Analytical Redundancy Techniques for Fault Detection in an Active Heavy Vehicle Suspension

B.P. Jeppesen and D. Cebon
Cambridge University Engineering Department

Department of Engineering, University of Cambridge,
Trumpington Street, Cambridge CB2 1PZ, UK.
Phone: +44-1223-332665
Fax: +44-1223-332662
Email: dc@eng.cam.ac.uk

This paper describes the design and implementation of an intelligent fault monitoring system for a prototype heavy vehicle active suspension system. A controller architecture has been designed, to facilitate safe start up and shutdown of the active control. This paper develops a method for monitoring the complete system, by dividing it into submodels which are more straightforward to analyse. The aim is to define a practical method for detecting faults, taking into account the nonlinearities in the real vehicle.

Keywords: heavy vehicle, active roll control, active suspension, fault detection, fault diagnosis.

1. INTRODUCTION

1.1 CVDC Experimental Vehicle

Cambridge Vehicle Dynamics Consortium (CVDC) has constructed an experimental heavy vehicle to test the benefits of active roll control [1] and semi-active ride control [2]. The vehicle is based on a 3-axle tanker semi-trailer with independent trailing arm suspension, and a standard 2-axle tractor unit. The semi-trailer has an active roll control system, actuated by hydraulic cylinders attached between the anti-roll bars (ARBs) and chassis. Valve-actuated continuously variable dampers are used to achieve semi-active ride control. The tractor will be fitted with similar systems. Fig. 1 shows the experimental vehicle. Fig. 2 is a CAD model of one axle of the active suspension, demonstrating the method of actuation.

The moment generated by the anti-roll bars is measured by load cells in line with the hydraulic cylinders. Suspension roll angle is measured by displacement transducers located between each trailing arm and the chassis. Speed, steering, and lateral acceleration (on both tractor and trailer) are amongst other sensors used.

1.2 Control Architecture

A distributed network of controllers is used, communicating over a CAN network [1]. One *Global Controller* (GC) is used to coordinate four *Local Controllers* (LCs). There is one LC each for roll and ride control on tractor and trailer units.

1.3 Safe Design

Various features have been added to the basic feedback control system to enhance safety. They follow guidelines produced by the Health & Safety Executive (HSE) [3]. A basic design premise of 'fail-safety' has been used for both hardware and software design, and has been augmented by an independent monitoring and safe-shutdown system [4].

Fig. 1 CVDC experimental vehicle

1.4 Sensor Monitoring

Approximately 70 sensors monitor the complete vehicle, with a similar number of control valves in the dampers and suspension hydraulics. It is a considerable challenge to monitor the health of all these components, but it is vital to detect safety-critical faults in order to operate the safe-shutdown system. Various levels of checking are used on sensor inputs. A simple open-circuit test is supplemented by knowledge of the range of values expected from the sensors during normal operation. However, potentially dangerous faults such as changing gains, offsets or actuator faults may not be detected using these methods.

This paper demonstrates how 'analytical redundancy' techniques [5] can be used to make fault detection more sensitive, and facilitate reliable fault diagnosis. Previous work [4] using the kinematics of the vehicle in the roll plane is extended by adding relations describing the vehicle roll dynamics, controller action and internal suspension forces.

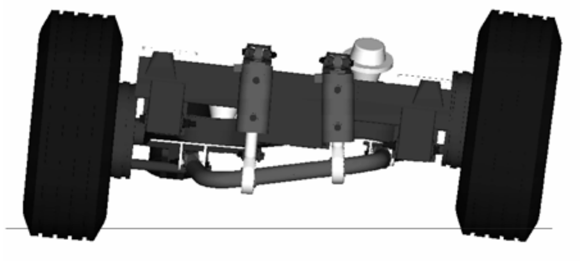

Fig. 2 CAD model of the CVDC active suspension.

2. LITERATURE REVIEW

2.1 Fault Detection and Isolation (FDI)

A number of methods have been developed for fault detection and isolation [5-8]. All methods of fault detection work by designing *residual* functions. The residual represents the difference between an estimated value and a measured one, which should be zero during normal operation, but large in the presence of faults.

