

Multi-Objective Optimisation of Distributed Active Magnetic Bearing Controllers

P. Schroder*, A. J. Chipperfield*, P. J. Fleming*, N. Grum†

* Dept. Automatic Control and Systems Eng., The University of Sheffield, Mappin Street, Sheffield, S1 3JD, UK.
Tel: +44 114 222 5250, Fax: +44 114 273 1729,
Email: P.Schroder@Sheffield.ac.uk.

† Rolls-Royce and Associates Limited, PO Box 31, Derby, DE24 8BJ, UK.
Tel: +44 1332 661 461 x5959,
Fax: +44 1332 622 948.

Abstract

A non-linear model of a Rolls-Royce turbo machine supported by active magnetic bearings (AMBs) is presented. A multi-objective genetic algorithm (MOGA) is used as a search and optimisation tool for designing distributed AMB controllers. The MOGA is used to select the controller structure as well as its parameters. The objective domain is comprised of measures of the controllers' dynamic performance in terms of disturbance and noise rejection, its efficiency and complexity. The optimisation is performed directly on the non-linear model to ensure that realistic controllers are produced.

Keywords: magnetic bearings, distributed control, optimal control, multi-objective genetic algorithms.

1. Introduction

Active magnetic bearings are used to suspend a rotating machine in a magnetic field. There is no contact between the rotor and the bearings, so there is no wear and little maintenance is required. There is also no need for lubrication systems which can be costly and may lead to oil contamination of the plant. AMBs can control the rotor position very accurately and predictably, so are useful for high precision and high speed machinery [1]. As accurate information about the rotor's position is required for the bearings to operate, this can be used for plant monitoring and fault diagnostics. AMBs can also be used to control the vibration caused by unbalanced rotors. However, AMBs are inherently unstable and require an active control scheme. This characteristic makes the overall system complex, and as there is typically a large software component, reliability and safety are expensive to ensure. Such 'mechatronic' problems in constructing AMB systems have led to a slow take up of the technology by industry [2].

This paper demonstrates a straight-forward approach to the design of AMB control systems. The design process is followed from the plant specification and modelling to controller design and selection. The application presented here is that of a Rolls-Royce marine turbo machine's rotor supported by AMBs. The system is required to operate in the

presence of measurement noise and large external disturbances. The aim of the control system designer is therefore to develop a controller that will maintain satisfactory control under these conditions. In systems of this type it is desirable to limit the order of the controller as a high sampling rate is generally required. This requirement motivated the structure selection distributed control approach described here, as previous work using the H_{∞} loop-shaping design procedure resulted in the generation of comparatively high order controllers [3]. Decentralised controllers have also been shown to be effective in controlling AMB systems, [4].

A number of conflicting design requirements have to be satisfied, disturbance rejection, noise rejection, efficiency and steady state error all have to be optimised. A multi-objective genetic algorithm is therefore used to search across a number of different controller structures and parameters [5]. The performance specifications are used as targets for the MOGA, which also optimises the control effort and the controller's order. The performance is measured directly from simulation of each controller on a non-linear plant model to increase the accuracy of the information obtained. When the search is terminated, the designer is presented with many different controllers all of which satisfy the specification. A choice is then made by the designer as to which controller offers the best overall performance for the application.

2. Active Magnetic Bearing Model

It is intended to support a turbo machine's rotor on active magnetic bearings to improve the system's reliability and performance. The turbo machine is intended for marine use where it will sit on a resiliently mounted platform which is designed to help absorb impulse disturbances. A model of the rotor-bearing system has been implemented in Simulink, a non-linear dynamic simulation environment. The model represents a stiff rotor supported by two active magnetic bearings, each constructed of six electromagnets. The limiting transient for the bearings is assumed to be a 30kN sinusoidal disturbance at 4 Hz. A schematic diagram of the rotor is shown in Fig. 1.

