Using servo-hydraulics to assess heavy vehicle suspensions for road wear

J. DE PONT, BSc, BE, ME, PhD, MSAE, Research Engineer, DSIR Industrial Development, Auckland, New Zealand

It is well established that different heavy vehicle suspensions generate varying levels of dynamic wheel force under the same vehicle speed and road roughness conditions. Although the quantitative relationship between dynamic wheel forces and pavement wear is still being debated it is generally accepted that higher levels of dynamic loading generate increased pavement wear. Thus, there are benefits to pavement managers in encouraging the use of suspensions which reduce the dynamic loads generated. In order to do this it is necessary to identify and quantify "road-friendly" suspensions. This paper describes an experimental programme aimed at developing a technique for assessing suspensions using a two post servo-hydraulic shaker facility.

1. INTRODUCTION

The major result of the AASHO road test of the late 1950s was the now well known "fourth power law" which essentially says that the wear to a pavement caused by an axle passing over it is proportional to the fourth power of the static axle load. Although this relationship has been the subject of considerable debate over the years, it is the basis of most pavement design codes. The AASHO test effectively averaged dynamic effects. Normal road vehicles were used and so dynamic loadings occurred but no attempt was made to account for them. From the late 1960s onwards a number of researchers began to investigate dynamic wheel forces. Whittemore et al (ref. 1) developed a transducer for measuring dynamic wheel forces and compared it with strain gauging the vehicle axles. Both techniques were found to give satisfactory results. Sweatman (ref. 2) used the wheel force transducer to measure the dynamic wheel forces generated by nine different vehicle-suspension configurations at various speeds over a range of road roughnesses. To characterise the dynamic wheel force behaviour, Sweatman uses a measure called the dynamic load coefficient (dlc) which is defined as

$$dlc = \frac{\textit{standard deviation of wheel force}}{\textit{mean wheel load}}$$

There is some variation in interpretation of the mean wheel load term in the denominator. Sometimes the mean for the test is used and sometimes the global mean or static load is used. The differences between these are small and reflect limitations in the load-sharing of the suspension as well as factors such as road camber. Sweatman found that the dlc varied significantly from suspension to suspension, and that it was dependent on vehicle speed and road roughness. Woodrooffe et al (ref. 3), as part of the Roads and Transportation Association of Canada (RTAC) Vehicle Weights and Dimensions Study (which spawned these symposia) used strain gauged axles and accelerometers to compare the dynamic behaviour of a number of different suspensions. As with Sweatman's work significant differences in dlc values for different suspensions were observed. Dynamic wheel force measurements have also been undertaken in other research programmes, for example, in Sweden (ref. 4), Germany (ref. 5) and the United Kingdom (ref. 6-7). Although it is not possible to compare the results of these tests directly because conditions were not identical, they are consistent. Thus, it has been shown conclusively that there are significant differences in the level of dynamic wheel forces generated by different suspensions. A number of relationships between dynamic wheel forces and pavement wear have been postulated but they are still subject to considerable debate and none is universally accepted.

Taking a very simplistic view, it could be argued that the dynamic wheel forces are a variation about the static load and the effects of high loadings are counterbalanced by the low loadings, thus having no effect. However, extrapolating the AASHO "fourth power law" to apply to dynamic loads implies that the additional wear effects of the above average loads are greater than the reduced wear of below average loads giving a wear attributable to the dynamics. Assuming that the distribution of dynamic wheel forces is random and Gaussian and that the "fourth power law" relationship between wheel forces and pavement wear applies, Eisenmann (ref. 8) developed a road stress factor

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

where P_{stat} = *mean axle load*

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

K = *constant*

Heavy vehicles and roads: technology, safety and policy. Thomas Telford, London, 1992.

A dynamic road stress factor which represents the damage due to dynamic effects can then be defined as

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

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

Although the dynamic wheel force measurements have been shown to be approximately Gaussian, the assumption of randomness is clearly not correct. Experimentally measured wheel forces show a high degree of repeatability, and, in fact, because heavy vehicle suspension characteristics are, to some extent, similar, the dynamic wheel forces of different suspensions over the same pavements at the same speeds exhibit similarities. To reflect this lack of randomness Sweatman (ref. 2) proposes an alternative dynamic road stress factor based on the 95th percentile of the wheel force distribution. His argument is essentially that the peak dynamic wheel forces will tend to occur at roughly the same locations for all vehicles and that these peak loads cause the damage. Assuming a Gaussian distribution, the 95th percentile impact factor is

