Adaptive Control of a
High Redundancy Actuator
Using The Geometric Approach

Thomas Steffen, Argyrios Zolotas, Roger Dixon, Roger Goodall

Control Systems Group,
Department of Electronic and Electrical Engineering,
Loughborough University, Loughborough, LE11 3TU, UK,
http://www.lboro.ac.uk/departments/el/research/scg/
t.steffen, r.dixon, r.m.goodall, a.c.zolotas@lboro.ac.uk

Abstract: : The High Redundancy Actuator project deals with the construction of an actuator using many redundant actuation elements. Whilst this promises a high degree of fault tolerance, the high number of components poses a unique challenge from a control perspective, especially when actuation elements are used in series.

This paper describes how an adaptive control scheme can be used to deal with faults in a High Redundancy Actuator. This is based on previous results leading to a simplified model of the HRA with serial elements. In case of the fault, the parameters change, but the otherwise the deviation from the simplified model is minimal.

This approach has two benefits. For one, it can restore the original system dynamics even after a fault has occurred. The parameter estimate can also be used for health monitoring purposes, because it reflects the number of effective faults in the HRA.

Keywords: high redundancy actuator, fault-tolerant control, active fault tolerance, fault accommodation, adaptive control, parameter estimation, health monitoring, geometric approach.

1. HIGH REDUNDANCY ACTUATION

High Redundancy Actuation (HRA) is a new approach to fault tolerant actuation, where an actuator comprises a large number of actuation elements (see Fig. 1). Faults in the individual elements can be accommodated without resulting in a failure of the complete actuator system.

The concept of the High Redundancy Actuation (HRA) is inspired by the human musculature. A muscle is composed of many individual muscle cells, each of which provides only a minute contribution to the force and the travel of the muscle. The aim of this project is to use the same principle of co-operation with existing actuation technology to provide intrinsic fault tolerance.
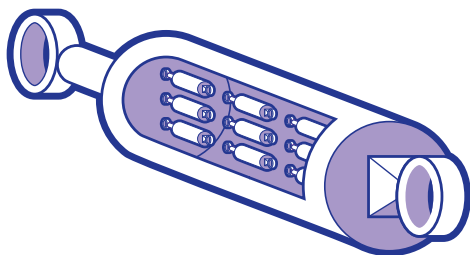

Fig. 1. Configuration of a High Redundancy Actuator

An important feature of the High Redundancy Actuator is that the actuator elements are connected both in parallel and in series. This makes it possible to deal with the two main fault models: lock-up and loose elements. In case of a lock-up fault, the available travel is slightly reduced, but the elements in series can still move. If an element fails loose, this reduces the maximum force, but the elements in parallel are still effective.

Responding to faults in the HRA can still be a challenge, because of the complexity of the system. While parallel elements can be considered as one mass, serial elements have to be modelled as several masses, and this leads to high order system. This paper shows how this complexity can be avoided by using an adaptive control scheme. Instead of detecting the faults explicitly, this scheme can adjust the control parameters implicitly, using a significantly simpler model. This is possible because under certain circumstances, the HRA can be controlled using a simplified model, and faults are reflected in parameter changes (not structural changes). Unlike a robust controller (Steffen et al., 2008b), the adaptive approach (Astolfi et al., 2008) can compensate the effect of faults on the closed loop behaviour. So the original behaviour can be restored without the complexity inherent in more detailed implementation of fault detection and reconfiguration schemes (Davies et al., 2008a). Despite the simplicity, the parameter estimation also can be used for health monitoring of the HRA.
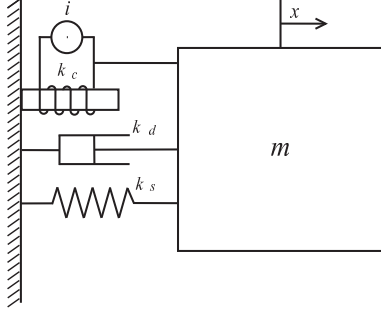
Fig. 2. Electromagnetic Actuation Element

Section 2 introduces the simplified model of the HRA. Sections 3 and 4 describes the design of the geometric controller and the parameter estimator. The results of this approach are compared with non-adaptive control in Section 5, leading to the conclusion and outlook in Section 6.

