ASSESSING HEAVY VEHICLE SUSPENSION FOR ROAD WEAR

Transfund New Zealand Research Report No. 95

ASSESSING HEAVY VEHICLE SUSPENSIONS FOR ROAD WEAR

JOHN DE PONT
Transport Engineering Research New Zealand Limited

Transfund New Zealand Research Report 95

ISBN 0-478-11053-7
ISSN 1174-0574

PO Box 2331, Lambton Quay, Wellington, New Zealand
Telephone (04) 473-0220; Facsimile (04) 499-0733

de Pont, J.J. 1997. Assessing heavy vehicle suspensions for road wear.
Transfund New Zealand Research Report 95. 102pp.

Keywords: suspension assessment, road-friendliness, pavement wear, dynamic wheel loads

AN IMPORTANT NOTE FOR THE READER

The research detailed in this report was commissioned by Transit New Zealand when it had responsibility for funding roading in New Zealand. This funding is now the responsibility of Transfund New Zealand.

While this report is believed to be correct at the time of publication, Transit New Zealand, Transfund New Zealand, and their employees and agents involved in preparation and publication cannot accept any contractual, tortious or other liability for its content or for any consequences arising from its use and make no warranties or representations of any kind whatsoever in relation to any of its contents.

The report is only made available on the basis that all users of it, whether direct or indirect, must take appropriate legal or other expert advice in relation to their own circumstances and must rely solely on their own judgement and such legal or other expert advice.

The material contained in this report is the out put of research and should not be construed in any way as policy adopted by Transit New Zealand or Transfund New Zealand but may form the basis of future policy.

CONTENTS

ACKNOWLEDGEMENTS

The research work described in this report involved the efforts of a number of people. The author would specially like to acknowledge the contributions of Peter Baas, Mihai Costache, Kailash Thakur, Gary Bastin, Tom Field, and Ron Ter Huurne. A parallel research programme funded through the New Zealand Public Good Science Fund provided useful input to this work and also benefited from it.

EXECUTIVE SUMMARY

The AASHO road test of the late 1950s led to the development of the fourth power law, which states that the amount of pavement wear generated by the passage of an axle is proportional to the fourth power of the static axle load. This is a heuristic rule based on curve fitting to the observed data. Consequently it implicitly includes dynamic loading, environmental factors and anything else that contributed to pavement wear during the tests. The assumption was that all these factors were identical for the different pavements and loads involved.

Experimental research programmes in the 1970s and 1980s showed significant variations in the levels of dynamic loading generated by different suspension systems. Although there is some debate over the validity of the fourth power law and in particular, whether four is the correct value for all forms of pavement design and distress, it is generally accepted that some form of power relationship is appropriate for relating pavement wear to applied axle loads. With any power greater than unity, higher dynamic loading would be expected to lead to increased pavement wear. For this reason it is in the interests of pavement managers to encourage the use of suspensions that generate lower levels of dynamic loading. These are often referred to as road-friendly suspensions.

In order to be able to offer incentives for using road-friendly suspensions, it is necessary to have a reliable assessment procedure for identifying them. The research described in this report was aimed at developing assessment procedures using a small-scale two-post servohydraulic shaker facility. This type of facility, particularly the larger ones with four or six actuators are sometimes called road simulators.

Dynamic wheel forces have previously been measured in a number of research studies. However, the instrumentation used was either costly or time-consuming to install or both. None of the methods used are suitable for a routine assessment procedure. The concept behind the method used in the first part of this research was that relatively simple instrumentation would be fitted to the vehicle and the suspension deflection and acceleration responses would be measured during road trials. The vehicle would then be mounted on the shaker facility and using a sophisticated software control system the shaker excitations required to obtain the same response at the two wheels being tested would be determined. The corresponding wheel forces could then be measured from instrumentation on the shakers. The assumption being that if the suspension deflections were the same the wheel forces would be too.

To test this hypothesis a test vehicle fitted with steel leaf spring suspension was comprehensively instrumented with both the simple instrumentation required for the proposed assessment procedure and the more complex instrumentation required to measure wheel forces during road trials. A series of road tests were undertaken with this vehicle at different speeds on roads of varying roughness. The vehicle was then returned to the laboratory for shaker trials. Considerable work was required to develop the shaker control software to the point where the required excitations could be determined reliably and rapidly. Eventually the system was developed to the point where measured road response could be replicated in the laboratory with acceptable fidelity in a reasonable time (less than an hour).

To ensure that these results were not specific to the suspension being tested, the vehicle was modified and fitted with air suspension and the testing repeated. A complete further set of tests was undertaken with a different load distribution to determine what effect this had on the dynamic behaviour of the vehicle. These two series of tests generated several important results. First although in most other research studies, air suspensions generated lower dynamic loads than steel suspensions, this was not the case for this particular air suspension. Its performance was similar to that of the steel suspension it replaced. Modal testing showed that although the suspension was significantly softer than the steel suspension it replaced it had relatively poor damping. This was a function of the design. The suspension was new and in good working order. The shock absorbers were removed and tested and found to be up to specification. This has important ramifications. Some regulators in offering incentives to road-friendly suspensions have given a blanket approval to air suspensions. This example shows that this is not justified.

Shaker testing this suspension configuration also identified a fundamental problem with the assumptions underlying the assessment method. Although the shaker control software was able to replicate the on-road suspension deflection behaviour very well, the resulting vehicle motions and wheel forces were very much higher than they had been on the road. After much investigation it was found that this was because the vehicle had considerable auxiliary roll stiffness which was not provided by the springs. This roll stiffness generated substantial reaction forces without the corresponding suspension deflections. This problem was overcomes by modifying the shaker rig so that each actuator excited a whole axle rather than just a wheel. This eliminated all roll behaviour from the test vehicle during the laboratory trials and gave acceptably good results.

In parallel with this project a programme of accelerated pavement testing to investigate the influence of dynamic loading on pavement performance is being undertaken at the Canterbury Accelerated Pavement Testing Indoor Facility (CAPTIF). To enhance the linkages between the two projects a series of shaker tests was carried out on one of the Simulated Loading and Vehicle Emulator (SLAVE) unit from CAPTIF. This is the vehicle that applies the loading at CAPTIF and essentially behaves like a quarter truck. Although this vehicle is a considerably simpler dynamic system than a complete truck, the match between the measured responses at the CAPTIF facility and those on the shaker were not as good as had been anticipated. Although the match in suspension response was good, there was an additional structural vibration mode present during the shaker testing, which had not been observed during CAPTIF operations. The frequency of this mode was similar to that of the axle hop response of the vehicle and so complicated the measured accelerations. Although not a general finding, this does provide a caution against using simple acceleration based instrumentation for wheel force estimation particularly with flexible vehicle structures.

Simple linear pitch-plane models were developed and used to investigate options for enhancing the shaker-based assessment procedures. In the first instance, the model parameters were tuned to match the test vehicle behaviour as closely as possible. This tuned model was then used to calculate how road profile data should be modified to generate the shaker excitations required to elicit the same suspension response. This eliminates the need for a road test prior to shaker testing and for any vehicle instrumentation. The results of this work were very encouraging. The second application of the linear models was to determine how a vehicle should be configured for testing so that the suspension rating obtained was as independent as possible of the vehicle to which it was fitted. Again the results were very encouraging leading to a recommended strategy of concentrating the load as much as possible over the suspension under test and minimising its pitch inertia.

During the course of this work a number of other research groups around the world were also investigating suspension assessment methods. Some of the most important findings are reviewed. These together with the findings of this project are combined to develop the following recommendations for suspension assessment procedures.

- Although there is some argument for assessing complete vehicle-suspension combinations, this is unlikely to be practical given the complexity of the testing and the large number of vehicles. It is therefore recommended that the assessment procedure should rate suspensions in a vehicle-independent way if possible.
- Design prescriptive ratings such as "all air suspensions qualify as road-friendly" should not be used. If an assessment regime is instituted, then to qualify for the road-friendliness rating a suspension should be tested.
- There are three basic methods of suspension assessment test that appear to give usable results.
- The simplest is based on the "drop test" specified in the EC directive on road-friendly suspensions. (The other two methods specified by the EC are not suitable for testing axle groups). The EC test, which specifies a single drop height, does not adequately take into account suspension non-linearities. If this approach is used it is recommended that the test be undertaken at three drop heights. This test provides information on two suspension parameters, natural frequency and damping. Suitable limit values have been proposed.
- (Woodrooffe 1997) proposed a more complicated approach in his report for the OECD DIVINE project. He suggests a low amplitude (1 mm) sinusoidal sweep excitation. This identifies both the sprung mass and unsprung mass resonance modes and values of frequency and damping can be obtained. Thus this is a more complete characterisation of the suspension system, although only for small excitation amplitudes.
- A more complete characterisation can be achieved by using the shaker facility to measure the wheel forces for virtual road profiles at different roughness levels and at various speeds. Using regression analysis these measured wheel forces could be reduced to characteristic values at some nominal roughness and speed. The advantage of this approach is that the testing directly reflects the operational situation.
- For all three testing approaches the vehicle should be configured to minimise the influence of the rest of the vehicle on the suspension rating. A strategy for doing this has been developed.
- Although all three methods could be undertaken on the shaker rig, simpler test rigs could be designed and built for the first two procedures.
- A formula for combining individual suspension ratings into a whole vehicle rating is proposed. This would be easily applied to the results of the third test method but more difficult for the first two where the outcome is likely to be of the form pass/fail.
- In addition to the initial rating it is important to ensure that suspensions maintain their performance in service. In this respect damper performance is critical.
- Before implementing an assessment regime a trial set of assessments should be undertaken to establish reasonable limits for the rating results.

In summary, three suspension assessment techniques have been found that are practical and consistent with current international opinion. The servohydraulic shaker facility can be used for all three methods and in this case there is relatively little extra cost in conducting all three tests at once. Cheaper purpose-designed facilities could be used for the two simpler procedures. Although a test is likely to cost several thousand dollars, only one of each suspension design would need to be tested and so the total cost would not be excessive. Assuming that perhaps fifty suspensions would be tested the total cost would be of the order of \$100,000 - \$150,000 which would be spread over all vehicles using those suspensions. Compared to the likely associated savings in pavement wear of at least \$16M p.a. (see section 1.3) this is a very small amount. If incentives for using "road-friendly" suspensions are introduced, it is important to ensure that suspensions maintain their performance in-service. Damper performance is critical to this issue and a regime of damper testing is proposed for this purpose. As these tests would be required for every vehicle with a road friendliness rating on a regular basis the costs of this regime would be greater than the costs of the initial assessments. These costs need not all be attributed to the "road-friendliness" rating procedure as there are safety benefits associated with ensuring that dampers are maintained. Nevertheless the design of any incentive scheme would need to take these costs into account and canvass operator attitudes to determine its viability.

ABSTRACT

This report describes a substantial research project to determine procedures for assessing the road-friendliness of heavy vehicle suspension systems. The initial aim was to develop a procedure for replicating the on-road dynamic wheel forces generated by a heavy vehicle suspension in the laboratory on a simple two-post servohydraulic shaker facility. The first series of tests were undertaken on a test vehicle fitted with steel suspension. Dynamic wheel forces were measured during a set of road trials. These were followed by a set of shaker trials, which aimed to match the on-road suspension dynamics. The vehicle was modified and fitted with air suspension and the testing repeated. With appropriate modifications to the shaker control software and the vehicle support rig a reliable and reasonably fast method for replicating the suspension dynamics in the laboratory was developed. Additional shaker trials were undertaken on the loading vehicle from the Canterbury Accelerated Pavement Testing Indoor Facility (CAPTIF).

Simple linear computer simulation models were developed and used to refine the testing procedure. In the first instance a tuned model was used to calculate the required shaker excitations directly from the road profiles. This eliminates the need for a road test to provide the input data for the shakers. The second use of the linear models was to develop testing strategies that minimise the effect of other vehicle parameters on the suspension's response. In this way a vehicle independent suspension rating could be obtained. The findings of this study were then combined with those of other researchers to present a set of recommendations for suspension assessment procedures.

1 INTRODUCTION

1.1 Overview

This document is the final report of a long and extensive program of testing and modelling aimed at developing techniques for assessing suspension performance with respect to "road-friendliness". Much of the work has been reported previously in conference papers and publications. However, for completeness those results will be included in this report.

The research was conducted in conjunction with a New Zealand Public Good Science Fund project that was concerned with the theoretical aspects of developing the suspension assessment procedure. This included the development of the servo-hydraulic shaker control algorithms and computer modelling of vehicle responses. The two projects provided considerable cross-linkages to each other with considerable benefits to both. The report incorporates findings from both projects.

During the time that this work was proceeding a number of other researchers around the world were also addressing this issue. The author has close linkages with these people and a number of the developments reported here came about through these interactions.

The remainder of this introductory chapter will focus on presenting the background information that led to this research program, followed by an outline of the research undertaken. The work described in this report focuses on developing a performance test rating procedure using a small-scale relatively low cost, general-purpose servohydraulic facility.

1.2 Vertical Dynamic Wheel Forces

Analysis of the results of the AASHO road test (Anon 1962) of the 1950s generated a widely used relationship between axle loads and pavement wear which has become known as the fourth power law. This states that the decrease in "pavement serviceability" generated by a heavy vehicle axle passing over a pavement is proportional to the fourth power of the static axle load. Although the fourth power law was developed empirically, mathematical models of fatigue behaviour and laboratory tests on pavement specimens also produce a power relationship between applied load and reduction in service life. Thus there is a physical justification for this type of relationship.

More recently, numbers of studies (Hahn 1985; Mitchell and Gyenes 1989; Sweatman 1983; Woodrooffe et al. 1986) have measured the dynamic wheel forces generated by heavy vehicles and compared the performance of various suspensions. These studies have all shown significant differences in the level of dynamic wheel force generated between different suspensions under the same test conditions. The most widely used measure of dynamic wheel force is called the dynamic load coefficient (DLC) which is defined as

$$DLC = \frac{standard\ deviation\ of\ wheel\ force}{static\ wheel\ load}$$

Some authors use the mean axle load rather than the static axle load as the denominator for normalising this expression. With perfect load sharing the two are identical. In practice, there are differences but the effect of these is small. Although the results of the different studies are not directly comparable there are clear common trends. In general, DLC increases with road unevenness and with vehicle speed and the difference in DLC between suspensions becomes more pronounced. At low speeds on smooth roads relatively little difference between suspensions is observed. In almost all cases air suspensions produced lower DLC values than steel suspensions under identical conditions. A notable exception was a torsion bar suspension (Sweatman 1983) which performed as well as the best air suspensions. As DLC is dependent not only on the suspension but also on speed and road unevenness it is not possible to put absolute values on the performance of a particular suspension. (Sweatman 1983), proposes a reference test condition which is equivalent to travelling at highway speed over a road of moderately high unevenness. At this test condition he recorded DLC values ranging from 0.13 to 0.27. This range is typical of that observed by the other researchers.

1.3 Pavement Wear Implications

To convert dynamic wheel force to an estimate of induced pavement wear, a number of assumptions are made. The most widely used approach is that developed by (Eisenmann 1975), which assumes that the wheel force distribution is Gaussian and random and that the fourth power law applies. On this basis a road stress factor, Φ, can be defined as

$$= KP_{stat}^{4}\ [1 + 6\bar{s}^{2} + 3\bar{s}^{4}]$$

$$where\ \ P_{stat} = mean\ axle\ load$$

$$\bar{s} = coeff\ of\ variation\ of\ dynamic\ wheel\ load \equiv DLC$$

$$K = constant$$

A dynamic road stress factor that reflects the contribution of vehicle dynamics to road stress can then be defined as

$$\nu = \frac{\Phi}{KP_{stat}^{4}}$$

$$ie\ \ \nu = 1 + 6\bar{s}^{2} + 3\bar{s}^{4}$$

Applying this formula to the range of DLC values given previously results in dynamic road stress factors in the range 1.10 to 1.45. This suggests that, under those operating conditions, the worst suspensions would generate over 30% more road wear than the best. The assumptions underlying the Eisenmann formula are, however, not correct. Although the wheel force distribution is approximately Gaussian and can be regarded as random in time, the measurement experiments have shown good repeatability indicating that, for a given vehicle, the peak wheel forces occur at the same spatial position along the pavement every time. There is some evidence (Cebon 1987) that even for different heavy vehicles, the peak loads tend to occur at the same approximate positions along the pavement. More recently the OECD DIVINE programme conducted an experiment using an array of Weigh-in-Motion (WIM) sensors on a section of the French National highway system which recorded a pattern of spatial repeatability (O'Connor et al. 1996). (Sweatman 1983) argued that it is the peak loads that are important in causing pavement failure and so proposed an alternative road stress factor using the 95th percentile loads. Applying the fourth power law this factor is given by

$$\Phi_{95th} = (1 + 1.645\,DLC)^4$$

Using this equation with the DLC range presented previously leads to dynamic road stress factors ranging from 2.17 to 4.35. That is, based on this theory the worst suspensions generate 100% more pavement wear than the best.

For all pavement managers, maintenance of the infrastructure is a major cost. A significant proportion of this maintenance requirement is attributable to load-related wear produced primarily by heavy vehicles. The road damage calculations above are based on dynamic wheel force levels for a "standardised" combination of road roughness and vehicle speed as presented by (Sweatman 1983). This corresponds to 80 km/h on a pavement of moderate roughness (113 NAASRA counts/km) or 100 km/h on a smoother road (72 NAASRA counts/km) or alternatively lower speeds on rougher roads. Under less aggressive operating conditions, that is smoother roads or lower speeds or both, the difference between suspensions reduces. However, even under these conditions and using the more conservative estimate of the pavement wear effect given by the Eisenmann dynamic road stress factor, the potential cost reductions from increased use of "road friendly" suspensions are substantial.

The New Zealand Land Transport Pricing Study (Ministry of Transport 1997) attributes road maintenance costs of $321M p.a. to heavy vehicles. Even taking a very conservative estimate of a 10% reduction in pavement wear due to the adoption of road-friendly suspensions and assuming only half the fleet (in terms of weight and distance travelled) is affected still gives annual savings of over $16M.

1.4 Assessing "Road-Friendliness"

In order to encourage the use of "road-friendly" vehicles it is necessary to be able to identify them and thus a method of assessment is required. One of the fundamental issues that needs to be resolved is whether it is satisfactory to rate suspensions independently of the vehicles to which they are fitted or whether it is necessary to rate complete vehicle-suspension configurations. The first approach has practical advantages, while the latter is clearly more rigorous, although in the extreme it also involves testing all possible loading configurations. It is the view of this author, that assessing suspensions independently is the only viable practical approach. Methods for combining individual suspension ratings into a whole vehicle rating are discussed later in this report (section 7.3). However, the suspension assessment techniques developed in this study could also be used for whole vehicle rating.

A number of approaches to rating suspension performance have been used and in some countries are the bases of regulations. These can be categorised into three basic forms; type based, parameter based, and performance based.

The simplest way of rating the "road-friendliness" of a suspension or vehicle-suspension configuration is based on approving generic types. This has been used for a number of years in Europe where, for certain applications, axles fitted with air suspensions have higher maximum allowable axle loads than those fitted with other suspensions. For example, in the UK, a three axle bogie fitted with air suspension has a maximum permitted bogie load of 24 tonnes while for other suspensions the maximum is 22.5 tonnes. The rationale behind this approach is that most of the research studies have shown air suspensions to be relatively road friendly. While this type of regulation has encouraged the use of particular forms of road friendly suspension, it offers nothing to other well performing suspensions such as the torsion bar type measured by (Sweatman 1983) or to encourage innovation by vehicle designers. Several studies, including this one (section 4.4), have shown that a poorly damped air suspension does not perform well in terms of road-friendliness. Under the type approval approach it would still be classed as road-friendly.

To overcome these deficiencies, parametric rating tests have been developed. A correlation between the level of dynamic wheel loads and the natural frequency and damping of the vehicle body bounce mode has been observed. The European Communities (Council of the European Communities 1992) have used this to develop a rating to define suspensions that are "equivalent-to-air" for road-friendliness. The criteria used are that the natural frequency should be below 2 Hz and that the damping should be greater than 0.2 of critical damping with more than 50% of the damping being provided by hydraulic dampers. Three tests are specified to measure these parameters. They involve measuring the response after either, driving slowly off a standard step, or dropping the vehicle a specified height, or compressing the suspension and releasing it. At present, these criteria and tests are limited to single drive axles, primarily because they do not give consistent results when applied to axle groups. The advantage of this approach is that the test can be straightforward, well defined, repeatable and transferable. The main disadvantage is that there is no direct link between the parameters measured by this test and road friendliness. It should be noted that this rating is a combination of type approval and parameter based because air suspensions are not required to be tested.

The third type of rating procedure is to use a performance test based on conditions which reflect the actual operating environment. This could be done by testing the suspension as fitted to a vehicle over some designated test pavements and measuring its response. Its performance could then be compared with that of reference suspensions. This approach has some inherent difficulties. It is dependent on the test pavements which will change in profile over time and which will vary with test location. Also the wheel force measurement techniques used in the research studies are time consuming and expensive and not suited for routine assessment use.

Attempts have been made to overcome some of these problems. Load sensing mats (Cole and Cebon 1989) have been used to measure wheel forces directly without vehicle instrumentation. These are laid over normal pavement sections and are sufficiently flexible to transmit the road profile through to the vehicle. The intention behind these mats was that they could be easily moved from test site to test site and thus although they are expensive (when a realistic length of pavement is covered) they could be highly utilised. However, experiments with their use have shown that in order to obtain accurate results the mats need to be bonded to the pavement rather than just pinned as initially envisaged. This clearly reduces their transportability.

The use of an artificial road of specified profile has been proposed (Gyenes et al. 1992) to standardise the testing. This could have an exaggerated roughness and be traversed at slow speed. Some adjustment would need to be made for the tyre-enveloping effect, which cannot be simply scaled by speed.

Another option is to simulate the vehicle behaviour either physically using a road simulator or mathematically using a computer simulation model. Road simulators consist of four or six vertical servohydraulic actuators topped with platforms that enable the vehicle to be mounted on its tyres. Excitations corresponding to the road profiles can then be applied. There are some differences between the dynamic behaviour of a rolling and a stationary tyre and some allowances for this may need to be made. A number of vehicle manufacturers have these facilities primarily for vehicle and component endurance testing. Facilities were developed at the National Research Council of Canada and at the Federal Highways Administration in Virginia (USA) during the course of this project and were used for investigating vehicle dynamics and pavement interaction. Results from these efforts will be discussed in chapter 8.

Numerous computer simulation programs that model heavy vehicle dynamic behaviour have been developed. These can be used to predict the vehicle's response to arbitrary road profiles. Although some model validation has been undertaken comparing measured and predicted behaviour, further work is required before the results of applying simulation modelling to an arbitrary vehicle configuration are generally accepted.

1.5 The Research Programme

The ultimate purpose of the research was to develop a procedure for assessing the road-friendliness of suspensions using a relatively small-scale and inexpensive two-post servohydraulic shaker facility. The first step was to determine whether the shaker facility could be used in conjunction with some relatively sophisticated control software as a road simulator. That is, whether it could get a vehicle to replicate its on-road dynamic wheel forces at the suspension under test.

The first step in the experiment was to fully instrument a test vehicle, which had steel suspensions and conduct a series of road tests to measure its vertical dynamic response. This is described in detail in chapter 2.

The vehicle was then returned to the laboratory and mounted on the servohydraulic shaker facility where a series of shaker tests were undertaken to attempt to replicate its on-road behaviour. This phase of testing is covered in chapter 3.

In order to ensure that the results were not specific to the particular suspension type, the vehicle was fitted with air suspension and the previous two sets of tests were repeated. Changes in load distribution were also investigated. These experiments are described in chapters 4 and 5.

During this project, New Zealand became involved in the OECD DIVINE programme, in particular with Element 1, which was an accelerated pavement test carried out at the Canterbury Accelerated Pavement Test Indoor Facility (CAPTIF). In order to link this project with the findings of DIVINE, which was very much focused on the same issues, a series of shaker tests were conducted on the CAPTIF vehicles. These tests are described in chapter 6.

As part of a Foundation for Research Science and Technology Public Good Science Fund funded project, which was conducted in parallel to this work, a number of simple computer models, were developed to assist with developing assessment techniques and procedures. These are described in chapter 7.

Chapter 8 covers the work on suspension assessment, which was being undertaken at other institutions during the time this project was under way and describes how the findings fit in with those of this project.

Finally chapter 9 draws conclusions and presents some options and recommendations for useable suspension assessment procedures.

2 ROAD TRIALS WITH STEEL SUSPENSION

2.1 Introduction

The test vehicle used throughout this research was a liquid tanker trailer with a single front axle and a rear tandem set. Initially this vehicle was fitted with a standard steel leaf spring suspension system. In this configuration the gross laden weight of the vehicle when filled with water was 17.5 tonnes. This level of loading is typical of the New Zealand heavy vehicle fleet. Figure 1 shows the test rig.

Figure 1. The test vehicle during road trials.

The vehicle was extensively instrumented with 26 parameters being measured. With this instrumentation in place the vehicle was driven over five test sections of pavement at three different speeds. The sites chosen included an urban section, a motorway section and three rural roads and covered a range of roughnesses. The results of these experiments provide data on heavy vehicle dynamic behaviour in the New Zealand context as well as input for the next stage of the programme.

2.2 Instrumentation and Measurement Systems

The instrumentation consisted of four set of transducers at each of the six wheel positions. These were strain gauges to measure the shear forces, an accelerometer on the axle to measure the dynamic behaviour of the unsprung mass, an accelerometer on the chassis to measure the behaviour of the sprung mass, and a linear voltage displacement transducer (LVDT) to measure suspension deflections. In addition there was a "fifth wheel" to measure vehicle velocity and elapsed distance. These 26 transducers were connected through signal conditioning and filters to a Hewlett Packard HP3852 data acquisition system. The signals were sampled digitally at 100 Hz per channel. As the principal body modes of the vehicle were expected to be between 2 and 5 Hz and the modes of the unsprung masses in the range 9-15 Hz this sampling rate was considered sufficient to obtain a satisfactory representation of these responses. The selection of a sampling rate is a trade-off between a high value to maximise the detail available and a low value to keep the data volumes manageable. Even at 100 Hz sampling each test run produced a 625kb data file.

A number of issues arose during the instrumentation and calibration, which will be detailed below. Vertical wheel forces were to be measured using strain gauges applied to the axles between the spring hanger and the wheel. Two options were available here and other researchers have used both in the past. One was to use gauges on the top and bottom faces of the axle to measure the bending strains. The alternative was to use gauges along the neutral axes of the axle to measure the shear strains. These gauge locations for these two options are illustrated in figure 2.

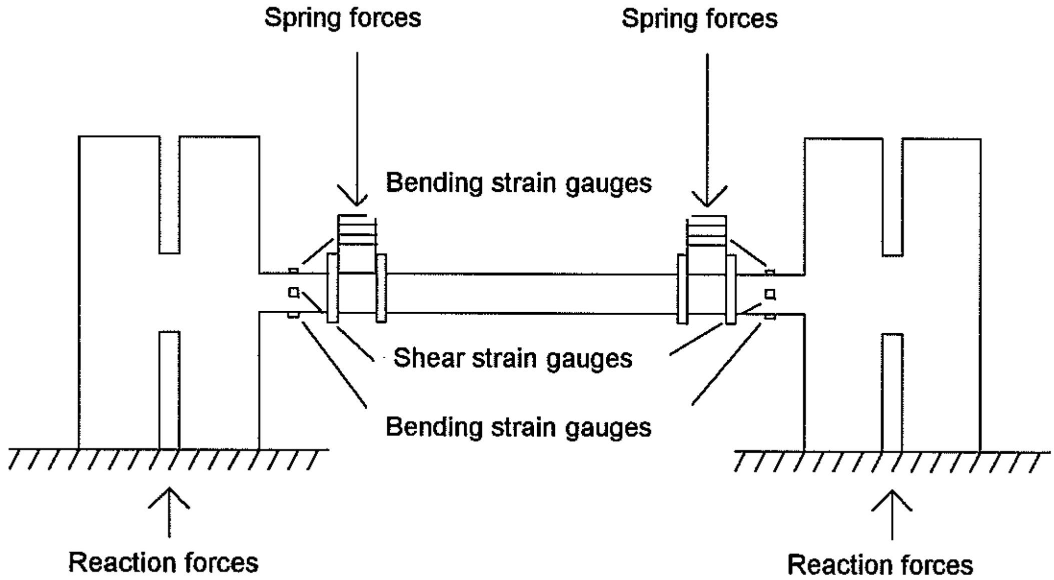

Figure 2. Strain gauge locations.

As the length of axle between the wheel and the spring hanger is relatively short and the axle is quite stiff the strain levels that will occur are quite small. The advantage of using bending gauges is that the bending strains should be larger than the shear strains and so will produce higher signal levels. This will result in a better signal to noise ratio. However, because the distance from the wheel hub to the gauges is necessarily small compared to the radius of the wheel, lateral forces applied to the wheel at the road interface will contribute significantly to the bending moment as measured by the gauges. This will result in errors in the computed vertical forces. Using shear gauges avoids this problem but at the expense of a reduced signal level.

To investigate these two alternatives one axle was instrumented with sets of both types of gauges. These were calibrated statically using the servo-hydraulic rams to apply loads to the wheels. Although approximate calculations before testing suggested that the signal level from the bending gauges might be up to ten times larger than that from the shear gauges, the actual factor measured was less than three. Because of their other advantages it was decided to use the shear gauges. This proved to be the correct decision. Although at the time this instrumentation was installed most other researchers were using bending gauges, the problems with these are now well recognised and the more recent studies all favour shear gauges.

Initially the strain gauges were calibrated statically using the servo-hydraulic shakers to apply the loads. However, this proved to be very time consuming because it was not possible to move the full vehicle off the shaker support rig in its initial configuration. Between calibrations on different wheels the vehicle had to be emptied, moved, and then refilled. Consequently a calibration procedure, which involved removing the wheel and applying a hydraulic jack through a load cell directly onto the hub to apply the loads, was developed.

The accelerometers used were of the piezoelectric type. Thus their frequency response is linear only down to a particular limit frequency which in this case was 1 Hz. However, as the primary modes of the vehicle were expected to be between 2 and 5 Hz for the sprung mass and 10-15 Hz for the unsprung mass, this lack of low frequency response was not expected to cause problems. The calibration factors for the accelerometers were given in their specification sheets and tests with a reference shaker indicated that these were correct. The accelerometers were attached to the trailer by a screw mounting into aluminium blocks, which were attached to the vehicle using epoxy glue. Figure 3 shows the accelerometers and strain gauges.

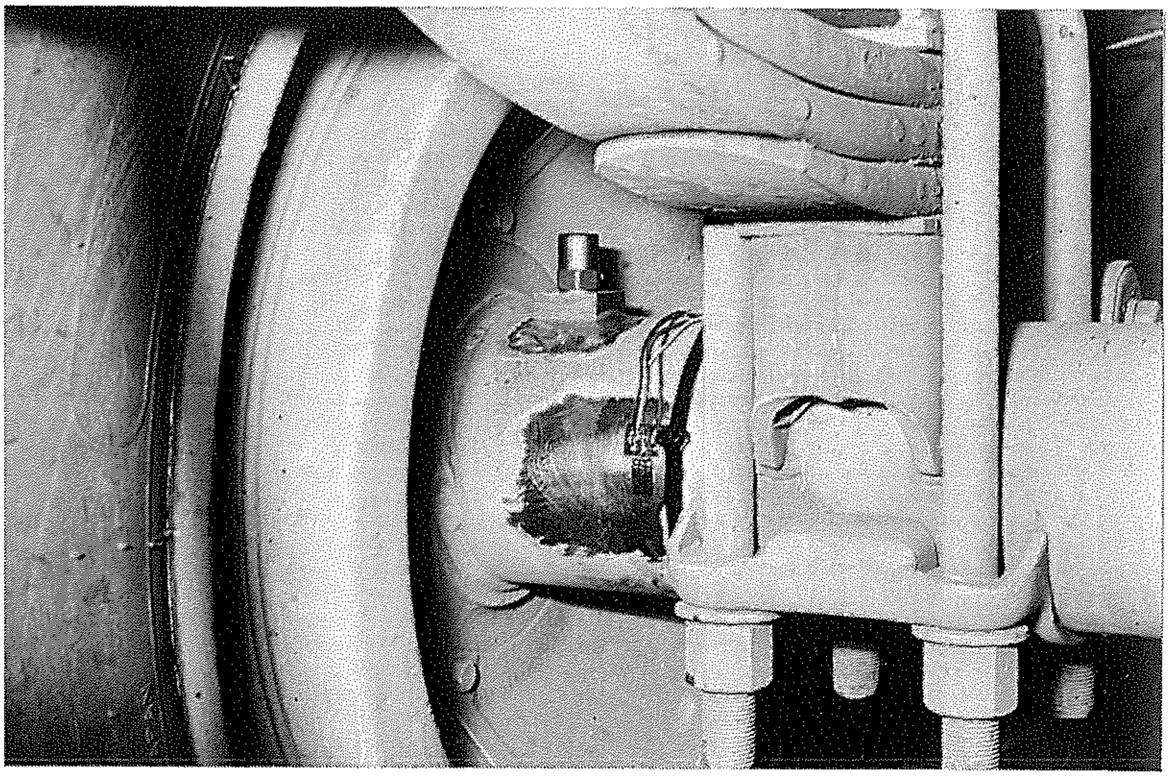

Figure 3. Strain gauges and accelerometers.

The LVDTs were attached to the vehicle using purpose built fittings, which were attached to the trailer structure with machine screws and hose clips. Calibration was done in situ by undoing the slider from the vehicle and setting it against gauge blocks.

All the calibrations outlined are static. It can easily be shown that the force of the wheel on the road is given by the shear force as measured by the strain gauges plus an inertial component from the portion of the unsprung mass that is outboard of the strain gauges. That is

$$F_{wheel} = F_{shear} + ma$$

A derivation of this relationship can be found in (Woodrooffe et al. 1986). Thus it is relatively simple to calculate the dynamic wheel forces for a test from the shear gauge signal and the axle accelerometer signal.

The mass term in this expression was initially determined quite crudely from manufacturer supplied data. The mass of the axle assembly was given as 335kg. Between the strain gauge locations the axle is a tube 1.077m long with an outer diameter of 125mm and inner diameter of 100mm. From these dimensions the mass of this section was calculated to be 37kg. Thus the outboard mass associated with each wheel position is (335 - 37)/2 = 149kg. The wheels were weighed and had a mass of 92kg each, giving a total mass outboard of each wheel of 333kg. During the shaker testing program, which is described in chapter 3, this calculated value was checked by dynamically calibrating the system. The previous equation can be rewritten as

$$m = \frac{F_{wheel} - F_{shear}}{a}$$

The wheel force term was measured with load cells on the servo-hydraulic platforms. Simultaneously, the shear gauge signal and the accelerometer signal were also recorded. Initially this was done while the servo-hydraulic shakers were excited with white noise. However, it was found that clearer results could be obtained by exciting the rams with a sinusoidal signal of a frequency close to the axle resonances of the vehicle (15 Hz was used). This is because around those frequencies the inertia correction term has its maximum effect and so the response signal levels are greater and accuracy is better. The mass as determined by this method was also 333kg. While it is remarkable that the two values were exactly the same it should be noted that even relatively large variations in this term have only a minor effect on the resultant wheel forces.

2.3 The Test Sites

2.3.1 Location

The five test sites were chosen to include a mixture of urban, rural and highway roads as well as a range of roughness. In order to measure only the variations in dynamic wheel forces caused by pavement condition it was a requirement that the test sections be straight (to avoid cornering forces) and reasonably flat. In addition they had to be sufficiently long to be able to obtain data at a relatively constant road speed for a reasonable duration, also allowing for space to accelerate up to this speed and for braking at the end. In practice, this meant about 1.7 km for an open road section and somewhat less for an urban road. These conditions proved surprisingly restrictive but five sites were found. They are:

Church St, Penrose	from Neilson St intersection to Mays Rd intersection.
Southern Motorway	from south of the Manurewa off-ramp to the Takanini off-ramp.
Airfield Rd, Takanini	from Porchester Rd intersection to Mill Rd intersection.
Mill Rd, Takanini	from Alfriston school to Airfield Rd intersection.

Takanini-Clevedon Rd from Alfriston-Ardmore Rd to Papakura-Clevedon Rd.

2.3.2 Roughness measurement

The roughness values at these sites were not known at the time they were selected. Although it had been arranged for the Australian Road Research Board (ARRB) to bring their laser profilometer to New Zealand to measure the road profiles and determine the roughness, there were a number of delays and these measurements did not take place until well after the road tests were completed. The original estimates of roughness were achieved by subjective estimation from ride quality.

