Modelling of a semi-active hydropneumatic spring–damper unit

N.J. Theron* and P.S. Els

Dynamic Systems Group,
Department of Mechanical and Aeronautical Engineering,
University of Pretoria, Pretoria 0002, South Africa
E-mail: nico.theron@up.ac.za   E-mail: schalk.els@eng.up.ac.za
*Corresponding author

Abstract: The mathematical modelling of a suspension unit is considered. The unit comprises a hydraulic cylinder connecting the vehicle body to the unsprung mass, two nitrogen-filled accumulator springs and two damper ports. The model takes the deflection rate as input and iteratively employs simple fluid dynamics theory to calculate the flow-rates from each accumulator to the cylinder. It calculates the pressure in the accumulators by time-integrating the flow rates to determine the gas volumes and then invoking ideal gas theory. This renders the dynamic force of the unit as output. Model predictions are compared with measurements.

Keywords: semi-active suspension; hydropneumatic; twin accumulator spring–damper; vehicle dynamics; off-road; mathematical model; simulation; model verification.

Biographical notes: Nicolaas Johannes Theron obtained a Master's Degree from the University of Stellenbosch, South Africa, in 1985, and a PhD from Rensselaer Polytechnic Institute in 1994, both in Mechanical Engineering. He was involved in helicopter structural dynamics for a number of years. He also taught for five years at the University of Stellenbosch, and was appointed as an Associate Professor of Mechanical Engineering, at the University of Pretoria in 2000. He is interested in structural dynamics and control. He is a member of the South African Institution for Mechanical Engineering and the American Helicopter Society. He is registered as a Professional Engineer by the Engineering Council of South Africa (ECSA).

Pieter Schalk Els obtained a Master's Degree in Mechanical Engineering in 1993 from the University of Pretoria, South Africa. He worked in the off-road and military vehicle industry for five years as a test and design engineer, before joining the Department of Mechanical and Aeronautical Engineering at the University of Pretoria as a Senior Lecturer in 1999. He is currently lecturing on dynamics, vehicle engineering and vehicle dynamics. He has published seven papers in peer-reviewed international journals, contributed to the proceedings of 13 international conferences and obtained four patents. He is a member of the South African Institution of Mechanical Engineering and the Society of Automotive Engineers (SAE), and is registered as a Professional Engineer by the Engineering Council of South Africa (ECSA).

1 Introduction

The design of wheeled vehicle suspension systems always involves a compromise between ride comfort and handling. For good ride comfort, a compliant suspension system is normally required, while good handling demands a stiff suspension system to control body roll. In the case of off-road vehicles (as used by the military), it is very difficult to achieve a good compromise because these vehicles are also used on highways at high speeds. With a normal passive suspension system, the characteristics of the springs and dampers are fixed at the design stage and cannot be changed afterwards. By using controllable springs and dampers, these characteristics can be changed while the vehicle is moving. It therefore becomes possible to have soft settings for good ride comfort whilst travelling in a straight line on a good road, while the suspension can be changed to a hard setting moments later to give good handling when the vehicle has to change direction as required for lane changing or even accident avoidance. It also becomes possible to adapt the suspension characteristics to different terrains. Controllable suspension systems can, therefore, reduce or even eliminate the ride comfort vs. handling compromise.

Wallentowitz and Holdman (2001) give a frequency domain analysis of the effect of spring and damper stiffness on the transfer function of the suspension. They conclude that two spring stages are sufficient to overcome the compromise associated with passive systems. A soft spring is required to optimise ride comfort while a stiff spring is only used during cornering and braking, when the soft spring will result in an unacceptable body roll and pitch. Karnopp and Margolis (1984) discuss the effects of a change in spring and damper rates on the transfer function of a single degree of freedom suspension system. It is said that changing the damping alone is not a very good way of stiffening or softening a suspension system. Therefore, a system that enables both the spring and damper rates to be adjusted is proposed. For this purpose a system with two air volumes separated by control valves is proposed. Decker et al. (1988) describe a prototype-adjustable air spring developed by BOSCH of which the spring characteristic can be changed between several values by fast (25 ms) switching of different air volumes. The adjustable spring is used in conjunction with a fast (4 ms) semi-active damper. Limited simulation results are presented.

Hine and Pearce (1988) discuss an industrialised version of a semi-active suspension, developed by Armstrong. A two or three state adjustable damper is combined with an air or oleopneumatic spring that is said to offer both height and spring-rate control. Citroën has been applying hydropneumatic suspension systems to their passenger cars for many years. The semi-active suspension unit fitted to the XM model consist of three spheres and four dampers for both the front and rear suspensions. The system can be switched to a low-spring and low-damping state (three spheres and four dampers) or high-spring and high-damping state (two spheres and two dampers). Another system fitted to Citroën's Activa 2 research prototype car is described by Birch et al. (1990). The system is an upgrade of that used for the XM and adds an active antiroll system that can double the roll stiffness almost instantly to counter body roll.

An Electronic Controlled Suspension (ECS), as fitted to the 1984 Mitsubishi Galant, is discussed by Mizuguchi et al. (1984). A two-stage spring is constructed using an air spring in parallel with a conventional metal coil spring. The air spring consists of two chambers connected by a valve. The valve is closed to activate the stiff spring rate. The suspension is set either to 'off' (soft spring and soft damper) or to 'on' (hard spring

and damper). A system very similar to that offered on the Galant is used on the 1986 model Toyota Soarer, as described by Hirose et al. (1988). This system changes both spring and damper characteristics using direct current electric motors. The air spring uses main- and sub-air-chambers connected by a disc valve to change the gas volume and therefore the spring characteristic. The spring and damper rates are changed simultaneously by a single electric motor.

We are currently involved in a research project with the aim of developing a suspension unit comprising a two-stage, semi-active, hydropneumatic spring, combined with a two-stage semi-active damper, for an off-road vehicle as typically used by the military. Although the concept of a two-stage spring is not new, no mathematical models of similar systems have been found in the literature. Furthermore, existing systems are all applied to road vehicles, while the current application is for an off-road vehicle that also requires good on-road capabilities. The aim of the research reported on in this paper was to develop for this suspension unit a mathematical model that can be used in vehicle dynamic simulations.

2   Physical description of the suspension unit

The suspension unit is developed for installation in a specific model of sport utility vehicle of two tonne gross mass. The main strut of the suspension unit is fixed between the vehicle body and live axle, replacing both the coil spring and telescopic damper of the standard production model. A vehicle will typically be equipped with four of these suspension units, one for each wheel.

The basic layout of the suspension unit is given in Figure 1. The main strut chamber is hydraulically connected to two piston accumulators via a valve block. All the control valves, hydraulic damper valves, control ports and channels are accommodated inside the valve block, which forms an integral part of the upper structure of the main strut.

Figure 1   Basic layout of the hydropneumatic suspension unit
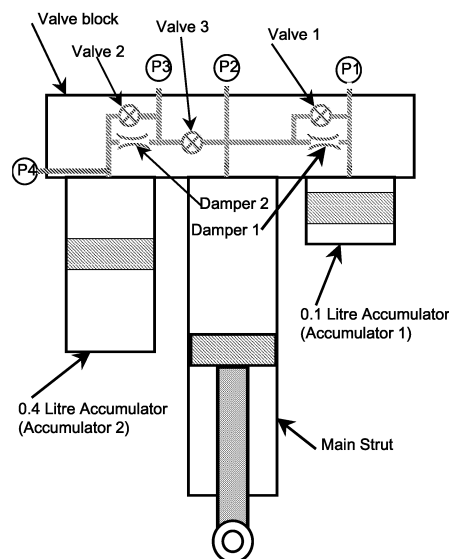
The accumulators are charged with nitrogen gas while the suspension unit is subjected to a compressive force equal to the static weight that the unit will see installed in the vehicle. At static equilibrium, the pressure in the unit is about 4 MPa. Accumulator 1 is typically charged with 0.1 litres and accumulator 2 with 0.4 litres of gas. The low spring rate is achieved by compressing the combined gas volume of both accumulators. By sealing off accumulator 2, a smaller gas volume is compressed and a higher spring rate is achieved. Accumulator 2 is sealed off with two parallel electrohydraulic valves that are in combination referred to as valve 3. (Two valves are used merely to reduce losses when open. The two valves are always switched on simultaneously.) Spring rates can be individually tailored by changing the two gas volumes during charging.