In practice, there is a distinction between the detection of fast-acting, possibly safety-critical faults, and faults which are non-safety-critical and slower to develop, for example due to wear. The former are most likely to be detected by state-estimation and instantaneous comparison of prediction with measurement, while the latter are detected using parameter estimation techniques which require a certain time window and excitation of the system. Probability analysis can be used to judge, from the residual values, when a fault or change has taken place [6, 7]. This paper is concerned primarily with detection of fast-acting faults, detected via state estimation.

Isolation, in the literature, means diagnosis of the faulty component. If faults are allowed to occur simultaneously, then for a diagnosis, at least as many independent residual functions as faults considered are required. In practice, it is usually assumed that only one fault occurs at a time, which facilitates more robust fault diagnosis [8].

2.2 Techniques for linear systems

Ideally, an online model of the system would be accurate enough that open-loop prediction of measured sensor values would be all that is required to generate residual functions. However, due to uncertainties in the model or unknown inputs, the open-loop state estimates may be inaccurate.

Estimation techniques borrowed from the control literature [6], namely Luenberger observers or Kalman filters [9, 10] can be used to feed measured sensor values back into the model to correct any drift in the predictions. In order to preserve a structure for fault diagnosis logic, a bank of such 'observers' is then required, each based on a different subset of sensors. Suggested configurations are the *Dedicated Observer Scheme*, where each observer is based on one measurement only, and the more robust *Generalised Observer Scheme*, where each observer is based on all but one of the measurements [6, 8].

As an alternative to introducing feedback, robustness can be improved by extending residuals based on state estimation, from the consideration of the current sensor readings only, to those for a finite window of time [8]. This type of residual function is known as a *parity relation*. The conceptual similarities to parameter estimation have been explored mathematically [11].

3. FDI IN EXPERIMENTAL VEHICLE ACTIVE SUSPENSION

The control system in the experimental vehicle aims to reduce the effect of steering inputs on the lateral load transfer of the vehicle, and hence its tendency to roll over. Load transfer is not measured directly by the control system however.

3.1 Divide the full vehicle model into subsystems

Attempting fault diagnosis by using a 'bank' of online observers, each based on a full articulated vehicle model, would be too expensive computationally. Instead, the full vehicle model is divided into several subsystems, each of whose inputs and outputs are measured (Fig. 3). As well as reducing the order of each model to be simulated, the approach separates vehicle yaw dynamics, which are a nonlinear, function of speed, from the roll dynamics, which are assumed to be linear. There is some unmodelled coupling between the two due to inertial effects, but this is assumed to be small.

Vertical (ride) motion is not modelled, and again, its coupling with roll motion is assumed to be small. This includes the Controller action, since the valve demand d is really the *difference* between demands across the axle. Closed-loop control of the vertical position of the anti-roll bar is also used, but not modelled since it is not safety-critical.

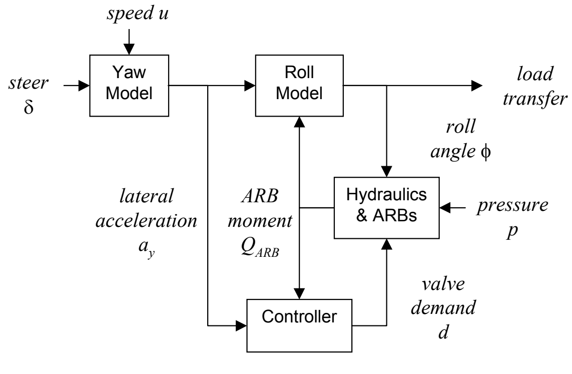

Fig. 3 Vehicle as a set of subsystems

This paper concentrates on modelling the trailer roll motion. A reliable estimate of trailer lateral acceleration is already available. Three observers are used, and are described below. The term 'observer' does not automatically imply the use of corrective feedback, and the use of this is discussed in each case.

