Multi-Agent Control of High Redundancy Actuation

J. Davies * **T. Steffen** * **R. Dixon** * **R.M. Goodall** *

* *Control Systems Group, Loughborough University, Loughborough, LE11 3TU, UK, http://www.lboro.ac.uk/departments/el/research/scg*

Abstract: The High Redundancy Actuator (HRA) project investigates the use of a relatively high number of small actuation elements, assembled in series and parallel in order to form a single actuator which has intrinsic fault tolerance. Both passive and active methods of control are planned for use with the HRA. This paper presents progress towards a multiple model control scheme for the HRA applied through the framework of multi-agent control.

1. INTRODUCTION

1.1 Fault Tolerant Control and Actuator Redundancy

A fault may be defined as a defect or imperfection that occurs in the hardware or software of a system. Faults in automated processes will often cause undesired reactions which could manifest as failures, where an expected action is not completed by the overall system. The consequences of failures could include damage to the plant, its environment, or people in the vicinity of that plant [Blanke et al., 2001]. Fault tolerant control aims to prevent failures and achieve adequate system performance in the presence of faults.

The majority of research to date has concentrated on sensor faults. Significant advances have been made in this area, but most of these strategies are not applicable to actuator faults. This is attributable to the fundamental differences between actuators and sensors. Sensors deal with information, and measurements may be processed or replicated analytically to provide fault tolerance. Actuators, however, must deal with energy conversion, and as a result actuator redundancy is essential if fault tolerance is to be achieved in the presence of actuator faults. Actuation force will always be required to keep the system in control and bring it to the desired state [Patton, 1991].

The common solution for critical systems involves straightforward parallel replication of actuators. Each redundant actuator must be capable of performing the task alone and possibly override the other faulty actuators. This over-engineering incurs penalties as cost and weight are increased and subsequently, efficiency is reduced.

1.2 High Redundancy Actuation and Multi-Agent Control.

The High Redundancy Actuator (HRA) concept is a novel approach to actuator fault tolerance, inspired by musculature. Muscles are composed of many individual cells, each of which provides a minute contribution to the force and the travel of the muscle. These properties allow the muscle, as a whole, to be highly resilient to individual cell damage.

The HRA project aims to use the same principle of co-operation to provide intrinsic fault tolerance using existing technology. To achieve this, a high number of small actuator elements are assembled in parallel and series to form one high redundancy actuator (see Figure 1).

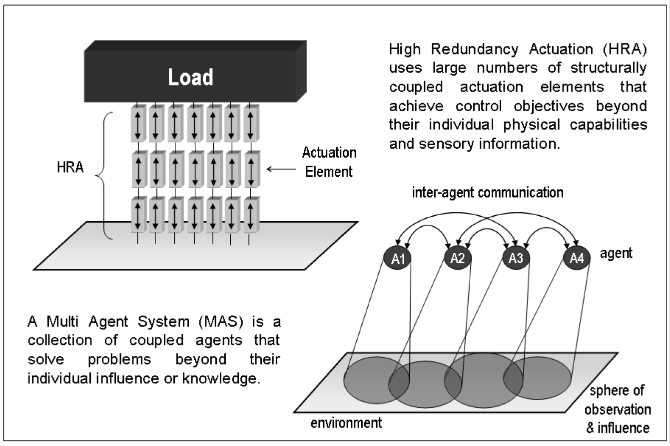

Figure 1. HRA and MAS.

Faults in elements will affect the maximum capability, but through control techniques, required performance can be maintained.

The HRA is an important new approach within the overall area of fault tolerant control. When applicable, it can provide actuators that have graceful degradation, and that continue to operate at close to nominal performance even in the presence of multiple faults in the elements.

The main focus of the HRA project thus far has utilised robust control techniques. These techniques have been shown to be theoretically viable for fault tolerant control of low levels of redundancy [Du et al., 2007], and the practical testing of these results on a two-by-two electromechanical rig is in its final stages [Du et al., 2008].

Electro-magnetic actuation is now being considered as a candidate element within the HRA, the modelling of which in both nominal and fault condition has been detailed in [Davies et al., 2008]. Research is ongoing into the robust control of these elements at higher levels of redundancy

[Steffen et al., 2007]. Results to date suggest that robust control should be a satisfactory method of achieving fault tolerant control of these structures. Indeed, the robust, passive [1] control approach is attractive, as its simplicity and constancy mitigate many of the associated problems with active control methods. However, research into more intelligent, active approaches is also an objective of the HRA project, to ascertain the levels of fault tolerance and nominal performance attainable in comparison to passive methods.