2. SIMPLIFIED MODEL OF THE HRA

The HRA considered here uses direct electromagnetic actuation (voice coil principle). Other technologies are also possible, but they may lead to a slightly more complex model.

Single Element

A single actuation element (Fig. 2) behaves like a spring-damper combination

$$m_1\ddot{x} = k_1 u_1 - d_1\dot{x} - r_1 x \quad , \qquad (1)$$

where m_1 is the moving mass, k_1 is the input coefficient, d_1 is the damping factor (accounting for mechanical and electrical damping), r_1 is the elasticity of the spring, u_1 is the input signal, $\ddot{x}_1$ is the speed and x_1 the position of the mass. For the state space model, the system state

$$\mathbf{x} = (\dot{x}_1 \ x_1)^T$$

is used. The parameters for the model are

$$m_1 = 0.2\,\mathrm{kg}$$
$$d_1 = 10\,\frac{\mathrm{Ns}}{\mathrm{m}}$$
$$r_1 = 1\,\frac{\mathrm{N}}{\mathrm{m}}$$
$$k_1 = 10\,\frac{\mathrm{N}}{\mathrm{V}} \quad .$$

Further details on the modelling of the actuator can be found in (Davies et al., 2008b).

HRA Model

As soon as elements are used series, each element becomes its own moving masses, and needs to be modelled individually. So the detailed model contains further states (Fig. 3), leading to a high order model of the form

$$\ddot{x}_i = f(\dot{x}_{i-1}, x_{i-1}, \dot{x}_i, x_i, \dot{x}_{i+1}, x_{i+1}) + k_i u_i + k_i' u_{i+1} \quad .$$

However, as shown in (Steffen et al., 2008a), it is possible to align the dynamics of all masses, so that they move at strictly related speeds and distances according to

$$\frac{\dot{x}_i}{\alpha_i} = \frac{\dot{x}_j}{\alpha_j}$$

and

$$\frac{x_i}{\alpha_i} = \frac{x_j}{\alpha_j} \quad .$$

This is achieved by increasing the force create by the bottom elements, to compensate for the slightly higher load, because they have to move the top elements in addition to the load.

In the frequency domain, this leads to the cancelling of all but two poles with input-decoupling zeros. Then the connecting masses have no influence on the input-output behaviour, and only the behaviour of a single mass system (dominated by the load mass) remains. So Equation (1) can be applied to the whole system, if the parameters are adjusted accordingly:

$$m\ddot{x} = ku - \frac{1}{n}(d\dot{x} - rx)$$

where n denotes the number of elements. For the example application, $n = 4$ and intermediate masses of $m_1 = m_2 = m_3 = 0.2\,\mathrm{kg}$ are used, together with a load mass of $m_4 = 1\,\mathrm{kg}$. The effective mass is slightly larger than the load, but this is compensated by the parameter tuning. The overall spring constant and damping are just one quarter of those of an individual element, because each element contributes only one quarter of the load speed and distance. This leads to the SISO transfer function

$$G_0(s) = \frac{1}{(s + 2.4)(s + 0.104)} \qquad (2)$$

for the nominal system.