The road profiles of the test sites were measured using the ARRB laser profilometer. This device consists of an aluminium beam mounted to the front of a standard passenger car (in this case, a 1985 Holden Commodore station wagon) as illustrated in Figure 4. Mounted on this beam were three laser transducers for measuring height off the ground - one in each wheel path and one in the centre. For this work a special beam was made to give a track width of 1.83m corresponding to the normal width of a truck with dual tyres. In line with each laser, an accelerometer was mounted to measure the vertical motion of the beam. By subtracting the vertical displacement calculated from the accelerometer signal from the distance to the pavement measured by the lasers the profile of the pavement could be calculated. A transducer for measuring the distance travelled was attached to the right front wheel of the vehicle and was used to trigger the laser measuring system. Measurements were taken at a rate of 20 per metre (i.e. every 50mm). These could then be extracted as a profile data file or used to calculate pavement roughness measured in either NAASRA counts or using the International Roughness Index (IRI). A more detailed description of the device can be found in (Prem 1987)

Figure 4. ARRB laser profilometer.

Two practical difficulties arose in using the profilometer data. Firstly the start of a measurement run was manually triggered by a keystroke on the computer doing the data acquisition. This was typically accurate to about 2-3 metres, which is more than adequate for its normal typical application, which was to measure road roughness and for the calibration of road roughness meters. However, in this work we wished to reference the pavement profile to the vehicle behaviour and thus a more accurate location point was desirable. This was resolved by placing a 200mm wide, 19mm thick plank in the left-hand wheel path at the start point. The profile of this plank showed up quite clearly in the data. The second problem concerned some rather large low frequency components in the profile data. The most obvious of these was a single rise and fall of typically 2-3 metres along the test section when to all intents and purposes the road was flat. These low frequency components were not repeatable from run to run. In calculating the profile, the acceleration data was integrated twice to determine the vertical movement of the beam, which was then used to correct the laser signal. A small error in the zero of the accelerometer signal generates a large parabolic error after two integrations. For the purposes of calculating roughness measures inaccuracies in these very long wavelength components are not important. A rise and fall of 2-3 metres along a one-kilometre stretch of road does not excite any dynamic response from a vehicle suspension at normal speeds. Even for this work these components were of no real importance other than that they swamped the true profile information and made it more difficult to check the repeatability of the data and to compare it with the shaker driving signals. The obvious solution was to apply some form of high pass filter to the data. However, the problem is that the magnitude of the components to be removed was considerably greater than that of the true data. Using one of the common filters such as a four-pole Butterworth does not provide sufficient attenuation. ARRB filter the signal by applying a moving average (which is a low pass filter) to determine the low frequency components which are then subtracted from the original signal. This does indeed make the signals repeatable from one run to the next. However, from the frequency response of this filter, which is shown in figure 5, it can be seen that the characteristics of this filter are not ideal as frequencies above the notional cut-off are significantly affected. With digital filtering a much more idealised filter is achievable. Thus a filter with the characteristics shown in figure 6 was applied with the cut-off frequency at 0.01 cycles/m (i.e. a wavelength of 100m).

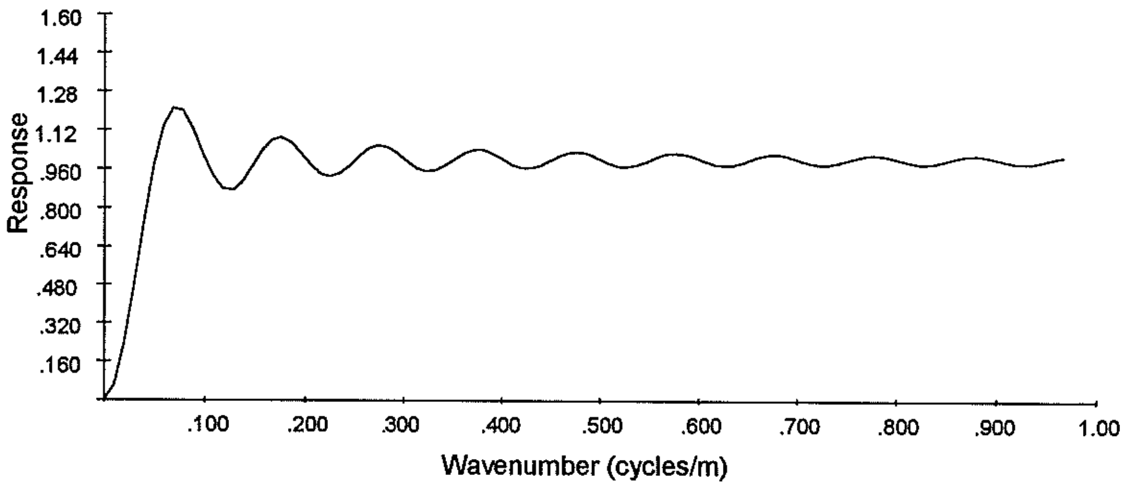

Figure 5. ARRB filtering process response.

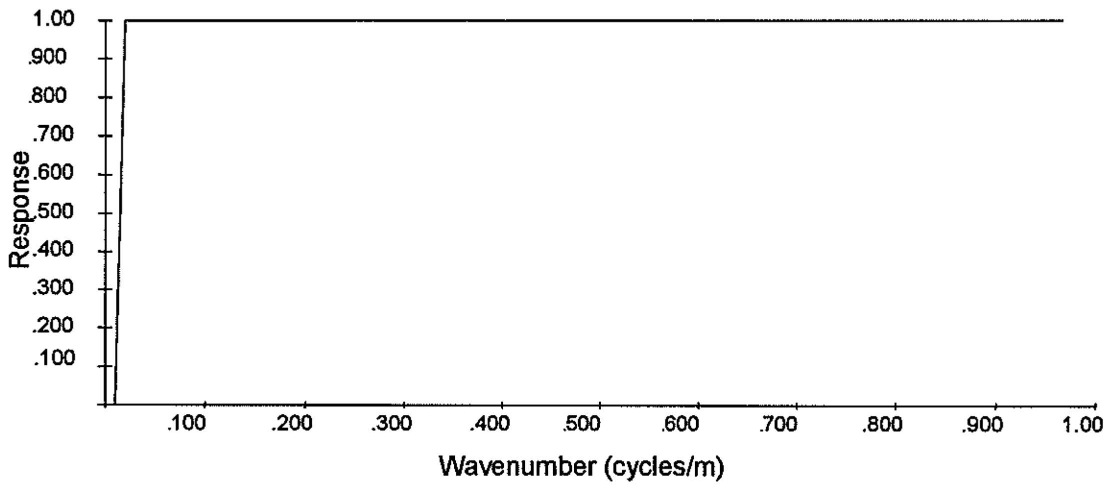

Figure 6. Idealised high pass filter.

The profile data produced were traces of vertical height against horizontal distance. To relate these to the results of the road tests and the laboratory tests they needed to be converted to traces of vertical height against time. As each of the test runs was done at a constant speed, a trace of horizontal distance against time could be produced for each. This is very simple if constant speed is assumed. However, it is slightly more accurate to use the "fifth wheel" speed data, which are readily available. At each time point corresponding to a sampling of the transducers, a value of horizontal distance was known. For this distance value, a height on the profile curve could be extracted. This could be done by linear interpolation between the two closest points. However, because the sampling rate for the profiles was significantly higher than that of the other transducers and as the change in height value between consecutive profile points was small, it was sufficient to use the closest profile point.

These road profile data were used to calculate roughness values for each site as a whole and also when divided into 200m subsections. These results are summarised in Table 1.

Table 1. Road roughness of the test sites in NAASRA counts/km.

Site	**Average Roughness**	**Range of Roughness for 200m sections**
Church St	92	71 – 112
Motorway	32	27 - 37
Airfield Rd	98	68 - 140
Mill Rd	87	45 - 190
Takanini-Clevedon	77	48 - 95

These roughness values were averaged over several repeat measurements with only small variations between runs noted. (Sweatman 1983) used sites with roughness values ranging from 22 - 194 NAASRA counts/km. Although the average values for sites in this study do not cover the same range, the values for the shorter sub-sections do. Furthermore, even with quite extensive searching, no sites were found in the Auckland region, which met the other requirements and had roughness values outside the range of the test sites.

A more detailed analysis of the roughness of the sub-sections and the relationship between these and the dynamic loads generated by the vehicle was undertaken as part of another Transit New Zealand research project (de Pont 1994).

2.4 The Test Program

At each site, tests were undertaken at several target vehicle speeds: 45, 50 and 55 kph for Church St, 75; 80, 85 and 100 kph for the motorway; and 75, 80 and 85 kph at the other three sites. In fact, the Airfield Rd site was not long enough to achieve and maintain 85 kph over the test pavement. These speed ranges represent a slightly different approach to that used by other researchers who have tended to cover a wider range of speeds (for example, (Sweatman 1983) and (Woodrooffe et al. 1986) both used 40, 60 and 80 kph). The reason for this is twofold. Firstly there was no need to investigate the general relationship between dynamic wheel forces and vehicle speed as this is quite well established. Secondly by choosing a narrow range there was an opportunity to consider the proposition (Heath 1987) that for a particular piece of pavement some critical speeds will cause far higher wheel forces than others. This is based on the idea that at these critical speeds the particular unevennesses in the pavement will couple with the natural frequencies of the vehicle to magnify the effect. While this work is not an investigation of this hypothesis, there is more chance of observing this phenomenon with relatively closely spaced speeds.

The vehicle was filled with water and weighed axle by axle on a weighbridge to measure the static loads. These are shown in Table 2. It is notable that although the rear and middle axles were part of a tandem set which is "load sharing" they were not carrying equal loads. The variation was, however, within the 10% allowed by the regulations.

Table 2. Static Axle Loadings.

Axle Position	Weight(kg)
Rear	5800
Middle	5230
Front	6260
Total	17290

Tyre pressures were all set to 690 kPa, which is typical for normal heavy vehicle operations. This was important as it turned out that the tyre stiffness was of similar magnitude to the spring stiffness thus formed an important part of the suspension system. A change in tyre pressure has a significant effect on dynamic response.

2.5 Results

2.5.1 Dynamic wheel forces

Before commencing road tests the acquisition system was triggered to run through a dummy test with the vehicle totally stationary. These measurements provided an estimate of the level of noise that could be expected. The measured data are the sum of the true signal and the noise. Assuming these two are independent the variance of the data is equal to the sum of the variance of the signal and the variance of the noise. As the DLC is the normalised square root of the variance of the wheel forces it can be adjusted to correct for the noise on this basis. If the signal to noise ratio is high the error in not correcting is small.

As mentioned previously the piezoelectric accelerometers used had a linear frequency response down to 1 Hz with a declining response below this frequency. As the literature suggested that the body modes of the vehicle should be around 2 - 4 Hz these were expected to be adequate. However, a spectral analysis of the wheel forces for a typical run showed a significant mode at about 0.5 Hz as in Figure 7. It was thought that this mode corresponded to side to side sloshing of the water in the tank. The amount of baffling in the tank was limited. If we look at the theory of standing waves in lakes, called seiches, as described by (Barber 1969) we find that the period for a cross-section shape of the approximate dimensions of the trailer is about 1.8 seconds. This matches well with the liquid slosh hypothesis. The axle-mounted accelerometers will under-respond at this frequency. However, as this mode principally involves roll of the vehicle body very little axle motion will occur. In any case, the inertial correction term overall was typically very small being of the order of 10% of the wheel force (some studies have ignored it). Thus the error from the under response of the accelerometers at the very low frequencies was likely to be very small.

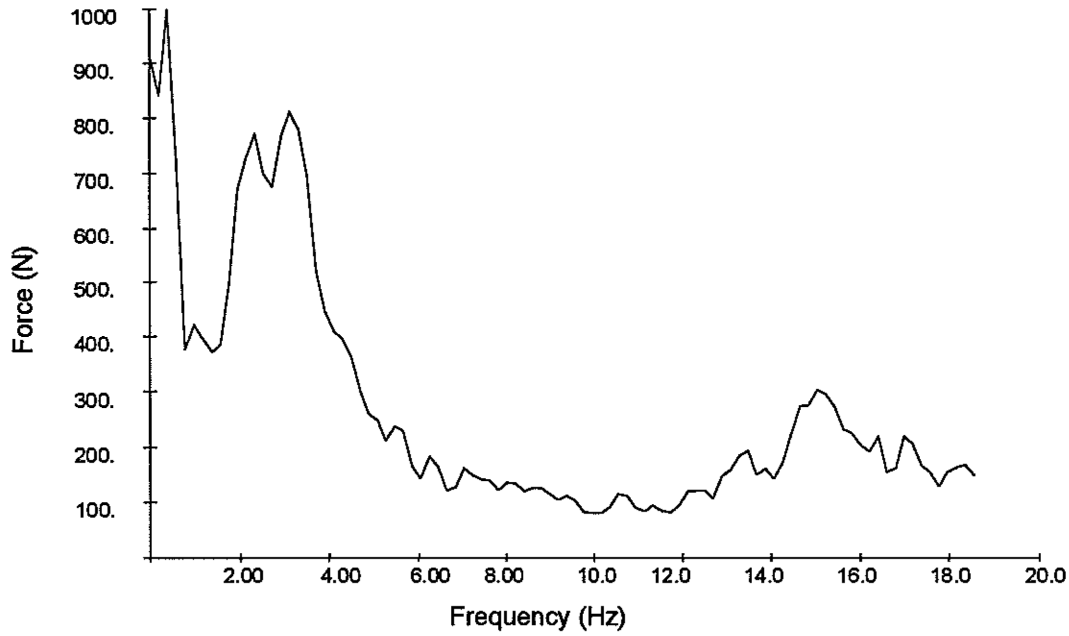

Figure 7. Spectral analysis of wheel forces.

Dynamic wheel forces are usually characterised used DLC as described in section 1.2. Initially the DLC values were calculated using the static loads as obtained from the weighbridge for each wheel as the denominator terms. However, this approach resulted in the forward axle of the tandem set consistently having a higher DLC than the rear axle. Assuming that the difference between the axles on the weighbridge was caused by stiction in the suspension and that, on average, the system did load share between the two axles, each axle can be assigned a static load equal to half the load on the system. Calculating the DLCs on this basis gives a greater similarity between the two axles of the same group. A summary of the DLCs obtained is presented in Table 3.

Table 3. DLC values from road test.

Site	Average speed	Left rear wheel	Left middle wheel	Left front wheel
Church St	45	0.122	0.132	0.137
Church St	50	0.135	0.145	0.127
Church St	55	0.148	0.162	0.151
Motorway	75	0.066	0.076	0.066
Motorway	80	0.071	0.072	0.075
Motorway	85	0.073	0.079	0.104
Motorway	100	0.097	0.091	0.101
Airfield Rd	75	0.134	0.120	0.113
Airfield Rd	80	0.152	0.134	0.136
Mill Rd	75	0.156	0.146	0.161
Mill Rd	80	0.157	0.130	0.143
Mill Rd	85	0.178	0.138	0.169
Takanini-Clevedon Rd	75	0.127	0.127	0.123
Takanini-Clevedon Rd	80	0.144	0.140	0.137
Takanini-Clevedon Rd	85	0.153	0.154	0.160

Considering the wheels in the tandem axle group it can be seen that dynamic load coefficients varied between 0.07 for the motorway at 75 kph and 0.18 for Mill Rd at 85 kph. The Mill Rd section included a concrete bridge with quite an uneven approach, which would provide a severe test for any suspension. However, DLCs of 0.16 were recorded on Church St at 55 kph. This range of values is similar to that reported by (Sweatman 1983) for the same type of suspension.

2.5.2 Repeatability of tests

At four of the sites a repeat run was undertaken at one of the test speeds. For Mill Rd and Takanini-Clevedon Rd, the DLC values calculated for a repeat run were within 2% of the original. For Church St and Airfield Rd the DLC values from the repeat runs were up to 10% different from the original values. The speed and lane position control used in the testing was crude. The driver used a digital display of the fifth wheel speed output and did his best to maintain a constant reading. Track repeatability was simply achieved by the driver attempting to keep to the centre of the lane. However, even for Church St and Airfield Rd, the measured fifth wheel data showed only small speed variations between repeat runs and so this was probably not the cause of the dynamic load differences. These two sites also had a greater variability in road roughness values between repeat runs of the profilometer which suggest a greater sensitivity to track position.

Figure 8 shows the wheel force signals for the left rear wheel at the beginning of the Church St test for the 50 kph run and the repeat run. In Figure 8a wheel force is plotted against time and it can be seen that the two signals are starting to diverge slightly in time even after only 3 seconds or so. However, the magnitude of this divergence is only 0.04 seconds. Based on a vehicle speed of 50 kph this means that after travelling 40 metres the vehicle position is 0.5 metres different. This is well within the consistency of speed expected. In Figure 8b, wheel force is plotted against distance as determined by the fifth wheel signal. This improves the correlation between the two signals. Overall the two wheel force signals appear to show good repeatability.

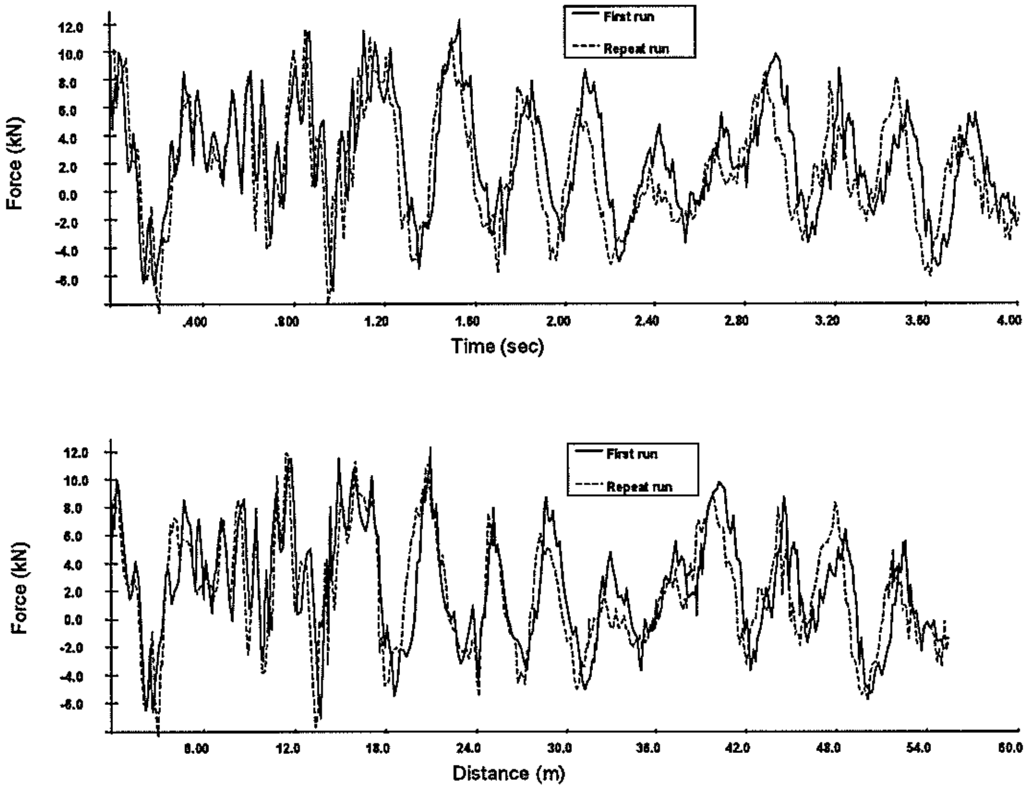

Figure 8. Comparison of wheel forces from repeat test runs.

2.5.3 Speed and roughness effects

In general, dynamic load coefficients increased with vehicle speed as expected. The main exception to this was at Mill Rd where the 75 km/h run produced slightly higher DLC values than the 80 km/h run. However, the behaviour of the vehicle on this section was complicated by the bridge, which generated relatively high dynamic loads. Considering only the data after the bridge gave a lower DLC for the 75 km/h run than for the 80 km/h one. There was little evidence in the data to support Heath's proposition, mentioned earlier, that lower speeds might induce higher dynamic wheel forces because the forcing function generated by the pavement might excite the natural resonance modes of the vehicle. The experiments were not specifically designed to test this hypothesis and it has by no means been disproved either.

Similarly dynamic loads increased with pavement roughness. This issue was analysed in more depth by (de Pont 1994) where the test sites were broken down to 200m sections and the DLC values calculated for each of these. Using this approach a much greater range of roughness values were present. This analysis showed that the dynamic loads were linearly related to the roughness, which contrasts with (Sweatman 1983) who found that dynamic loading was proportional to the square root of roughness.

2.5.4 Dynamic load sharing

Axles in a group are generally configured to equalise the load so that each carries an equal share. (In some cases, for example, where the axles have different tyre configurations, the load sharing is not equal but there is still a mechanism for re-distributing the loads). Static load sharing is a legal requirement but dynamic load sharing is not. The importance of load sharing comes about because of the fourth power relationship between wheel loads and pavement wear. If we consider two axles in a group with perfect load sharing causing one unit of pavement wear each, the group will generate two units of pavement wear. If the two axles do not share the load well and one axle is carrying 1.1 times the mean load while the other is carrying 0.9 times the mean load, then the resulting pavement wear will be $1.1^4 + 0.9^4 = 2.12$ units of pavement wear, that is, a 6% increase. Changing the axle load variations to 1.2 and 0.8 of the mean load leads to an increase in pavement wear of 24%. Thus static and low speed (quasi-static) dynamic load sharing are clearly very important. When considering full dynamic response it is important to reduce the level of dynamic loads. Dynamic load sharing will tend to contribute to this reduction by averaging the forces between the axles. However, the primary factors are still the suspension performance and the magnitude of the excitation, which is determined by road roughness and vehicle speed.

The issue of dynamic load sharing has been raised by a number of authors but does not appear to have genuinely been measured. (Sweatman 1983) compares the average of the dynamic wheel forces with the expected static load as a measure of dynamic load sharing. This is not truly a measure of dynamic load sharing but of the average load sharing behaviour. If, at any instant, the force from a wheel is divided by the average for the set of wheels in its axle group at that same instant, the resulting value reflects the dynamic load sharing of the group. Clearly the average of this signal is expected to be unity and its standard deviation gives a measure of the dynamic load sharing capability of the suspension. As this measure is already dimensionless it can be called the dynamic load-sharing coefficient. Mathematically this can be expressed as follows:

$$DLS_i = \frac{nF_i}{\sum_{i=1}^{i=n} F_i}$$

where DLS_i = dynamic load sharing at axle i

F_i = instantaneous wheel force at axle i

and n = number of axles in the group

The dynamic load-sharing coefficients are the standard deviations of the dynamic load sharing functions DLS_i. If there are only two axles in the group then $DLS_1 = 1 - DLS_2$ and the dynamic load-sharing coefficient is the same for both axles.

Table 4 shows the values of this measure.

Table 4. Dynamic load-sharing coefficients.

Site	45kph	50kph	55kph	75kph	80kph	85kph	100kph
Church St	0.067	0.079	0.100				
Motorway				0.030	0.030	0.030	0.033
Airfield Rd				0.053	0.052		
Mill Rd				0.070	0.056	0.077	
Takanini-Clevedon Rd				0.055	0.058	0.070	

As can be seen this coefficient varies between 0.03 and 0.10 and approximately follows the trends of the DLC values, which were presented in Table 3, with, in general, a higher DLC implying a higher load-sharing coefficient. The load-sharing coefficient appears to be typically a little under half the DLC.

An alternative measure of the load sharing capability of a tandem suspension is to consider the difference between the wheel force signals from the two axles in the group. If the two signals were independent (i.e. no load sharing at all) the variance of this combined signal would be equal to the sum of the variances of the two components. If the load sharing were perfect, the variance of this signal would be zero. Thus taking the square root of half the variance of this difference signal and normalising by the static load gives a measure which can be compared to the DLC. Mathematically this can be written as:

$$\text{Load difference coefficient} = \frac{\sqrt{\text{Variance}\,(F_1 - F_2)}}{2\,(\text{static wheel load})}$$

If this measure is equal to the DLC there is no load sharing while a low value implies good load sharing. Measurement uncertainty ensures that a zero value is not possible. The values of this measure are given in

Site	45kph	50kph	55kph	75kph	80kph	85kph	100kph
Church St	0.091	0.105	0.107				
Motorway				0.043	0.042	0.042	0.046
Airfield Rd				0.076	0.074		
Mill Rd				0.089	0.074	0.090	
Takanini-Clevedon Rd				0.075	0.077	0.088	

Table 5. Load difference coefficient.

Site	45kph	50kph	55kph	75kph	80kph	85kph	100kph
Church St	0.091	0.105	0.107				
Motorway				0.043	0.042	0.042	0.046
Airfield Rd				0.076	0.074		
Mill Rd				0.089	0.074	0.090	
Takanini-Clevedon Rd				0.075	0.077	0.088	

These values are typically a little over half the DLC values for the same test indicating that a significant amount of dynamic load sharing is occurring. Figure 9 shows a plot of the force magnitude spectra of this difference function compared to that of one of the individual wheel forces for a test run. From this it can be seen that dynamic load sharing is occurring at the low frequencies where the difference function is considerably lower than the wheel forces. At higher frequencies there is no reduction.

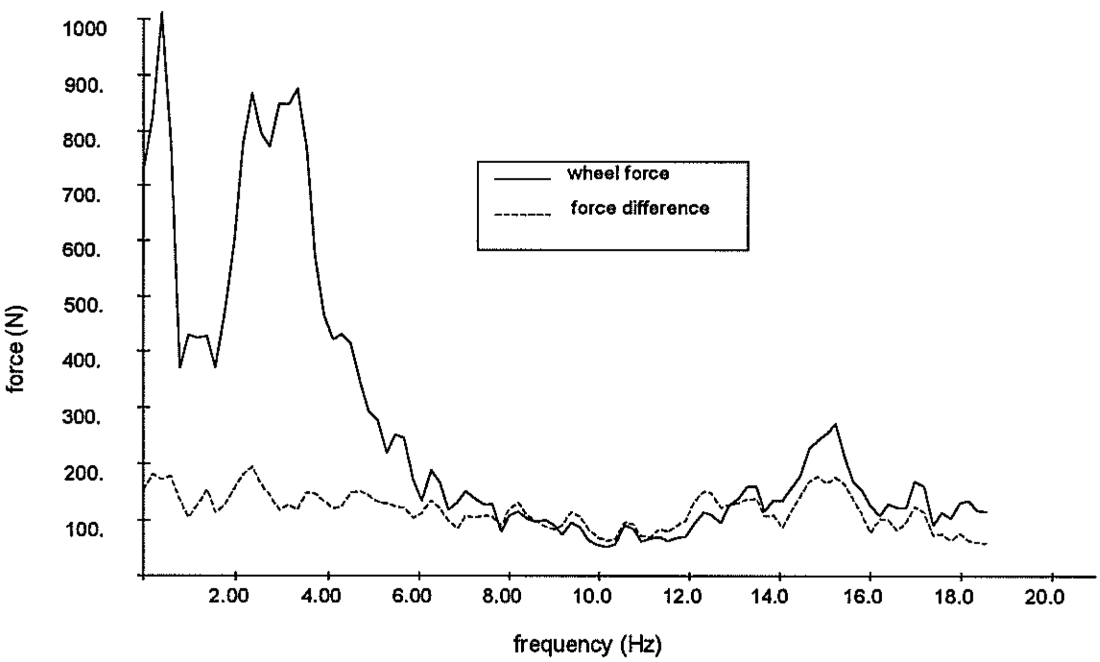

Figure 9. Comparison of load difference with wheel forces.

2.5.5 Alternative estimates of dynamic load

At this point it is worth considering whether the dynamic wheel forces can be determined by any other methods from the instrumentation that was on the trailer. The simplest and most obvious method is to take the sum of the chassis accelerometer signal multiplied by the sprung mass component for that wheel and the axle accelerometer signal multiplied by the unsprung mass component. That is,

Wheel Force = sprung mass x chassis acceleration + unsprung mass x axle acceleration

Although this method is attractive because of its simplicity it is not theoretically correct. In the first place it ignores any body-flexing behaviour of the trailer (bending and twisting of the vehicle structure). In this case, with a tanker vehicle, these effects are minor and the errors in ignoring them are negligible. However, with a more flexible vehicle this may not be so. Secondly, there are a number of whole body modes of the vehicle including bounce, pitch and roll. A modal analysis undertaken in the laboratory on this trailer found four whole body modes of the vehicle that would contribute to vertical forces. These were a longitudinal translation motion with a pitching component, bounce, pitch and roll. The first of these may have occurred because in the laboratory the park brake of the trailer was activated and thus the vehicle could rock to and fro on its brakes. Of the others only the forces associated with a pure bounce mode can be calculated by multiplying the acceleration by the mass. The vertical accelerations associated with pitch and roll modes would have to be converted to angular accelerations and then be multiplied by the appropriate moment of inertia to calculate the restoring torques generated which could then be converted to equivalent vertical forces. These would only be equal to those calculated from the mass multiplied by acceleration formula if the suspension was attached to the vehicle at the respective radii of gyration. With this particular trailer and its uniformly distributed load this was approximately so. Even then this is a simplification because, in general, the three modes are not "pure". The bounce mode has some associated pitch motion and so on.

A spectral analysis of the chassis accelerometer signals showed that, as well as the vehicle body modes at around 3 Hz, there were higher frequency components of significant amplitude. Although, some response at these frequencies was to be expected from the sprung mass reacting to the motion of the unsprung mass, the measured amplitude of these components was too high. It was not possible for the whole sprung mass to be moving at this amplitude and frequency. It appeared that flexing in the chassis member on which the accelerometer was located generated this signal. Therefore these frequencies were filtered from the signal before calculating the wheel forces by this method. This, of course, also eliminated the information on the sprung mass motion at these frequencies.

Figure 10 shows a comparison of the wheel force signals calculated by the two methods for a typical run. The upper trace shows the time domain comparison. Clearly the two signals match quite well at the lower frequencies which correspond to the sprung mass response of the vehicle but the match at the higher frequencies is not as good. Comparing the two wheel forces in the frequency domain as shown in the lower graph illustrates the differences more clearly. This shows that, as expected, the accelerometer method underestimated the 0.5 Hz component found with the strain gauges method. However, it overestimated the response at both the sprung mass and the unsprung mass frequencies with the result that the overall DLC value was similar. For the lower frequencies the overestimation came from treating all the body modes as if they were bounce. At the higher frequencies the error came from filtering the high frequency components out of the chassis accelerometer signals. This is equivalent to rigidly fixing the sprung mass for those frequencies. Clearly the sprung mass has some response to the motion of the unsprung mass and this motion reduces the wheel forces generated

A second more complex approach to calculating the wheel forces from the accelerometer signals is to try to overcome one of the deficiencies of the previous method by separating the different components of body motion and multiplying each by appropriate mass or inertia values. By linear combination of the four chassis accelerometers on the corners of the vehicle the vehicle motion can be separated into bounce, pitch and roll. The bounce acceleration is multiplied by the sprung mass component for the wheel, while the pitch and roll accelerations are multiplied by their respective moments of inertia and divided by the distance from the centre of motion to the suspension attachment point. As with the previous method the force component from the unsprung mass is then added to this to give the total wheel force.

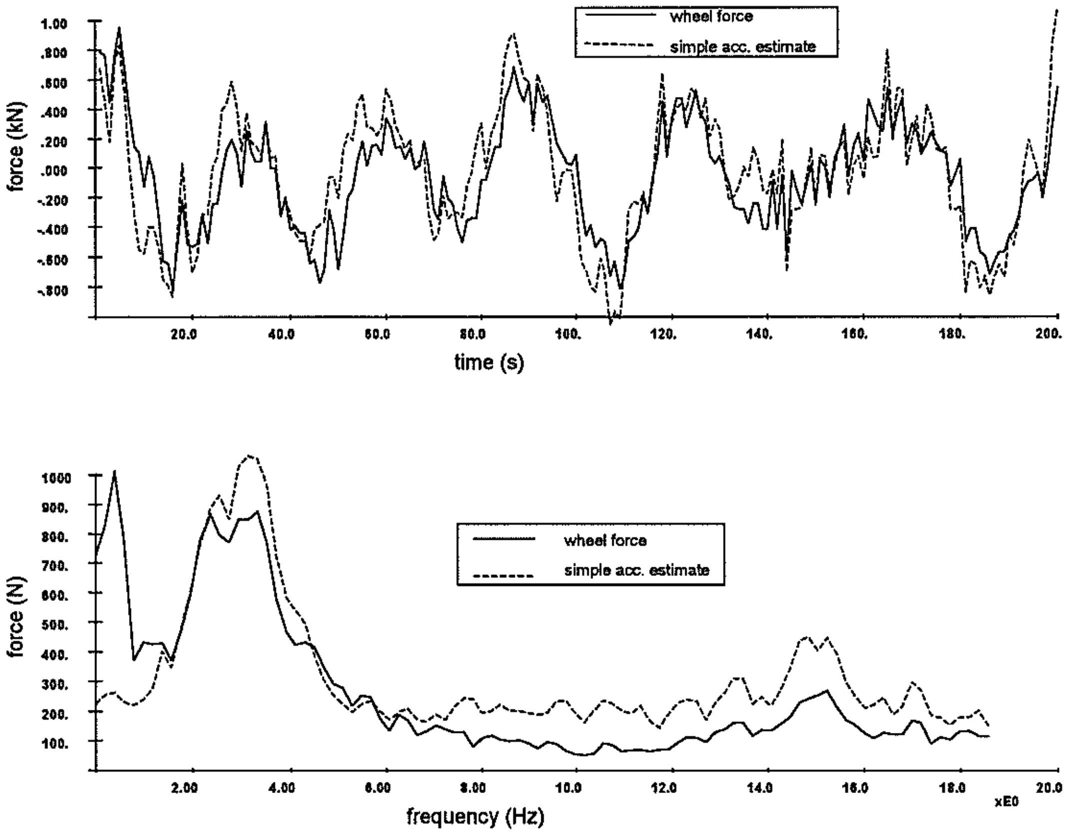

Figure 10. Wheel forces – strain gauges versus simple accelerometer method.

The roll and pitch inertia values needed for this method were calculated by assuming the trailer consisted of uniformly distributed mass (the body and load) and some lumped masses (the axle assemblies). The masses of these components were estimated from the tare weight and the tare and laden axle loads. The length, height and width dimensions were measured. An example of the wheel forces measured by the strain gauges compared with those calculated from the accelerometers in this way is illustrated in Figure 11. The lower graph, which provides a comparison in the frequency domain, shows a significantly better match at the sprung mass frequencies (2-5 Hz) than the simple accelerometer method. The match at the higher frequencies is similar. However, this method has the additional complication that the moments of inertia and the locations of the centres of pitch and roll need to be determined.

The DLC values for the left rear wheel obtained by all three methods are compared in Table 6. It can be seen that for most runs the values from both accelerometer methods were quite close to the strain gauge based values, although in the most extreme cases the differences were up to 20%. As a general trend the simple accelerometer method tended to overestimate the DLC value while the complex method tended to underestimate it. The complex method appeared to be slightly better than the simple one in that its maximum errors were lower.

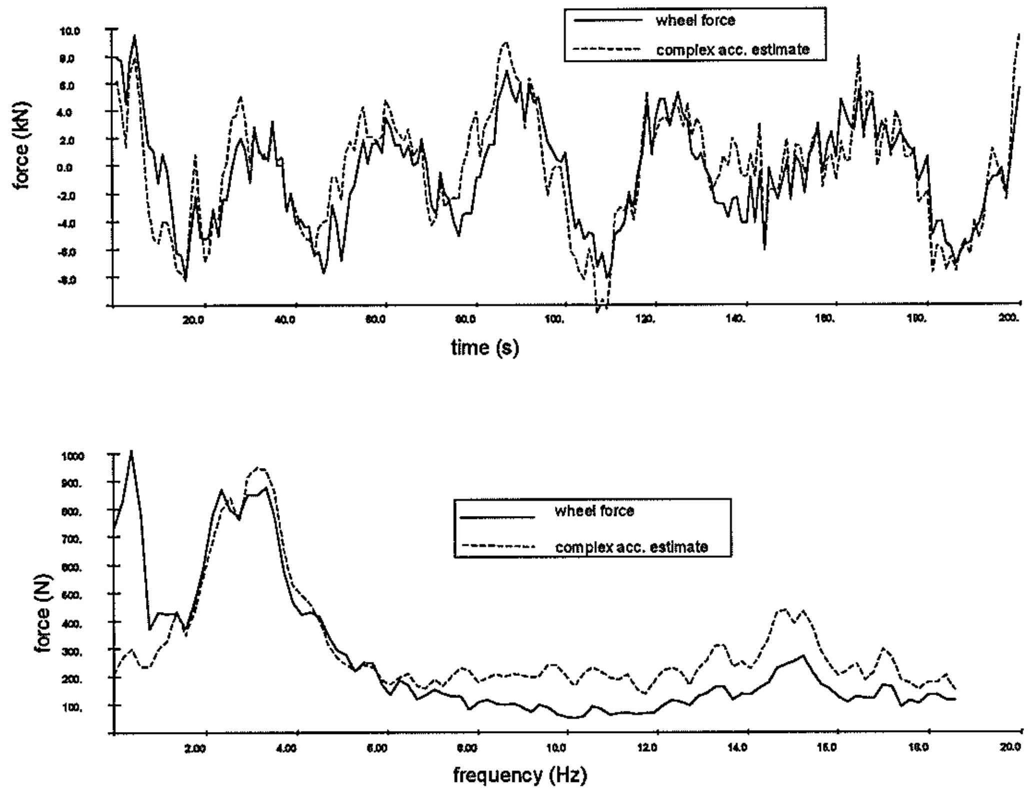

Figure 11. Wheel forces – strain gauges versus complex accelerometer method.