Multi-Agent Systems (MAS) are the focus of this active fault tolerance scheme. MAS was chosen as an intelligent approach to controlling the HRA as the two concepts are strongly related (Figure 1). Both are inspired by natural mechanisms which utilise vast numbers of relatively simple cells/processes to form complex structures/behaviours.

1.3 Overview

This paper presents the concepts and objectives of applying MAS to an electro-magnetic HRA with an example of Multi-Agent Control (MAC) applied to a parallel-series (PS) HRA. Section 2 briefly introduces agent concepts and discusses the rationale behind MAC of HRA. The current MAC scheme is described in Section 3, from both an agent and agency architecture perspective. Section 4 details simulation results of a MA controlled 4x4 HRA in comparison to a global control scheme with the injection of faults into the system The further development of the multi-agent scheme is discussed in Section 5.

2. MULTI-AGENT CONTROL OF A HIGH REDUNDANCY ACTUATION (MACHRA)

2.1 Concepts of Multi-Agent Systems

An agent is a physical or virtual entity situated in its environment, which acts autonomously and flexibly within its purview to achieve goals in a real-time manner [Jennings et al., 1998]. A MAS therefore, is a collection of agents that are socially coupled and collaborate to achieve some objective, which in the case of MAC is the control of a system.

These agent characteristics resemble the concept of closed-loop control, which achieves objectives through sensing and acting. However, there are important differences within the agent concept. The most obvious difference is the social interaction and negotiation that exists between agents. Also, the agent philosophy is strongly associated with localisation, a point emphasised by [Ferber, 1999].

According to [Weiss, 1999], agent concepts are most beneficial and applicable in applications that have one or more of the following attributes:

Modularity/Decentralisation - A physically or functionally modular system is naturally identifiable with an agent structure. Similarly, agents are useful in a decentralised system, as they may be associated with distributed subsystems and their pro-active capabilities allow low-level decisions to be made locally, facilitating the management of large systems.

Changability/Ill-structure - The modular and decentralised nature of agents allow the structure of the agency, or agents themselves, to be changed with minimum impact to the system, providing a robust adaptable solution. This is important in systems that are likely to change frequently, or ill-structured systems where the domain structure is not completely specified or static.

Complexity - A complex system, with many interacting elements and behaviours, can be served well by an agent approach, as problems may be solved in a more efficient and timely manner.

2.2 Why Take a Multi-Agent Approach to HRA Control?

In addition to the inherent similarities of HRA and MAS, the HRA is likely to benefit from a multi-agent approach as it displays many of the properties in Section 2.1.

The key rationale for combining MA concepts with HRA, however, is the structuring of both of these concepts. The HRA, viewed as a whole is a complex system, but if viewed as a collection of simpler, similar (if not identical), physically distributed modules, the complexity and changeable nature of the system's dynamics and structure can be handled at a local level, allowing objectives to be met with greater speed and efficiency. MASs facilitate the control of such decompositions, and due to their communicative and flexible qualities, potentially provide greater robustness and adaptability in fault situations.

The structuring of control is often neglected in the field of control engineering as the problem is stated in the form of a single plant model [van Breemen and de Vries, 2000]. The HRA is a complex, highly structured system, with well defined interactions between simple elements. An unstructured approach will have difficulties dealing with this complexity. MAC can replicate the structure of the HRA, which should simplify the individual control algorithms. The process industry acknowledges that the structuring of control is an important issue when applied to a decomposed system, thus it is given more attention in this field and numerous MACS have been proposed in this application area e.g. [Wang and Wang, 1997] .

The actual control technique implemented in the agent is peripheral to the MAC structure. Classic or modern designs based on multiple model approaches can be implemented, with the MA concept providing the mechanism for intelligently deciding which controller to employ locally. Adaptive controllers could also be applied, again with the agents providing the decomposition of the problem.

Essentially, agent methods provide a framework to apply localised active control, fault detection and health monitoring to the HRA, whilst avoiding some of the issues associated with active control. Multiple model control schemes often have one active global controller, and a supervisor that decides which controller should be active. This centralisation can create problems with bumpless transfer, as large control signal changes can occur when switching between schemes and the supervisor becomes a single point of failure, increasing the systems reliance upon

[1] In this context, passive refers to a static control structure and algorithm.

fault detection. In addition to this, a global view on the system can make faults more complex to diagnose. These centralisation issues are negated by MAC, as are issues associated with adaptive control.