Behaviour with Faults

This paper only deals with lock-up faults, because they have a more significant influence on the dynamics than other fault modes. While the effect of a loose fault can be compensated locally by using parallel elements, a lock-up fault invariably changes the structure of the system, and therefore of the model. This is because a lock-up of an element means that the two masses it links move at the same speed, so they can be modelled as one mass (Fig. 4). For example, a lock-up of the second element leads to

$$G_{F2}(s) = 10\frac{s + 134.5}{s + 134.1}\frac{s + 41.6}{s + 41.1}\frac{1}{(s + 3.3)(s + 0.103)} \quad , \tag{3}$$

and a double fault in elements 2 and 3 leads to the transfer function

$$G_{F23}(s) = 10\frac{1}{(s + 5)(s + 0.102)} \quad . \qquad (4)$$

There are several changes introduced by the faults. The main difference is the change of the faster pole from 2.4 to 3.3 or 5. This is expected, because with one element immobilised, the HRA is less flexible, leading to higher damping and stiffness. The other changes are an incompletely pole-zero cancellation, a slight change of the slower pole, and
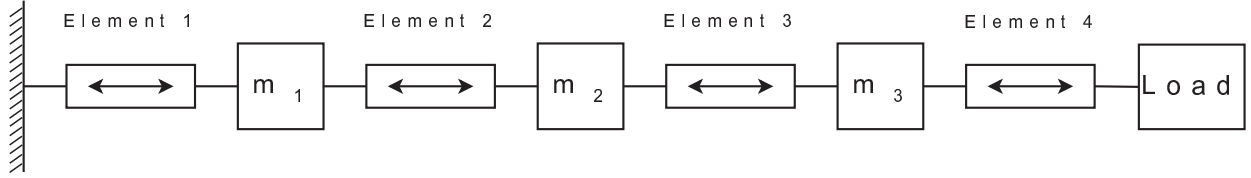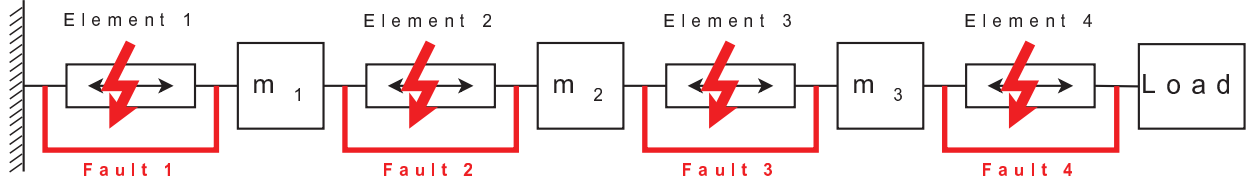
Fig. 3. Four Actuation Elements in Series

Fig. 4. Lock-up faults

a different amplification for some fault cases. These have only a minimal effect on the open loop behaviour, and are negligible in the closed loop behaviour.

All transfer functions were derived by linearisation from a detailed MATLAB/Simulink model of the system, and then verified manually. Unfortunately the process is too lengthy for inclusion in this paper.

Parametric Model

For the design of the adaptive controller, the minor changes in the transfer function will be disregarded. Instead, the adaption focuses on the change in the faster pole. The following parametrised system model is used:

$$\ddot{x} = f(\mathbf{x}, u, \theta) = \frac{1}{m}\left(ku - \theta(d\dot{x} + rx)\right) \quad . \qquad (5)$$

This is an approximation of the models shown above, when $\frac{1}{\theta}$ represents the number of working elements in series. So the nominal value is $\theta = \frac{1}{4}$, and it increases to $\frac{1}{3}$, $\frac{1}{2}$ etc with each lock-up fault.

3. ADAPTIVE CONTROLLER DESIGN

The design of the adaptive controller design follows the geometric approach proposed in (Astolfi et al., 2008; Ortega et al., 2003; Astolfi and Ortega, 2003). It consists of two main steps: the assignment of target dynamics to the system, and the choice of an attractive and invariant manifold for the parameter to estimate. Both steps are independent, so they can be performed in any order. Here, the target dynamics are handled first.

