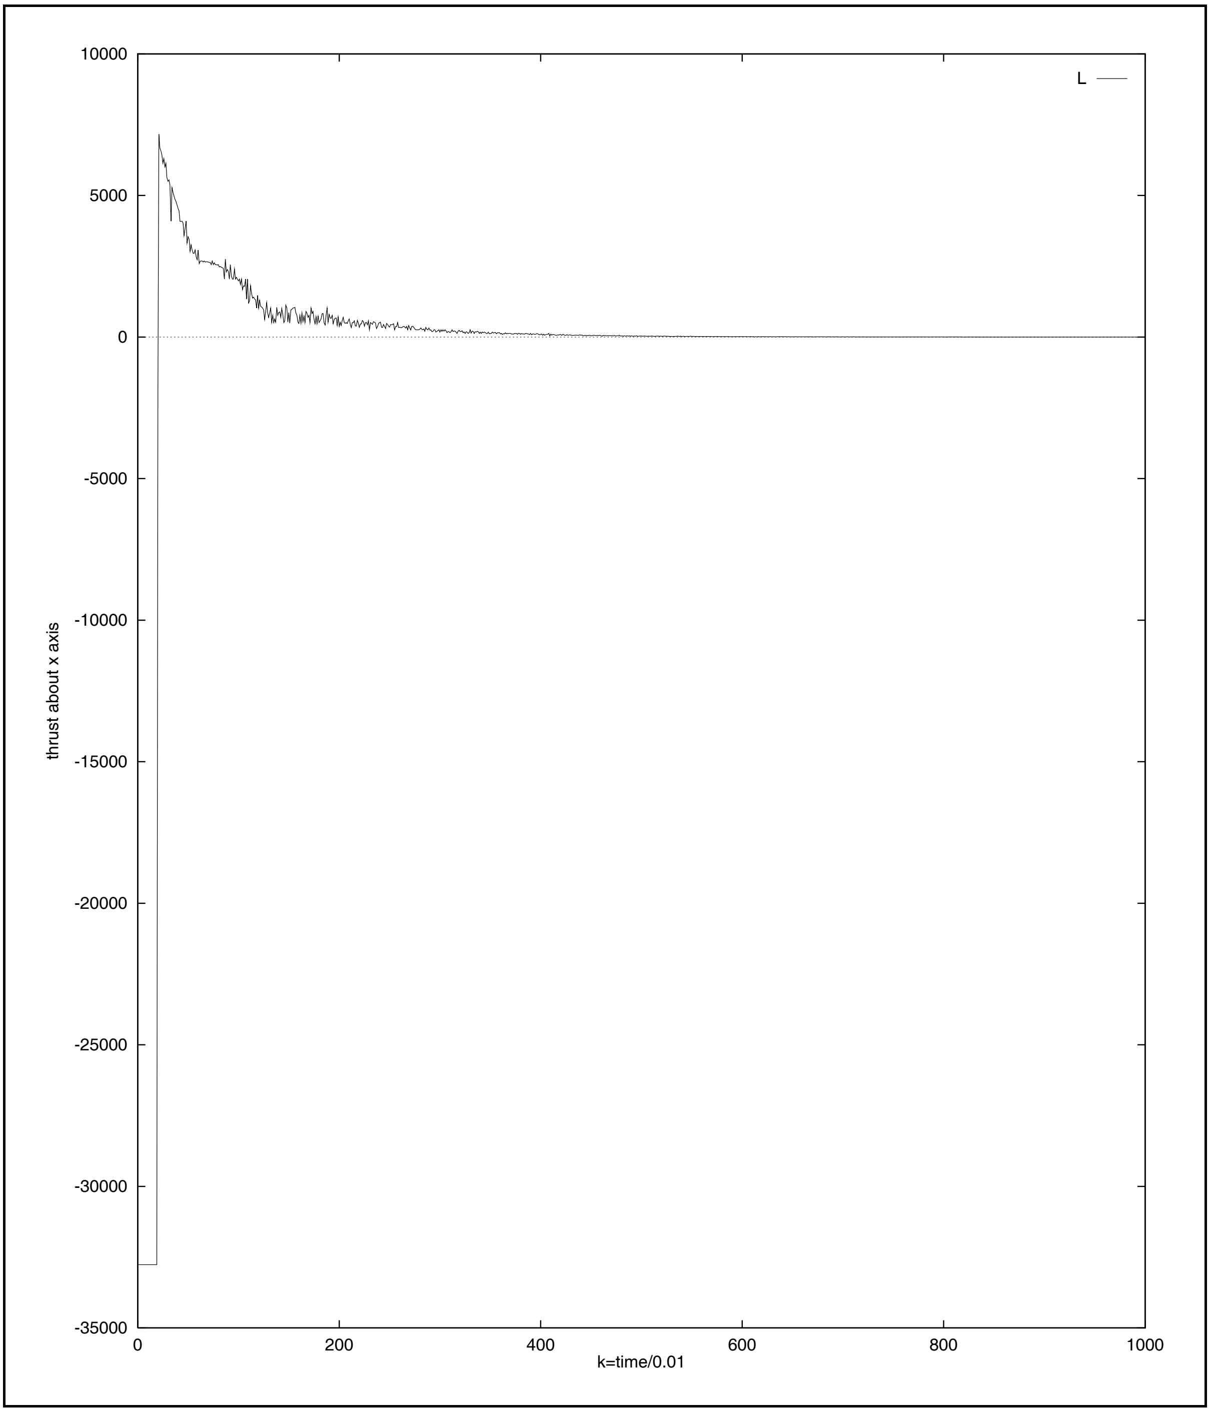

Figure 12 Applied thrust G_1 about the x body axis during the application of the genetic controller. Initial conditions $(\omega_1, \omega_2, \omega_3) = (3, 2, 1)$ and $(\varphi, \theta, \psi) = (2, 1, 3)$. Stable target.

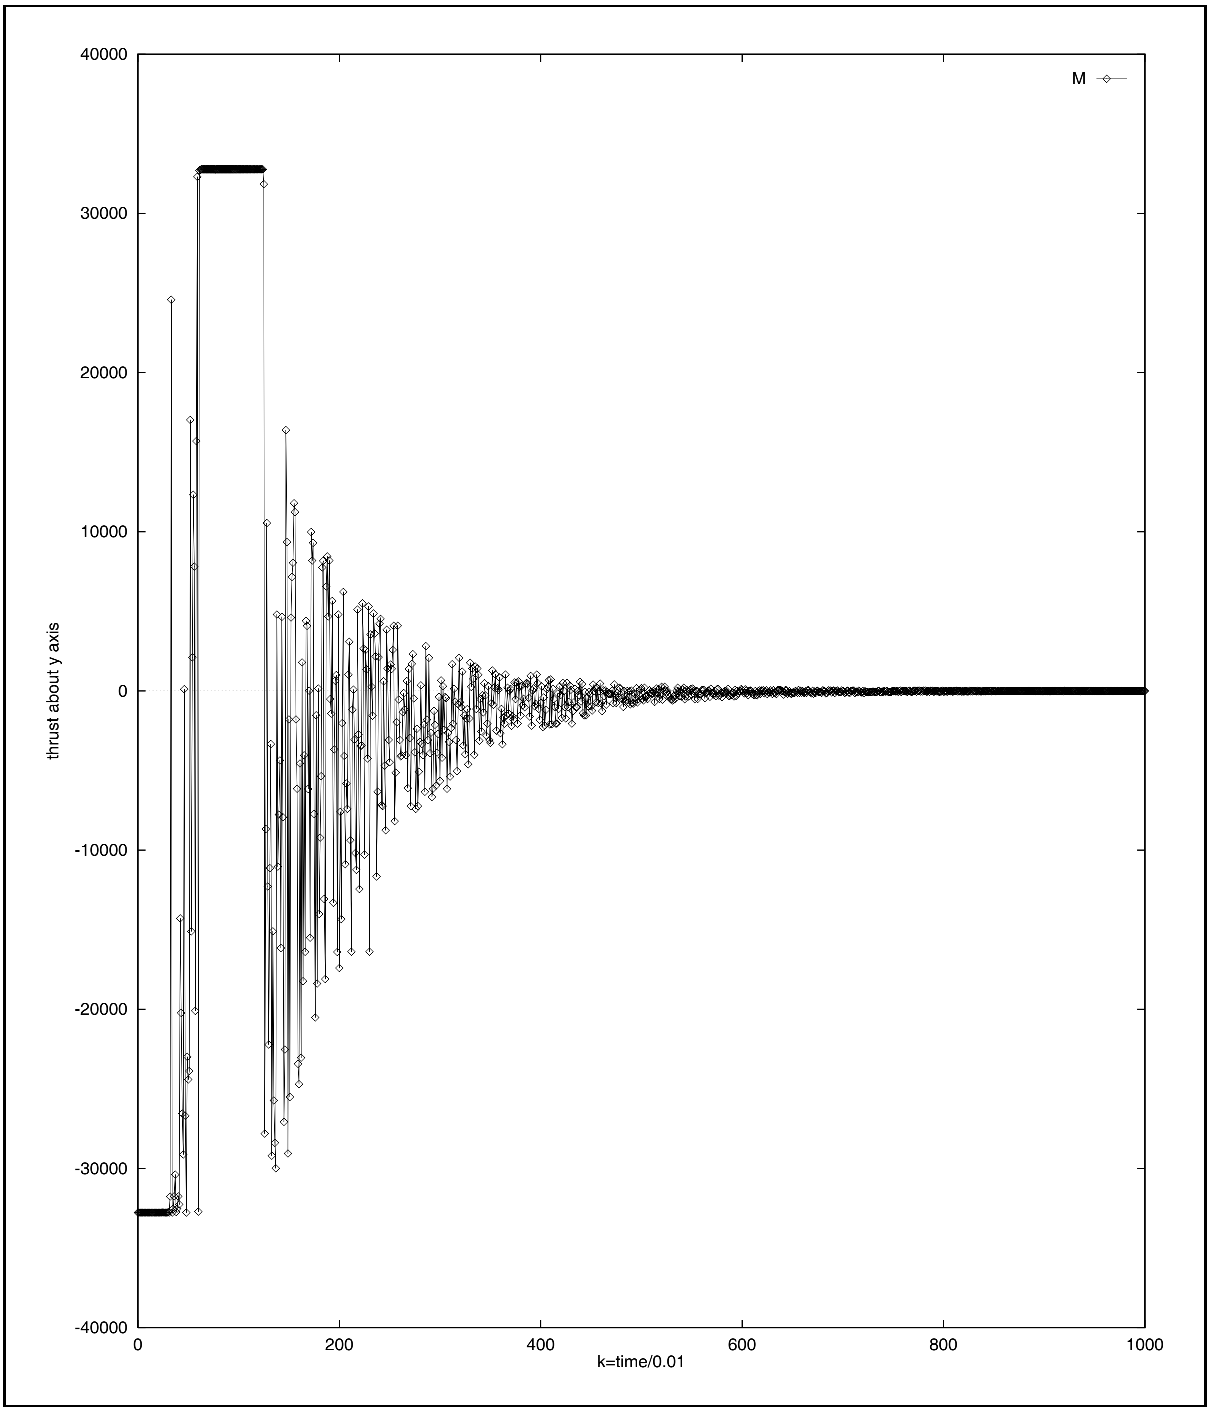

Figure 13 Applied thrust G_2 about the y body axis during the application of the genetic controller. Initial conditions $(\omega_1, \omega_2, \omega_3) = (3, 2, 1)$ and $(\varphi, \theta, \psi) = (2, 1, 3)$. Stable target.