The unpredictability of centralised adaptive control schemes should be alleviated somewhat by the decentralisation MAC offers. Undesirable changes within modules will have less affect on the system as a whole, perhaps even with other agents adapting to counter-balance the unwanted behaviours. Localisation of control may also improve on response speed issues associated with adaptive control.

Nonetheless, there are a number of potential issues associated with MASs that require careful attention such as deliberation, communication and negotiation delays, agent non-consensus and communication failure.

2.3 MACHRA Objectives

The objectives for the use of MAC in this project include those made for the control of the HRA with robust techniques, namely:

- Control of the elements resulting in a unified dynamic for the HRA.
- Nominal or acceptable behaviour of the HRA in element fault conditions.
- Graceful degradation of the HRA as fault levels increase beyond their critical point.
- Health and capability monitoring of the elements for maintenance/operator use.

If the inclusion of intelligence within the control scheme is to be justified then the MA controlled HRA must achieve tangibly more in comparison to passive methods. Thus, the objective for MAC of an HRA also include:

Increased reliability - Robust techniques can be limited in the number of faults or fault types they can accommodate. The structure of the HRA alleviates this problem somewhat, as the number of elements reduces the overall affect of faults on the system. Nevertheless, a more intelligent scheme, such as multi-agents, may accommodate even greater fault levels and fault types.

Improved nominal performance - Passive fault accommodation methods require the controller design to be robust enough to produce adequate performance during faulty conditions. This can lead to conservative performance in nominal conditions. An active control scheme can offer an increase in nominal performance as the control action can be changed in fault situations. Agent schemes may also provide performance enhancement due to their potential to pre-empt situations.

3. MACHRA SCHEME

The MACHRA scheme is currently in the investigative stage, concentrating on parallel in series (PS) configurations with lock-up and loose faults. Initial agent architectures and agency structures have been designed and simulated.

At present, Matlab/Simulink is used to create and simulate HRA assemblies, details of which can be found in [Davies et al., 2008]. Stateflow is used to simulate the inner rule-based logic of the agents and their communication. This provides a fast prototyping tool of the agents for use with Matlab/Simulink.

3.1 Agency Architecture

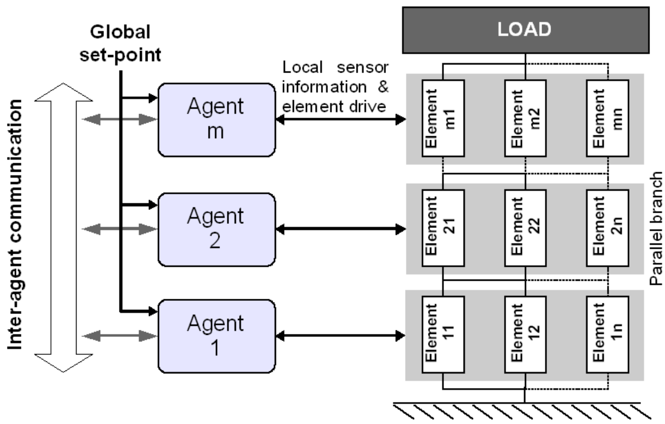

Figure 2. MACHRA agency architecture

The term agency architecture refers to the configuration of multiple agents on the macro scale. Figure 2 displays the MACHRA scheme's agency architecture for a $m \times n$ HRA PS configuration. There is an agent per parallel branch of elements, each of which is responsible for the control and detection of faults within its elements. In this configuration, lock-up faults will render the parallel branch inoperative, adding that branch's weight to the load of its neighbouring branches. However, loose faults will not affect the travel capability of that branch, as long as one operational element remains: PS assemblies have inherent fault tolerance to this fault type. Different configurations will provide more or less inherent tolerance to these faults, hence the configuration must be chosen to suit the application.

All agents within this scheme are identical and peers, consistent with the spirit of MAC where no hierarchy should exist. A global set-point for the whole HRA is given to each agent, as well as local position sensory input from its branch. Communication between agents is broadcasted via a bus. However, agents only consider messages from structural neighbours. If lock-up faults occur, the agent's structural neighbours will change and thus different messages become relevant.