For low damping, the hydraulic dampers (dampers 1 and 2) are short circuited by opening the electrohydraulic bypass valves (valves 1 and 2). For high damping, these valves are closed and the hydraulic fluid is forced through the dampers resulting in a high damping force.

A first and, later, a second prototype unit have been built and were tested and characterised to obtain all the parameters required for mathematical modelling as well as vehicle dynamics simulation. On both prototypes, provision is made for four pressure transducers (P_1–P_4), to measure the four important pressures P_1–P_4 in the system. These pressures are defined below.

3   Mathematical modelling of the suspension unit

In developing a mathematical model, a tension force in the unit is considered positive while a compressive force is negative. Any extension of the unit relative to a reference state is considered as a positive (relative) displacement and compression of the unit as negative displacement. An extensional speed is considered positive and a compression speed as negative.

For the purposes of vehicle dynamics simulation, a mathematical model of this unit is required, which calculates the combined spring–damper force for a certain set of valve settings and a given state of displacement and speed. One may, therefore, consider the force of the suspension unit as the output of the model and the valve settings of the three valves and the displacement and speed of the unit as the inputs to the model, where the model calculates the output for given inputs. This calculation is typically performed within a time step in a simulation run and is repeated for each time step.

The output force of the unit is essentially directly related to the pressure P_2 in the main strut cylinder. This pressure depends on the pressures in the two accumulators, the flow through and corresponding pressure drops over the two dampers and corresponding channels and the valve switching. The pressure in the accumulators depends on the volume of oil in the accumulators, which is related to the displacement of the suspension unit and the state of valve 3. An alternative way of looking at the volume of oil in the accumulators is to realise that this is determined by the flow history, i.e., these volumes may be determined by integrating the flow rates in the two main branches of the system. Using this approach makes the mathematical model independent of the displacement of the unit as an input. This is indeed the approach that was used in modelling the unit. The input to the model of the suspension unit is therefore, in addition to the three valve switch signals, only the extensional speed $\dot{x}$ of the unit. From this, the volume flow rate

$q = A\dot{x}$ into the main strut cylinder of cross-sectional area A can directly be calculated. The flow rates in the two branches are taken as q_i, $i = 1, 2$ for the branch associated with accumulator i, positive in the direction from the accumulator towards the main strut.

3.1 Pressure-dependent valve switching

It is assumed that the electric signals with which the various valves are switched changes instantaneously from low to high values, or vice versa. When this happens, valve and other dynamics prevent immediate pressure and flow changes. These dynamic effects are not currently modelled mathematically, but are taken into account empirically. The valve response time is defined as the time that it takes for the pressure at the low-pressure valve port to rise to 95% of the pressure difference between the low- and high-pressure ports after the valve was switched open electrically. It is assumed for now that the valve response time is the same, irrespective of the valve port to which the higher pressure is applied. This, of course, is not generally true. The valve response times of a typical valve used in the suspension unit was measured at various pressure drops across the valve. It was found that the measured valve response times are approximated fairly well by a parabolic curve that was fitted to the data using a least squares approach. This parabolic curve was subsequently employed in the mathematical model with respect to all three valves. It was also found for this valve that the time it takes for the pressure at the low-pressure port to rise to 5% of the pressure difference was typically half the measured valve response time at all pressure drop test points.

Wherever the state of the valve is taken into account in the model, a fraction f_i between zero and one is used, where the subscript $i = 1, 2, 3$ indicates the valve number. For switching on the valve (electric signal going from low to high, valve going from closed to open), $f_i = 0$ before the electrical signal switches, $f_i = 0.05$ at half the valve response time after the electrical signal switches, $f_i = 0.95$ at the valve response time and $f_i = 1$ after 1.5 times the valve response time. In between these time points, a piecewise cubic Hermite interpolation is used to calculate the fraction. For switching off the valve, the same type of interpolation is used on the reversed sequence.

3.2 Pressure drop over dampers and valve 3

Owing to the complexity of possible dampers that may be used in the suspension unit, it was decided to use table look-up techniques to get the pressure drop over the damper for a given flow rate through the damper. Quite often, the pressure-flow characteristics display significant hysteresis. For now, the table look-up procedure employed does not provide for possible hysteresis. The pressure drop over the damper was measured, with the bypass valve both open and closed, for various positive and negative flow rates in a practically realistic range. This measured data was used to establish a high-damping and a low-damping damper curve, corresponding to the bypass valve being closed and open, respectively. These curves are used in the table look-up procedure for both dampers 1 and 2, since they currently are identical and their bypass valves are also identical. The fact that the internal passages in the valve block for the two dampers at this time are not identical is neglected for now.

For a certain flow rate q_i, the pressure drop over damper i, $i = 1,2$, is calculated as $\Delta P_{di} = f_i \Delta P_{doi} + (1 - f_i)\Delta P_{dci}$, where ΔP_{dci} is the pressure drop interpolated at q_i from the high damping graph of damper i, while ΔP_{doi} is the pressure drop interpolated at q_i from the low damping graph of damper i.

When valve 3 is fully open ($f_3 = 1$), the pressure drop over valve 3, ΔP_{v3}, is calculated using an experimentally determined loss factor and the flow q_2. When the valve is opening ($0 < f_3 < 1$, f_3 increasing), a value ΔP_{v3o} is calculated in exactly the same way as ΔP_{v3} above, but the actual pressure drop over the valve is taken as $\Delta P_{v3} = f_3 \Delta P_{v3o} + (1 - f_3)\Delta P_{v3i}$, where ΔP_{v3i} is the pressure drop over the valve before the switching started. When the valve is closing ($0 < f_3 < 1$, f_3 decreasing), on the other hand, the actual pressure drop over the valve is taken as $\Delta P_{v3} = f_3 \Delta P_{v3o} + (1 - f_3)\Delta P_{v3e}$, where ΔP_{v3e} is the pressure drop over the valve calculated for the scenario where all variables are at their current values except $q = q_1$ and $q_2 = 0$, i.e., as if valve 3 is fully closed.

3.3 Flow and pressure calculation

The mathematical model is essentially based on the assumption that the hydraulic fluid is incompressible. In the simulation, however, the compressibility of the fluid is taken into account as a refining correction in the calculation of the gas volumes in the accumulators. This correction is based on the various major volumes of fluid in the system, each at its respective pressure, and the bulk modulus of the hydraulic fluid.

Whenever valve 3 is closed, the system can be modelled as a third-order non-linear state space system; otherwise a fourth-order non-linear state space system with an algebraic constraint is obtained. These two alternative situations will now be considered separately.

3.3.1 Valve 3 closed

When valve 3 is closed, $q = q_1$ and $q_2 = 0$, owing to the assumed incompressibility of the hydraulic fluid. Let the volume of gas in accumulator i be V_{gi}. The rate of change in the gas volume in accumulator 1 is

$$\dot{V}_{g1} = q_1. \tag{1}$$

The pressure P_{accui} in the accumulator i is calculated using the ideal gas law

$$P_{accui} = m_i R T_i / V_{gi} = R T_i / v_i \tag{2}$$

where m_i is the mass of gas with which the accumulator is charged, $R = 296.797$ is the gas constant for nitrogen, $v_i = V_{gi}/m_i$ is the specific volume and T_i is the absolute temperature of the gas in the accumulator. ($P_{accui} = P_1$ for accumulator 1 and $P_{accui} = P_4$ for accumulator 2.) T_i is calculated by solving the following differential equation, as suggested by Els and Grobbelaar (1999):

$$\dot{T}_i = \frac{T_{i0} - T_i}{\tau_i} - \frac{T_i}{c_v}\left(\frac{\partial P_{accui}}{\partial T_i}\right)_v \dot{v}_i \tag{3}$$

where T_{i0} is the initial gas temperature, in this taken as the ambient temperature, τ_i is the thermal time constant of the accumulator and c_v is the specific heat at constant volume of the gas. The thermal time constant is taken at experimentally determined values of 4.8 seconds for both accumulators. Calculating the gas temperature in this way means that if the gas is suddenly compressed, the model calculates the pressure rise along an adiabatic compression curve, while the temperature rises. However, if the gas is subsequently allowed to cool down, the model allows the pressure to drop to the value indicated by the isothermal compression curve.