$$IF_{95th} = 1 + 1.645dlc$$

and the associated road stress factor, based on the fourth power law is

$$\Phi_{95th} = (IF_{95th})^4$$

Both these road stress factors have similar characteristics in that with zero dlcs they reduce to unity, and they both give the same ranking of suspensions. However, the relative magnitudes of these factors vary substantially. For example, in Sweatman's work the best and worst suspensions have dlcs of 0.13 and 0.27 respectively under arbitrary identical test conditions. The corresponding Eisenmann dynamic road stress factors are 1.10 and 1.46 while the 95th percentile road stress factors are 2.14 and 4.37. That is, the Eisenmann factor implies that the worst suspension does approximately 30% more damage than the best, whereas the 95th percentile factor implies over 100% more damage. Cebon (ref. 9) argues that using the temporal wheel force distribution is flawed as it is the spatial distribution of wheel forces along the pavement that determine wear. He proposes that the spatial distributions for each of a vehicle's wheels in a given wheel path should be accumulated to give a whole-vehicle pavement loading distribution. Based on this distribution he developed five pavement wear criteria based on models of pavement response. These criteria vary in complexity and also in relative magnitude, although the ranking of suspensions tends to be unaffected. Cebon tested his criteria using computer simulation models of heavy vehicles (validated with some measured data). For a particular set of operating conditions (ie a vehicle with steel four spring suspension at highway speed on a "good" road) the values calculated for his different pavement wear factors varied from 1.1 to in excess of 4. As with the previous road wear factors, unity represents the zero dynamics situation.

Reviewing this from the viewpoint of a pavement manager then, it can be seen that, some suspensions have better performance in terms of pavement wear than others and thus the use of "road-friendly" suspensions should be encouraged. In order to operate a sensible policy to do this it is necessary, firstly, to be able to assess the performance of a suspension or vehicle/suspension combination in terms of the dynamic wheel forces it generates and secondly, to be able to relate this assessment to its contribution to pavement wear.

The work described in this paper addresses itself principally to the first problem, assessing the dynamic wheel force behaviour of a suspension. The second issue of converting this assessment to a pavement wear factor is primarily of importance in developing incentives for "road-friendly" suspensions. However, the two are closely linked and so in developing an assessment procedure, its end use must be kept in mind.

2. CURRENT STATE OF SUSPENSION ASSESSMENT

There are a number of requirements for an ideal procedure for assessing the dynamic wheel force behaviour of a heavy vehicle suspension. These may or may not be achievable in practice. Clearly the most important requirement is that the assessment reflect the contribution of the suspension to pavement wear. The procedure must be economically viable, and relatively straightforward to undertake. Also it should be repeatable and transferrable.

As described in the previous section a number of researchers have investigated the dynamic wheel force behaviour of heavy vehicle suspensions. The techniques used in these studies for measuring wheel forces are generally not well suited for routine use in an assessment procedure. The wheel force transducers are relatively expensive and require hub adaptors for each axle type, while the strain gauging approach is too time-consuming. It has also been shown that the dynamic wheel force behaviour is dependent on vehicle speed and road roughness. This variability should, if possible, be removed to achieve repeatability and transferability.

EC regulations allow a drive axle with twin tyres and air suspension to operate at an axle load of 11.5 tonnes instead of 10.5 tonnes for an axle not so equipped. The basis for this regulation is that all research to date has shown that air suspensions are relatively "road-friendly". This approach is clearly unsatisfactory as it stifles technical development and penalises any suspension which is not air regardless of performance. To resolve this EC directive 90/486/EEC on heavy vehicle characteristics defines a simple suspension assessment process to qualify suspensions as "equivalent-to-air". "Equivalent-to-air" is defined as having a natural frequency of below 2 Hz and a damping ratio of at least 20% of critical with less than 50% of the damping being frictional. To test for compliance with these criteria three alternative procedures are defined. These are:-

i) a pull down test, where the chassis is pulled down till the axle load is 1.5 times its static value and then released suddenly. The resulting oscillation is monitored with a displacement transducer and analyzed to find the frequency and damping.
ii) a drop test, where the chassis is lifted by 80 mm and then dropped suddenly. As for the previous test the resulting oscillation is analyzed.
iii) a bump test, where the vehicle is driven at low speed over a ramp with an 80 mm drop at the end. The resulting oscillation is analyzed.

Tests at MAN trucks (ref. 10) have shown that these three tests give similar results for a single drive axle but that there are significant differences with tandem axle configurations. This assessment process has advantages in that it is relatively straightforward and has good repeatability and transferability. Its drawbacks are that the link between the test criteria and pavement wear is tenuous and that it cannot be applied successfully to tandem axle configurations.