Target Dynamics

The control objective is fast convergence without overshoot. The target dynamics reflect this by using a double pole pair at -20:

$$G_{\text{target}}(s) = \frac{1}{\left(\frac{1}{20}s + 1\right)^2} \quad .$$

In state space, this leads to

$$\ddot{x} = f_t(\mathbf{x}) = (40 \ 400)(\mathbf{x}_{\text{ref}} - \mathbf{x}) + \ddot{x}_{\text{ref}}$$

where $\mathbf{x}_{\text{ref}}$ and $\ddot{x}_{\text{ref}}$ are state and acceleration of the reference trajectory. The target dynamics can be turned into a controller by solving the system model

$$f(\mathbf{x}, u, \theta) = f_t(\mathbf{x})$$

for u:

$$u = \frac{1}{k}\left(\theta(d\dot{x} + rx) + mf_t(\mathbf{x})\right) \quad .$$

This is a state feedback controller, so it requires knowledge of the speed $\dot{x}$ in addition to position x. The speed can be measured or observed, but in the latter case it is important that the observer does not depend on the unknown parameter θ. One way to achieve this is to measure the acceleration $\ddot{x}$. Then $\dot{x}$ can be derived from x and $\ddot{x}$ using either two complementary filters or a state observer, without resorting to the difficult to implement integration or differentiation operations.

Reference Trajectory

If matching trajectories are used such that $\ddot{x}_{\text{ref}}$ is the derivative of $\dot{x}_{\text{ref}}$, which is a derivative of x_{ref}, the system will follow these exactly. This is a significant advantage over conventional controllers, because it means that the disturbance response and the input response can be designed separately, leaving a lot of freedom in the design the reference trajectory. The only requirement is that the second derivative $\ddot{x}_{\text{ref}}$ in bounded, and within the physical capability of the system.

For the example application, a piecewise defined reference trajectory will be used for each step, consisting of three phases:

(1) 25 ms of constant acceleration,
(2) 100 ms of constant speed, followed by
(3) 25 ms of constant deceleration.

With this trajectory, the acceleration remains within reasonable bounds, the final position is reached, and the system settles after only 150 ms (see Fig. 5).

4. ESTIMATOR DESIGN

The second step in the design process is the design of the parameter estimator. The estimator is an inherently non-linear system, and therefore it will be performed exclusively in state space.
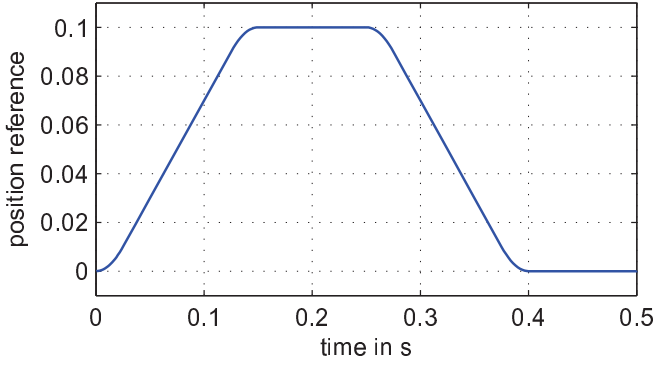
Fig. 5. Reference Trajectory

Parameter Estimate

The key idea of the geometric approach for adaptive control is that the parameter is not estimated directly. Instead, a non-linear transformation is applied to the parameter, and the resulting value is used as a state. With θ_e being the estimate, the transformed value is

$$\hat{\theta} = \theta_e - \beta(\mathbf{x}) .$$

There are several possible choices for the function $\beta(\mathbf{x})$. For this application, a quadratic form

$$\beta(\mathbf{x}) = \mathbf{x}^T \mathbf{Q} \mathbf{x} ,$$

is used, because it is easy to analyse. The symmetric matrix $\mathbf{Q}$ is the central design parameter for the estimation. Experience shows that a quadratic function works well for mechanical systems, but other choices are also possible. The parameter estimate can be easily calculated by inverting the transformation:

$$\theta_e = \hat{\theta} + \mathbf{x}^T \mathbf{Q} \mathbf{x} . \tag{6}$$

Estimator Invariance

The goal is of course for the estimate θ_e to approach the real parameter θ, and then to stay there. Both conditions are considered separately, starting with the later one. It is equivalent to $\theta_e = \theta$ being an invariant manifold. Because the real parameter θ is assumed to be constant, this can be achieved by finding parameters such that

$$\dot{\theta}_e = 0$$

as long as

$$\theta_e = \theta \quad .$$

This condition can be satisfied by taking the derivative of (6) and solving for

$$\dot{\hat{\theta}} = -\beta'(\mathbf{x}) \begin{pmatrix} f_t(\mathbf{x}) \\ \dot{x} \end{pmatrix}$$

$$= -2\mathbf{x}^T \mathbf{Q} \begin{pmatrix} f_t(\mathbf{x}) \\ \dot{x} \end{pmatrix} .$$

Combining the two differential equations leads to the extended system

$$\ddot{x} = f(\mathbf{x}, u, \theta)$$
$$\dot{\hat{\theta}} = -2\mathbf{x}^T \mathbf{Q} \begin{pmatrix} f_t(\mathbf{x}) \\ \dot{x} \end{pmatrix} .$$

Inserting the parameters and functions leads to the complete model

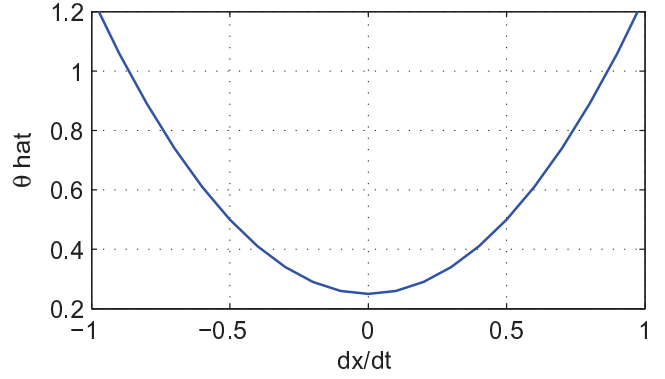Fig. 6. $\hat{\theta}(\mathbf{x})$ over $\dot{x}$ for $\theta = \frac{1}{4}$

$$\ddot{x} = \ddot{x}_{\text{ref}} + \frac{\hat{\theta} - \theta + \mathbf{x}^T \mathbf{Q} \mathbf{x}}{m} (d\ r)\mathbf{x} + (40\ 400)(\mathbf{x}_{\text{ref}} - \mathbf{x})$$

$$\dot{\hat{\theta}} = -2\mathbf{x}^T \mathbf{Q} \begin{pmatrix} \ddot{x}_{\text{ref}} + (40\ 400)(\mathbf{x}_{\text{ref}} - \mathbf{x}) \\ \dot{x} \end{pmatrix} .$$

Estimator Convergence

The second condition is that θ_e converges to θ, which is equivalent to the manifold $\theta_e = \theta$ being attractive. This can be studied using the estimator error

$$z = \theta_e - \theta = \hat{\theta} - \theta + \mathbf{x}^T \mathbf{Q} \mathbf{x}$$

and its derivative

$$\dot{z} = 2\mathbf{x}^T \mathbf{Q} \begin{pmatrix} z \\ 0 \end{pmatrix} (d\ r)\mathbf{x} .$$

$$= 2\mathbf{x}^T \mathbf{Q} \begin{pmatrix} d\ r \\ 0\ 0 \end{pmatrix} \mathbf{x} z$$

A choice of

$$\mathbf{Q} = -k \begin{pmatrix} d\ r \\ r\ 0 \end{pmatrix}$$

with $k > 0$ ensures that the product

$$\mathbf{x}^T \mathbf{Q} \begin{pmatrix} d\ r \\ 0\ 0 \end{pmatrix} \mathbf{x} = -(\dot{x} Q_{11} + x Q_{21})(d\dot{x} + rx)$$

$$= -k(d\dot{x} + rx)^2$$

is always non-positive. (The right half of $\mathbf{Q}$ does not contribute to the convergence, so Q_{22} can be set to zero to simplify the implementation.) This means that z is stable and converging to zero everywhere except on the line $d\dot{x} = rx$, where it is constant.