From equation (2) it follows that

$$\frac{\partial P_{accui}}{\partial T_i} = \frac{R}{v_i}. \tag{4}$$

Substituting this in equation (3) renders

$$\dot{T}_i = \frac{T_{i0} - T_i}{\tau_i} - \frac{T_i R}{c_v V_{gi}} q_i = f_{Ti}(T_i, V_{gi}, q_i), \tag{5}$$

where the $f_{Ti}(T_i, V_{gi}, q_i)$ on the right hand side indicates that $\dot{T}_i$ is a function of the variables T_i, q_i and V_{gi}. Equation (5) is non-linear owing to the appearance of the product of these variables.

Since $q_2 = 0$, there is no change in the gas volume in accumulator two. The pressure in this accumulator may, however, still change, as the gas temperature may change. The third-order system is thus defined by the three differential equations, equations (1) and (5) for $i = 1, 2$. Within a simulation time step, in addition to these three differential equations, various other variables are calculated (for example, the accumulator pressures with equation (2)). There are, however, no algebraic equations that need to be solved simultaneously with the three differential equations, and the solution is therefore fairly straightforward.

Once q_1 for the current time step has been calculated, the pressure P_2 in the main strut cylinder is calculated by calculating ΔP_{d1}, as described in Section 3.2, and then $P_2 = P_1 - \Delta P_{d1}$. With P_2 known, the output of the model is simply calculated by multiplying this pressure with the negative of the main strut cross-sectional area.

3.3.2 Valve 3 open, opening or closing

When valve 3 is partially or fully opened, the flow rate q_2 is no longer zero. Owing to the assumed incompressibility of the hydraulic fluid, $q = q_1 + q_2$. The rate of change in gas volume in accumulator 1 is still given by equation (1), while the rate of change in the gas volume in accumulator 2 is

$$\dot{V}_{g2} = q_2 = q - q_1. \tag{6}$$

In this case, however, an additional algebraic equation needs to be solved simultaneously with the differential equations. This equation may be considered as a constraint that needs to be satisfied, i.e., the pressure P_2 in the main strut cylinder calculated along two different paths must be the same. Let P_{21} be the pressure in the main strut cylinder, calculated along the branch connecting this to accumulator 1, as outlined in Section 3.3.1 above (which, for a given flow rate q_1, is also valid in this case). P_{21} is therefore a

function of the flow rate q_1 and the pressure P_1. The pressure P_1, by equation (2), is a function of T_1 and V_{g1}. Therefore, $P_{21} = P_{21}(q_1, T_1, V_{g1})$. In a similar way the pressure P_{22} in the main strut cylinder, calculated along the branch connecting this to accumulator 2, may be calculated, by first calculating ΔP_{d2} and ΔP_{v3} at the flow rate $q_2 = q - q_1$, as described in Section 3.2. Then, $P_3 = P_4 - \Delta P_{d2}$ and $P_{22} = P_3 - \Delta P_{v3}$. The pressure P_4 is a function of T_2 and V_{g2}; therefore, $P_{22} = P_{22}(q, q_1, T_2, V_{g2})$. The algebraic constraint may then be written as:

$$0 = P_{21}(q_1, T_1, V_{g1}) - P_{22}(q, q_1, T_2, V_{g2}). \tag{7}$$

Also, whereas equation (5) is still valid for accumulator 1, for accumulator 2, the flow rate q_2 needs to be substituted with $q - q_1$, so that the system dynamics may be summarised in the following non-linear state space representation:

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \dot{T}_1 \\ \dot{T}_2 \\ \dot{V}_{g1} \\ \dot{V}_{g2} \\ \dot{q}_1 \end{bmatrix} = \begin{bmatrix} f_{T1}(T_1, V_{g1}, q_1) \\ f_{T2}(T_2, V_{g2}, q, q_1) \\ q_1 \\ q - q_1 \\ P_{21}(q_1, T_1, V_{g1}) - P_{22}(q, q_1, T_2, V_{g2}) \end{bmatrix}, \tag{8}$$

with input variable q and state variables T_1, T_2, V_{g1}, V_{g2} and q_1. The flow rate q_1 is not truly a state variable, but it is convenient to consider it as such, in order to write the four differential equations and the algebraic constraint in a single equation as above.

The matrix on the left of equation (8) is often called a mass matrix. This equation is an example of the so-called differential–algebraic equation, as the mass matrix is singular. This singularity is clearly caused by the algebraic constraint.

3.4 Implementation in SIMULINK®

As mentioned above, the aim of this research was to develop a mathematical model of the suspension unit, to be used in vehicle dynamic simulations. It was decided earlier to use the ADAMS® program for the vehicle dynamics simulation. A very convenient way to interface a mathematical model like that of the suspension unit, as described above, with an ADAMS model of a larger system (in this case the vehicle and its suspension system components other than the suspension units) is to implement the mathematical model in the SIMULINK environment. ADAMS can be linked to MATLAB® SIMULINK subprograms. For this reason, the mathematical model was implemented in SIMULINK.

MATLAB provides a solution scheme for differential–algebraic equations, and as a consequence SIMULINK has the ability to model algebraic constraints. Solution of the differential–algebraic equation, equation (8), using this functionality has been unsuccessful thus far. The mathematical model was, however, implemented successfully in SIMULINK by, within each time step, first calculating the valve fractions f_i, $i = 1,2,3$, based on the pressure drops over the valves at the end of the previous time step and then enforcing the algebraic constraint using a Newton–Raphson-type iteration to find the values of q_1, q_2, P_2 and P_3. After these values have been calculated, T_1, T_2, V_{g1} and V_{g2} are calculated by solving the four first order differential equations contained in equation (8). Lastly P_1 and P_4 are calculated using equation (2). During the Newton–Raphson-type iteration, the values of P_1 and P_4 at the end of the previous time step are used.

This iteration is performed in a MATLAB s-function that is called by the SIMULINK program. Once P_2 is calculated, the output force of the suspension unit for the current time step may be calculated as $F = - P_2 A$ and the program may move on to the next time step. It should be noted that the friction between the piston and the cylinder walls and the piston rod and its bushing is neglected in the calculation of F.

4   Validation of the mathematical model

The model of the suspension unit has been validated by comparing its predicted force output with forces measured on the second prototype unit. Figure 2 indicates the prototype suspension unit mounted in the test rig. A 100 kN Schenck Hydropuls actuator, fitted with a three-stage valve, is used under displacement control.

Figure 2   The second prototype suspension on the test rig
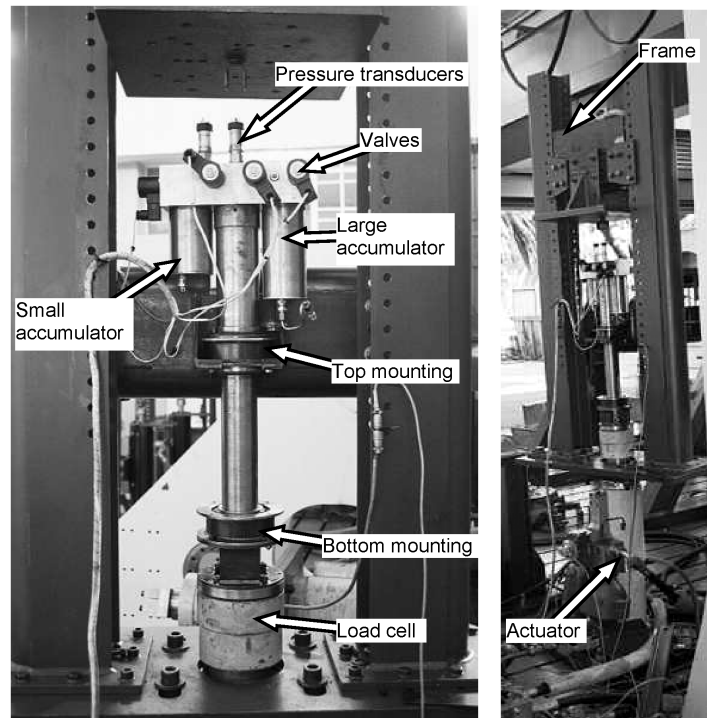
During testing on the hydrodynamic testing machine, the displacement feedback signal and resulting force, as measured with a load cell, were recorded. In addition to these two signals, the signals from the four pressure transducers measuring pressures P_1–P_4 and the electric command signals for switching the valves were also recorded. All these signals were filtered, to prevent aliasing, digitised and stored on disc.