Bearing A is situated at the left of Fig. 1, and its electromagnets have a designed maximum load of 6kN. Its

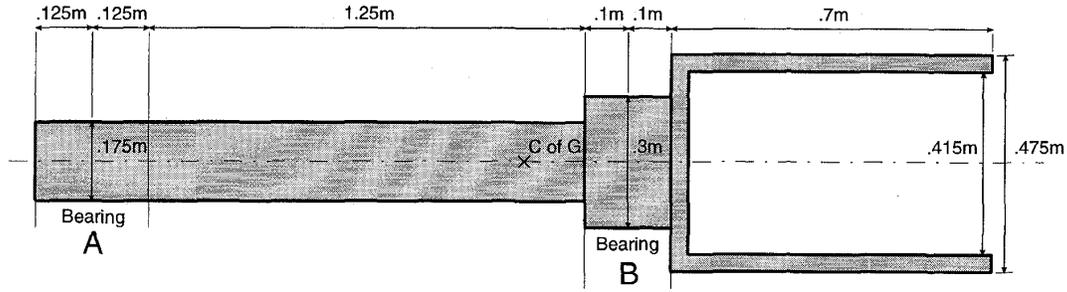


Figure 1: Turbo machine's rotor

purpose is mainly to keep the rotor steady whilst most of the disturbance forces are absorbed by bearing B which is situated close to the rotor's centre of gravity and has a designed maximum load of 30kN. Both magnetic bearings have the structure shown in Fig 2.

The magnetic bearings modelled here consist of three main stages, a power amplifier model, a model of the coil and a model of the magnet. The power amplifier is essentially a voltage amplifier with a high gain; its voltage range is -300V to +300V and it has a maximum output current of 5A. The maximum magnetic flux generated by the magnets is limited to 1.8T for the bearing B and 1.5T for bearing A. The current in the magnet coil is made to track the control signal by use of unity gain negative feedback of the current in the magnet coil. This AMB configuration is known as current control [6]. A bias current is maintained in each magnet in the bearing. The control signal is added to this for the upper magnets, labelled '+' in Fig. 2 and subtracted from it for the lower magnets, labelled '-'. The linearity of the system is increased by treating the magnets in pairs in this way.

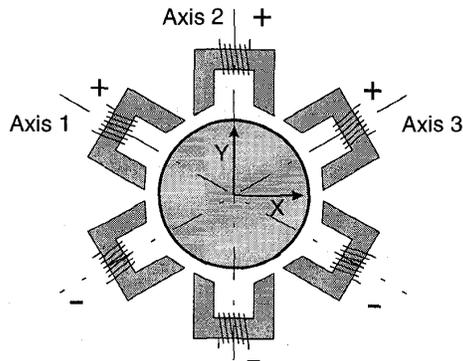


Figure 2: The magnetic bearing structure

The inductance model receives voltages from the amplifier model and converts them to the currents in each magnet coil by the relationship,

$$V = Ri + L \frac{di}{dt},$$

where R = coil resistance,
 L = coil inductance,

i = current in the coil and
 V = voltage across the coil.

The magnet model takes the current in the coils and the gaps between the magnets and the rotor and calculates the forces developed on the rotor. The force generated by the magnets is given by:

$$F = k \frac{i^2}{d^2},$$

where F = force exerted on the rotor,
 d = gap between rotor and magnet, and,

$$k = \text{magnet constant} = \frac{1}{4} \mu_0 N^2 A.$$

Here, μ_0 = magnetic field constant of a vacuum,
 N = number of turns on the magnet coil, and
 A = cross sectional area of the magnet.

This model is linearised for stability analysis by removing the limiting elements from the power amplifier model and using a Taylor series approximation of the magnet model.

3. Multi-objective GA Structure

A multi-objective genetic algorithm is used to find a set of controller structures and parameters. It is implemented using the GA Toolbox for MATLAB [7], developed in house, with additional extensions to accommodate multi-objective ranking, sharing and mating restrictions [8]. The salient features of this MOGA are shown in Fig. 3.

