Efficient Gradient-Based Optimisation of Suspension Characteristics for an Off-Road Vehicle

by

Michael John Thoresson

Submitted in partial fulfilment of the requirements for the degree

Philosophiae Doctor (Mechanical Engineering)

in the Faculty of

Engineering, Built Environment and Information Technology (EBIT)

University of Pretoria,
Pretoria

July 2007

Efficient Gradient-Based Optimisation of Suspension Characteristics for an Off-Road Vehicle

Michael John Thoresson

Supervisor:	Dr. P.E. Uys
Co-Supervisor:	Dr. P.S. Els
Department:	Mechanical and Aeronautical Engineering
Degree:	Ph.D. (Mechanical Engineering)

Abstract

The efficient optimisation of vehicle suspension systems is of increasing interest for vehicle manufacturers. The main aim of this thesis is to develop a methodology for efficiently optimising an off-road vehicle's suspension for both ride comfort and handling, using gradient based optimisation. Good ride comfort of a vehicle traditionally requires a soft suspension setup, while good handling requires a hard suspension setup. The suspension system being optimised is a semi-active suspension system that has the ability to switch between a ride comfort and a handling setting. This optimisation is performed using the gradient-based optimisation algorithm Dynamic-Q.

In order to perform the optimisation, the vehicle had to be accurately modelled in a multi-body dynamics package. This model, although very accurate, exhibited a high degree of non-linearity, resulting in a computationally expensive model that exhibited severe numerical noise. In order to perform handling optimisation, a novel closed loop driver model was developed that made use of the Magic Formula to describe the gain parameter

for the single point driver model's steering gain. This non-linear gain allowed the successful implementation of a single point preview driver model for the closed loop double lane change manœuvre, close to the vehicle's handling limit.

Due to the high levels of numerical noise present in the simulation model's objective and constraint functions, the use of central finite differencing for the determination of gradient information was investigated, and found to improve the optimisation convergence history. The full simulation model, however, had to be used for the determination of this gradient information, making the optimisation process prohibitively expensive, when many design variables are considered. The use of carefully chosen simplified two-dimensional non-linear models were investigated for the determination of this gradient information. It was found that this substantially reduced the total number of expensive full simulation evaluations required, thereby speeding up the optimisation time.

It was, however, found that as more design variables were considered, some variables exhibited a lower level of sensitivity than the other design variables resulting in the optimisation algorithm terminating at sub-optimal points in the design space. A novel automatic scaling procedure is proposed for scaling the design variables when Dynamic-Q is used. This scaling methodology attempts to make the n-dimensional design space more spherical in nature, ensuring the better performance of Dynamic-Q, which makes spherical approximations of the optimisation problem at each iteration step.

The results of this study indicate that gradient-based mathematical optimisation methods may indeed be successfully integrated with a multi-body dynamics analysis computer program for the optimisation of a vehicle's suspension system. Methods for avoiding the negative effects of numerical noise in the optimisation process have been proposed and successfully implemented, resulting in an improved methodology for gradient-based optimisation of vehicle suspension systems.

Keywords : gradient-based mathematical optimisation, vehicle suspension, spring and damper characteristics, Dynamic-Q, ride comfort, handling.

Doeltreffende Gradiëntgebaseerde Optimering van die Suspensiestelsel van 'n Veldvoertuig

Michael John Thoresson

Studieleier: Dr. P.E. Uys
Mede-leier: Dr. P.S. Els
Departement: Meganiese en Lugvaartkundige Ingenieurswese
Graad: Ph.D. (Meganiese Ingenieurswese)

Opsomming

Voertuigvervaardigers besef al hoe meer hoe belangrik doeltreffende optimering van suspensiestelsels is. Die hoofdoel van hierdie proefskrif is om 'n metode te ontwikkel om 'n veldvoertuig se suspensiestelsel vir ritgemak en hantering te optimeer, deur gebruik te maak van gradiëntgebaseerde wiskundige optimering. Tradisionele voertuigsuspensiestelsels het 'n sagte suspensiekarakteristiek nodig vir goeie ritgemak, en 'n harde suspensiekarakteristiek vir hantering. Die suspensiestelsel wat geoptimeer word is 'n semi-aktiewe suspensiestelsel, wat die vermoë het om te skakel tussen 'n ritgemak- en 'n hanteringskarakteristiek. Dié karakteristieke word bepaal deur gebruik te maak van die gradiëntgebaseerde wiskundige-optimeringsalgoritme Dynamic-Q.

Om die suspensiestelsel te optimeer is die voertuig akkuraat gemodelleer in 'n multi-liggaamdinamika sagteware pakket. Dié model het, as gevolg van sy lewensgetrouheid, hoë geraasvlakke en is berekeningsintensief as gevolg van die nie-lineariteit van die stelsel. Vir die suksesvolle optimering van hantering, moes 'n unieke stuurbeheermodel geimplementeer word. Hierdie model maak gebruik van die towerformule, normaalweg gebruik vir die modellering van bande, om die stuurinsetaanwinsfaktor vir die enkelpuntstuurbeheerder te verkry. Deur die stuurinset met 'n nie-lineêre aanwins te modelleer, kon 'n enkelpuntstuurbeheerder suksesvol

geïmplementeer word, om die voertuig naby aan sy hanteringslimiete deur die dubbelbaanverandering te stuur.

As gevolg van die hoë geraasvlakke wat in die doel- en begrensingsfunksies teenwoordig is, is sentrale eindige verskille ondersoek vir die berekening van gradiëntinligting. Daar is vasgestel dat die optimeringsgeskiedenis verbeter kan word deur gebruik te maak van sentrale eindige verskille vir gradiëntinligting. Die berekening van hierdie gradiënte is duur, omdat die volledige berekeningsintensiewe simulasiemodel gebruik word. Die gebruik van goed gekose nie-lineêre vereenvoudigde modelle vir die bepaling van gradiëntinligting is ondersoek. Daar is bevind dat die vereenvoudigde modelle goed werk vir die verkryging van gradiëntinligting en dat die optimeringstyd heelwat verminder kan word.

Wanneer meer ontwerpsveranderlikes in ag geneem word by die optimering, word bevind dat die doel- en begrensingsfunksies nie dieselfde sensitiwiteit het vir al die veranderlikes nie. Die gevolg is dat die optimeringsproses termineer op suboptimale punte in die ontwerpsruimte. 'n Unieke skalingsmetode is voorgestel vir gebruik met Dynamic-Q om aan die veranderlikes gelyke sensitiwiteit te gee. Die skalingsmetode maak die n-dimensionele ontwerpsruimte meer sferies van aard. Dynamic-Q, wat gebruik maak van sferiese subprobleme, kan dus beter benaderings tot die sub-probleme maak, en as gevolg daarvan is die optimering meer suksesvol.

Die resultate van hierdie werk, bevestig dat gradiëntgebaseerde optimeringsalgoritmes suksesvol met 'n multi-liggamdinaamika analise-program geïntegreer kan word om die suspensiekarakteristieke van 'n veldvoertuig te bepaal. Metodes om die nadelige effekte van numeriese geraas te oorkom, is voorgestel en suksesvol geïmplementeer, vir gradiëntgebaseerde optimering van voertuigsuspensiestelsels.

Sleutelwoorde : wiskundige optimering, voertuigsuspensie,
veer- en demperkarakteristieke,
Dynamic-Q, ritgemak, hantering.

Contents

Acknowledgments

- My study leaders Dr. P.E. Uys and Dr. P.S. Els for their continued support and enthusiasm throughout this research.
- Prof. J.A. Snyman, for his continual support and enthusiasm with regards to the optimisation throughout this research.
- Optimisation related investigations were performed under the auspices of the Multi-disciplinary Design Optimisation Group (MDOG) of the Department of Mechanical and Aeronautical Engineering of the University of Pretoria, funded by the NRF Optimisation for Industry research grants.
- National Research Foundation (NRF) for continued financial support of this research under Thutuka Contract No. TTK-2004-081-0000-43 and TTK-2005-080-1000-34.
- The vehicle dynamics simulation for the design of the controllable suspension system is based upon work supported by the European Research Office of the US Army under Contracts N68171-01-M-5852, N62558-02-M-6372 and N62558-04-P-6004.
- Land Rover South Africa for supplying technical information on vehicles.
- My colleagues and fellow students in the Dynamic Systems Group and the Department of Mechanical and Aeronautical Engineering for the interesting discussions and assistance in improving my Afrikaans communication skills.
- My parents for their continual support and encouragement throughout my studies.

Preface

This thesis is part of a group effort at the University of Pretoria, to improve the safety, comfort and handling of Sport Utility Vehicles (SUV's), through the use of an intelligent suspension system. I have included here an article I wrote for the Saturday Star, which in a light hearted manner introduces the reader to the project.

I am your status, favourite toy, your freedom and your recreation. But while I am good for all these things, I still have many faults, just like you! You might not think so! My design gives you a sense of security and power. While you are really dying to know who I am, I have to keep you in suspension a little longer, as this is the key.

The secret is almost out! We are relaxing in the veld, and you are marvelling at my abilities, spraying mud everywhere, ploughing the sand, climbing the boulders. Like mountain goats nothing can stop us.

Party-time is over. Back in the week, running late as always. The cursed fat passenger, stubborn as an axe, will not buckle up. We go faster than a cheetah about to kill, along the motorway, my big powerful horses really galloping. BUT, just in front, out pops a ghostly pedestrian. No time to even curse, your arms frantically turning the wheel. Wow, just missed him! But my height catches you out and, roll, roll, roll your pride, roughly down the roadside.

Both looking as though battered by a ram, you emerge to the frozen motorway, gawkers everywhere. The vultures like sumo wrestlers, fighting for the first feed. The pedestrian? Well, gone like the wind. And finally the sound of sirens, help on it's way.

This is the fate of many a SUV (Sport Utility Vehicle) on the roads. Many owners of these vehicles are fooled into thinking they are safe as they are so high above the other road users. It is this height coupled to a soft suspension for good off-road manœuvrability, which leads to a dangerous package when performing sports car manœuvres, typically occurring during accident conditions.

What are the solutions? First of all always wear your safety belt. Then buy a sports car for the road and a 4×4 *for off-road only. No this is too*

expensive, what else can be done? Remove the pedestrian. We would love to, but the pedestrian could be anything in the road, an animal, a pothole, or a tree if you're drunk. No, the vehicle needs to be driven with less steam. And in the mean time, some more work is being done on the vehicle's suspension, which is exactly what we are busy doing. We at the University of Pretoria, are developing a controllable suspension system, for this type of vehicle. See, I told you! Keeping you suspended is the key.

The suspension is the link between the tyres and the vehicle body. By controlling this link intelligently, with the help of Newton's laws of physics, the accident could have been better avoided. Yes, you read right, the laws of physics apply to everything especially vehicles. Our unique suspension has two different spring settings and two different damper settings per wheel, in order to better eliminate the compromise between off-road and on-road driving.

With the help of intelligence we switch between the spring and damper settings while the driver is doing his normal task of driving. We thus have a setting for severe handling manoeuvres, and one for comfort and off-road driving. The controller uses it's knowledge of physics to switch the suspension within 50 milliseconds.

Let's take a quick look at the accident again. Driver driving fast, suspension in comfort mode, pedestrian appears in road. Driver frantically turns steering wheel, suspension switches to handling spring and adjusts the damping every 50 milliseconds. Driver misses the pedestrian, overcorrects a little, vehicle stays upright, and driver continues driving. The intelligent suspension did not panic, and kept the vehicle upright better than the driver did. This is because it remembers the laws of physics at all times.

