Chapter 5

Simplified Vehicle Models

The work discussed in Chapter 4 suggested that central finite differences rather than forward finite differences should be used for gradient calculation. This, however, implies more function evaluations per iteration. To circumvent additional costs related to more function evaluations, and the high level of noise present, the feasibility of using simplified models for gradient evaluation is investigated in this chapter. Proposed is the use of carefully chosen simplified numerical models of the vehicle dynamics for computing gradient information, and a detailed vehicle model for obtaining objective function values at each iteration step. It is proposed that a non-linear pitch-plane model, be used for the gradient information, when optimising ride comfort. When optimising for handling, the use of a non-linear bicycle model, that includes roll, is suggested. The gradients of the objective function and constraint functions are obtained through the use of central finite differences, within Dynamic-Q, via numerical simulation using the proposed simplified models. The importance of correctly scaling these simplified models is emphasised. The models are validated against the full simulation model.

5.1 Optimisation Procedure

The use of simplified numerical models of the full vehicle model, for the determination of gradient information, is investigated. Although the Dynamic-Q optimisation method is used, the principle can be applied to any gradient-based optimisation method. For the determination of the required

first order gradient information central finite differencing is used. Central finite differencing was found to significantly improve the gradient based optimisation process, as discussed in Chapter 4. If the simplified vehicle models can be used for the determination of the gradient information, the number of numerically expensive simulations of the full vehicle model can be reduced to one per iteration, as it is only required to obtain the objective and constraint function values. This has the advantage that the total optimisation time can be greatly reduced, as the analysis of the simplified models take approximately 10% of the simulation time of the full vehicle model. Traditionally the use of central finite differences would have resulted in $2n+1$ full simulations per iteration, where n is the number of design variables. In this case the optimisation takes effectively, in terms of computational time, $2n$ times 0.1 for the gradient evaluation and 1 for the objective function evaluation resulting in an equivalent $0.2n+1$ function evaluations per iteration.

5.2 Definition of Optimisation Parameters

Before the optimisation can be performed, the design variables, objective functions, and constraints need to be defined and scaled. These need to be defined before the simplified models can be developed.

5.2.1 Definition of Design Variables

As before the assumption is made that the left and right suspension settings will be the same, but that front and rear settings may differ. The design variables chosen for optimisation are therefore the static gas volume of the accumulator (Figure 5.1), and damper force scale factor (Figure 5.2), on both the front and rear axles. Thus there are two design variables per axle.

For this initial study the standard rear damper force characteristic is multiplied by a factor which constitutes the damping design variable (Figure 5.2). The general shape and switch velocities of the damper are thus kept the same. This chapter only considers the cases of two and four

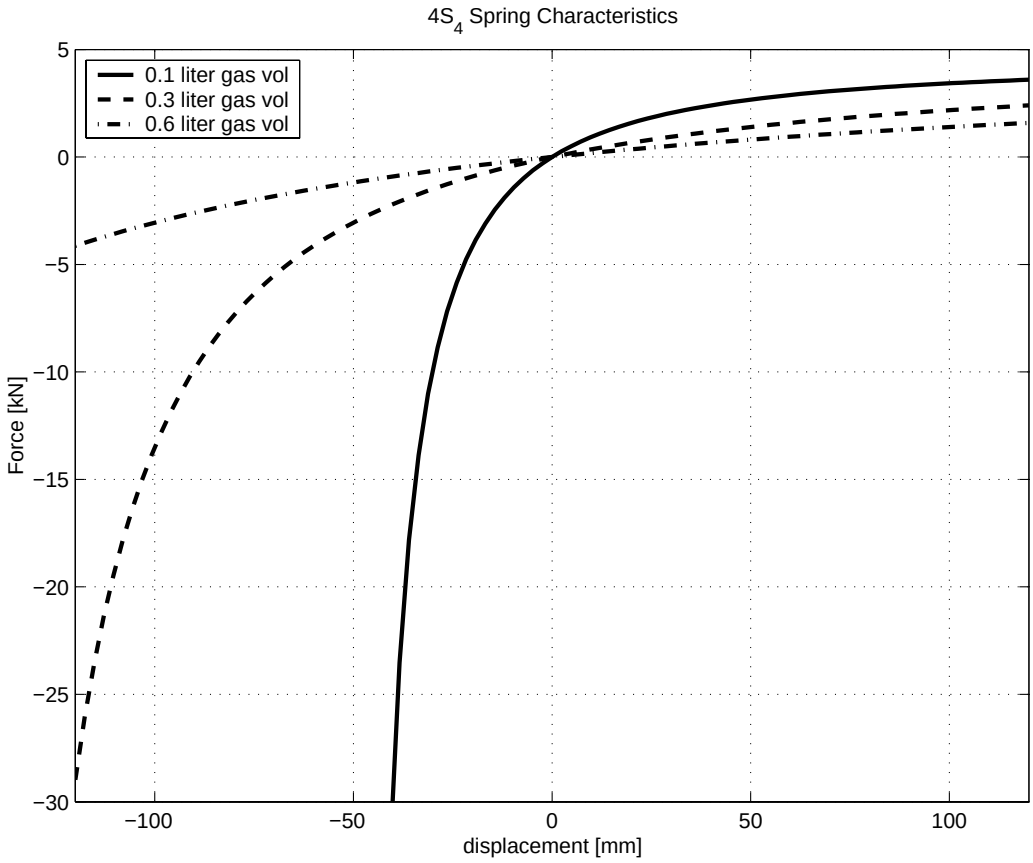

Figure 5.1: Definition of $4S_4$ spring characteristics for various gas volumes

design variables, which respectively corresponds to the case where the spring and damper characteristics are identical for the front and rear axles (two design variables), and where they may differ for front and rear (four design variables).

The static gas volume of the accumulator is denoted by $gvol$, and allowed to range from 0.1 to 0.6 $liters$. The range is dictated by the smallest and largest gas volumes that are possible with the current $4S_4$ unit. The damper force scale factor is denoted by $dpsf$, and allowed to range from 0.1 to 3. The range is again determined by the current design limits of the $4S_4$ unit (see paragraph 1.2). The design variables are normalised to allow a range from 0.001 to 1 in magnitude, which are accordingly chosen as upper and lower bounds. The normalisation of the design variables is generally sound optimisation practice, to ensure that the problem to be solved by the optimisation algorithm, is not poorly scaled. Poor scaling results in optimisation difficulties, and poor

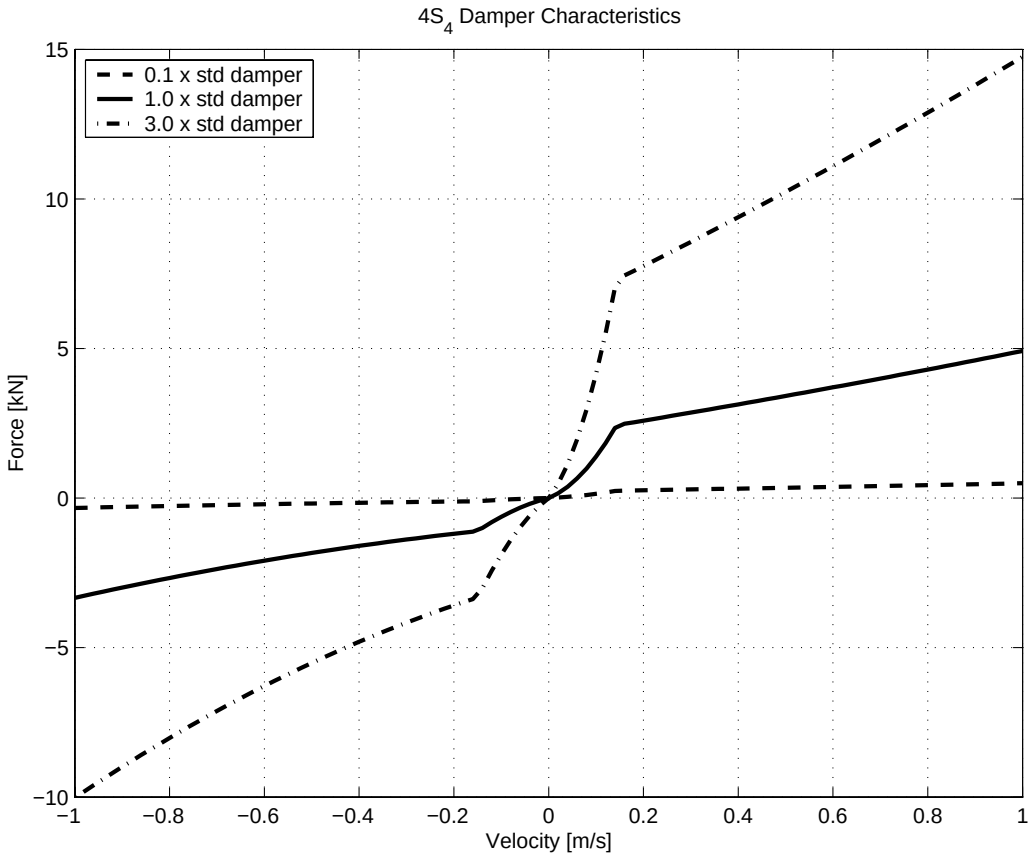

Figure 5.2: Definition of $4S_4$ damper characteristics for various damper scale factors

convergence, and could be a reason for the difficulties encountered in Chapter 4. The i^{th} design variable x_i is defined as a ratio of: the parameter's current value $v_{current}$, the lowest permissible value v_{low}, and the highest permissible value v_{high}, as follows:

$$x_i = \frac{v_{current} - v_{low}}{v_{high} - v_{low}} \tag{5.1}$$

The design variables are then explicitly defined as follows:

$$x_1 = \frac{dpsf - 0.1}{3 - 0.1}, \qquad x_2 = \frac{gvol - 0.1}{0.6 - 0.1} \tag{5.2}$$

with bounds

$$0.001 \leq x_i \leq 1, \qquad i = 1, 2 \tag{5.3}$$

For the four design variable problem the front and rear settings are uncoupled, meaning that there are separate front and rear damper scale factors and

front and rear spring static gas volumes. This results in two design variables describing the front and two describing the rear, giving four design variables in total.

The front damper scale factor is denoted by $dpsff$, the front static gas volume by $gvolf$, the rear damper scale factor by $dpsfr$, and the rear static gas volume by $gvolr$. These design variables are also allowed to range from 0.001 to 1 in magnitude. Thus the design variables are defined explicitly as follows:

$$x_1 = \frac{dpsff-0.1}{3-0.1}, \qquad x_2 = \frac{gvolf-0.1}{0.6-0.1}$$
$$x_3 = \frac{dpsfr-0.1}{3-0.1}, \qquad x_4 = \frac{gvolr-0.1}{0.6-0.1} \tag{5.4}$$

with bounds

$$0.001 \leq x_i \leq 1, \qquad i = 1, ..., 4 \tag{5.5}$$

5.2.2 Definition of Objective Functions

For ride comfort the motion of the vehicle is simulated for travelling in a straight line over the local Belgian paving, and the sum of the driver a_{zRMSd} and passenger a_{zRMSp} frequency weighted (according to British Standard BS6841 1987) root mean square (RMS) vertical accelerations are used for the objective function. This was found to be a sufficiently representative measure of passengers' subjective comments by Els (2005). The Belgian paving test track used, is located at the Gerotek Test Facilities (Gerotek 2006), and has a ISO8608 (1995) roughness coefficient G_{do} of 1×10^{-4} $m^2/(cycles/m)$, and a terrain index ω of 4 (Thoresson 2003).