The factor k is a design parameter that determines the speed of the convergence. Bigger values of k lead to a faster convergence. For the example used here, $k = 0.1$ would be a conservative choice, and $k = 1$ leads to near instant convergence.

The resulting function $\hat{\theta}(\mathbf{x})$ is a parabola in x and $\dot{x}$. Only the dependency on $\dot{x}$, is shown in Fig. 6, assuming $x = 0$ and $\theta = \frac{1}{4}$. In this case, the parabola is symmetric, based at $\hat{\theta} = \frac{1}{4}$. For other values of x, the base point is different, but the function always intersects with $\hat{\theta} = \theta$ at $\dot{x} = 0$.

Intuitive Interpretation

Now that the adaptive control scheme is complete, it is possible to give it an intuitive interpretation of the non-linear dynamics. For this purpose, only a single state is considered, so that the vector field of state and estimate remains two-dimensional. As the state $\mathbf{x}$ changes according to the target dynamics and the reference input, the transformed estimate $\hat{\theta}$ attempts to track $\theta - \mathbf{x}^T \mathbf{Q} \mathbf{x}$. The change required for $\hat{\theta}$ depends on the change of the state $\mathbf{x}$. This is anticipated according to the model $f(\mathbf{x}, u, \theta_e)$, which again depends on the estimate θ_e itself.

If the estimate θ_e is too high, the speed of the system returning back to the origin is over-estimated. So the system is slower when returning to the origin than expect, which is when $\hat{\theta}$ is falling. This means that $\hat{\theta}$ falls for longer than anticipated, reducing $\hat{\theta}$ towards the correct value. On the other hand, movement away from the original is faster than expected, giving $\hat{\theta}$ little opportunity to rise. This means θ_e will be reduced with every movement of $\mathbf{x}$, until it converges towards θ. The opposite is also true: if θ_e is too low, the system is faster moving away from the origin than towards it, and θ_e will increase.

5. SIMULATION RESULTS

To test the ability of the adaptive controller to deal with faults, a number of simulations are presented. The reference is a step sequence that makes the system move back and forth twice every second. The reference is modified to be trackable using the changes discussed in Section 3. After one second (two periods), a lock-up fault in the second element is simulated, followed by a fault of the third element a second later.

The system response with only the geometric controller, i.e. without parameter adaptation, is shown in Fig. 8. It is easy to see that each fault changes the dynamics of the system further from the model, and this leads to a mismatch between the planned and the realised trajectory. The influence is small in terms of the amplitude (less than 5% of the travel), but the convergence of this deviation is slow. The control loop would eventually reduce the deviation significantly, but the steady state is not reached within the short periods of 150 ms during which the reference remains constant. So for a positioning application, this would be an unacceptable result. Tighter control could reduce the deviation faster, but it cannot be eliminated completely.

The system response with an adaptive controller is shown in Fig. 7. A conservative estimator parameter $k = 0.1$ is used here, leading to slow adaptation. This simulation again shows a deviation from the reference for the first movement after a new fault occurs (at $t = 1\,$s and $t = 2\,$s). After this movement, the controller is adapting, and the system no longer shows any noticeable deviation from the nominal behaviour. The graph of the estimate θ_e confirms that the behaviour of the system is reflected correctly. The nominal values of $\theta = \frac{1}{3}$ after the first fault and $\theta = \frac{1}{2}$ after the second are reached reliably.