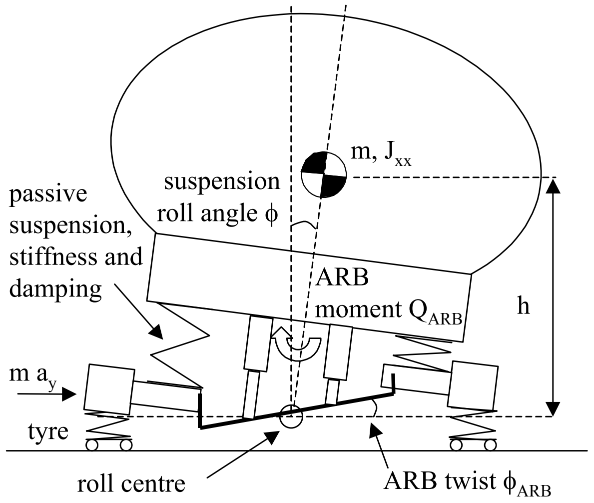

Fig. 4 Trailer roll model

3.2 Observer 1: ϕ (a_y, Q_{ARB})

Observer 1 models trailer roll motion with a single axle (Fig. 4). In practice, comparisons between adjacent axles would be used to identify which axle is the source of a fault. Motion of the unsprung mass is not considered.

The overturning moment due to measured lateral acceleration, plus the countering active control torque are the inputs to the model, with suspension roll angle being the output. The passive model includes roll stiffness and damping, with the sprung mass centre of mass (CoM) at a height *h* above the roll centre. The governing equation is:

$$(J_{xx} + m h^2)\, \phi'' + L\, \phi' + K\, \phi = Q_{ARB} \qquad (1)$$

The parameters were initially estimated from trailer mass, dimensions and suspension data. By comparing the model with measured responses during transient manoeuvres, it was found necessary to increase *K* and *L* by 10% to account for the constraint imposed by the tractor unit and other unmodelled stiffnesses.

Table 1 Roll model parameters

m	sprung mass	31 000 kg
Jxx	moment of inertia about sprung mass CoM	20 000 kg m^2
h	height of sprung mass CoM above roll centre	1.8 m
K	passive roll stiffness (excluding ARBs)	2.2 MNm/rad
L	passive roll damping (excluding hydraulics)	560 kNms/rad
K_{ARB}	passive stiffness of ARBs only	3.1 MNm/rad
L_{ARB}	passive damping of hydraulics in 'float' mode	260 kNms/rad

3.3 Observer 2: Q_{ARB} (d, ϕ)

Observer 2 models the hydraulic system only, with the following assumptions:

1. Hydraulic fluid is incompressible.
2. Flow is proportional to valve position *d*; hence, the angle of ARB to chassis (= $\phi + \phi_{ARB}$) is proportional to the integral of *d.*
3. Flow is proportional to nominal pump pressure *p* (which is known, and constant).

2 and 3 are simplifications since flow is limited by the flow rate of the hydraulic pump and leakage. However, combining the assumptions gives:

$$Q_{ARB} = K_{ARB}\, [\, K_d \int d\, dt - \phi\,] \qquad (2)$$

K_d depends on pressure, and was found by experiment to be 2.3 deg/(Vs). Since (2) involves direct integration, it may be expected that this observer will require feedback to correct drift.

3.4 Observer 3: Q_{ARB} (a_y)

Observer 3 models the *Roll Model*, *Hydraulics, ARBs*, and *Controller*. It is not attempted to simulate the controller calculations separately due to the computational effort required. Monitoring the desired complete system performance gives an accurate indication of safety and allows a simpler model to be used; malfunction of the controller can still be detected logically in combination with Observers 1 and 2.

The ARB controller aims to generate a torque Q_{ARB} proportional to a_y. Assuming that the main limiting factor is the speed of response of the proportional valves used (20ms, from manufacturer's data), a simple first order response is assumed as an approximate model for FDI purposes,

$$T\, Q_{ARB}' + Q_{ARB} = K_{ay}\, a_y \qquad (3)$$

with T = 20ms. K_{ay} is defined by the control system software to be 44 kNm/(m/s^2) in the tests conducted here. Since Q_{ARB} is the subject of closed-loop control, and the model (3) is stable, there should be no need for corrective feedback in the observer.

3.5 Fault diagnosis by combining observers

Three separate *residual* functions, r_1 to r_3, are generated by taking the absolute difference between the predicted and measured values for each observer. Each measures the compatibility of a different subset of measurements. Table 2 shows how 4 different example faults can be diagnosed by the combination of residuals which detect them.