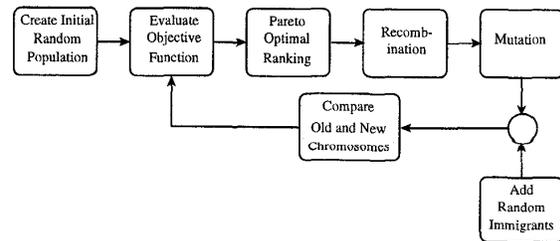


Figure 3: The MOGA

Multi-objective ranking is based upon the dominance of an individual, i.e. how many individuals out-perform it in the objective space. This kind of ranking is non-unique, a

number individuals may be ranked 0, ie. non-dominated. Ranking may also be combined with goal and/or priority information to discriminate between non-dominated solutions. For example, a solution in which all the goals are satisfied may be considered superior, or *preferable*, to a non-dominated one in which some components go beyond the goal boundaries. All the preferred individuals thus achieve the same fitness, however the number of actual offspring may differ due to the stochastic nature of the selection mechanism. Thus, an accumulation of the imbalances in reproduction can lead the search into an arbitrary area of the trade-off surface. This phenomenon is known as genetic drift and can drastically reduce the quality and efficiency of the search. Proposed as a solution to genetic drift, fitness sharing penalizes the fitness of individuals in popular neighbourhoods in favour of more remote individuals of similar fitness [9].

Recombining arbitrary pairs of non-dominated individuals can result in the production of an unacceptably large number of unfit offspring, or lethals. A further refinement to the MOGA is therefore to bias the manner in which individuals are paired for recombination, often termed mating restriction [10]. This restricts reproduction to individuals that are within a given distance of each other. The population diversity is maintained by adding random genetic information at each generation as well as mutating existing individuals.

The MOGA keeps a record of all the non-dominated individuals. This is updated each generation so that stored individuals that become dominated are deleted. Individuals that have been evaluated in the previous generation are not evaluated again. This reduces the optimisation time by between 10% and 20%.

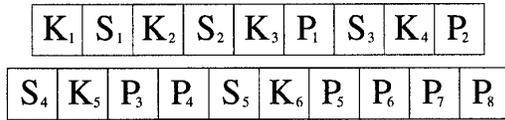


Figure 4: Structure of chromosome section

The controller parameters are represented in a chromosome constructed of 2 large sections, the structure of one of which is shown in Fig. 4. These sections represent the compensator structure for each bearing. It is assumed that each axis of control on a single bearing will use the same controller. It is therefore only necessary to evolve two compensators per individual (one for each bearing) and replicate both of these three times (one for each axis of control). The full compensator can then be expressed as follows:-

$$\text{Compensator} = \text{Parallel} \{ C_j \},$$

$$j = 1, 2, 3 - \text{Bearing A},$$

$$j = 4, 5, 6 - \text{Bearing B},$$

where C denotes an individual loop numbered j and *Parallel* indicates that the compensators are stacked with separate

inputs and outputs. An individual loop consists of the sum of a number of possible dynamic elements,

$$C_i = \sum \left\{ \begin{array}{l} K_1 \\ S_1 = 1, K_2 \frac{1}{s} \\ S_2 = 1, K_3 \frac{s}{P_1 s + 1} \\ S_3 = 1, K_4 \frac{P_2 s + 1}{s} \\ S_4 = 1, K_5 \frac{s + P_3}{s + P_4} \\ S_5 = 1, K_6 \frac{s^2 + P_5 s + P_6}{s^2 + P_7 s + P_8} \end{array} \right.$$

Here, the parameters S_i denote whether the transfer function element associated with it is active in the compensator. For example, a PID compensator would have $S_1 = S_2 = 1$, $S_3 = S_4 = S_5 = 0$, and would use the gains K_{1-3} and the parameter P_1 (representing the position of the differentiator's pole). The information is stored as a binary chromosome, with the S_i parameters represented by 1 bit and the others consuming 14 bits, each using a gray coded logarithmic representation. The chromosome is therefore comprised of 402 binary elements.

The search space contains a large number of possible controller structures, ranging from simple proportional control, through PID to a sum of first and second order transfer function elements. As a single compensator is constructed for each bearing and replicated for each axis of control for that bearing, possible controllers can take orders of any multiple of three states from zero to 36. In this way the MOGA is able to select both the structure and the parameters of the compensator.