If this deviation in the first movement is not acceptable, it is possible to make the estimate converge faster by
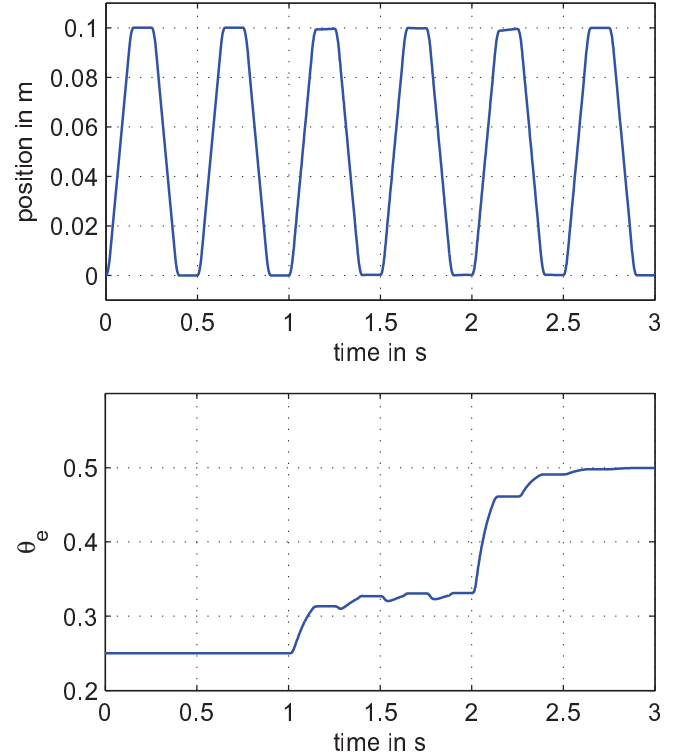

Fig. 7. Simulation With Slow Adaption

choosing $k = 1$. The results of this are shown in Fig. 9. This simulation demonstrates a very fast adaptation of the controller, which settles to the new parameter within under 100 ms. This means that at the end of the first movement after a fault occurs, the controller is already adapted to the new situation, and the trajectory settles nominally. So this controller effectively restores the original behaviour completely. The downside is that the parameter estimate is very sensitive to the measurements and to unmodelled dynamics. Because only a simplified model was used for the design of the controller, these unmodelled dynamics are present in the simulation, and they lead to the moderate dynamic errors in the estimate after the first fault. This does not seem to have any impact on the dynamics of the system, but it shows that a practical application needs to be instrumented and tested carefully. The same does not happen after the second fault, because this cancels out most of the unmodelled dynamics.
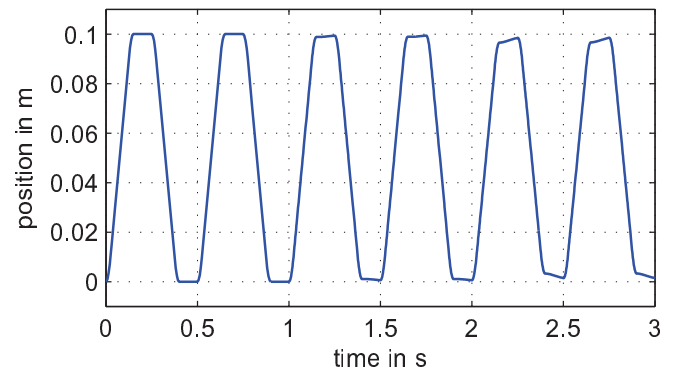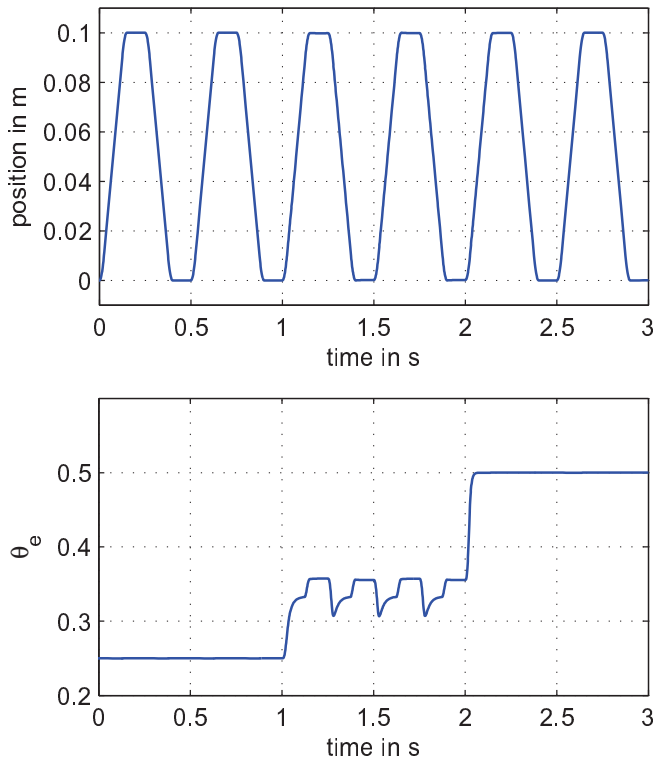