Following sound optimisation practice the objective function is also scaled as for the design variables to range between zero and one (equations 5.2 to 5.5). This is done by assuming that the maximum and minimum objective function values will lie on one of the corners of the design space. The four corners for the two design variable case were evaluated. The maximum vertical RMS acceleration was found to be 4.4 m/s^2, and the minimum to be 0.7 m/s^2. RMS accelerations are then scaled so that the expected maximum

and minimum values lie between zero and one. The ride comfort objective function $f_{ride}(x)$, is defined as the sum of the scaled driver and passenger accelerations divided by two, as follows:

$$f_{ride}(x) = \frac{\sum(\frac{a_{zRMSd}-0.7}{4.4-0.7}, \frac{a_{zRMSp}-0.7}{4.4-0.7})}{2} \tag{5.6}$$

The handling objective function is defined as the sum of the normalised first peak value of the body roll angle $\varphi_{1^{st}peak}$ for the first lane change (Els and Uys 2003) of the ISO3888-1 (1999) double lane change manœuvre, and the normalised RMS roll velocity $\dot{\varphi}_{RMS}$ for the whole double lane change manœuvre. The RMS roll velocity is now used in addition to the roll angle, so as to have a measure of the transient stability of the vehicle in roll, which was previously not considered in Chapter 4. The handling objective function $f_{hand}(x)$ is defined as the sum of these normalised parameters divided by two, as follows:

$$f_{hand}(x) = \frac{\sum(\frac{(\dot{\varphi}_{RMS}-0.8)0.9}{5.7-0.8} + 0.1, \frac{(\varphi_{1^{st}peak}-1.4)0.9}{12.2-1.4} + 0.1)}{2} \tag{5.7}$$

5.2.3 Definition of Inequality Constraint Functions

Tyre hop effects need to be considered when optimising for ride comfort, as the damping design variables tend to be sensitive to tyre hop (Uys et al. 2006b). In the preliminary study discussed in Chapter 4, it was found that the optimal ride comfort was found at the expense of vehicle stability on the road, thus necessitating the consideration of tyre hop. The requirement was introduced, that the tyre could only be permitted to loose contact with the ground for 10% or less of the simulation time, when considering typical off-road and rough terrain. The time the tyre has lost contact with the ground was determined by observing when the tyre's vertical force F_{ztyre_i} is equal to zero. The tyre hop effect is added as inequality constraints for each individual tyre i as follows:

$$g_i(x) = 10(\frac{\sum t(F_{ztyre_i} = 0)}{t_{total}} - 0.1) \leq 0 \tag{5.8}$$

The factor of 10 was used to better scale the tyre hop constraint between minus one and one.

Suspension working space was not included as an inequality constraint as the non-linear bump and rebound stops are included in the simulation models. Thus the simulation models will restrict the suspension working space.

5.3 Simplified Vehicle Models

The need for simplified models to obtain smoother (less noisy) gradient information, is justified by the high amplitude noise inherently present in the MSC.ADAMS simulation model, as illustrated in Figure 5.3. This figure reflects the change in the ride comfort objective function value for a change in only the front damper design variable x_1. This was performed at the center of the design space. It can be seen that the noise in relation to the objective function value is severe, especially when considering the tyre hop constraint values. Figure 5.4 represents the objective and constraint values for changes in the front damper design variable, for the simplified ride comfort vehicle model, discussed in detail in paragraph 5.3.2. It can be seen that the noise present in the objective function is greatly reduced, although no significant benefit is observed when considering the constraint functions. It is speculated, that this is attributed to the low tyre damping, which results in unstable tyre dynamics.

5.3.1 Handling Model

For the simplified vehicle handling model it is assumed that the vehicle drives on a smooth surface, and uses exactly the same steering input as the MSC.ADAMS model for that iteration. The model consists of two parts, namely the lateral and yaw dynamics, and then the resulting roll dynamics of the body. For the formulation of the equations of motion for the simplified handling model, Figures 5.5 and 5.6 are considered. The model is simplified so that only three degrees of freedom are considered, namely: body roll φ, vehicle yaw ψ and vehicle lateral displacement y. The assumption will be made that the vehicle will drive at a constant longitudinal velocity $\dot{x}$ along

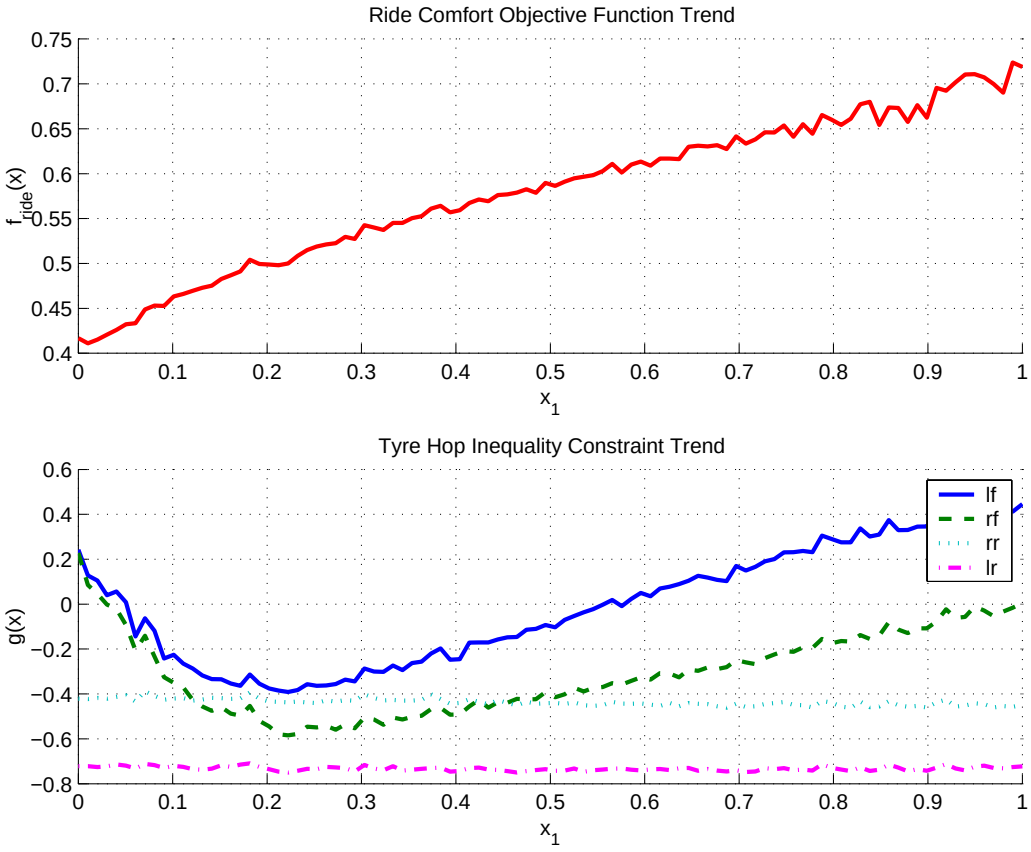

Figure 5.3: Level of inherent numerical noise in objective function and inequality constraints, for change in front damper design variable x_1, for full vehicle MSC.ADAMS model

the vehicle's x-axis. Looking at the top view of the vehicle (Figure 5.5) the overall yaw and lateral equations of motion can be formulated. For yaw:

$$\sum M_z = I_z\ddot{\psi} = a(F_{y1} + F_{y2}) - b(F_{y3} + F_{y4}) \tag{5.9}$$

where it is assumed that the steer angle δ is small (i.e. $F_{yi}cos(\delta) \approx F_{yi}$). Thus the full lateral tyre force F_{yi} acts along the y-axis. Also the longitudinal component of the lateral tyre force is low in magnitude and can be ignored. For the lateral direction:

$$\sum F_y = m_v\ddot{y}_v = F_{y1} + F_{y2} + F_{y3} + F_{y4} \tag{5.10}$$

Similarly by considering Figure 5.6 the equation of motion for the body roll about the body *cg* can be formulated as follows:

$$\sum M_x = I_x\ddot{\varphi} = (f_{4S_{4l}} - f_{4S_{4r}})\frac{t_s}{2} + h_{cg}(F_{yl} + F_{yr})\frac{m_b}{m_v} \tag{5.11}$$

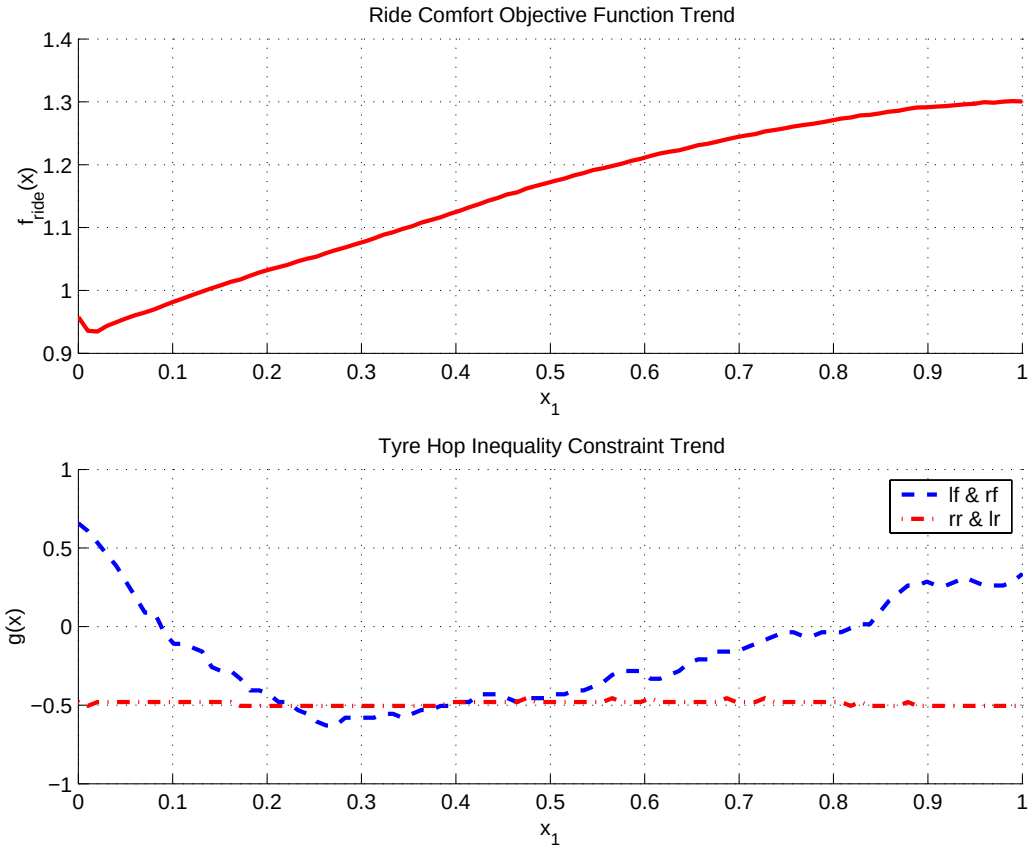

Figure 5.4: Level of inherent numerical noise in objective function and inequality constraints, for change in front damper design variable x_1, when considering the simplified MATLAB model