With both of these methods the 0.5 Hz vehicle response was not found because of limitations of the accelerometer response below 1 Hz. This problem could easily be overcome by using different accelerometers. Secondly, there appears to be flexing of the chassis members on which the accelerometers monitoring the sprung mass behaviour were mounted giving excessively high signals at the unsprung mass natural frequencies. This was countered by filtering out these high frequency components from these signals. However, this also filters out the true high frequency response of the sprung mass, which introduces other errors. Locating the chassis accelerometers differently could eliminate this difficulty and improve the accuracy of accelerometer-based methods of estimating wheel force. As these effects, which are not accurately modelled by these methods, were relatively small the overall estimates of DLC were reasonably accurate, though with the changes mentioned a significant improvement should be possible.

However, there were a number of factors related to this vehicle which would enhance the likely success of accelerometer based methods which would not always be the case. As a tanker vehicle it's body is relatively stiff reducing the likelihood of flexure of the vehicle structure occurring and affecting the measured accelerations without corresponding wheel force effects. The wheel forces compared in Table 6 come from an axle, which is part of a tandem group. Thus there is an implicit assumption in the two accelerometer methods that the forces generated by the sprung mass bounce behaviour are equally shared. As was shown previously the particular suspension used has quite good load sharing characteristics at low frequencies and so this assumption is valid. This may not be true of other suspension types.

Table 6. Comparison of DLC values for left real wheel by the three methods.

Site	Average speed	Strain gauge method	Simple accelerometer method	Complex accelerometer method
Church St	45	0.122	0.146	0.134
Church St	50	0.135	0.153	0.143
Church St	55	0.148	0.155	0.145
Motorway	75	0.066	0.070	0.066
Motorway	80	0.071	0.072	0.064
Motorway	85	0.073	0.079	0.070
Motorway	100	0.097	0.116	0.098
Airfield Rd	75	0.134	0.123	0.115
Airfield Rd	80	0.152	0.151	0.131
Mill Rd	75	0.156	0.158	0.147
Mill Rd	80	0.157	0.160	0.144
Mill Rd	85	0.178	0.181	0.165
Takanini-Clevedon Rd	75	0.127	0.128	0.118
Takanini-Clevedon Rd	80	0.144	0.146	0.133
Takanini-Clevedon Rd	85	0.153	0.166	0.151

2.6 Conclusions

The dynamic wheel force measurements generated results that are very much in line with those of other researchers for similar suspensions. Dynamic wheel forces increase with increasing speed and increasing road roughness. Repeat runs at the same speed over the same road showed a high degree of repeatability in the wheel forces generated.

Dynamic load sharing was investigated to an extent not previously reported in other research. Two measures for characterising dynamic load sharing were defined. The first of these, the dynamic load-sharing coefficient quantifies the actual dynamic load sharing. It has a value of zero for perfect load sharing but the interpretation of non-zero values in absolute terms is not so clear cut. The second measure, the load difference coefficient is derived from the standard deviation of the difference between the wheel forces on wheels in the same axle group. As with the load-sharing coefficient it has a value of zero for perfect load sharing. However, it also has an upper limit in that with no load sharing it will be equal to the DLC value. These two measures show very similar trends. For the particular vehicle under test there was a significant level of dynamic load sharing. A spectral analysis on the load difference function shows that this was primarily for the low frequencies corresponding to the body modes of the vehicle.

Two methods for estimating the dynamic wheel forces using the chassis and axle accelerometer signals were investigated. For the particular vehicle used in these tests the results were reasonable. However, without the benefit of knowing the actual wheel forces as measured by the strain gauges it is difficult to be sure of the level of uncertainty in these estimates.

3 SHAKER TRIALS WITH STEEL SUSPENSION

3.1 Introduction

The vehicle was then brought into the laboratory and put on a two post servo-hydraulic shaker facility. The control system for the actuators and the output from the instrumentation were then used to replicate the LVDT responses measured during the road tests on the two wheels being shaken. The support platforms built for the trailer were fitted with load cells to measure the wheel forces directly during these shaker tests. As with the road tests, the axle accelerometers were used to monitor the behaviour of the unsprung masses. By correcting for the differences between the road and laboratory behaviour of these masses it was intended to relate the wheel forces as measured in the laboratory back to those measured on the road. If successful this would enable us to measure the dynamic wheel forces by a road test using only an LVDT and an accelerometer at each wheel of interest followed by a shaker test in the laboratory. As this instrumentation is significantly easier and faster to apply than strain gauges, this could be a cost-effective assessment technique.

3.2 The Shaker Test Facility

On completion of the road test program the vehicle was brought into the laboratory for testing on the servo-hydraulic shaker facility. The servo-hydraulic facility consists of two 5 tonne rams which are capable of a frequency response of 0 - 100 Hz. For testing the vehicle a special rig as illustrated in Figure 12 was built. One of the key features of the support rig is the use of air bags, which were designed to support the static weight of the vehicle. Thus the rams are only required to provide the dynamic component of wheel force. Load cells were fitted under each platform so that the wheel forces could be measured directly from the rig.

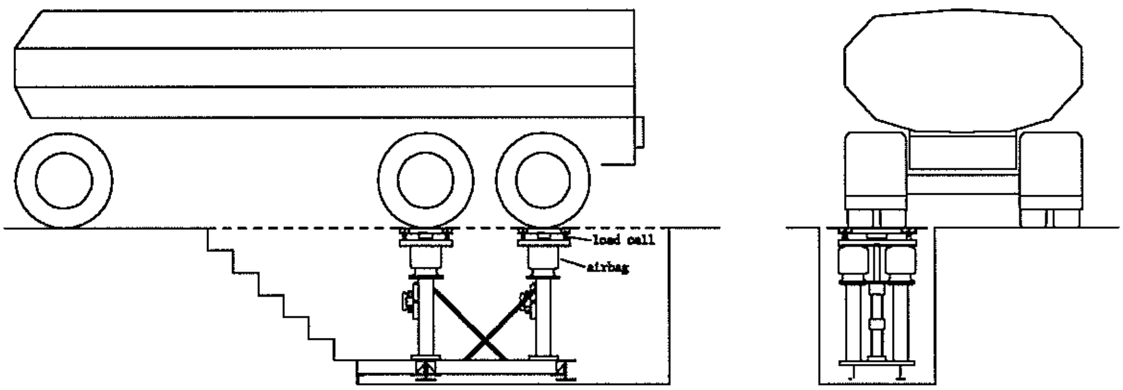

Figure 12. Vehicle support for shaker rig.

Each ram was controlled by a microprocessor system, which could generate sine waves, square waves and triangle waves or match an externally supplied arbitrary force or displacement signal. The Hewlett Packard HP3852 data acquisition system, which was used in the road tests, also has a waveform generating module which was used to provide the input signal for the rams. Control of the entire system was via a personal computer attached to the data acquisition unit.

3.3 The SYSCOMP Shaker Control Software

For any mechanical system being shaken by the rams the response of the system is not the same as the input. In many applications the aim is to generate a particular response at some location in the system rather than to send a particular input. Even in the simple case where the target response is at the interface between the ram and the test object, the mass and stiffness of the object and supporting structures influence the system response so that the response is not identical to the demand input signal. In order to determine the input signals needed to achieve a desired response the following iterative algorithm was developed and encoded into a software package known as SYSCOMP (Wester 1989).

- An initial input signal (usually noise of an appropriate bandwidth) is sent to the ram and the response of the system is measured.
- The transfer function between the input and the response is calculated and then inverted.
- The inverse transfer function is applied to the target response to generate a new input signal to the ram.
- The new input signal is applied and the response is measured and compared to the target. If the match is satisfactory the process is complete, otherwise the previous two steps are repeated.

If the dynamic system is linear the process is complete after one cycle, as the transfer function is totally independent of the shape of the target response. Most systems have a degree of non-linearity (the rams themselves are non-linear) and it will be necessary to iterate several times before the desired response is achieved. With both rams operating on a coupled system the problem becomes more complicated as there as two inputs, two responses, and four transfer functions but the principles are the same.

This algorithm proved to have a number of problems when applied to this situation, which are described in the following sections. During the testing significant modifications were made to the algorithm to improve its stability speed and performance. These will be described as part of the description of the testing which follows.

3.4 The First Shaker Tests

The trailer was mounted as shown in Figure 13 with the left-hand side of the rear axle group on the rams with the same loading and tyre pressures as for the road tests. The target response functions used were the suspension deflections as measured during road tests by the LVDTs. One test run from each site was used. The iteration process was applied until a reasonable match between the target and the response signals was achieved. At this point the data acquisition system was used to monitor the strain gauges on the two wheels being shaken, the axle accelerometers and the LVDTs on those wheels, the chassis accelerometers at the four corners of the vehicle, the load cells on the platforms and the ram displacements.

Although the process of finding the inverse transfer functions to generate the inputs required for particular target responses appears straightforward, it proved to be surprisingly difficult for the trailer suspension. The theory behind inverting the transfer functions requires that the two input signals are statistically independent. However, the road test situation being modelled effectively has both wheels being excited by the same forcing function with only a small time delay between them. Therefore, if the inputs are calculated from the target responses from the road tests, the statistical independence condition no longer holds. Furthermore, the highly non-linear nature of the suspension caused significant problems. The friction in the springs resulted in a stiffness that was much higher at small displacements than at large ones. The effect of this was that if the first iteration generated a response with amplitudes below the target the transfer functions based on the high stiffness would lead to calculated excitations which are substantially too high. This caused increasingly large oscillations in the levels of excitation required and the procedure did not converge. Also the inverse transfer functions were, for this vehicle, very specific to the target signals and hence had to be recalculated for each test.

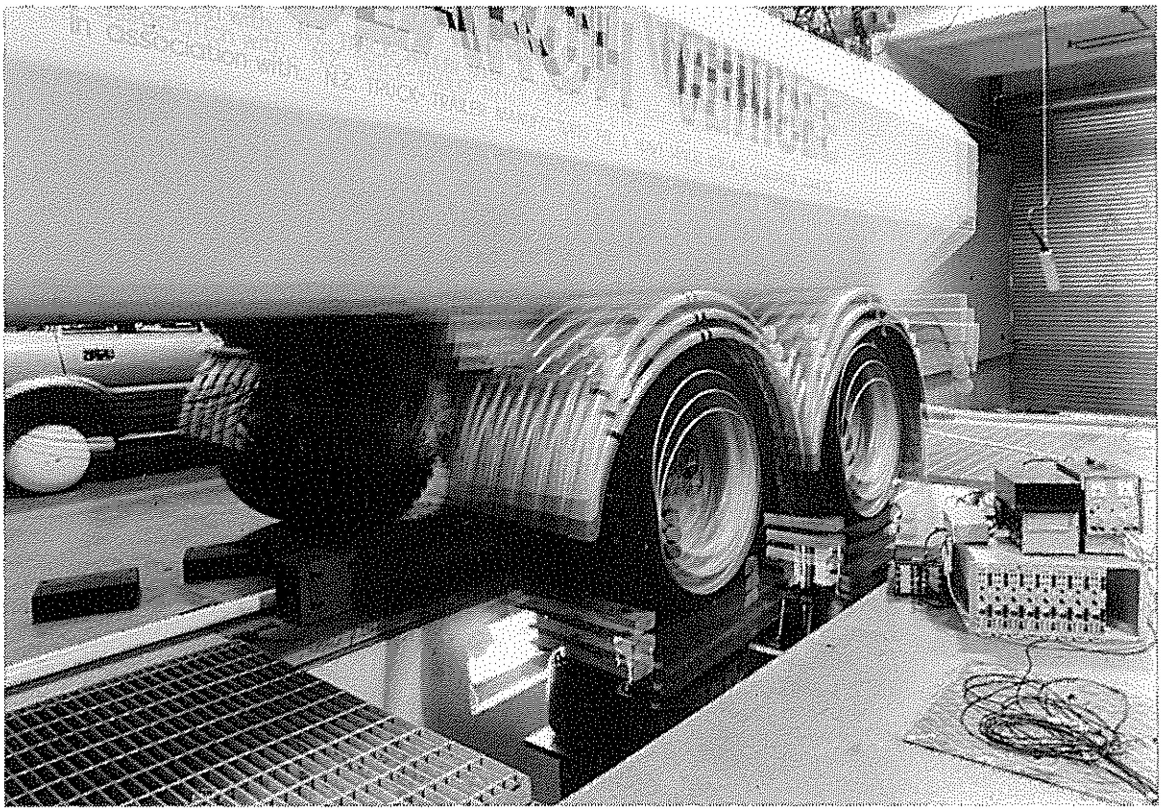

Figure 13. Shaker testing of trailer.

These difficulties were overcome by starting with pseudo target signals that were random but with the same spectral characteristics as the road test targets but statistically independent. The procedure for calculating the inverse transfer functions was modified so that, for each iteration, the new inverse transfer function was a linear combination of the old one and the newly calculated one. This slowed the rate of change of these functions and allowed the procedure to converge. Once the rams successfully reproduced the pseudo targets, the calculated inverse transfer functions were applied to the target signals taken from the road test measurements and the process continued. As the road inputs that the servo-hydraulics were trying to emulate were not statistically independent, the automated procedure did not work very well at this stage. Convergence was achieved by manually editing the inverse transfer functions over a number of iteration cycles. With practice a particular road test could be simulated in about four to five hours. In terms of the aims of the project, which were to develop a cost-effective servohydraulic shaker testing procedure, this is much too long.

One test run from each of the five sites was used for shaker testing. The match between the target signals and the response achieved on the shakers, while not perfect, is acceptable as shown by the examples in Figure 14. Applying a Fourier transform to these signals and comparing them in the frequency domain shows very good correlation as illustrated in Figure 16.

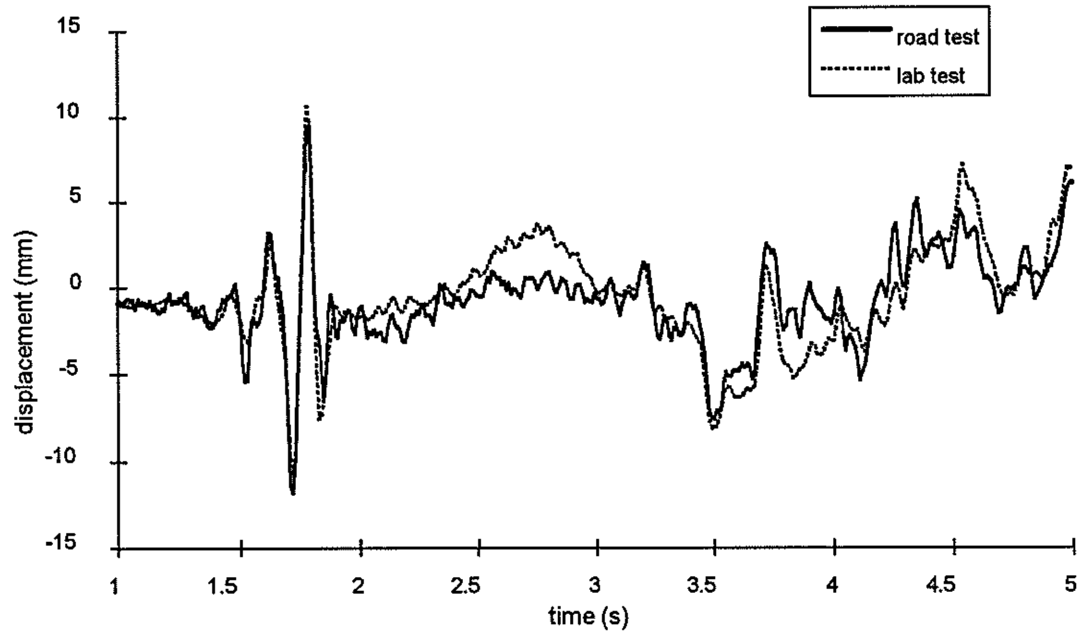

Figure 14. Rear wheel suspension displacements.

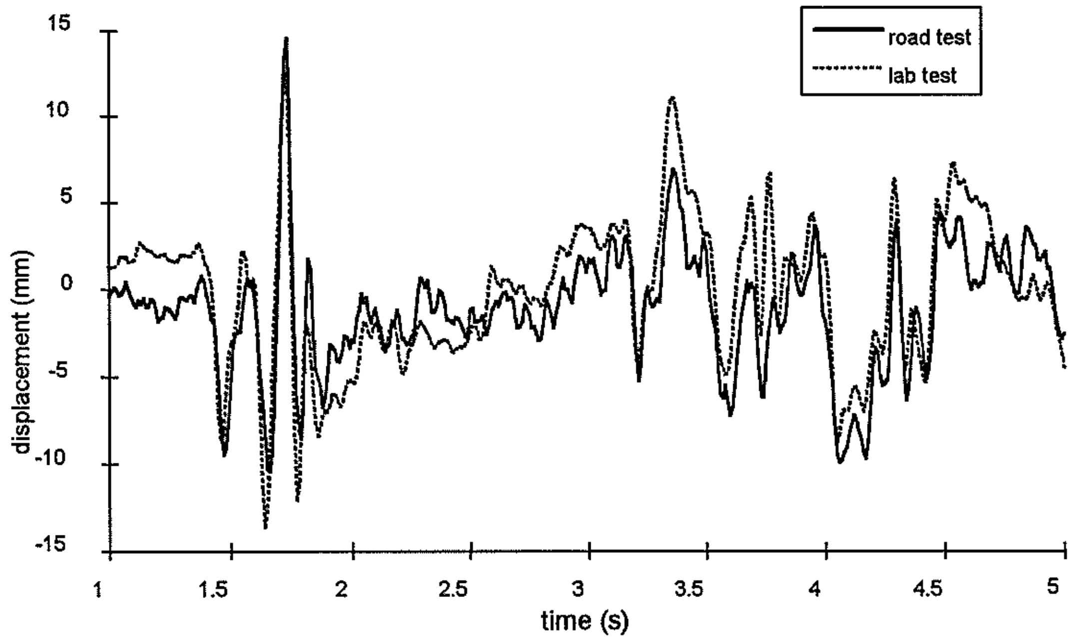

Figure 15. Middle wheel suspension displacements.

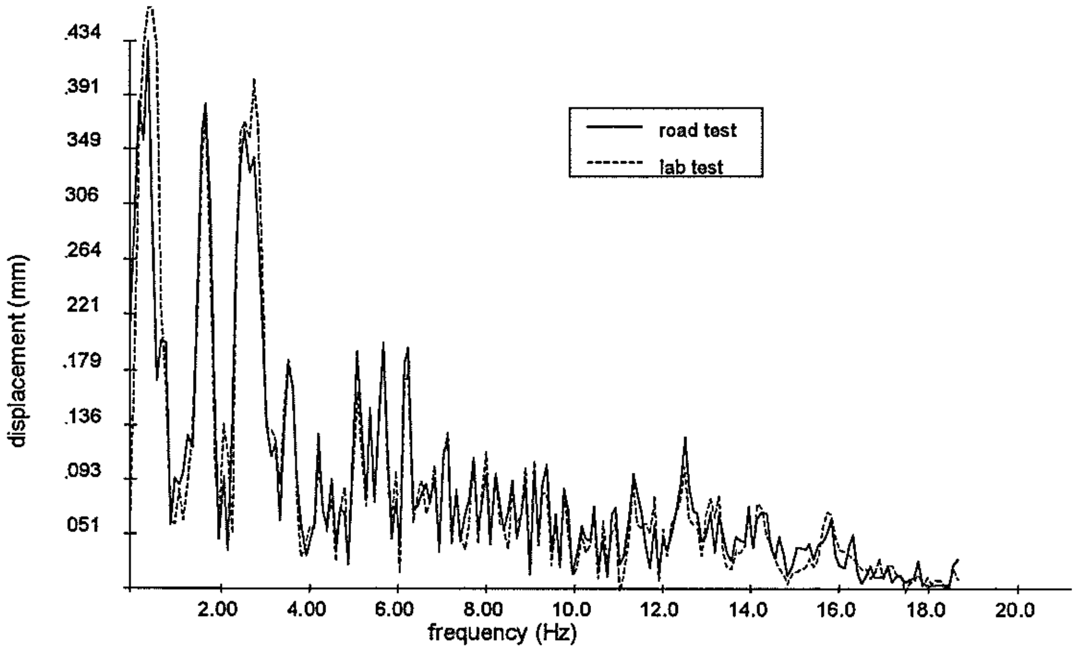

Figure 16. Suspension displacement spectra for road and lab tests.

The DLCs calculated from the wheel forces measured during the servo-hydraulic experiments are listed with the values from the comparable road test data in Table 7.

Table 7. Comparison of DLC vlues between road and shaker tests.

Site	Left rear wheel		Left middle wheel	
	Road test	Shaker test	Road test	Shaker test
Church St	0.155	0.154	0.170	0.151
Motorway	0.075	0.087	0.081	0.080
Airfield Rd	0.140	0.141	0.132	0.129
Mill Rd	0.160	0.162	0.132	0.144
Takanini-Clevedon Rd	0.157	0.151	0.160	0.132

Most of these values are very close (within 2%). For some of the runs there are larger discrepancies (up to 15%), often on only one of the two wheels. Its is likely that this match could be improved by taking more time still on additional iterations of the process to obtain a better fit. Comparing the actual wheel force signals shows a reasonable good match between the road and laboratory test results. This is illustrated by the trace shown in Figure 17, which, as in the previous two figures, comes from tests from the Mill Rd site.

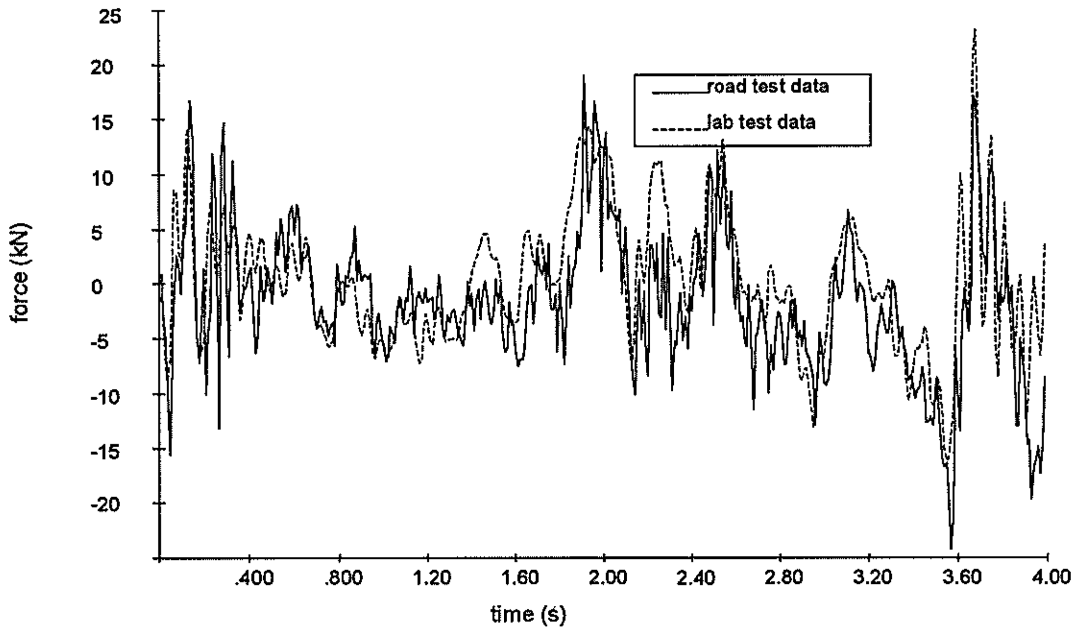

Figure 17. Comparison of wheel forces between lab and road tests.

As with the suspension displacements a spectral analysis of the wheel forces, as illustrated by the example in Figure 18, shows a very good correlation.

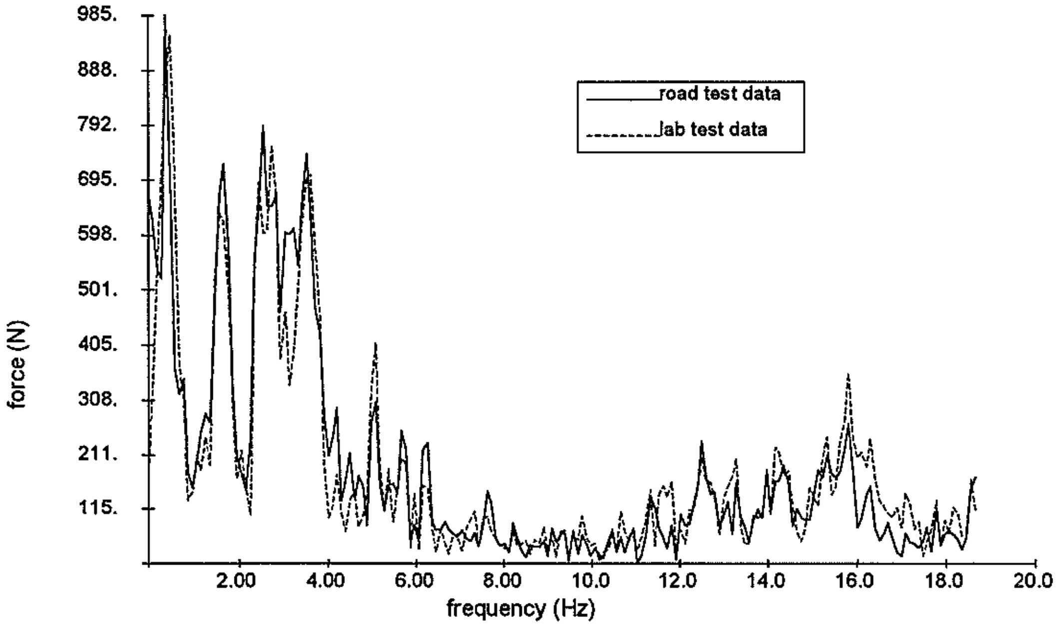

Figure 18. Wheel force spectra for lab and road tests.

There was speculation prior to the testing that by exciting the wheels on only one side of the vehicle its body roll mode would be excited to an extent that was not present during road tests. There was no evidence of this mode being excited to any greater extent during the laboratory trials. The reason for this is that the shaker excitation signals were tuned to elicit the same suspension motions as were measured. They were not directly based on the road profiles. Therefore if there was a tendency to excite the body roll behaviour of the vehicle during laboratory tests it would be offset by an appropriate reduction in the shaker excitations.

3.5 The Shaker Displacements and the Road Profiles

The approach used in section 3.4 had been to iteratively calculate the shaker excitations required to generate the same suspension responses as were measured during the road tests. Throughout this first set of tests the control of the excitation of the servohydraulic rams allowed them to act independently of each other connected only in that they were acting on a single vehicle. As the vehicle was being excited using only two wheels in an attempt to match the behaviour (albeit only at those two wheels) on the road where all six wheels were excited it seemed appropriate to maintain maximum flexibility. These ram displacements were expected to be related to the road profiles but because of the difference in the excitation mechanism during the road tests they would not be identical. In this section we review the shaker excitations and their relationship to the road profiles.

Figure 19 shows the solution for shaker displacements for the two rams for one of the tests described in section 3.4. It is clear that these two signals are neither independent of each other nor identical. The match between them at lower frequencies is good, while at higher frequencies there are differences although these are not very clearly visible on this plot which spans more than 40 seconds.

The relationship between these two shaker excitation signals is perhaps more clearly shown by their power spectral density[1] (PSD) functions, which are plotted in Figure 20. The magnitude of the excitation is greater for the middle wheel shaker than for the rear wheel.

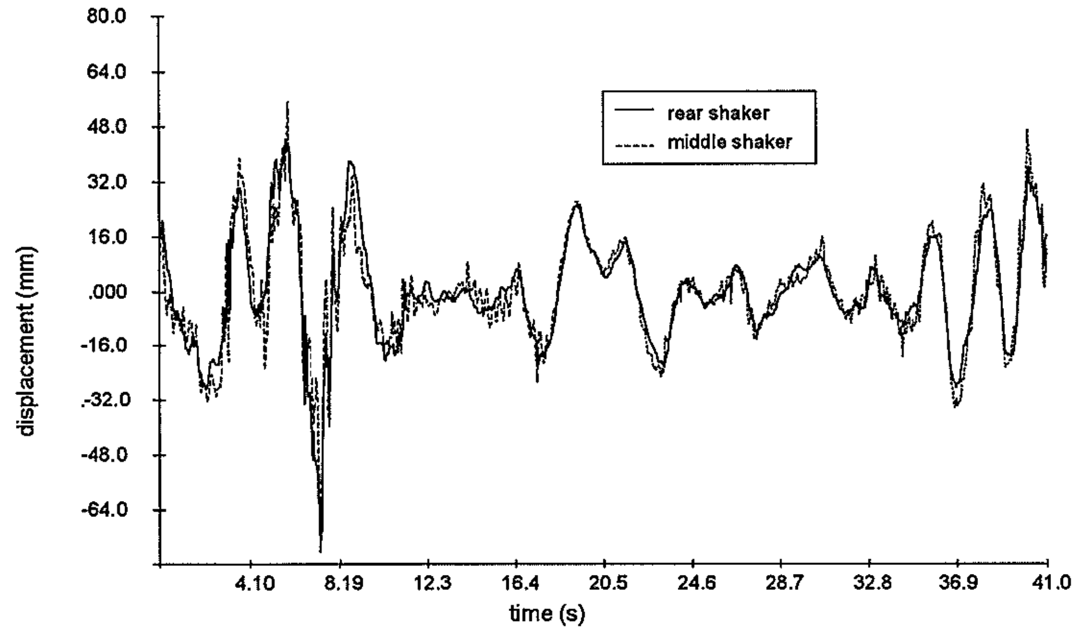

Figure 19. Servohydraulic ram displacements for Mill Rd test.

[1] The power spectral density is calculated using a Fourier transform of the time domain signal. It gives a measure of the "power" of the different frequency components in the signal.

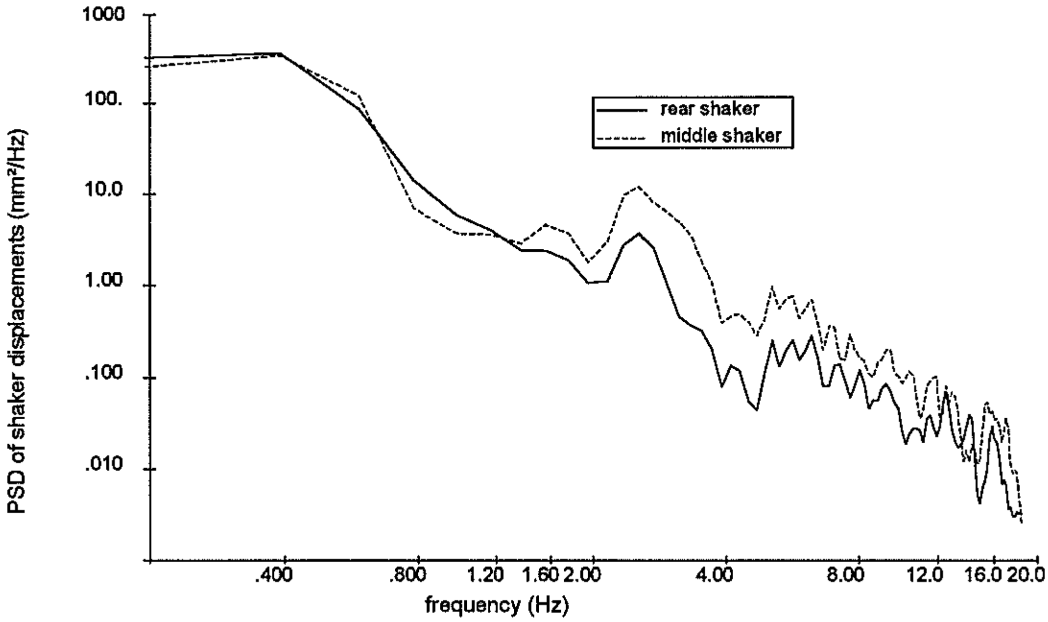

Figure 20. PSD functions of the two shaker displacements.

The other tests gave similar relationships between the two rams with good correspondence at low frequencies and some differences at higher frequencies. However, there is no obvious pattern to these differences. It appears that there may well be a range of combinations of ram displacement signals that will generate the correct response from the vehicle. Retrospectively, this is not unreasonable as the principal sources of dynamic wheel forces are the body modes of the vehicle and these can be excited by either of the rams.

Also worth noting is that there is a "bump" in the curves around the body mode frequencies of the vehicle (3-5 Hz). The PSD functions of the ram displacements have a similar shape to those of road profiles apart from these "bumps". Again, retrospectively this is not surprising. During the road test all six wheels of the vehicle excite the body modes of the vehicle, while in the laboratory tests similar excitation must be achieved using only two wheels. Thus it is necessary to supply additional energy at those frequencies.

To compare these PSD functions directly with the road profile PSD it is necessary to convert between spatial frequency and temporal frequency. By using the fifth wheel data from the road tests it is possible, for each road test run, to convert the road profile information, which has the form of vertical displacement versus horizontal distance, to a form of vertical displacement versus time. These data can then be compared with the shaker displacements during the laboratory tests. As the two rams were not undergoing the same displacements during a test, the choice of ram to compare with the road profile is arbitrary. Figure 21 compares the PSD of the road profile with that of one of the rams for comparable tests.

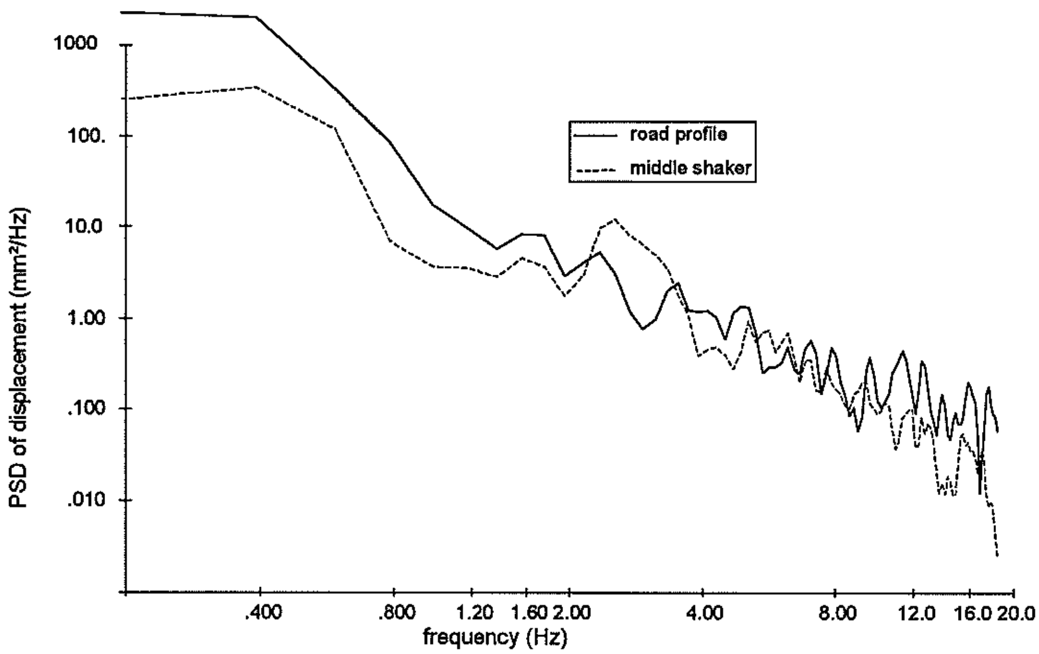

Figure 21. PSD of road profile and shaker displacements.

The shape of the two curves are similar with the road profile having more low frequency response and the ram having the "bump" at 3-5 Hz. Comparing the two signals in the time domain does not show our understanding of their relationship because they are both dominated by low frequency components, which have only limited influence on the vehicle response.

3.6 Modifications to the Shaker Control Software

As described in section 3.4 it was possible to obtain a reasonably good match in the wheel force response between the measured on-road behaviour and the laboratory shaker tests. In fact, compared to the later reported findings of other researchers (Gyenes and Mitchell 1996) with much larger scale and more complex shaker facilities, who used road profile data as the source of the input signals, the match was extremely good.