This approach may initially appear to use a rather inefficient method of storing the information, as any one individual is likely to contain parameters that are not used. However it has the advantage that the genetic information associated with elements that are not currently active is preserved into future generations. In order to ensure that this preservation occurs over many generations, only the bits representing active parameters are subjected to the mutation operator. For example, if $S_1 = 0$ the the bits representing K_2 are not allowed to mutate and the probability of mutating the other bits (that are active) is recalculated such that the overall probability of the individual mutating is 0.5. The bits controlling the individuals' structure are always allowed to mutate.

4. The Optimisation Objectives

In order for the MOGA to rank the possible controllers, a function is required to evaluate the controllers' performances against a number of decision objectives. This multi-objective function first constructs the controller in state space and uses

this in conjunction with the state space model of the plant to analyse the stability of the closed loop system. Controllers that exhibit satisfactory stability characteristics are simulated on the non-linear model to evaluate the performance and efficiency objectives (1-8 in Table 1). Initial experiments used a discontinuous objective function to encourage the search to find solutions with stable eigenvalues. However, when a single stable solution was found it had a tendency to overwhelm the population within a few generations. This is because the probability of finding other stable solutions in this period is comparatively small due to the large size of the search space. Thus the stable solution faces no real competition for several generations during which time a great deal of population diversity is lost.

The solution to this problem adopted here is to introduce two extra objectives to encourage the optimisation to converge on a diverse set of solutions. The first extra objective is chosen to be the maximum real part of the eigenvalues of the closed loop system, λ_{\max} . The goal for this objective is set to be zero, to ensure stability. Compensators with closed loop eigenvalues very much greater than zero are not simulated on the non-linear system (the objective function returns infinity for objectives 1-9 and 11). Solutions with their eigenvalues close to, but greater than zero are simulated so that their performance on the non-linear system governs their ranking, not just the proximity of their maximum eigenvalues to zero. This non-linear system performance is recorded as the reciprocal of the length of simulation time, τ^{-1} , before the rotor moves unacceptably far from its desired position. The MOGA was configured to treat τ^{-1} and λ_{\max} as constraints, not optimisation objectives, so no attempt is made to reduce them beyond their constraint thresholds, $1/5.5$ for τ^{-1} and 0 for λ_{\max} . Any solution which fails to complete the simulation is ranked by these objectives alone. In this way, the search is encouraged to maintain a diverse population whilst moving to stable regions with good disturbance rejection characteristics.

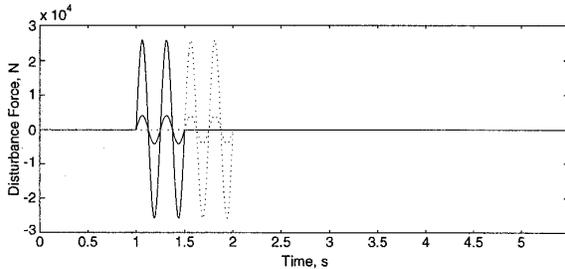


Figure 5: Disturbance forces

The controller's performance is evaluated by simulating the non-linear model of the system subjected to the disturbance signal shown in Fig. 5. This disturbance is composed of four signals, representing forces acting on the rotor at each bearing, in the X (dotted) and Y (solid) directions. The signal applied at bearing B is proportionally larger than that at A, as

B is intended to absorb most of the disturbance forces. This load distribution equates to a 30kN peak disturbance of a similar shape applied to the rotor's centre of gravity. The system is also subjected to a half second burst of sensor noise at 4s. This is composed of the first six harmonics of mains power transmission (the most prevalent noise source in the machine's operating environment), Fig. 6. This noise is applied equally to all the sensors and roughly equates to an uncertainty in the rotor's position of 5 microns.