Currently there are three research programmes underway investigating assessment procedures. Cebon and his group at Cambridge University (ref. 11) have developed load sensing mats for this purpose. These mats are made of a flexible material and contain capacitive load measuring transducers. By attaching these mats sequentially along a test section of pavement and monitoring the transducers as the vehicle passes over it the spatial distribution of wheel forces can be measured. The mats are sufficiently flexible to allow the components of the pavement profile which excite the dynamic response of the vehicle to be transmitted. This approach fits in with Cebon's previous work on pavement wear criteria (ref. 9) which are based on the spatial distribution of the wheel forces and the cumulative effects of all the axles on a vehicle. The original intent was that the system would be highly transportable. However, in order to get reliable wheel force measurements, it has been found necessary to bond the mats to the pavement thus reducing the portability. The method is easy to apply but it does rely on actual road profiles to provide the vehicle excitation which may limit the repeatability and transferability of the results. Currently a prototype of this system is being trialed.

Woodrooffe and his team at National Research Council, Canada are developing a generic test trailer for suspension assessment. This trailer has a long light drawbar which is intended to minimize the influence of the towing vehicle on its response. It has been designed to allow any suspension to be fitted and has purpose built strain gauged axles for measuring the wheel forces. It is intended that measurements be undertaken on ordinary pavements at normal vehicle speeds. The concept is that suspensions would receive a rating in a vehicle independent environment. This contrasts with Cebon's approach of whole vehicle assessment. It has the same difficulties in that using actual pavement for testing limits the repeatability and transferability of results.

Sweatman (ref. 2) showed that the dlc for a particular suspension was proportional to $VR^{0.5}$ where V is the vehicle speed in km/hr and R is the road roughness in NAASRA counts/km. Thus he proposed a test condition of $VR^{0.5} = 850$ which corresponds to highway speed on a moderately rough road. Woodrooffe (ref. 12) suggests that rather than trying to match this type of test condition exactly, a number of tests should be undertaken with conditions close to the specified test condition. A linear regression fit can be applied to the results and then evaluated at the test condition value. This would reduce the variability of the procedure. This approach is more appropriate to the test trailer assessment method than the load sensing mats.

A complication is this type of test condition specification is that there are different road roughness measures in use in different countries and there is no exact equivalence between them. More importantly they are generally based on passenger car response to the road profile and thus it is conceivable that two pavements of equal road roughness, by whatever measure, might evoke different heavy vehicle responses.

The third research programme looking at suspension assessment is being carried out at the Department of Scientific and Industrial Research (DSIR) Industrial Development Division in Auckland, New Zealand and is the subject of this paper. It is based around using a relatively small general purpose servohydraulic shaker facility to simulate on-road behaviour.

3. USING SERVOHYDRAULICS FOR ASSESSMENT

3.1 Basic concept

Initially the idea of this assessment technique was that the vehicle would be instrumented with linear voltage differential transformers (LVDTs) between the axle and the chassis and with axle mounted accelerometers adjacent to the suspension units being assessed. It was felt that these transducers could be fitted relatively quickly and easily. The vehicle would then be submitted to a series of road tests and the vehicle response as measured by the transducers would be recorded. Following this the vehicle would return to the laboratory and mounted on the servohydraulic shakers which would then be used to replicate the measured on-road responses for the wheels being excited. Load cells in the vehicle support platform would be used to measure the resulting wheel forces. With this scenario there is still a requirement for a road test component to the test procedure and consequently it has the same repeatability and transferability limitations as the other approaches. However, there is some scope for generating the servohydraulic excitation signals theoretically and eliminating the road test requirement. This would significantly enhance the procedure's repeatability and transferability.

3.2 The servohydraulic facility
The servohydraulic facility used for these trials consists of two rams, each with a rated capacity of 5 tonnes, a displacement range of ± 125 mm and a frequency response of about 0 - 100 Hz (although the displacement range achievable at higher frequencies is limited).

The rams were built locally to a design obtained from National Engineering Laboratory (NEL) in Scotland. The servohydraulic system was designed in-house using commercially available valves, accumulators etc. The pump driving the system operates at a working pressure of 210 bar and can supply oil at 120 l/min. Bladder accumulators on the supply line enable an instantaneous oil flow rate of 180 l/min. Control of the rams is provided by two Instron 2165 series controllers with both force and position control being possible. Although the controllers can generate the standard waveforms (square, triangle and sine) internally, most applications use externally supplied waveforms. A Hewlett Packard HP3852 data acquisition unit with the arbitrary waveform generator module is used for this purpose.

The facility was developed for general purpose testing work and has been used for a wide variety of applications. For the vehicle testing programme two ram support platforms (figure 1) were designed and built.