Comparing the load cell force and the pressure P_2 measurements clearly showed that the error made in the model by neglecting the friction on the sliding parts of the unit and taking the output force of the unit as $-P_2 A$ is not insignificant but generally quite small. There is a second reason, other than friction, for the difference between the load cell

force and $-P_2A$, especially in situations of oscillation at high frequency. During vehicle simulation, the inertial properties of the piston and piston rod should be combined with those of the unsprung mass, so that the associated dynamic effects are taken into account by the ADAMS model, rather than the SIMULINK model. The output of the SIMULINK model should, therefore, be the suspension unit output force before the inertial effect of the piston and piston rod have been taken into account. The load cell, however, measures the suspension unit net output force after accelerating this mass. It is therefore prudent, in the comparison of the mathematical model with the measured results, to compare the output of the model in terms of measured and calculated $-P_2A$ values. In the discussion that follows, all reference to measured force should be understood to mean force calculated from the measured pressure P_2.

Since the mathematical model does not accept a displacement time history as input, but rather the extensional speed time history, the measured displacement signal first had to be differentiated with respect to time. It was always possible to bring the displacement signal back to its initial value at the end of a test run. The differentiation was therefore performed by transforming the whole displacement time history of a test run to the frequency domain using a Fast Fourier Transform (FFT), then multiplying the resultant double-sided complex spectrum with $j\omega$, setting all values corresponding to frequencies above a chosen low-pass filter cut-off frequency and below the negative of this cut-off frequency to zero and, lastly, back transforming the signal to the time domain using the inverse FFT. (In this, $j = \sqrt{-1}$ and ω is the circular frequency.) This procedure not only performs the differentiation, but also realises a low-pass filter with very sharp cut-off properties and no magnitude and phase distortion below the cut-off frequency. During vehicle simulation, this differentiation of the displacement is not required, since the ADAMS model directly calculates the required speeds.

To first test the spring properties without the influence of the dampers, the suspension unit was cycled through a triangular wave displacement at low speed, as indicated in Figure 3. This figure also shows the output force of the suspension unit, as calculated from the measured pressure P_2, for the case of stiff spring and low-damping properties. Figure 4 shows the comparison between the measured and SIMULINK-calculated time histories for this case, for the pressure in the active accumulator, P_1, and the main strut, P_2. Even though the nominal gas volume of accumulator 1 at the static wheel load was designed to be 0.1 litres, during this simulation, it was adjusted to 0.135 litres, in order to obtain what was considered an acceptable correlation between the measured and calculated results. This adjustment is to some extent justified owing to the fact that the volume calculation during design did not take into account some small cavities and screw thread inside the accumulator, and it was also determined that it is rather difficult to fill the suspension unit with oil, without trapping small pockets of air inside the unit. The volume of accumulator 1 could have been adjusted to an even higher value, to get an even closer correlation between the measured and calculated results at the peak at 100 seconds in Figure 4; but during the tests, other evidence was found that valve 3 was prone to leak at a high-pressure differential, which may have caused a reduced pressure P_1 and P_2 during the measurement. The force–displacement graph obtained during the test depicted in Figure 4 is shown in Figure 5, once again comparing the measured and calculated results. This test was repeated with a high-damping setting and essentially the same results were obtained, as expected, since the very slow speed renders a very small damping force.

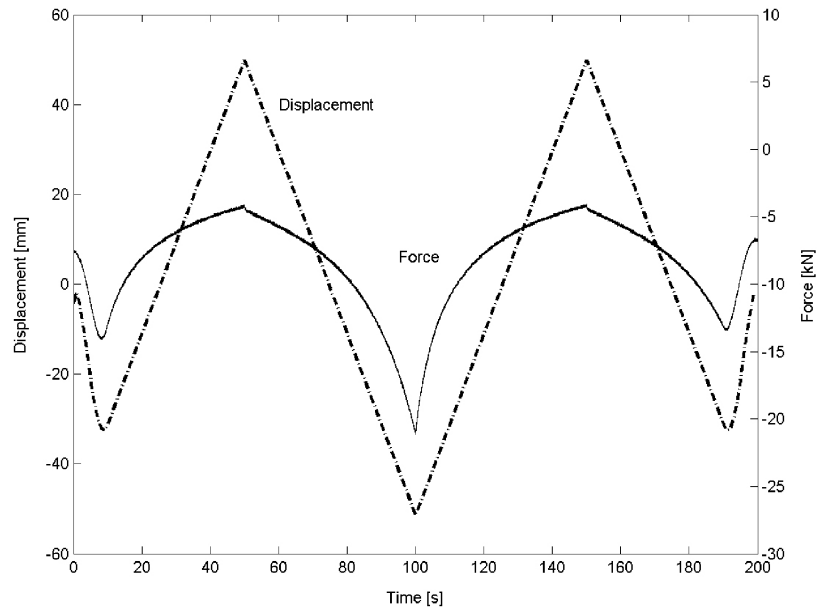**Figure 3**   Measured input and output: stiff spring and low damping at low speed

Figure 4   Comparison between measured and calculated values of P_1 and P_2: stiff spring and low damping at low speed
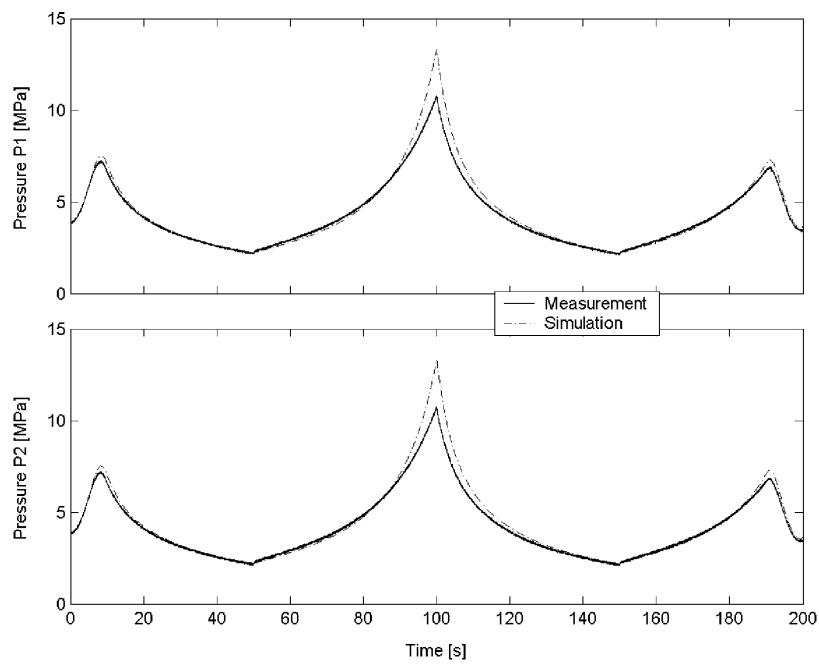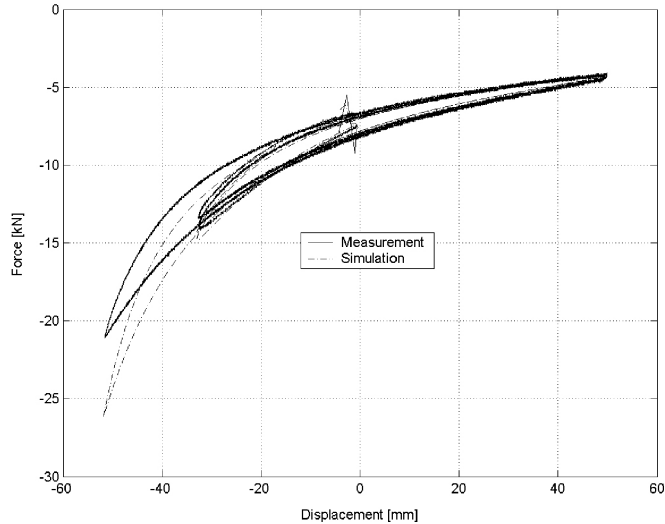
Figure 5   Comparison between measured and calculated force–displacement curve: stiff spring and low damping at low speed