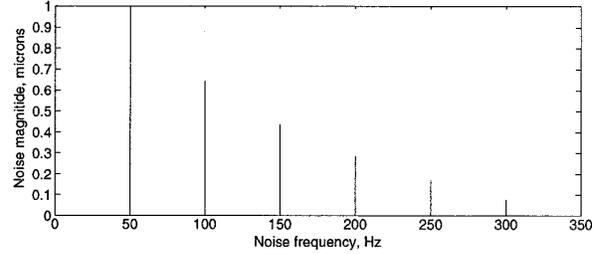


Figure 6: Sensor noise spectrum

The performance of a bearing is conventionally evaluated by measuring its stiffness. As the MOGA is a minimising optimiser, the compliance (the reciprocal of stiffness) was used for the actual optimisation objectives. This is calculated by measuring the peak displacement of the rotor at the disturbed portion of the response and dividing this by the peak force applied. A compliance of $6.66e-9$ m/N means that under a 30kN disturbance the rotor will move 0.2mm. This is the minimum standard of performance acceptable for this system. A similar measure is used for evaluating the noise rejection, η is the maximum rotor displacement divided by the maximum noise amplitude.

Objv. No.	Objective Description	Tag
1	Bearing A steady state error	e_{ssa}
2	Bearing B steady state error	e_{ssb}
3	Bearing A compliance at 4 Hz	C_a
4	Bearing B compliance at 4 Hz	C_b
5	Bearing A maximum current	I_{maxa}
6	Bearing B maximum current	I_{maxb}
7	Bearing A noise susceptibility	η_a
8	Bearing B noise susceptibility	η_b
9	Controller complexity (order)	N_k
10	Max. real closed-loop eigenvalue	λ_{\max}
11	1 / (length of simulation)	τ^{-1}

Table 1: The optimisation objectives

Table 1 shows the design objectives for the controlled rotor - bearing system, objectives 1 - 8 and 11 are determined by simulating the non-linear model in Simulink. The steady state

errors were measured after allowing the simulated system to settle for 1s after the sensor noise disturbance was applied. The maximum currents developed in the coils during the disturbances were used to give a measure of the efficiency of the controllers at each bearing. The complexity, N_k is simply the order of the complete compensator.

5. Results

Fig. 7 shows a typical trade-off graph for the AMB system. The x -axis shows the design objectives and the y -axis shows the objective domain performance of the controllers.

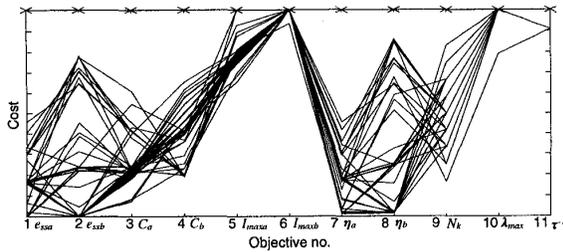


Figure 7: Sample trade-off graph

Trade-offs between adjacent objectives result in the crossing of lines whereas concurrent lines represent non-competing objectives. The 'x' marks represent the optimisation targets. For example, it can be seen that there is a sharp overall trade-off surface between the compliance performances of the two bearings, objectives 3 and 4, although there are some solutions that do well in both. Interestingly, objectives 5 and 6, the maximum currents in each bearing, do not appear to compete with each other as much as might be expected. This is because the majority of controllers cause the current in bearing B to saturate at maximum load. This is to be expected because the magnets were designed to just tolerate this load. The controllers shown here are of orders ranging from 6 to 30, and a wide variety of different structures, even controllers that have the same number of states, N_k , can be constructed completely differently. Objectives 10 and 11, the maximum real eigenvalue, λ_{\max} and the reciprocal of the simulation time, τ^{-1} are not relevant for the analysis. It is necessary, but sufficient, that they both meet their goals.