Fig. 8. Simulation Without Adaption

Fig. 9. Simulation With Fast Adaption

6. CONCLUSIONS AND FURTHER RESEARCH

The simulations show that the adaptive controller produces a very constant and precise control of an HRA in case of faults. The results are significantly superior to those achievable with a robust controller. So the adaptive control would be useful in application areas where precise movements are required, such as pick and place in manufacturing or the control of a sensitive position.

The complexity of the adaptive controller is only slightly higher than a normal PID controller: the three separate modules for trajectory generation, target dynamics and parameter estimation consisting of only a few mathematical operations each. So the adaptive controller could be easily implemented even in a very simple digital processor.

One of the big promises of this control scheme is that will scale easily to higher order systems, such as 10×10 HRA configurations. The complexity of the adaptive controller does not increase with the number of elements, because only the simplified 2 state system model is used for the design.

Further research will go into several separate directions. One goal is to add the adaptation of several parameters at the same time, to accommodate different fault modes and to respond to constant disturbances (similar to a PI controller). The other challenge is to test the results on an experiment test bed currently under construction. Finally, it is planned to integrate the estimated parameter into a comprehensive health monitoring system for the HRA.

7. ACKNOWLEDGEMENTS

REFERENCES

A. Astolfi and R. Ortega. Immersion and invariance: A new tool for stabilization and adaptive control of nonlinear systems. *IEEE Transactions on Automatic Control*, 48 (4):590–606, 2003.

A. Astolfi, D. Karagiannis, and Romeo Ortega. *Nonlinear and adaptive control with applications.* Springer, London, 2008. ISBN 1848000650; 9781848000650; 1848000669; 9781848000667.

J. Davies, T. Steffen, R. Dixon, and R. M. Goodall. Multi-agent control of a 10x10 high redundancy actuator. In *Proceedings of the 23rd IAR Workshop on Advanced Control and Diagnosis*, 2008a.

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. In *Proceedings of the IFAC World Congress 2008*, Jul 6-11 2008b.

R. Ortega, L. Hsu, and A. Astolfi. Immersion and invariance adaptive control of linear multivariable systems. *Systems and Control Letters*, 49(1):37–47, 2003.

T. Steffen, R. Dixon, R. M. Goodall, and A. Zolotas. Multi-variable control of a high redundancy actuator. In *Actuator 2008 – International Conference and Exhibition on New Actuators and Drive Systems – Conference Proceedings*, pages 473—476. HVG, 2008a. ISBN 3-933339-10-3.

T. Steffen, R. Dixon, R. M. Goodall, and A. C. Zolotas. Robust control of a high redundancy actuator. In *Proceedings of the UKACC Control Conference 2008*, Sep 2-4 2008b. Tu09.06/p149.