In the previous section several problems with the original shaker control algorithm were identified. These were overcome with crude fixes, which allowed the method to be used. These fixes required significant operator intervention with the effect that it took a relatively long time to achieve a successful laboratory test. As the primary objective of this testing programme was to develop an efficient and economical method for assessing suspensions in the laboratory, there was clearly a need to improve the algorithm and overcome these deficiencies.

The most significant problem with the original algorithm arose because of the lack of statistical independence between the excitation signals. These were obviously related to the on-road excitations, which, apart from a time lag, were identical for the two wheels being considered. In physical terms, the body motions of the vehicle could be excited by either or both of the two shakers. As there is no unique solution, at some frequencies solutions will occur that require very little input from one or other of the shakers to achieve the output. This gives a value near zero in the transfer function, which becomes a "spike". The manual adjustment of the inverse transfer functions outlined in section 3.4 consists mainly of removing these "spikes". At the start of the iteration process the excitation signals could be selected to be completely independent and there was no problem. As the iteration process proceeded and the excitation signals approached the required solution, they became less and less independent of each other and consequently the problem of "spikes" became worse.

It is also possible that some of the difficulties in getting the servohydraulic system and software to find a solution for shaker displacement signals that generated the target vehicle response were caused by there being multiple solutions. Although it is not totally obvious from Figure 19, analyses of the two shaker displacement signals from each test show that they are quite well correlated but with a time delay between them. This time lag is approximately equal to the axle spacing divided by the vehicle speed for the particular road test. Thus it seems possible that a single driving signal could be used for the rams with an appropriate time delay between the two rams.

The SYSCOMP software was modified to implement this feature as an option. With two target signals and effectively only one excitation the transfer function matrix contains only two terms rather than four. This simplifies the problem considerably. The inverse transfer function matrix has only two terms and the singularity problems are eliminated. However, in attempting to match two targets with only one excitation the solution will be a compromise, but it was hoped that because of the way the target signals were generated the solution would be close. The procedure now converged relatively quickly, taking less than an hour. The match between actual response and target response at this point was reasonably good as illustrated by the examples shown in Figure 22 and 23. The match for the forward wheel of the tandem set was consistently better than that for the rear wheel. This is somewhat surprising as the algorithm is based on finding a compromise solution, which provides the best fit to both targets. No explanation has been found for this phenomenon.

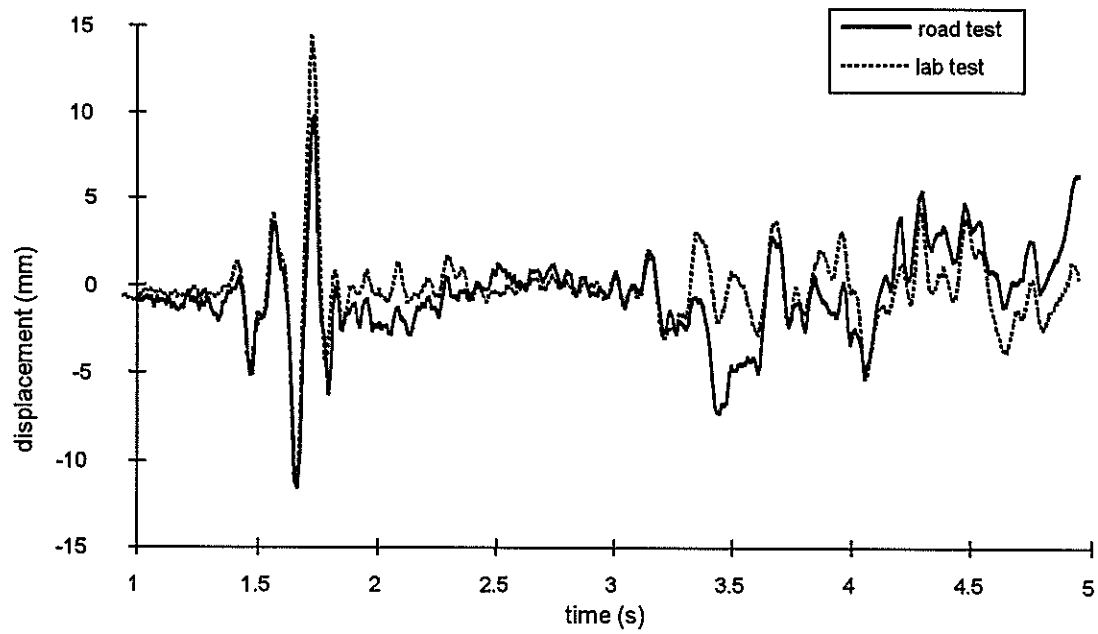

Figure 22. Rear wheel suspension displacements for single signal method

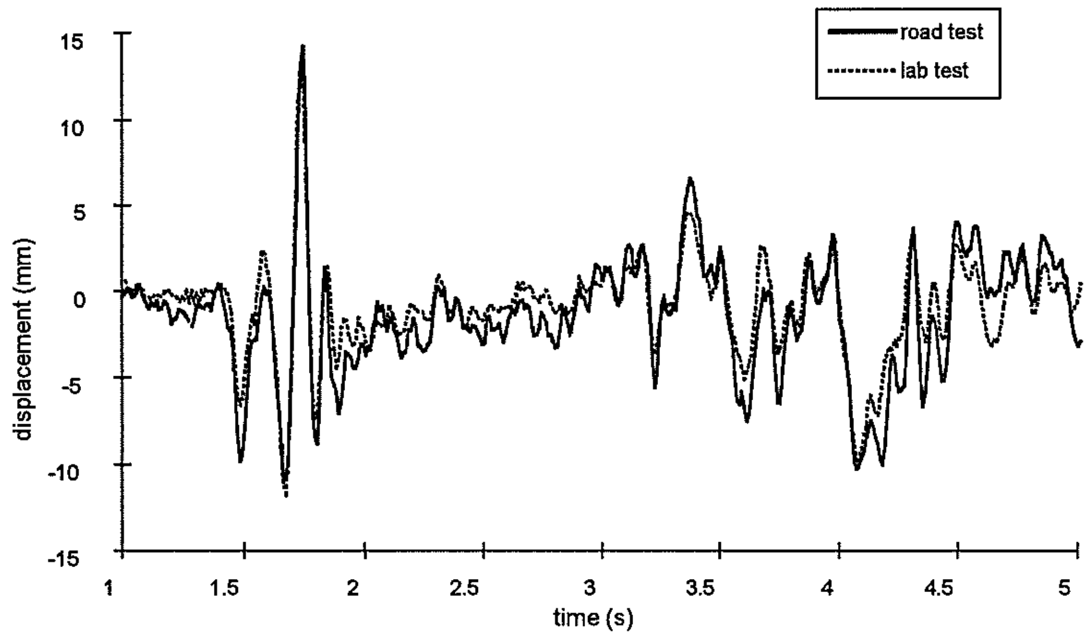

Figure 23. Middle wheel suspension displacements for single signal method.

Although the approach of using a single excitation signal solved the problem of singularities in the transfer functions and so improved the speed of the process significantly, this resulted in an inferior solution. Therefore a new algorithm with a different approach was developed. Rather than calculating new inverse transfer functions at each iteration and using these to calculate the new excitation signals, excitation corrections are calculated using the differences between the actual responses and the targets and an inverse transfer function, which was determined at the beginning of the test. Adding the corrections to the previous excitations then generates the new excitations. As the inverse transfer functions are calculated only once at the start of the process, the excitation signals used for this can be designed to be independent and of sufficient length to ensure a good estimate. With the previous two algorithms the length of the target signals determines the length of the signals used for the transfer function calculation, which may be less than ideal. This new approach worked very well. Using only one excitation signal with a time shift as with the previous algorithm achieved a convergence slightly more quickly with the final match between target and response being of similar quality. With two independent excitations this procedure still achieved convergence in less than an hour with the match between response and target clearly superior to those previously obtained. Figure 24 and 25 show the results for the same road test data as used in the previous figures.

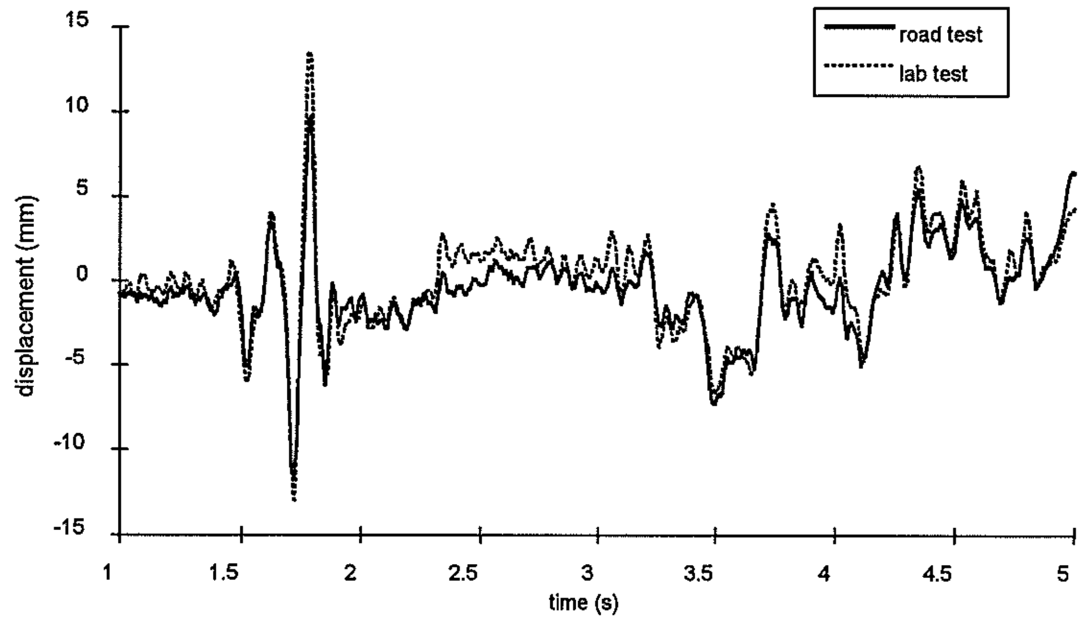

Figure 24. Rear wheel suspension displacements for improved two signal method.

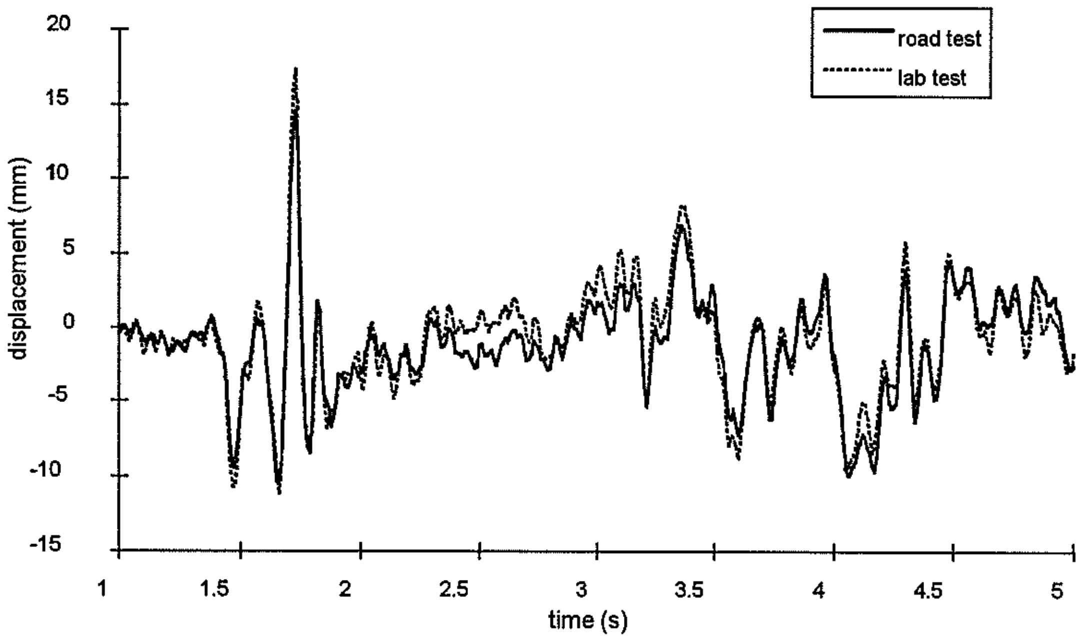

Figure 25. Middle wheel suspension displacements for improved two signal method.

3.7 Conclusions

A series of shaker trials were undertaken to attempt to replicate the measured on-road wheel forces in the laboratory. A computer-based control algorithm was developed to calculate the shaker excitations required to produce the same suspension displacements as were measured on the road. The hypothesis was that, with some minor corrections, the wheel forces measured in the laboratory would be the same as those measured during the road trials and could be used for rating the suspension system. As measuring the wheel forces in the laboratory is much more straightforward than on the road this could form the basis of a cost-effective suspension assessment procedure.

In its original form the algorithm produced satisfactory results with quite a good match in suspension displacements and a reasonable match in wheel forces. However, it had two types of numerical problems. The first was a lack of convergence, which was overcome using a relaxation factor. The second was singularities in the transfer functions which caused problems when inverting. This was solved by manually editing the inverse transfer functions to adjust the problem values. The effect of this was that the process was too slow to be viable as a low-cost test.

The first stage of improving the algorithm was to modify it to allow the use of a single excitation signal to both shakers with a time lag separation between them. This solved the singularity problem and improved the speed of the process significantly but with a small loss of quality in the solution. The second stage involved developing a completely new algorithm, which applied piecewise corrections to the excitation signals. This proved to be fast and stable and produced good results.

4 ROAD TRIALS WITH AIR SUSPENSION

4.1 Introduction

The testing described in the previous two chapters all relates to a particular vehicle fitted with a four-spring steel leaf suspension. To ensure that these results were not specific to the suspension type used the same vehicle was modified and fitted with an air suspension with viscous damping for repeat trials. In general, based on the measurements of other researchers, the steel suspension is considered rather "road-unfriendly" while the air suspension is regarded as "road-friendly".

As well throughout the previous tests, the vehicle was loaded with water. This provides a relatively uniform load distribution and a particular relationship between the body bounce and body pitch resonances. To determine whether the load distribution had any influence on the vehicle responses, the vehicle was also loaded with lead ingot loaded as much as possible in the ends of the tank to give the maximum possible pitch inertia.

4.2 The Test Sites

The test sites and speeds used were the same as those used for the steel suspension tests. As there was a significant time interval between the two series of tests, the road roughness values of the test sites were measured using a NAASRA roughness meter to assess the extent to which the road profiles had changed. A comparison of the results of these measurements with those in the first test is summarised in Table 8 below.

Table 8. Road roughness of the test sites in NAASRA counts/km.

Site	**Measurements from steel suspension tests**		**Measurements from air suspension tests**	
	Average Roughness	**Range of Roughness for 200m sections**	**Average Roughness**	**Range of Roughness for 200m sections**
Church St	92	71 – 112	120	97 - 162
Motorway	32	27 - 37	32	29 - 37
Airfield Rd	98	68 - 140	101	65 - 142
Mill Rd	87	45 - 190	71	41 - 113
Takanini-Clevedon	77	48 - 95	69	49 - 85

From these data we see that, in general, the profiles had not changed greatly. The most significant changes were at Church St where the average roughness appears to have increased noticeably and at Mill Rd where the roughness has decreased. The situation at Mill Rd is readily explained. The site includes a short span bridge, which originally had quite a rough approach. Between the two sets of measurements maintenance work was undertaken and this bridge approach was repaired. This had a significant effect both in reducing the average roughness and the peak roughness, which occurred at this location. The change at Church St is more difficult to explain. The roughness values for the steel suspension test were calculated from the actual measured profiles using the IRI quarter car model which traverses the section at 80 km/h. The values form the air suspension test were measured using a NAASRA roughness meter, which, at the Church St site, was operating at 50 km/h. At all the other sites the NAASRA meter was run at 80 km/h. Operating the IRI quarter car model at lower speeds does increase the calculated roughness value. Therefore it is possible that the apparent increase in roughness at Church St is not due to profile changes but to an inadequate compensation for the lower operating speed of the NAASRA meter vehicle on urban roads.

4.3 The Test Program

As mentioned previously, the same vehicle that was used in the previous tests, a three-axle liquid tanker trailer was modified and fitted with air suspension. As with the steel suspension tests the vehicle was extensively instrumented. This included shear strain gauges to measure vertical forces, axle mounted accelerometers to measure unsprung mass accelerations, chassis mounted accelerometers to measure sprung mass accelerations, LVDTs to measure suspension deflections and a "fifth" wheel to measure vehicle speed.

For the first series of tests the vehicle was filled with water and road tests were conducted at the same five test sites each at three speeds. Because the towing vehicle used for these tests had less power than the vehicle used for the steel suspension tests, the highest speeds were not attainable at some of the test sites. Therefore there are some minor differences between the speeds used for these tests and those used previously. However, at each site there were common speeds for all test configurations. After the tests with water loading, the trailer was loaded with lead ingots through the inspection hatches at either end of the tank. The lead was stacked as close as possible to the ends of the tank to maximise the pitch inertia. Each end of the tank was loaded to give approximately the same axle loads as the water loading. The complete set of tests was repeated with the lead loading.

4.4 Results

The axle loads for these tests were slightly higher than for the previous tests with the steel suspension. This is primarily because the suspension is heavier leading to a higher tare weight. In both cases where water loading was used the tanker was filled. Table 9 shows the measured static axle loads for both the water and lead loading cases.

Table 9. Static axle loads.

Axle Position	Weight (kg) with water load	Weight (kg) with lead load
Rear	6360	6520
Middle	6020	6320
Front	6800	6280
Total	19180	19120

Considering first the case where the tanker was loaded with water, we find that, surprisingly, the DLC values obtained from these tests were quite similar in magnitude to those obtained previously with the same vehicle on steel springs. A comparison using one set of tests from each site is shown in Table 10. In general, air suspensions are considered "road-friendly" and steel suspensions "road-unfriendly". Therefore, we would have expected the DLC values of the air suspension to be significantly lower than those of the steel. This is clearly not the case. If anything, the air suspension appears to generate slightly higher dynamic loads. The reason for this is the inadequate damping of the air suspension. This is discussed in more detail later in this section.

Table 10. A comparison of DLC values between steel and air suspension.

Site	Suspension type	Rear wheel	Middle wheel	Front wheel
Church St	Steel	0.135	0.145	0.127
Church St	Air	0.147	0.136	0.140
Motorway	Steel	0.073	0.079	0.104
Motorway	Air	0.106	0.093	0.097
Airfield Rd	Steel	0.134	0.120	0.113
Airfield Rd	Air	0.221	0.144	0.168
Mill Rd	Steel	0.157	0.130	0.143
Mill Rd	Air	0.206	0.159	0.166
Takanini-Clevedon Rd	Steel	0.127	0.127	0.123
Takanini-Clevedon Rd	Air	0.125	0.135	0.157

Similarly, we can compare the dynamic wheel forces generated by the vehicle with water loading to those with lead loading under the same test conditions. These are presented in Table 11.

Table 11. Comparison of DLC values for lead and water loading.

Site	Loading type	Rear wheel	Middle wheel	Front wheel
Church St	Water	0.147	0.136	0.140
Church St	Lead	0.128	0.065	0.127
Motorway	Water	0.106	0.093	0.097
Motorway	Lead	0.115	0.078	0.118
Airfield Rd	Water	0.198	0.137	0.149
Airfield Rd	Lead	0.179	0.122	0.165
Mill Rd	Water	0.206	0.159	0.166
Mill Rd	Lead	0.160	0.123	0.158
Takanini-Clevedon Rd	Water	0.125	0.135	0.157
Takanini-Clevedon Rd	Lead	0.137	0.114	0.126

There does not appear to be a significant difference in the dynamic wheel force generated by the two different forms of loading. It appears that the lead loading case may have slightly lower wheel forces possibly because of the elimination of the liquid slosh mode. However, this change is very small.

Figure 26 and 27 show the amplitude spectra of the wheel forces for one of the road tests (Mill Rd at 75 km/h) with the vehicle loaded with lead and water respectively. The mode at approximately 0.5 Hz is clearly visible in the case with water loading and not in the case of the lead loading. This supports the hypothesis presented in section 2.5 that this is a liquid sloshing mode.

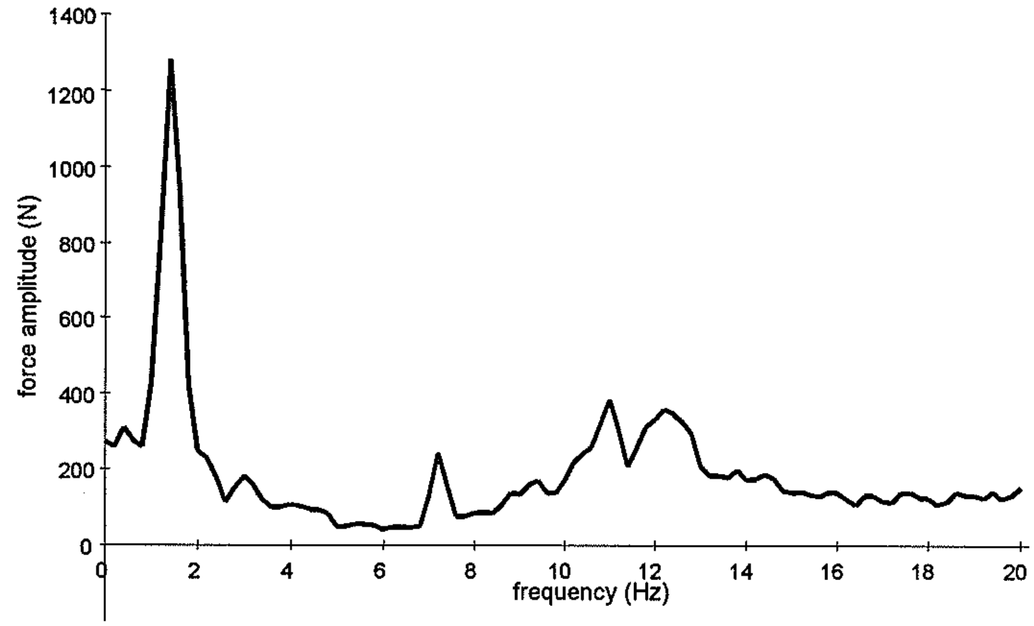

Figure 26. Example force amplitude spectra for lead loading.

There is an unexpected resonance at approximately 7.5 Hz. This frequency corresponds to the wheel rotation frequency at the test speed and so this resonance is probably caused by an imbalance in the wheel. When the spectra for test runs at different speeds are compared the frequency of the resonance changes as expected.

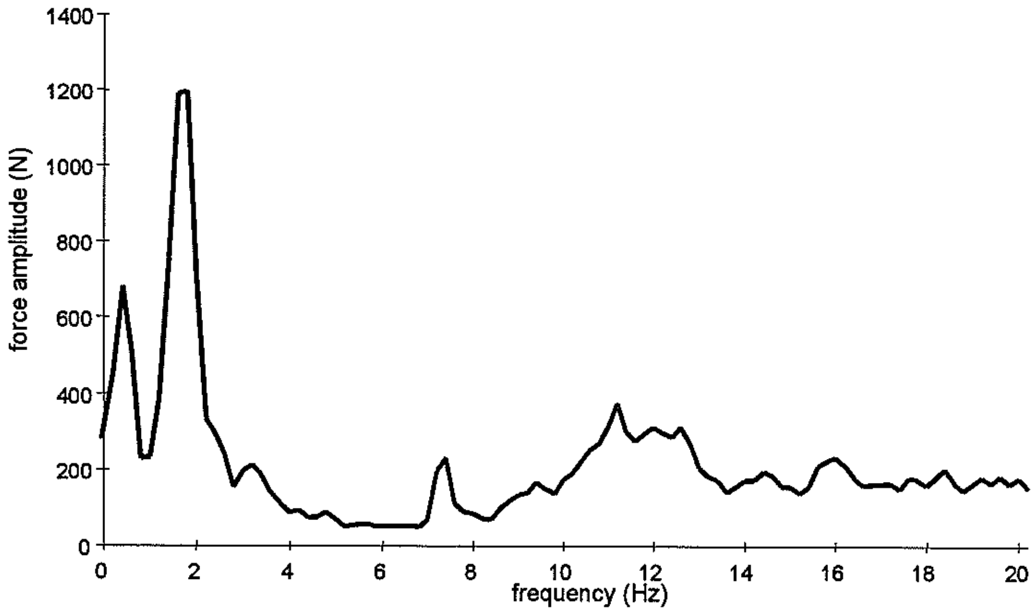

Figure 27. Example force amplitude spectra for water loading.

Comparing the spectra in these two figures with that in Figure 7 for the steel suspension case leads us to a number of observations. The natural frequencies of the dominant wheel force modes were substantially lower than with the steel springs (1.8 Hz compared to 2.5 - 3.5 Hz) and with more clearly defined modes as would expected from a softer more linear spring. However, in spite of this, the air suspension is not more "road-friendly" than the steel previously used. This appears to be because of inadequate damping in the air suspension. The sprung mass resonance represented by the peak at 1.8 Hz in the spectra is narrow and high indicating poor damping. A modal analysis of the vehicle gave damping values for the vehicle bounce and pitch modes of 5-6 %, which are very low. Further investigations have found that the suspension configuration and fitting were in accordance with the manufacturer's specifications and that the dampers were functioning correctly on a test machine. However, the damper-mounting angle used in this suspension results in relatively small damper motions and low damping forces. These measurements illustrate quite clearly the flaws in a design approval approach to rating road-friendliness. This suspension would qualify as "road-friendly" under a policy of blanket approval of air suspensions (as is the case in the EC). However, it offers no real benefit to the pavement over the steel suspension, which it replaced. It is worth noting that this particular suspension is widely used in New Zealand.

It can also be seen that, for the air suspension, the unsprung mass modes also occur at lower frequencies than for the steel (around 12 Hz instead of 15 Hz). The magnitude of these is similar.

4.5 Conclusions

The road tests, which were previously undertaken for the vehicle fitted with steel suspension, were repeated for the same vehicle, which had been fitted with air suspension. In addition to running the tests with the vehicle loaded with water, a complete set of tests was undertaken with the vehicle loaded with lead ingots distributed to give markedly different pitch inertia.

Rather surprisingly, this air suspension did not generate lower levels of dynamic loading than the steel suspension and thus cannot be rated as more "road-friendly". Although the suspension was softer, it was inadequately damped. This highlights the critical importance of damping on the performance of softer suspensions and shows the weaknesses of a design-based approach to "road-friendliness" rating. The poor damping performance of this suspension was completely attributable to poor design. All the components of the suspension were in good condition and performing to specification.

The lead loading had minimal effect on the level of dynamic wheel forces generated by the vehicle. A resonance mode at about 0.5 Hz which was present in all the tests with water loading was absent in the test with lead loading. It had previously been hypothesised that this mode was generated by a liquid sloshing phenomenon. Its absence from the test with lead loading supports this theory.

5 SHAKER TRIALS WITH AIR SUSPENSION

5.1 Introduction

As with the steel suspension tests the vehicle was mounted on the servo-hydraulic shakers facility. The SYSCOMP software system was used to iteratively determine the excitation signals required to produce the same suspension deflection responses in the laboratory as had previously been recorded during the road tests. The purpose of these trials was to demonstrate that the procedure for assessing suspensions developed with the steel suspension vehicle was also valid for an air suspension.

In the first series of tests this was found not to be the case and modifications to the test rig and the procedure were necessary. Once these had been made further shaker tests were carried out.

5.2 The First Series of Shaker Tests

In the first set of tests the vehicle was mounted with the two wheels on the left-hand side of the rear tandem axle group mounted on the two actuators. The iterative process of determining the excitation signals required to obtain the same suspension displacements at the two wheels was commenced. The first set of road data used was from the Mill Rd site with a vehicle speed of 80 km/h. Although this site is of only moderate roughness, the vehicle motions obtained during this process became very large and for safety reasons, it was not possible to complete this test.

Shaker trials had been successfully completed for this test site at this speed when the steel suspension was fitted. As the on-road measured wheel forces for the vehicle when fitted with air suspension were of similar magnitude, the method was clearly not working correctly for the air suspended vehicle.

To see if it was possible to find a solution for any site, road test data from the smoothest site, the Southern Motorway, were used. With these, a successful match between the road and laboratory suspension displacements was achieved. Figure 28 shows an example of this match for the left rear wheel from an 85 km/h test run. The match obtained is better than previously achieved on the steel springs. This is probably due to having larger suspension displacements (better signal-to-noise ratio) and a more linear spring response. However, when the wheel forces were compared, it was found that the wheel forces recorded in the laboratory were 4-5 times greater than those measured on the road. Figure 29 shows the wheel forces corresponding to the Figure 28 displacements. Figure 30 shows a comparison of the wheel force power spectral density functions for the same wheel. It can be seen that all the frequency components that were measured on the road are also present in the shaker test measurements but at greater magnitudes.

This seriously undermines the whole concept behind this proposed assessment procedure. The assumption was that if the suspension deflections were the same, the reaction forces they generate would be the same and so, after correcting for differences in the unsprung mass inertia forces, the wheel forces would be the same. This was clearly not so. A substantial contribution to wheel forces was coming from some other source

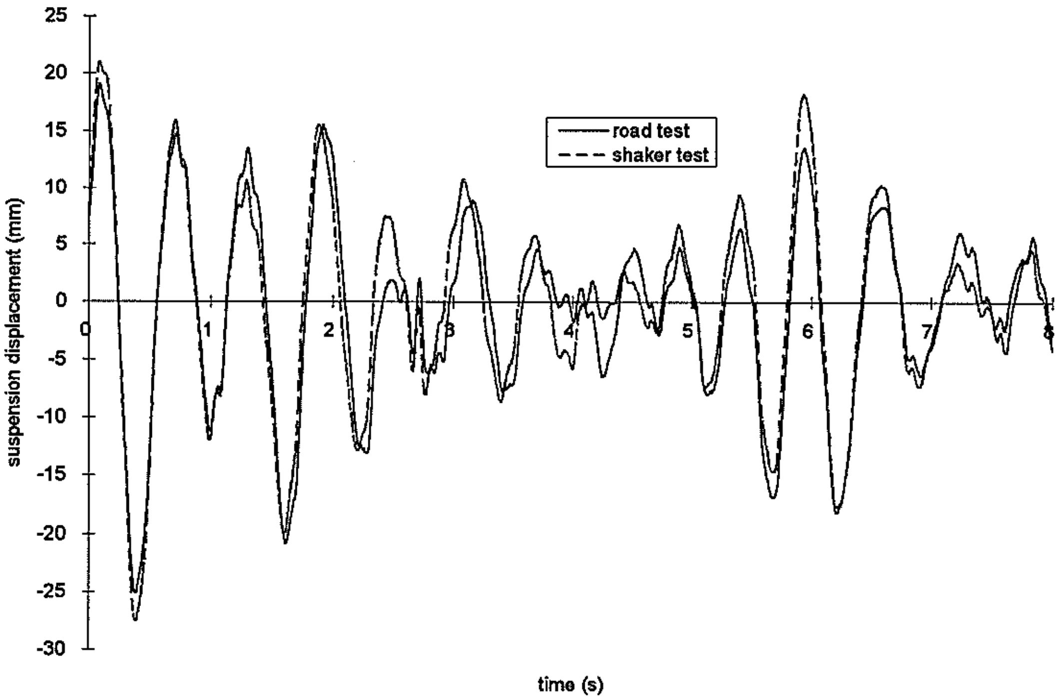

Figure 28. Comparison of suspension displacements between road and shaker tests.

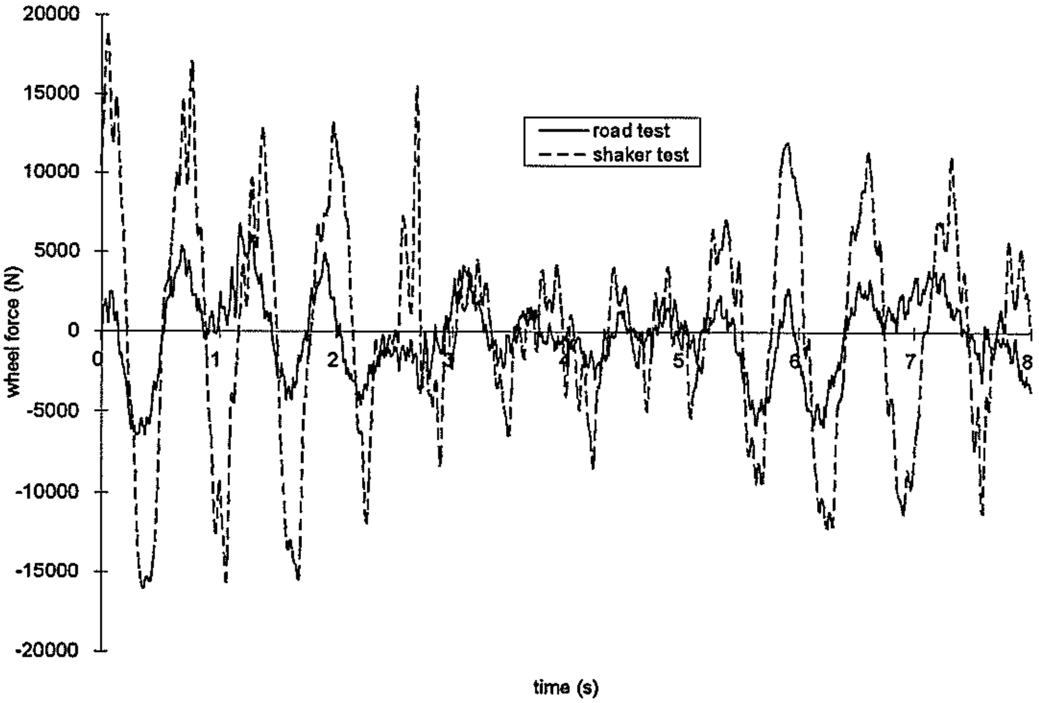

Figure 29. Comparison of wheel forces between road and shaker tests.

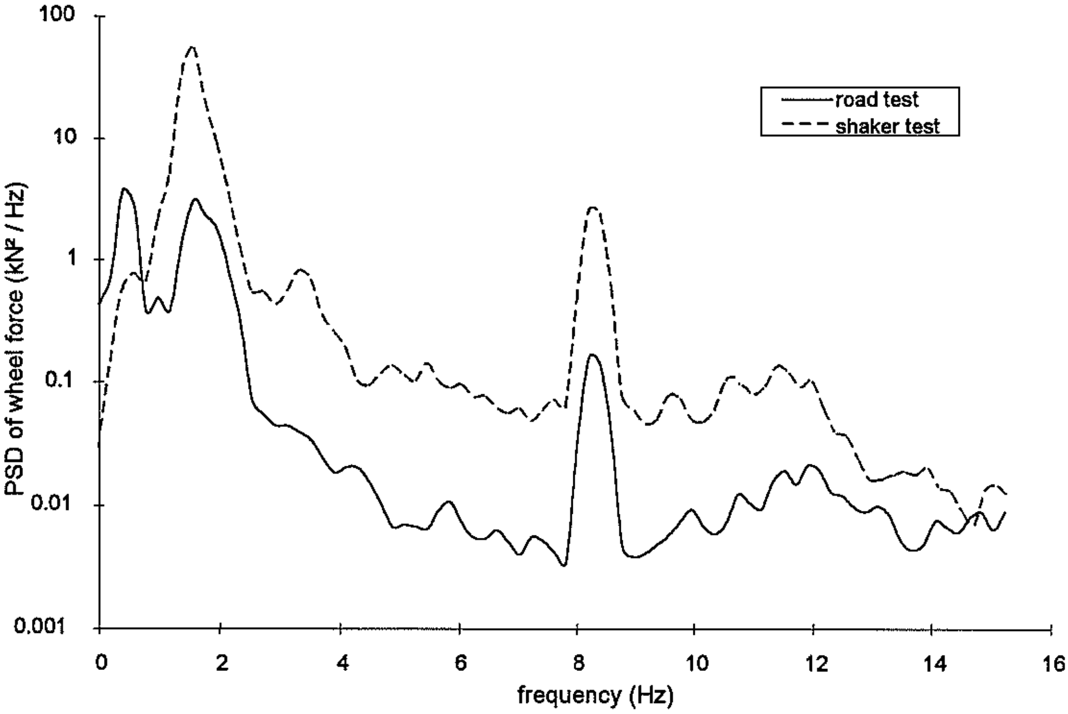

Figure 30. Comparison of wheel force PSD between road and shaker tests.