The arrangement of objectives shown in Fig. 7 is not conducive to useful analysis of the conflicts between performance and complexity, for example. Fortunately, the visualisation tool allows the objectives to be swapped around so that further insight into the trade-off surface may be gained. Fig. 8 shows a selection of objectives arranged symmetrically, with objectives relating to bearing A on the left, bearing B on the right and complexity in the middle. Objectives 3 and 4, the bearings' compliance, are the main measure of the bearings' dynamic performance and are therefore of most interest. Fig. 8 demonstrates that there is little correlation between compliance and the overall system's complexity, N_k , although the solutions with the best

compliance tend to be the most complex ones. There is little trade-off between compliance and the maximum current used by bearing A, objectives 3 and 5. This is attributed to the rate limit on the build up of current in the coils causing the rotor's motion under solutions that have good compliance being closer in phase to the limit of the power amplifier's performance, and therefore requiring less power overall to reject the disturbance. Comparing the noise rejection objectives 7 and 8 to the steady state error objectives 1 and 2 it can be seen that the noise causes a slight worsening of the steady state error in most cases, as might be expected.

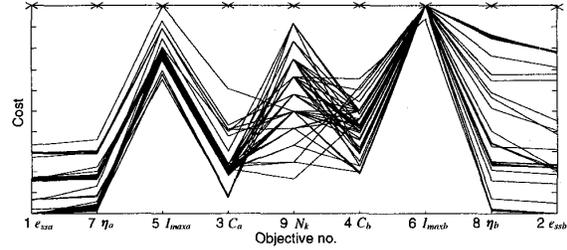


Figure 8: Rearranged trade-off surface

The task of the control systems designer is then to select a solution that best fits the requirements of the system under construction. For example, if it is desirable to minimise the complexity of the controller, some performance and efficiency may be sacrificed. This is the case for the controller, the rotor responses of which are shown in Fig. 9. There are four response traces super-imposed in the figure, the response of the rotor at bearing A, dotted, bearing B, solid, in both the vertical and horizontal directions. The disturbance is first applied in the vertical direction, then the horizontal, so the responses can be distinguished. This controller has 6 states and conforms to a conventional PD structure for each axis of each bearing. It has a compliance of approximately $6 \times 10^{-9} \text{m/N}$, and a steady state error of just over 0.1mm for the larger bearing. It rejects the noise at 4 seconds successfully, however the large time constant and steady state error for bearing B would rule it out in practice.

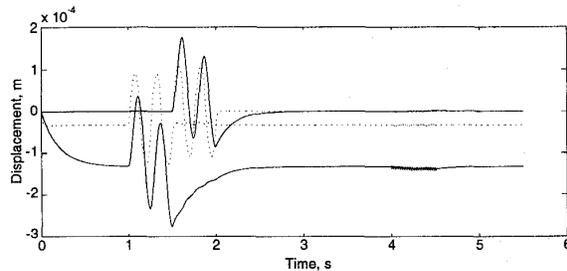


Figure 9: Rotor responses under a low order controller

To investigate the solutions with good steady state error characteristics, the goal settings for these objectives were reduced to 1 micron and the scales on the objective graph adjusted accordingly, Fig. 10. This reveals that no controllers under 18 states achieve this steady state performance and

satisfy the other objective goals. Fig. 11 shows the response of an 18 state controller with zero steady state error. It has a compliance of $5 \times 10^{-9} \text{m/N}$, good noise rejection and a fast settling time. It offers a practical solution for the application.