Next, a similar test was conducted but at considerably higher speeds, to generate a significant damping effect. The input displacement and output force for a stiff spring and low-damping setting is shown in Figure 6. The comparison between the measured and calculated time histories for this case, for P_1 and P_2, are shown in Figure 7 and the force–displacement graph obtained during this test in Figure 8. The correlation between measurement and calculation displayed in Figure 7 is generally good, except at the high-pressure peaks. The calculated force–displacement graph shows an interesting Figure 8 shape, which was not observed in the measurement or in any other simulation result. When evaluating the force–displacement graphs generated by the simulation, one needs to bear in mind that the model does not yet provide for hysteresis in the damping properties. This may account for the strange curve calculated and displayed in Figure 8.

Figure 6   Measured input and output: stiff spring and low damping at high speed
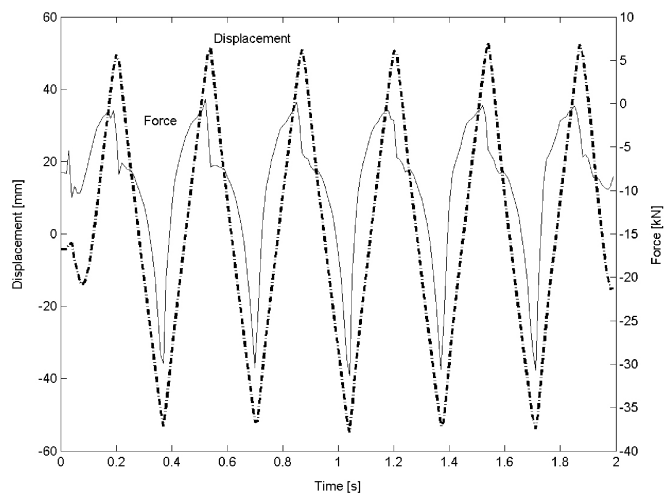
Figure 7 Comparison between measured and calculated values of P_1 and P_2: stiff spring and low damping at high speed. Only a part of the time history is shown, as it is typical of the whole time history
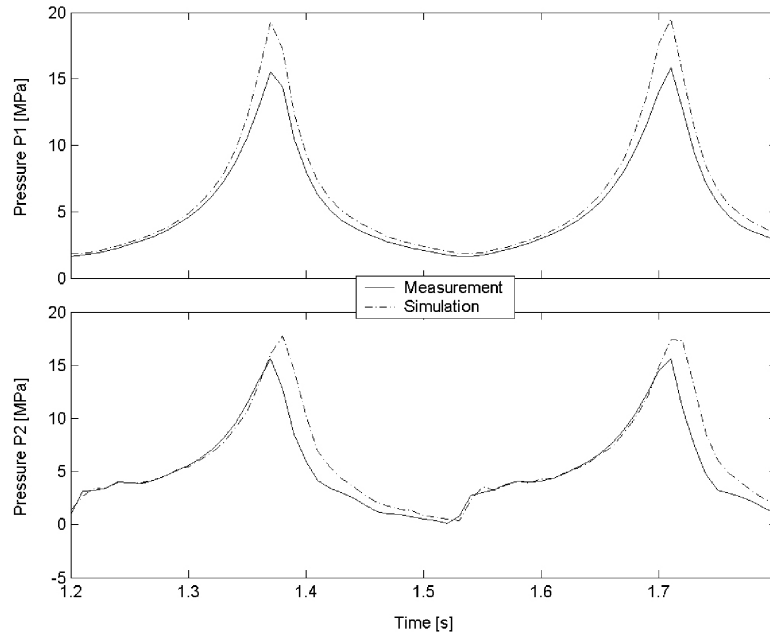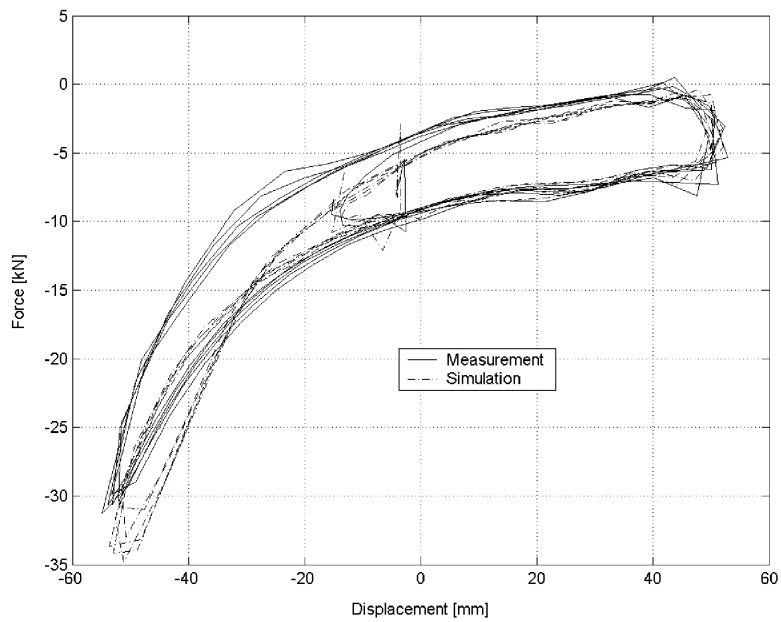

Figure 8 Comparison between measured and calculated force–displacement curve: stiff spring and low damping at high speed

Next, the same kind of test as that shown in Figure 6 was performed, only now with high damping (i.e., high stiffness and high damping, triangular displacement excitation at high speed). The input displacement and measured force time histories are shown in Figure 9. It is clear that the output force is clipped at about zero Newton, and the reason for this is that the pressure cannot drop very far below zero (atmospheric pressure), because at lower pressures the oil starts to boil, preventing further pressure drop. In any case, the pressure cannot drop below zero absolute, which would correspond to a positive output force of merely 196 N. The time histories of the pressures P_1 and P_2 are shown in Figure 10. The SIMULINK model has been constructed such that pressure P_2 will only drop to zero. It is seen that while P_2 is dropping, the model follows the measurement quite well into the saturation at zero. The model, however, recovers from this more quickly than the actual physical unit, causing the calculated pressure to start rising significantly earlier on the compression stroke than the measured pressure. After a delay, the calculation and the measurement meet up again with good correlation until this is repeated in the next cycle. One possible explanation for this delay is that in the physical unit, some boiling of the oil at low pressure occurs, a phenomenon that is not provided for in the SIMULINK model. Oil vapour caused by boiling and suspended in the oil is expected to cause a delay in pressure rise on compression. In this case, during the low-pressure part of the P_1 cycle, the correlation between simulation and measurement is not as good as that observed in the results discussed earlier. This may be related to the suspected boiling of the oil. The poor correlation in both the P_1 and P_2 results is not of serious concern, as the situation where the suspension unit is subjected to a prescribed high-speed rebound that can cause P_2 to drop to zero, even though easy to create on a test bench, is highly unlikely with the unit installed in a vehicle, even under rough road conditions. There is simply no downwards pull on the wheel available to cause such a condition. The force–displacement graph generated for this test is shown in Figure 11.