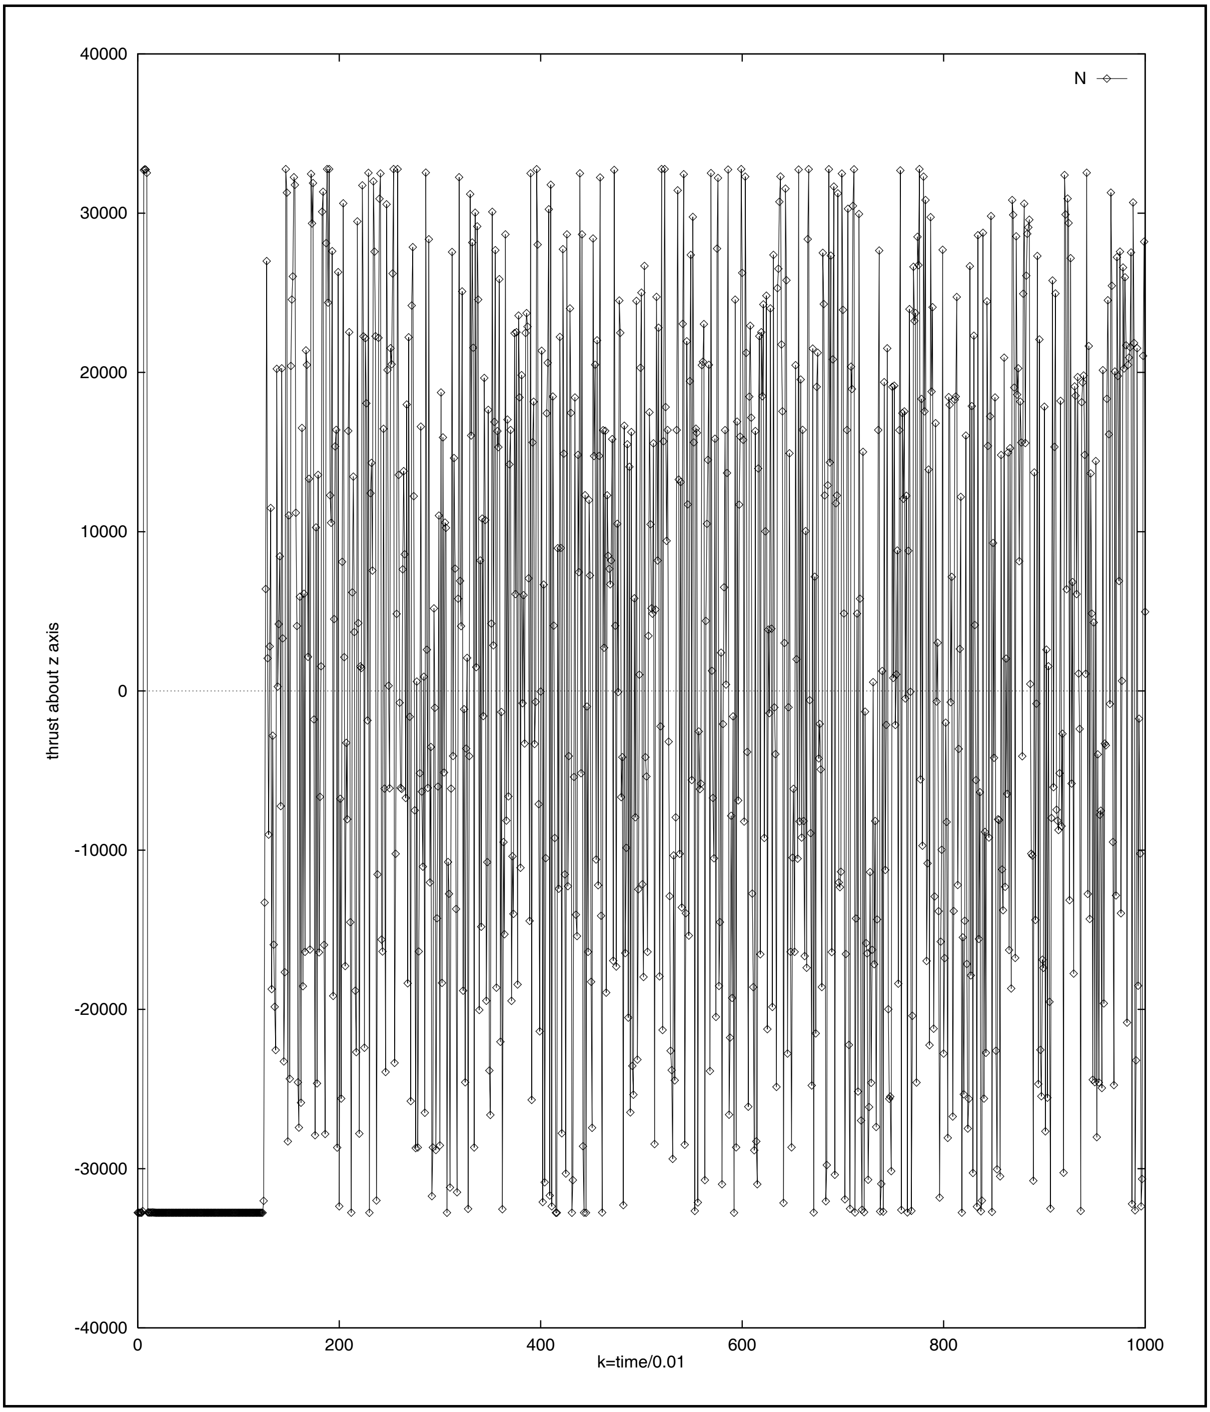

Figure 14 Applied thrust G_3 about the z body axis during the application of the genetic controller. Initial conditions $(\omega_1, \omega_2, \omega_3) = (3, 2, 1)$ and $(\varphi, \theta, \psi) = (2, 1, 3)$. Stable target.

Spin stabilization of a satellite about an unstable axis.

In this section the task of stabilizing a satellite (rigid body) about an unstable axis is considered. The moments of inertia of the satellite are $I_x = 1160$, $I_y = 24000$ and $I_z = 23300$. Simple rotation about the *z* body principal axis is *unstable* by the 'Tennis Racquet' theorem, since it corresponds to the intermediate principal axis $I_x < I_z < I_y$ [Arnold 1978], [Bender 1978]. Thus in this case if the body is rotating about its *z* axis, a very small disturbance may produce a very great change in the subsequent motion [Routh 1892], [Bondi 1986].

The goal in this simulation is to spin stabilize the satellite about its *z* (unstable) axis. Thus the target state of the system is set to $(\omega_1, \omega_2, \omega_3) = (0, 0, 1)$ and $(\varphi, \theta, \psi) = (0, 0, 0)$. According to this a fixed spin of $\omega_3 = 1.0$ is specified. The initial conditions are:

$$(\omega_1, \omega_2, \omega_3) = (1.7494, 2.9092, 1.3276) \text{ and } (\varphi, \theta, \psi) = (2, 1, 3).$$

The objective function for the Genetic controller is specified to be:

$$F(G_1, G_2, G_3) = |\varphi| + |\dot{\varphi} + \varphi| + |\theta| + |\dot{\theta} + \theta| + |\dot{\psi} - 1.0| \qquad (11)$$

Figure 15 shows the evolution in time of the angular velocities about the *x, y, z* body axes, respectively, during the application of the genetic controller. The controller soon leads two of them to the prespecified value of zero and the third to the prespecified value of one. Once this is achieved the controller maintains these angular velocities.

In Figure 16 the reorientation of the satellite for the angles φ, θ during the application of the controller is shown. Whilst these are becoming zero, the satellite is rotating about its *z* axis (Figure 17). We observe that the controller not only reorients the satellite, but it maintains this reorientation.

The applied thrusts during the genetic control of the satellite are shown in Figure 18 - Figure 20.

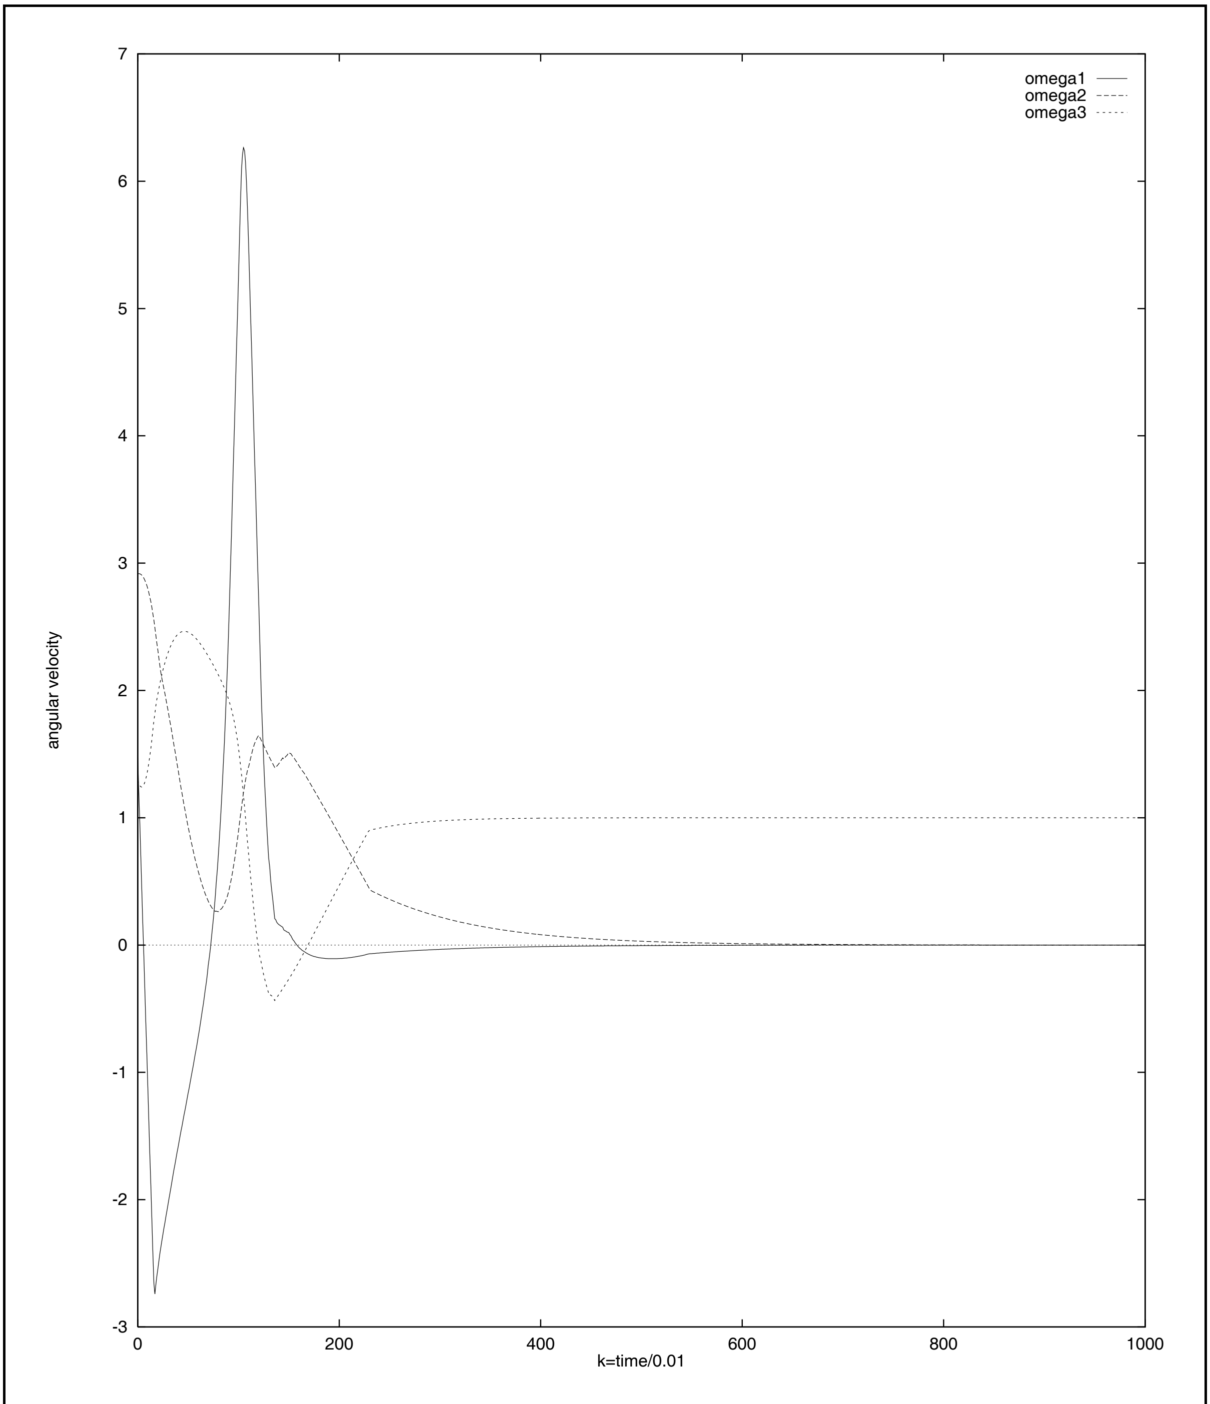

Figure 15 Angular velocities (ω_1, ω_2, ω_3) during the application of the genetic controller. Spin stabilization about and unstable axis.