3.2 Agent Architecture

The current agent architecture is illustrated in Figure 3. This architecture has similarities with subsumption, first introduced by [Brooks, 1986], that uses behaviours layered in order of abstraction to produce more complex emergent behaviors in a reactive time-frame. This reactivity is key in the HRA as, due to the fast dynamics of the electromagnetic elements, a purely deliberative architecture may not provide the response times needed.

The most reactive, basal behaviors are situated on the bottom layer, in this case the Control Module (CM), which provides the drive signal to the element based on the global

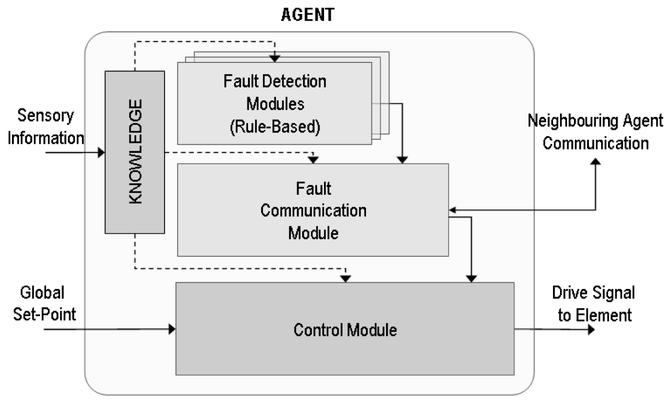

Figure 3. MACHRA agent architecture

set-point. A multiple model control scheme is employed within this design as the CM contains a look-up table with simple classical control designs based on the number of active agents structurally surrounding the agent e.g. more control action is required if an element's neighbour locks.

The Fault Detection Module (FDM) is the most abstracted layer, and thus affects those below it. As its name suggests, the FDM detects faults within its element. At the current stage of development, one fault type (lock-up faults) is detected. Future agents will have more than one module, arranged either as peers in a single layer or as separate layers ordered by the severity of the fault type. The module contains rule-based logic which determines the fault status of the element based on sensory information and internal knowledge.

The middle layer is the Fault Communication Module (FCM). This module communicates the fault status and global capability estimate to other agents, receiving messages of the same content from other agents. This information is also passed to the CM, where it is used to choose a controller and decide what portion of the overall global set-point to make its objective. In the absence of this information, the CM assumes nominal conditions if no previous communication with the FCM has been made, and last known conditions in the presence of communication history.

One shortcoming of the subsumption architecture, according to [van Breemen and de Vries, 2000], is its inability to combine information from different layers. This problem is negated in the MACHRA agent structure, as information is stored within an inter-accessible knowledge module. This module contains both knowledge given to the agent on start-up and that deduced within the individual modules.Another commonly cited inadequacy of the subsumption scheme, is the lack of consideration of previous events. This is not the case within the MACHRA architecture, as the modules have previous state-based behaviours, making the agent not purely reactive, but hybrid.

4. MAC VS GLOBAL CONTROL SIMULATIONS

An example that illustrates the potential of MAC of HRA, in comparison to centralised passive control, is provided in this section. The HRA system chosen is a 4x4 PS configuration with an overall travel control objective. The

Table 1. Fault Cases

Case	Description	HRA State
Nom.	All elements are healthy	Healthy & capable
F1	Branch nearest load locked	Faulty, but capable
F2	2 branches nearest load locked	Faulty and critical
F3	3 branches nearest load locked	Graceful degrade

Table 2. Requirements

Performance Requirements	
Travel Window	$\pm 0.015\text{m}\,(2 \times element\,travel)$
Overshoot	$<2\%$
Rise Time	<0.5s
Settling Time	<0.75s

elements work upon a load that is twice as large as the inter-element masses and for the purposes of this example it is assumed that the HRA is over-dimensioned by a factor of two i.e. the maximum required travel is twice the travel of a single element [2] .

As the PS assembly is naturally tolerant to loose faults in terms of travel control, they will not be considered here. However, element lock-ups immobilise the parallel branches, and thus will be considered. Theoretically, a 4x4 system, of this dimensioning, may incur up to eight lock-up faults and still be capable of meeting its travel requirement. However, in a worst-case scenario, where single lock-ups occur in different branches, two lock-ups will bring the travel capability to critical point. Hence, faults will be injected in this worst-case manner, as described in Table 1.