Considering the differences between the road and laboratory vehicle responses, it was clear that the vehicle was being subjected to much greater roll motion in the laboratory. From (Fancher et al. 1986) figures were obtained for typical values of suspension vertical stiffness and composite roll stiffness. Typical values for vertical stiffness were 1,400-2,300 kN/m per axle for steel suspensions and 175-1,220 kN/m per axle for air suspensions. Based on a spring base of 1.0 m, these suspension units would provide a roll stiffness of 350-575 kN m/radian for steel and 44-305 kN m/radian for air. The composite roll stiffness figures given by Fancher are 422-649 kN m/radian for steel suspensions and 182- 584 kN m/radian for air suspensions. In general, a lower vertical stiffness corresponds to lower roll stiffness for each suspension type. For the steel suspensions, therefore, most of the roll stiffness is provided by the vertical suspension stiffness with the auxiliary roll stiffness 20% or less of that provided by the springs. For the air suspension, the auxiliary roll stiffness is actually greater than the roll stiffness of the springs and may be as much as three times as great. This leads us to the hypothesis that the increased wheel forces measured in the laboratory are generated as reaction forces to the auxiliary roll stiffness which is resisting the additional roll motion generated in the laboratory. These forces do not occur during road tests because the roll motion is not present and they do not cause the suspension to deflect.

5.3 The Second Series of Tests

To test this theory, the vehicle was mounted transversely across the shakers so that both wheels of one axle could be excited. The servo-hydraulic system was set up so that both actuators produced the same excitation response and SYSCOMP software was used to determine the excitation required to match the suspension deflection at one axle (the rearmost). By applying the same excitation to both wheels of an axle no vehicle roll behaviour occurs. However, by only shaking one axle in the tandem set the load transfers between axles are not replicated accurately.

As before, it was possible to accurately replicate the on-road suspension deflections for the wheels under test. A sample of these is shown in Figure 31, which covers the same road test data as Figure 28.

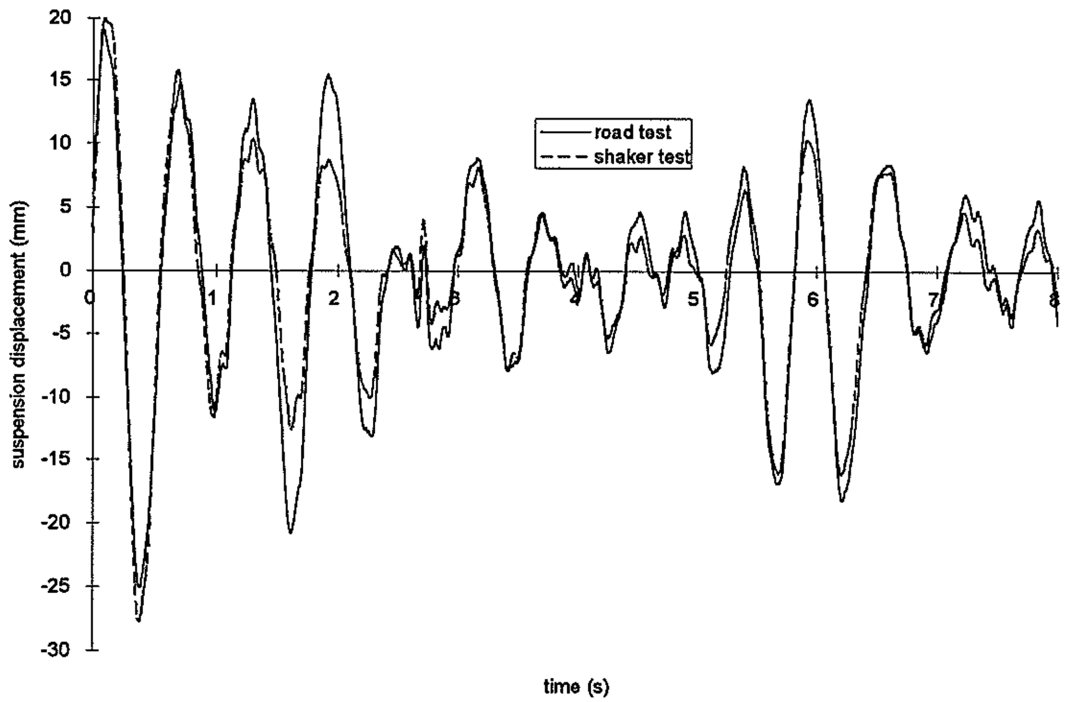

Figure 31. Comparison of suspension deflections between shaker and road tests.

As can be seen from Figure 32 the corresponding wheel force signals are still substantially larger than those measured on the road but only by a factor of 1.5 - 1.7.

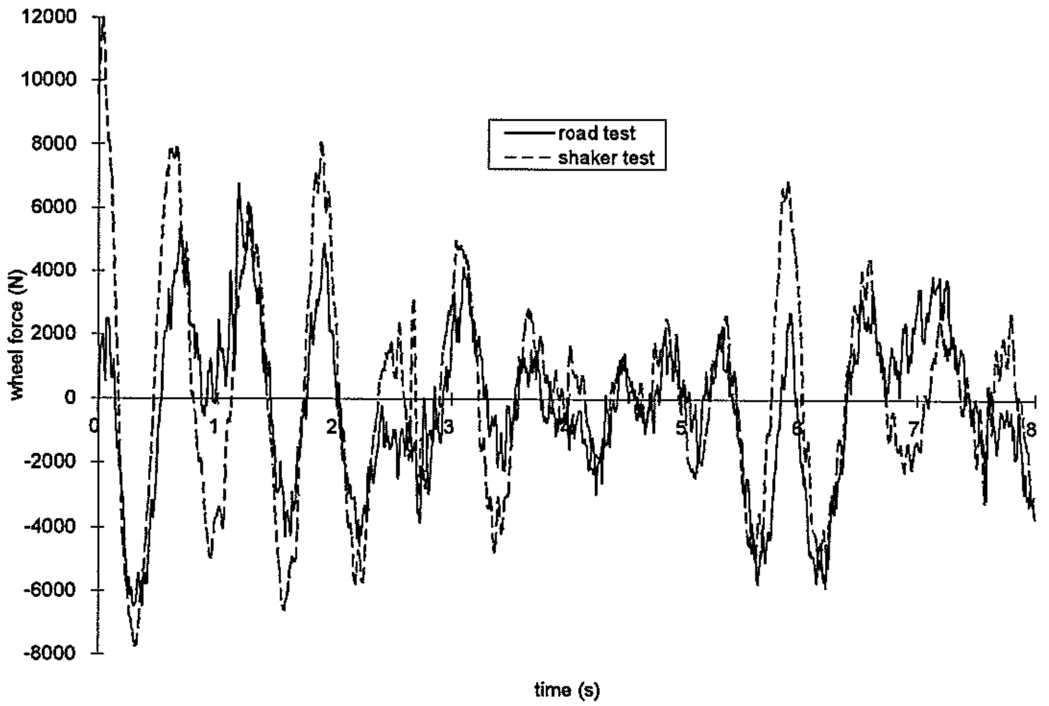

Figure 32. Comparison of wheel forces between road and shaker tests.

Thus, the auxiliary roll stiffness appears to have been a major contributor to the original discrepancy. It is assumed that the remaining difference in the forces was caused because only one axle was excited. This eliminates the load sharing which would occur if both axles were excited. All the input force needed to generate the appropriate sprung mass motion has to be supplied through this one axle.

5.4 Modifications to the Shaker Test Facility

Having demonstrated that the probable source of the problem was the auxiliary roll stiffness of the suspension, a new vehicle support rig as illustrated in Figure 33 was designed and built. With this configuration each actuator excites both wheels of an axle. Thus, both axles of a tandem suspension can be excited but it is not possible to reproduce any vehicle roll. It has been reported (OECD 1992) that, except on very rough roads, vehicle roll is not a major contributor to dynamic loads and hence this approach is acceptable.

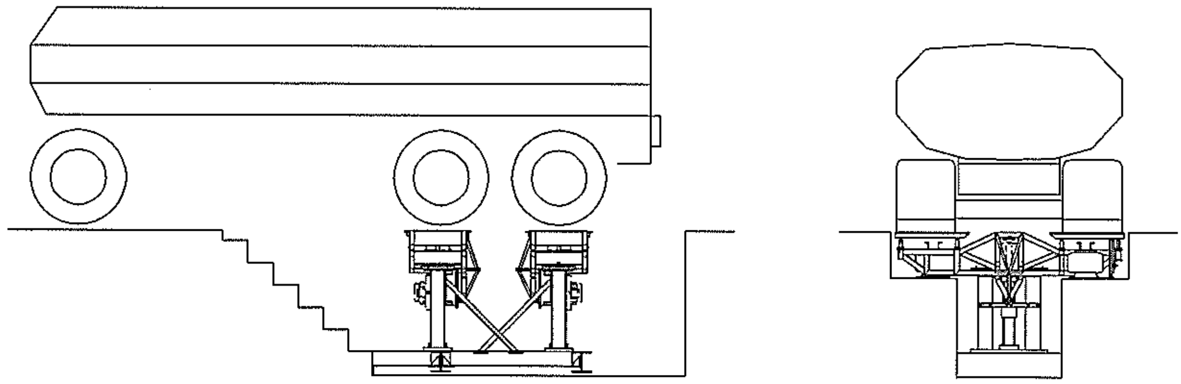

Figure 33. Modified shaker support rig.

5.5 The Third Series of Tests

The shaker tests were repeated on this rig and, as before, an excellent match of suspension displacements was achieved as illustrated in Figure 34. This covers the same road test data as Figure 28 and 31. The corresponding wheel force comparison is shown in Figure 35 and the power spectral densities of the wheel forces are compared in Figure 36. These match well at the fundamental vehicle modes between 1.5 and 2 Hz. However, in the road measurements there is a substantial wheel force component at 0.5 Hz, which is not matched in the laboratory tests. As mentioned previously side-to-side liquid sloshing in the tank generates this. With the new shaker support rig, which provides no roll excitation, this cannot be re-created in the laboratory. In the earlier tests with the steel suspension where only one side of the vehicle was excited, roll behaviour could be induced. However, because this mode involves a moving liquid therefore has a very non-linear response it was still not possible to match this behaviour exactly. Liquid sloshing is peculiar to this type of loading and therefore could be ignored in a general suspension assessment procedure. The road tests with the same vehicle loaded with lead ingots instead of water demonstrated that the liquid causes this motion. To illustrate the match in forces without this 0.5 Hz component, Figure 37 shows the same comparison as Figure 35 with a 1 Hz high pass filter applied to the road measurements.

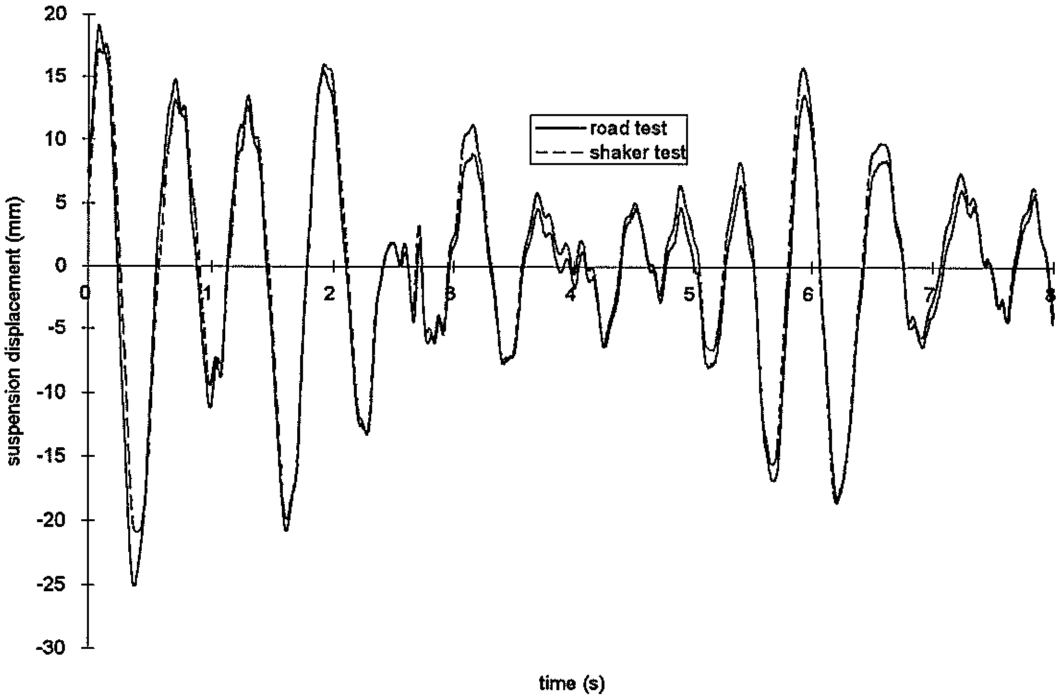

Figure 34. Comparison of suspension deflections for road and shaker tests.

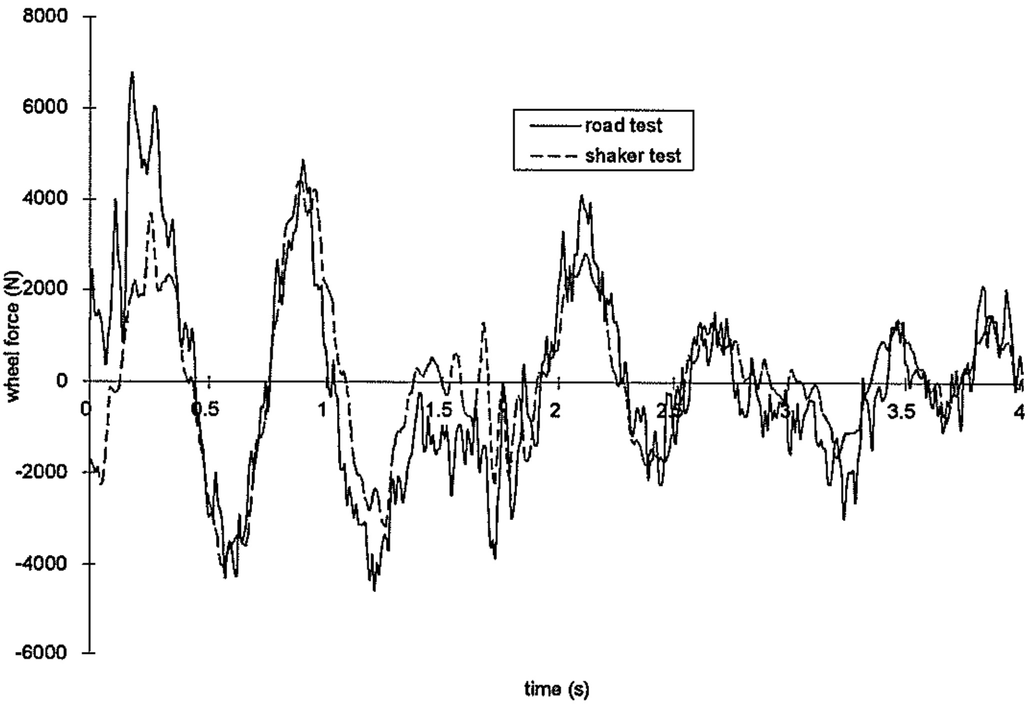

Figure 35. Comparison of wheel forces between road and shaker tests.

The laboratory measured wheel forces also do not exhibit the full amplitude at higher frequencies. It is not clear why this is so. It is possible that differences in dynamic characteristics between a rolling tyre and a stationary one influence this behaviour. Alternatively it may be that the resolution using suspension displacements as the target response results in the detailed force response at the higher frequencies being filtered. The forces depend on accelerations, which have proportionately larger amplitudes at high frequencies than displacements.

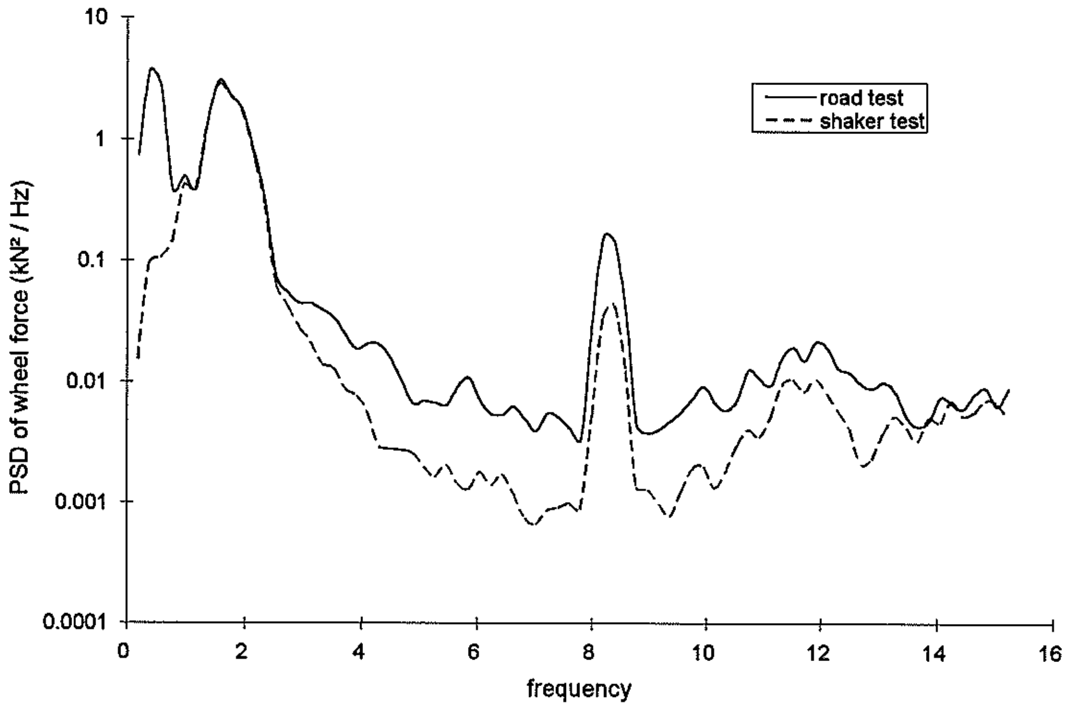

Figure 36. Comparison of wheel force PSD between road and shaker tests.

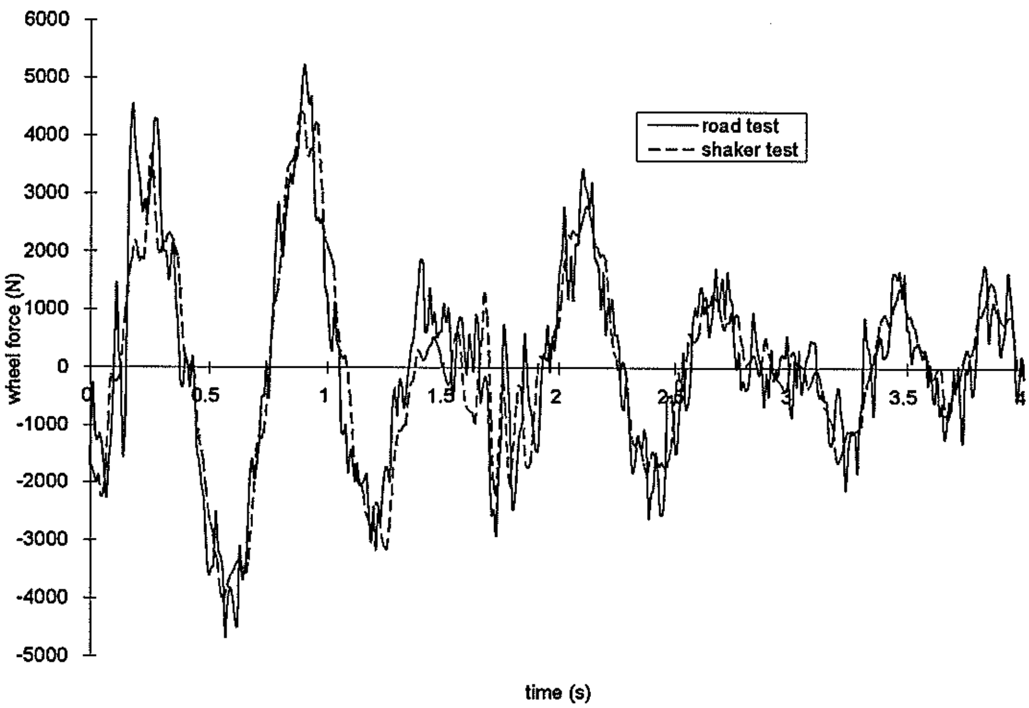

Figure 37. Comparison of wheel forces - high pass filtered at 1 Hz.

The peak in the wheel force spectra at 8 Hz is interesting. It was generated by an imbalance in the wheels. This was checked by calculating the wheel rotation frequency and verified from changes in the frequency of this mode with changes in speed in the various road tests. During the shaker trials there was, of course, no wheel rotation but because the shaker control algorithm was trying to match the measured on-road behaviour it generated appropriate excitations to obtain a response at this frequency.

5.6 Conclusions

The previously developed technique for simulating on-road suspension behaviour using a small-scale servohydraulic facility, which had been validated on a vehicle fitted with steel suspension, was applied to the same vehicle fitted with air suspension. Although the method replicated suspension displacements well, the resulting vehicle motions and wheel forces were excessive. As the aim of this technique is to estimate wheel forces it could not be used satisfactorily without modification.

It was postulated that these high wheel force levels were generated by the vehicle's auxiliary roll stiffness responding to roll excitations by the shakers. This hypothesis is supported by data on typical vehicle-suspension parameters and by shaker tests with no roll excitation.

The vehicle support structure on the servohydraulic shaker facility was changed so that no roll excitations are applied. With this modification the technique was successful in accurately replicating the wheel forces associated with the sprung mass motions of the vehicle which represent the major contribution to dynamic pavement loads.

At higher frequencies corresponding to the vehicle's unsprung mass motions the technique produces the correct responses but at lower amplitudes than during the road tests which are being simulated.

6 SHAKER TRIALS ON THE CAPTIF SLAVE UNITS

6.1 Introduction

During the course of the research, another project to investigate pavement performance under different levels of dynamic loading was initiated at the Canterbury Accelerated Pavement Testing Indoor Facility (CAPTIF). This project consists of a series of accelerated pavement tests where the two Simulated Loading and Vehicle Emulator (SLAVE) units are fitted with different suspensions and run on independent wheel paths on the same pavement from construction through to failure. By monitoring the pavement and the vehicles throughout the tests, the influence of vehicle dynamics on pavement performance can be determined.

The first test in this series was on a typical New Zealand pavement design with a thin surface coating on a structural layer of crushed rock laid over a silty clay subgrade. One of the SLAVE units was fitted with a traditional multi-leaf steel suspension with no other damping while the other was fitted with a more modern parabolic leaf spring and a viscous damper.

The second test in the series was undertaken for the OECD as Element 1 in the DIVINE international co-operative research programme. For this experiment the pavement used consisted of a thicker asphalt layer (85 mm) again over a layer of crushed rock on a silty clay subgrade. For this test one of the SLAVE units was again fitted with the multi-leaf steel suspension while the other was fitted with an air suspension and viscous damper. These suspensions tend to represent the typical extremes in "road-friendliness".

At least one further test is planned but has yet to commence. For this test the suspensions used for the OECD test will be used on a more typical New Zealand pavement design.

As the issue of pavement performance under dynamic loading is very closely interrelated to the issue of assessing suspensions for pavement wear, it was decided to link the two research projects by undertaking a series of shaker tests on the CAPTIF SLAVE units.

6.2 Modifications to the Shaker Test Facility

An elevation view of the CAPTIF facility is shown in Figure 38.

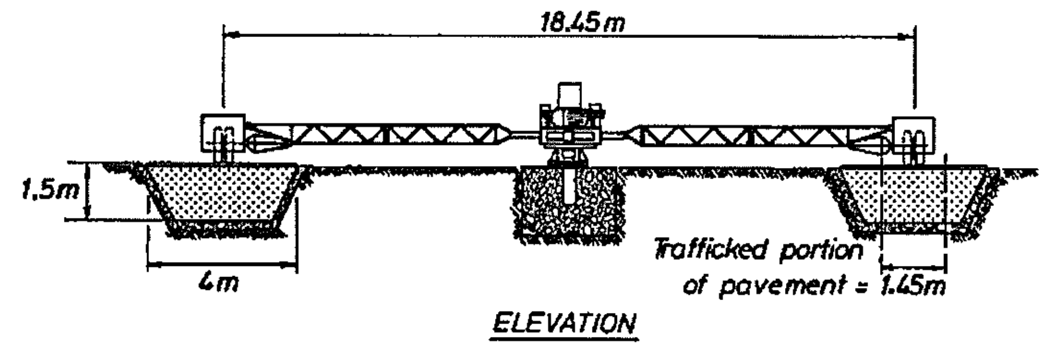

Figure 38. Elevation view of the CAPTIF.

As can be seen, the SLAVE units are attached to a frame mounted on the central pod by space frame arms. The connections between these arms and the central frame are through pinned joints that are free to rotate about a horizontal axis. Thus, the SLAVE units can move freely in the vertical direction. The SLAVE units, illustrated in Figure 39, were made using, as much as possible, standard heavy vehicle components so that the masses and suspension characteristics are similar to real vehicles.

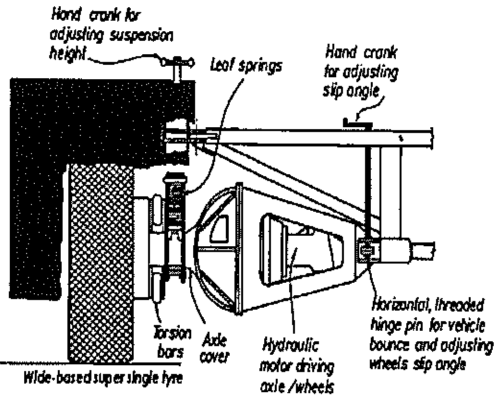

Figure 39. Schematic of CAPTIF SLAVE unit.

For the shaker trials one of the SLAVE units and its arm were removed from the CAPTIF and shipped to the shaker facility in Auckland along with suspension options, loading plates and both wide-based single and dual tyres.

The shaker support rig was in the modified configuration previously illustrated in Figure 33. As the SLAVE unit consists has a single wheel position rather than a complete axle it was not possible to mount on the shaker rig as it was. However, two large steel channel sections were used to span the two wheel supports of one of the shaker rigs so that the SLAVE axle could be positioned directly above the servo-hydraulic actuator. These channel sections were welded to the wheel support plate to ensure a rigid connection.

A rigid frame with pin joints at the correct height, similar to those on the central pod at CAPTIF, was constructed. This was rigidly bolted to a strong floor at the appropriate distance from the shaker rig. The SLAVE arm was then connected to this frame so that, apart from not moving horizontally, the geometry and configuration was as similar as possible to that at CAPTIF.

6.3 Instrumentation and Measurement Systems

The SLAVE units were instrumented with a chassis mounted and an axle mounted accelerometer and an LVDT. The accelerometers measure the sprung and unsprung mass accelerations respectively. It was assumed that because of the simplicity of the SLAVE units these two accelerations could be combined to calculate the wheel forces. In its simplest form this relationship is:

$$\text{Force} = M_{sprung} . a_{sprung} + m_{unsprung} . a_{unsprung}$$

Because of the geometry of the unit the actual relationship is more complex than this. The true motions of the components of the SLAVE units are not simple vertical movements but rotations about the pivot points. Taking this into account it is possible using simple mechanics to derive a relationship between the two accelerometer signals and the wheel forces. This has the same basic form as given in the previous except that the mass terms are not the actual masses of these components but mass factors which depend not only on the masses but also their position and distribution. The LVDTs were mounted between the chassis and the axle of the SLAVE units to measure the suspension deflections.

The platforms on top of the shaker rig were fitted with load cells that measure the wheel forces applied during testing. All the instrumentation was connected to a Hewlett-Packard HP3852 data acquisition system. Sampling rates were set to match those used during the DIVINE test at CAPTIF.

6.4 The Test Program

The excitation signals used for this test were selected from the road profiles measured at the CAPTIF during the OECD DIVINE experiment. To cover the complete range of roughness values experienced during the test, profiles from 20,000, 500,000 and 1,600,000 load cycles, and from both the inner and outer wheel paths, were chosen. These loading intervals correspond to the start, the middle and the end of the test. (The first 20,000 load cycles were taken up with pavement conditioning laps and zero measurements and thus the test proper commenced at 20,000 load cycles).

For the tests the SLAVE unit was in turn fitted with the multi-leaf steel and the air suspension and a wide based single tyre. With these two suspension-tyre combinations, the unit was tested using each of the six road profile excitations and a band limited noise excitation at three different amplitudes, which were used to investigate suspension non-linearity. With both suspensions, dual tyres were also fitted and the three noise excitations and two of the road profile excitations were applied.

Although this testing programme should have been relatively straightforward, there were some unexpected results, which led to quite extensive re-testing to try to identify their source. This is discussed in the next section.

6.5 Results

The SYSCOMP software was used to match the actual response to the target excitation signal. In this case the target and response signals related to the displacement of the shakers themselves rather than responses measured somewhere on the vehicle. This is a considerably simpler situation than the vehicle tests undertaken previously. It was found that the transfer function between excitation and response for this situation was approximately linear with magnitude decreasing with increasing frequency. Applying the inverse of this transfer function to the target signals resulted in excitations that produced a very good match to the target response. No iteration was necessary to improve the match.

With the instrumentation described, the wheel forces can be determined in two independent ways. The first is by calculation from the accelerometer signals, which is the method used during testing at CAPTIF, while the second is by direct measurement from the load cells on top of the shaker rig.

The first series of tests were with the air suspension, the wide based single tyre and the measured road profile from CAPTIF at the start of the DIVINE experiment. The suspension deflections on the shaker compared very well with those previously measured at CAPTIF as shown in Figure 40. It should be noted that these suspension deflections were not the target response as in the previous shaker testing. The excitation was generated directly from the road profile.

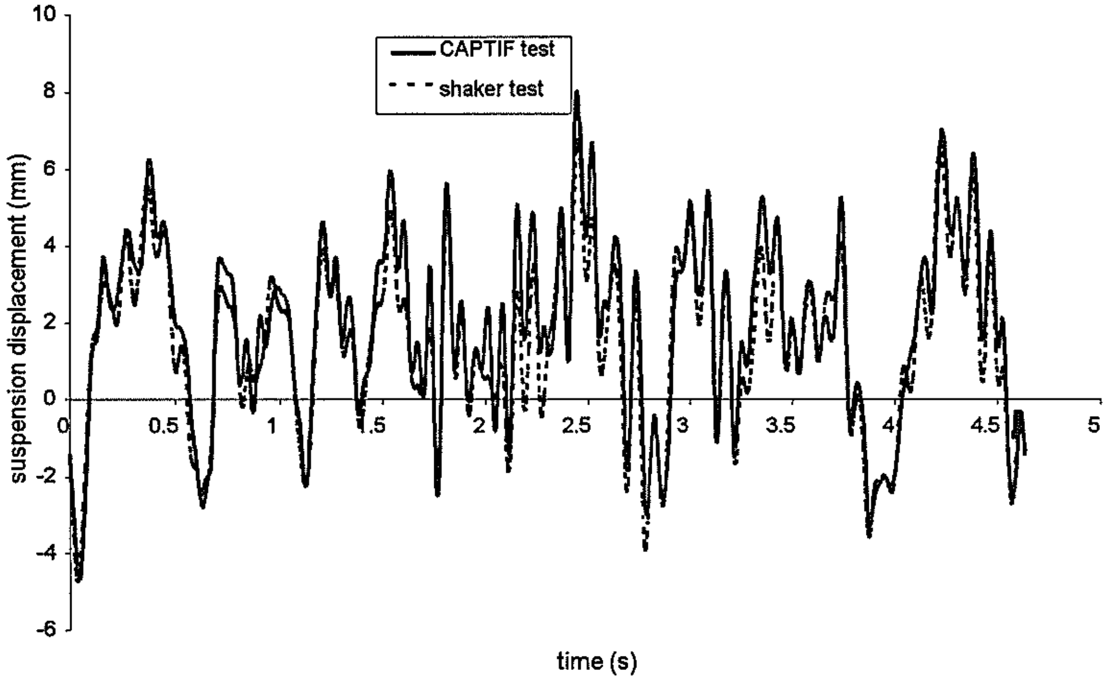

Figure 40. Comparison of suspension deflections between CAPTIF and shaker tests.

However, the chassis accelerations compared very poorly as shown in Figure 41. The match at lower frequencies is reasonable but at higher frequencies the amplitudes on the shaker were much higher. However, apart from this the signals appear to follow each other quite well.

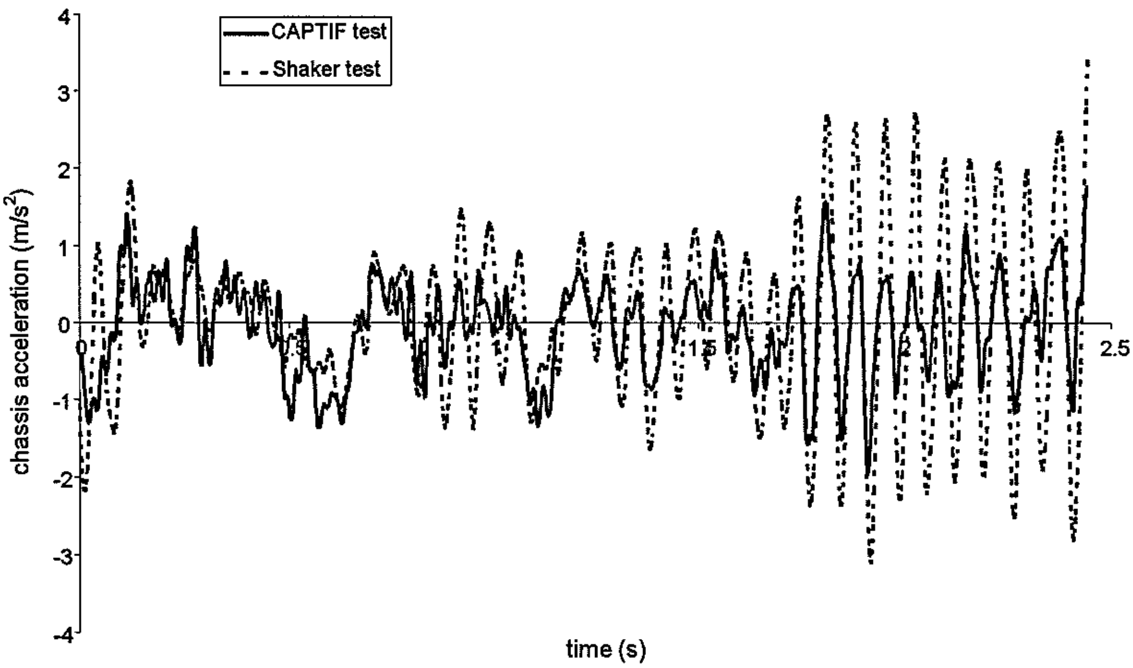

Figure 41. Comparison of chassis accelerations between CAPTIF and shaker tests.

This situation was baffling. The problem was clearly not one of noise because the signals tracked each other quite well. It was also not simply a matter of a calibration error or a gain setting because the differences in magnitude were frequency dependent. Looking at the axle accelerations, shown in Figure 42, we see a similar pattern to the chassis accelerations but the discrepancy in magnitudes is much less.

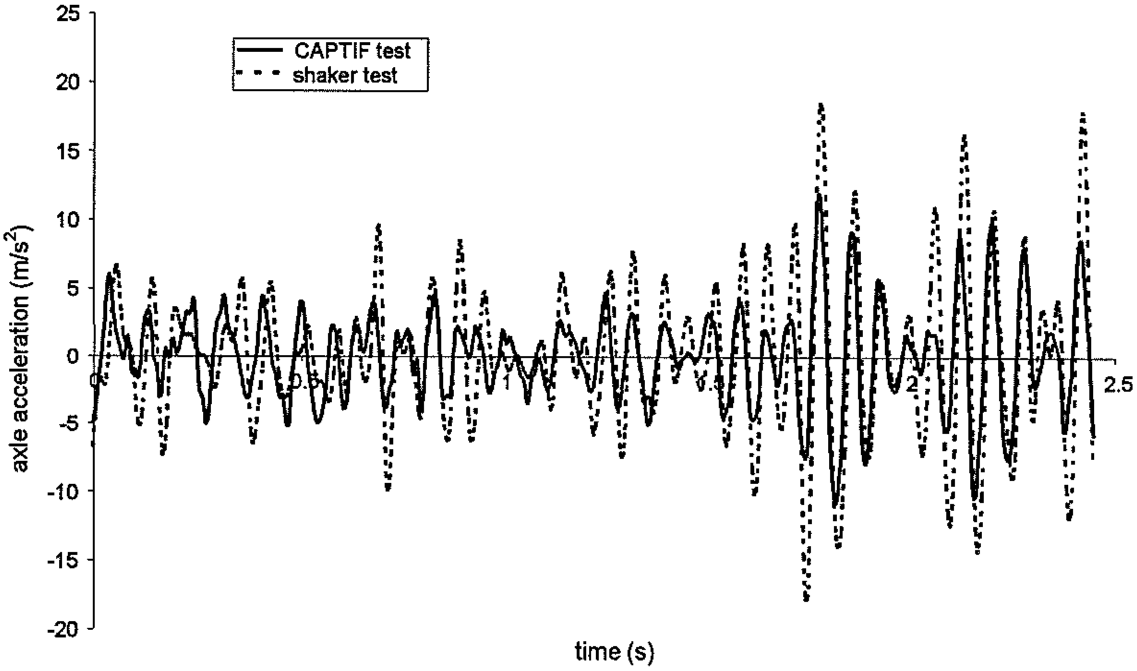

Figure 42. Comparison of axle accelerations between CAPTIF and shaker tests.