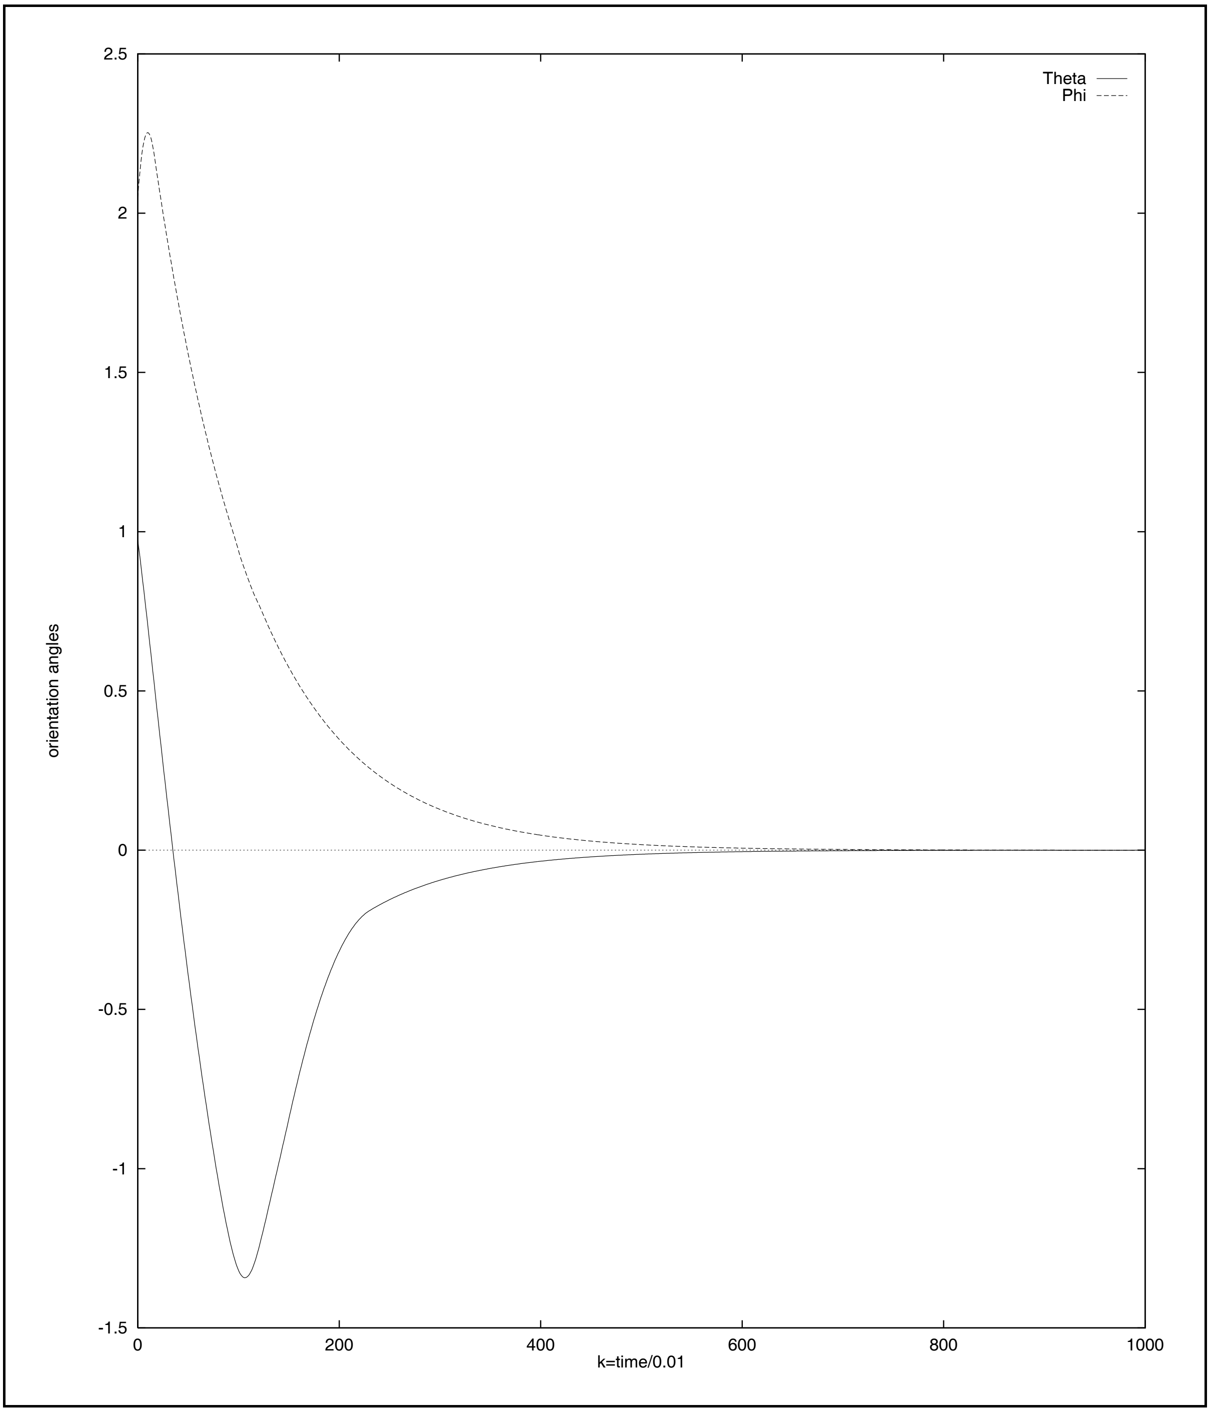

Figure 16 Orientation angles φ, θ during the application of the genetic controller. Spin stabilization about an unstable axis.

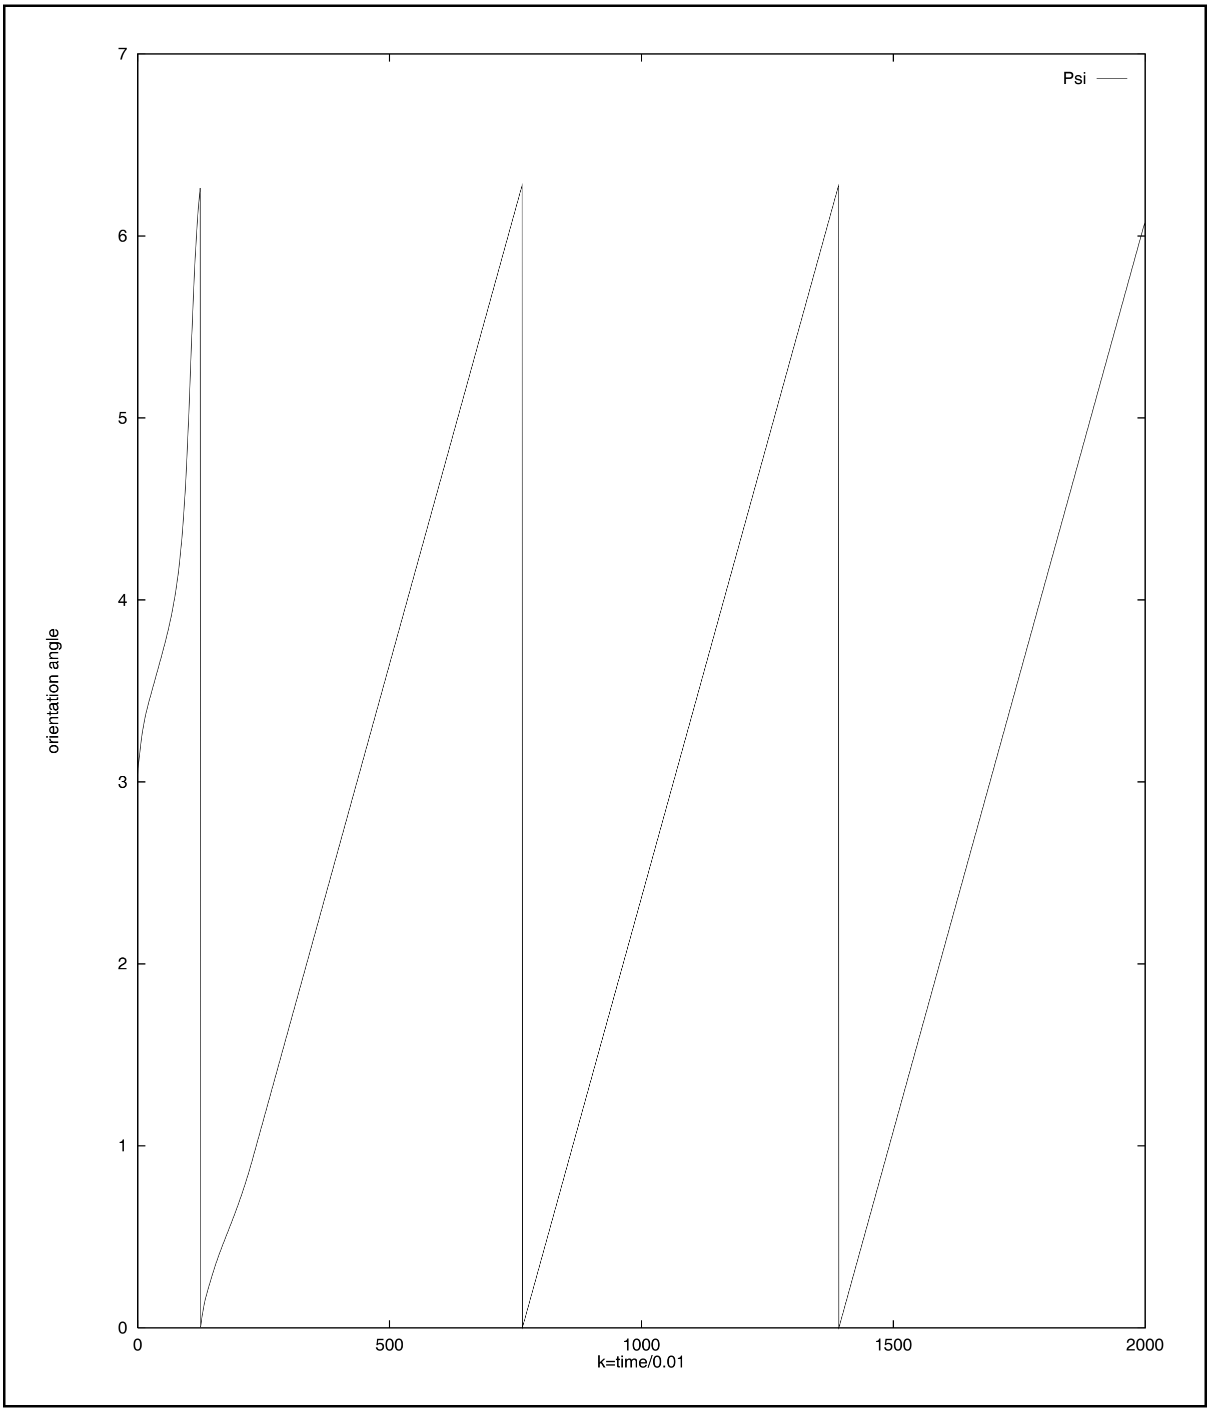

Figure 17 Orientation angle ψ during the application of the genetic controller. Spin stabilization about an unstable axis.

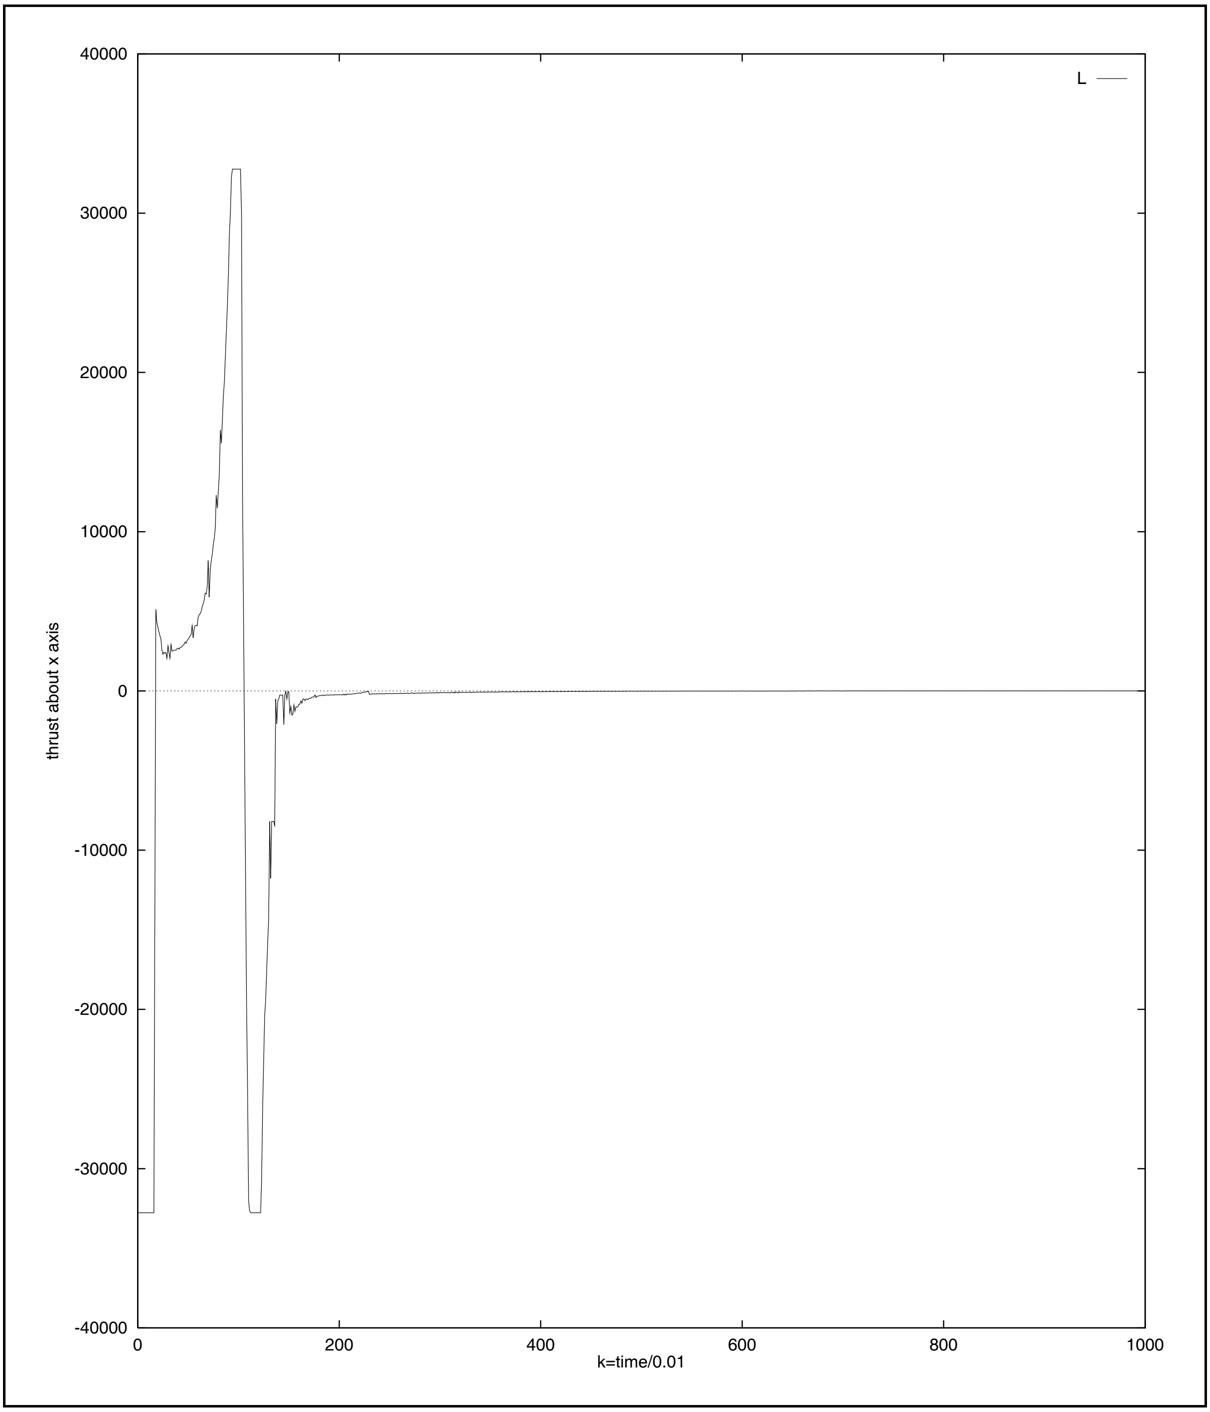

Figure 18 Applied thrust G_1 about the *x* body axis during the application of the genetic controller. Spin stabilization about an unstable axis.