The mass ratio $\frac{m_b}{m_v}$ is introduced so that the the tyres' lateral force effect on the vehicle body can be uncoupled from the axles and wheels, as the body motion is what our suspension can control. This was done so as to decrease the number of degrees of freedom to be calculated, helping to speed up simulation time. The left $f_{4S_{4l}}$ and right $f_{4S_{4r}}$ suspension forces are the sum of the suspension forces on the respective side. Similarly the left F_{yl} and right F_{yr} lateral forces are the sum of the lateral tyre forces for the respective side. The lateral forces are calculated by taking the vertical load and slip angle for the tyre, as inputs to the 'Magic Formula' Pacejka'89 (Bakker et al. 1989) tyre model using the same coefficients as for the full vehicle simulation model. For this model the following simplifications have been applied:

- The tyre lateral force produces a minimal longitudinal component that

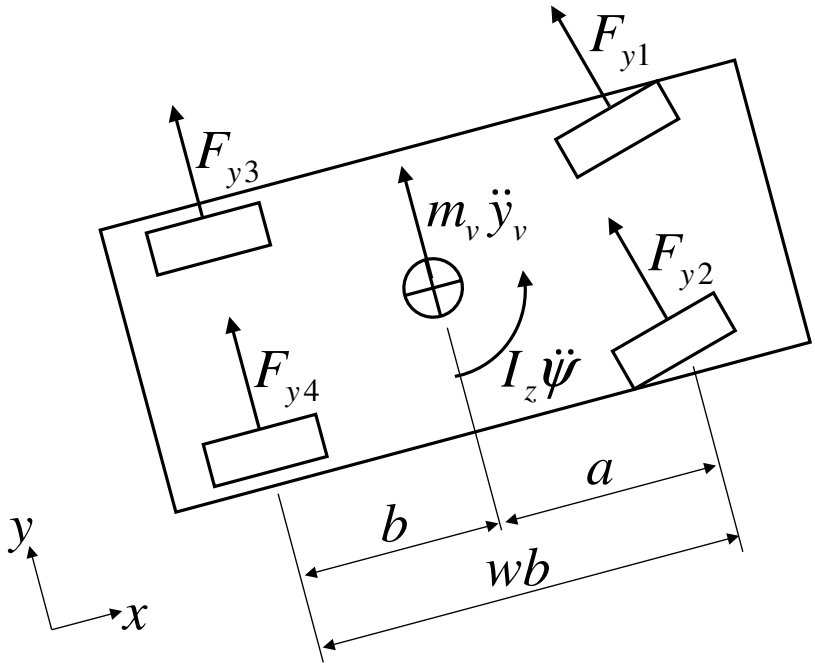

Figure 5.5: Top view of vehicle in handling manœuvre

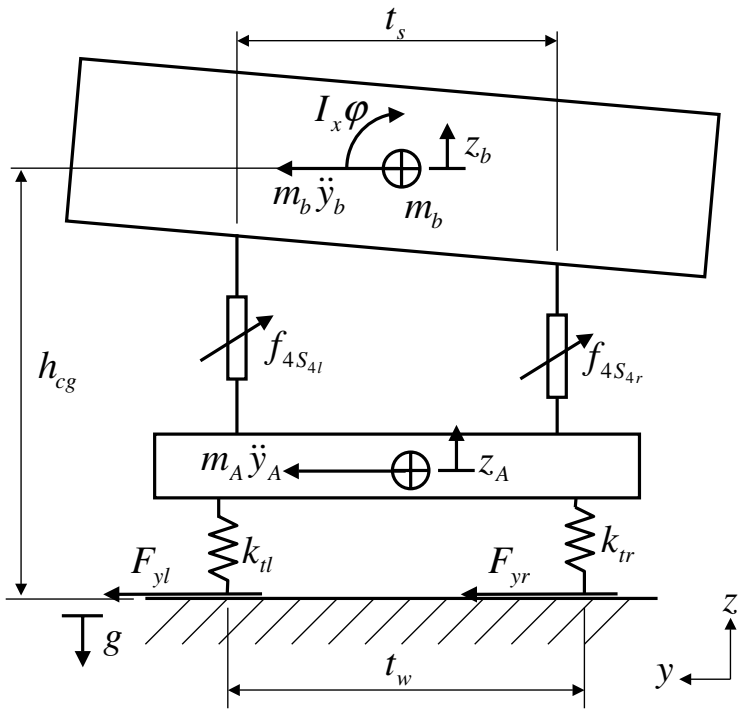

Figure 5.6: Rear view of vehicle indicating body roll

is taken up by the longitudinal driving force and can be ignored.

- No longitudinal effects except vehicle speed are considered.
- Nothing can be done about the tyre deflection and the angle that the axle makes with respect to the ground, for this reason the axle roll effects, due to tyre deflection are ignored.
- The MSC.ADAMS calculated steer angle is used as the input steer angle for the MATLAB simulation.
- Vertical tyre forces are taken as being the same mass proportion front to rear as the static case, of the side suspension force. (i.e. no longitudinal load transfer)

The simplified handling model is thus a significant simplification of the actual vehicle dynamics. It will be shown to still return very good trends when compared to the full vehicle simulation model.

5.3.2 Ride Comfort Model

For the simplified ride comfort vehicle model a simple pitch plane vehicle model, similar to that used by (Eberhard et al. 1995, Etman et al. 2002, Naudé and Snyman 2003a) and many other's, is used. The measured rough road profile seen by the full vehicle model's wheels is averaged left and right to give an effective centerline profile. The pitch plane model then follows the averaged path using a point follower tyre model. The basic layout of the simplified model is indicated in Figure 5.7. The equations describing the vehicle behaviour are derived as follows. Consider the forces acting on the front unsprung mass m_{tf}, as a result of the road disturbance input z_{rf}. The summation of vertical forces on the unsprung masses leads to:

$$\sum F_z = m_{tf}\ddot{z_3} = 2k_{tf}(-z_3+z_{rf}+\delta_{stat})+2c_{tf}(-\dot{z_3}+\dot{z_{rf}})-m_{tf}g-2f_{4S_{4f}} \quad (5.12)$$

for the front, and similarly for the rear:

$$\sum F_z = m_{tr}\ddot{z_4} = 2k_{tr}(-z_4+z_{rr}+\delta_{stat})+2c_{tr}(-\dot{z_4}+\dot{z_{rr}})-m_{tr}g-2f_{4S_{4r}} \quad (5.13)$$

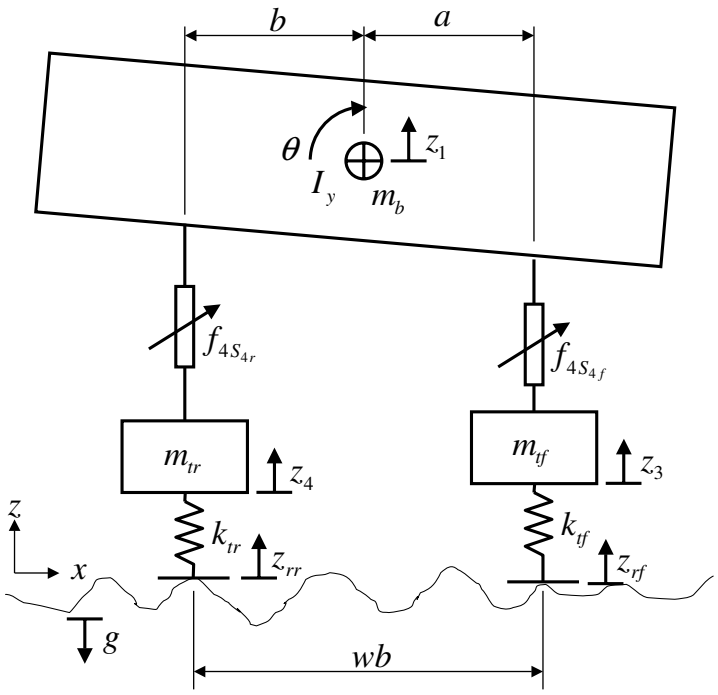

Figure 5.7: Simple pitch-plane vehicle model

where the 2 relates to the fact that there is both a left and a right $4S_4$ strut. It is taken that $g = 9.81m/s^2$. The m_{tf} is the total front axle unsprung mass including the two tyres. And $f_{4S_{4f}}$ is the $4S_4$ front suspension force which is a function of the displacement of the vehicle body m_b and the unsprung mass:

$$f_{4S_{4f}} = f(z_3 - z_1 + \theta a, \dot{z}_3 - \dot{z}_1 + \dot{\theta} a) \tag{5.14}$$

The rear suspension force $f_{4S_{4r}}$ can similarly be defined as:

$$f_{4S_{4r}} = f(z_4 - z_1 - \theta a, \dot{z}_4 - \dot{z}_1 - \dot{\theta} a) \tag{5.15}$$

The tyre spring stiffness and damping are only active while the tyre is in contact with the ground thus the following if statement also applies:

$$\begin{aligned} if \quad & z_3 - z_{rf} - \delta_{stat} < 0 \\ then \quad & k_{tf} = k_t \qquad c_{tf} = c_t \\ else \quad & k_{tf} = 0 \qquad c_{tf} = 0 \end{aligned} \tag{5.16}$$

For the sprung mass m_b two equations of motion are applicable, first for vertical motion:

$$\sum F_z = m_b \ddot{z}_1 = m_b g - 2f_{4S_{4f}} - 2f_{4S_{4r}} \tag{5.17}$$

and then for pitch motion:

$$\sum M_y = I_y\ddot{\theta} = a2f_{4S_{4f}} - b2f_{4S_{4r}} \tag{5.18}$$

These equations of motion can be manipulated as follows:

$$\begin{aligned}
-m_b\ddot{z_1} &= -m_b g + 2f_{4S_{4f}} + 2f_{4S_{4r}} \\
-I_y\ddot{\theta} &= -a2f_{4S_{4f}} + b2f_{4S_{4r}} \\
m_{tf}\ddot{z_3} + 2k_{tf}z_3 + 2c_{tf}\dot{z_3} &= 2k_{tf}(z_{rf} + \delta_{stat}) + 2c_{tf}\dot{z_{rf}} - m_{tf}g - 2f_{4S_{4f}} \\
m_{tr}\ddot{z_4} + 2k_{tr}z_4 + 2c_{tr}\dot{z_4} &= 2k_{tr}(z_{rr} + \delta_{stat}) + 2c_{tr}\dot{z_{rr}} - m_{tr}g - 2f_{4S_{4r}}
\end{aligned} \tag{5.19}$$

This results in a clear set of matrices for mass $\mathbf{M}$, stiffness $\mathbf{K}$, damping $\mathbf{C}$, and force F, which correspond with the formula:

$$\mathbf{M}\ddot{z} + \mathbf{K}z + \mathbf{C}\dot{z} = F \tag{5.20}$$

The above differential equations can be re-arranged, in order to be solved with a numerical integration scheme, as follows:

$$\begin{Bmatrix} \dot{z} \\ \ddot{z} \end{Bmatrix} = \begin{bmatrix} \mathbf{O} & \mathbf{I} \\ -\mathbf{M}^{-1}\mathbf{K} & -\mathbf{M}^{-1}\mathbf{C} \end{bmatrix} \begin{Bmatrix} z \\ \dot{z} \end{Bmatrix} + \begin{Bmatrix} O \\ \mathbf{M}^{-1}F \end{Bmatrix} \tag{5.21}$$

The modelling units of the models are meters and radians. For the execution of the numerical integration of the simplified models, the built-in MATLAB ode15s (Mathworks 2000b) solver is used with a relative tolerance of 1.5 mm and a maximum time step of 0.05 seconds. These simplified models solve in approximately 1 minute depending on design variables chosen while the average MSC.ADAMS model takes at least 10 minutes to solve, on a Pentium 4, 1.8 GHz processor with 1 G RAM.