List of Abbreviations

$4S_4$	4 State Semi-Active Suspension System
AD	Automatic Differentiation
ADAMS	Advanced Dynamic Analysis of Mechanical Systems
admsgrad	ADAMS Model Gradient
ANN	Artificial Neural Network
apg	Aberdeen Proving Ground
Ascl-DynQ	Automatic Scaling Dynamic-Q
ascl	Automatic Scaling
BFGS	Broyden-Fletcher-Goldfarb-Shanno
BS	British Standard
CCD	Central Composite Design
CFD	Computational Fluid Dynamics
cfd	Central Finite Difference
cg	Center of Gravity
DFP	Davidon-Fletcher-Powell
dof	Degree(s) of Freedom
dpsf	damper scale factor
ETOP	Euler-Trapezium Optimiser
ETOPC	Euler-Trapezium Optimiser for Constrained Problems
FEM	Finite Element Method
ffd	Forward Finite Difference
GA	Genetic Algorithm
gvol	static gas volume
lf	Left Front
inf	Infeasible Starting Point

lr	Left Rear
ISO	The International Organisation for Standardisation
LFOP	Leap-Frog Optimiser
LFOPC	Leap-Frog Optimiser for Constrained Problems
MAM	Multipoint Approximation Method
MDOG	Multi-disciplinary Design Optimisation Group
min	Minimum
matgrad	Matlab Model Gradient
max	Maximum
NRF	National Research Foundation
ode	Ordinary Differential Equation
PID	Proportional Integral Derivative
PSD	Power Spectral Density
RMS	Root Mean Square
RSM	Response Surface Methodology
SLP	Sequential Linear Programming
SQP	Sequential Quadratic Programming
std	Standard
SUMT	Sequential Unconstrained Minimisation Technique
SUV	Sports Utility Vehicle
VDV	Vibration Dose Value

List of Symbols

$\mathbf{A}$	Dynamic-Q Approximated Hessian Matrix
a	Curvature
a	Pseudo Arctangent Coefficient
a_i	Quadratic Constant of Polynomial Approximation
a_{RMS}	Root Mean Square Acceleration
a_{zRMSd}	Driver Root Mean Square Vertical Acceleration
a_{zRMSd}	Passenger Root Mean Square Vertical Acceleration
a_w	Weighted Acceleration
a_{0-13}	Magic Formula Coefficients
B	Stiffness Factor
b	Pseudo Arctangent Coefficient
b_i	Linear Constant of Polynomial Approximation
C	Shape Factor
C_d	Drag Factor
$\mathbf{C}$	Damping Matrix
c_t	Tyre Damping
D	Peak Factor
$dpsf$	Damper Scale Factor
$dpsff$	Damper Scale Factor Front
$dpsfr$	Damper Scale Factor Rear
dx_k	Perturbation in Design Variable k
E	Curvature Factor
F	Force
F_{dmp}	Damper Force
$F(\mathbf{x})$	Multi-Variable Function

F_{drive}	Driving Force
F_{ztyre}	Vertical Tyre Force
F_y	Lateral Force
F_z	Vertical Force
f	Front
f_{4S_4}	Force in $4S_4$ Suspension
$f(\mathbf{x})$	Objective Function
$\tilde{f}(\mathbf{x})$	Approximated Objective Function
$f^*(\mathbf{x})$	Optimum Objective Function Value
$f^*(\mathbf{x})_{hand}$	Optimum Handling Objective Function Value
$f_{hand}(x)$	Ride Comfort Objective Function
$f_{ride}(x)$	Ride Comfort Objective Function
G_{do}	Road Roughness Coefficient
g	Gravity
$gvol$	Static Gas Volume
$gvolf$	Static Gas Volume Front
$gvolr$	Static Gas Volume Rear
$g_j(\mathbf{x})$	Inequality Constraint Functions
$\hat{g}_{il}(\mathbf{x})$	Design Variable Lower Limit Inequality Constraint
$\check{g}_{iu}(\mathbf{x})$	Design variable Upper Limit Inequality Constraint
h	Finite Difference step Size
h_{cg}	Height of Center of Gravity
$h_j(\mathbf{x})$	Equality Constraint Functions
H_k	Hessian Matrix at Iteration k
$\mathbf{I}$	Identity Matrix
I_x	Mass Moment of Inertia About x-Axis
I_y	Mass Moment of Inertia About y-Axis
I_z	Mass Moment of Inertia About z-Axis
i	Index
j	Index

$\mathbf{K}$	Stiffness Matrix
k	Constant
k	Index
k_t	Tyre Vertical Stiffness
$\hat{k_i}$	Design Variable Lower Limit
$\check{k_i}$	Design variable Upper Limit
l	Counter
l	Left
m_A	Mass of Axle
m_b	Mass of Body
m_t	Mass of Tyre
m_v	Mass of Vehicle
$\mathbf{M}$	Mass Matrix
M_z	Moment About z-Axis
n	Number of Design Variables
$\mathbf{O}$	Zero Matrix
O	Zero Vector
$\mathbf{P}[i]$	Approximated Sub Problem
r	Right
r	Rear
sf	Scale Factor
S_h	Horizontal Shift
S_v	Vertical Shift
t	Time
t_i	Design Variable Scale Factor
t_s	Start Time
t_s	Track Width Suspension
t_{total}	Total Time
T	Transposed
u	Counter

v	Velocity
$v_{current}$	Current Variable Value
v_{low}	Variable Lower Limit
v_{high}	Variable Upper Limit
x	Relative Displacement
x	Magic Formula Slip Term
x_i	Design Variable i
x_k	Design Variable k
x_n	Design Variable at Iteration n
x_{1-14}	Design Variables
$\dot{x_{act}}$	Instantaneous Speed
$\dot{x_d}$	Desired Speed
$\dot{x}$	Vehicle Speed
$\mathbf{x}$	Vector of Design Variables
$\mathbf{x}^*$	Vector of Optimum Design Variables
X	Tyre Slip Angle
$\mathbf{X}$	Vector of Rescaled Design Variables
X_i	Rescaled Design Variable
y	Vehicle Lateral Displacement
$\ddot{y}_v$	Lateral Acceleration
Y	Tyre Lateral Force
wb	Wheelbase
z	Vertical Displacement
$\mathbf{z}$	Vector of Scaled Design Variables
z_i	Scaled Design Variable Between Zero and One
z_{rf}	Road Disturbance Input Front
z_{rr}	Road Disturbance Input Rear
α	Tyre Slip Angle
δ	Move Limit Magnitude
δ	Steering Angle

$\dot{\delta}$	Steering Rate
δ_{stat}	Static Spring Displacement
δt	Time Step
∂	Partial Derivative
γ	Camber Angle
∇	Gradient Vector
μ	Penalty Multipliers
μ_j	Penalty Multiplier
ω	Terrain Index
ψ	Vehicle Yaw Angle
ψ_d	Desired Yaw Angle
ψ_a	Actual Yaw Angle
$\dot{\psi}_a$	Actual Yaw Rate
$\ddot{\psi}$	Yaw Acceleration
τ	Preview Time
θ	Pitch Angle
$\dot{\theta}$	Pitch Rate
$\ddot{\theta}$	Pitch Acceleration
φ	Roll Angle
$\dot{\varphi}$	Roll Rate
$\ddot{\varphi}$	Roll Acceleration
ϑ_0	Current Rotational Angle
ϑ_p	Predicted Rotational Angle
$\dot{\vartheta}$	Current Rotational Velocity
$\ddot{\vartheta}$	Rotational Acceleration

Chapter 1

Introduction and Background

In today's competitive world, the need to develop a vehicle in the most efficient manner is of utmost importance. In particular, the need exists for robust and efficient optimisation algorithms for determining the optimal spring and damper characteristics of a vehicle for both ride comfort and handling. This optimisation is difficult to perform because of two reasons. First of all the objective and constraint functions used in the optimisation are determined via computationally expensive numerical simulations. Secondly, due to the need to include non-linear effects in the numerical model to accurately simulate reality, serious numerical noise may be present in the objective function. Both these factors, namely computational expense and the presence of noise, have seriously inhibited the general use of mathematical programming methods in the optimal design of mechanical systems. This research aims to provide the reader with an efficient methodology for optimising an off-road vehicle's suspension characteristics for ride comfort and handling.

1.1 Ride Comfort vs. Handling

Throughout the history of the modern motor vehicle, the suspension system design has been a compromise between ride comfort, handling and driver control. In newer passenger vehicles this compromise has been reduced by the addition of stiff anti-roll-bars, this allows for a soft suspension setup for vertical motion, associated with ride comfort, and a stiff suspension for roll

motion, typically handling manœuvres. Off-road vehicles and sports utility vehicles (SUV's) inherently have soft suspension characteristics, for good off-road manœuvrability, with the spin-off being good ride comfort, however, they are very unstable when handling is considered. Stiff anti-roll-bars are generally infeasible as they result in limited wheel travel, affecting the off-road manœuvrability.

Els (2006) investigated this compromise between ride comfort and handling in off-road vehicles. A four state semi-active suspension system, to be known as $4S_4$, was developed and tested. The unique feature of this system is that it can switch not only between different damper characteristics but also different spring characteristics. Els developed a control algorithm for this unique system that has the ability to automatically switch from the ride comfort mode to the handling mode, using no physical input from the driver. A prototype vehicle was fitted with the $4S_4$ system. Large improvements were achieved in terms of handling over the baseline vehicle, with large improvements in ride comfort when in the ride comfort setting, over the handling mode setting. This system thus eliminates the traditional compromise between ride comfort and handling, as it operates in ride comfort mode when driving in a relatively straight line, but should the driver begin a handling type manœuvre the system switches to the handling suspension mode. The handling mode's suspension characteristics are optimised for optimal handling and the ride comfort mode's suspension characteristics for optimal ride comfort, thereby eliminating the compromise associated with traditional suspension systems. The work presented in this document, discusses the optimisation of the suspension settings of the $4S_4$ system.

1.2 Development of the $4S_4$

The suspension unit currently under development, has the unique feature that it incorporates two damper packs (fitted with bypass valves) and two gas accumulators, effectively giving two damper characteristics and two spring characteristics in a single suspension unit. This unit will be referred to as

the '4-State Semi-Active Suspension System', or $4S_4$ (Theron and Els 2005).

The suspension consists of two settings, namely ride comfort and handling. The handling spring setting is achieved by the compression of a small gas volume, resulting in a stiff spring stiffness. The ride comfort spring setting is achieved by the compression of both the small gas volume and a larger gas volume resulting in a soft overall spring stiffness. In addition to the variable spring settings, the damping can be varied for each spring setting. The low damping setting, desirable for ride comfort, is achieved by the pressure drop, as a result of the flow through the by-pass valves to the spring accumulators. High damping is achieved, with the by-pass valves in the closed position, by the pressure drop, as a result of the oil flow through the damper pack for the desired spring.

Switching between the two spring and damper characteristics is achieved by solenoid valves as illustrated in Figure 1.1. Valve switching times vary between 50 and 100 milliseconds depending on system pressure. Spring and damper characteristics can be taken as design variables, to be optimised for both ride comfort and handling respectively. It is assumed that the suspension system will switch between the ride comfort and handling option, to suite the operating conditions, provided an intelligent control system can distinguish between the two different operating conditions, and switch the suspension system to the correct setting. Each operating setting is expected to have different optimum values for the spring and damper characteristics. This suspension has the ability to eliminate the traditional ride comfort vs. handling compromise.

1.3 The need for Optimisation

With the off-road vehicle's suspension system already a complex compromise between ride comfort and handling, the addition of additional complexity in the form of variable spring settings and damper settings, with associated control, the use of a few hit-and-miss hand calculations will not permit

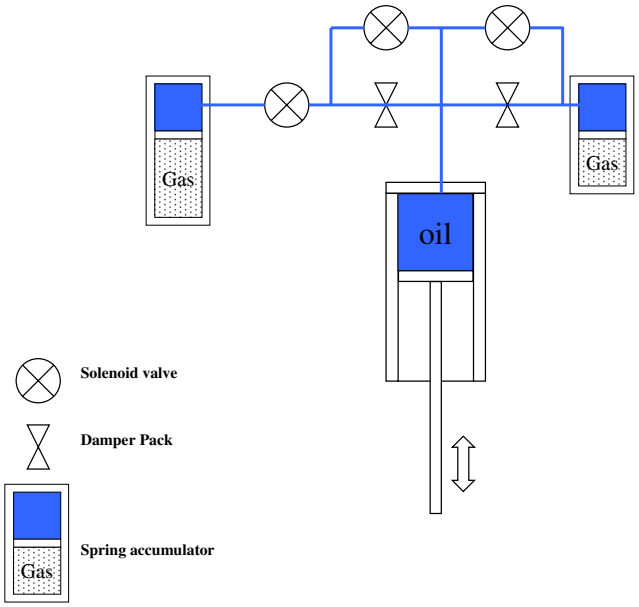

Figure 1.1: $4S_4$ Suspension Unit

the developed suspension system to live up to it's perceived qualities. To accurately model the vehicle for analysis purposes, of the new suspension system, requires the modelling of many highly non-linear components, like suspension characteristics, bushings, bump and rebound stops, and most importantly a very non-linear tyre. As a result of this complexity necessary to obtain accurate models, the design space that is to be investigated is dramatically large, non-linear and noisy. Where numerical noise in this thesis will be defined as: for small perturbations in the design variables sent to the full simulation model in MSC.ADAMS relatively large perturbations in the objective function values are noted. It has however also been suggested that this could be referred to as high sensitivity.