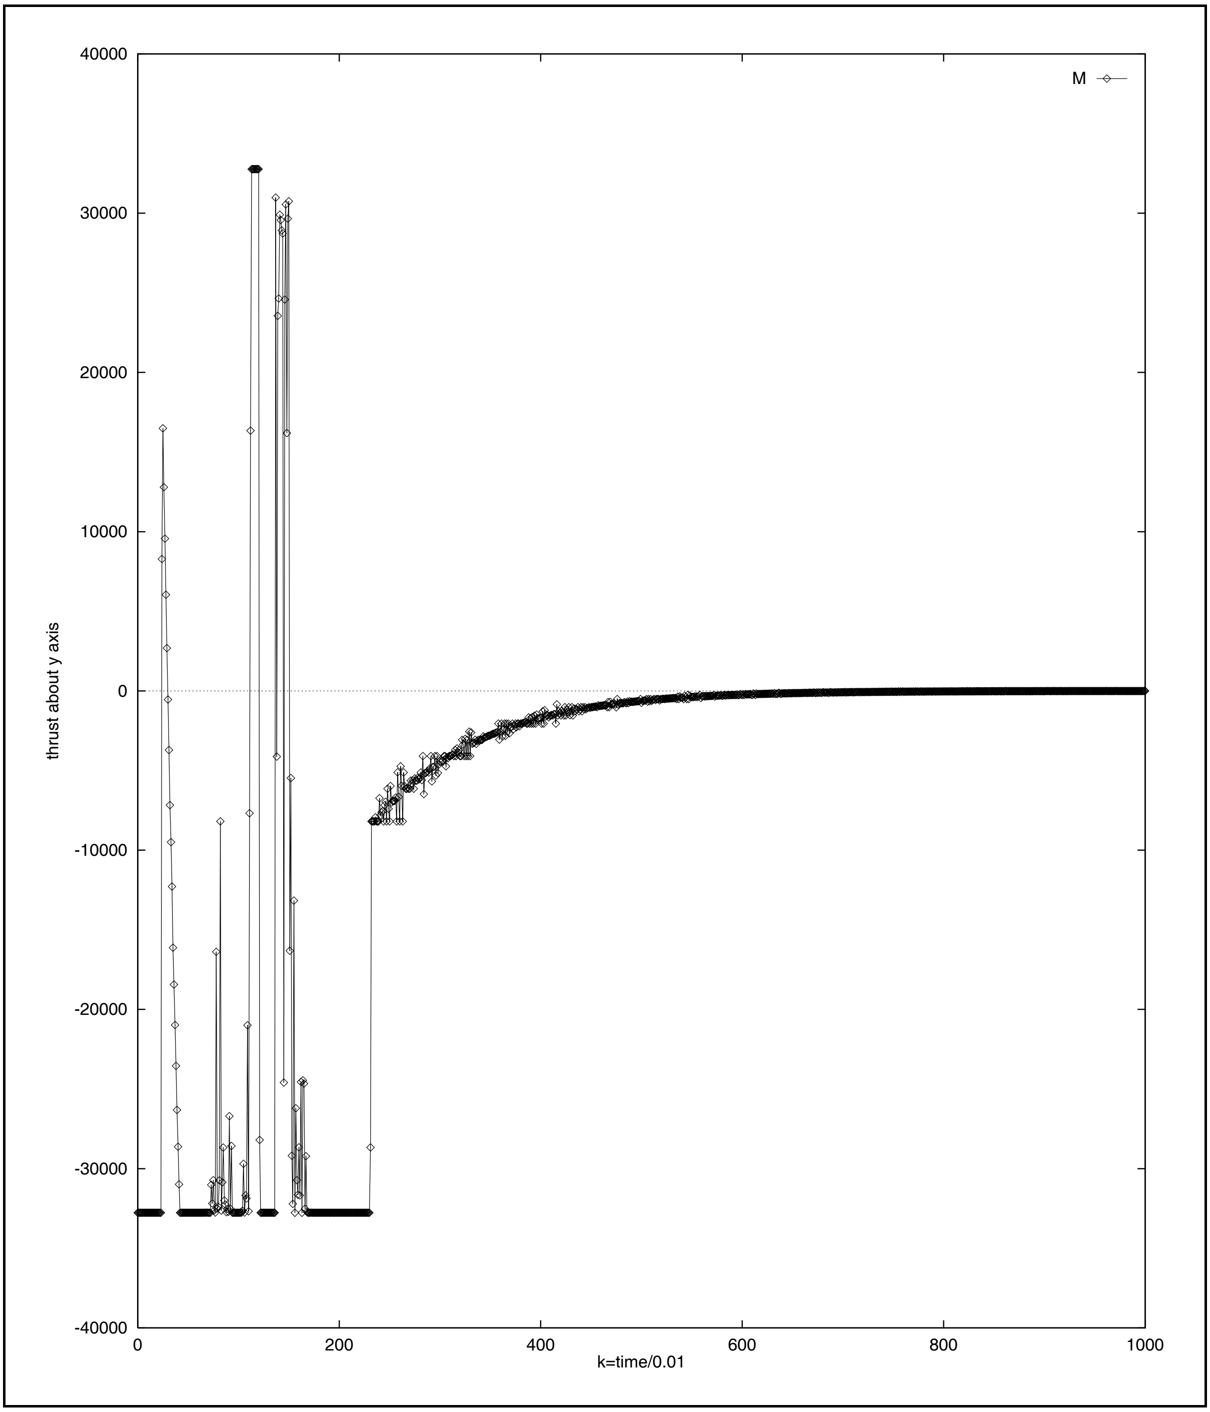

Figure 19 Applied thrust G_2 about the y body axis during the application of the genetic controller. Spin stabilization about an unstable axis.

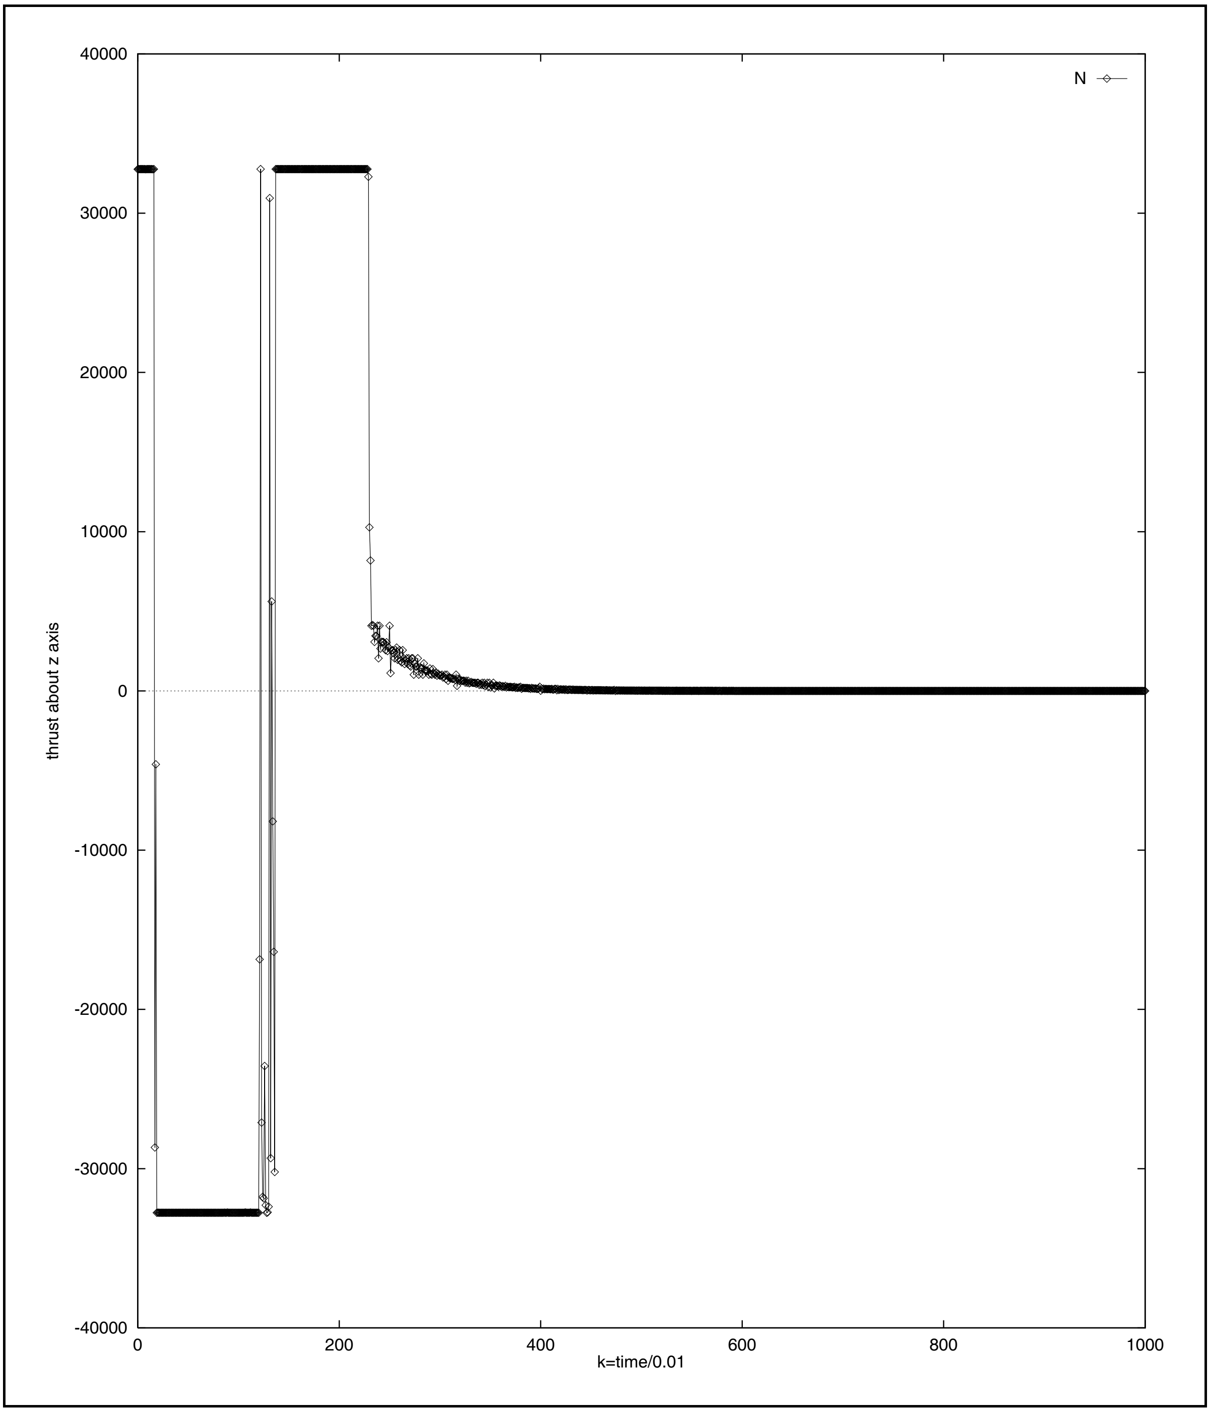

Figure 20 Applied thrust G_3 during the application of the genetic controller. Spin stabilization about an unstable axis.

Spin stabilization of a satellite subject to sensor noise.

Assume, that for some unknown reason (damage), a satellite with specified dynamics, changes its characteristics so that the moments of inertia become $I_x = 1160$, $I_y = 23300$ and $I_z = 24000$.

During the period where the system dynamics change, unknown forces lead it to the state $(\omega_1, \omega_2, \omega_3) = (1.8, 2.7, 1.4)$, $(\varphi, \theta, \psi) = (2, 1, 3)$. The goal is to spin stabilize the satellite about the *z* body axes, which is stable target state since the spin is about the axis having largest moment of inertia. Thus the target state of the system is set to $(\omega_1, \omega_2, \omega_3) = (0, 0, 1)$, $(\varphi, \theta, \psi) = (0, 0, 0)$.

Noise of 10% of the current sensor values (following a uniform distribution) is added to the sensor values (produced by the simulator), i.e. to the angular velocities $(\omega_1, \omega_2, \omega_3)$ about the body and the inertial orientation angles (φ, θ, ψ). Thus at any given time the controller has only imprecise knowledge of the actual system state.