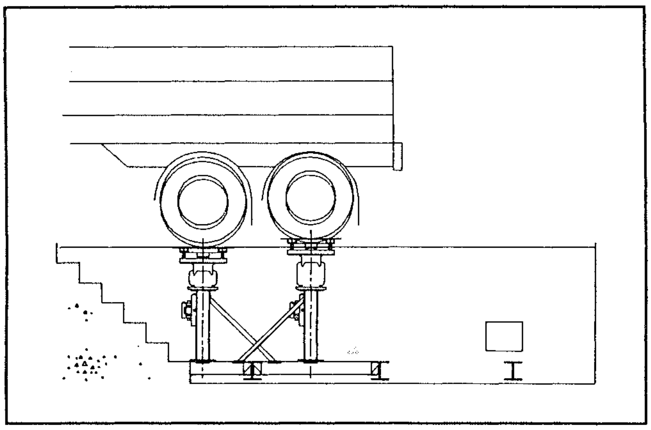

Figure 1. Vehicle support for shaker rig.

The rig is mounted in a pit so that the vehicle can be rolled onto it and all wheels are level. Airbags are used to support the static load of the vehicle. Thus the rams are used solely to provide the dynamic forces. Without this feature the 5 tonne ram capacity would not be sufficient to cope with a fully loaded vehicle.

Remote parameter compensation (RPC) software was developed locally in association with the testing programme and will be described in more detail in the following sections. The purpose of this software is to adjust the ram excitation signals to compensate for the system's effects on the response.

3.2 Road tests
To validate this assessment technique a series of road trials using an extensively instrumented liquid tanker trailer were conducted. This same vehicle was then used for the servohydraulic shaker tests.

Each wheel of the trailer was instrumented with:- strain gauges to measure the shear strain in the axle, an axle mounted accelerometer, an LVDT to measure the displacement between the axle and the chassis, and a chassis mounted accelerometer. Additionally a trailing "fifth" wheel was used to monitor the vehicle speed and distance travelled. As the trailer has six wheels this is a total of 26 channels of data. These 26 channels were monitored using a HP3852 data acquisition unit sampling at 100 Hz per channel.

Five test sites were chosen encompassing a range of roughnesses. Sites were selected to be relatively level and straight for approximately a kilometre with sufficient clear road before and after to enable the test to be done at constant highway speed. The purpose of these criteria was to, as much as possible, ensure that the vertical wheel forces were generated by the vehicle-pavement interaction alone. Each of the sites was traversed at three different speeds, with one of the runs being repeated as a check.

The speeds chosen were site dependent. Three of the sites were rural roads and the speeds used were 75, 80 and 85 km/hr. One of the sites was a section of motorway and an additional run at 100 km/hr was done while the final site was an urban road which was tested at 45, 50 and 55 km/hr. The reason for selecting such a narrow range of speeds was to investigate Heath's hypothesis (ref. 13) that the interaction of a vehicle's resonances with the pavement profile could result in peaks in dynamic response at specific speeds and so the relationship between dlc and speed would not be linear as found by Sweatman (ref. 2). No evidence was found to support Heath's hypothesis. However, this investigation was secondary to the principal aim of the project and so the tests were not designed specifically to resolve this issue. Thus the hypothesis was not disproved either.

To complete the picture on these tests, the Australian Road Research Board's (ARRB) laser profilometer was used to measured the pavement profile at all the test sites and to calculate the roughnesses. Average site roughness ranged from 32-98 NAASRA counts/km, though when evaluated in 200 m sections the range increases to 27-190 NAASRA counts/km.

Using the strain gauge signals and the axle accelerometers to adjust for the inertial effects of the mass outboard of the gauges, the wheel force signals can be calculated. From these, measures such as dlc and 95th percentile impacts, can be determined.

The instrumentation also allows the calculation of wheel forces by summing a chassis accelerometer signal multiplied by a sprung mass value with the axle accelerometer signal multiplied by an unsprung mass value. There are theoretical flaws in this approach because firstly, the body modes of the vehicle include bounce, pitch and roll, each of which is based on a different effective sprung mass contribution and secondly, the tandem axle suspension redistributes the sprung mass contribution via its load sharing mechanism. From these tests the importance of these flaws can be investigated.

The dlc values measured ranged from 0.06 to 0.18 depending on vehicle speed and road roughness. This is very similar to Sweatman's (ref. 2) values for a similar suspension. The linear relationship between dlc to $VR^{0.5}$ as observed by Sweatman did not appear to fit the data well. However, it was obvious that the cause of the poor fit was that the results obtained from the urban road test site gave higher dlc values than would be expected. Without these data a good fit was achieved. The reason for this is not clear. Without the urban road tests the range of speeds is quite small and so the influence of speed on the values of $VR^{0.5}$ is also small. That is, even if the relationship were wrong a good fit might be achieved. It is also possible that the nature of the road profile at the urban site generated a proportionately greater response from the heavy vehicle suspension than from the passenger car suspension model used to determine roughness.