5.3.3 Handling Model Validation

Figures 5.8 and 5.9 illustrate the comparison between the full vehicle MSC.ADAMS model and the simplified model for the handling objective function parameters, where it should be noted that the colours are for easier visualization purposes only. It can be seen that the simplified model does not display all the information of the full vehicle model, but the global optimum and maximum are the same. In general the trends are very similar, while only

varying in absolute values. The MATLAB handling model is thus scaled so as to give a better approximation of the MSC.ADAMS full vehicle model. For the scaling of the MATLAB simplified models, the two design variables were considered and 30 function evaluations were performed over the design space using the full MSC.ADAMS simulation model and the simplified model. The results for the simplified model were then scaled so that the surfaces coincided over most of the design space.

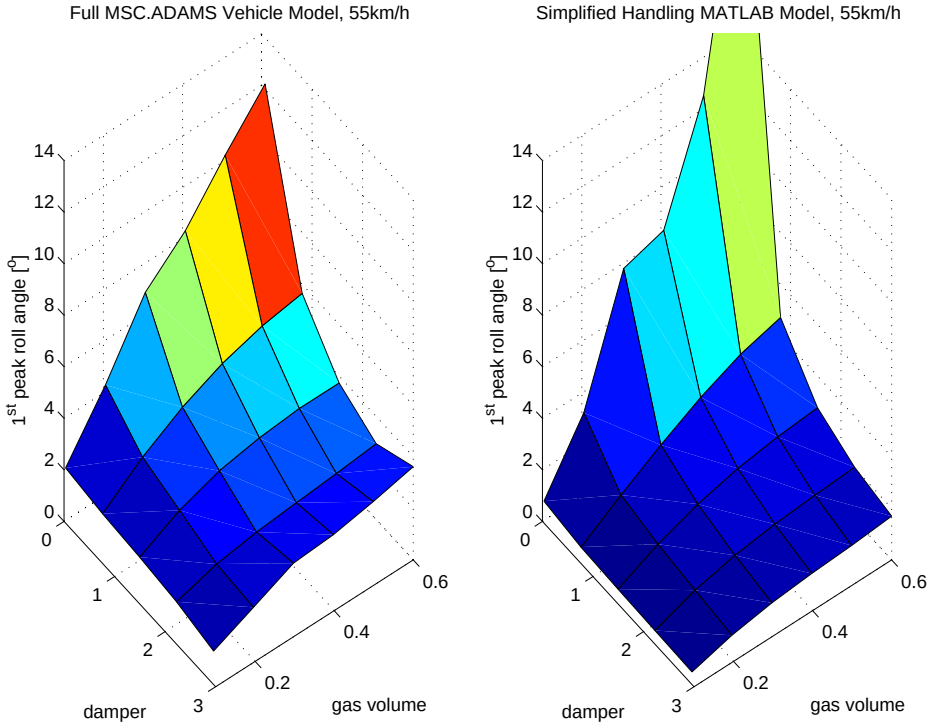

Figure 5.8: Validation of 1st peak roll angle over design space, for double lane change.

5.3.4 Ride Comfort Model Validation

The simplified MATLAB model for ride comfort was evaluated against the full MSC.ADAMS vehicle model to investigate whether the gradient closely matched that of the MSC.ADAMS model. The sum of the vertical weighted accelerations was normalised in both cases so that the objective function value would range from zero to one. Figures 5.10 to 5.12 illustrate the close correlation achieved when observing the effect of the design parameters on the objective function and the tyre hop effect.

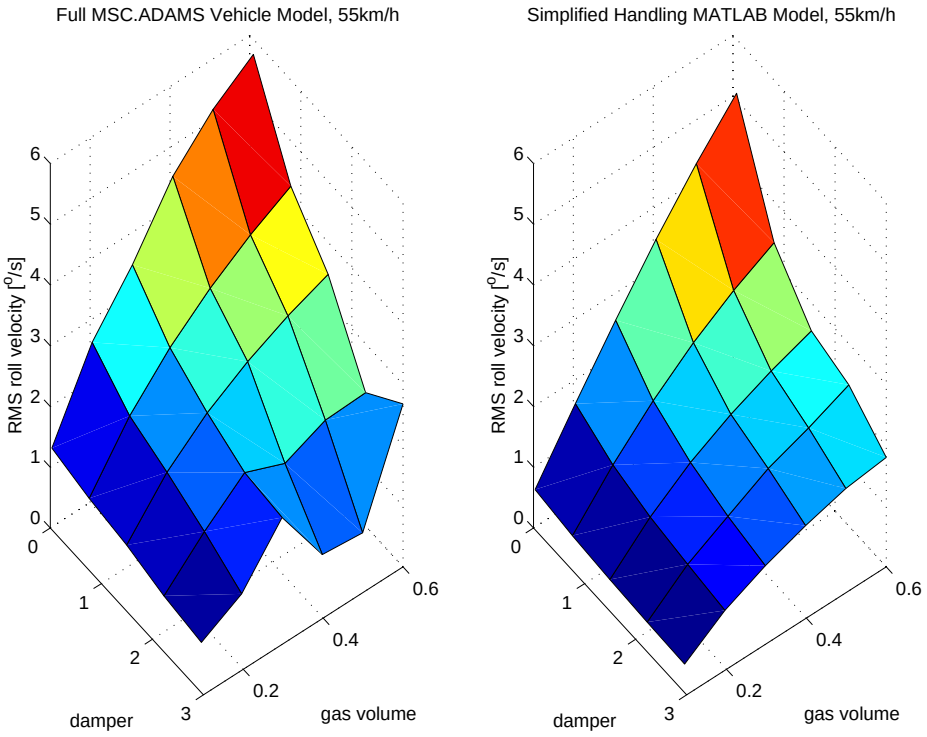

Figure 5.9: Validation of RMS roll velocity over design space, for double lane change.

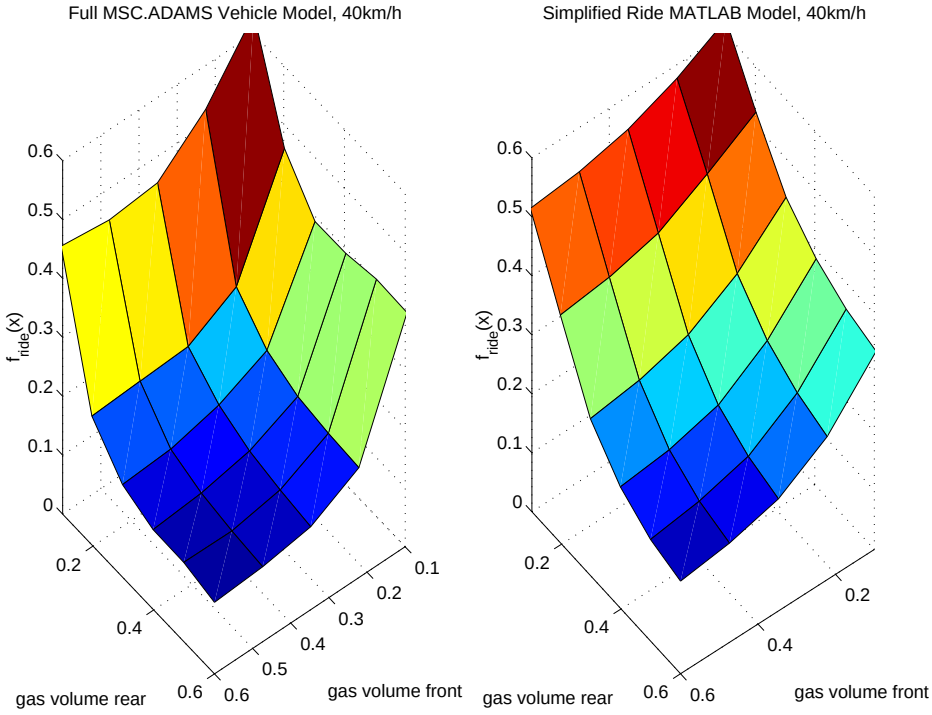

Figure 5.10: Model validation of ride comfort for differing front and rear gas volumes

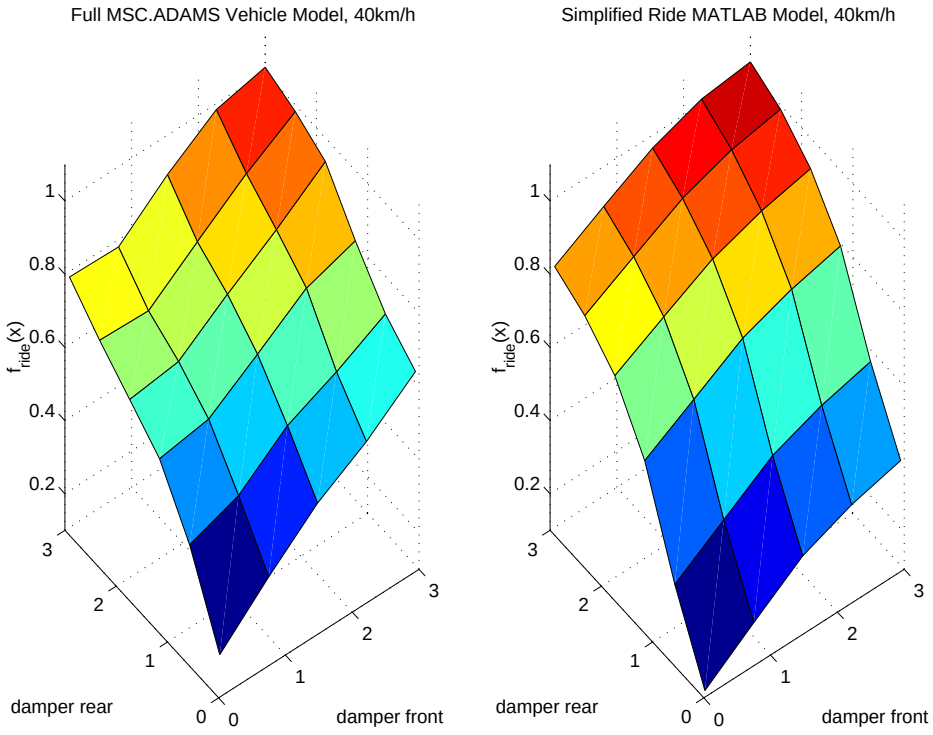

Figure 5.11: Model validation of ride comfort for differing front and rear damper scale factors