The choice of the objective function needs some care. Since noise is present, when the system is near the target state much of the error will be due to noise and we should like to reduce 'hunting' (i.e. over-energetic control torques) and thus the long term energy expenditure of maintaining the target state. Consequently we want to smooth the angular acceleration near the target state. One way (amongst many) to achieve this is to introduce a smoothing term which forces $\ddot{\psi} \sim -(\dot{\psi} - 1)$ near the target state. Thus our objective function is chosen to be

$$F(G_1, G_2, G_3) = |\varphi| + |\dot{\varphi} + \varphi| + |\theta| + |\dot{\theta} + \theta| + |\dot{\psi} - 1| + 0.01\ |\ddot{\psi} + \dot{\psi} - 1| \qquad (12)$$

at the next sensor sample. This presupposes additional sensors for angular acceleration and would involve adding an extra output $\ddot{\psi}$ to the LPN predictor. For the present purpose, assuming the hypothetical thrusts we take a value of $\ddot{\psi}$ provided by the simulator. Introduction of terms like $|\dot{\varphi} + \varphi|$ was motivated by two factors. First, if the angular velocities are initially high these terms will dominate, and the emphasis of the corresponding control actions will be mostly aimed at reducing them. Second, as the systems approaches the target state these terms will act as *damping* so as to avoid overshooting the target state. Many other choices which accomplish the same goals are possible. The coefficient 0.01 for the higher order smoothing term makes this term small when the system is far from the target state, and thus speeds up acquisition of the target state; effectively the angular acceleration is not considered near the start of the simulation. As the system approaches the target state the first five terms in the objective function become small and the last term, the smoothing term, becomes significant.

In all the simulations the GA will attempt to minimise the objective function by maximising the fitness function (10). The application of the genetic adaptive control architecture, described by Figure 3, with the objective function given by (12), leads to the situation described by Figure 21 - Figure 26.Figure 21 shows the evolution in time of the angular velocities $\omega_1, \omega_2, \omega_3$ about the *x, y, z* body axes respectively. The genetic controller soon leads both ω_1 and ω_2 to the prespecified value of zero and ω_3 to one. Figure 22 shows the reorientation of the satellite for the angles φ, θ during the application of the genetic controller. While these are becoming zero, the satellite is rotating about its *z* axis, Figure 23. It should be noted that the controller not only leads the system to a desired state, but it maintains this state afterwards.

The applied thrusts during the genetic control of the satellite are shown in Figure 24 - Figure 26. We observe that during the time that the angular velocities are large the thrusts vary very rapidly, so that in some situations they can be seen as applying a kind of *bang-bang control*. As soon as the angular velocities obtain small values, the required and applied thrusts G_1, G_2, G_3 become small and somewhat smoother. Not surprisingly the presence of significant noise reduces the smoothness of the control actions near the target compared with the previous example.

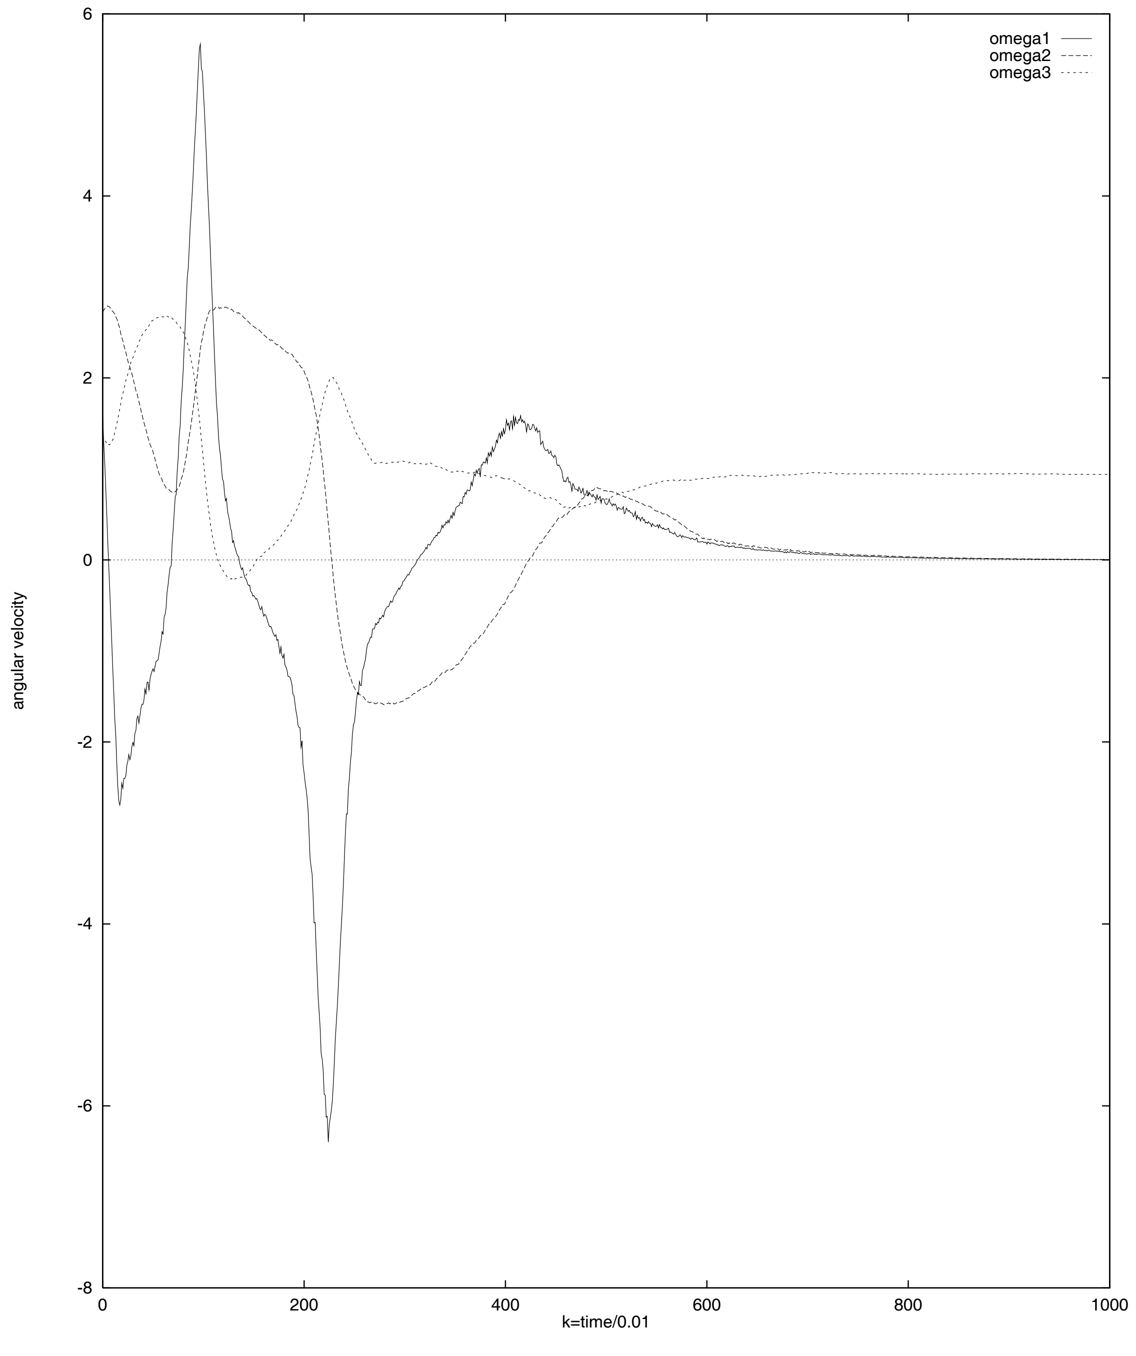

Figure 21 Angular velocities ω_1, ω_2, ω_3 during the application of the genetic controller in the presence of 10% sensor noise.

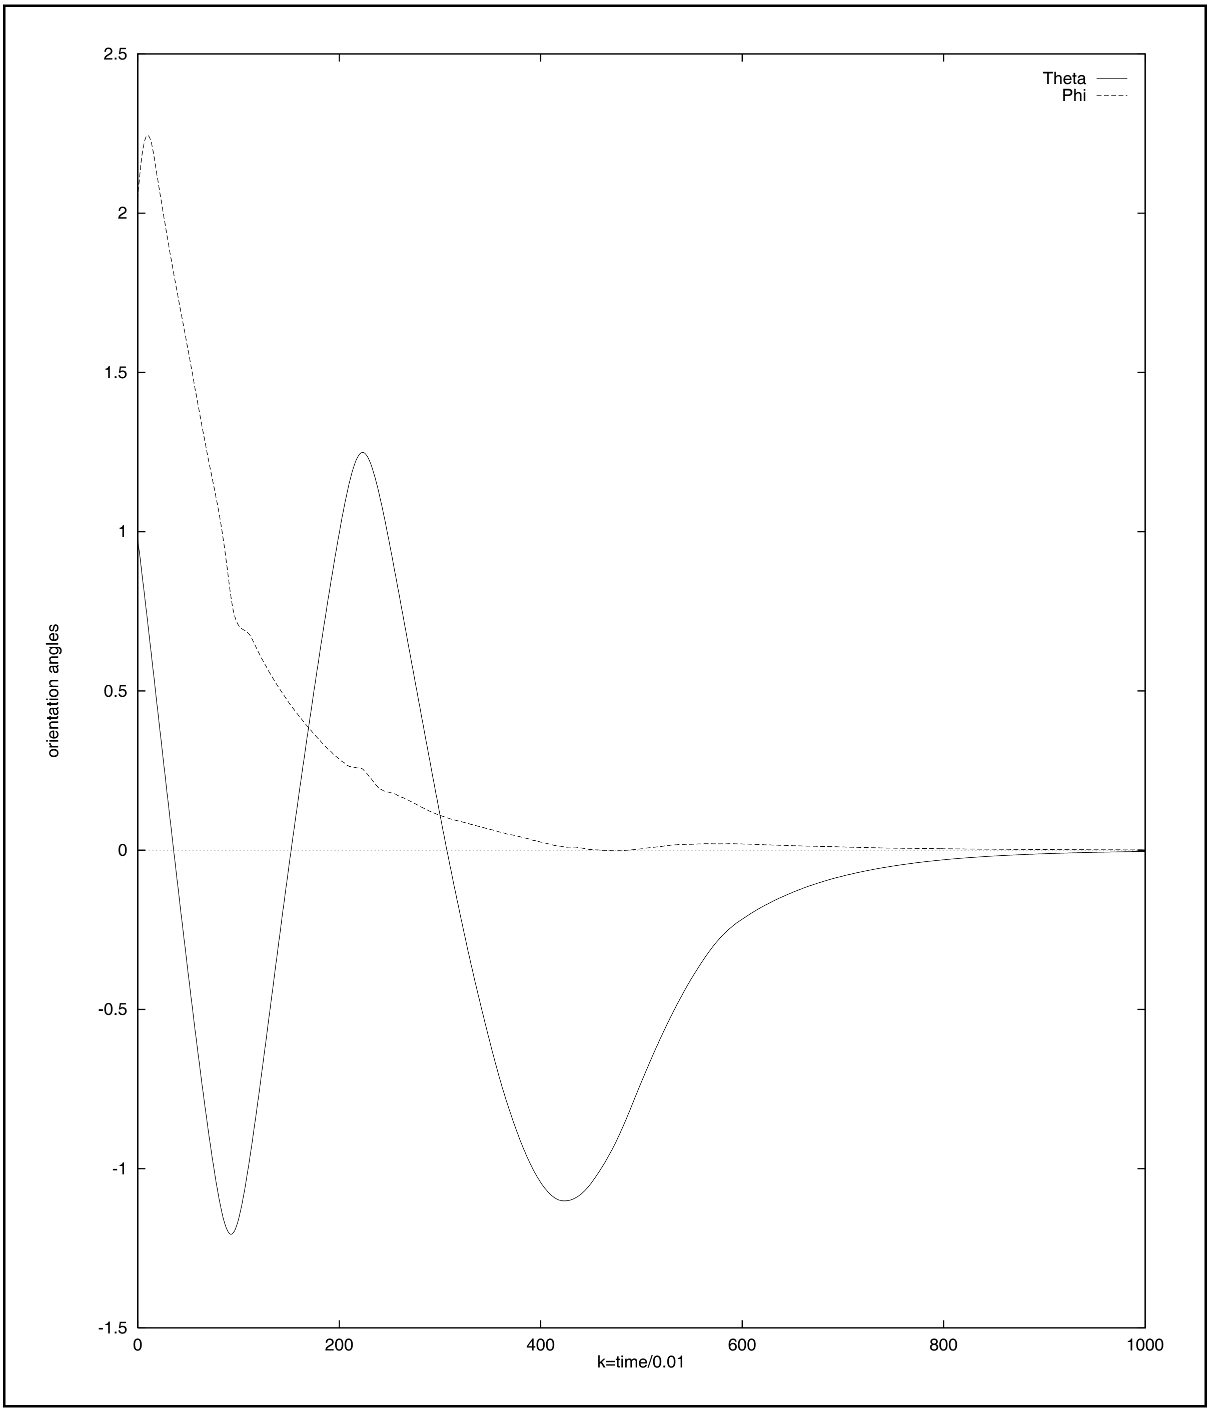

Figure 22 Orientation angles φ, θ during the application of the genetic controller in the presence of 10% sensor noise.