The problem was a concern because these accelerometer signals are used to calculate the wheel forces at CAPTIF. During the shaker tests the instrumentation on the shakers could also be used to calculate the wheel forces, which gives us a basis for comparison. These results were also surprising.

Figure 43 shows a comparison of the wheel forces calculated from the accelerometer signals with that from the load cells. The large magnitude response at the higher frequencies in the accelerometer signals results in corresponding components in the calculated wheel force signals. However, these components do not appear to the same extent in the wheel forces reported by the load cells. That is, these components in the accelerometer signal do not appear to be coming from vertical motions of the whole unit. To an extent this is expected. It is hard to imagine the whole 5 tonne mass of the SLAVE unit undergoing high frequency vibrations of this magnitude. The shaker simply cannot provide that much energy.

If we compare the load cell wheel forces with the accelerometer calculated ones from the CAPTIF tests, as in Figure 44 we find, somewhat surprisingly, that the match is better than with those calculated from the accelerometers during the shaker test. This is reassuring in that it indicates that the wheel forces calculated during the CAPTIF trial appear to be of approximately the correct magnitude but it does not explain the high levels of acceleration recorded during the shaker trials.

The SLAVE unit was then excited with sinusoidal excitations of various frequencies to see if the source of the problem could be identified. It was found that the first mode of the arm of the unit occurred at about 14 Hz. This mode causes a rotation of the chassis of the SLAVE about an axis in line with the tyre support. As the chassis accelerometer was mounted in line with the suspension rather that the tyre it registers this rotation as vertical acceleration. Thus it was hypothesised that the excitation of this mode was causing the higher than expected acceleration signals at these frequencies.

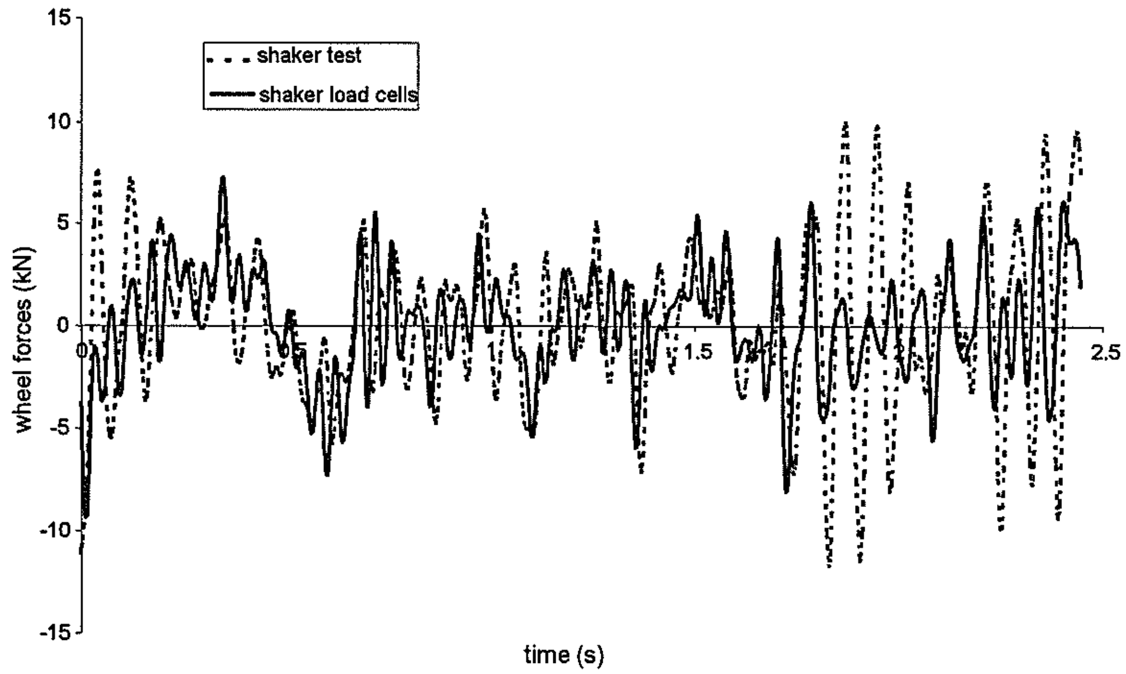

Figure 43. Comparison of wheel forces from accelerometers with load cells.

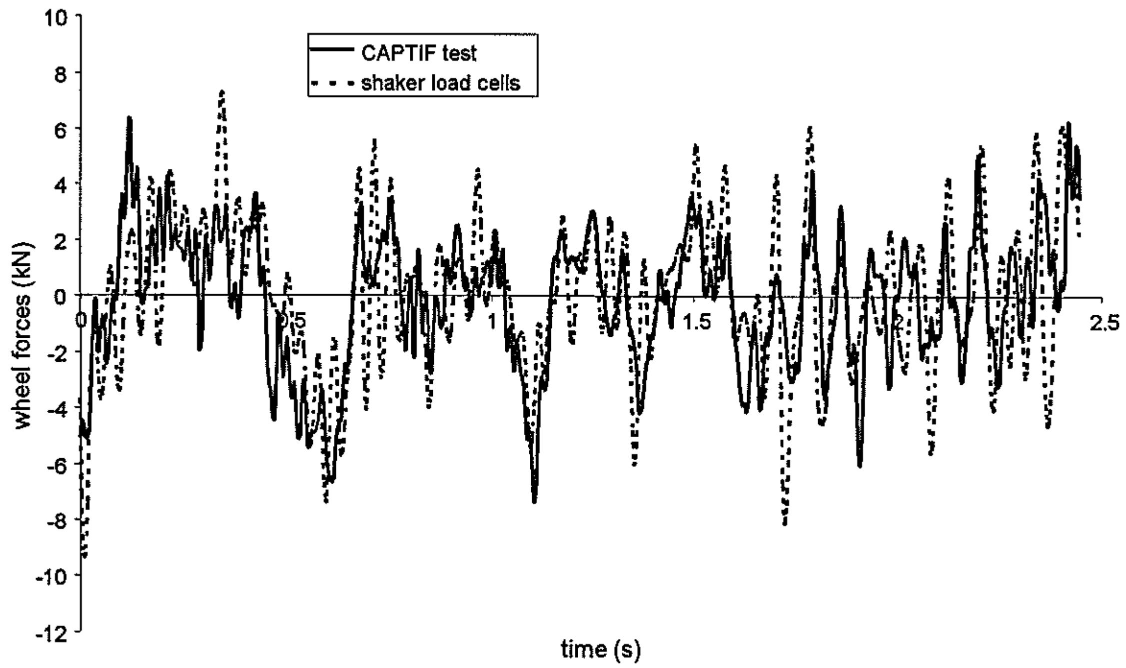

Figure 44. Comparison of wheel forces between CAPTIF test and shaker load cells.

To test this theory the chassis accelerometer was moved to a position in line with the centre of the wheels. Based on the original assumptions that the SLAVE units moved as a rigid body about the pivot at the end of the arms, this accelerometer signal should have been substantially the same as the previous one with possibly only a slight increase in magnitude. However, as shown in Figure 45, these two acceleration signals differ substantially and the one in line with the tyre has significantly less of the high frequency components. This supports the hypothesis that resonance modes in the arm are generating the problem high frequency signals. Unfortunately the frequency of these resonance modes is very close to that of the unsprung mass modes and so it is not possible to just filter them out.

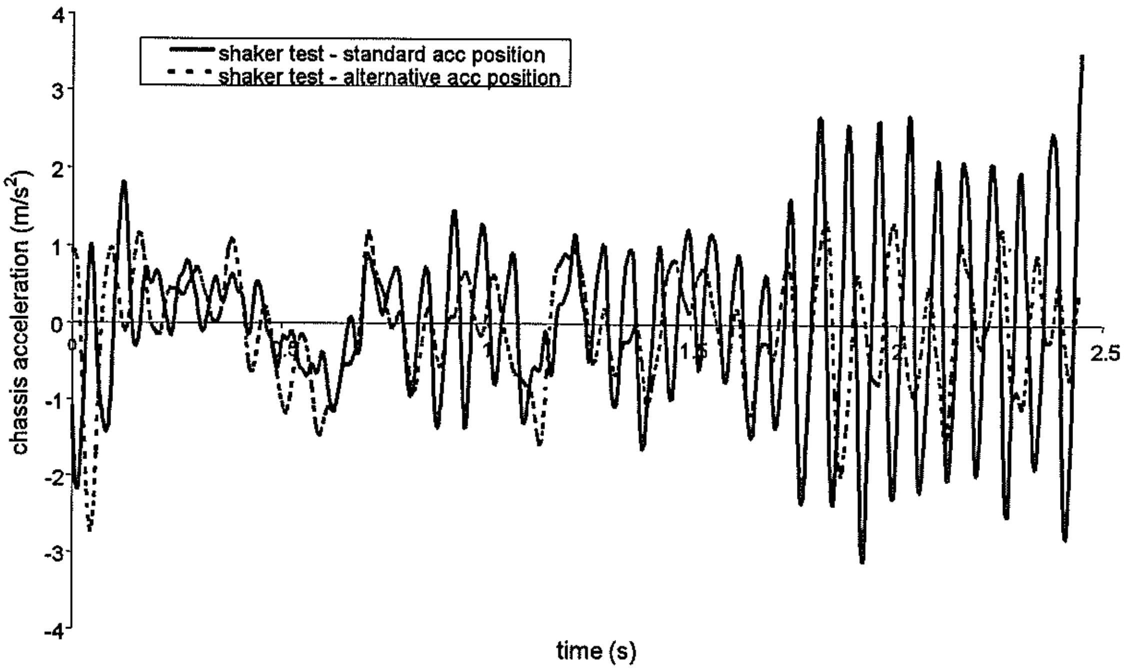

Figure 45. Comparison of two chassis accelerometer positions.

The other question is why these resonance modes were not a problem at CAPTIF. If we analyse the differences between the configuration for the shaker trials with that at CAPTIF we can postulate two probable reasons for this. The first is that at CAPTIF the vehicle was moving at 45 km/h while on the shaker it was stationary. At 45 km/h the vehicle experiences centrifugal accelerations of approximately 1.7g, which imparts high tensile forces on the arm. This will significantly increase the frequency of the arm resonance modes and moves them out of the range of interest. The second is that on the shakers the arm was attached to a support that was mounted rigidly to the floor. At CAPTIF the arm was attached to a frame mounted on the central pod. However, this frame is not restrained vertically but is held in place by its own weight. Thus there is a degree of vertical compliance, which will tend to damp out any resonance modes in the arm.

For the steel suspension, this problem of resonance of the arm was much less of a factor because the magnitude of the low frequency chassis bounce was more significant and dominated the response. In fact, some higher frequency components in the chassis, which were present in the CAPTIF signals, did not appear in the shaker test results.

The excitations with a noise signal at various amplitudes show, as expected, a linear response from the air suspension. That is, the measured signals from the three tests are virtual replicas of each other with the amplitudes changing in proportion to the excitation amplitudes. The behaviour of the steel suspension is much more complex with the signals from the three tests bearing only vague resemblance to each other. While the magnitude of the responses increased with increasing excitation the changes were not proportional and there were also frequency changes in the response.

Changing from wide single tyres to dual tyres had only minimal effect when the vehicle was fitted with air suspension. The chassis accelerations were very similar, while the axle acceleration showed some increase in the amplitudes of the response. There was no change in the frequencies of these responses. With the steel suspension changing tyres had a more significant effect. At low amplitude excitations the frequency of the main chassis acceleration response increased noticeably. At the higher amplitude excitations there was much less change. This is because of the stiction behaviour of steel leaf spring suspensions. At low excitation amplitudes the leaf spring remains locked and the vehicle bounces on its tyres. Thus changing the tyre stiffness has a significant impact. At higher excitation amplitudes the leaf spring begins to move and the response is controlled by its stiffness more than that of the tyres.

6.6 Conclusions

The results of the shaker trials on the CAPTIF SLAVE units were not as good as originally anticipated. Some of the difficulties normally encountered with replicating on-road response using road simulators should have been ameliorated in this case. Usually there are difficulties in synchronising the road profile data with the measured vehicle response because the vehicle's longitudinal and transverse position on the pavement with time is not accurately known. In this case the vehicle speed and transverse position are relatively accurately controlled by the system and referenced to the same clearly defined start point. When coupled with a relatively short track length this significantly reduces the synchronisation problem. Furthermore the SLAVE unit is a very simple dynamic system being effectively a "quarter truck" rather than a complete vehicle. In spite of these favourable attributes the response of the SLAVE showed significant differences between the shaker and the CAPTIF tests.

The major cause of these differences was the fundamental resonance mode of the connecting arm, which, in the shaker trials, had a natural frequency close to that of the axle hop mode. This generated quite high acceleration components at the chassis accelerometer in particular, without the corresponding wheel forces. This undermines the basis of the wheel force measuring system, which assumes that the accelerometers are recording the rigid body vertical accelerations of the sprung and unsprung masses. Fortunately this resonance of the arm does not appear to be present during the CAPTIF tests. It appears that the substantial centrifugal forces generated when CAPTIF is operating (1.7g at 45 km/h) may change the frequency of this mode out of the range of concern.

However, even though the shaker-based SLAVE responses did not match those from the CAPTIF well, some useful results were obtained from the trials:

- The match of the suspension deflections observed during the shaker trials with those recorded at CAPTIF was very good. This shows that the shakers did generate the correct suspension response.
- The natural modes of vibration of the structure generated acceleration signals during the shaker trials that were not directly related to the wheel forces. In section 2.5.5 it was shown that, for the test trailer, accelerometer signals could be used to give a reasonable estimate of wheel force. These tests provide a clear example where this is not the case. In

practice a similar situation will occur with relatively flexible vehicles such as flat deck trailers.

- Increasing tyre stiffness by changing from wide-based single tyres to dual tyres had very little effect with the air suspension, which was relatively soft and would be classed as "road-friendly". It did have an effect on the vehicle when fitted with the multi-leaf steel suspension because much of the time the steel suspension is effectively locked due to inter-leaf friction and the tyre is the main spring element.

7 SIMPLE COMPUTER MODELS

7.1 Introduction

As part of the New Zealand Public Good Science Fund research being conducted in association with this work, some relatively simple dynamic models were developed and used to investigate two important aspects of the problem of assessing suspension performance for pavement wear. The first of these was developing a method for eliminating the need for a road test as part of the assessment procedure. The second was developing a strategy for configuring a test vehicle so that a suspension assessment undertaken on that vehicle would be as vehicle-independent as possible. Although this modelling was not part of this research project as such, the findings of these two investigations have implications for the conclusions of this work and so they will be outlined in turn in this chapter.

7.2 Eliminating the Road Test

7.2.1 Introduction

The experiments that have been described in the previous chapters of this report form the basis of a potential suspension assessment procedure. Using relatively simple instrumentation the vehicle's suspension response to particular sets of operating conditions can be recorded. With the vehicle mounted on the servohydraulic shaker rig, the control software can then be used to determine the shaker excitations necessary to obtain the same suspension response in the laboratory. The corresponding wheel forces can then be recorded from the shaker instrumentation and the suspension's performance can be assessed.

The shaker tests described in chapters 3 and 5 have shown that these calculated excitation signals are similar to the road profile data (converted from the space domain to the time domain using vehicle speed) but have modifications to compensate for the differences in the method of excitation. The analysis outlined in the following sections, which has been described in detail by (de Pont et al. 1995), develops a technique for modifying the road profile to calculate the required shaker excitations. In this way the need for a road test can be eliminated. There are two major benefits from this. The first is that the test procedure would be simpler and faster reducing the cost. The second is that any arbitrary road profile data can be used for the testing making the procedure completely repeatable and transferable.

7.2.2 The Pitch Plane Model

For this analysis a simple pitch plane model as shown in Figure 46 was used. This model assumed only one spring and one damper at each axle position. Roll behaviour was ignored. This simplification was used in the last shaker tests on the vehicle with air suspension where the two post shaker rig used was set up so that each actuator excited a complete axle and so no roll motion was induced. This is justifiable because, except on very rough roads, roll is a negligible contributor to dynamic wheel forces (OECD 1992).

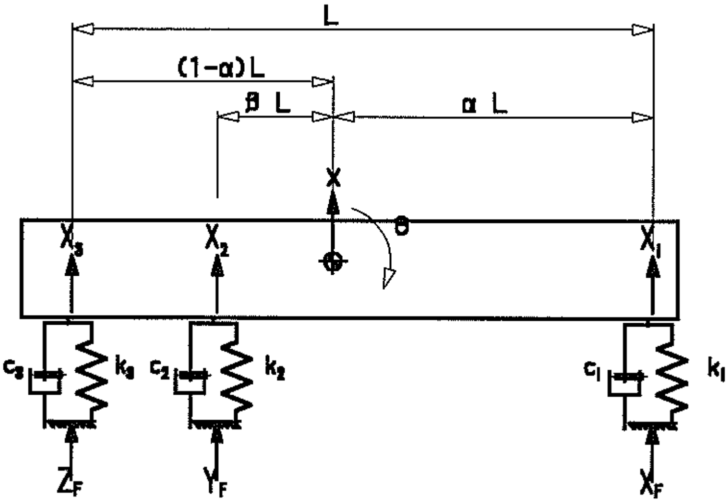

Figure 46. Simple pitch-plane model.

The model also ignores the unsprung mass behaviour of the vehicle. The equations of motion are:

$$m\frac{d^2x}{dt^2}+\sum_{i-1}^{3}(c_i\frac{dX_i}{dt}+k_iX_i)=0 \tag{1}$$

$$\frac{I}{L}\frac{d^2\theta}{dt^2}+\beta(k_2X_2+c_2\frac{dX_2}{dt})+\gamma(k_3X_3+c_3\frac{dX_3}{dt})-\alpha(k_1X_1+c_1\frac{dX_1}{dt})=0 \tag{2}$$

where the three suspension motions, X_i are related to the whole vehicle motions by

$$X_1=x-\alpha L\theta-X_F \tag{3}$$

$$X_2=x+\beta L\theta-Y_F \tag{4}$$

$$X_3=x+\gamma L\theta-Z_F \tag{5}$$

Although it might appear from the figure that the system has three degrees of freedom, the constraints described by equations (3) to (5) restrict the system to only two degrees of freedom. Thus the two equations (1) and (2) can be solved to give the following general solution.

$$y_i(t) = Y_i\exp[\{-\xi_i\omega_i + j\omega_i\sqrt{(1 - \xi_i^2)}\}t] \tag{6}$$

for i = 1, 2 and where the natural frequency of vibration and critical damping ratio are given by

$$\omega_i=\sqrt{\frac{(K_3+p_iI_rK_2)}{I_rm}};\quad \xi_i = \frac{\lambda\omega_i}{2} \tag{7}$$

These y_i are generalised variables which are a mixture of pitch and bounce motion as shown in equation (8)

$$y_i(t) = L\theta(t) + p_i x(t) \qquad (8)$$

and K_2, K_3 and I_r are composite stiffness and inertia terms. For a more detailed explanation the reader is referred to (de Pont et al. 1995).

In general, both natural modes of vibration are a mixture of bounce and pitch, which depend on the geometry of the vehicle. If I_r is less than a critical value the bounce mode is more prominent at the lower of the two frequencies and the pitch mode is more prominent for the higher. The frequencies of these resonance modes depend on the values of K_2, K_3, I_r, p_i and m, although these are not all independent.

The model's fundamental parameters are mass, pitch inertia, length, axle locations, Cg position and the three stiffness and damping values. Equations (1) and (2) can both be divided through by the mass, m, so that the inertia, stiffness and damping terms are all defined as quantities per unit mass. The actual value of the mass is then no longer needed. In the solution given by equations (6), (7) and (8) it was assumed that the damping is proportional to the stiffness for each axle. By weighing the vehicle axle by axle and measuring the axle positions, the mass, and Cg position can be calculated. Using modal testing (Ewins 1984) the natural frequencies, damping and mode shapes of the vehicle can be measured. (de Pont et al. 1995) show how these measurements can be used to calculate the vehicle parameters needed to evaluate the model.

Table 12. Vehicle parameters.

Mass (m)	Wheelbase (L)	α	β	γ
19535 kg	5.06m	0.58	0.173	0.42

Applying the method now to the three-axle tanker trailer fitted with air suspension, we measured and calculated the parameters listed in Table 12.

Using modal testing we found a natural frequency at 1.73 Hz but had great difficulty in finding the other frequency and mode shape we expected with confidence. Eventually we reached the conclusion that, in fact, the two natural frequencies were so close together at about 1.73 Hz (effectively both the same) that the modal analysis was not separating them. This necessitated some modifications to the method for calculating the model parameters. With these all the remaining parameters needed to evaluate the model were calculated and are listed in

Table 13.

Table 13. Vehicle dynamics parameters.

Inertia ratio (I/mL^2)	k_1/m N/kg.m	k_2/m N/kg.m	k_3/m N/kg.m	c_i/k_i λ
0.182	39.95	39.09	39.09	0.011

7.2.3 Response of the model to excitations

Having defined and calibrated the model we can use it to study the response of the vehicle due to external excitations applied to the three wheels, i.e. X_F (front wheel), Y_F (middle wheel), and Z_F (rear wheel). If these excitations are generated by a road profile they can be regarded as corresponding to a single excitation disturbance applied to each of the three axles with appropriate time lags which depend on the velocity of the vehicle. Using this form of excitation, the model equations can be solved to give the frequency response functions for the generalised co-ordinates, which can then be transformed to any of the model's physical co-ordinates. These could be the body bounce and pitch motions or the suspension displacement at any of the wheel positions.

Similarly we can solve for the shaker testing case where two independent excitations are applied to the two rear axles of the vehicle and no excitation is applied to the front axle. Again the model equations can be solved to give the frequency response functions for any of the model's physical co-ordinates.

By equating the suspension response at the two rear axles for the road profile excitation to that of the shaker excitation we can calculate the relationships needed between these two forms of excitation in order to generate the same response at the middle and rear axles. If we consider the excitations in the frequency domain these relationships can be regarded as transfer functions where a gain and phase shift applied to the road profile at each frequency generate the shaker excitations. As the units of the road profile and the shaker excitations are the same these transfer functions are dimensionless. The magnitudes of these transfer functions for the model of the trailer with air suspension for two different road speeds are shown in Figure 47 and 48.

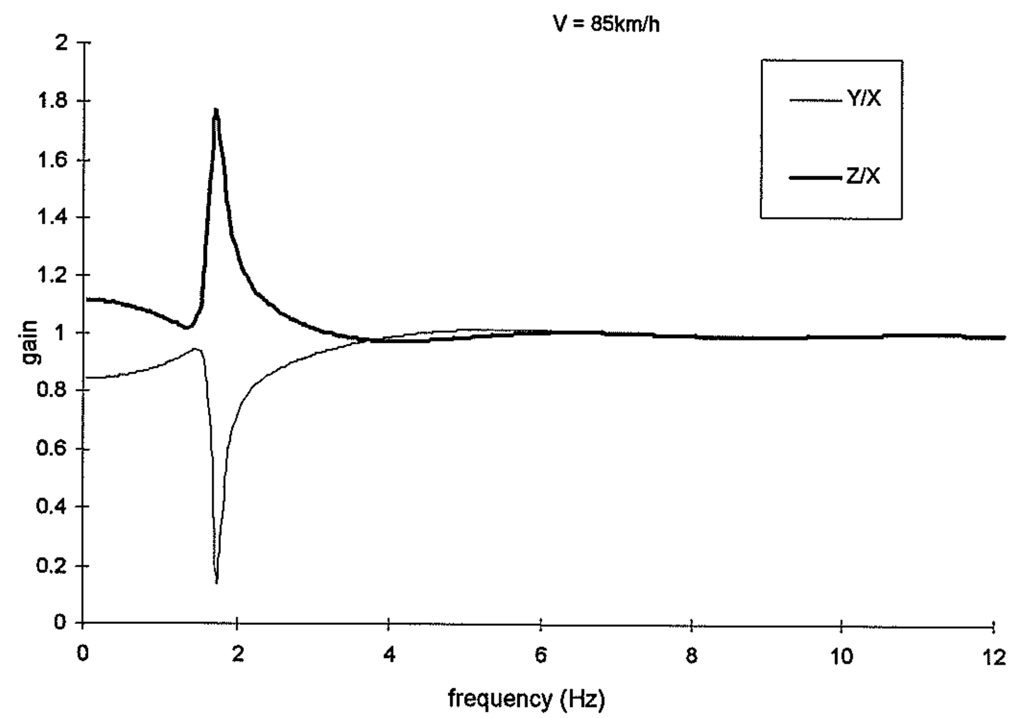

Figure 47. Transfer function magnitudes for v = 85 km/h

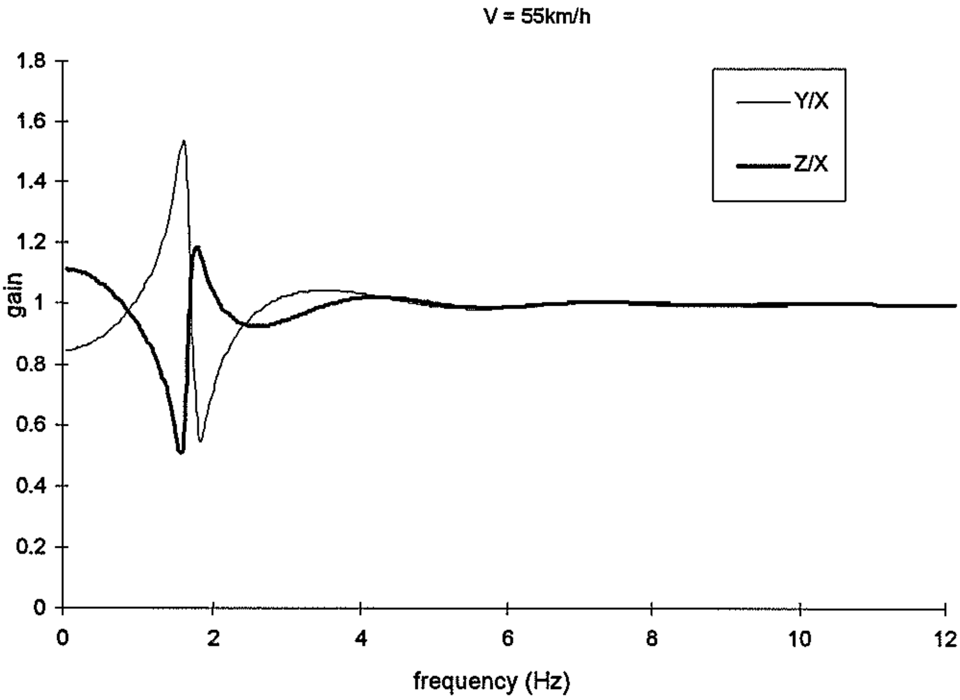

Figure 48. Transfer functions for v = 55 km/h

Note that these functions match the suspension displacement response of the model at middle and rear axles only. There are not sufficient degrees of freedom to match the behaviour at the front axle.

As can be seen these transfer functions are dependent on the vehicle speed. The effect of wheelbase filtering changes significantly with speed. The general pattern is that, at low frequencies, the simulated laboratory excitations are slightly modified versions of the road profile which compensate for the front axle excitation which occurred on the road. At the vehicle modes the compensation required is more complicated and speed-dependent and at higher frequencies the transfer functions tend towards unity i.e. the road profile excitations are applied unchanged.

7.2.4 Calculating laboratory excitations from pavement profiles

The transfer functions developed in the previous section can be used to calculate the shaker excitations to be applied at the middle and rear axles from measured road profile data.

Figure 49 shows the raw pavement profiles measured for the Church St site by the Australian Road Research Board laser profilometer. As the profile computation involves the double integration of an accelerometer signal (to compensate for vehicle motions), a small offset error in the measured acceleration signal results in a large parabola being superimposed over the signal. Thus we get the situation shown where there is an apparent difference in elevation of up to 5m between the left and right wheel paths, which is clearly impossible. As this parabola is of much greater magnitude than the underlying signal it is not easy to filter it out. The best method we found was to regression fit a second order polynomial and subtract this from the data. With this approach a good repeatability between repeat measurement sets was achieved. Figure 50 shows the processed profiles, which have been converted to the time domain based on the speed from a particular test run.

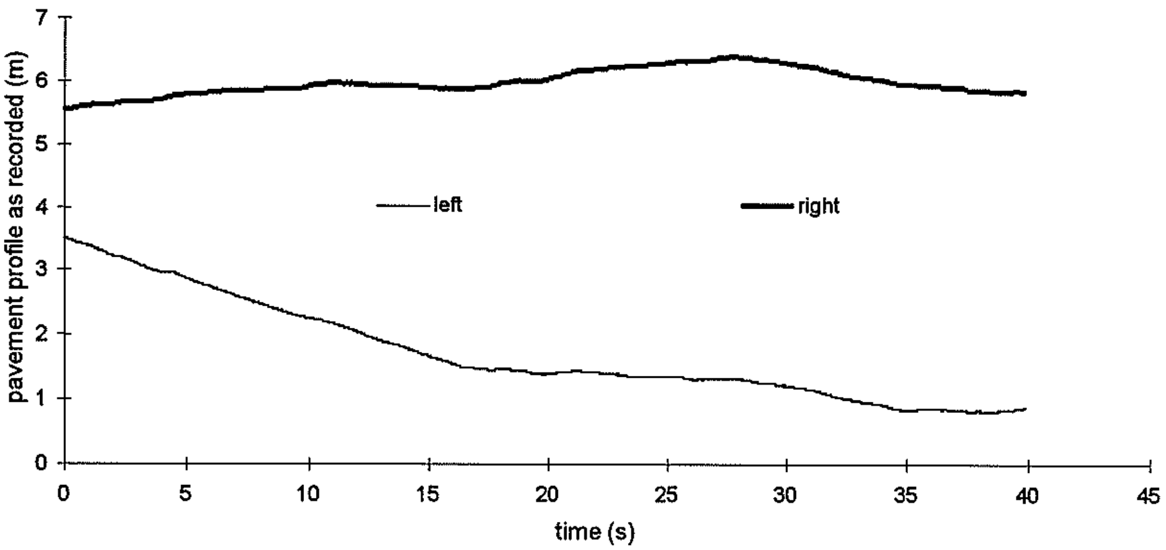

Figure 49. Measured road profiles.

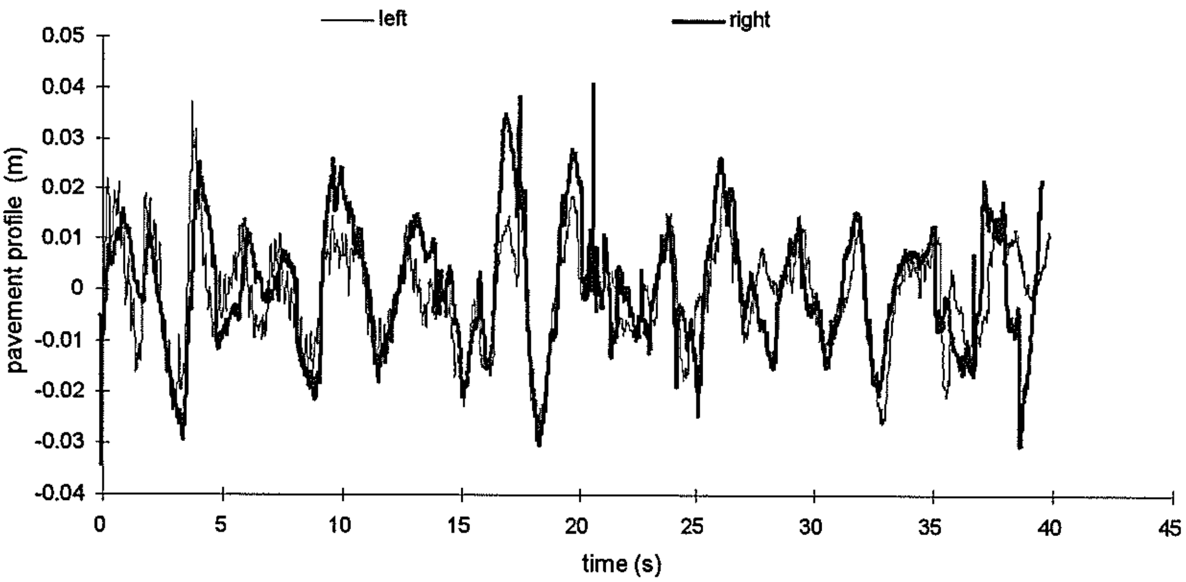

Figure 50. Filtered road profiles.

These pavement profile data were converted to the two equivalent shaker excitations using the transfer functions described above. Figure 51 and 52 show a comparison between these computed shaker excitation signals and those generated by the SYSCOMP software package in matching the measured on-road behaviour of the actual vehicle. The shaker excitations generated by SYSCOMP result in a very good match between the suspension displacements measured in the laboratory and those measured on the road (de Pont 1996a).

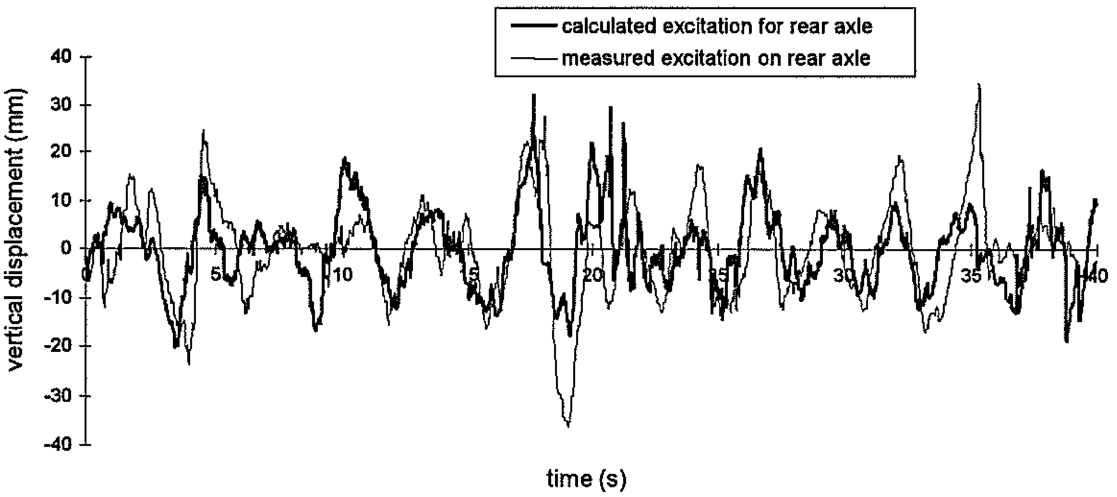

Figure 51. Measured vs 'computed from profile' rear excitations.

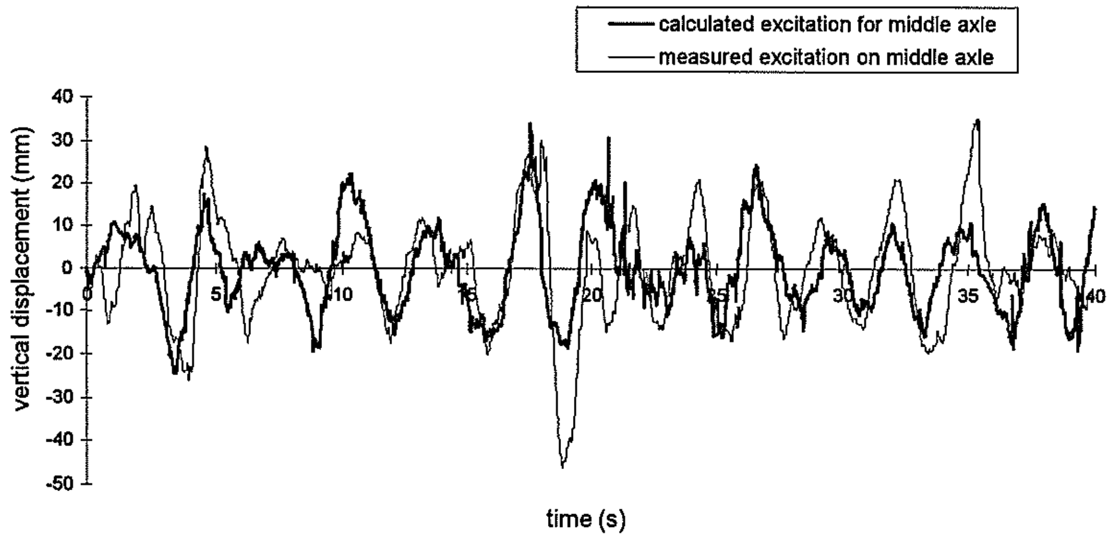

Figure 52. Measured vs 'computed from profile' middle excitations.