Using the chassis accelerometers to calculate wheel forces presented some difficulties. It was clear that the magnitude of the high frequency content of these signals was too great to represent motion of the sprung mass and was generated by flexing of the chassis members to which the accelerometers were attached. The suspension is attached to the same members. This problem was resolved by filtering these components out but, of course, this also filters out the true response at those frequencies.

Two methods of calculating the wheel forces from the chassis accelerometers were used. The simpler method uses only the signal of the chassis accelerometer directly above the wheel under consideration and uses the static load to calculate the sprung mass value. The more complex method uses the four chassis accelerometers at the corners of the chassis to separate the bounce, pitch and roll motions. Each of these motions is then multiplied by an "effective mass" value which is determined from the appropriate calculated moment of inertia.

Dlcs calculated by this method were generally, though not always, close to those calculated from the strain gauge signals with the complex method being a little better than the simple method. Typically differences were about 5-10% though worst case errors were about 15%.

Figures 2 and 3 below show a comparison of wheel forces calculated by these two methods compared in each case with the strain gauge calculated values. The second graph in each figure shows the Fourier transforms of the signals and so compares them in the frequency domain.

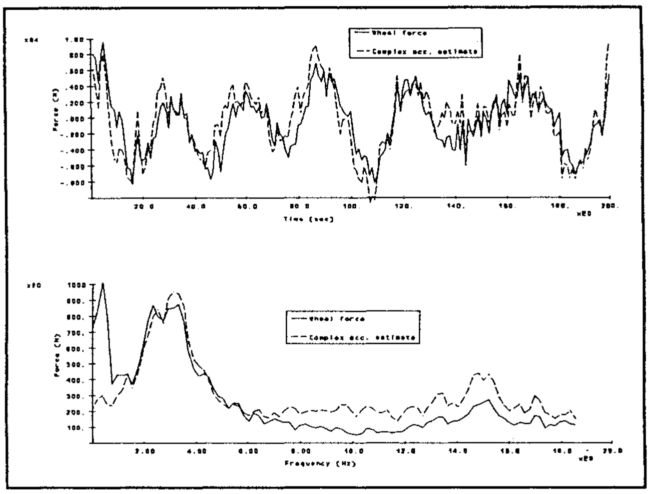

Figure 2. Wheel force by complex accelerometer method.

3.4 Servohydraulic tests

The aim of the servohydraulic trials was to provide shaker driving signals at the two rams which produced the same suspension deflections (at the excited wheels) as measured during the road tests by the LVDTs. The forces required to produce these deflections should be identical and are the sum of the force at the wheel and the inertia of the unsprung mass. By monitoring axle accelerometers during the shaker trials it is possible to compensate for any differences in this inertia component.

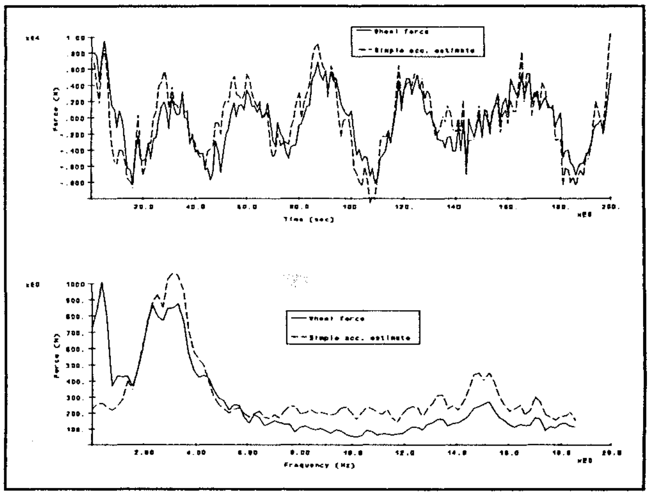

Figure 3 Wheel force by simple accelerometer method

Clearly applying the displacements measured by the LVDTs as excitation signals at the rams will not elicit the correct response and RPC is needed. The underlying concept of RPC for a single channel is very simple. An excitation signal is applied and the response is measured. From these two signals the transfer function can be calculated. Using the inverse of this transfer function and a target response signal, the required excitation signal can be determined and applied. The new response is then compared to the target. If not satisfactory, the excitation is adjusted and the process is repeated. A number of iterations may be needed.

Applying the process to two excitations with two responses increases the complexity but the principles are the same. Instead of a single transfer function, there is a 2x2 transfer function matrix. This matrix must be invertible and this means that the two excitations must be independent. There is a further underlying assumption that the dynamic system is linear which, because of the steel leaf springs, is clearly not correct.