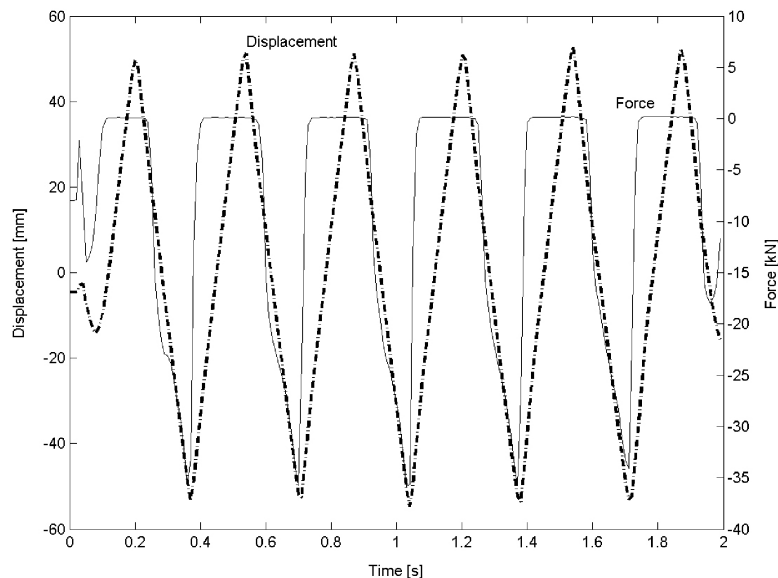**Figure 9**   Measured input and output: stiff spring and high damping at high speed

14

Figure 10 Comparison between measured and calculated values of P_1 and P_2: stiff spring and high damping at high speed. Only a part of the time history is shown, as it is typical of the whole time history
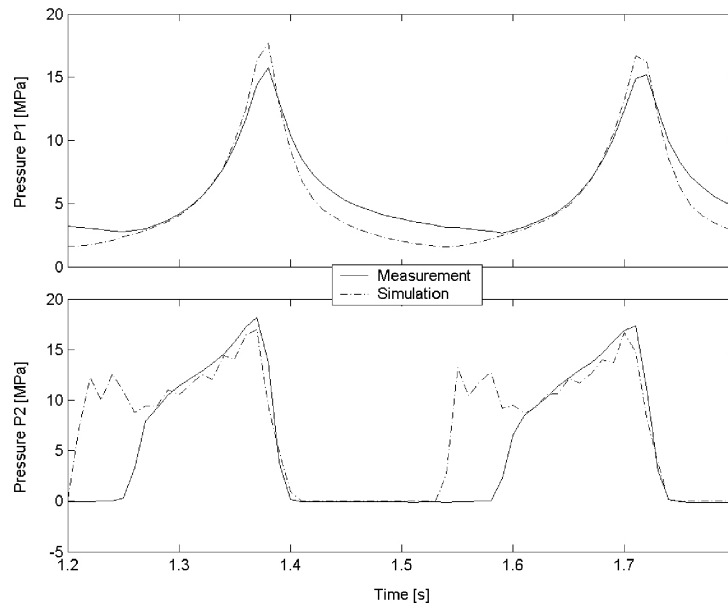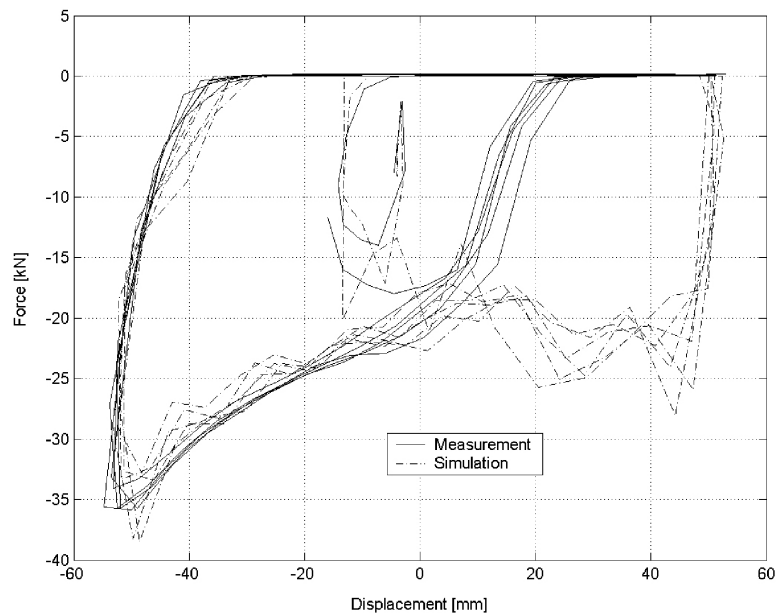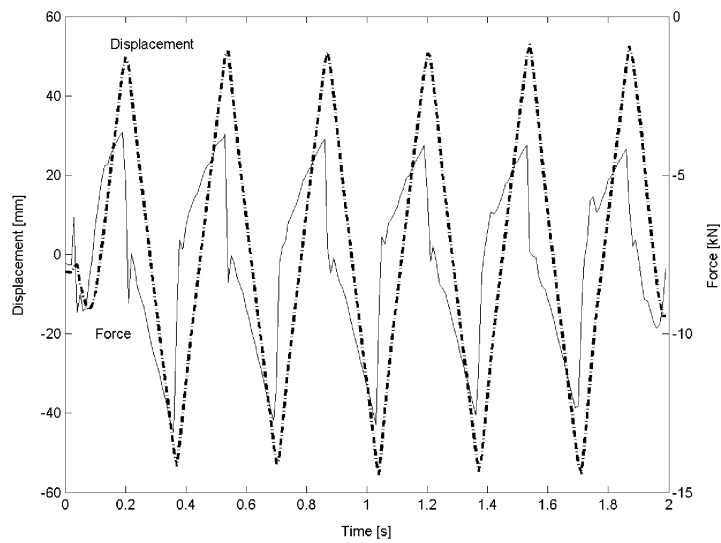

Figure 11 Comparison between measured and calculated force–displacement curve: stiff spring and high damping at high speed

Whereas all the results discussed above pertain to stiff spring scenarios, with valve 3 closed, the more complicated part of the model corresponds to the soft spring scenario. A test was done with soft-spring and low-damping settings, at low speed. The correlation

15

between the measurement and the simulation was generally good, comparable with the results obtained with the stiff spring discussed earlier.

Next, the above test was repeated at the high speed, the displacement input and measured force output shown in Figure 12. The comparison of the measured and calculated time histories of P_1 and P_4 for this case is shown in Figure 13 and that of P_2 and P_3 in Figure 14, with the force–displacement curve in Figure 15. Once again, the correlation is generally good.