Table 2 Detection of example faults

Fault	r_1 (Q_{ARB}, ϕ, a_y,)	r_2 (Q_{ARB}, ϕ , d)	r_3 (Q_{ARB}, a_y)
Roll sensor	Yes	Yes	No
Moment sensor	Yes	Yes	Yes
Hydraulics	No	Yes	Yes
Controller	No	No	Yes

4. VALIDATING-MODELS FOR SUBSYSTEMS

4.1 Test performed

In order to test the controller performance and that of each observer systematically, a simulated a_y input was fed to the controller while the experimental vehicle was stationary. The demand was changed, in steps of 1m/s^2, from zero up to 5 m/s^2 and back again.

Fig. 5 shows measured and estimated ϕ over the whole test, while Fig. 6 shows the response of Q_{ARB} to the controller over a smaller time window, for clarity.

In order to reduce controller effort, a 'deadband' was introduced into the closed-loop control, which is the reason for the pattern in measured Q_{ARB}, and the 'pulsed' valve demand signals.

4.2 Observer 1: ϕ (a_y, Q_{ARB})

Compared to the prediction of the observer with no feedback, measured ϕ appears to have smaller initial response and then a period of 'creep' (Fig. 5).

Since the vehicle parameters have already been validated, the limited initial response in ϕ is likely to be due partly to the stop/start nature of the test itself. Part of the measured moment may only be acting against friction in the suspension which limits the roll angle achieved.

The secondary movement in the measured response of ϕ is described as 'creep' since Fig. 6 shows that Q_{ARB} is simultaneously kept constant. The rate of creep also increases with applied Q_{ARB}, so that the effect is analogous to a slow roll damper in series with the passive suspension stiffness. This can be explained by the vehicle air suspension, which is connected across the vehicle in order to be operated by a single height sensor. The system is designed to allow slow equalisation of airbag pressure, with a time constant of 20s.

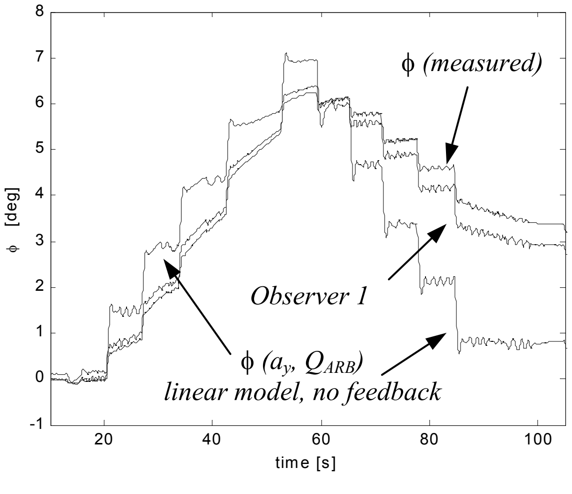

Fig. 5 Roll angle response to simulated a_y (Observer 1)

Although these two effects are not desirable, they are not dangerous and should not cause false alarms from the error in predicted ϕ . Some feedback is therefore used in Observer 1, keeping the prediction much closer to the measured value (Fig. 5).

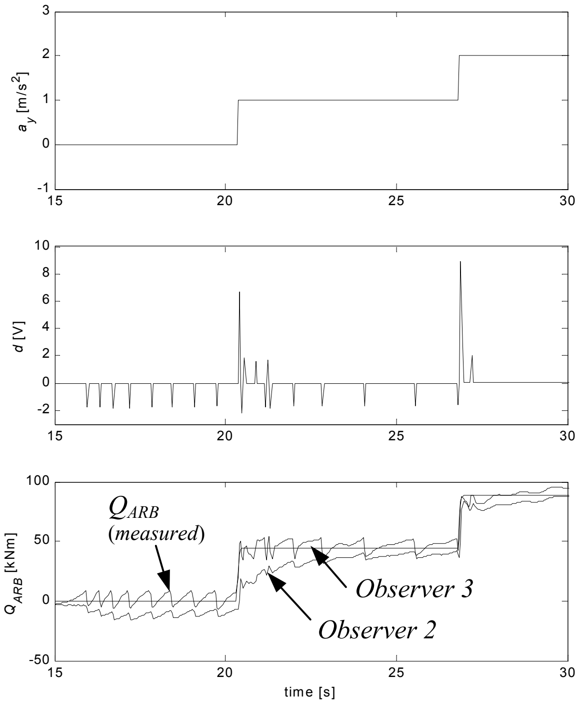

Fig. 6 Simulated a_y input (top); valve demand d (middle); moment response Q_{ARB}, Observers 2 and 3 (bottom)