To accurately define the damper and spring characteristics for front and rear suspension setups requires at least 14 design variables. With such a large number of design variables, it is impossible to visualise the effect of each design variable on the ride comfort or handling, to select the optimal

configuration. Additionally, the vehicle can travel at various speeds, over various terrains, and under various load conditions. The only way to take all these aspects into account is to make use of mathematical optimisation techniques. However, due to the sheer complexity of the problem to be solved, there are many aspects that need to be considered before mathematical optimisation will successfully determine the optimal suspension characteristics for the vehicle in question. The primary aim of this work is to propose a methodology for the efficient implementation of gradient-based mathematical optimisation for the optimisation of the off-road vehicle's suspension system.

1.4 Summary

In the author's masters degree dissertation (Thoresson 2003) the use of SQP and Dynamic-Q were investigated for vehicle suspension optimisation. It was found that the use of central finite differencing as opposed to forward finite differencing for the determination of gradient information for use within the Dynamic-Q optimisation algorithm, resulted in an improved optimisation convergence history. This is as a result of the central finite differences helping to reduce the undesirable effects of numerical noise on gradient determination. However, this came at the cost of additional expensive objective and constraint function evaluations per iteration. These additional expensive objective and constraint function evaluations result in a prohibitively expensive optimisation process when more design variables are considered.

The main aim of this work is the use gradient-based optimisation to efficiently optimise the off-road vehicle's suspension system for ride comfort and handling. In order to do this many steps have to be completed along the way.

This document describes the use of mathematical optimisation for vehicle suspension design, a summary into the investigation of the SQP and Dynamic-Q methods, followed by the advantages achieved when using central

finite differences for gradient information, the development of accurate models to describe the vehicle dynamics, the validation of simplified models for gradient information, implementation of the simplified models for 2 and 4 variable optimisation, complications encountered with numerous design variables, a proposed automatic scaling of design variables, application of the process to 14 design variable optimisation of ride comfort and handling, and the optimisation of the compromise suspension setup.

The following original contributions to the application of gradient-based optimisation for vehicle suspension design are presented in this Thesis. Firstly the application of multi-fidelity optimisation to vehicle suspension design, in which a detailed simulation model is used for the evaluation of the objective and constraint functions and simplified models for the evaluation of the finite difference gradients. Secondly automatic scaling of the design variables with respect to the topology of the objective functions, to improve the convergence of the optimisation algorithm for the problems considered here. Thirdly the development of a robust steering driver model based on the nonlinear Pacejka Magic Formula for the description of the steering gain factors.

Chapter 2

Mathematical Optimisation

In this chapter, the use of mathematical optimisation for vehicle suspension characteristics is discussed. The general properties of gradient-based and stochastic algorithms are evaluated. The optimisation algorithms that were selected for the investigation of the problem at hand are defined, and a methodology for their implementation is defined.

2.1 The Use of Mathematical Optimisation

The use of mathematical optimisation techniques for the improvement of the engineering design process, is rapidly gaining acceptance. There is great debate in the optimisation world as to whether gradient-based approximation techniques or stochastic-based methods, like genetic algorithms, are more efficient and suited to engineering design. Stochastic techniques generally require a large starting population, in order to achieve a sufficiently feasible solution. This makes the stochastic methods computationally expensive, when expensive numerical models, of the physical system are to be optimised. Most researchers have to utilise costly multiple processing systems, as the desktop computer can take days or even weeks to arrive at a solution. On the other hand, gradient-based optimisation techniques tend to be heavily dependent on the initial starting point, and require accurate gradient information for the iterative approximation of the design space. The determination of this gradient information, is costly when many design variables are considered. The gradient calculation is also severely

affected by numerical noise that is normally inherent in complex numerical simulation models, e.g. full vehicle models. Research, with reference to vehicle suspension optimisation, is now briefly discussed.

Dahlberg (1977, 1979), investigated the optimisation of a vehicle's suspension system for ride comfort and working space, subject to a random road input. A 1-degree of freedom (dof) model, was optimised using the Sequential Unconstrained Minimisation Technique (SUMT) (Fiacco and McCormick 1968). This was then expanded to a linear 2-dof model, to investigate the speed dependence of the optimal suspension settings. It was found that for a small suspension working space, the optimal spring and damper settings are heavily dependent on vehicle speed, while for a large working space the optimum is not really dependant on vehicle speed. It is suggested that active suspension systems be considered when small suspension working spaces are available.

Eberhard et al. (1995) successfully used a gradient based optimisation method (a sequential quadratic programming, or SQP, algorithm) to optimise a simple pitch-plane vehicle model's non-linear damper characteristics for ride comfort. The non-linear damper characteristic is modelled with piecewise Hermite splines. The Hermite splines, however, require difficult to handle constraints in order to ensure feasibility of the optimised damper characteristic. Nevertheless, satisfactory results were obtained. Boggs and Tolle (2000) provide an introduction to the SQP method and discuss recent developments for large scale non-linear applications.

Etman et al. (2002) designed a stroke dependent damper, for the front axle of a truck, using Sequential Linear Programming (SLP), a gradient based optimisation algorithm. They use a 2-dof quarter car model, for the initial investigation of the desired non-linear damper characteristics. Ride comfort is optimised using discrete road obstacles. The non-linear damper characteristics are modelled using an empirical piecewise quadratic approximation. Finally a full vehicle model is used for the ride comfort

optimisation, for one discrete road obstacle. Bump-stop contact is ignored, to remove numerical noise and lessen computational expense. Difficulties were experienced due to poor finite difference approximations of the gradients, and with multiple feasible optima being found.

Naudé and Snyman (2003a, 2003b) and Naudé (2001) make use of a pitch-plane vehicle model to optimise the piecewise linear damper characteristics of an off-road military vehicle, for ride comfort. The 'Leap-Frog' (LFOPC) optimisation algorithm (Snyman 2000) was used, and although taking many iterations to reach the optimum, the optimisation was completed within a few seconds, because the vehicle model code was specially written for the vehicle being investigated.

Baumal et al. (1998) compared the efficiency of a Genetic Algorithm (GA) to a gradient-based optimisation method (gradient projection method) for a pitch-plane vehicle model, that was computationally efficient. The GA converged to an optimum that was only a 4% improvement over the gradient based method, but, required thousands more objective function evaluations.

Eberhard et al. (1999) investigate the use of a stochastic optimiser (simulated annealing) and a gradient-based (deterministic) optimiser (a SQP algorithm) for the optimisation of a full linear vehicle model's ride comfort. The four design variables considered are the linear spring and damper coefficients, the distance of the body center of gravity (cg) between the axles and the track width of the wheels. They conclude that deterministic optimisation approaches offer rapidly converging algorithms that often get stuck in local minima, when optimising multi-body dynamic systems. Nevertheless, the global optimum may be obtained by these methods if used within a multi-start strategy. They also find that simulated annealing is useful in avoiding local minima. It does, however, require substantially more function evaluations in order to locate the global optimum. Thus both methods are successful in locating the global optimum. They consequently suggest a hybrid combination of stochastic and deterministic algorithms for optimisation. They state,

however, that the switching strategy is and will continue to be a challenging task.

Eriksson and Friberg (2000) optimised the linear spring and damper characteristics of the engine mounting system on a city bus, for ride comfort. Use was made of a linear finite element method (FEM) model to simulate the response of the bus to a given road input, with three passenger positions used for the ride comfort objective function. A 7 % improvement in ride comfort was achieved and it was found that the local minima, to which the gradient based algorithm (form of SQP algorithm, with gradients determined by forward finite differencing) converged to, were heavily dependent on the initial starting point. Eriksson and Arora (2002) investigated the use of three continuous global optimisation methods for the ride comfort optimisation of the city bus. It was found that the modified zooming method in terms of number of objective function evaluations (464) is most efficient in locating the global optimum.

Gobbi et. al. (1999, 1999) use a back-propagated Artificial Neural Network (ANN) of the full vehicle simulation model, coupled with a genetic algorithm for the optimisation of ride and handling of a sedan vehicle. Suspension non-linearities are modelled as piecewise linear approximations. The full simulation model has been verified against test data. The ANN was used for function evaluations within the genetic algorithm optimisation process. However, this methodology requires an extensive number of function evaluations, of the expensive full simulation model, to sufficiently train a representative ANN, making it infeasible for stand-alone workstations.

Schuller et al. (2002) optimised the comfort and handling of a BMW sedan using a simplified vehicle model composed of transfer functions. Because of the nature of the vehicle model the suspension design parameters were only allowed to have a small variance of 15% over the current vehicle design. This process thus aims to refine an already feasible design for the next model launch. The numerical model solves faster than real-time, making the use

of genetic algorithms feasible. Only open loop handling manœuvres were considered for the optimisation process.

Andersson and Eriksson (2004) optimised the non-linear damper and spring characteristics of a full city bus vehicle model, that was validated against test data. The model consists of non-linear bushings, bump-stops, springs, dampers and a non-linear 'Magic Formula' tyre model. The ride comfort of the bus was optimised for three discrete road obstacles, with a 23 % improvement achieved. The handling was optimised using a single lane change manœuvre at 40 and 80 km/h, with a 6 % improvement achieved. The handling objective function is defined as a combination of the yaw rate gain and yaw rate time lag, with an inequality constraint limiting the maximum body roll angle to less than 1.3 *degrees*. The built-in MSC.ADAMS SQP method was used, and the optima were reached after approximately 145 function evaluations. An attempt was made at the combined optimisation of handling and ride comfort, and it was found that the result is heavily dependent on the weights assigned to the various performance objectives.

Gonsalves and Ambrósio (2005) make use of a vehicle model consisting of a flexible vehicle body and linear spring and damper characteristics, to perform optimisation of the suspension characteristics for ride comfort and handling of a sports car. The ride comfort objective function consists of the ride index, which is the summed contributions of the vibration dose values for different positions in the vehicle. The handling objective function consists of the time taken to reach steady state lateral acceleration and the overshoot of the roll angle for an open loop manœuvre. The optimisation algorithm used is the Modified Method of Feasible Directions of Vanderplaats (1992), with improvement in ride comfort and handling achieved.

Els et al. (2006) compared the efficiency of the Dynamic-Q optimisation algorithm to the SQP method for vehicle suspension optimisation. They found that the use of central finite differencing for the determination of gradient information improved the convergence of the Dynamic-Q

optimisation algorithm towards a feasible optimum within fewer objective function evaluations, when compared to SQP or Dynamic-Q with forward finite differencing. The objective functions exhibited severe noise. It appeared, however, that using central finite differencing with relatively large steps in computing gradient information, was successful in smoothing out the effect of the noise in the optimisation.

Bandler et al. (2004) and Koziel et al. (2005) introduced to the engineering optimisation world the theory of 'Space Mapping', which makes use of a coarse simple model (surrogate model) and a detailed fine model for the optimisation process. The Space Mapping technique involves the matching and updating of the coarse model to more accurately describe the fine model. This has been successfully applied to the structural optimisation of a vehicle for crash safety, by Redhe and Nilsson (2004). In their research the coarse model was constructed using linear Response Surface Methodology (RSM) with the optimisation converging within 14 iterations, and using a total of 26 expensive function evaluations. However, the RSM model must be trained.