Figure 12 Measured input and output: soft spring and low damping at high speed

Figure 13 Comparison between measured and calculated values of P_1 and P_4: soft spring and low damping at high speed. Only a part of the time history is shown, as it is typical of the whole time history
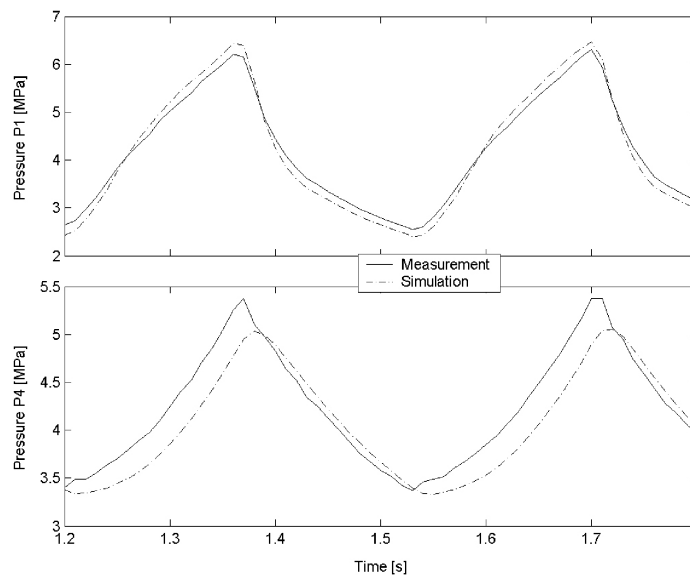
Figure 14 Comparison between measured and calculated values of P_2 and P_3: soft spring and low damping at high speed. Only a part of the time history is shown, as it is typical of the whole time history
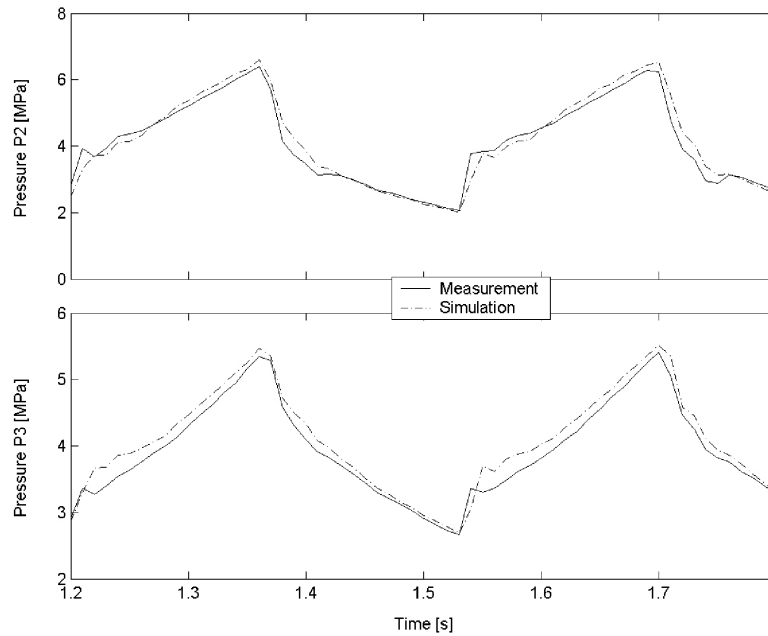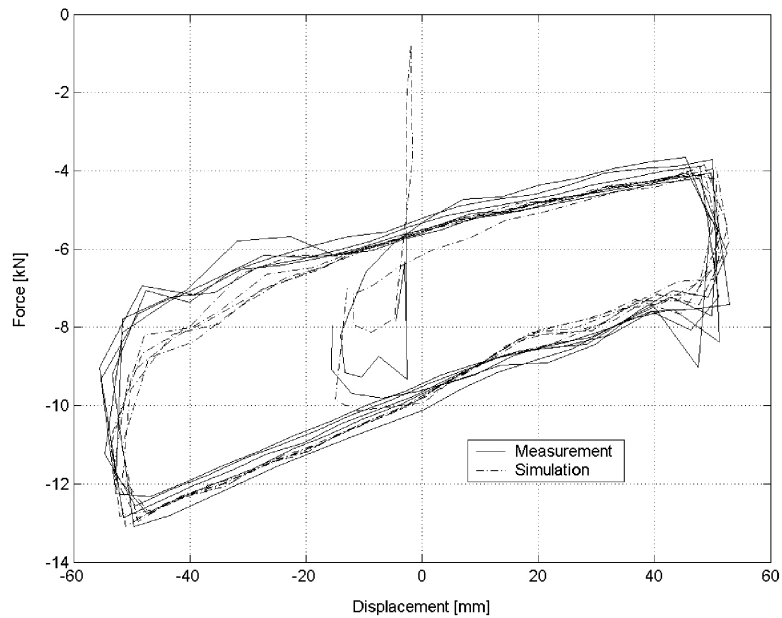

Figure 15 Comparison between measured and calculated force–displacement curve: soft spring and low damping at high speed

Lastly, a test was performed on the suspension unit wherein it was compressed for some distance in the stiff spring mode, then kept at this displacement for a while, after which valve 3 was opened and the pressures in the system allowed to equalise. Valve 3 was then closed again and the unit was then further compressed. This was repeated twice after which the unit was extended in a similar stepwise manner. This procedure, referred to herein as the incremental compression test, is well illustrated in Figure 16, which shows the time histories of the input displacement, the measured output force and the switch signal for valve 3. The switch signal is not plotted against a specific scale; it merely indicates when the valve is open (high) or closed (low). This whole test was conducted with a low-damping setting. The measured and calculated time histories of the pressure in the two accumulators are shown in Figure 17, while the time histories of P_2 and P_3 are shown in Figure 18. In this case, the gas volumes of accumulator 1 and 2 during the simulation were taken as 0.111 litres and 0.4 litres, respectively. The change in the volume of accumulator 1 may be justified by the fact that the suspension unit was emptied of both gas and oil, and then refilled, between this test and the test described earlier. With these settings, the only correlation that does not seem good is that between the measured and calculated time histories of P_4. It should, however, be realised that while valve 3 is closed, P_3 and P_4 should practically be identical, as there is now flow through damper 2 or its bypass valve. If the measured values of P_4 and P_3 from Figures 17 and 18 are compared, during the first second, when valve 3 is indeed closed, it is seen that the P_4 pressure transducer reads a pressure slightly higher that the P_3 transducer, by the same amount as the difference in the measured and calculated P_4 values in Figure 17. If based on this observation, it is assumed that an offset was present in the P_4 measurement, the correlation between the measurement and the simulation result may be considered as very good. The measured and calculated force–displacement graphs for this test are shown in Figure 19. This figure also shows a very good correlation between measurement and simulation.

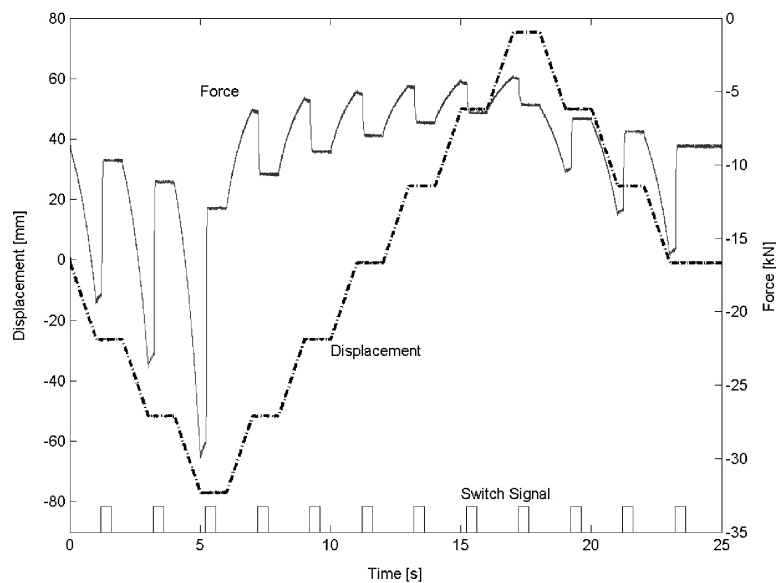**Figure 16** Measured input and output, and valve 3 switch signal: incremental compression test with low damping

Figure 17 Comparison between measured and calculated values of P_1 and P_4: incremental compression test with low damping
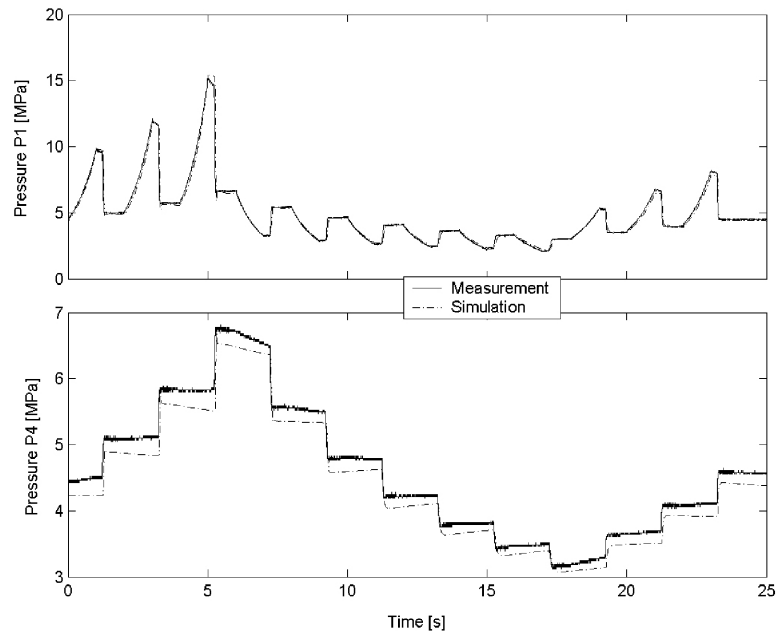