The values of the excitations computed from the road profiles using the model are in reasonably good agreement with those generated by SYSCOMP except for some time shifts at some points. This can be attributed, at least in part, to speed fluctuations during the road test. The profile data were converted to the time domain assuming a constant vehicle speed. Although the tests were nominally undertaken at steady speed some variations are inevitable. A further source of variability is that the measurements of vehicle response and pavement profile were not carried out simultaneously. There was an interval of over a year between these measurements and thus it is likely that some profile changes occurred in this time. Bearing these two factors in mind the match between the calculated excitations and the measured ones is satisfactory.

7.2.5 Calculating excitations from the response

The transfer functions can also be used to calculate shaker excitation signals from the measured suspension displacements. These are the excitations needed to elicit these responses from the model and not necessarily those required for the actual vehicle. The results of this computation are shown in Figure 53 and 54 where they are compared with the laboratory measured excitations.

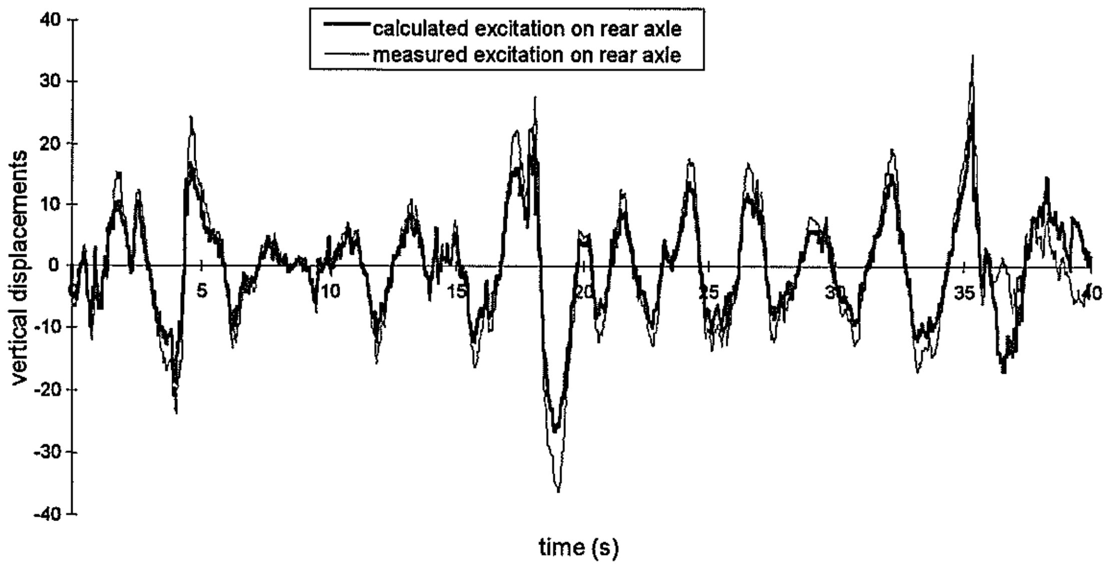

Figure 53. Measured vs 'computed from response' rear excitations.

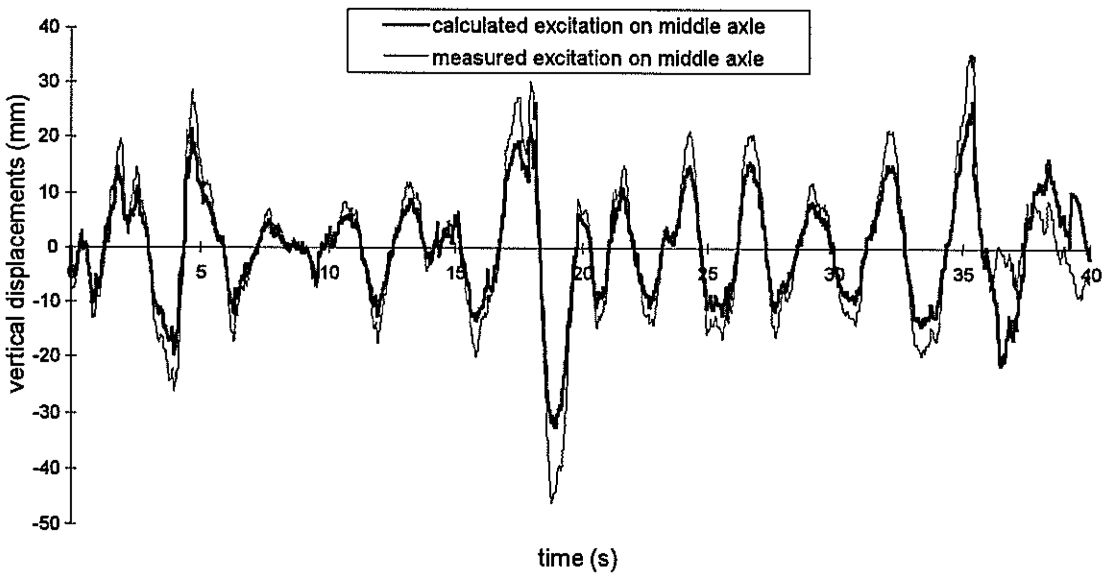

Figure 54. Measured vs 'computed from response' middle excitations.

The agreement between computed and exact data is excellent and in fact much better than the comparisons in Figure 51 and 52. This indicates that the model represents the sprung mass behaviour of the vehicle quite well and that the differences between the excitations calculated from the road profiles and the measured ones are probably caused by errors in the profile rather than inadequacies of the model.

It is possible to use this fit between the model response and the actual vehicle response to optimise the values of the model parameters. This is an alternative approach to the modal analysis undertaken previously. If we do this we can recalculate the excitations using the optimised model parameters. The results of doing this are shown in Figure 55 and 56.

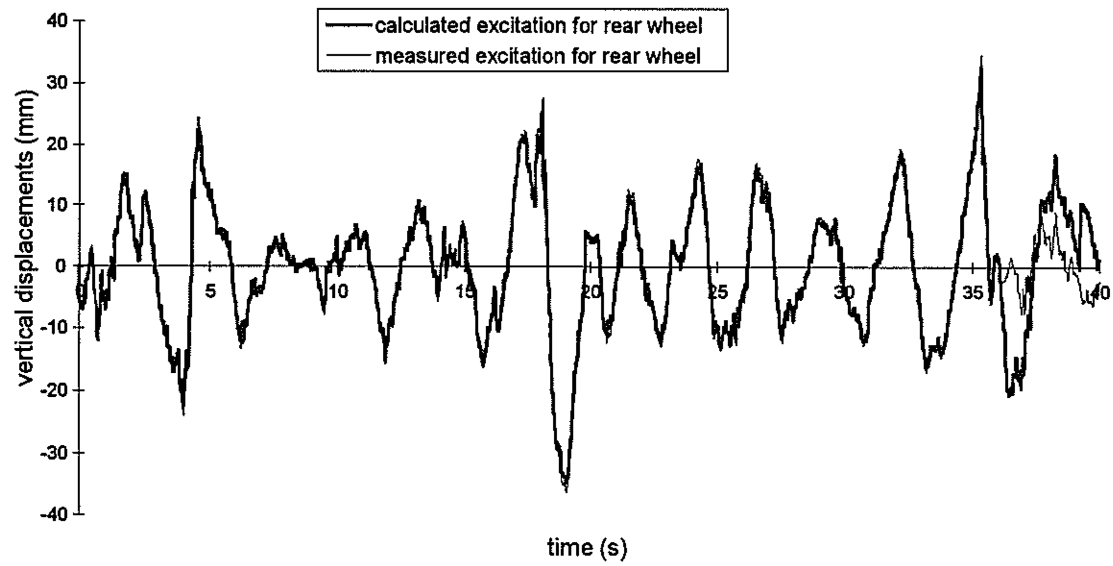

Figure 55. Measured vs 'optimised computed from response' rear excitations

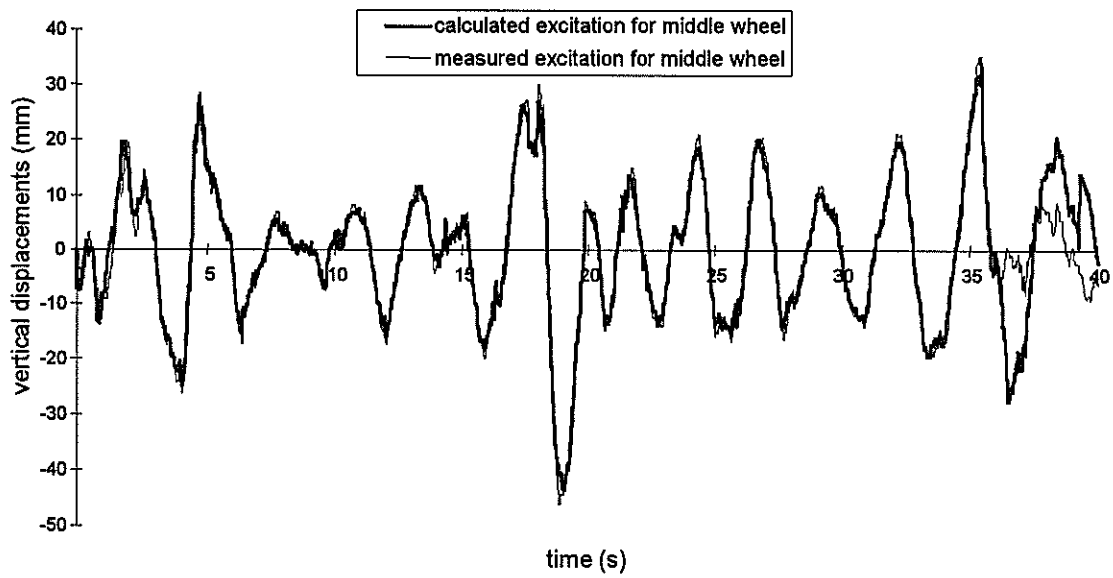

Figure 56. Measured vs 'optimised computed from response' middle excitations.

The fit is almost perfect. There is a minor discrepancy at the end of the time trace which comes about because the Fourier transform procedure used assumes that the signals are circular which is not correct. The original vehicle response was measured during a test on a straight section of road. There is no reason why the end of the recorded response should match up with the beginning. The optimised model parameters used for this fit are given in Table 14.

Table 14. Optimised vehicle dynamics parameters.

Inertia ratio (I/mL^2)	k_1/m N/kg.m	k_2/m N/kg.m	k_3/m N/kg.m	c_i/k_i λ
0.140	44.32	43.37	43.37	0.010

Referring back to

Table 13, it can be seen that the changes in parameters are small. The spring stiffnesses have increased by about 11% and the inertia ratio is much closer to the value estimated from the vehicle geometry. The damping ratio is virtually unchanged.

7.2.6 Higher frequency modes

The simplified model described above simulates the sprung mass response of the vehicle and when used to calculate shaker excitations from road profiles only modifies the low frequency components. The higher frequency components of the road profile are unchanged and also part of the shaker excitations. If we consider the typical dynamic behaviour of heavy vehicles we see that this is perfectly reasonable. At higher frequencies the modes of vibration are primarily unsprung mass modes, that is, vibrations of the axle assembly. These will be stimulated by direct excitation at the axle concerned and largely unaffected by inputs at the other axles. For this reason, the excitations required in the laboratory to produce the same behaviour as on the road will be very similar.

7.2.7 Summary

A simplified model was developed for converting the in-service dynamic excitations applied from passing over a road profile into an equivalent pair of shaker excitations to produce the same suspension response at the two axles being excited. This model was not intended to realistically model the vehicle dynamics in detail.

The model simulates the sprung mass response of the vehicle only and uses linear spring and damper elements. We have developed an algorithm for using modal testing measurements on the vehicle to estimate the model's parameters. Using the model it was possible to calculate the required laboratory excitations from road profile measurements. These excitations were compared with those generated using the shaker control software. The match obtained was reasonable given that the input road profile data were not ideal. (They were recorded a considerable time before the vehicle response measurements and the conversion to time co-ordinates using vehicle speed was rather simplified). Using the vehicle response measurements as the input to the excitation calculations gave a significantly improved match. Using the measured response to optimise the model parameters gave even better results. This suggested a modified modal analysis procedure, which used the shakers as excitations and suspension displacement measurements as the response.

7.3 Vehicle Configurations and Procedures for Testing

7.3.1 Introduction

For practical reasons it is better to have an assessment procedure which rates suspensions in a vehicle independent way than to have to rate each vehicle-suspension combination. Ideally, for simplicity and cost reasons, this rating should be able to be done with the suspension fitted to an actual vehicle rather than to some special test rig. This analysis, which is presented in detail in (de Pont 1996b), investigates the effect of changing various vehicle parameters on the response of the suspension under test. This has two purposes. One is to determine whether it is reasonable to assess a suspension in isolation or whether the whole vehicle effects are too important. The other is to establish vehicle configurations and test procedures that minimise the influence of the rest of the vehicle on the results for the suspension under test.

7.3.2 The simple pitch-plane model

The basis of this investigation is the pitch-plane vehicle model shown in Figure 57, which is even simpler than that used in section 7.2. This model represents only the sprung mass resonance modes of the vehicle. It is assumed that whole vehicle effects will not significantly affect the unsprung mass resonance modes.

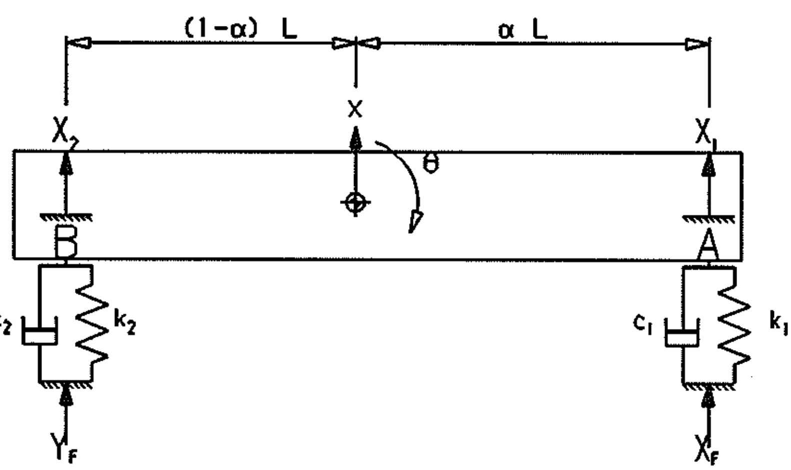

Figure 57. Two degree of freedom pitch plane model.

This model can be regarded as representing either a rigid truck or a full trailer. It can also approximate a semitrailer or tractor unit but it does not model the effect of the dynamic behaviour of the other vehicle in such a combination. The suspensions could represent a single axle or an axle group.

For this analysis it is assumed that the front of the vehicle is position A and that the suspension being assessed is that at position B. Thus the parameters of this suspension were held constant and the influence on its behaviour of varying the other vehicle parameters was investigated.

The equations of motion for this model are simplified versions of equations (1)-(5) in section 7.2.2 as the third axle position does not exist. As before these equations can be manipulated to eliminate the body bounce and rotation variables leaving the body motions at the two axle positions as the only two variables.

In this analysis, two cases were considered for the excitation displacements, X_F and Y_F. The first case simulated the normal roadway excitation with the two excitations identical but separated by a time lag. The second case simulated a suspension assessment test situation where an excitation is applied to the suspension being evaluated and nothing is applied to the other suspension.

In both cases, the equations of motion were solved analytically to determine the frequency response functions for the displacements and from these, the frequency response functions for the wheel forces. Typical road profiles have displacement power spectral density functions that are inversely proportional to the frequency squared. Thus the amplitude spectra of typical road profiles are proportional to the inverse of frequency and so weighting the wheel force frequency response function by the inverse of frequency will give a weighted wheel force frequency response function which is proportional to a typical on-road response.

Dynamic wheel forces are usually characterised with a measure known as the dynamic load coefficient (DLC) which is defined as the standard deviation of the wheel forces divided by the static wheel load. This can also be calculated by taking the square root of the integral of the power spectral density function of the wheel forces and dividing that by the static load. Therefore the weighted wheel force frequency response function can be used to calculate a function which is proportional to the portion of DLC generated by the sprung mass resonance modes. To distinguish this from true DLC we will call it the wheel force coefficient (WFC).

7.3.3 A reference model

The model and the measures that will be derived from it have now been defined. Before investigating the influence of the various whole body effects on the model's behaviour, a reference set of values must be determined. The parameter values for this reference model have been selected somewhat arbitrarily but with the aim of representing a credible vehicle and suspension configuration.

As stated previously it was assumed that the suspension being analysed is at position B. The static axle load at that position was fixed, arbitrarily, at 10 tonnes. The wheel base, L, was set at 6 m and the position of the centre of gravity at 3.75 m from position A (ie. $\alpha = 0.625$) which gives a total vehicle mass of 16 tonnes. The moment of inertia, I, can be written as mr^2 where r is the radius of gyration. Therefore, a quantity called the inertia ratio is defined as $I_r = (r/L)^2$. The inertia coefficient in the matrix form of the equations of motion then becomes mI_r. For the reference model I_r is set to 0.1. (If the mass were a lumped mass at the centre of gravity, I_r would be zero, if it were two lumped masses located at the suspension positions, I_r would be equal to 0.234.) The vehicle speed was arbitrarily set at 90 km/h or 25 m/s. The spring stiffness values and damping rates were chosen to reflect what would be regarded as "road-friendly" suspensions, ie. $k_1 = 600$ kN/m, $k_2 = 1,000$ kN/m, $c_1 = 24$ kN/m/s, and $c_2 = 40$ kN/m/s. The relative values at the two ends of the vehicle are proportional to the axle loads with, as a result, a natural frequency in bounce of 1.59 Hz and a damping ratio of 20 %.

Using these parameters the magnitude of the displacement frequency response function at position B for a roadway excitation is as shown in Figure 58. Both the bounce and pitch resonance modes are clearly visible. There also appears to be a wheelbase filtering effect, which shows up as waviness at the higher frequencies. At low frequencies the response is unity, that is, the vehicle body moves with the road profile and the suspension does not compress. At high frequencies the response tends towards zero. That is, the road unevenness is nearly all being absorbed by suspension motion with very little movement of the vehicle itself.

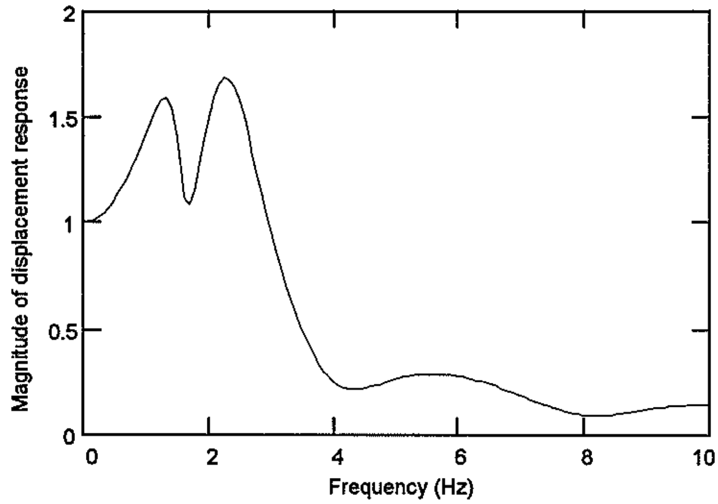

Figure 58. Displacement frequency response function for road excitation

For a suspension test excitation where the other end of the vehicle is not excited the corresponding displacement magnitude response function is as shown in Figure 59

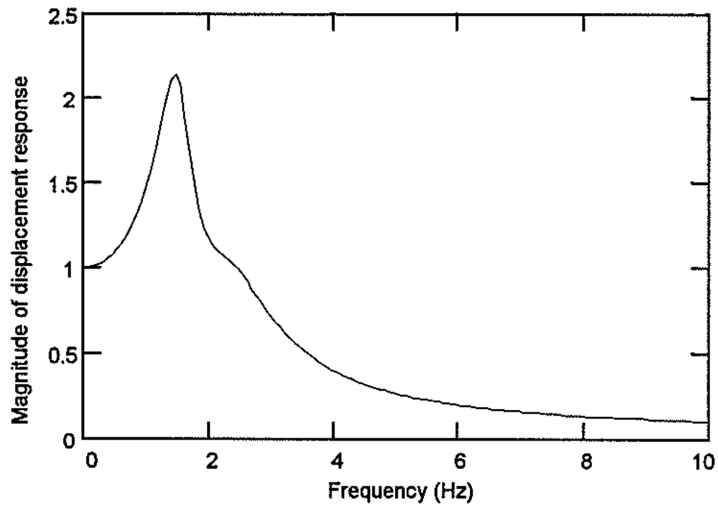

Figure 59. Displacement frequency response function for suspension test excitation

In this situation the response is dominated by the bounce resonance mode with the pitch mode appearing only as a slight bump in the curve. There are, of course, no effects from the interaction of the wheelbase with the vehicle speed.

The corresponding weighted wheel force response functions for the two excitations are shown in Figure 60 and 61 respectively. The forms of both these are very similar to those of the displacements. Calculating the WFC values for these two cases gives values of 0.721 and 0.761 respectively. The absolute values of these are not meaningful because the weighting function was not scaled to have realistic amplitudes. However, it is interesting that the two forms of excitation give results that are only approximately 5% different. The roadway excitation value of WFC is used as the baseline for the parameter assessments that follow.

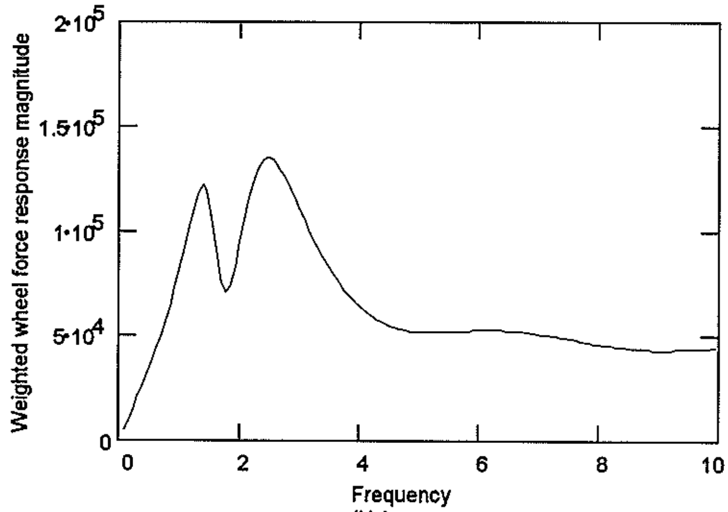

Figure 60. Weighted wheel force response magnitude for road excitation

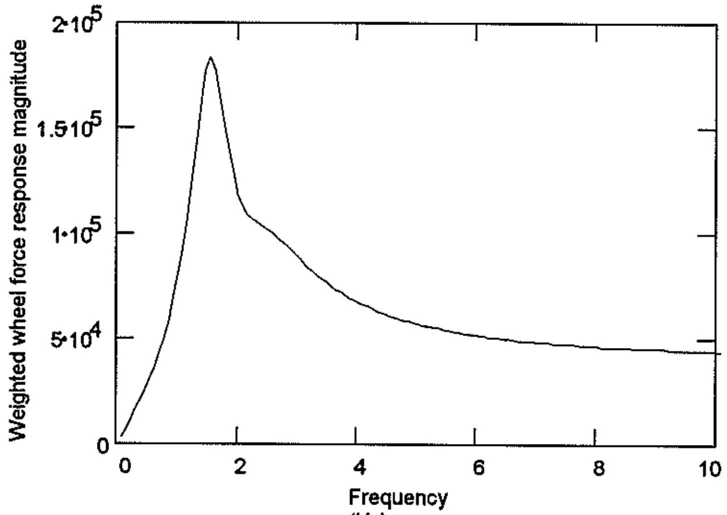

Figure 61. Weighted wheel force response magnitude for suspension test excitation

7.3.4 The effects of changing vehicle parameters

The model was then used to evaluate the effects of changing various vehicle parameters on the response of the suspension with both a road-type excitation and a shaker-type excitation. The parameters considered were: the centre of gravity (Cg) position, the vehicle speed, the pitch inertia, and the stiffness and damping of the other suspension on the vehicle, varied both independently and together. Each parameter was varied through a range of possible values.

The results of analysing all these parameter variations are presented in detail in (de Pont 1996b) and will not be repeated in full here. The analysis of one parameter, Cg position, will be used as an example of what was done. In varying the Cg position there are two logical options for the corresponding variation in the mass. One is to keep the axle load at position B constant and allow the total vehicle mass to vary. The alternative is to keep the total vehicle mass constant and adjust the axle loads proportionately. In this analysis because it was assumed that we are trying to assess the suspension at position B, the axle load at this position is kept constant.

Figure 62 shows the weighted wheel force responses with the roadway excitation for a selection of Cg positions. Although it is theoretically possible for the parameter α to vary from 0 to 1, even the range shown in the figure (0.375 - 0.75) is probably extreme for the relative suspension stiffness values because of the resulting static loads at position A.

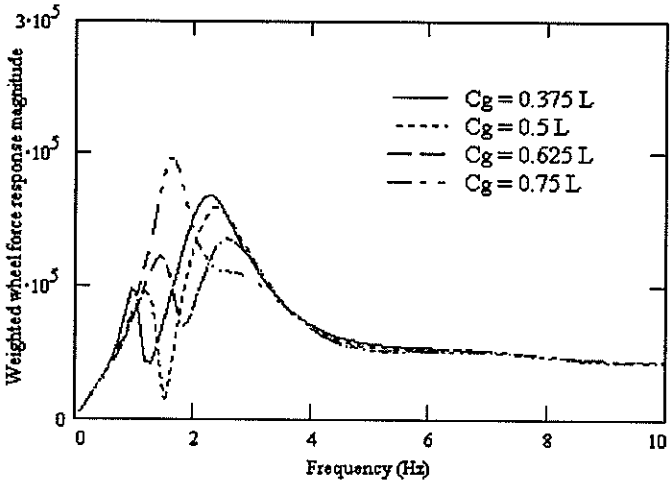

Figure 62. Effect of Cg position on weighted wheel force response to road excitation

With the lowest value of α, the response is dominated by the "pitch" resonance mode, which occurs at the higher of two frequencies. As the Cg position moves further to the rear of the vehicle the response of the "bounce" resonance increases and the response of the "pitch" mode decreases. The frequencies of both these resonance modes also increase. This is because as the Cg moves to the rear, the total vehicle mass required to provide the correct axle load at position B decreases.

Figure 63 shows the weighted wheel force responses when the suspension test excitation is applied. The trends are identical and more clearly defined with very uniform shifts in magnitude and frequency. In the roadway excitation case it appears that wheelbase-filtering effects may be influencing the responses of the two vibration resonance modes.

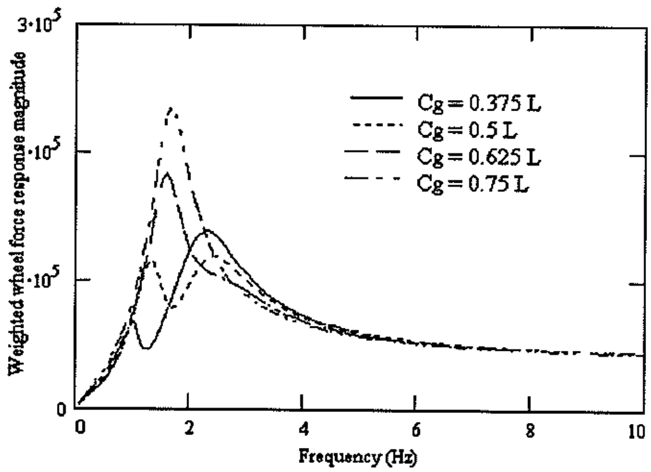

Figure 63. Effect of Cg position on weighted wheel force to suspension test excitation

The effect on the WFC value as a ratio of the reference value for the two excitation situations are shown in Figure 64.

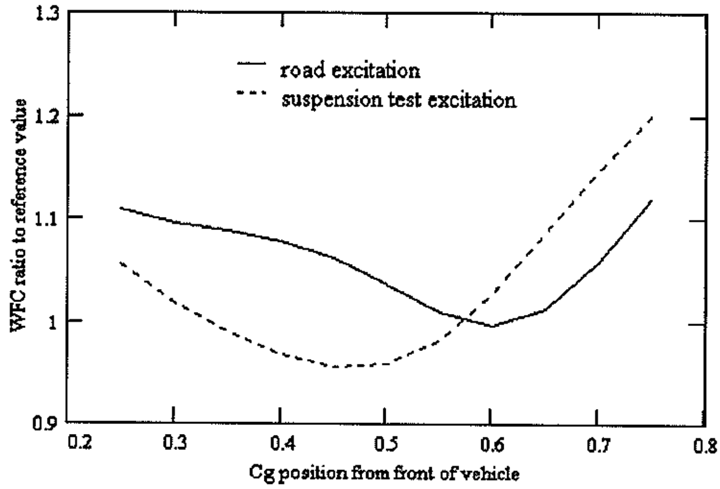

Figure 64. Effect of varying Cg position on WFC

For the roadway excitation, the minimum value for WFC occurs when the Cg is approximately at the reference value. In this situation the two suspensions are carrying loads in proportion to their stiffness. With the suspension test type of excitation the minimum occurs when the Cg is nearer the other end of the vehicle (this also increases the axle load at that end of the vehicle).

Vehicle speed only affects the road-type excitation response. These effects are quite pronounced as the wheelbase filtering causes the resonance modes to be reinforced or suppressed depending on the speed. However, when the fact that the road appears rougher to the vehicle as speed increases is taken into account this is the dominant factor and wheelbase effects are minor. These results were consistent with the measurements of dynamic wheel forces undertaken by a number of researchers (Hahn 1985; Mitchell and Gyenes 1989; Sweatman 1983; Woodrooffe et al. 1986).

The pitch inertia is primarily controlled by the load distribution. If the vehicle load were to consist of a point mass at the Cg the inertia would be zero. As the pitch inertia increases the frequency of the pitch resonance will decrease. When the inertia is equivalent to that of two point masses located over the suspension positions the pitch resonance merges with the bounce resonance to give a single peak in the response function.

The effects of pitch inertia change on WFC for the road and suspension test excitations are very similar except at very high pitch inertia values. The suspension test excitation generally gives slightly higher WFC values but the relationship between the two is consistent.

We then considered how the stiffness of the suspension at the other end of the vehicle influences the behaviour of the suspension being assessed. As this stiffness was increased the natural frequencies of the resonance modes increased. For the roadway-type of excitation the relative magnitude of the two response modes also changed with the lower frequency bounce mode increasing and the higher frequency pitch mode decreasing.

Figure 65 shows the effect on WFC of varying the spring stiffness at the other end of the vehicle for the two excitation cases. Both of these have minima at some intermediate value of stiffness. For the road excitation case, this minimum is at approximately 500 kN/m, which is a little below the reference value. However, if the spring is softened further the WFC increases. That is, making the other suspension more "road-friendly" increases the wheel forces from the test suspension. Note that the damping was not changed as this stiffness was varied. For the suspension test excitation, the pattern is similar but the location of the minimum is at a much lower stiffness (approximately 300 kN/m).

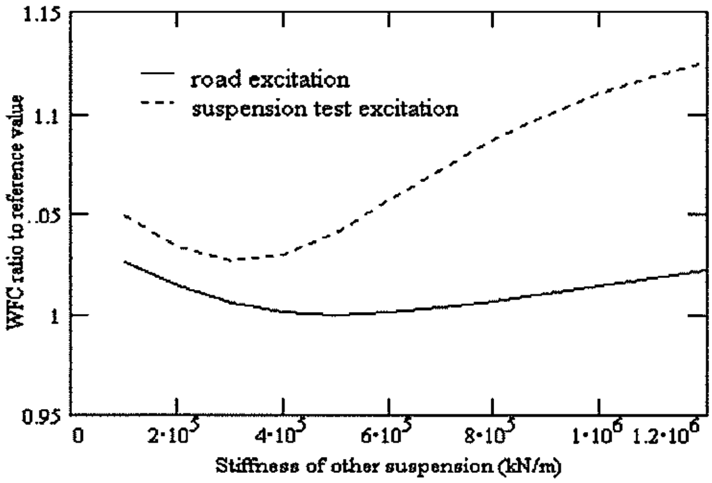

Figure 65 Effect of varying stiffness on WFC

As with the stiffness, the damping values for the other suspension were varied from zero to double the reference value. The results were entirely as expected. Increasing the damping had virtually no effect on the resonant frequencies. The magnitudes of the responses decreased with increased damping and broadened so that the delineation between the two resonance modes was less clear. The effect on WFC for the two excitation cases showed similar trends except at high damping levels when the suspension test excitation showed a rise in WFC. In general, increasing the damping at the other end of the vehicle results in lower dynamic wheel forces at the test suspension.

The analysis of variations in the stiffness and damping of the suspension at the other end of the vehicle assumed, in both cases, that the other parameter (damping and stiffness) was kept constant. It is probably more reasonable to assume that the ratio between stiffness and damping would stay constant as they were varied. In fact, the behaviour was very similar to that observed when just the stiffness was varied. With the lower stiffness values there was also less damping, which had the effect of delineating the two resonance modes more clearly and increasing their response magnitudes slightly. Similarly the high stiffness case was more damped and the reverse effect was observed. These differences were relatively small.

The effects on WFC of these parameter variations were also similar to those for varying the stiffness independently. Again it was demonstrated that improving the road-friendliness of the other suspension could result in an increase in the dynamic loadings of the suspension under consideration. However, it should be noted that this also results in a large reduction in the dynamic loading of the suspension that has had its stiffness reduced. If the wheel force responses from the two suspensions are combined to give the total vehicle wheel force response applied to the pavement a whole vehicle WFC can be calculated. This increases monotonically with increasing suspension stiffness.

7.3.5 Implications for road friendliness assessment

The purpose of this analysis was to address two issues. Firstly, how significantly do the whole vehicle effects influence the response of a suspension and does this mean that a whole vehicle approach to road friendliness rating is required. Secondly, if a suspension test for road friendliness can be justified what strategies should be implemented to minimise the whole vehicle effects on the results.

Changes in vehicle speed have two effects. The first is that coupled with wheelbase certain speeds reinforce or suppress the resonance modes of the vehicle under a roadway type of excitation. (Speed is not a parameter in a suspension test type of excitation.) The second effect is that as speed increases the apparent roughness of the road, as seen by the vehicle increases resulting in higher levels of dynamic loading. This second effect is dominant with the first effect providing minor perturbations to the overall trend. This roughness effect is not related to whole vehicle behaviour and can be accommodated in a suspension test.

The question of changes in Cg position and pitch inertia is related primarily to the location and distribution of the payload. These are very much whole vehicle factors. However, for many vehicles they can both vary from trip to trip. It is difficult to see how a whole vehicle assessment test could encompass all possible configurations. In general it appears that if, as in the reference case, the suspension stiffness values are matched to a particular Cg position then varying that Cg position will result in increases in the dynamic wheel forces. However, for a practical range of Cg positions, these increases are small (<10% in the case of the model). For the reference model with a roadway excitation increasing pitch inertia increases dynamic wheel forces. A similar trend is present when a suspension test form of excitation is applied.

The third category of parameters evaluated was the stiffness and damping of the suspension at the other end of the vehicle. This is the main area where whole vehicle assessment could have some justification. The effect of changing the damping of the other suspension was as expected. Reducing the damping results in higher dynamic wheel forces at the suspension under consideration although a very substantial change in damping is needed before the increase in wheel forces is significant. Increasing the damping has very little effect. Note that this was relative to the reference model which had both suspensions well damped. Varying the stiffness (or the stiffness and damping together) of the other suspension shows the effect described by (Cebon 1993) i.e. that a reduction in the stiffness of the other suspension can result in an increase in the dynamic loads of the suspension being considered. However, there was a reduction in the total dynamic loads applied by the vehicle. Thus, although improving the road friendliness of one of the suspensions may result in a less road friendly performance from the other suspension the overall road friendliness of the vehicle is improved. Thus this factor does not justify the need for whole vehicle assessment rather than a suspension assessment.