4.3 Observer 2: Q_{ARB} (d, ϕ)

Due to the integration of valve position used in Observer 2, an observer using feedback is required to estimate Q_{ARB} from d and the roll angle response ϕ. Fig. 6 (bottom) shows that the resulting observer works well, though it accumulates errors due to the controller corrections at low Q_{ARB}. These corrections occur because the proportional valves are difficult to centre precisely and some small flow takes place in one direction at demanded 'zero'.

4.4 Observer 3: Q_{ARB} (a_y)

Fig. 6 (bottom) gives a comparison between this first-order response observer and the controlled Q_{ARB}. It shows that Observer 3 is a good first order fit to the controller response (measured Q_{ARB}), with no need for additional corrective feedback.

5. FAULT DETECTION USING RESIDUALS BASED ON SUBSYSTEMS

5.1 Detection of sensor failure

Fig. 7 shows measured a_y and ϕ during a manoeuvre, together with the three residual functions. A fault has been introduced into the ϕ measurement, of an intermittent short circuit to zero. This type of fault would be difficult to detect by simple checks on the range of ϕ, since zero roll angle is at the centre of the working range.

As predicted by Table 2, r_1 and r_2 detect the fault, but r_3 does not. The initial sharp response and subsequent decay in the innovations is due to the use of feedback in Observers 1 and 2. The innovations will reflect the form and duration of the added fault signal more accurately for small gain values, but this must be balanced against the need for feedback to correct drift in estimation.

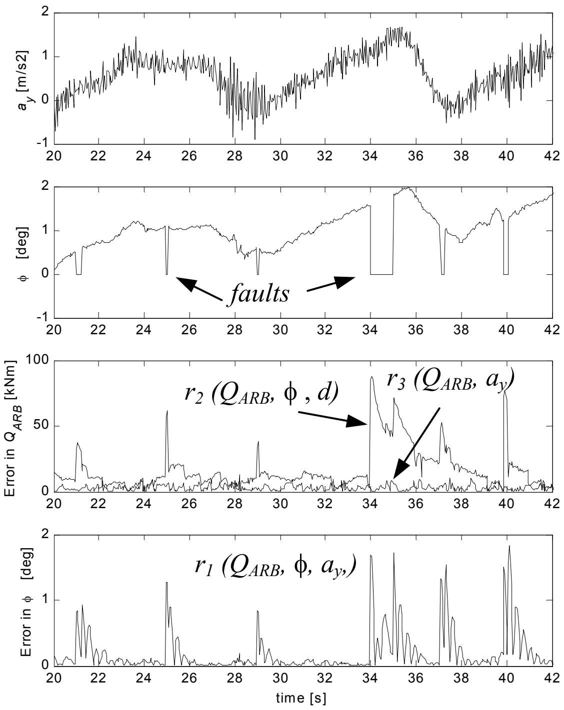

Fig. 7 Effect on residuals of fault in ϕ measurement

5.2 Detection of loss of ARB stiffness

If valves in the hydraulics fail, the ARBs become disconnected hydraulically from the vehicle, and it is no longer possible to control Q_{ARB}. This type of failure was simulated by switching the ARBs manually into a 'floating' mode and then 'locking' them again while driving the vehicle and noting the results in the observers. No controller action was attempted during the test to avoid damaging equipment. r_3 is therefore not reproduced, as it simply shows an error whenever a_y is nonzero. The test is also relevant to passive vehicles which experience a loss in suspension stiffness.