Space Mapping is also refered to as multi-fidelity optimisation, which is also defined as the use of a high-fidelity model (fine model in space mapping speak), and a medium or low fidelity model (coarse model), for optimisation. Balabanov and Venter (2004) made use of a greatly simplified finite element method (FEM) model of a full FEM model for the determination of gradient information for structural optimisation, with success. Gobbi et al. (2005) suggest the use of neural networks, or piecewise quadratic function approximations of the full simulation model, when optimising a vehicle's dynamics. van Keulen and Toropov (2006) investigate the use of the Multipoint Approximation Method (MAM) for a FEM structural problem that exhibits numerical noise. The basic idea is to replace the noisy optimisation problem with a succession of noise-free approximations at each iteration. This noise-free approximation is then optimised, and the optimum used for the next iteration point. van Keulen and Toropov (2006) also suggest the use of mechanistic approximations, to be used for

the optimisation process, where the simplified numerical model is based on a prior knowledge of the physical system.

Papalambros (2002) suggests constructing surrogate models for optimisation, by making use of the computationally expensive simulation model for 'computational experiments'. With this experimental data curve-fitting techniques are applied to represent the original functions with acceptable accuracy. The problem with this method, however, is that the correct underlying form of the fit needs to be chosen, and higher accuracy requires increased sample points, resulting in increased computational cost.

The concept of Automatic Differentiation (AD) is a novel way of obtaining gradient information with one function evaluation (Tolsma and Barton 1998, Bartholomew-Biggs et al. 2000). This methodology was evaluated by Bischof et al. (2005) for the shape optimisation of an airfoil, with the objective function being evaluated by a software chain. Although AD provides more accurate gradient information than forward finite differences, the evaluation of the objective function was approximately 16 times slower than the original code for eight ($n = 8$) design variables. Using forward finite differences would have used the original code $n + 1$ times, equating to a cost of nine times the cost of one function evaluation of the original code. The other downside of AD is that access to the original source code is necessary, and it is normally not available when commercial simulation software, such as MSC.ADAMS is used.

Snyman (2005a) introduced a new implementation of the conjugate gradient method (Euler-trapezium optimiser for constrained problems, ETOPC) that overcomes the problem of severe numerical noise superimposed on a smooth underlying objective function. Snyman introduces a novel gradient-only line search, that requires two gradient vector evaluations per search direction, and no explicit function evaluations. It is also found that the computation of the gradients by central finite differences with relatively large perturbations, allowed for smoothing out of the inherent numerical noise.

The principal aim of this work is to promote the use of gradient-based optimisation algorithms for vehicle suspension optimisation. In order to do this, the complications associated with computational cost and inherent numerical noise have to be investigated. For this reason this work investigates the use of the Sequential Quadratic Programming (SQP) method and the locally developed Dynamic-Q method, for the optimisation of the suspension problem.

2.2 The SQP Method

The Sequential Quadratic Programming (SQP) optimisation algorithm is well known and is considered the industry-standard gradient-based method for constrained optimisation problems if the number of variables is not too large. The version used here is found in Matlab's Optimisation Toolbox (Mathworks 2000a). SQP makes use of successive quadratic approximations of the objective and constraint functions at each iteration step. In constructing these approximations second order differential information is required, in the form of the Hessian matrix. The Hessian matrix is approximated by making use of the Broyden-Fletcher-Goldfarb-Shanno (BFGS) approximation. The BFGS method relies on forward finite differences to approximate the gradient of the objective function. The Hessian matrix does, however, require updating if the problem behaves poorly, requiring an extra $n + 1$ function evaluations per iteration. SQP makes use of line searches to find the solution of the approximate subproblem, this solution is then the next iteration point.

2.3 The Dynamic-Q Method

Complications associated with computational cost and inherent numerical noise have to be investigated in this study, for this reason the locally developed Dynamic-Q optimisation algorithm is used. Having direct access to the code allows more freedom to investigate the effects of different optimisation concepts. Dynamic-Q has also proved to be a feasible algorithm for vehicle

suspension optimisation by Els and Uys (2003). The Dynamic-Q method has been developed to address the general optimisation problem:

$$\underset{w.r.t.x}{minimize} \quad f(\mathbf{x}), \quad \mathbf{x} = [x_1, x_2, .., x_n]^T \in R^n \tag{2.1}$$

subject to the inequality constraints:

$$g_j(\mathbf{x}) \leq 0, \qquad j = 1, 2, .., m \tag{2.2}$$

and the equality constraints:

$$h_j(\mathbf{x}) = 0, \qquad j = 1, 2, .., r \tag{2.3}$$

where $f(\mathbf{x})$, $g_j(\mathbf{x})$ and $h_j(\mathbf{x})$ are scalar functions of $\mathbf{x}$. In this formulation $\mathbf{x}$ is the vector of design variables, $f(\mathbf{x})$ is the objective function, $g_j(\mathbf{x})$ the inequality constraint functions, and $h_j(\mathbf{x})$ the equality constraint functions.

The Dynamic-Q algorithm is defined as: 'Applying a *Dynamic* trajectory optimisation algorithm to successive spherical *Quadratic* approximations of the actual optimisation problem' (Snyman and Hay 2002). This algorithm has the major advantage that it only needs to do relatively few function evaluations of the original expensive objective function to construct a simple quadratic approximate function. This new approximate sub-problem's objective and constraint functions can then be evaluated cheaply and the optimum point of the approximate sub-problem may be found economically, using the robust dynamic trajectory method LFOPC (Snyman 2000). At this new approximate optimum point, a new quadratic approximate sub-problem of the objective and constraint functions is constructed, that is again optimised. This procedure is iteratively repeated until convergence is obtained. This method is very efficient for optimising objective and constraint functions that require an expensive computer simulation for their evaluation. In standard form Dynamic-Q makes use of forward finite differences to obtain gradient information required for the generation of the approximations. The details of the method can be found in the publications by Snyman and Hay (2002), and Els and Uys (2003) where it was applied to a similar vehicle as in this study, and formed the building block for this work. A basic outline

of the algorithm is set out below.

A sequence of approximate sub-problems ***P****[i] i = 0,1,2,...* are generated by constructing successive spherically quadratic approximations to the objective and constraint functions, at successive points $\mathbf{x}_i$. The approximation to the objective function, for example, is as follows:

$$\tilde{f}(\mathbf{x}) = f(\mathbf{x}_i)(\mathbf{x} - \mathbf{x}_i) + \nabla^T f(\mathbf{x}_i)(\mathbf{x} - \mathbf{x}_i) + \frac{1}{2}(\mathbf{x} - \mathbf{x}_i)^T \mathbf{A}(\mathbf{x} - \mathbf{x}_i) \tag{2.4}$$

The Hessian matrix $\mathbf{A}$ takes on a simple diagonal matrix form:

$$\mathbf{A} = a\mathbf{I}; \tag{2.5}$$

This form of Hessian matrix indicates that the approximate subproblems are spherically quadratic in nature. The curvature a takes on a value of zero for the first subproblem $i = 0$. Thereafter it is defined by:

$$a = \frac{2[f(\mathbf{x}_{i-1}) - f(\mathbf{x}_i) - \nabla^T f(\mathbf{x}_i)(\mathbf{x}_{i-1} - \mathbf{x}_i)]}{\|\mathbf{x}_{i-1} - \mathbf{x}_i\|^2} \tag{2.6}$$

The approximate constraint functions are constructed in a similar manner. If the gradient vectors ∇f, ∇g, and ∇h are not known analytically they may be approximated by first order finite differences, traditionally forward finite differences are used.

Additional side constraints of the form $\hat{k}_i \leq x_i \leq \check{k}_i$ are normally imposed on the design variables. Because these constraints do not exhibit curvature properties they are treated as linear inequality constraints. These constraints thus take on the form:

$$\hat{g}_{il}(\mathbf{x}) = \hat{k}_i - x_i \leq 0, \qquad l = 1, ..., r \leq n, \tag{2.7}$$

$$\check{g}_{iu}(\mathbf{x}) = x_i - \check{k}_i \leq 0, \qquad u = 1, ..., s \leq n, \tag{2.8}$$

To obtain stable and controlled convergence of the solutions of successive approximate sub-problems, a move limit is set which takes on the form of an inequality:

$$g_\delta(\mathbf{x}) = \left\|\mathbf{x} - \mathbf{x}^{i-1}\right\|^2 - \delta^2 \leq 0 \tag{2.9}$$

where δ corresponds to the specified maximum magnitude of the move limit. The approximate subproblem at $\mathbf{x}_{i-1}$ can now be solved using the dynamic trajectory 'Leap-Frog' optimisation algorithm for constrained optimisation LFOPC. This solution is taken as $\mathbf{x}_i$, the point at which the next approximate sub-problem is constructed. This process is continued until convergence is obtained. The process is illustrated in a simplified form in Figure 2.1, where f represents the approximated subproblem at each iteration step, and x_n the x value obtained at each iteration step. The x_1 value was limited by the allowable move limit.

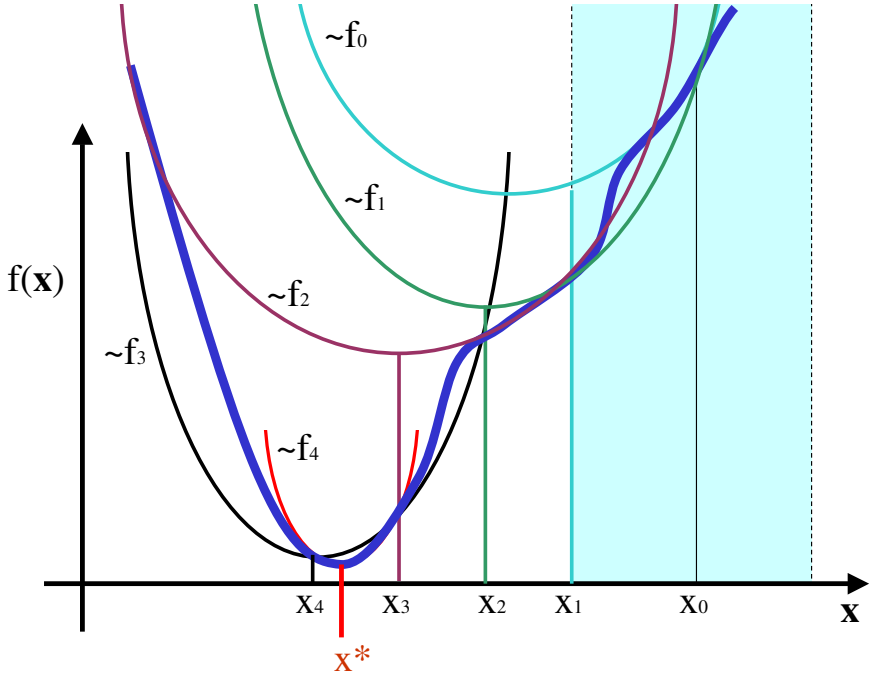

Figure 2.1: Simplified illustration on how Dynamic-Q progresses with optimisation iterations

2.4 Gradient Approximation Methods

Most gradient-based optimisation algorithms require the determination of the first and/or second order gradient information of the objective and constraint functions with respect to the design variables. In most engineering optimisation problems this gradient information is not analytically available.

The only information available to the designer is the values of objective and constraint functions obtained via expensive simulations. This paragraph investigates the use of forward and central finite differences in the Dynamic-Q optimisation algorithm, for the determination of the first order gradient information.

2.4.1 Forward Finite Difference (ffd)

This is the simplest and most economic method for approximating the gradients of the objective and constraint functions, required by gradient-based mathematical optimisation algorithms. This method approximates the first order gradient information of a multi-variable function $F(\mathbf{x})$, by evaluating the change in the function $F(\mathbf{x})$ for a small change dx_k in each of the design variables x_k, $k = 1, 2, ..., n$, as illustrated in Figure 2.2. Thus, in order to carry out the full gradient vector evaluation, a total number of $n+1$ function evaluations are required for each iteration, where n is the total number of design variables. The forward finite difference approximation to the k^{th} component of the gradient at $\mathbf{x}$ is defined as follows:

$$\frac{\partial F}{\partial x_k} = \frac{F(x_1, x_2, ..., x_k + dx_k, ..., x_n) - F(\mathbf{x})}{dx_k} \tag{2.10}$$

for $k = 1, 2, ..., n$. Noisy objective functions, however, severely limit the accuracy of the forward finite difference gradient approximation, as is apparent from Figure 2.2. This can be partly overcome by using larger step sizes dx_k or by considering instead, central finite differences.