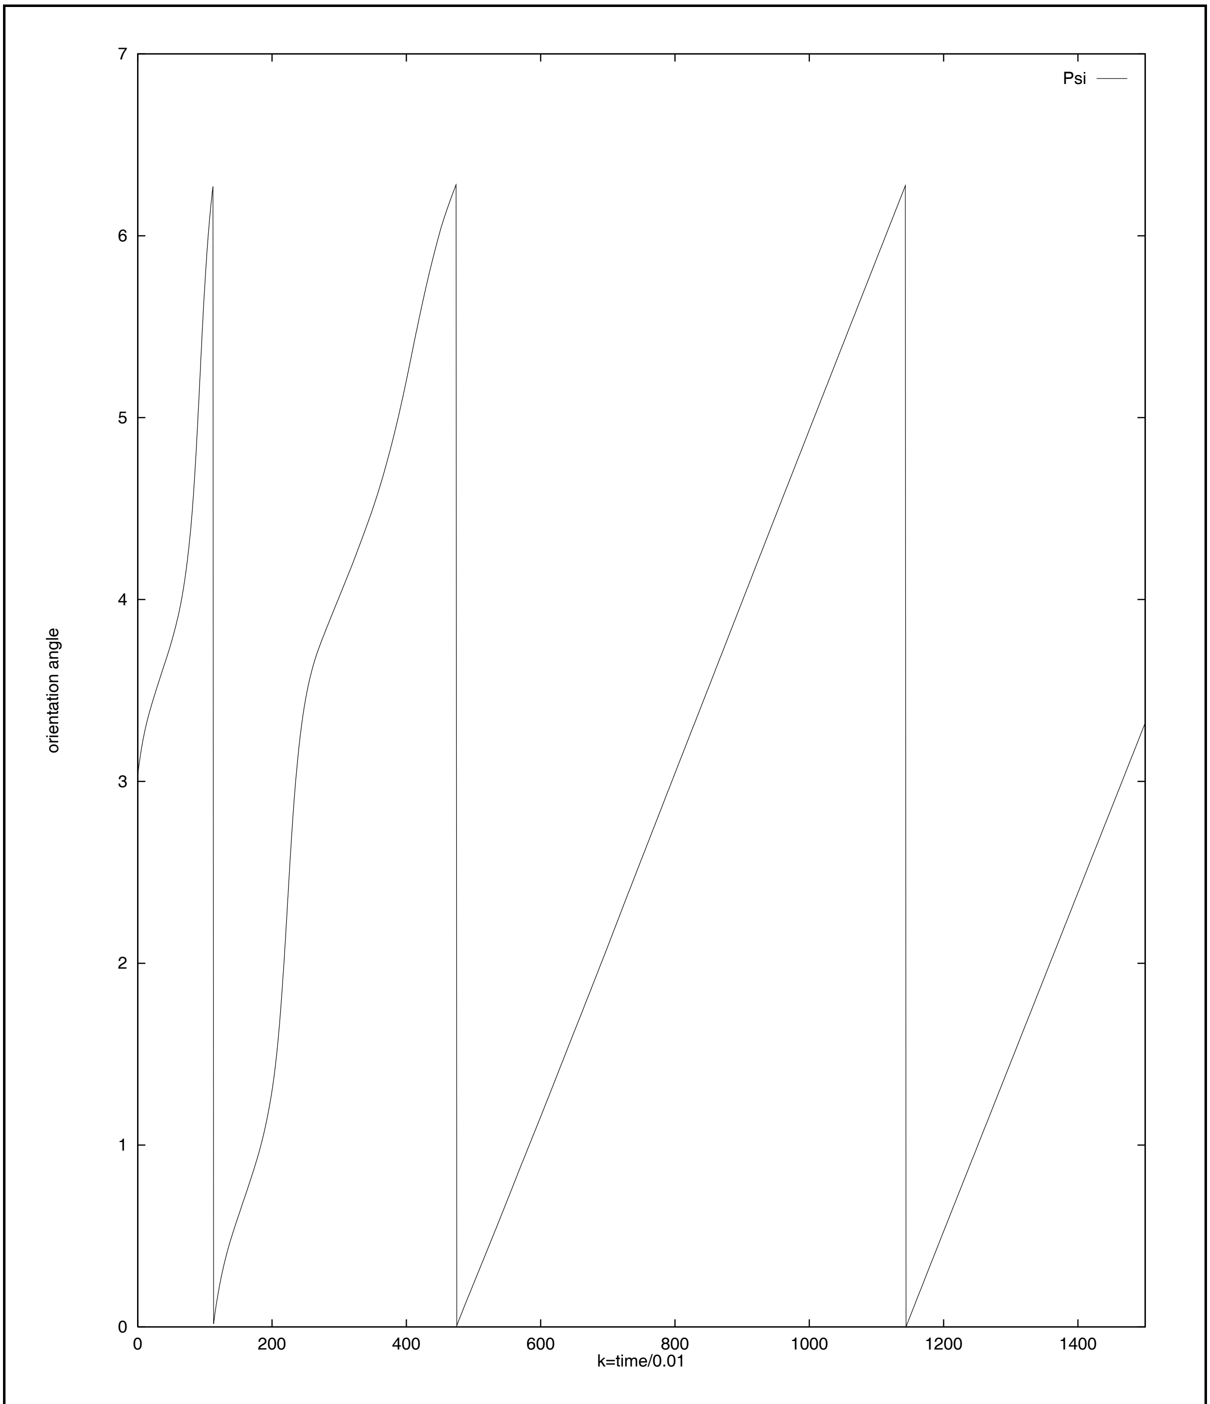

Figure 23 Orientation angle ψ during the application of the genetic controller in the presence of 10% sensor noise.

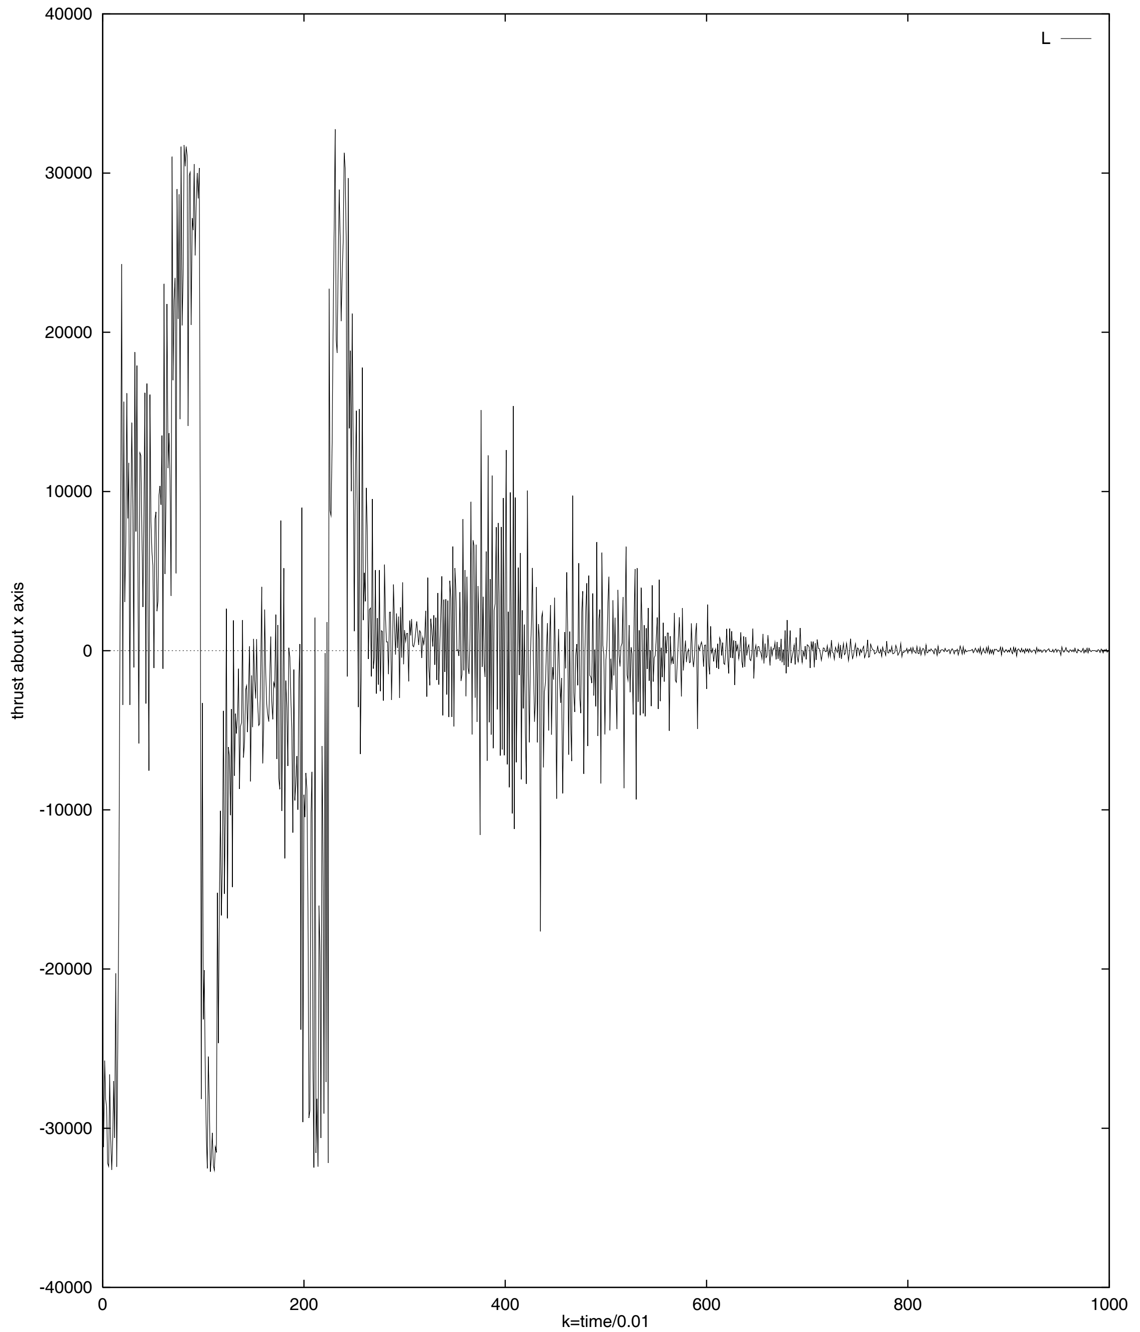

Figure 24 Applied thrust G_1 during the application of the genetic controller in the presence of 10% sensor noise.