In order to aid diagnosis of this particular fault, an additional residual was designed to model effect of the fault, which causes the ARB to act as an additional roll damper as it 'floats' in the hydraulic circuit:

$$r_4 = | Q_{ARB} + L_{ARB} \phi' | \qquad (4)$$

r_4 will be small when the ARBs are 'floating', but large when the ARB is 'locked' during normal operation. The value of L_{ARB}, the equivalent roll damper, was estimated by noting the response of Q_{ARB} to roll rate during transient manoeuvres while the ARBs were in 'float mode'.

r_4 and r_2 now represent alternative models of the behaviour of the ARB, only one of which should be true at a time as during this test. This is an example of *Multi-Hypothesis Testing* [6]. The technique generally uses a set of residuals, each one based on the system model containing a different fault, and one containing no fault. Probability analysis is used to choose the most likely candidate model from the set. This is most useful for diagnosing faults which change the system dynamics, but in a predictable way assuming the fault is known.

Fig. 8 shows the results. When the vehicle is driving along the straight, none of the residuals register. This is because detecting loss of stiffness is essentially parameter estimation, and requires sufficient excitation of the suspension. The increased roll motion while the ARBs are 'floating' is detected clearly by r_2, which represents the error in the model assuming the ARBs are 'locked'. Conversely, r_4 is small because the ARBs are now 'floating', and this model is valid during this period. r_1 is small throughout the motion, indicating that the passive part of the vehicle's suspension is fault-free.

6. CONCLUSIONS AND FURTHER WORK

6.1 Conclusions

1. A method has been demonstrated for detecting faults in complex control systems, by dividing the complete system into simpler subsystems.
2. Weak coupling between subsystems may be allowed if the resulting errors are small.
3. The use of feedback in observer-based residuals is a compromise between the need to correct model inaccuracies or unknown inputs, and the desire to identify the form and duration of the fault.
4. The different residuals should be designed to facilitate simple fault diagnosis logic, each using as few measured signals as possible.
5. The number of possible diagnoses is limited by the number of independent residuals designed, which in turn depends on the number of available measurements.

6.2 Further work

The fault detection problem here has been considered up to the point of residual generation, with the aim being to design a set of residuals which will have sufficient accuracy and information content to diagnose a fault.

How to generate accurate diagnoses from the residuals requires further investigation. In particular, how to use probability theory or pattern recognition techniques to improve decision making and avoid false diagnoses when different residuals have different response times.