2.4.2 Central Finite Difference (cfd)

Central finite differences make use of a function evaluation on either side of the current iteration point $\mathbf{x}$, resulting in a better approximation to the gradient of the underlying smooth function in the presence of noise. Although this method requires $2n + 1$ function evaluations per gradient vector evaluation, it may result in fewer optimisation iterations to obtain a minimum.

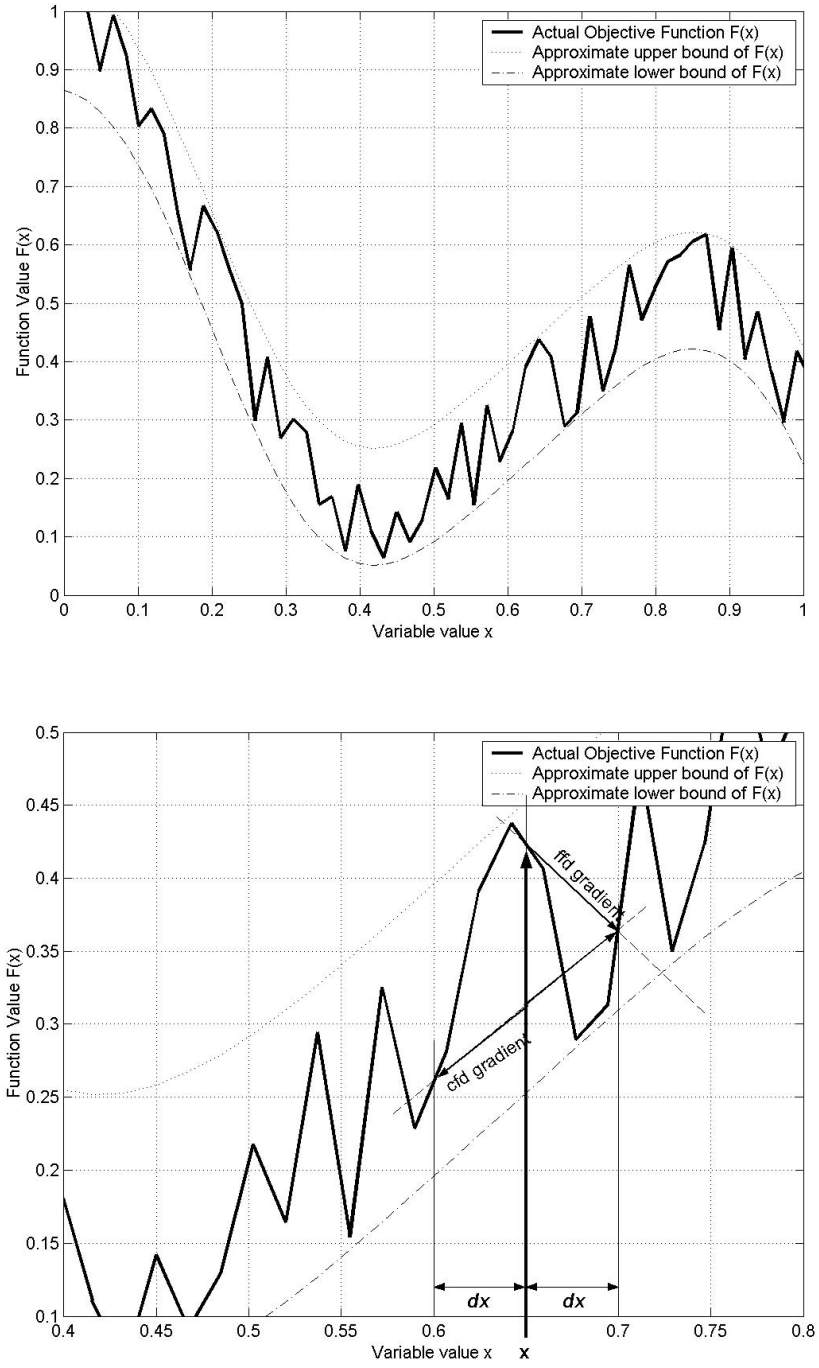

Figure 2.2: Finite difference gradient approximation methods

The central finite difference procedure is defined as follows:

$$\frac{\partial F}{\partial x_k} = \frac{F(x_1, x_2, ..., x_k + dx_k, ..., x_n) - F(x_1, x_2, ..., x_k - dx_k, ..., x_n)}{2dx_k} \quad (2.11)$$

for $k = 1, 2, ..., n$. In this way the gradient is evaluated by looking at information behind and ahead of the current iteration point, while the forward finite difference only looks ahead of the current iteration point. This results

in a more accurate approximation to the function gradient, when noise is present, as illustrated for the case depicted in Figure 2.2. The effects of noise cannot be completely eliminated by this method, but it certainly yields gradient approximations that are superior to that given by forward finite differences.

2.4.3 Higher Order Gradient Information

The Sequential Quadratic Programming (SQP) method (Mathworks 2000a, Vanderplaats 1999) and other Quasi-Newton optimisation algorithms such as the Davidon-Fletcher-Powell (DFP) method uses, in addition to first order gradient approximations, also second order curvature information. This information is very costly to obtain, as it corresponds to a partial derivative of a partial derivative. This information is stored in a n x n square matrix, commonly known as the Hessian matrix. The Broyden-Fletcher-Goldfarb-Shanno (BFGS) approximation to the Hessian matrix is used in Matlab's implementation of SQP. The Hessian matrix is approximated and updated at iteration $k+1$, $k = 0, 1, 2, ...$ by:

$$H_{k+1} = H_k + \frac{\mathbf{q}_k \mathbf{q}_k^T}{\mathbf{q}_k^T \mathbf{s}_k} - \frac{H_k^T \mathbf{s}_k^T \mathbf{s}_k H_k}{\mathbf{s}_k^T H_k \mathbf{s}_k} \tag{2.12}$$

where

$$\mathbf{s}_k = \mathbf{x}_{k+1} - \mathbf{x}_k \tag{2.13}$$

and

$$\mathbf{q}_k = \nabla f(\mathbf{x}_{k+1}) - \nabla f(\mathbf{x}_k) \tag{2.14}$$

and

$$\nabla f(\mathbf{x}_k) = [\frac{\partial f}{\partial x_1}, \frac{\partial f}{\partial x_2}, ..., \frac{\partial f}{\partial x_n}] \tag{2.15}$$

At the start of the optimisation procedure, (i.e. at iteration $k = 0$) most algorithms set H_0 equal to any positive definite symmetric matrix, normally the identity matrix I. Thereafter the approximation is updated at every iteration via equations 2.12 - 2.14.

2.5 Conclusions

This chapter looked at vehicle suspension optimisation research, and defined the optimisation methods to be used for the rest of this work.

The primary aim of this work is the promotion of gradient-based optimisation algorithms for vehicle suspension optimisation, due to the minimal number of function evaluations they require over stochastic based methods to arrive at a feasible optimum. The decision was thus taken that the SQP method, with it's strong industry presence, and the locally developed Dynamic-Q method will be used.

The successful implementation of gradient-based methods, is strongly dependent on good gradient information. Finite differencing is, however, necessary for the determination of gradient information when the objective and constraint functions are determined via numerical simulations. Forward and central finite differencing will be investigated for it's efficiency in determining gradient information.

Chapter 3

Full Vehicle Simulation Model

Two different versions of the full vehicle simulation model of the test vehicle will now be described. The models are validated against experimental results. A unique steering driver model is proposed and successfully implemented. This driver model makes use of a non-linear gain, modelled with the Magic Formula, traditionally used for the modelling of tyre characteristics.

3.1 Initial Vehicle Model

A Land Rover Defender 110 was initially modelled in ADAMS View 12 (MSC 2005) with standard suspension settings as a baseline. The ADAMS 521 interpolation tyre model is used, because of its ability to incorporate test data in table format. The tyre's vertical dynamics and load dependent lateral dynamics are thus considered in this model. In order to keep the model as simple as possible, yet as complex as necessary, longitudinal dynamic behaviour of the tyres and vehicle is not considered here. The anti-roll bar and bump stops are left unchanged. Only the spring and damper characteristics are changed for optimisation purposes. This study builds on current research into a two-state semi-active spring-damper system. The semi-active unit has been included in the ADAMS model and replaces the standard springs and dampers. The inertias of the vehicle body were determined by scaling down data available for an armoured prototype Land Rover 110 Wagon, and were considered to be representative of the lighter vehicle.

The complete model consists of 16 unconstrained degrees of freedom, 23 moving parts, 11 spherical joints, 10 revolute joints, 9 Hooke's joints, and one motion defined by the steering driver. The vehicle direction of heading is controlled by a simple single point steering driver, adjusting the steering wheel rotation according to the difference of the desired course from the current course at a preview distance ahead of the vehicle.

3.2 Refined Vehicle Model

A refined model of the Land Rover Defender 110 is also modelled in MSC.ADAMS View (MSC 2005) with standard suspension settings, as a baseline. For this model, the non-linear MSC.ADAMS Pacejka 89 tyre model (Bakker et al. 1989) is fitted to measured tyre data, and used within the model. This tyre model was selected as it was found that the 5.2.1 tyre model could not handle tyre slip angles larger than 3 degrees. The Pacejka 89 tyre model was used with a point follower approximation for rough terrain, to speed up the simulation speed, and as a result of limited tyre and test track data available at the time. As in the initial model, the tyre's vertical dynamics and load dependent lateral dynamics are also considered in this model. In order to keep the model as simple as possible, yet as complex as necessary, longitudinal dynamic behaviour of the tyres and vehicle is again not considered here. The vehicle body is modelled as two rigid bodies connected along the roll axis at the chassis height, by a revolute joint and a torsional spring, in order to better capture the vehicle dynamics due to body torsion in roll. The anti-roll bar is modelled as a torsional spring between the two rear trailing arms to be representative of the actual anti-roll bar's effect. The bump and rebound stops, are modelled with non-linear splines, as force elements between the axles and vehicle body. The suspension bushings are modelled as kinematic joints with torsional spring characteristics that are representative of the actual vehicle's suspension joint characteristics, in an effort to speed up the solution time, and help decrease numerical noise. The baseline vehicle's springs and dampers are modelled

Table 3.1: MSC.ADAMS vehicle model's degrees of freedom

Body	Degrees of Freedom	Associated Motions
Vehicle Body	7	body torsion
(2 rigid bodies)		longitudinal, lateral, vertical
		roll, pitch, yaw
Front Axle	2	roll, vertical
Rear Axle	2	roll, vertical
Wheels	4 x 1	rotation

with measured non-linear splines. The vehicle's center of gravity (cg) height and moments of inertia were measured (Uys et al. 2005) and used within the model. Only the spring and damper characteristics are changed for optimisation purposes. The $4S_4$ unit has been included in the MSC.ADAMS model, using the MSC.ADAMS Controls environment to include the Simulink model, and replaces the standard springs and dampers. Due to the fact that different suspension characteristics are being included the first two seconds of the simulation have to be discarded, while the vehicle is settling into an equilibrium condition. Figures 3.1 and 3.2 indicates the detailed kinematic modelling of the rear and front suspensions. The complete model consists of 15 unconstrained degrees of freedom, 16 moving parts, 6 spherical joints, 8 revolute joints, 7 Hooke's joints, and one motion defined by the steering driver. The degrees of freedom are indicated in Table 3.1.

The vehicle's direction of heading is controlled by a carefully tuned yaw rate steering driver, adjusting the front wheels' steering angles according to the difference of the desired course from the current course at a preview distance ahead of the vehicle (see paragraph 3.4). The longitudinal driver is modelled as a variable force attached to the body at wheel height depending on the difference between the instantaneous speed and desired speed (see paragraph 3.3). This MSC.ADAMS model is linked to MATLAB (Mathworks 2000b) through a Simulink block that requires as inputs the spring and damper design variable values, and returns outputs of vertical accelerations, vehicle body roll angle and roll velocity.