4.1 Control Schemes

Figure 4 portrays the global and MAC control schemes. In the global scheme, a single, phase advance controller is designed to meet arbitrary transient requirements, displayed in Table 2, with good stability margins

Two potential MAC control schemes are included in this paper. The first scheme, MACS1 has set-point redistribution via gain scheduling only. Each agent has a phase advance controller [3] designed to meet the requirements in the nominal case. In fault conditions, the reduction in capability is communicated to other agents, and the extra travel required is distributed amongst the remaining active elements, minus the lock-up position of the faulty element. The control algorithm in each agent remains unchanged.

The second scheme, MACS2 utilises the same control algorithms as MACS1 under nominal conditions. The set-point redistribution described in MACS1 is also present. However, in addition to this gain scheduling, another controller with a faster dynamic per fault condition is designed and implemented when faults occur. Hence, this scheme illustrates set-point redistribution with a multiple model control scheme, as the controllers are based upon a series of locked fault models with increased load.

[2] This is an unrealistically low element mass-to-load ratio and high level of over-dimensioning for a HRA due to the relatively low level of redundancy used in this example.

[3] The controllers in each agent are identical, as this reduces design time and aids verification for high integrity applications.

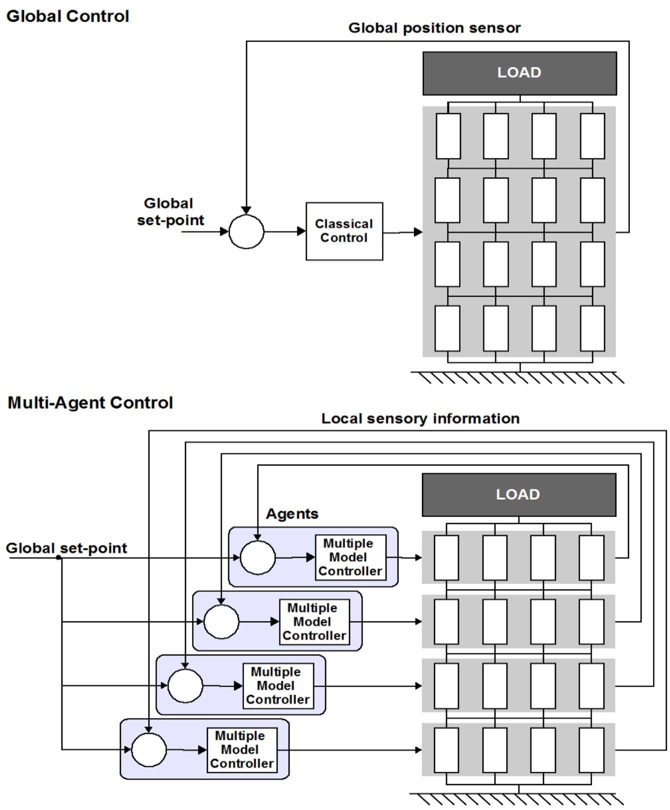

Figure 4. Global and agent control schemes

Table 3. Global control, MACS1 and MACS2 simulations

Fault Case	Overshoot	Rise Time	Settling Time
Glob. Nom.	1.80%	0.30s	0.45s
Glob. F1	0.37%	0.40s	0.62s
Glob. F2	0%	0.63s	1.14s
Glob. F3	0%	1.38s	2.48s
MACS1 Nom.	1.83%	0.30s	0.44s
MACS1 F1	0.95%	0.30s	0.45s
MACS1 F2	0%	0.30s	0.54s
MACS1 F3	0%	0.36s	0.73s
MACS2 Nom.	1.83%	0.30s	0.44s
MACS2 F1	1.01%	0.30s	0.45s
MACS2 F2	0.42%	0.30s	0.46s
MACS2 F3	0%	0.30s	0.49s

4.2 Comparison of Control under Fault Conditions

All faults were injected at t=0 and the set-point was an attainable travel in the worst fault case. Figure 5 shows a step response of the three control schemes under these conditions and their characteristics are summarised in Table 3. The simulations show that as faults occur, the increasing load slows the response.

Comparing the global scheme to MACS1, it can be observed that the transient performance of the globally controlled HRA degrades to a greater extent than that of the MA controlled system with input redistribution only. The rise time and settling time increase significantly with each fault in the global case, whereas very little change is observed until three of the four branches are locked in the MAC case. This result illustrates that the localisation of control in this manner is favourable in comparison to the global approach.