Several approaches to adjusting the excitation signals were tried. The first was to calculate a new inverse transfer function at each iteration. This was aimed at overcoming the response non-linearity by determining pseudo transfer functions which were "correct" at the desired response levels. Two difficulties arose with this method. Firstly, because the road behaviour being simulated was generated by the two wheels passing over the same profile (with a small time delay), the two excitation signals required are not likely to be independent. This lack of independence manifests itself as "spikes" in the calculated inverse transfer functions. An editing procedure in the software was used to eliminate these "spikes". Secondly, the procedure overcompensates and is unstable. This is not surprising when the physical system is considered. At low excitation levels the stiction of the steel leaf spring generates very low response. However, as the excitation level increases this stiction is overcome and the spring becomes effectively much less stiff and the response is proportionately much greater. To avoid this over-compensation, a "relaxation factor" was introduced so that at each iteration, the inverse transfer function was adjusted by only a portion of the newly calculated one. With these modifications to the procedure it was possible to achieve a reasonable match to the on-road response. However, because the editing of the inverse transfer functions at each iteration was a slow process, about five hours testing was required to match a single road test. This clearly is too slow for a routine test procedure. Figures 4 and 5 show an example of the fit achieved with this algorithm.

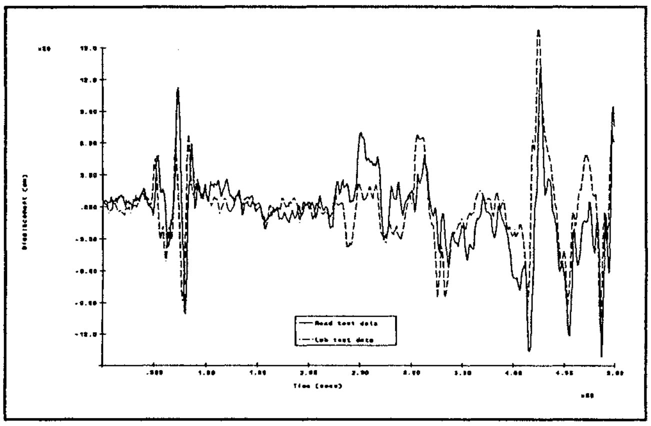

Figure 4 Sample suspension deflection signal for rear wheel

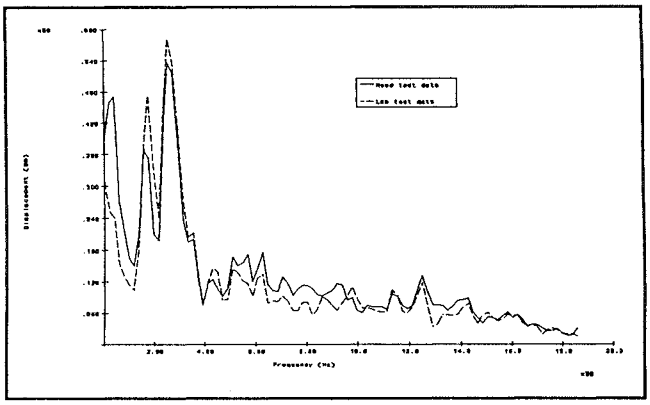

Figure 5 Spectral distribution of rear suspension deflections

As the on-road suspension response was generated by the wheels passing over the same road profile at the same speed with only a time delay between them, it was felt that it might be possible to generate the laboratory response in the same way. Thus the software was modified to allow both rams to be sent the same excitation signal with a time delay corresponding to vehicle speed between them. As there were still two response target signals, the inverse transfer function was represented by a 2x1 matrix and the calculated excitation signal is a compromise between the two targets. This does, however, eliminate the "spikes" problem experienced with the previous method and consequently the time-consuming editing. As before, it was necessary to use "relaxation" for the iteration process to be convergent. Using this method a "solution" can be found in less than an hour. Figures 6 and 7 show the match achieved.

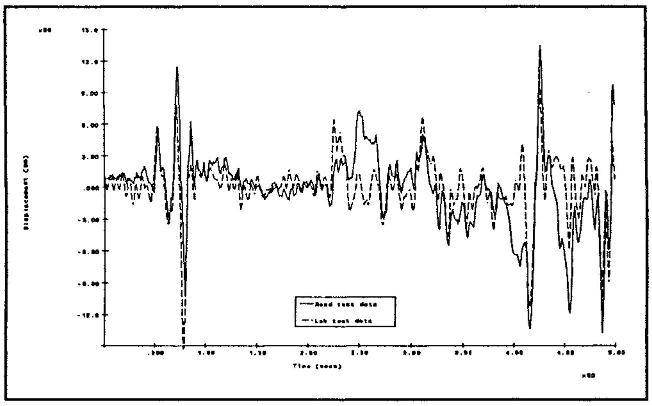

Figure 6 Suspension deflection signal for rear wheel.