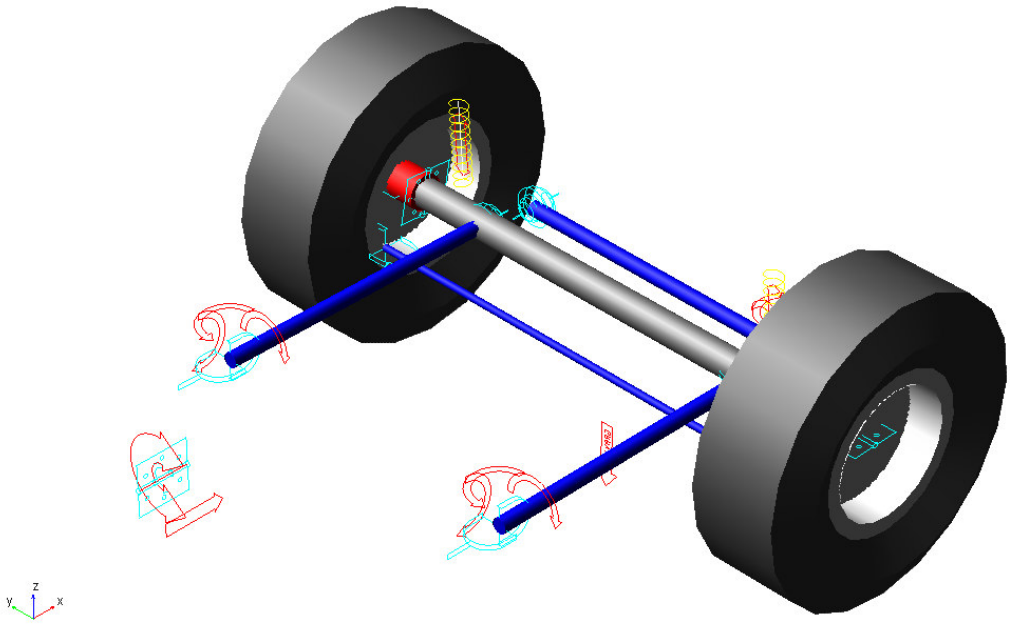

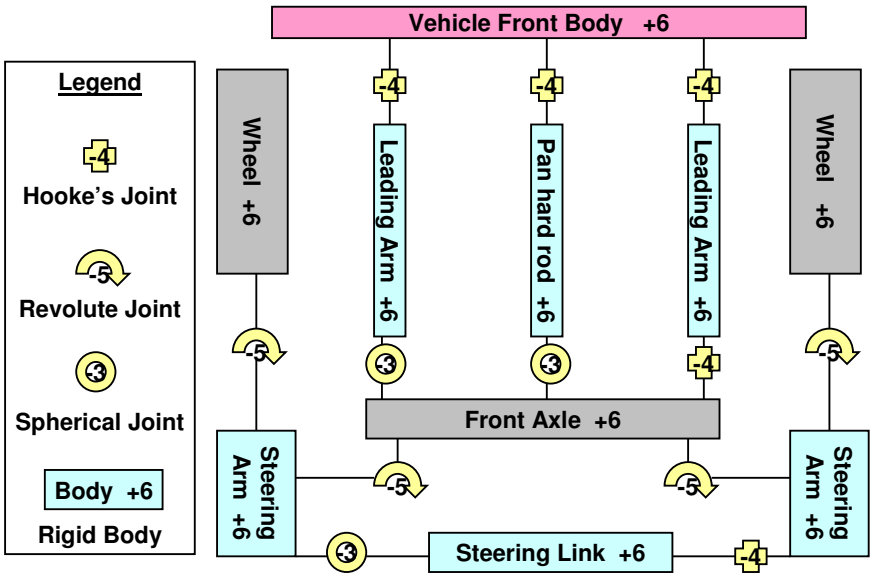

Figure 3.1: Modelling of the full vehicle in MSC.ADAMS, front suspension

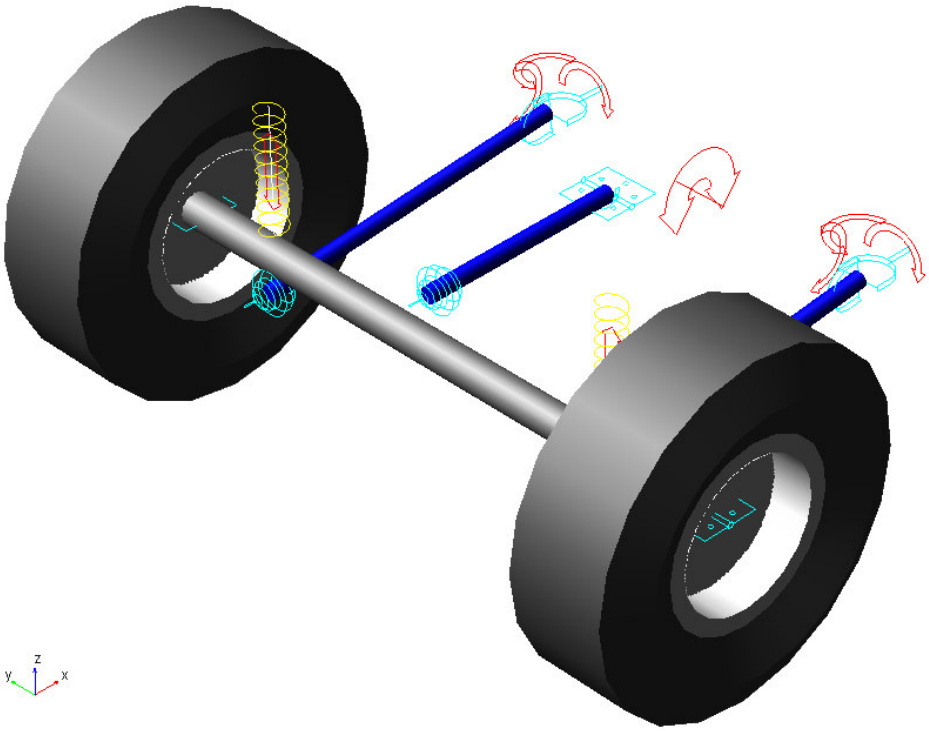

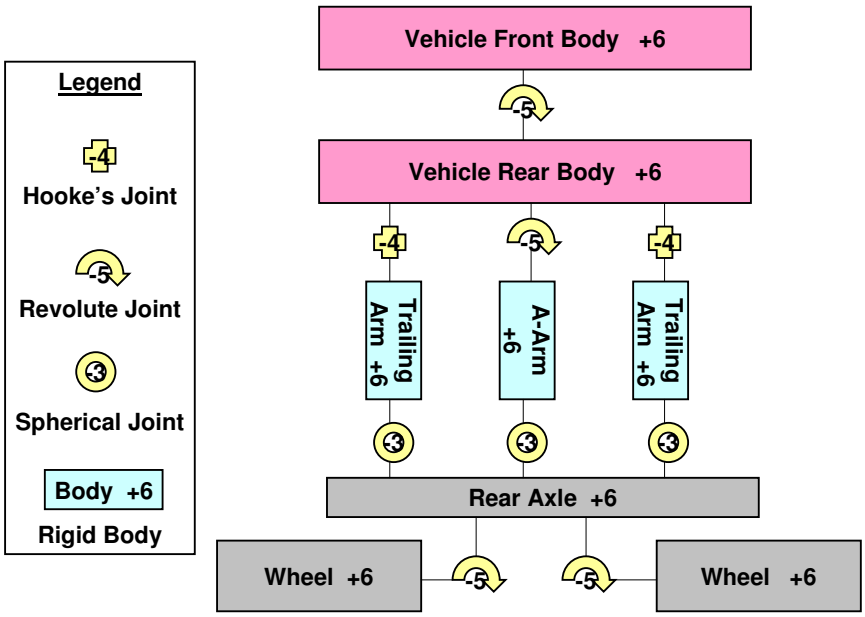

Figure 3.2: Modelling of the full vehicle in MSC.ADAMS, rear suspension

3.2.1 Validation of Full Vehicle Model

The MSC.ADAMS full vehicle model is validated against measured test results performed on the baseline vehicle. The measurement positions are defined by Figure 3.3 and Table 3.2. The correlation results are presented in Figure 3.4 for the baseline vehicle travelling over two discrete bumps to evaluate vertical dynamics, and in Figure 3.5 for the vehicle performing a double lane change manœuvre at 65 km/h. From the results it is evident that the model returns excellent correlation to the actual vehicle. It is, however, computationally expensive to solve and exhibits severe numerical noise due to all the included non-linear effects.

Table 3.2: Land Rover 110 test points

channel	point	position	measure	axis
1	B	center of gravity	velocity	longitudinal
2	G	left front bumper	acceleration	longitudinal
3				lateral
4				vertical
5	C	rear passenger	acceleration	longitudinal
6				lateral
7				vertical
8	I	right front bumper	acceleration	vertical
9	A	steering arm	displacement	relative arm/body
10	D	left rear spring	displacement	relative body/axle
11	E	right rear spring		
12	F	left front spring		
13	H	right front spring		
14	B	center of gravity	angular velocity	roll
15				pitch
16				yaw

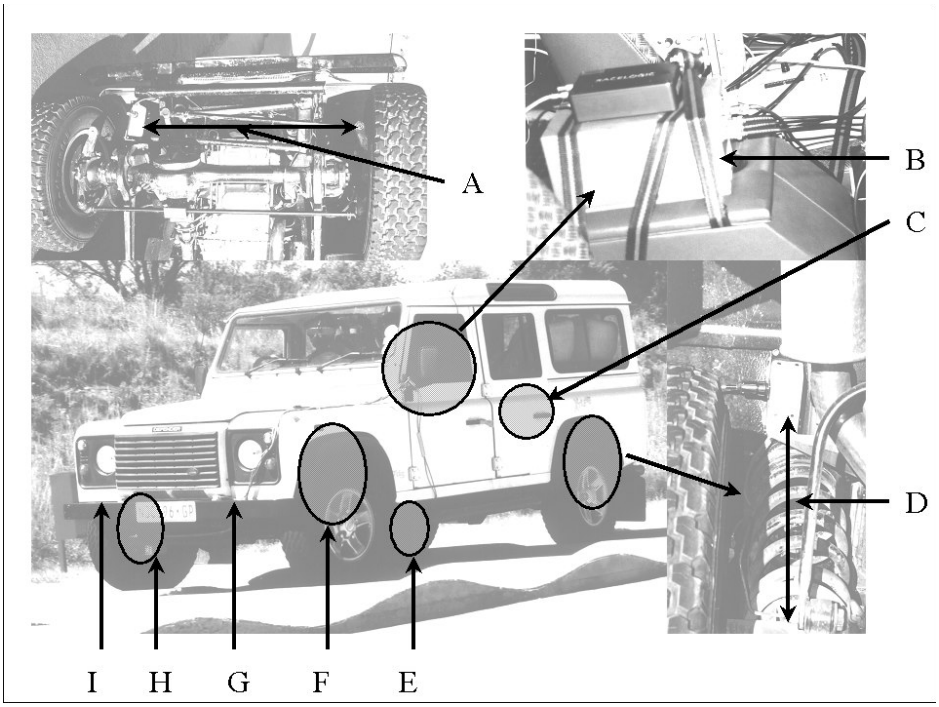

Figure 3.3: Test vehicle indicating measurement positions

3.3 Vehicle Speed Control

The speed control is modelled as a variable force F_{drive} attached to the body at wheel center height. The magnitude of this force depends on the difference between the instantaneous speed $\dot{x}_{act}$ and desired speed $\dot{x}_d$. Because the vehicle is a four-wheel drive with open differentials, the vertical tyre force $F_{z_{tyre}}$ is measured at all tyres (1 to 4). If a tyre looses contact with the ground, the driving force to the vehicle is removed until all wheels are again in contact with the ground. The driving force is thus defined as:

$$\begin{aligned} if \quad & F_{z_{tyre}1\rightarrow4} = 0 \\ & F_{drive} = 0 \\ & else \\ F_{drive} &= 4\frac{min(1200, 1200(\dot{x}_d - \dot{x}_{act}))}{0.4} \\ & end \end{aligned} \tag{3.1}$$