Figure 18 Comparison between measured and calculated values of P_2 and P_3: incremental compression test with low damping
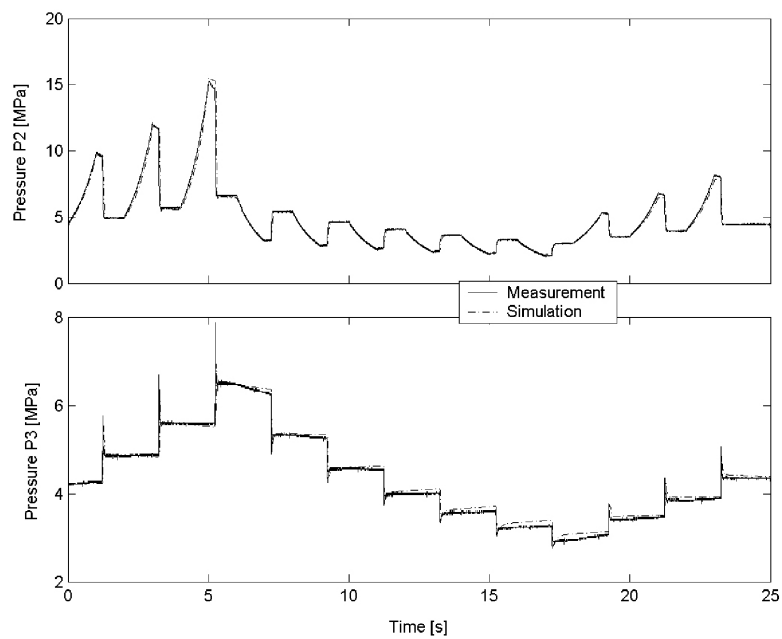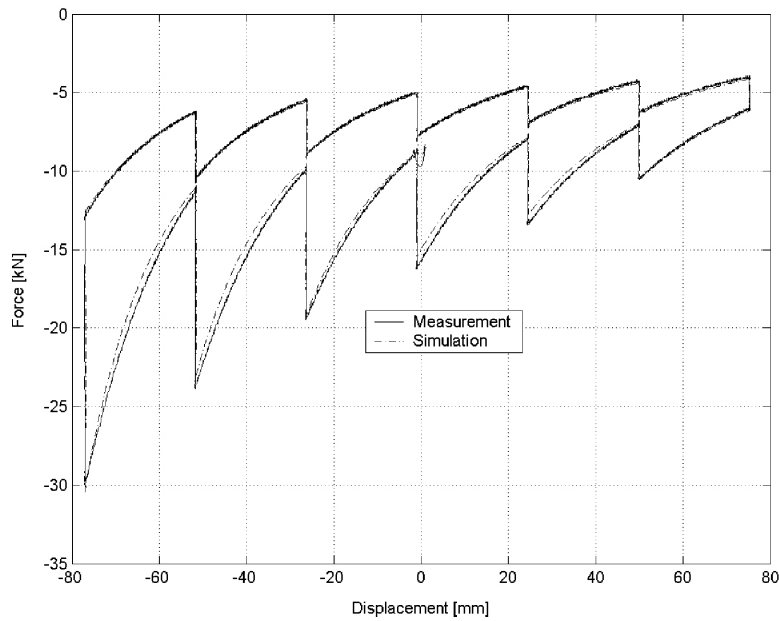
Figure 19 Comparison between measured and calculated force–displacement curve: incremental compression test with low damping

It is also worth noting that the slow drop in pressure P_1 right after achieving the local peaks at the end of the compression strokes in Figure 17, just before valve 3 is opened, is predicted quite well by the model. Since the displacement input does not vary in this period, it is evident that the cooling of gas in accumulator 1 causes this pressure drop. This effect is captured adequately in the model by the use of equation (3).

Since the displacement of the suspension unit is not taken as an input in the mathematical model, it is necessary to check that the displacement of the unit that would be mandated by the solution of the differential equations like equation (1) does, in fact, correspond to the actual displacement experienced by the unit. During all the tests described in this section, this was, in fact, checked and the correlation was exceptionally good. At this time, it is proposed that a similar check should be incorporated in an implementation of the SIMULINK model within an ADAMS simulation of vehicle dynamics.

To date, no measurement was done to specifically validate the way that the pressure-dependent valve switching was implemented in the mathematical model. This issue will be addressed during the development of strategies for controlling the suspension unit in a vehicle.

5   Conclusion

The mathematical model of a suspension unit consisting of a two-stage, semi-active, hydropneumatic spring, combined with a two-stage, semi-active damper was presented. This model was implemented in SIMULINK. Both a first and second prototype of this unit were built and tested, and some of the results measured on the second prototype were

compared with results predicted by the mathematical model. Agreement between the model predictions and the measurements was generally good. Some aspects where the model or the quantifying of its parameters needs improvement were identified. In particular, the tests to date clearly identified the need for a better way of quantifying the mass of gas loaded into the two accumulators.

Further work will be done on testing the model within simulations of first a quarter car and finally a full vehicle equipped with these suspension units. These results will be compared with road test measurements on a corresponding vehicle. Before this can become a reality, though, a suitable control strategy for the whole suspension system comprising four of these suspension units need to be developed. In the mean time, a SIMULINK model comprising four of these units was successfully incorporated into an ADAMS vehicle dynamics model of the sport utility vehicle in question, and the model was in simulation driven over a stretch of rough road at various valve settings. In all cases, the simulation runs were completed successfully and the results seemed plausible.

Acknowledgements

References

Birch, S., Yamaguchi, J. and Demmler, A. (1990) 'Tech briefs – concepts – Citroën's Activa 2', *Automotive Engineering*, Vol. 98, No. 2, pp.55, 56.

Decker, H., Schramm, W. and Kallenbach, R. (1988) 'A practical approach towards advanced semi-active suspension systems', *Proceeding of the ImechE – Advanced Suspension*, 24–25 October, London, pp.93–99.

Els, P.S. and Grobbelaar, B. (1999) 'Heat transfer effects on hydropneumatic suspension systems', *Journal of Terramechanics*, Vol. 36, pp.197–205.

Hine, P.J. and Pearce, P.T. (1988) 'A practical intelligent damping system', *IMechE*, Vol. C436, No. 88, pp.141–147.

Hirose, M., Matsushige, S., Buma, S. and Kamiya, K. (1988) 'Toyota electronic modulated suspension system for the 1986 soarer', *IEEE Transactions on Industrial Electronics*, Vol. 35, No. 2, pp.193–200.

Karnopp, D. and Margolis, D. (1984) 'Adaptive suspension concepts for road vehicles', *Vehicle System Dynamics*, Vol. 13, pp.145–160.

Mizuguchi, M., Chikamari, S., Suda, T. and Kobayashi, K. (1984) *Electronic Controlled Suspension (ECS)*, SAE Technical Paper 845051, Society of Automotive Engineers, Warrendale.

Wallentowitz, H. and Holdman, P. (2001) *Hardware and Software Demands on Adjustable Shock Absorbers for Trucks and Passenger Cars*, http://www.ika.rwth-aachen.de/ueberuns/publikat/ ph-hdt/hdt-e-95.htm, Accessed 30 April, 2003.