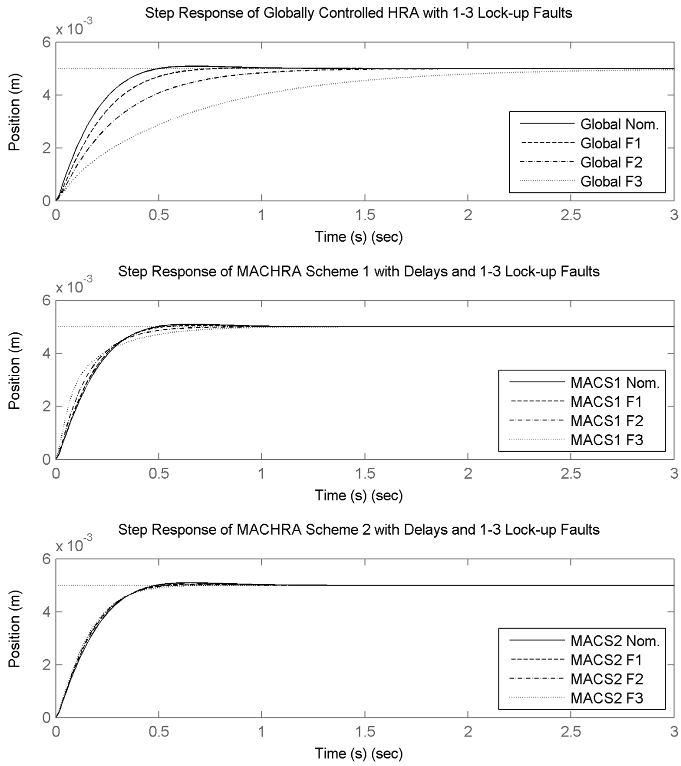

Figure 5. Step response for global control, MACS1 and MACS2

The results shown in Table 3 for MACS2 are an improvement on those for MACS1. The rise time does not lengthen under any fault condition, and the change in settling time is significantly reduced. In addition to these marginal improvements, it can be seen from Figure 5, that the transient response envelope is tightened in general when the control algorithm is changed under fault conditions.

The MACHRA schemes simulated in this example offer significantly improved performance under fault conditions in comparison to a simple global control technique. However, it must also be acknowledged that the fault levels and over-dimensioning present in this example are much higher than those conceived by the HRA concept, and less distinction between the control schemes will be present with higher order configurations. Ultimately, the necessity for inclusion of active control strategies such as MAC will be dictated by the stringency of the requirements for a specific application.

4.3 Fault Detection and Control Reconfiguration Delays

The results discussed in Section 4.2 assumed that fault detection and control reconfiguration in the MA schemes was instantaneous, which is unrealistic. If Stateflow is used to simulate the multi-agent type fault detection, communication and control reconfiguration then a delay is introduced, providing a more realistic representation.

Figure 6 is a pulse train simulation of the previous MAC scheme's, simulated with Stateflow. As previously, all faults were injected at t=0 and an attainable travel demand was made.