The gain value of 1200 was determined to be sufficiently large to ensure fast stable acceleration of the vehicle model from rest up to the desired simulation speed. This force is multiplied by 4 as there is one force acting on

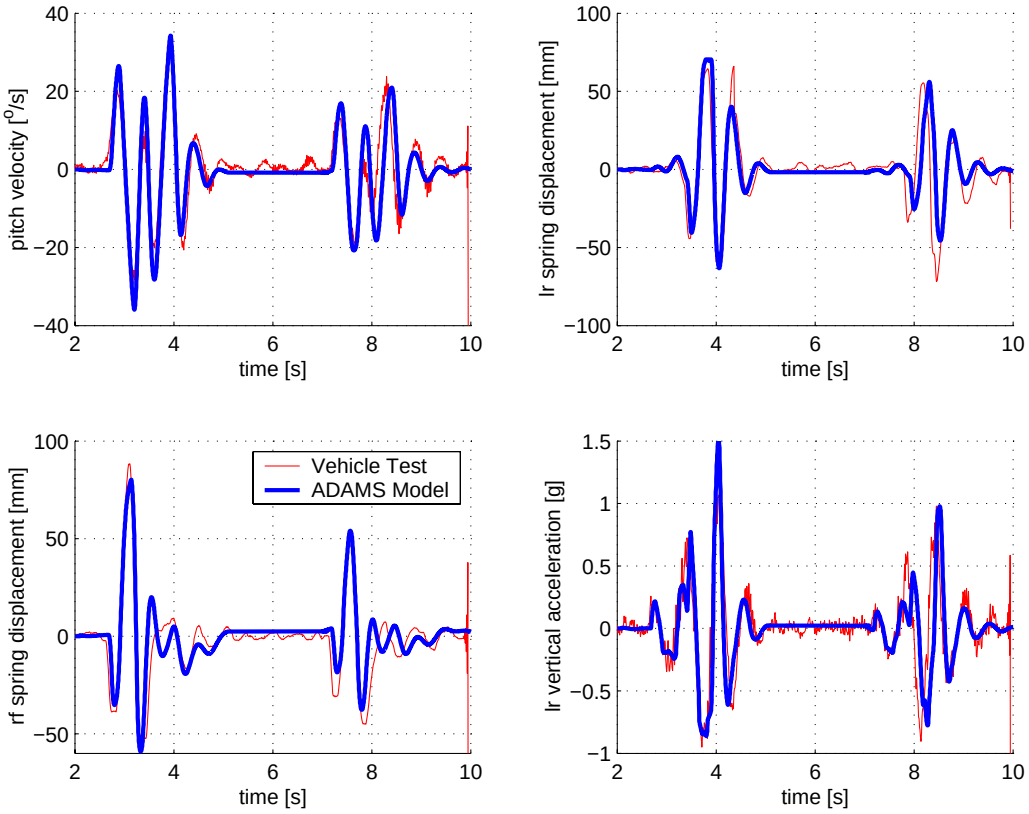

Figure 3.4: Discrete bumps, 15 km/h, validation of MSC.ADAMS model's vertical dynamics

the vehicle representative of the torque applied to the four wheels, and 0.4 meters is the radius of the tyres. The MSC.ADAMS model is then linked to the Simulink (Mathworks 2000b) based driver model that returns as outputs the desired vehicle speed and steering angle, calculated using the vehicle's dynamic response.

3.4 Driver Model For Steering Control

The use of driver models for the simulation of closed loop vehicle handling manœuvres is vital. However, great difficulty is often experienced in determining the gain parameters for a stable driver at all speeds, and vehicle parameters. A stable driver model is of critical importance during mathematical optimisation of vehicle spring and damper characteristics for handling, especially when suspension parameters are allowed to change over a wide range. The determination of these gain factors becomes

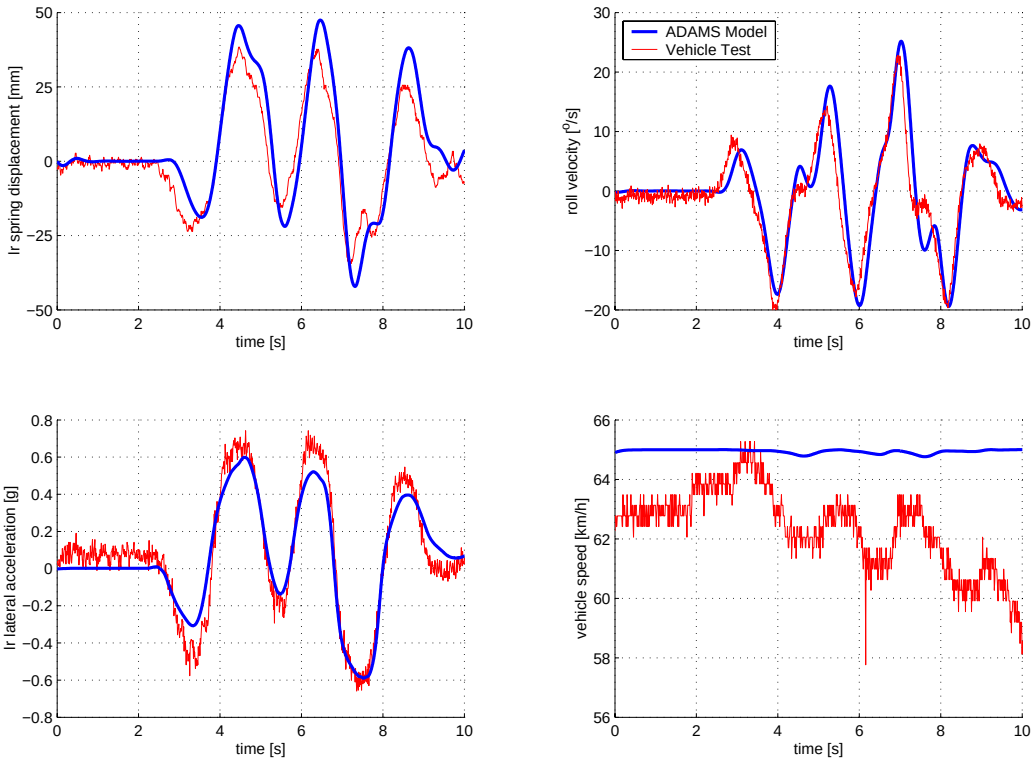

Figure 3.5: Double lane change, 65 km/h, validation of MSC.ADAMS model's handling dynamics

especially complex when accurate full non-linear vehicle models, with large suspension travel, are to be controlled. Single point preview models are normally unstable for such non-linear vehicle models. This paragraph investigates the relationship between vehicle yaw response and non-linear tyre characteristics. The non-linearity of the tyre characteristics is replicated for the steering gain parameter, ensuring the feasibility of single point preview models. This paragraph proposes the fitting of the Magic Formula, usually used for tyre modelling, to the non-linear response of the vehicle's yaw acceleration vs. steering velocity in terms of vehicle speed. The subject of the Magic Formula is reformulated, and used to determine the required steering input, for a given vehicle speed and desired yaw acceleration. The proposed steering driver is applied to the refined non-linear full vehicle model of a Sports Utility Vehicle (see paragraph 3.2), performing a severe double lane change manœuvre, and simulation results are compared to measured results. It is concluded that the proposed driver has definite merit, with

excellent correlation to test results.

The primary reason for requiring a driver model in the present study, is for the optimisation of the vehicle's suspension system. The suspension characteristics are to be optimised for handling, while performing the closed loop ISO3888-1 (1999) double lane change manœuvre. The driver model thus has to be robust for various suspension setups, and perform only one simulation to return the objective function value. Thus steering controllers with learning capability will not be considered, as the suspension could be vastly different from one simulation to the next. Only lateral path following is considered in this preliminary research, as the double lane change manœuvre is normally performed at a constant vehicle speed.

Previous research into lateral vehicle model drivers, was conducted amongst others by Sharp et al. (2000) who implement a linear, multiple preview point controller, with steering saturation limits mimicking tyre saturation, for vehicle tracking. The vehicle model used is a 5-degree-of-freedom (dof) model, with non-linear Magic Formula tyre characteristics, but no suspension deflection. This model is successfully applied to a Formula One vehicle performing a lane change manœuvre. Also Gordon et al. (2002) make use of a novel method, based on convergent vector fields, to control the vehicle along desired routes. The vehicle model is a 3-dof vehicle, with non-linear Magic Formula tyre characteristics, but with no suspension deflection included. The driver model is successfully applied to lane change manœuvres.

The primary similarity between these methods is that vehicle models with no suspension deflection were used. The current research is, however, concerned with the development of a controllable suspension system for Sports Utility Vehicles (SUV's). The suspension system thus has to be modelled, and the handling dynamics simulated for widely varying suspension settings. The vehicle in question has a comparatively soft suspension, coupled to a high center of gravity, resulting in large suspension deflections when performing the double lane change manœuvre. This large suspension deflection, results

in highly unstable vehicle behaviour, eliminating the use of driver models suited to vehicles with minimal suspension deflection. Steady state rollover calculations also show that the vehicle will roll over before it will slide.

Proköp (2001) implements a PID (Proportional Integral Derivative) prediction model for tracking control of a bicycle model vehicle. The driver model makes use of a driver plant model that is representative of the actual vehicle. The driver plant increases in complexity to perform the required dynamic manœuvre, from a point mass to a four wheel model with elastokinematic suspension. This model is then optimised with the SQP (Sequential Quadratic Programming) optimisation algorithm for each time step. This approach, however, becomes computationally expensive, when optimisation of the vehicle's handling is to be considered.

For the current research several driver model approaches were implemented, but with limited success. Due to the difficulty encountered with the implementation of a driver model for steering control, it was decided to characterize the whole vehicle system, using step steer, and ramp steer inputs, and observe various vehicle parameters. This lead to the discovery that the relationship between vehicle yaw acceleration vs. steering rate for various vehicle speeds appeared very similar to the side force vs. slip angle characteristics of the tyres. With this discovery it was decided to implement the proposed novel driver model, with the non-linear gain factor modelled with the Pacejka Magic Formula, normally used for tyre data.

3.4.1 Driver Model Description

To investigate the relationship between vehicle response and steering inputs, simulations were performed for various steering input rates (Figure 3.6, where t_s is the start of the ramp when the vehicle has reached the desired speed), at various vehicle speeds. It was found that there existed a trend very similar to the tyre's lateral force vs. slip angle at various vertical loads, (Figure 3.7) with the vehicle's yaw acceleration vs. steering rate at different vehicle

speeds (Figure 3.8). Because of this relationship it was postulated that the vehicle could be controlled by comparing the actual yaw acceleration to the desired yaw acceleration, and adjusting the steering input rate.