Although it was anticipated that this method would find a compromise solution for the two wheels, in fact, the solution for the front wheel of the tandem set (not shown) was consistently better than that for the rear wheel (shown in figures). No explanation for this phenomenon was found. Overall, the match to the road behaviour was slightly inferior to that achieved by the previous method but the process was much quicker and required much less operator intervention.

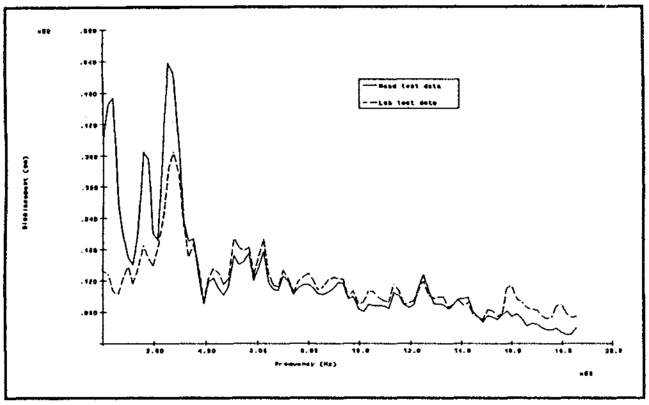

Figure 7 Spectral distribution of rear suspension deflections.

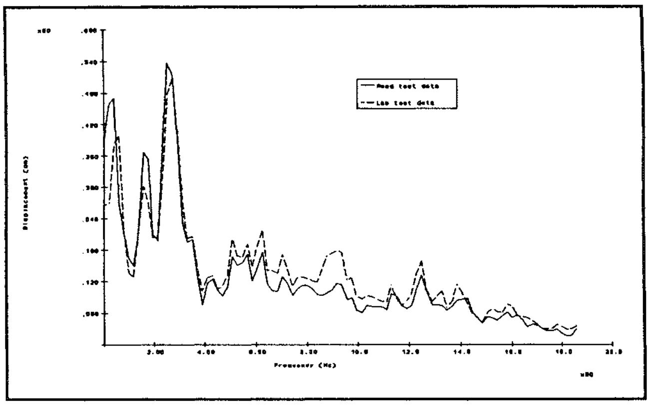

Figure 9 Spectral distribution of rear suspension deflections.

Finally a fundamentally different approach was tried. The inverse transfer functions were calculated once at the beginning of the process using random excitations. These were then used to estimate the excitations needed to produce the target responses as before. Each subsequent iteration considered the difference between the actual responses and the target responses and calculated the changes in the excitations needed. No recalculation of the inverse transfer functions was done and consequently the process was very straightforward.

Using this method with a single excitation signal with a time shift applied to both rams gave similar results to the previous method with a slight gain in speed. However, allowing the two excitations to be independent, resulted in a solution which was clearly superior to those previously obtained and which could be achieved in less than an hour. Figures 8 and 9 illustrate the results for the same data as shown previously.

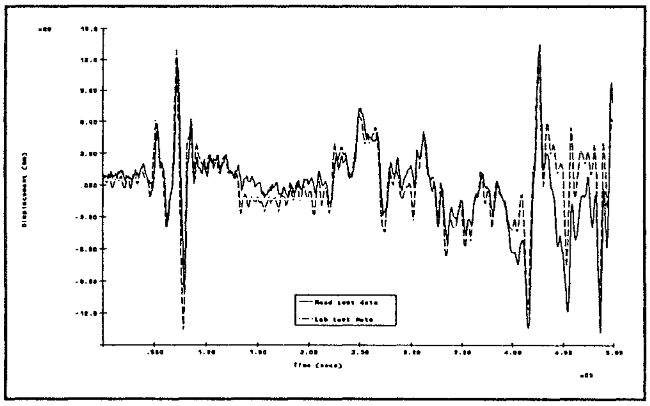

Figure 8 Suspension deflection signal for rear wheel

All the discussion so far has centred on matching suspension deflections between the laboratory and road tests when the aim is to match wheel forces. Figure 10 compares the wheel forces as calculated from the strain gauge signals on the road with those measured by the ram load cells. As can be seen the match is quite good.

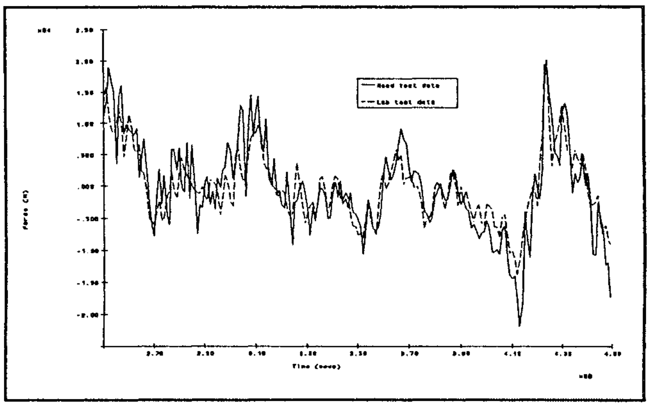

Figure 10 Force signal for rear wheel.