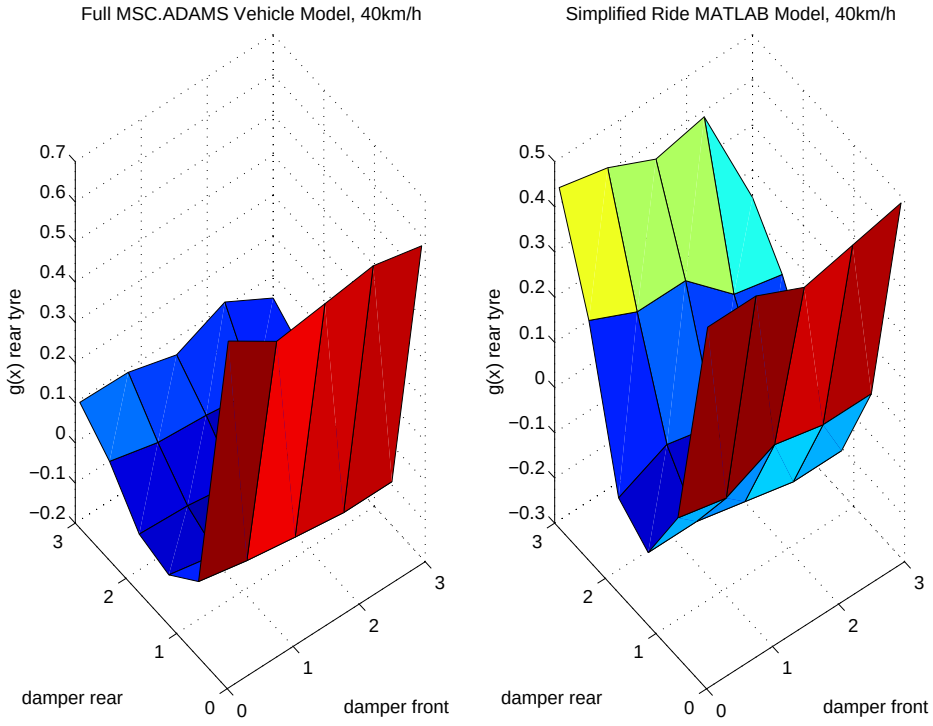

Figure 5.12: Model validation of ride comfort for differing front and rear damper scale factors: effect on rear tyre hop

5.4 Conclusions

In this chapter the combined use of simplified numerical vehicle models and computationally expensive full vehicle simulation models in gradient-based optimisation algorithms, for vehicle suspension optimisation was investigated.

In particular the specific optimisation methodology to be used is described and the objective functions and design variables are defined. The full vehicle, modelled in MSC.ADAMS, returns excellent correlation with measured results as presented in Chapter 3. However, this model is computationally expensive and exhibits severe numerical noise.

In order to help overcome the problems associated with high computational cost and numerical noise in the optimisation process, the use of simplified models of the vehicle is suggested. These models exhibit very similar trends to the full vehicle simulation model, however, the absolute values are not the same. It is also important to note that the constraints, especially the tyre hop constraints, do not necessarily cross the zero axis at the correct points, even though the gradient trends are very similar. The required scaling of the simplified models to be more representative of the full vehicle model is presented. The cost of this scaling must be taken into account when optimising. Here 30 expensive full vehicle model simulations per simplified model were performed. The simplified model's objective functions were suitably scaled, to be representative of the full simulation model's objective function values. Once scaled, the simplified models are representative of the full vehicle simulation model, but exhibit significantly less numerical noise, and solve significantly faster.

Chapter 6 investigates the implementation of the simplified models in the optimisation procedure. The optimisation results using the full simulation vehicle model throughout, will be compared to that obtained using the simplified models for computation of the gradient information. The use of the simplified models for optimisation information as well as the full simulation

model is known as multi-fidelity optimisation, and further discussed in Chapter 6.

Chapter 6

Multi-Fidelity Optimisation

Chapter 5 proposed a methodology for the efficient determination of gradient information, when optimising for a vehicle's suspension characteristics. The non-linear full vehicle model, and simplified models for gradient information have been discussed, and validated. Chapters 2 and 4 presented a brief history of vehicle suspension optimisation, the general problem of numerical noise, and computationally expensive simulation models. Proposed is the use of simplified mathematical models for calculating gradient information, and the full simulation model for determining the objective function value when optimising an off-road vehicle's suspension characteristics. Although this application uses the gradient-based optimisation algorithm Dynamic-Q, the principle can be applied to any gradient-based optimisation algorithm.

In this chapter, the simplified models presented in Chapter 5 are used for gradient information simulations, in the optimisation of the vehicle's suspension characteristics, for ride comfort and handling. The simplified vehicle models for handling and ride comfort, as described in Chapter 5, are used to decrease the computational complexity of the full vehicle simulation model, while still capturing the trends over the design space. The convergence histories of the optimisation are compared to those obtained when only the full, computationally expensive, vehicle model is used. For illustration of the proposed gradient-based optimisation methodology, up to four design variables are considered in modelling the suspension characteristics.

The proposed methodology is found to be an efficient alternative for the optimisation of the vehicle's suspension system. The undesirable effects associated with noise in the gradient information is effectively reduced, in the optimisation process. Substantial benefits are achieved in terms of computational time needed to reach a solution.

6.1 Optimisation Procedure

This chapter compares the optimisation results when using the full vehicle simulation model for objective function value and gradient information (*admsgrad*), as traditionally used in gradient-based optimisation, to the use of the full vehicle model for only objective function value, and the simplified models for gradient information (*matgrad*). Central finite differences, at a computational cost of $2n + 1$ function evaluations per iteration (where n is the number of design variables), is used for the determination of the gradient information. The use of central finite differences for gradient information, was found to improve optimisation convergence in the presence of severe numerical noise by Els et al. (2006), and discussed in Chapter 4.

The use of only the MSC.ADAMS full vehicle model in the optimisation (*admsgrad*) has a computational cost of $2n + 1$ computationally expensive simulations per iteration. The use of the MSC.ADAMS full vehicle model for only the objective function value, and the simplified MATLAB vehicle models for gradient information (*matgrad*), has a computational cost of one computationally expensive simulation per iteration, and $2n$ computationally inexpensive simulations per iteration. The simplified MATLAB models solve in approximately 10% of the full vehicle model's simulation time. Sufficient gradient information is obtained, after the simplified models have been scaled at a once-off cost of 30 computationally expensive simulations.

With the proposed methodology, more starting points or design variables can be efficiently considered, in less computational time, making gradient-based

approximation methods for optimisation of vehicle suspension systems more feasible. The simplified models also exhibit less numerical noise than the full simulation model, resulting in smoother gradient information. For the optimisation, the same normalised design variables as discussed in Chapter 5 and the same normalised objective and constraint functions are used.

6.2 Handling Optimisation Results

Presented in the following subsections is the handling optimisation results for two and four design variables. This is considered first as up to this stage a reasonably firm feeling of the problem has been built with which to test the results. This is to demonstrate the concept before considering the full optimisation design variables.

6.2.1 Two Design Variable Optimisation

The results for the comparison between the *admsgrad* and the *matgrad*, when optimising handling for two design variables, are illustrated in Figure 6.1. It can be seen that the use of the simplified model for the gradient information (*matgrad*) converged to an optimum after 12 iterations and 13 expensive function evaluations. The use of the computationally expensive full vehicle model, for gradient information (*admsgrad*), converged to the same optimum point within 15 iterations, but took 80 computationally expensive function evaluations of the full vehicle model. The simplified model solves in approximately 10% of the solution time of the full MSC.ADAMS vehicle model. Central finite differences is used for the gradient determination, at a cost of $2n + 1$ function evaluations per iteration, where n is the number of design variables. When using only the MSC.ADAMS model for gradient and objective function evaluation (*admsgrad*), one iteration of two design variables costs the equivalent of 500% of the computational time of one MSC.ADAMS model simulation. When using the simplified models for gradient information, and only one full MSC.ADAMS simulation for the objective function value, the cost of one iteration is equivalent to 100% +

2x2x10%, which is the equivalent of 140% of the computational time of one MSC.ADAMS simulation. The use of the simplified models for the determination of gradient information, is approximately 3.5 times faster than using only the MSC.ADAMS model, when considering two design variables. This highlights the advantages in terms of simulation time achievable for just two design variables. It is also observed that the use of the simplified model for gradient information does not introduce instabilities in the optimisation convergence history. The simplified model produces sufficiently accurate gradient information to drive the optimisation to the same optimum.

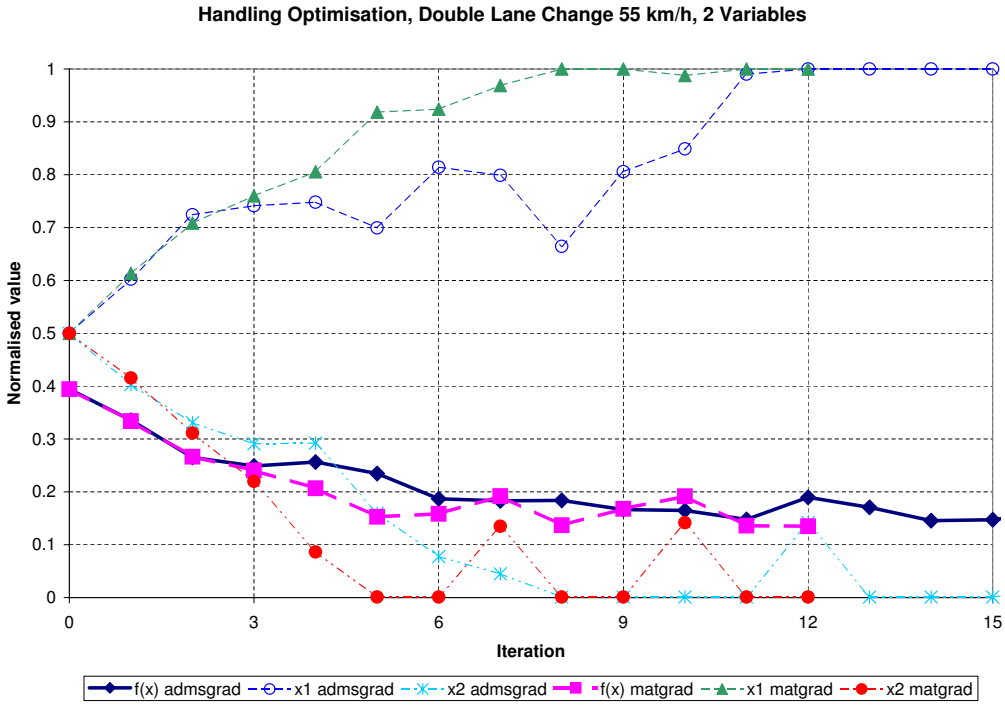

Figure 6.1: Handling optimisation convergence histories for full MSC.ADAMS model, and using the simplified MATLAB model for gradient information, 2 design variables

6.2.2 Four Design Variable Optimisation

With the successful results obtained for the two design variable handling optimisation, the problem was expanded to four design variables, thus allowing the front and rear suspension characteristics to be independent of each other. It is believed that the four design variable problem

will exhibit more local minima, and the use of the simple model for gradient information needs to be tested for robustness. The results of the four design variable optimisation, where the full MSC.ADAMS model was used for gradient information are presented in Figure 6.2. From the figure it can be seen that the optimisation converged to a minimum identical to that for two design variables, considering the noise levels present in the numerical model. It is noted from the optimisation convergence history, that there are repeated equal local minima at iterations five, eight, and ten. It can be seen that the design variable $x1$ (front damper) takes on a value around 0.9, and $x3$ (rear damper design variable) takes on a value of 1. It is also evident that design variable $x4$ (rear gas volume design variable) moved to the boundary, and should be at the lowest value. However, interestingly the front gas volume, design variable $x2$ takes on a value around 0.27 (iteration 5), but can also take on a value around 0.07 (iteration 10).