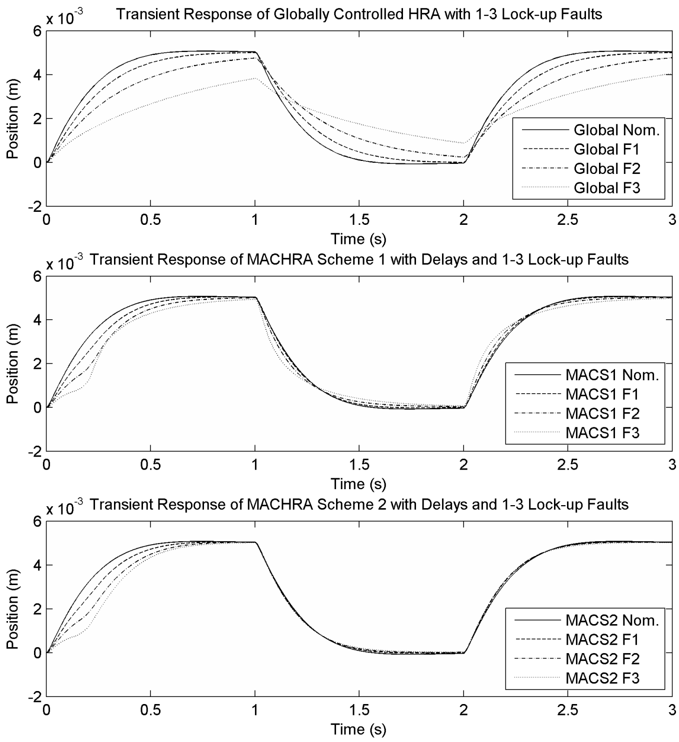

Figure 6. Pulse train for global control, MACS1 and MACS2 with delays

The reconfiguration of the MACSs can be observed in the first pulse rise, resulting in a slower response. This is particularly pronounced where three faults need to be detected and communicated. However, in the remaining operational period, where no faults occur, there are no delay affects, and as such the response obtained in the previous section applies. This ephemeral performance degradation on fault injection may be considered acceptable in a real application, as faults are likely to be an infrequent event.Planned Developments

Many developments are planned for the MACHRA. These include the improvement and extension of current fault detection schemes, and control techniques. The configurations to which the MAC is applied must also be extended to series-parallel and mixed configurations.

Changes to the architecture of the MAC scheme may be investigated and robust testing procedures are required. Issues regarding feasibility at very high levels of redundancy must also be addressed.

The practical testing of MAC on a experimental electro-magnetic HRA is also planned, which should give an indication of such a scheme's performance in a real-world situation.

5. CONCLUSIONS

MAC potentially provides an active fault tolerant solution to controlling the HRA that improves on the nominal and faulty performance of passive schemes, whilst negating some of the issues associated with control reconfiguration. The simulation of a suggested agent/agency architecture was included within an example that compared MAC with a more traditional global scheme. This example illustrated that requirements can be met under greater fault levels with the assistance of MA concepts, using input redistribution alone or in conjunction with localised multiple model control.

Many further developments for the MACHRA are planned as well as the practical testing of MAC on a experimental electro-magnetic HRA, which will help give an indication a MAC schemes feasibility for this application.

ACKNOWLEDGEMENTS

This project is a cooperation of the Control Systems group at Loughborough University, the Systems Engineering and Innovation Centre (SEIC), and the actuator supply SMAC UK Ltd.. The project is funded by the UK's Engineering and Physical Sciences Research Council (EPSRC) under reference EP/D078350/1.

REFERENCES

M. Blanke, M. Staroswiecki, and N. E. Wu. Concepts and methods in fault-tolerant control. *American Control Conference, Proceedings of the 2001*, 4, 2001.

R.A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2(1): 14–23, 1986.

J. Davies, T. Steffen, R. Dixon, R. M. Goodall, A. C. Zolotas, and J. Pearson. Modelling of high redundancy actuation utilising multiple moving coil actuators. *IFAC World Congress, Seoul*, 2008.

X. Du, R. Dixon, R. M. Goodall, and A. C. Zolotas. Lqg control of a highly redundant actuator. In *Conference for Advanced Intelligent Mechatronics*, 2007.

X. Du, R. Dixon, R.M. Goodall, and A.C. Zolotas. Assessment of strategies for control of high redundancy actuators. *Actuator 08, Bremen*, 2008.

J. Ferber. *Multi-agent systems: an introduction to distributed artificial intelligence.* Addison-Wesley, United States, 1999.

N. R. Jennings, K. Sycara, and M. Wooldridge. A roadmap of agent research and development. *Autonomous Agents and Multi-Agent Systems*, 1(1):7–38, 1998.

R. J. Patton. Fault detection and diagnosis in aerospace systems using analytical redundancy. *Computing and Control Engineering Journal*, 2(3), 1991.

T. Steffen, J. Davies, R. Dixon, R. M. Goodall, and A. C. Zolotas. Using a series of moving coils as high redundancy actuator. In *IEEE Conference for Advanced Intelligent Mechatronics*, Zurich, 2007.

A. van Breemen and T. de Vries. An agent-based framework for designing multi-controller systems. *International Conference on The Practical Applications of Intelligent Agents and Multi-Agent Technology, Manchester, 2000*, 2000.

H. Wang and C. Wang. Intelligent agents in the nuclear industry. *IEEE Computer*, 30(11):28–34, 1997.

G. Weiss. *Multiagent Systems A Modern Approach to Distributed Artificial Intelligence.* The MIT Press, United States of America, 1999.