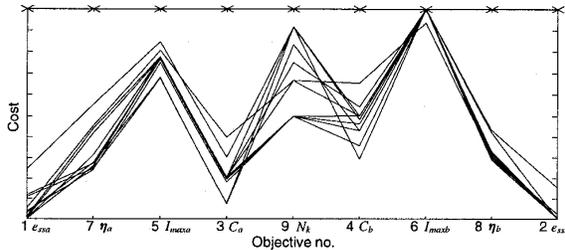


Figure 10: Trade-off surface for 1 micron steady state error

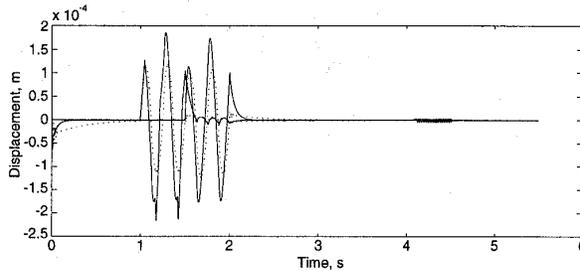


Figure 11: Rotor responses under an 18 state controller

6. Concluding Remarks

A multi-objective genetic algorithm is used as a design tool for generating distributed active magnetic bearing controllers. The MOGA is used to search the structure and parameter space of a distributed controller configuration. The optimisation is performed directly on a non-linear model of the complete system, simulated under worst case disturbance conditions. Several measures of AMB performance are used as objectives for the optimisation, from these a great deal of information about the limiting characteristics of the many possible controllers can be inferred. A large collection of satisfactory controllers is generated, from which it is possible to select an efficient controller with good disturbance and noise rejection characteristics. Some difficulty was encountered encouraging the MOGA to converge, due to the large size of the search space. As a practical design technique, better results could be achieved by deciding on a single controller structure before running the optimisation as this dramatically reduces the size of the search space.

The accuracy of the trade-off surface is increased by direct simulation on a non-linear system model, so a more informed choice of controller can be made. However this slows down the search and an appropriate compromise between accuracy of information and speed is desirable. In the end, it is the designer who makes the decision about what controller structure and parameters are to be used. The MOGA is simply used as an efficient way to explore the possibilities offered by each alternative.

7. Acknowledgements

The authors gratefully acknowledge the support of this research by Rolls - Royce & Associates Limited and UK EPSRC grants on "Evolutionary Algorithms in Systems Integration and Performance Optimisation" (GR/K 36591) and "Multi-objective Genetic Algorithms" (GR/J 70857). The authors also wish to thank Dr. Carlos Fonseca for the multi-objective extensions to the GA Toolbox.

8. References

- [1] V. Iannello - Synchrony Inc., "Magnetic bearing systems for gas turbine engines", *Mag'95*, Alexandria, Virginia, 1995.
- [2] G. Schweitzer, "Magnetic bearings - applications, concepts and theory", *JSME Int. J. Series III, Vol. 33, No. 1*, 1990.
- [3] P. Schroder, A. J. Chipperfield, N. Grum and P. Fleming, "Robust multivariable control of active magnetic bearings", *The European Control Conference*, Brussels, July, 1997.
- [4] C. Delprete, G. Genta and S. Carabelli, "Control strategies for the decentralised control of active magnetic bearings", *4th International Symposium on Magnetic Bearings*, ETH Zürich, pp. 29-34, Aug, 1994.
- [5] N. V. Dakev, J. F. Whidborne and A. J. Chipperfield, " H_∞ design of an EMS control system for a maglev vehicle using evolutionary algorithms", To appear *Proc. I. Mech. E. Part 1: Journal of Systems and Control*, 1997.
- [6] G. Schweitzer, H. Bleuler and A. Traxler, "Active Magnetic Bearings, basics, properties and applications", *ETH, Zürich*, 1994.
- [7] A. J. Chipperfield, P. J. Fleming and H. P. Pohlheim, "A genetic algorithm toolbox for MATLAB", *Proc. Int. Conf. Systems Engineering*, Coventry, UK, pp. 200-207, Sept., 1994.
- [8] C. M. Fonseca and P. J. Fleming, "An overview of evolutionary algorithms in multi-objective optimisation" *Evolutionary Computation*, Vol. 1, No. 1, pp 25-49, 1993.
- [9] D. E. Goldberg and J. Richardson, "Genetic algorithms with sharing for multimodal function optimisation", *Proc. 2nd Int. Conf. on Genetic Algorithms*, J. J. Grefenstette (Ed.), pp. 41-49, 1987.
- [10] K. Deb and D. E. Goldberg, "An investigation of niche and species formation in genetic function optimisation", *Proc. 3rd Int. Conf. on Genetic Algorithms*, J. D. Schaffer (Ed.), pp. 42-50, 1989.
- [11] S. Skogestad and I. Postlethwaite, "Multivariable feedback control, analysis and design", *John Wiley & Sons Ltd*, Chichester, England, 1996.
- [12] K. Nonami, W. He and H. Nishimura, "Robust control of magnetic levitation systems by means of H_∞ control / μ -synthesis", *JSME Int. J. Series C, Vol. 37, No. 3*, 1994