A possible method for incorporating suspension road friendliness ratings into a vehicle road friendliness rating is as follows. If each suspension, i, has a rating r_i which is proportional to the level of dynamic loading it generates (in particular circumstances) and has a static wheel load, m_i, then the whole vehicle rating R could be calculated using the formula:

$$R = \frac{\sqrt{\sum_i m_i^2 . r_i^2}}{\sum_i m_i}$$

If the distributions of dynamic wheel force from the different suspensions were independent, this would be the correct statistical formula for the DLC of the average dynamic wheel forces. However, although the distributions are not independent, their dependence is not deterministic as it varies with operating conditions. This formula tends to favour having all suspensions on a vehicle with similar performance and discourage configurations that include an unfriendly suspension. From the previous analysis this appears to be a desirable characteristic.

The second issue to consider is how to design a suspension assessment test with the suspension fitted to an actual vehicle so that the influence of the rest of the vehicle is minimised. As mentioned previously, (Woodrooffe 1997) addresses this issue by using a dedicated purpose built vehicle. However, if an appropriate loading and testing strategy is used it may be possible to test a suspension while it is fitted to any vehicle.

Reviewing the results of the analysis leads to some reasonably obvious strategies. The primary one is to, as much as possible, apply the load over the suspension being tested. This reduces the influence of the pitch resonance mode and consequently that of the other suspension. For example, if the Cg position is at 0.2 of the wheelbase length from the axis of the suspension being tested, the effect on the WFC of varying the stiffness and damping of the other suspension is very small (less than 4% variation over quite large changes in the suspension parameters). This loading pattern also reduces the influence of the pitch inertia to less than half what it was on the reference vehicle. In practical terms the range of pitch inertia values which are achievable while using the recommended loading strategy is probably quite small though it may be desirable to specify some minimum value.

The other major issue for a testing strategy is to adequately test suspension non-linearities. This requires that the amplitudes and frequencies of excitation used are selected appropriately. This issue was not addressed in this analysis but will be discussed in chapter 8.

7.3.6 Summary

A simple linear dynamic model was built to investigate the influence of the whole vehicle behaviour on the response of the suspension at one location on the vehicle. The purpose of the investigation was twofold. Firstly to investigate whether a suspension-based assessment of road friendliness was adequate or whether the influence of whole vehicle effects meant that only a whole vehicle assessment was meaningful. Secondly if a suspension test for road-friendliness were conducted with the suspension fitted to a real vehicle what strategies could be applied to minimise the influence of whole vehicle effects on the results.

The analysis showed that some whole vehicle effects did influence the response of the suspension significantly. Several of these effects, vehicle speed, Cg position and pitch inertia, vary with operating conditions. Thus a whole vehicle assessment in a particular loading configuration will give results which are specific to that configuration. These are no more useful than the results of a suspension-by-suspension assessment.

The stiffness and damping of the suspension at the other end of the vehicle did influence the response of the suspension under consideration. In fact, an improvement in the road friendliness of the other suspension can result in an increase in dynamic loads from the suspension being assessed. However, the dynamic loads of the vehicle as a whole reduced. Again there was no reason to expect that a combination of the assessments of the individual suspensions would not provide an appropriate ranking. A method for combining these assessments was proposed.

Finally a review of the findings were used to deduce a strategy for suspension testing which would minimise the influence of whole vehicle effects on the results. This strategy is straightforward. The vehicle should be loaded so that the Cg is as close as possible to the suspension being tested. A minimum value for the inertia ratio, which is a function of the pitch inertia, may need to be set. Some measurements of actual vehicles are needed to determine appropriate practical numerical values for these parameters.

7.4 Conclusions

Simple linear pitch-plane models have been developed and used to investigate aspects of the problem of assessing real heavy vehicle suspensions, which are, in general, non-linear and much more complex.

In the first application of these simple models, techniques were developed for tuning the model parameters so that the model responses matched the body modes of the actual vehicle under test as closely as possible. This "tuned" model was then used to convert the road profiles into the shaker excitations required to elicit the same response from the suspension. This process compensates for the case where there are insufficient shakers in the facility to excite every wheel on the vehicle. Although the match between the shaker excitations calculated in this way and those generated by the SYSCOMP software during actual testing was only fair, this appears to be largely because of the difficulty in relating the measured road profile information to the actual profile seen by the vehicle during road testing.

The second application of the simple models was to investigate the influence of whole vehicle factors on the response of a suspension fitted to the vehicle. The purpose of this was to determine whether it is reasonable to rate suspensions in isolation from the vehicle they are fitted to and if so to develop strategies for doing this assessment. It was found that while whole vehicle factors do have a significant influence on the wheel forces generated by a suspension these do not necessitate a whole vehicle approach to rating suspensions. A technique for combining individual suspension assessments to give a whole vehicle assessment is proposed. It was also shown that by loading the vehicle so that the load centre is as close as possible to the suspension being assessed and by keeping the pitch inertia of the load low it should be possible to minimise the influence of the vehicle factors on the rating obtained.

8 OTHER WORK ON SUSPENSION ASSESSMENT

8.1 Introduction

During the course of this research other researchers were also investigating suspension assessment for "road-friendliness" using a range of approaches. In this chapter some of most useful of this work is reviewed.

There are several philosophical differences in the approaches taken. The first major issue is the question of whether the assessment should rate individual suspensions or the whole vehicle-suspension configuration. Assessing individual suspensions and applying the resulting rating to vehicles fitted with that suspension in some way (for example, as proposed in section 7.3.5) has practical advantages. Only a fraction of the assessments would be required which would enable a more complete test to be undertaken. On the other hand, one of the main arguments for whole vehicle assessments is given by (Cebon 1993) who quotes an example from a simulation study where improving the suspension on one axle of the vehicle increased the dynamic wheel loads measured at the other axle. The mechanism for this is illustrated in section 7.3.4 and in particular Figure 65. In this case though the degradation in performance was less than the improvement at the suspension that was changed. Thus the overall effect was beneficial. An appropriate formula for combining individual suspension ratings into an overall vehicle rating would reflect this, at least to some extent. There are practical disadvantages to whole vehicle rating. The main one of these is the number of tests that might be required. As well as testing each vehicle model in service it would probably be necessary to test all configurations of combination vehicles. Even then there are issues of load distribution and the resulting changes in moments of inertia.

The other important issue of rating philosophy relates primarily to the complexity of testing required. Three types of approach have been used in the past. In order of increasing complexity these are design prescriptive, parameter specification and performance measures. Examples of each of these approaches appear in the following three sections.

8.2 The EC "Equivalent-to-Air" Test

In Europe there is considerable concern over the increasing volume of heavy vehicle traffic and the consequent effect on the infrastructure both in terms of congestion and wear. The potential for the increased use of "road-friendly" suspensions to reduce the level of infrastructure damage was recognised quite early on in some European countries and they instigated policies to encourage these suspensions. Initially higher axle load limits were allowed for axles fitted with air suspensions. This is a design prescriptive approach and was based on the fact that all the research studies on measuring dynamic wheel forces at that time had found that air suspensions generated lower dynamic loads.

As part of the European Union's efforts to harmonise heavy vehicle weight and dimension regulations, this policy was incorporated into the EC rules via a directive. As it was felt that the design prescriptive approach was unfair to other potentially well performing suspension types and would stifle innovation, the EC directive was broadened to allow other suspensions to qualify for this "road-friendly" axle load limit. Basically the directive states that for an axle to qualify for the allowance it must be fitted with dual tyres and have an air suspension, or the suspension must qualify as equivalent-to-air. The equivalent-to-air rating requires that the suspension have a natural frequency below 2 Hz, that it has a total damping greater than 20% of critical and that more than 50% of this damping is provided by a viscous damper. This is an example of the parameter specification approach. Three choices of test method for determining these parameters are specified. They are all based on measuring the suspension's response to a fairly large step perturbation. The first method involves the vehicle traversing, at crawl speed, a ramp of specified design that culminates in an 80mm drop. With the second method the chassis is lifted 80mm and released, while the third method requires the chassis to be pulled down until the axle load is 1.5 times the normal static value and then released. The directive is aimed at single drive axles and for these the three test methods give reasonably consistent results. However, for axle group this is not the case.

The EC directive is, in fact, a hybrid of two approaches to suspension assessment. Air suspensions qualify as "road-friendly" as of right so this part is design prescriptive. All other suspensions are required to meet certain parameter specifications using prescribed tests. There are flaws in this EC approach. It has been shown in this study and in others (Hahn 1985) that some air suspensions have poor damping and that without adequate damping air suspensions are no more "road-friendly" than typical steel suspensions. Yet according to the directive they automatically qualify for "road-friendly" status and are entitled to a higher axle load limit. The test procedures are all based on a fairly large perturbation. For suspensions with non-linear characteristics the response this test produces may be quite different to the response that would be generated by testing with a different perturbation amplitude. Because of the relatively high amplitude used in this test, the results probably give a better indication of rough road rather than smooth road performance

Although this approach has some merit in being easy to apply and in having results that are well understood, there is a concern over the extent to which the test results reflect in-service behaviour.

8.3 The Cambridge University WIM Mats

Cebon and his colleagues at Cambridge University (Cole and Cebon 1989) have developed a capacitive strip Weigh-in-Motion (WIM) sensor that they have embedded in tough polymer mats. Each mat was 1.2m long and contained three WIM sensors. These mats are claimed to be sufficiently flexible to follow the underlying pavement profile well at the wavelengths that excite heavy vehicle response. By placing a series of these mats end to end over a section of pavement it is possible to measure the dynamic wheel forces applied to that pavement by a vehicle as it passes over the mats.

Although a sufficiently large set of mats to undertake a wheel force assessment (say 50-100m test length) is relatively expensive, no vehicle instrumentation is required and the testing can be very quick. The method was particularly aimed at whole vehicle assessment. The initial concept was that the mats would be attached to the pavement by simply pinning the four corner of each mat. Thus the facility would be easily transportable and could be set up anywhere.

In practice there were difficulties. In order to get accurate results it was necessary to bond the mats to the pavement using an epoxy resin. This, of course, has a significant impact on the ease with which the system could be transported and set up. A trial installation of the mats has been used to investigate the spatial repeatability of dynamic wheel loads and some 3,000 vehicles were monitored with reasonable success.

These load mats represent the only attempt at a whole vehicle procedure, which has any credibility. They do, however, have limitations. The vehicle excitation is generated by the underlying pavement profile, which will vary over time and cannot be identical from test site to test site. Thus there is an issue of the portability and repeatability of the results. One possible way around this is the method proposed by (Woodrooffe et al. 1988) who suggested making an assessment at a number of pavement roughness levels and then interpolating the results to obtain a rating at some nominal pavement roughness. This would mean that several tests would be required for each vehicle assessment. There are several ways this could be achieved. One possibility is to construct a test pavement with particular roughness characteristics that span the range required and to place the mats over this. A second is to have mat installations on several pavements of different roughness and to run the vehicles over each of them.

Using the mats it would be feasible to construct an assessment facility to rate whole vehicle suspension configurations. This would be quite expensive to construct but individual tests could be undertaken relatively simply and quickly. Issues of portability of results and suitable loading configurations still have to be resolved.

8.4 The OECD DIVINE Tests

Element 3 of the OECD DIVINE (OECD 1997) programme was a series of road simulator tests undertaken at the National Research Council of Canada facility in Ottawa. This research element consisted of two parts. The first was to evaluate the ability of road simulators (multiple post servohydraulic shaker facilities) to simulate in-service heavy vehicle dynamics. The second part of the project was to determine methods for evaluating road-friendliness. (Woodrooffe 1997) has reported this work in depth.

The first part of the report (Woodrooffe 1997) reviews the EC drop tests and an experiment evaluating this test undertaken at the Transport Research Laboratory (TRL) in the UK. The TRL study found that of the suspensions they tested, only one air suspension met the road-friendly criteria. They also found that the perturbation amplitude had a significant influence on the results particularly for the damping ratio and recommended reducing this to 40mm.

The experimental programme for the OECD DIVINE project used a purpose-built semi-trailer, which consisted of a relatively long and light space frame structure connecting the kingpin to a sub-frame. This sub-frame was designed to allow a whole range of suspension and axle assemblies to be fitted. Loading was applied so that its centre of gravity was approximately in line with the axis of the suspension. This configuration was designed to minimise the external influences on the suspension's performance. Three different suspensions were fitted to the vehicle and a series of servohydraulic shaker trials were conducted using the NRC four-post road simulator.

- The tests were all conducted with identical signals applied to the left and right actuators on an axle. That is, no roll excitation. (It is worth noting that although the shaker facility at Industrial Research Limited has only two actuators, its vehicle support rig has been designed to apply each excitation to a whole axle. Thus it would be possible to repeat any or all of the NRC tests here). The tests undertaken included:

- simulated EC drop tests on one axle, with both axles in phase and with a phase lag corresponding to a 5 km/h traverse of the ramp,
- quasi-static load equalisation tests,
- sinusoidal sweep excitations with axles in-phase and anti-phase,
- road profile based excitations from measured data for three roughness levels,
- road profile based excitations generated numerically with appropriate spectral characteristics at three roughness levels.

The main findings of this project were as follows:
- Quasi-static load equalisation is important for road friendliness because of the power law relationships normally associated with pavement distress.
- With drop tests along the lines required by the EC directive it was found that the simultaneous drop of both axles in the group clearly differentiated between the suspension types in terms of the sprung mass response. The calculated level of damping was sensitive to drop height while the natural frequencies were relatively stable.
- Drop tests involving only a single axle of a group or with a phase lag corresponding to vehicle crawl speed did not produce results that were consistent across all suspension types and therefore are not suitable as an assessment procedure.
- Low amplitude (1mm) sinusoidal sweep excitations with both axles in phase clearly identified the main sprung and unsprung resonance modes of the suspension with their associated frequencies and damping.
- Anti-phase sinusoidal sweep excitations did not excite the sprung mass mode and so was not as usable.
- Larger amplitude sinusoidal sweep excitations did not produce as clear a distinction of suspension parameters.
- The road profile based excitations produced the most comprehensive picture of the relative suspension performance. Virtual road profiles are preferred because they are repeatable and transferable.
- The importance of damping for road-friendliness was again reinforced. All the air suspension tests were repeated with the dampers removed resulting in a significant deterioration in performance.

Based on these results it was suggested that three techniques could be used for suspension assessment. These were the simultaneous axle drop test, the in-phase sinusoidal sweep and the virtual road profile excitation. The capital cost of the equipment required for these three tests varies substantially, as does the level of detail generated in the results.

8.5 Summary and Results

A number of other researchers have also been investigating techniques for suspension assessment for "road-friendliness" over the period when this research was undertaken. Although there are some theoretical reasons for assessing the whole vehicle-suspension configuration rather than the suspension assembly this is not practicable as it would involve testing every vehicle model in the fleet and with combination vehicles possibly more than once. Of the assessment techniques investigated only the Cambridge University load mats show any promise for this scale of testing. Even then these have a number of drawbacks which limit their suitability. They are relatively expensive and need to be attached to an existing pavement. Thus the results will depend on the profile of the underlying pavement and are not readily transferable.

Assessing individual suspensions and applying the rating to vehicles fitted with these suspensions appears to be the most practical approach. Even though the design prescriptive approach such as rating all air suspensions as road-friendly has been used by some policymakers it is fundamentally flawed. Apart from stifling innovation from suspension designers, a number of counter examples have been found.

Parameter specification, while not ideal, offers more promise. The characteristics that make a suspension more road friendly are reasonably well understood: relatively soft springing (low natural frequency), a good level of viscous damping and good load equalisation. The difficulty lies in choosing an appropriate test method. The methods specified in the EC directive identify the sprung mass resonance modes of the suspension. However, the parameter values obtained depend on the amplitude of the excitation. Furthermore not all of the three methods are suitable for testing axle groups. The low amplitude in-phase sinusoidal sweep excitation proposed by Woodrooffe clearly identifies both the sprung mass and unsprung mass resonance modes together with frequency and damping values. The low amplitude (1mm) was chosen because it produced clear results for all three suspension types evaluated. However, part of the reason for this was that at these amplitudes the multi-leaf steel suspension is effectively locked and the vehicle bounces on its tyres, which are a relatively linear spring element. This reflects the situation on smooth roads. At greater amplitudes the steel spring overcomes this stiction and its apparent stiffness decreases and its performance improves but is relatively non-linear and difficult to analyse. This situation could become even more complicated with the advent of semi-active and active suspension systems, which are currently under development. Parameters such as stiffness and damping have no real meaning for these types of suspensions.

The final type of testing proposed was to excite the suspension on a servohydraulic shaker facility using virtual road profiles of various roughness levels. Although this form of testing requires relatively expensive equipment, it most closely replicates the in-service conditions and should provide the best rating values.

The issue of maintaining damper performance during the service life of the vehicles was raised in several of the studies discussed in this chapter. None of them presented any methodology for in-service testing of dampers.

9 CONCLUSIONS AND RECOMMENDATIONS

9.1 Summary of Findings

The aim of this study was to develop methods for assessing suspensions for road-friendliness based around the use of a small-scale (two-post) servohydraulic shaker facility. The first stage of the programme was to determine whether the shakers could be used to replicate the in-service wheel forces at the suspension being assessed. For this work a three axle tanker trailer with steel four-spring suspension was extensively instrumented and used in a series of road trials to measure the actual dynamic wheel forces generated on a range of test roads at a number of speeds. Two statistics for characterising the dynamic load sharing capability of the suspension were proposed. The values of this statistic for the suspension under test were about midway between no load sharing and perfect load sharing. Without testing a number of other suspensions it is not possible to use this for any comparative rating.

This vehicle was then returned to the laboratory and mounted on the servohydraulic shaker facility so that the shakers could excite two of its wheels. It was realised that because the excitations were only being applied to two wheels on one side of the vehicle, using the road profiles as input would not generate the correct response. Instead a sophisticated software algorithm was developed to calculate the excitations required to produce the same suspension deflection responses on the shakers as had previously been recorded during the road trials. The hypothesis was that if the suspension deflections were identical to those measured on the road, the corresponding reaction forces at the spring support would also be identical. The wheel forces are the sum of these reaction forces and inertia contributions from the unsprung masses outboard of the springs. The inertia corrections can be calculated using the measured accelerations of the unsprung masses. If this hypothesis is valid a suspension assessment could be undertaken by measuring the suspension deflections and unsprung mass accelerations during a road test and then using the shaker facility to replicate the suspension deflections and measure the wheel forces on the platforms and correcting these for the inertia differences.

The software system for calculating the shaker excitations underwent considerable development before a fast, reliable and stable algorithm was achieved. At this point the match between on-road response and shaker response for the wheels under test was good and, in fact, superior to that achieved by other researchers (Gyenes and Mitchell 1996; Stanzel and Preston-Thomas 1996) using large scale shaker facilities with road profile based excitations.

To ensure that these results were not specific to the suspension configuration used the vehicle was fitted with an air suspension and the road and shaker testing repeated. In addition to the water loading used in the first series of tests an alternative load distribution was applied using lead ingots to investigate the effects of changes in pitch inertia.

The road test results were rather surprising. In general the air suspension did not produce lower dynamic wheel loads than the steel suspension had. When the vehicle was subsequently mounted on the shaker facility and a modal analysis was undertaken it was found that although the air suspension was significantly softer than the steel and thus had sprung mass resonance modes with lower natural frequencies, the damping of these modes was very low ($<6\%$). This provided a graphic illustration of the critical role of damping in the performance of softer suspensions and particularly air. As the suspension was brand new this result was of some concern and so the suspension was checked out very carefully. However, there was nothing wrong with it other than its basic design, which was similar to other suspensions. It had been fitted totally in accordance with the manufacturer's specifications and the dampers were in good working order.

The effect on the level of dynamic loading of changing from water loading to lead loading was minimal. For some of the tests the dynamic loads were higher while for others they were lower. With water loading there was a side-to-side liquid slosh mode at about 0.5 Hz which disappeared in the lead loading case. However, this mode was not a significant contributor to dynamic wheel loads.

The shaker trials on this vehicle-suspension configuration identified a major flaw in the assumptions underlying this approach to suspension assessment. Although the shaker control software was able to match the suspension deflections measured on the road very well (at least for the smoothest test site), the corresponding vehicle motions and wheel forces were substantially greater. For the rougher test sites, the motions were so violent the tests were not completed for safety reasons. After considerable analysis and further testing it was found that the auxiliary roll stiffness inherent in this type of suspension design was the source of these additional large wheel force components. To overcome this the shaker support rig was modified so that each actuator excited an entire axle rather than just one wheel. This eliminated all roll behaviour from the shaker trials. With this change a good match between the on-road and laboratory responses was achieved for this suspension.

Shaker trials were undertaken on a SLAVE unit from the CAPTIF in order to link the results from this research to a related program investigating the effects of dynamic loading on pavement performance and life. Although this vehicle is a simpler dynamic system than a complete truck the match between shaker trials and in-service measurements was not as good as expected. The reason for this was a structural resonance of the unit, which was present in the shaker tests but not during the in-service tests. This highlights some of the potential problems with using simplified instrumentation systems to estimate wheel forces. The actual suspension responses as characterised by the deflections matched quite well.

Simple linear dynamic models were developed and used to improve the assessment technique. The first application of the simple model involved tuning the model parameters to match the vehicle's characteristics as closely as possible and then using this tuned model to calculate how the road profile inputs would need to be modified in order to obtain the same suspension responses on the shaker. The results of this were very encouraging. Applying the modifications to the measured road profiles and comparing these calculated excitations with those generated by the shaker software showed a fair rather than a good match. However, the primary source of the discrepancies was probably the difficulties in synchronising the measured profile with the actual excitation applied to the vehicle. The longitudinal and lateral position of the test vehicle and the lateral position of the profiler were not known to sufficient accuracy for this. Using the linear model to back calculate the excitation from the measured response showed a very good match indicating the potential usefulness of this approach. The second use of the linear models was to investigate the extent to which the vehicle parameters would influence the rating of a suspension fitted to that vehicle. This was then extended to determine whether, with a suitable testing strategy, a vehicle-independent suspension rating could be obtained by testing a suspension while fitted to an actual vehicle. Again the results of this were very encouraging.

Finally the report reviews some of the other work on suspension assessment which was undertaken by other researchers during the course of this project. Although most of the approaches were somewhat different to this work their findings are useful in developing recommendations for a suspension assessment procedure.

9.2 Recommendations for Suspension Assessment

From the results of this research a number of recommendations for suspension assessment procedures can be made.

The design prescriptive approach, for example rating all air suspensions as road-friendly, is flawed. The air suspension fitted to the test trailer in this research was no more road-friendly than the steel suspension it replaced. This was a common widely used suspension in good condition. If a suspension assessment regime is introduced with incentives for road-friendliness, then no suspension should qualify without being tested.

One of the fundamental issues is whether the assessment should be undertaken on the whole vehicle or just on the suspension. For practical reasons it appears to be more appropriate to assess suspensions and rate them independently and then to combine these in some way to give a vehicle rating. This is primarily because of the potentially large number of tests that would be required if all vehicle models had to be tested. It should be noted that the shaker testing procedures developed in this research can be used for both whole vehicle assessment and suspension only assessment.

Testing to measure suspension parameters is the simplest valid approach to rating suspensions but some care needs to be taken in defining the test procedures. The simplest of these approaches are the EC tests, which measure the natural frequency and damping of the fundamental body bounce of the vehicle. As shown by (Woodrooffe 1997) for axle groups the EC test needs to be done with all axles in-phase. This eliminates the crawl speed traverse of the defined ramp as a test method for axle groups. Woodrooffe also argues against the pull-down test option because the first motion of the suspension is rebound and for most shock absorbers the rebound damping is much higher than the bump damping. Thus this test will tend to overestimate the suspension's damping. This leaves the drop test as the only suitable test for axle groups. There are two further issues related to this test. The first is the drop height. This is an issue because with non-linear suspension behaviour it affects the results. The most satisfactory solution to this is to require the test to be undertaken at several heights, for example 20mm, 40mm and 80mm and require all result to meet the standards. The second issue relates to the vehicle independence of the results. The modelling work in section 7.3.4 and (de Pont 1996b) shows that the results can be influenced by vehicle parameters. Therefore it is recommended that the loading strategies suggested in section 7.3.5 be required. That is, the load should be concentrated as much as possible over the suspension being tested and should have low pitch inertia. Realistic numerical limits for these criteria have yet to be established.

This drop test can be undertaken using the shaker facility or a purpose built mechanical system could be designed. The suspension response can be measured using wheel force transducers in the rig or alternatively with suspension displacement transducers fitted to the vehicle. The suspension parameters obtained would be the natural frequency and damping of the suspension so the question is what are reasonable limits for road friendliness. The EC specifies 2 Hz for the natural frequency and Woodrooffe recommends 1.5 Hz. There is a degree of arbitrariness about these values. In general lower is better though with diminishing returns. Both sources suggest 20% as the minimum damping limit with at least half this damping provided by the viscous damper (to check this the tests need to be repeated without the dampers connected). The OECD DIVINE report (OECD 1997) shows a graph based on vehicle simulations relating dynamic wheel forces to damping. This shows that reducing damping below 10% leads to significant increases in dynamic wheel forces particularly on rougher roads while increasing damping above 20% leads to only minimal improvements. One option here would be to require a damping level of at least 20% on new suspensions and to set limits for in-service damper deterioration that ensure that the minimum actual in-service limit is greater than say 15%.

An alternative test for measuring suspension parameters is the in-phase low amplitude (1mm) sinusoidal sweep proposed by (Woodrooffe 1997). This test clearly identifies the natural frequency and damping of both the sprung mass and unsprung mass modes of the suspension and thus provides a more complete characterisation of the suspension than the EC drop test. Woodrooffe conducted his trials of this test using a purpose built trailer, which was designed to minimise the influence of whole vehicle effects on the results. From the modelling work described in section 7.3.5 and (de Pont 1996b) we would expect to be able to achieve the same effect by controlling the loading of the vehicle. This test could be undertaken on the two post shaker facility or a purpose-designed rig could be built. As this method uses a constant amplitude excitation this rig design would not be complicated. The result of this test is a trace, which is effectively the frequency response function of the suspension, clearly identifying the two natural frequency values and enabling the two corresponding damping values to be calculated. However, Woodrooffe makes no recommendations as to appropriate limits on these values for road-friendliness rating. If this test were to be adopted some preliminary testing on a range of vehicles would be required to establish benchmarks for the results. The excitation amplitude of this test will not be sufficient to overcome the stiction in most traditional multi-leaf steel suspensions, but this is a fair reflection of the situation on most of the highway network.

The final and most comprehensive approach to assessing suspensions is a performance-based test. This could be undertaken using the servohydraulic shaker facility to apply virtual road profiles to the suspension under test. Provided the vehicle loading were applied in accordance with the recommendations of section 7.3.5 and (de Pont 1996b) the profiles would not have to be modified. The virtual road profiles could be applied at several roughness levels and at different vehicle speeds and the corresponding dynamic wheel loads (DLC values) measured. From these there are many criteria available for road-friendliness rating. If DLC were plotted against road roughness and vehicle speed, one could set a limit plane that the measured values must stay below. Alternatively a series of planes could be set so that there was a range of ratings. A linear regression analysis could be applied and used to calculate a nominal DLC value at a specific speed and roughness, which would be the suspension's rating.

All the tests described so far have tested the dynamic performance of the suspension in some way. However, there are some other factors that come into play when considering road-friendliness. At the start of this section the argument for suspension testing rather than whole vehicle testing was presented. If this approach is taken there is a need to consider how individual suspension ratings might be combined to give an overall vehicle rating. In section 7.3.5 a method for combining individual suspension ratings was proposed. Although this method has not been evaluated in any depth it has some attractive characteristics, primarily in that it encourages all the suspensions on a vehicle to be of similar performance, which the models suggest is desirable. However, to apply this method each suspension needs to have a numerical rating preferably based on DLC. The EC drop test does not really provide this - it is designed to give a yes-no rating. One option here would simply be to require that for a vehicle to qualify as road-friendly all suspensions on the vehicle must qualify (possibly with the exception of steer axles). Similarly the sinusoidal sweep test does not result in a single rating number that can be used. However, there are options for condensing the results into a single statistic. These still need investigation. The results of testing using virtual road profile excitations on the shakers could easily be converted to a single number rating for use in a whole vehicle formula.

As well as the dynamic response of the suspension, its road-friendliness also depends on its quasi-static load equalisation ability. If this is poor one axle of a group may consistently carry more than its share of the load. With the generally accepted power law relationships between axle load and pavement wear this can lead to significant additional pavement wear. It should be noted that this characteristic is independent of dynamic performance. Walking beam suspensions which have generally poor dynamic wheel force behaviour have very good load equalisation characteristics. The OECD DIVINE report (OECD 1997) suggests a load equalisation rate of 0.3kN/mm as the criterion for road-friendliness. Measuring the quasi-static load equalisation is very straightforward on the shaker rig but would require a rig with some instrumentation if it were to be undertaken elsewhere.

The final issue is the question of in-service monitoring of damper performance. Many researchers have identified the critical role of dampers in the dynamic performance of softer more road-friendly suspensions and hence the need for monitoring to ensure that a suspension which has been rated as road-friendly stays road-friendly. How this in-service monitoring might be accomplished has not been addressed to any extent. If a sinusoidal sweep test device were built it could be used for in-service testing. However, the suspension assessment process assumed only testing one of each suspension type while in-service monitoring requires that every vehicle fitted with the suspension is tested, so many more testing facilities would be needed. Furthermore, this testing procedure averages the behaviour of all the dampers in the axle group. If only one damper in the set was faulty, it is be no means certain that this test would identify it. An alternative approach is to test the dampers off the vehicle. Some shock absorber suppliers have testing equipment, which produce force displacement curves for particular test conditions (typically a 50mm sinusoidal excitation at 80 cycles/min). If these curves were recorded for the dampers when new, then regular in-service tests could be used to ensure that their performance stayed within some tolerance of the original values. Vehicles with road-friendliness ratings would be expected to carry current test certificates for all their shock absorbers. Currently the testing cost is approximately $10 per shock absorber, although this would probably reduce as testing became commonplace. For a fleet operator with a number of vehicles fitted with the same suspensions this process would be manageable as shock absorbers could be exchanged for freshly certificated ones as part of the maintenance schedule and then tested and approved for use on the next vehicle. However, for single vehicle operators this may be more expensive as it may require owning two complete sets of shock absorbers. Of course having well-maintained shock absorbers has benefits to overall vehicle safety and handling, vehicle wear and cargo damage as well as to minimising road wear.

Overall three approaches to assessing suspensions have been proposed. They vary in complexity and cost and also in the level of detail in the information they provide. Before implementation of any of these schemes, trial tests should be undertaken to examine the range of parameter values that would result in order to set realistic benchmarks for what constitutes road-friendliness. The only test where some guidelines exist is the EC drop test.

All three techniques are practical and consistent with current international opinion. From an economic viewpoint, the servohydraulic shaker facility could be used for all three methods and in this case there is relatively little extra cost in conducting all three tests at once. Although cheaper purpose-designed facilities could be used for the two simpler procedures, these do not exist currently. A test is likely to cost several thousand dollars, but only one example of each suspension design would need to be tested and so the total cost would not be excessive. Assuming that perhaps fifty suspensions would be tested the total cost would be of the order of \$100,000 - \$150,000 which would be spread over all vehicles using those suspensions. Compared to the likely associated savings in pavement wear of at least \$16M p.a. (see section 1.3) this is a very small amount. However, the proposed regime of in-service damper testing to ensure that suspensions maintain their performance does incur additional costs. Although each test would be cheap, these tests would be required for each damper on every vehicle with a "road friendliness" rating on a regular basis. These costs need not all be attributed to the "road-friendliness" rating procedure as there are safety benefits associated with ensuring that dampers are maintained. Nevertheless the design of any incentive scheme would need to take these costs into account and canvass operator attitudes to determine its viability.

REFERENCES

Anon. (1962). The AASHO Road Test - Report 5: Pavement Research. *Special Report 61E*, Highway Research Board, Washington, D.C.

Barber, N. F. (1969). *Water waves*, Wykeham, London.

Cebon, D. (1987). Examination of the road damage caused by three articulated vehicles. *10th IAVSD Symposium on the Dynamics of Vehicles on Roads and on Tracks*, Prague, 65-76.

Cebon, D. (1993). Interaction between heavy vehicles and roads, SAE SP-951 930001.

Cole, D. J., and Cebon, D. (1989). Simulation and measurement of dynamic tyre forces. *2nd International Symposium on Heavy Vehicle Weights and Dimensions*, Kelowna, British Columbia.

Council of the European Communities. (1992). Annex III of the Council Directive 92/7/EEC amending Directive 85/3/EEC on the weights, dimensions and certain technical characteristics of certain road vehicle, Council of European Communities, Brussels.

de Pont, J. J. (1994). Road Profile Characterisation. *Research report 29*, Transit New Zealand, Wellington.

de Pont, J. J. (1996a). Experiences with simulating on-road heavy vehicle suspension behaviour using servo-hydraulics. *Heavy Vehicle Systems, Special Series, Int J of Vehicle Design*, 3(1-4), 127-139.

de Pont, J. J. (1996b). Suspensions or Whole Vehicles? Rating Road-Friendliness. *Fourth Engineering Foundation Conference on Vehicle-Infrastructure Interaction*, San Diego.

de Pont, J. J., Thakur, K. P., and Costache, M. B. (1995). Simulating In-Service Heavy Vehicle Suspension Dynamics. *4th International Symposium on Heavy Vehicle Weights and Dimensions*, Ann Arbor, Michigan.

Eisenmann, J. (1975). Dynamic wheel load fluctuations - road stress. *Strasse und Autobahn*, 4, 127-128.

Ewins, D. J. (1984). *Modal Testing: Theory and Practice*, Research Studies Press, Letchworth, Herts.

Fancher, P. S., Ervin, R. D., Winkler, C. B., and Gillespie, T. D. (1986). A factbook of the mechanical properties of the components for single-unit and articulated heavy trucks. *Report DOT HS 807 125*, US Dept of Transportation, Washington, DC.

Gyenes, L., and Mitchell, C. G. B. (1996). Measuring dynamic loads for heavy vehicle suspensions using a road simulator. *Heavy Vehicle Systems, Special Series, Int J of Vehicle Design*, 3(1-4), 191-221.

Gyenes, L., Mitchell, C. G. B., and Phillips, S.D. (1992). Dynamic pavement loads and tests of road - friendliness for heavy vehicle suspensions. *3rd International Symposium on Heavy Vehicle Weights and Dimensions*, Cambridge, UK.

Hahn, W. D. (1985). Effects of commercial vehicle design on road stress - vehicle research results, Institut fur Kruftfahrwesen, Universitat Hannover, Hannover.

Heath, A. N. (1987). Heavy Vehicle Design Affecting Road Loading. *ARRB/FORS Symposium on Heavy Vehicle Suspension Characteristics*, Canberra, 251-270.

Ministry of Transport. (1997). Options for the future. *Land transport pricing study.* Discussion document, New Zealand Ministry of Transport, Wellington.

Mitchell, C. G. B., and Gyenes, L. (1989). Dynamic pavement loads measured for a variety of truck suspensions. *2nd International Symposium on Heavy Vehicle Weights and Dimension*, Kelowna, British Columbia.

O'Connor, T., O'Brien, E. J., and Jacob, B. (1996). An experimental investigation of spatial repeatability. *Fourth Engineering Foundation Conference on Vehicle-Infrastructure Interaction*, San Diego.

OECD, R. T. R. P. (1992). Dynamic loading of pavements. *OECD Report IRRD No 847714*, OECD, Paris.

OECD, R. T. R. P. (1997). *Dynamic interaction of heavy vehicles with roads and bridges. OECD Report Draft final report*, OECD, Paris.

Prem, H. (1987). A laser-based highway-speed road profile measuring system. *10th IAVSD Symposium on the Dynamics of Vehicles on Roads and on Tracks*, Prague.

Stanzel, M., and Preston-Thomas, J. (1996). Road simulator testing. *Fourth Engineering Foundation Conference on Vehicle-Infrastructure Interaction*, San Diego.

Sweatman, P. F. (1983). A study of dynamic wheel forces in axle group suspensions of heavy vehicles. *Special Report 27*, Australian Road Research Board, Melbourne.

Wester, E. (1989). Syscomp - system response compensation program manual. Software manual, Auckland Industrial Development Division, DSIR, Auckland.

Woodrooffe, J. H. F. (1997). Heavy truck suspensions: Methods for evaluating road-friendliness. *Research Report C-96-4*, NRC & Roaduser Research, Ottawa.

Woodrooffe, J. H. F., LeBlanc, P. A., and LePiane, K. R. (1986). Effects of suspension variations on the dynamic wheel loads of a heavy articulated highway vehicle. *Vehicle Weights and Dimensions Study Vol 11*, Roads and Transportation Association of Canada.

Woodrooffe, J. H. F., LeBlanc, P. A., and Papagiannakis, A. T. (1988). Suspension dynamics - experimental findings and regulatory implications. *SAE technical paper series 881847*, Society of Automotive Engineers, Indianapolis, Indiana.