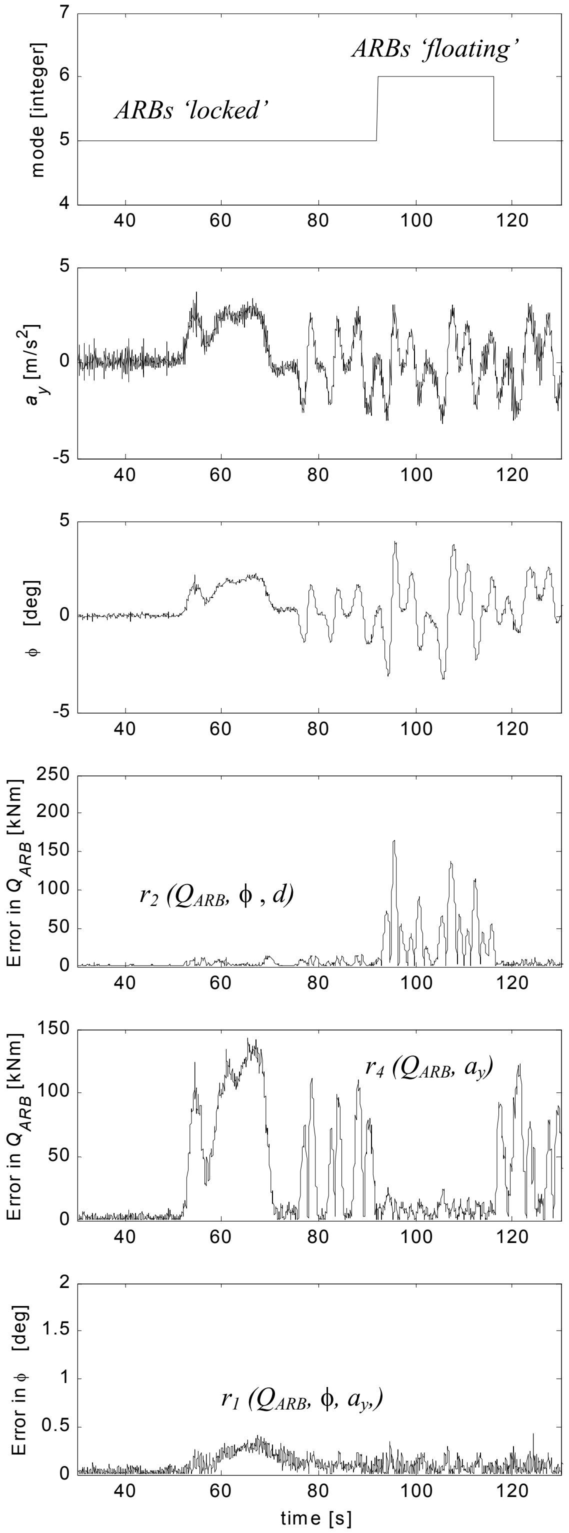

Fig. 8 Innovations from observers when ARBs are 'switched' off and on again

ACKNOWLEDGEMENTS

The authors wish to acknowledge the financial support of the Cambridge Vehicle Dynamics Consortium (CVDC) and the UK Engineering and Physical Sciences Research Council. CVDC currently comprises the Universities of Cambridge and Cranfield together with the following industrial partners : ArvinMeritor, Firestone Industrial Products, Fluid Power Design Ltd, General Trailers, Koni BV, Mektronika Systems Ltd, MIRA, QinetiQ Ltd, Shell UK Ltd, Tinsley Bridge Ltd and Volvo Trucks. The authors also wish to thank Richard Roebuck, Brian Jujnovich, Edwin Stone, Arnaud Miège, Samuel Leslie and the technical staff of Cambridge University Engineering Department for their help in building the experimental vehicle and performing the experiments described.

REFERENCES

[1] Sampson, D.J.M., B.P. Jeppesen, and D. Cebon "The Development of an Active Roll Control System for Heavy Vehicles", Proc. of 6th International Symposium on Heavy Vehicle Weights and Dimensions, 2000, Saskatoon, Canada.

[2] Roebuck, R.L., *et al.* "Developments in Semi-Active Heavy Vehicle Suspensions", Proc. of Proc. 6th International Symposium on Heavy Vehicle Weights and Dimensions, 2000, Saskatoon, Canada.

[3] HSE "Programmable Electronic Systems in Safety Related Applications, Parts 1 and 2", 1987, Sheffield: Health and Safety Exective, HMSO.

[4] Jeppesen, B.P. and D. Cebon "Real-Time Fault Identification in an Active Roll Control System", Proc. of 17th IAVSD Symposium, 2001, Copenhagen.

[5] Chow, E.Y. and A.S. Willsky "Analytical Redundancy and the Design of Robust Failure Detection Systems", IEEE Transactions on Automatic Control, 1984, AC-29(7): pp. 603-514.

[6] Patton, R., P. Frank, and R. Clark "Fault Diagnosis in Dynamic Systems, Theory and Application", 1989, Hemel Hempstead, UK: Prentice Hall International (UK) Ltd.

[7] Willsky, A.S. "A survey of design methods for failure detection in dynamic systems", Automatica, 1976, 12: pp. 601-611.

[8] Frank, P.M. "Fault Diagnosis in Dynamic Systems Using Analytical and Knowledge-based Redundancy - A Survey and Some New Results", Automatica, 1990, 26(3): pp. 459-474.

[9] Luenberger, D.G. "An introduction to observers", IEEE Transactions on Automatic Control, 1971, AC-16(6): pp. 596 - 602.

[10] Franklin, G.F., J.D. Powell, and M. Workman "Digital Control of Dynamic Systems", 3rd ed, 1998: Addison-Wesley.

[11] Gertler, J. and G. DiPierro "On the relationship between parity relations and paramater estimation", Proc. of Fault Detection, Supervision and Safety for Technical Processes (SAFEPROCESS '97), 1997, Kingston Upon Hull: IFAC.