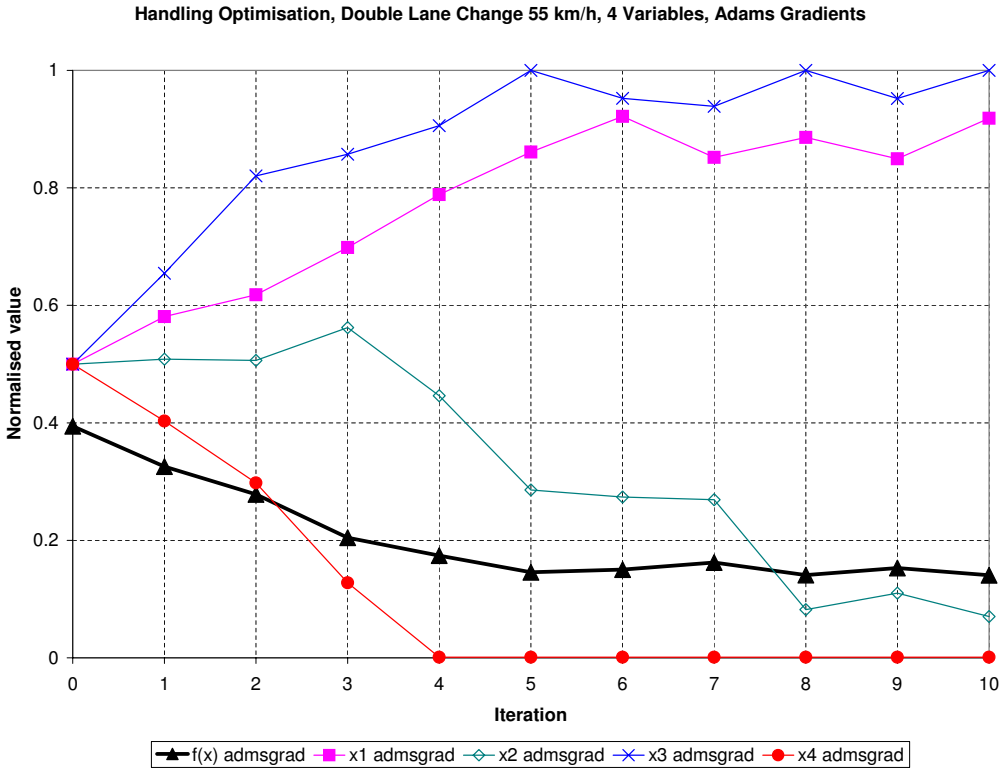

Figure 6.2: Handling optimisation convergence history using the full MSC.ADAMS model for gradient information, 4 design variables

Considering the optimisation convergence history, when the simplified model is used for the gradient evaluations (Figure 6.3), it can be seen that the optimisation process converges to a minimum identical to that for two design variables and four design variables using the MSC.ADAMS model for gradient information. The design variable values converge to different values, indicating the presence of multiple equivalent local minima. From the results it is clear that no difficulties are experienced in obtaining a feasible optimum and that both the solutions are equally feasible. The four design variable optimisation for seven optimisation iterations, using the simplified model, is approximately five times faster than using only the full MSC.ADAMS vehicle model.

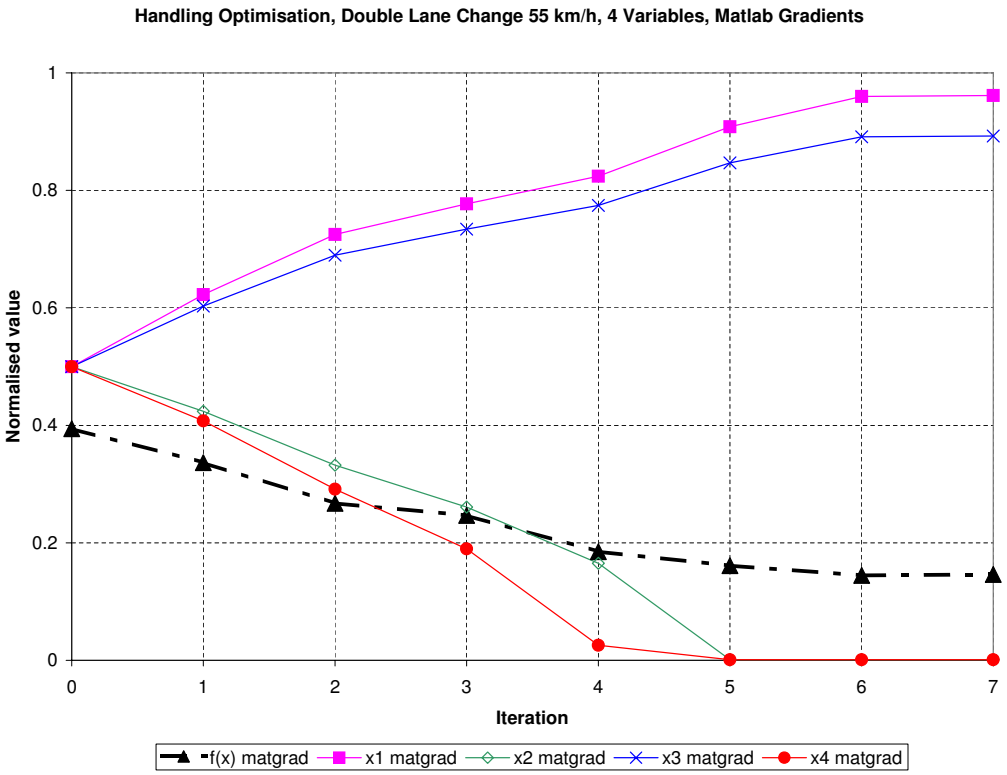

Figure 6.3: Handling optimisation convergence histories using the simplified MATLAB model for gradient information, 4 design variables

6.3 Ride Comfort Optimisation Results

For the ride comfort optimisation the implementation of the tyre hop had to be investigated before the optimisation could be performed. Once the tyre hop had been implemented the ride comfort was optimised considering two and four design variables.

6.3.1 Tyre Hop in the Optimisation Process

The ride comfort optimisation has to be performed considering tyre hop effects, as the vehicle can become unstable on the road should the tyres constantly loose contact with the road, as concluded in Chapter 4. The tyre hop constraints tend to exhibit a more prominent role, than the objective function, on the damping design variable's lower limit. An investigation was performed, to determine the most effective method of including the tyre hop effect within the optimisation process. The following conditions were considered:

- Constrained optimisation: (*constrained*) The objective function is defined as in equation (5.6). The tyre hop constraints are defined as: the individual tyre's vertical force F_{ztyre_i} may not be equal to zero for more than 10% of the total time t_{total}, when travelling on rough off-road terrain, and scaled as follows:

$$g_i(x) = 10(\frac{\sum t(F_{ztyre_i}=0)}{t_{total}} - 0.1) \leq 0, \qquad i = 1, ..., 4 \tag{6.1}$$

  Results are indicated in Figure 6.4.

- Unconstrained optimisation: The objective function is defined as in equation (5.6). The constraints, as defined in equation (6.1), are only monitored, but not considered by the optimisation algorithm, (*unconstrained*). The results are indicated in Figure 6.5.

The equivalent objective function $f(x)_{eq}$ values presented in Figures 6.4 and 6.5 is the ride comfort objective function defined by equation (5.6). The equivalent inequality constraint value $g(x)_{eq}$ is defined as:

$$g(x)_{eq} = max_{i=1,..4}(g_i(x)) \tag{6.2}$$

representing the maximum of the tyre hop constraint function of the four wheels. From the results it can be seen that the constrained optimisation (*constrained*, Figure 6.4), returns the lowest objective function value for the tyre hop inequality constraint being satisfied. In general the front tyres contributed most to the tyre hop, compared to the rear tyres, however, the rear tyres also contributed in the optimisation convergence history, making the inclusion of all tyres as constraints necessary. It was found that a tyre hop limit of 10% for the particular road in question is a reasonable constraint, as smaller limits tend to overconstrain the optimisation. It is thus decided that the tyre hop limit of 10% will be included as a constraint for all future ride comfort optimisation, when travelling over rough off-road terrain.