There are similarities between the road profile data and the shaker displacement signals. Comparing them spectrally indicates that the shaker displacements have additional energy at the body mode frequencies of the vehicle. This is not surprising as, in the laboratory, the body modes are excited by shaking only two wheels while on the road all six wheels are excited. The relationship between the road profiles and the shaker excitations is primarily a function of the vehicle as a dynamic system. Work is in progress to see whether it is possible, by characterizing the vehicle's dynamics, to find an estimate of this relationship. If successful this would enable the shaker excitations needed to simulate a particular pavement profile to be determined without a road test.

4. CONCLUSIONS

It has been established that the use of "road-friendly" suspensions reduces pavement wear. For the pavement manager, there are two primary issues involved in encouraging better suspensions. Firstly, how much reduction in pavement wear is achieved and consequently how much incentive should be offered. Secondly, how can a suspension be assessed for performance in relation to pavement wear. This paper addresses the second issue.

An assessment technique has been developed which uses two servohydraulic shakers to replicate the on-road behaviour of a heavy vehicle suspension in the laboratory. It involves fitting the suspension with LVDTs and axle mounted accelerometers and monitoring them during some on road tests. Shaker excitation signals which produce the same suspension responses are then developed and the wheel forces measured directly. It is expected that it would be possible to undertake a complete test involving three or four pavement sections in about a day.

It is possible that a relationship can be established between the shaker excitations, the pavement profile and the vehicle characteristics. If this can be done it will be possible to undertake the servohydraulic tests without first doing a road test. This would eliminate the vehicle instrumentation and reduce the time required to less than half a day. It has the additional advantage that the road profile used is constant, thus making the test repeatable and transferable. The investigation into this option is proceeding.

ACKNOWLEDGEMENTS

The research described is this paper has been funded by Transit New Zealand and by the New Zealand Foundation for Science, Research and Technology. The author is very grateful to these organisations for their support.

REFERENCES

1. WHITTEMORE A.P., WILEY J.R., SCHULTZ P.C. and POLLOCK D.E. Dynamic pavement loads of heavy highway vehicles. NCHRP Report 105, Highway Research Board. 1970.

2. SWEATMAN P.F. A study of dynamic wheel forces in axle group suspensions of heavy vehicles. Australian Road Research Board, Special Report No 27, 1983.

3. WOODROOFFE J.H.F., LEBLANC P.A. and LEPIANE K.R. Effects of suspension variations on the dynamic wheel loads of a heavy articulated highway vehicle. Roads and Transportation Association of Canada, Vehicle Weights and Dimensions Study, Vol 11, 1986.

4. MAGNUSSON G. Measurement of dynamic wheel load. VTI Report 279A, Linkoping, Sweden, 1987.

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

6. DICKERSON R.S. and MACE D.G.W. Dynamic pavement force measurements with a two axle heavy goods vehicle. TRRL Supplementary Report SR688. Transport and Road Research Laboratory, Crowthorne.

7. MITCHELL C.G.B. and GYENES L. Dynamic pavement loads measured for a variety of truck suspensions. Proceedings 2nd International Symposium on Heavy Vehicle Weights and Dimensions, Kelowna, British Columbia, June 1989.

8. EISENMANN J. Dynamic load fluctuations - road stress. Strasse und Autobahn 4.

9. CEBON D. An investigation of the dynamic interaction between wheeled vehicles and road surfaces. PhD dissertation, Cambridge University, 1985.

10. WYPICH P. "Road sparing design" -current status of test measurements being carried out at MAN. Paper presented to OECD Scientific Expert Group IR2, Munich, October 1990.

11. COLE, D.J. and CEBON, D. Simulation and Measurement of Dynamic Tyre Forces. Proceedings 2nd International Symposium on Heavy Vehicle Weights and Dimensions, Kelowna, British Columbia, June 1989.

12. WOODROOFFE J.H.F., LEBLANC P.A. and PAPAGIANNAKIS A.T. Suspension dynamics - Experimental findings and regulatory implications. SAE Truck and Bus Meeting, Nov 1988, SAE Publication 881847.

13. HEATH A.N. Heavy Vehicle Design Affecting Road Loading. Proc ARRB/FORS Symposium on Heavy Vehicle Suspension Characteristics, Canberra, March 1987, pp251-270.