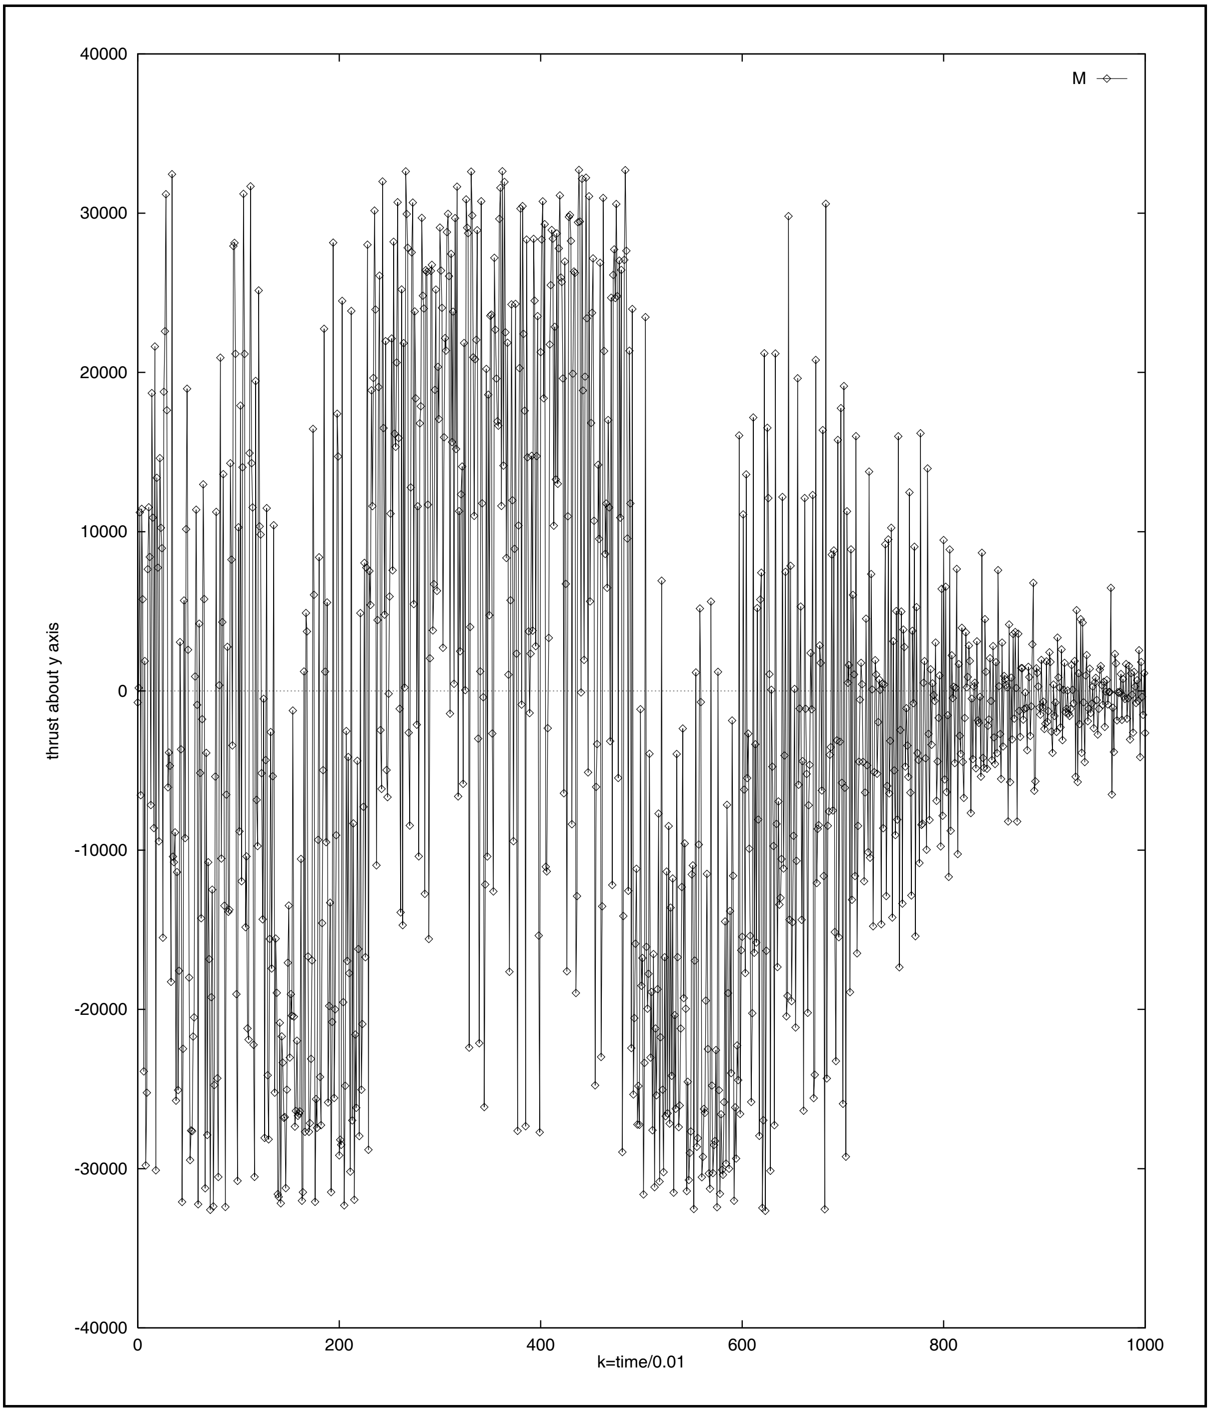

Figure 25 Applied thrust G_2 during the application of the genetic controller in the presence of 10% sensor noise.

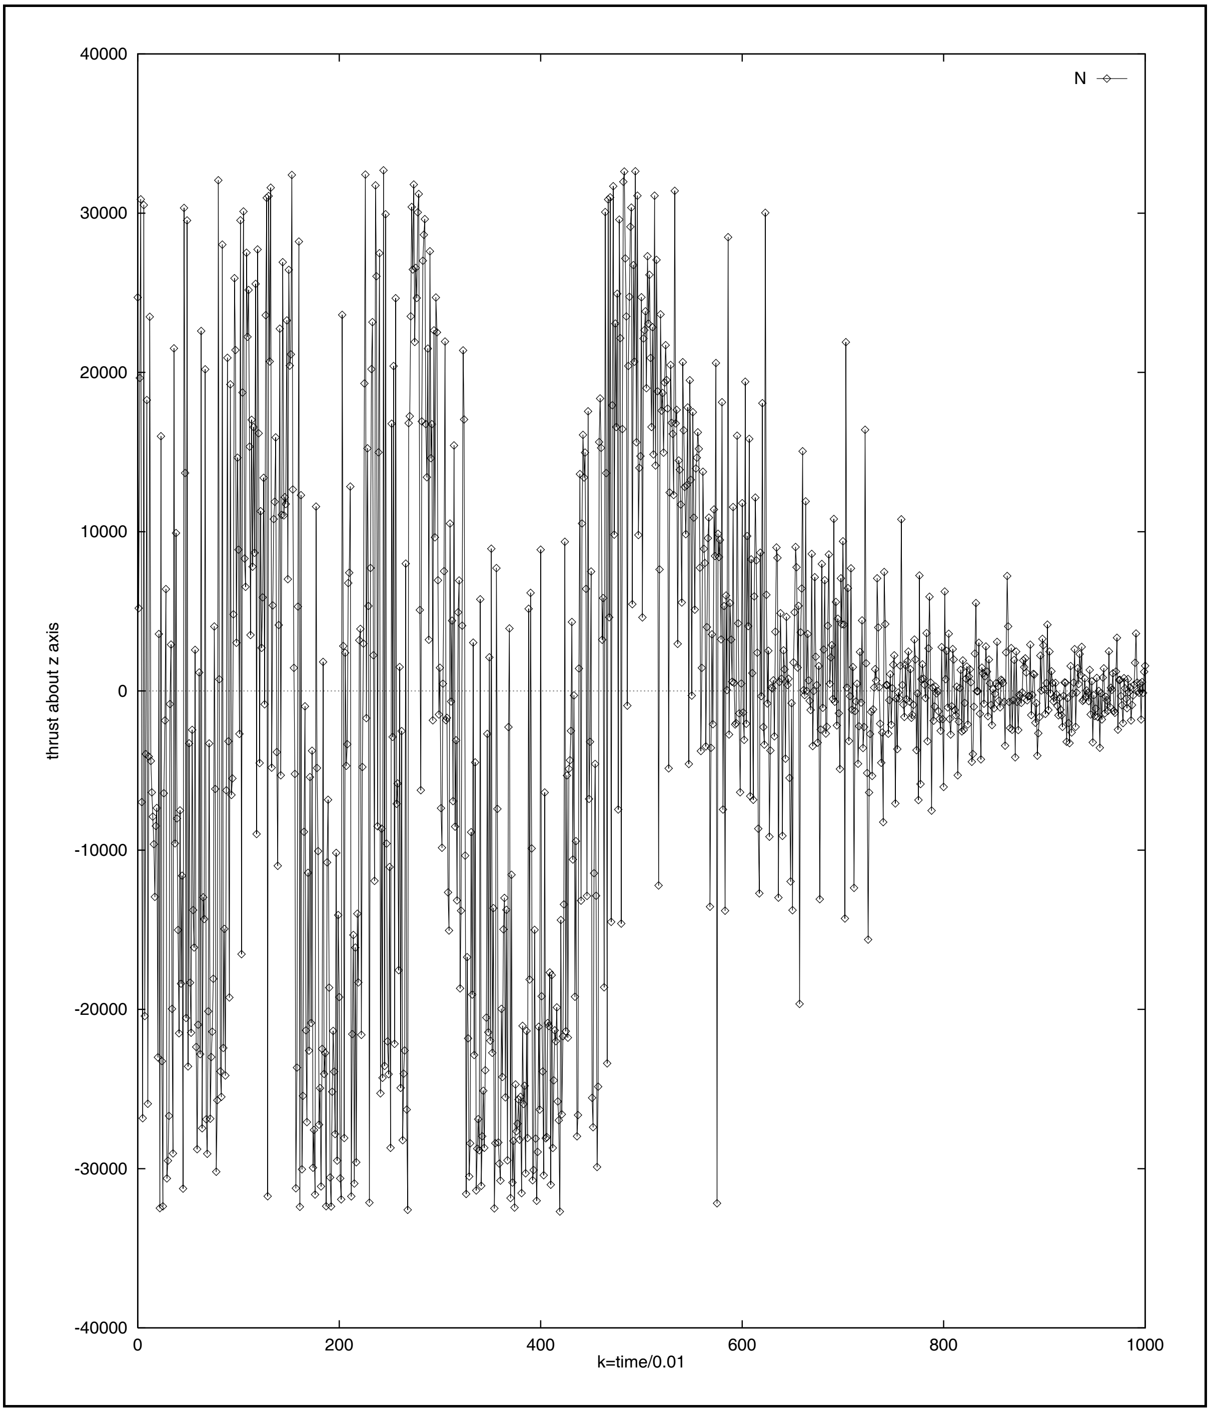

Figure 26 Applied thrust G_3 during the application of the genetic controller in the presence of 10% sensor noise.

Conclusions.

We have examined simulation experiments in which a speculative neuro-genetic architecture was applied to the adaptive attitude control problem for arbitrary transitions of system states. The experiments tested the ability of the system to acquire and maintain an arbitrary target state with no prior knowledge of the system dynamics. This was achieved for both dynamically stable and unstable target states. By themselves these two experiments are encouraging (especially since simple objective functions were used for the genetic algorithm) but are by no means definitive: real systems should be able to cope with a variety of complications one of which is inaccuracies in the sensors. In the third experiment the robustness of the architecture was tested by introducing 10% noise added to the sensor data. The system coped very effectively with this additional difficulty.

The dynamic system used in the experiments was low dimensional but exhibits highly non-linear characteristics typical of many real life control problems. The experiments suggest that given an appropriate hardware implementation the adaptive Neuro-Genetic architecture is feasible for a wide range of potential control applications. The precise hardware configuration required would depend on the complexity of the system to be controlled and the speed of events in the real world. A discrete logic implementation of the neural network and a gate array version of the genetic algorithm, with an additional processor and memory, would more than suffice for most applications.

The advantage of this particular proposal lies principally in the fact that the architecture can in principle be applied to non-linear control problems in which the inverse kinematics are ill posed and potentially good control solutions may even be discontinuous functions of time, i.e. it is perfectly general. A potential disadvantage is that, without suitable simulation experiments for a particular application domain, no guarantees can be offered that the method will be effective.

Our experiments raised some interesting questions. It seems likely that for the particular system considered there are control solutions which may acquire the target state in approximately the same time interval, but which are less discontinuous and more energy efficient. Of course, energy efficiency will usually involve a tradeoff against time-to-target. We think it probable that by suitably constraining the genetic algorithm search it may be possible to improve energy efficiency for the attitude control problem. In general the question of using a genetic algorithm for multi-objective optimization is an area requiring further research.

Acknowledgements.

This work was supported in part by *SERC 90800355*.

References.

Attitude control review

[Bloch 1992] A. M. Bloch, P. S. Krishnaprasad, J. E. Marsenden and G. Sanchez de Alvarez. *Stabilization of rigid body dynamics by internal and external forces*. Automatica. 28(4): 745-756. 1992.

[Branets 1988] V. N. Branets, D. M. Weinberg. V. P. Verestchagin. N. N. Danilov-Nitusov, V. P. Legostayev, V. N. Platonov. Yu. P. Semyonov, V.S. Semyachkin, B. E. Chertook and N. N. Sheremetyevsky. *Development experience of the attitude control system using single-axis control moment gyros for long-term orbiting space stations*. ACTA Astronautica 18:91-98. 1988.

[Calico 1983] R. A. Calico Jr. and G. S. Yeake. *Active attitude control of a spinning symmetrical satellite in an elliptic orbit*. Journal of Guidance, Control and Dynamics. 6(4):315-320. July-August 1983.

[Ceballos 1986] Décio Castelho Ceballos. *Compensating structure and parameter optimization for attitude control of a flexible spacecraft*. Journal of Guidance, Control and Dynamics 9(2):248-249. March-April 1986.

[Chavy 1991] S. Chavy and F. Fantar. *Design of attitude control system for injection of spinned payloads*. ACTA Astronautica 25(4):185-197. 1991.

[Chretien 1982] J. P. Chretien, C. Reboulet and P. Rodrigo. *Attitude control of a satellite with a rotating solar array*. Journal of Guidance, Control and Dynamics 5(6):580-596. November-December 1982.

[Crouch 1984] Peter E. Crouch. *Spacecraft attitude control and stabilization: Applications of geometric control theory to rigid body models*. IEEE Transactions on Automatic Control. AC-29:321-331. 1984.

[Daley 1987] S. Daley and K. F. Gill. *Attitude control of a spacecraft using an extended self-organising fuzzy logic controller*. Proceedings of the Institute of Mechanical Engineers 201(C2):97-106. 1987.

[Davis 1990] William R. Davis and David A. Levinson. *Attitude control of a satellite containing a spinning rotor with a nadir-pointing axis*. Journal of the Astronautical Sciences 38(1):1-19. January-March 1990.

[De Melo 1984] Jorge D. De Melo and Sahjendra N. Singh. *VTOL aircraft control output tracking sensitivity design*. IEEE Transactions on Aerospace and Electronic Systems. AES-20(2):82-93. March 1984.

[De 1992] Prasada K. De and Ashok Iyer. *Stable fraction approach and attitude control of a spinning satellite using a gyrotorquer*. International Journal of Systems Science 23(3):311-328. 1992.

[Dettleff 1986] G. Dettleff, R. D. Boettcher, C. Dankert, G. Koppenwallner and H. Legge. *Attitude control thruster plume flow modeling and experiments*. Journal of Spacecraft and Rockets 23(5):476-481. September-October 1986.

[Dodds 1991] S. J. Dodds and A. B. Walker. *Sliding-mode control system for the three-axis attitude control of rigid-body spacecraft with unknown dynamics parameters*. International Journal of Control 54(4):737-761. 1991.

[Dodds 1984] S. J. Dodds and S. E. Williamson. *A signed switching time bang-bang attitude control law for fine pointing flexible spacecraft*. International Journal of Control 40(4):795-811. 1984.

[Dodds 1981] S. J. Dodds. *Adaptive, high-precision, satellite attitude control for microprocessor implementation*. Automatica 17(4):563-574. 1981.

[ElMadany 1992] M. M. ElMadany and Z. Abduljabbar. *On the ride and attitude control of road vehicles*. Computers and Structures 42(2):245-253. 1992.

[Fenton 1981] J. Fenton and K. F. Gill. *Flexible spacecraft attitude control using a simple p+d algorithm*. Aeronautical Journal 85(844):185-180. May 1981.

[Furuta 1984] K. Furuta, T. Ochiai and N. Ono. *Attitude control of a triple inverted pendulum*. International Journal of Control. 39(6):1351-1365. 1984.