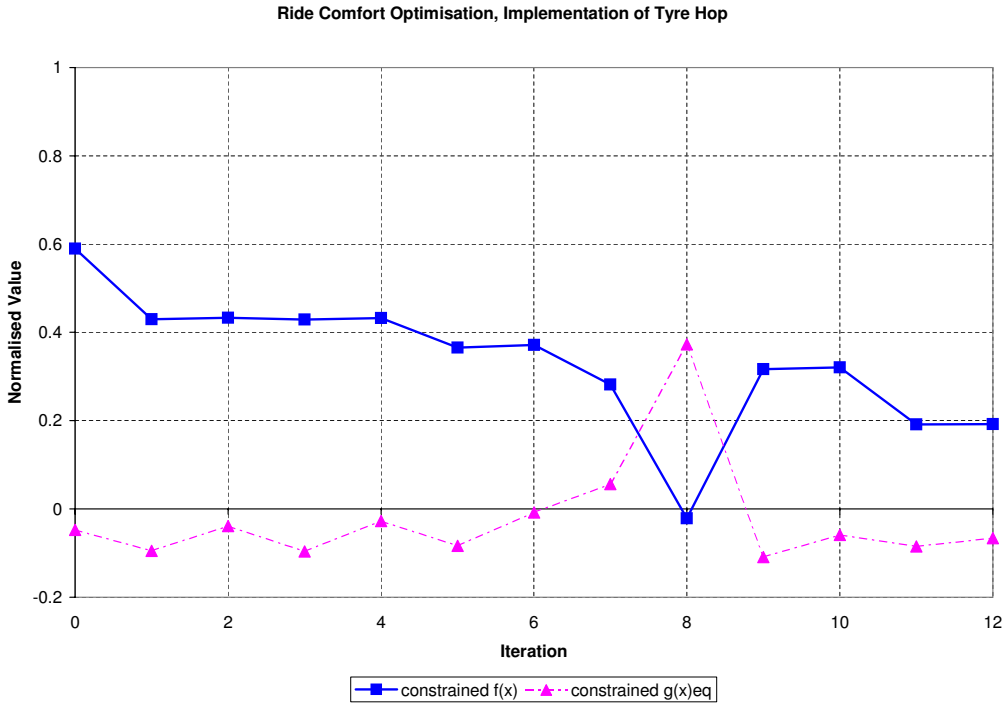

Figure 6.4: Implementing tyre hop as a constraint in ride comfort optimisation

6.3.2 Two Design Variable Optimisation

The vehicle suspension settings were optimised for ride comfort, for two design variables, with the tyre hop constraint included, as defined in equation (6.1). The results of the optimisation process, for using only the MSC.ADAMS model for gradient information, compared to using the

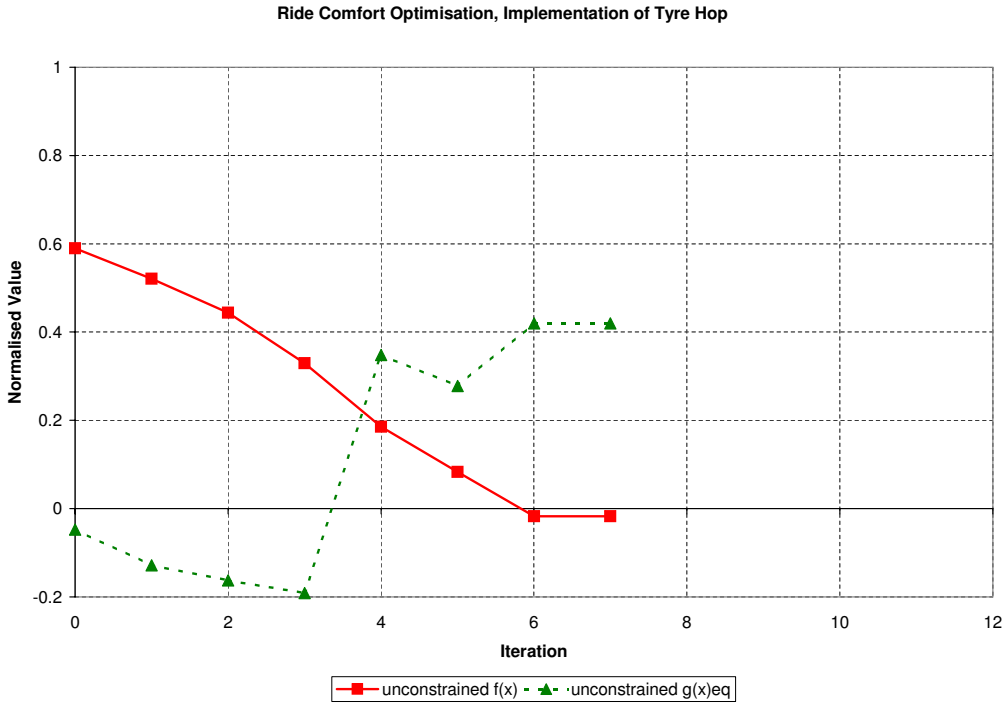

Figure 6.5: Observing tyre hop value while performing ride comfort optimisation

simplified pitch-plane model for gradient information are presented in Figure 6.6. It can be seen that the simplified gradients (*matgrad*) took approximately 24 iterations (25 expensive function evaluations) corresponding to an effective cost of 35 expensive function evaluations in terms of time, to reach an optimum. However, identical local minima, in terms of the objective function value, were repeatedly reached at iterations 10, 13, 17 and 20. The expensive gradients (*admsgrad*) effectively reached the optimum after 8 iterations at a cost of 45 expensive function evaluations, with an identical objective function value minimum repeated at iteration 19. Although the use of the simplified model for gradient information took more iterations, the total computing time was significantly less than using only the expensive numerical model for function values and gradient information. It is also apparent from the convergence histories that the use of the simplified model for gradient information, results in a much smoother convergence history, giving greater confidence in the computed results.

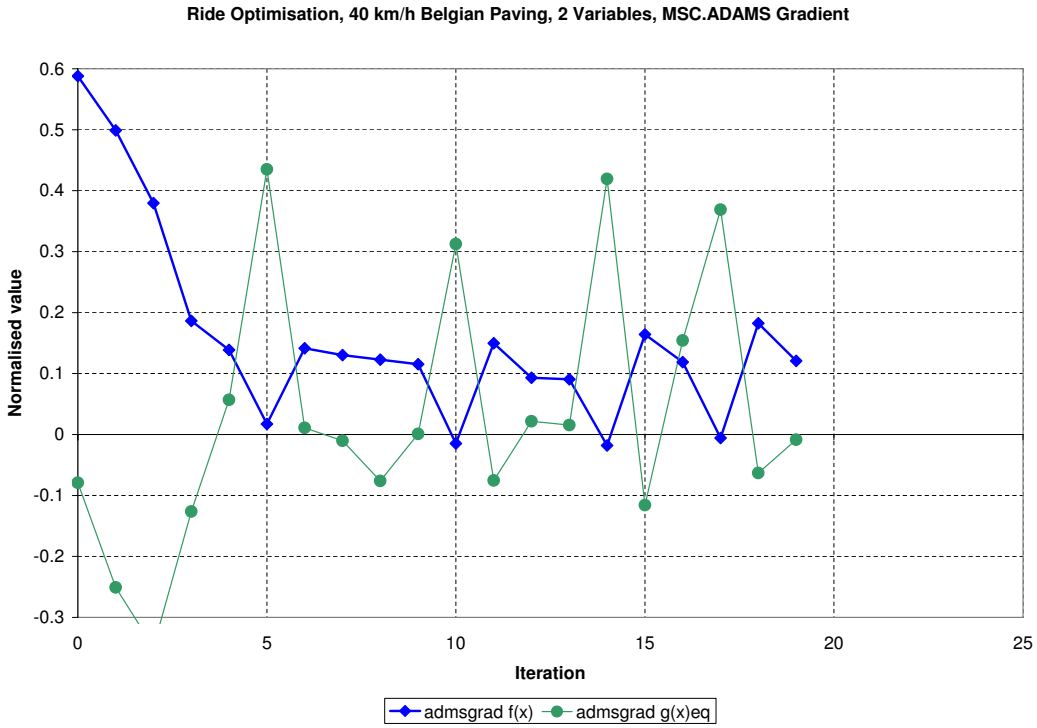

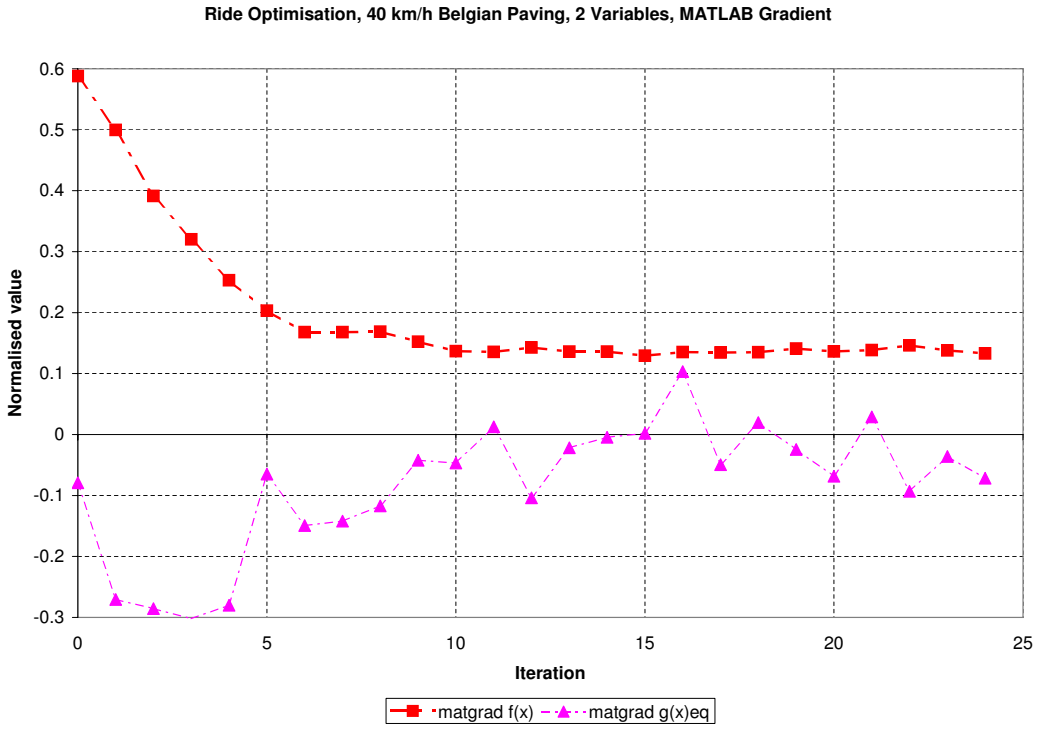

Figure 6.6: Comparison of the optimisation histories for the MSC.ADAMS gradient and simple MATLAB model gradient methods for 2 design variable ride comfort optimisation

6.3.3 Two Design Variable Optimisation, MATLAB Model Only

With such reasonable results obtained using the simple model for the computation of gradient information, it is necessary to justify the use of the complete MSC.ADAMS vehicle model for the function value in the optimisation process. The same optimisation was done as above but using only the simple Matlab model for the optimisation procedure. From the results in Figure 6.7 it can be seen that the function values are not the same as the MSC.ADAMS simulation values (calculated at iteration 5 and 25) and that the optimisation algorithm will converge to an infeasible point, when considering the constraints. Thus the use of the full MSC.ADAMS vehicle model is necessary in order to ensure the optimisation algorithm terminates at a feasible minimum. Although the simplified model has very similar trends, the absolute values are not always the same, especially when considering the tyre hop constraints. This explains why the converged solution may not be feasible.