[Gregory 1992] J. C. Gregory and P. N. Peters. *Measurement of the passive attitude control performance of a recovered spacecraft*. Journal of Guidance, Control and Dynamics 15(1):282-284. January-February 1992.

[Herman 1991] A. L. Herman and B. A. Conway. *Optimal spacecraft attitude control using collocation and non-linear programming*. Journal of Guidance, Control and Dynamics 15(5):1287-1289. 1991.

[Hermes 1980] H. Hermes. On a stabilizing feedback attitude control. Journal of Optimization Theory and Applications 31(3):373-384. July 1980.

[Hermes 1985] Henry Hermes and David Hogenson. *The explicit synthesis of stabilizing (time optimal) feedback controls for the attude control of a rotating satellite*. Applied Mathematics and Computation. 16:229-240. 1985.

[Hess 1984] Ronald A. Hess. *Analysis of aircraft attitude control systems prone to pilot induced oscillations.* Journal of Guidance, Control and Dynamics 7(1):106-112. January-February 1984.

[Hodgart 1982] M. S. Hodgart. *Attitude control and dynamics of UOSAT angular motion.* Radio and Electronic Engineer 52(8/9):379-384. August-September 1982.

[Humble 1990] Ronald W. Humble. *Two dimensional tethered satellite attitude dynamics.* Journal of the Astronautical Sciences. 38(1):21-27. January-March 1990.

[Iyer 1989] Ashok Iyer and Sahjendra N. Singh. *Mfds of spinning satellite and attitude control using gyrotorquers.* IEEE Transactions on Aerospace and Electronic Systems 25(5):611-620. September 1989.

[Jahanshahi 1982] M. H. Jahanshahi. *Simultaneous calibrations of VOYAGER celestial and interial attitude control systems in flight.* IEEE Transactions on Aerospace and Electronic Systems. AES-18(1):21-28. January 1982.

[Jahangir 1993] E. Jahangir and R. M. Howe. *Time-optimal control scheme for a spinning missile.* Journal of Guidance, Control and Dynamics. 16(2):346-353. March-April 1993.

[Junkins 1982] J. L. Junkins, S. Rajaram, W. A. Baracat and C. K. Carrington. *Precision autonomous satellite attitude control using momentum transfer and magnetic torquer.* Journal of the Astronautical Sciences. 30(1):31-48. January-March 1982.

[Kline-Schoder 1993] Robert J. Kline-Schoder and J. David Powell. *Precision attitude control for tethered satellites.* Journal of Guidance, Control and Dynamics. 16(1):168-174. January-February 1993.

[Lemke 1987] Larry G. Lemke, J. David Powell and Xiaohua He. *Attitude control of tethered spacecraft.* Journal of the Astronautical Sciences. 35(1):41-55. January-March 1987.

[Lin 1989] Y. Y. Lin and L. G. Kraige. *Enhanced techniques for solving the two-point boundary problem associated with the optimal attitude control of spacecraft.* Journal of the Astronautical Sciences 37(1):1-15. 1989.

[Long 1992] Theresa W. Long. *Satellite attitude control in the equivalent axis frame.* ACTA Astronautica. 26(5):299-305. 1992.

[Meyer 1966] George Meyer. *On the use of Euler's theorem on rotations for the synthesis of attitude control systems.* Technical Report TN D-3643. NASA 1966.

[Meyer 1971] George Meyer. *Design and global analysis of spacecraft attitude control systems.* Technical report TR R-361. NASA 1971.

[Muhlfelder 1986] L. Muhlfelder. *Evolution of an attitude control system for body-stabilized communication spacecraft.* Journal of Guidance, Control and Dynamics 9(1):108-112. January-February 1986.

[Nicosia 1992] S. Nicosia and P. Tomei. *Nonlinear observer and output feedback attitude control of spacecraft.* IEEE Transactions on Aerospace and Electronic systems. 28(4):970-977. 1992.

[Paielli 1993] Russell A. Paielli and Ralph E. Bach. *Attitude control with realization of linear error dynamics.* Journal of Guidance, Control and Dynamics. 16(1):182-189. January-February 1993.

[Pande 1982] K. C. Pande and R. Venkatachalam. *Optimal solar pressure attitude control of spacecraft - (i).* ACTA Astronautica 9(9):533-540, 1982.

[Parkinson 1990] Bradford W. Parkinson and N. Jeremy Kasdin. *A magnetic attitude control system for precision pointing of the rolling GP-B spacecraft.* ACTA Astronautica 21(6/7):477-486. 1990.

[Parlos 1992] Alexander G. Parlos and John W. Sunkel. *Adaptive attitude control and momentum management for* *large-angle spacecraft maneuvres.* Journal of Giudance, Control and Dynamics. 15(4):1018-1028. July-August 1992.

[Piper 1992] George E. Piper and Harry G. Kwatny. *Complicated dynamics in spacecraft attitude control systems.* Journal of Guidance, Control and Dynamics 15(4):825-831. July-August 1992.

[Pope 1991] Tonya Pope. *Real-time simulation of the space station Freedom attitude control system.* Simulation. 57(1):17-25. 1991.

[Rahasingh 1987] C. K. Rahasingh and Shashi K. Shrivastava. *Orbit and attitude control of a geostationary inertially oriented large flexible plate-like spacecraft.* ACTA Astronautica 15(11):823-832. 1987.

[Rhee 1992] Ihnseok Rhee and Jason L. Speyer. *Robust momentum management and attitude control system for the space station.* Journal of Guidance, Control and Dynamics 15(2):342-351. March-April 1992.

[Shao-hua 1981] Yu Shao-hua. *Satellite attitude control with decomposed controller.* Journal of Guidance, Control and Dynamics 14(6):584-585. November-December 1981.

[Sheela 1991] B. V. Sheela and Chandra Shekhar. *New star identification technique for attitude control.* Journal of Guidance, Control and Dynamics 14(2):477-480. March-April 1991.

[Singh 1993] Sahjendra N. Singh and Theodore C. Bossart. *Exact feedback linearization and control of space station using cmg.* IEEE Transactions on Automatic Control. 38(1):184-187. January 1993.

[Singh 1989] Sahjendra N. Singh and Ashok Iyer. *Non-linear decoupling sliding mode control and attitude control of spacecraft.* IEEE Transactions on Aerospace and Electronic Systems 25(5):621-633. September 1989.

[Singh 1988] Sahjendra N. Singh. *Flexible spacecraft maneuver: Inverse attitude control and modal stabilization.* ACTA Astronautica 17(1):1-9. 1988

[Singh 1987] Sahjendra N. Singh. *Robust non-linear attitude control of flexible spacecraft.* IEEE Transactions on Aerospace and Electronic Systems. AES-23(3):380-387. May 1987.

[Singh 1987] Sahjendra N. Singh. *Attitude control of a three rotor gyrostat in the presence of uncertainty.* Journal of the Astronautical Sciences 35(3):329-345. July-September 1987.

[Singh 1986] Sahjendra N. Singh. *Nonlinear attitude control of flexible spacecraft under distrubance torque.* ACTA Astronautica 13(8):507-514. 1986.

[Singh 1984] Sahjendra N. Singh. *Non-linear adaptive attitude control of flexible spacecraft.* IEEE Transactions on Aerospace and Electronic Systems. AES-20(2):94-103. March 1984.

[Skaar 1986] S. B. Skaar, L. Tang and I. Yalda-Mooshabad. *On-off attitude control of flexible satellites.* Journal of Guidance, Control and Dynamics 9(4):507-510. July-August 1986.

[Tcuchiya 1982] K. Tcuchiya, M. Inoue. N. Wakasuqi and T. Yamaguchi. *Advanced reaction wheel controller for attitude control of spcecraft.* ACTA Astronautica. 9(12):697-702. 1982.

[Thomas 1991] Segun Thomas. *Dynamics of spacecraft and manipulators.* Simulation. 57(1):56-72. 1991.

[Vadali 1992] R. S. Vadali and H. S. Oh. *Space station attitude control and momentum management. A non-linear look*. Journal of Guidance, Control and Dynamics 15(3):577-586. May-June 1992.

[Van Den Bosch 1986] P. P. J. Van Den Bosch, W. Jongkind and A. C. M. Van Swieten. *Adaptive attitude control for large slew-angle manoeuvres.* Automatica 22(2):209-215. 1986.

[Venkatachalam 1982] R. Venkatachalam and K. C. Pande. *Optimal solar pressure attitude control of spacecraft - (ii).* ACTA Astronautica 9(9):541-545. 1982.

[Wen 1991] John Ting-Yung Wen and Kenneth Kreutz-Delgado. *The attitude control problem.* IEEE Transactions on Automatic Control 36(10):1148-1162. October 1991.

[Wie 1990] Bong Wie, Auren Hu and Ramendra Singh. *Multibody interaction effects on space station attitude control and momentum management.* Journal of Guidance, Control and Dynamics. 13(6):993-999. November-december 1990.

[Woo 1988] Henry H. Woo, Henry D. Morgan and Eric T. Falangas. *Momentum management and attitude control design for a space station.* Journal of Guidance, Control and Dynamics 11(1):19-25. January-February 1988.

[Yuan 1986] Joseph S-C Yuan and Michael E. Stieber. *Robust beam-pointing and attitude control of a flexible spacecraft.* Journal of Guidance, Control and Dynamics 9(2):228-234. March-April 1986.

Other references

[Kristinsson 1992] Kristinn Kristinsson and Guy A. Dumont. *System Identification and Control Using Genetic Algorithms*. IEEE Transactions on Systems, Man and Cybernetics, 22(5):1033-1046, September-October 1992.

[Arnold 1978] V. I. Arnold. *Mathematical Methods of Classical Mechanics*. Springer-Verlag, 1978.

[Bender 1978] Carl M. Bender and Steven A. Orszag. *Advanced Mathematical Methods for Scientists and Engineers*. McGraw-Hill. 1978.

[Bondi 1986] Herman Bondi. *The rigid body dynamics of unidirectional spin.* Proc. Roy. Soc. Lond. A405:265-474, 1986.

[Bryson 1975] A. E. Bryson and Y. Ho. *Applied Optimal Control*. John Wiley, 1975.

[Chobotov 1991] Vladimir A. Chobotov. *Spacecraft Attitude Dynamics and Control*. Krieger Publishing Company, Malabar, FL. 1991.

[Conte 1980] S. D. Conte and Carl de Boor. *Elementary Numerical Analysis: An Algorithmic Approach*. McGraw-Hill. New York, Third Edition. 1980.

[Davis 1991] Lawrence Davis, editor. *Handbook of Genetic Algorithms*. Van Nostrand Reinhold, New York, 1991.

[Dracopoulos 1993] D. C. Dracopoulos and Antonia J. Jones. *Neuromodels of analytic Dynamic Systems*, Neural Computing & Applications, **1**(4):268-279, 1993.

[Gollub 1989] Gene H. Gollub and Charles F. Van Loan. *Matrix Computations*. John Hopkins University Press. Second Edition. 1989.

[Grefenstette 1989] J. Grefenstette. *A system for learning Control Strategies with Genetic Algorithms*. In: Proceedings of the Third International Conference on Genetic Algorithms, Ed. J. D. Schaffer, 183-189. Morgan Kaufmann, San Mateo, CA. 1989.

[Grefenstette 1986] John J. Grefenstette. *Optimization of control parameters for genetic algorithms.* IEEE Transactions on Systems, Man and Cybernetics. SMC-16:122-128, 1986.

[Jones 1993] *Genetic Algorithms and their Applications to the Design of Neural Networks*, Neural Computing & Applications, **1**: 32-45, 1993.

[Mendel 1970] J. M. Mendel and R. W. McLaren. *Reinforcement learning control and pattern recognition systems*. In Adaptive Learning and Pattern Recognition Systems, editors J. M. Mendel and K. S. Fu, Academic Press, 1970.

[Miller 1990] W. Thomas Miller, Richard S. Sutton and Paul J. Werbos, editors. *Neural networks for Control*. MIT Press, 1990.

[Narendra 1992] Kumpati S. Narendra. *Adaptive control of dynamical systems using neural networks*. In David A. White and Donald A. Sofge, editors. Handbook of Intelligent Control, Chapter 5. van Nostrand Reinhold, 1992.

[Odetayo 1989] M. O. Odetayo D. R. McGregor. *Genetic Algorithms for inducing Control Rules for a dynamic system*. In: Proceedings of the Third International Conference on Genetic Algorithms, Ed. J. D. Schaffer, 177-182. Morgan Kaufmann, San Mateo, CA. 1989.

[Routh 1892] E. J. Routh. *Dynamics of a System of Rigid Bodies: Part II*. Macmillan and Co., 1892.

[Shimada 1979] I. Shimada and T. Nagashima. *A mathematical approach to ergodic problem of dissipative dynamical systems.* Progress of Theoretical Physics, **61**:1605-1616, 1979.

[Werbos 1988] Paul J. Werbos. *Generalization of backpropagation with application to a recurrent gas market model*. Neural Networks 1:339-359, 1988.

[Werbos 1989] Paul J. Werbos. *Maximizing long term gas industry profits in two minutes in Lotus using neural network methods*. IEEE Transactions on Systems, Man and Cybernetics, 19(2):315-333, March-April 1989.

[Wertz 1978] James R. Wertz. *Spacecraft Attitude Determination and Control*. Ed. James R. Wertz, Kluwer Academic Publishers, 1978.

[Wilkinson 1965] J. H. Wilkinson. *The Algebraic Eigenvalue Problem*. Clarendon Press. 1965.