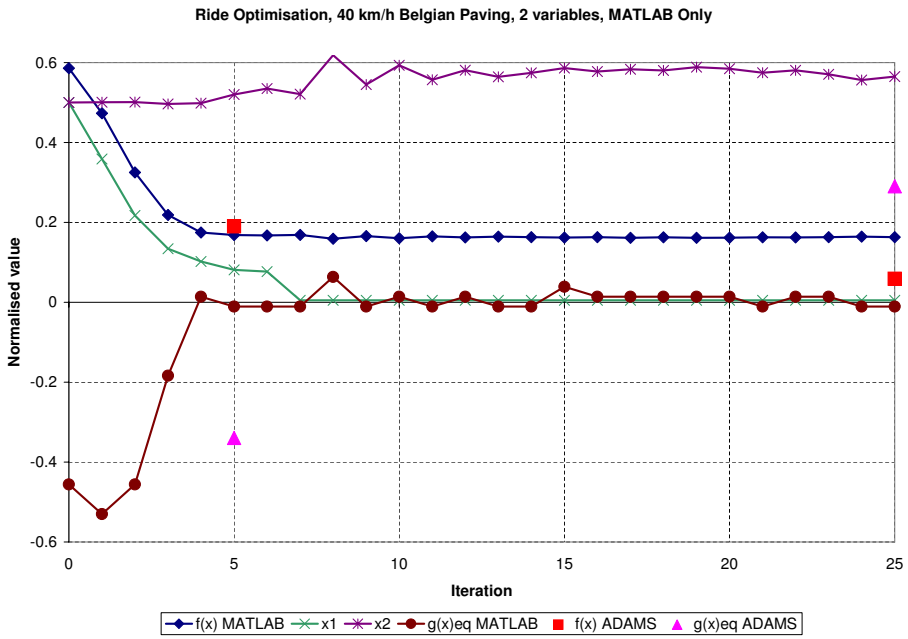

Figure 6.7: Ride comfort optimisation convergence history for using only the simple MATLAB based model, for objective function value, gradients, and tyre hop information.

6.3.4 Four Design Variable Optimisation

The four design variable ride comfort optimisation, was started from the optimum achieved from the two design variable optimisation. The optimisation process worked equally well as in the previously considered cases, although only small improvements are visible from the starting point, as can be seen from the MSC.ADAMS gradient history in Figure 6.8, and the Matlab gradient history in Figure 6.9. It is observed that although both methods converge to equally feasible solutions, the front and rear spring characteristics should differ in absolute value as can be seen by design variables $x2$ and $x4$. The result of this is that if the front gas volume is larger the front seated passengers will experience better ride comfort than the rear passengers, and the opposite if the rear spring gas volume is larger.

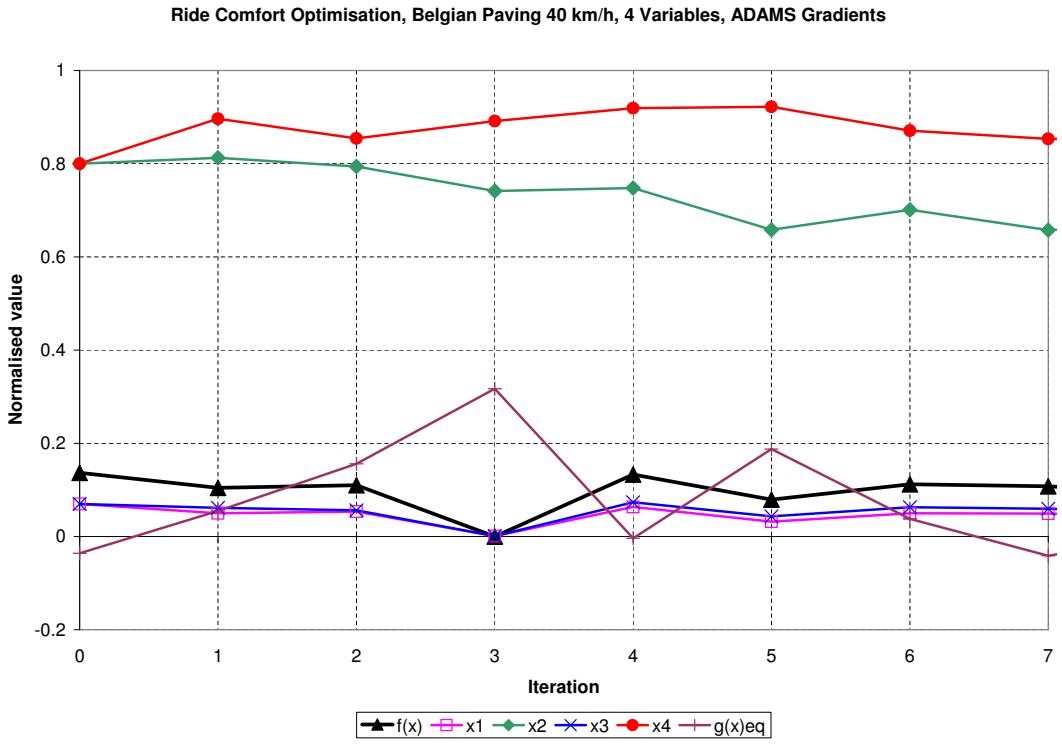

Figure 6.8: Ride Comfort optimisation convergence history for 4 design variables using the full MSC.ADAMS model for gradient information, starting at the optimum from two design variables

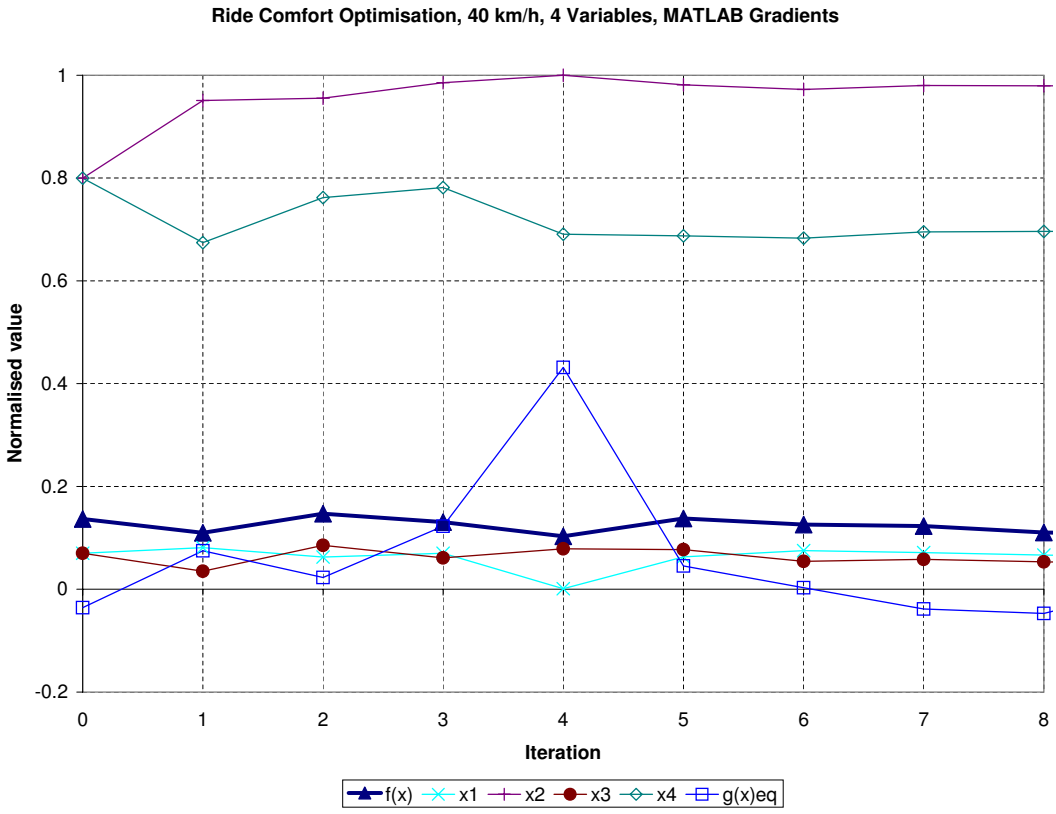

Figure 6.9: Ride Comfort optimisation convergence history for 4 design variables using the simple matlab model for gradient information, starting at the optimum from two design variables

From the above studies it is concluded that the optimisation process, making use of the simplified Matlab models for gradient information, produces equally feasible results in substantially less computational time. It will now be assumed that these models are sufficiently representative of the system for gradient information.

6.4 Summary of Results

Presented in Table 6.1 are the results for the optimisation runs. From the results it can be seen that the handling optimum suspension settings lie on the opposite corner of the design space to the ride comfort optima. If reasonable handling is to be achieved, then the ride comfort suffers, while if good ride comfort is to be achieved then the handling suffers. This is the

Table 6.1: Summary of Results for Optimisation Objectives

variables, opt. run	Fig.	# iter. (eq evals)	$f^*(\mathbf{x})$ ± 0.01	$\dot{\varphi}_{RMS}$ $[^o/s]$	φ_{peak} $[^o]$	a_{RMS_d} $[m/s^2]$	a_{RMS_p} $[m/s^2]$
Handling							
2, matgrad	6.1	12 (18.2)	0.15	0.57	3.0	-	-
2, admsgrad	6.1	15 (80)	0.15	0.57	3.0	-	-
4, admsgrad	6.2	6 (63)	0.15	0.54	3.2	-	-
4, matgrad	6.3	7 (14.4)	0.15	0.55	3.1	-	-
Ride							
2, matgrad	6.6	24 (35)	0.13	-	-	1.20	1.18
2, admsgrad	6.6	19 (100)	0.12	-	-	1.16	1.14
4, matgrad	6.9	9 (18)	0.11	-	-	1.14	1.08
4, admsgrad	6.8	7 (72)	0.11	-	-	1.10	1.10

traditional compromise, that the $4S_4$ suspension avoids due to the ability to switch between the optimum handling and ride comfort settings. The resulting optimal damping multiplication factors and spring gas volumes are presented in Table 6.2.

6.5 Conclusions

This chapter has shown that the use of simplified mathematical models, of the computationally intensive full simulation model, for use in computing gradient information, can significantly improve the optimisation process, when two and four design variables are considered. Firstly the optimisation process is significantly faster in terms of total optimisation time. Secondly the simplified models help to reduce numerical noise in the evaluation of the gradients, resulting in smoother convergence histories. Thirdly the simplified models are sufficiently representative of the vehicle system, when used for gradient information, although their absolute values may differ, and need to be properly scaled before use.

Table 6.2: Summary of optimum damper factors and gas volumes

opt. run	Fig.	$dpsff$	$gvolf$	$dpsfr$	$gvolr$
Handling					
2, matgrad	6.1	3.00	0.10	3.00	0.10
2, admsgrad	6.1	3.00	0.10	3.00	0.10
4, admsgrad	6.2	2.72	0.24	3.00	0.10
4, matgrad	6.3	2.89	0.10	2.69	0.10
Ride					
2, matgrad	6.6	0.30	0.51	0.30	0.51
2, admsgrad	6.6	0.29	0.54	0.29	0.54
4, matgrad	6.9	0.29	0.56	0.25	0.47
4, admsgrad	6.8	0.24	0.43	0.27	0.53

For the handling optimisation, it was found that the two methods gave identical optimum solutions, and that the optimal solutions lie along the maximum boundary of the damper design variable, and the lower boundary of the spring design variable.

For the ride comfort optimisation, the inclusion of the vehicle's tyre hop was investigated. It was found that the best results were achieved when including the tyre hop as an inequality constraint in the optimisation process. It was also found that the tyre hop tends to constrain the damping parameter from running towards its lower boundary constraint.

The methodology proposed is thus an efficient means of optimising a vehicle's suspension system for ride comfort and handling. This makes the use of deterministic gradient based optimisation algorithms most suitable, and competitive for suspension optimisation. More design variables will be incorporated and the combined optimisation of both ride comfort and handling considered in the following chapters.