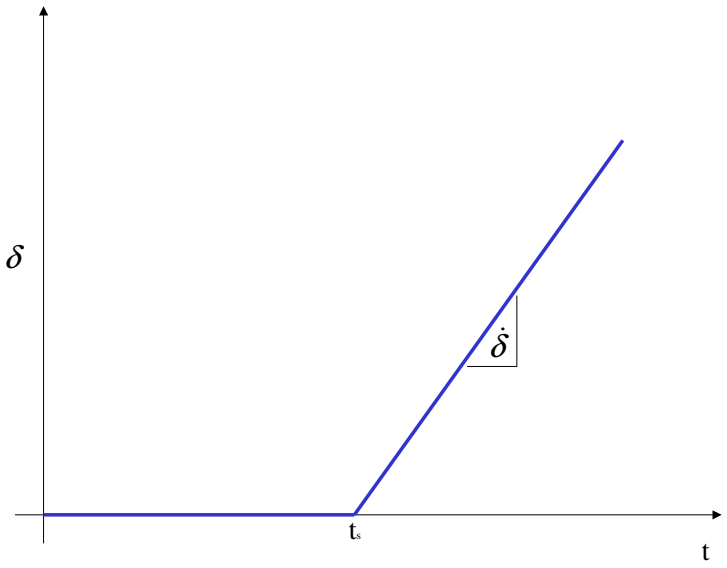

Figure 3.6: Vehicle characterisation steering input

From dynamics principles it is known that, for a rigid body undergoing motion in a plane, the rotational angle as a function of time is dependant on: the current rotational angle ϑ_0, the current rotational velocity $\dot{\vartheta}$, the rotational acceleration $\ddot{\vartheta}$, and the time step δt over which the rotational acceleration is assumed constant. If the rotational acceleration is not constant, but sufficiently small time steps are considered, the predicted rotational angle ϑ_p will be sufficiently well approximated. The predicted rotational angle can be determined as follows:

$$\vartheta_p = \vartheta_0 + \dot{\vartheta}\delta t + \frac{1}{2}\ddot{\vartheta}\delta t^2 \tag{3.2}$$

The above equation can be modified for a vehicle's yaw rotation motion by defining ϑ as the yaw angle ψ. Considering Figure 3.9, the driver model parameters can now be defined as:

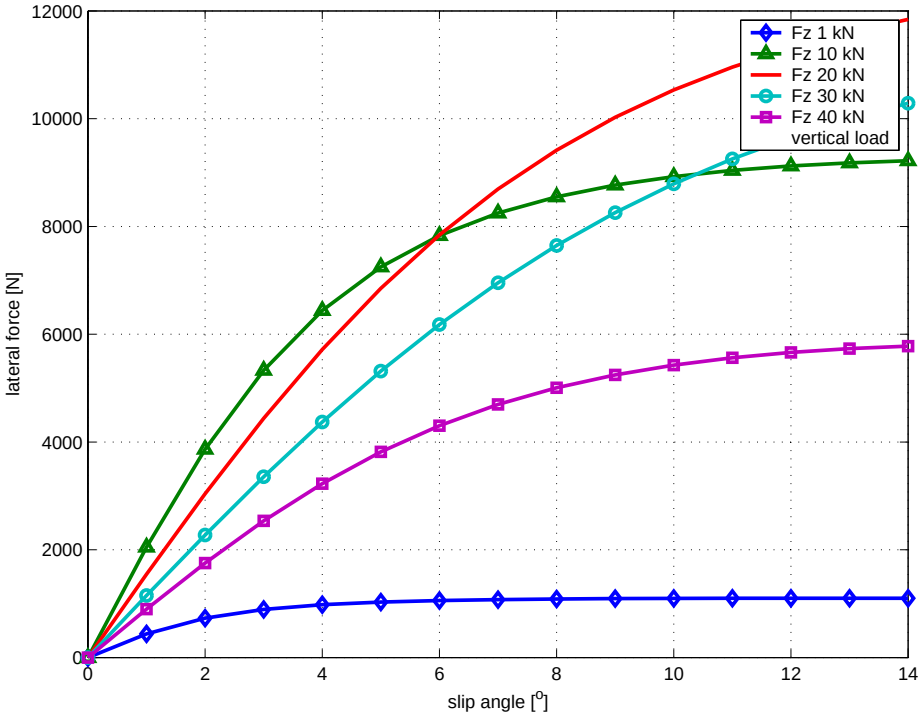

Figure 3.7: Tyre's lateral force vs. slip angle characteristics for different vertical loads

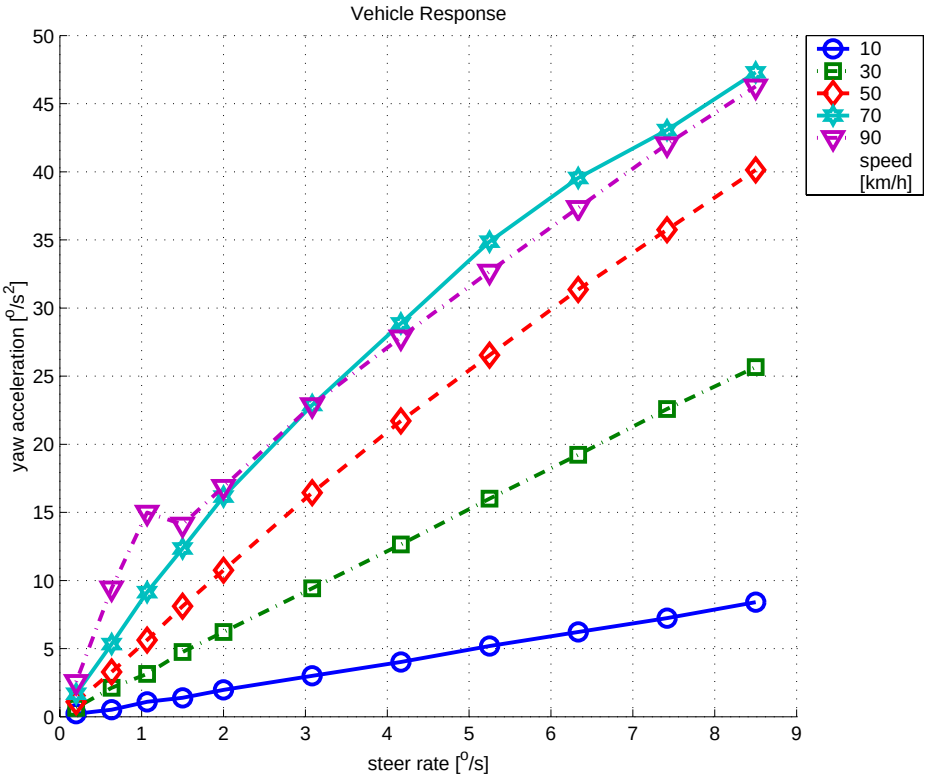

Figure 3.8: Vehicle yaw acceleration response to different steering rate inputs

- desired yaw angle of the vehicle ψ_d, equivalent to ϑ_p
- actual vehicle yaw angle ψ_a, equivalent to ϑ_0
- actual vehicle yaw rate $\dot{\psi}_a$, equivalent to $\dot{\vartheta}$
- response/preview time τ, equivalent to δt
- vehicle forward velocity $\dot{x}$

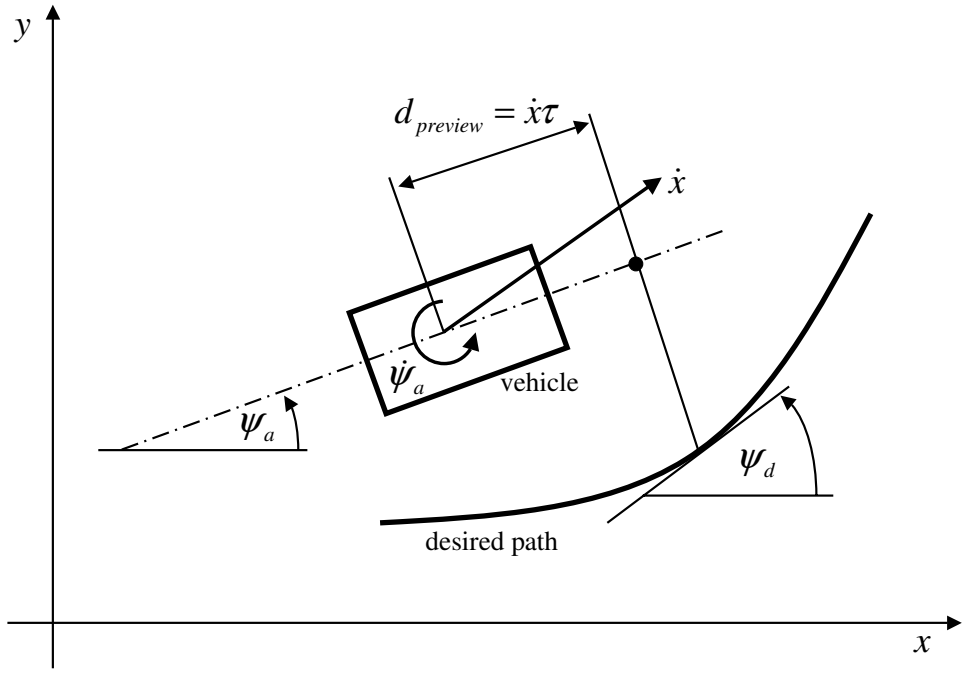

Figure 3.9: Definition of driver model parameters

The yaw acceleration $\ddot{\psi}$ needed to obtain the desired yaw angle is calculated from equation (3.2), substituting in the equivalent variables, as follows:

$$\ddot{\psi} = 2\frac{\psi_d - \psi_a - \dot{\psi}_a\tau}{\tau^2} \tag{3.3}$$

The vehicle's steady state yaw acceleration $\ddot{\psi}$ with respect to different steering rates $\dot{\delta}$, was determined for a number of constant vehicle speeds $\dot{x}$ and is presented in Figure 3.8. Where the vehicle's response did not reach steady-state, and the vehicle slided out, or rolled over, the yaw acceleration just prior to loss of control was used. This process is computationally

expensive as 11 different steering ramp rates, for each vehicle speed, were applied to the vehicle model and simulated. The steady state yaw acceleration reached was then used to generate the figure. When comparing Figure 3.8 to the vehicle's lateral tyre characteristics, presented in Figure 3.7, it appears reasonable that the Magic Formula could also be fitted to the steering response data. Therefore the reformulated Magic Formula, discussed below, is fitted to this data, and returns the required steering rate $\dot{\delta}$, which is defined as:

$$\dot{\delta} = f(\ddot{\psi}, \dot{x}) \tag{3.4}$$

As output, the driver model provides the required steering rate $\dot{\delta}$, which is then integrated for the time step δt to give the required steering angle δ.

The Magic Formula is fitted through the obtained data, as it is a continuous function over the fitted range. Normal polynomial curve fits would be discreet for the vehicle speed they are fitted to and an interpolation scheme would be necessary for in-between vehicle speeds. The Magic Formula is thus a continuous approximation described by 12 values, as opposed to multiple curve formulae, requiring intermediate interpolation.

3.4.2 Magic Formula Fits

The Magic Formula was proposed by Bakker et al. (1989) to describe the tyre's handling characteristics in one formula. In the current study the Magic Formula will be considered in terms of the tyre's lateral force vs. slip angle relationship, which directly affects the vehicle's handling and steering response. The Magic Formula is defined as:

$$\begin{aligned} y(x) &= Dsin(Carctan\{Bx - E(Bx - arctan(Bx))\}) \\ Y(X) &= y(x) + S_v \\ x &= X + S_h \end{aligned} \tag{3.5}$$

The terms are defined as:

- $Y(X)$ tyre lateral force F_y
- X tyre slip angle α

- B stiffness factor
- C shape factor
- D peak factor
- E curvature factor
- S_h horizontal shift
- S_v vertical shift

These terms are dependent on the vertical tyre load F_z and camber angle γ. The lateral force F_y vs. tyre slip angle α relationship typically takes on the shape as indicated in Figure 3.7, for different vertical loads. Considering the shape of Figure 3.8 presenting the yaw acceleration vs. steering rate for different vehicle speeds, the Magic Formula can be successfully fitted, with the parameters redefined as:

- vertical tyre load F_z is equivalent to vehicle speed $\dot{x}$
- tyre slip angle α is equivalent to steering rate $\dot{\delta}$
- tyre lateral force F_y is equivalent to vehicle yaw acceleration $\ddot{\psi}$

The Magic Formula for the vehicle's steering response can now be stated as:

$$\begin{aligned} y(x) &= Dsin(Carctan\{Bx - E(Bx - arctan(Bx))\}) \\ Y(X) &= y(x) + S_v \\ x &= X + S_h \end{aligned} \tag{3.6}$$

With the terms defined as:

- $Y(X)$ yaw acceleration $\ddot{\psi}$
- X steering rate $\dot{\delta}$
- B stiffness factor
- C shape factor
- D peak factor

- E curvature factor
- S_h horizontal shift
- S_v vertical shift

With the redefined parameters, the Magic Formula coefficients can be determined in the usual manner. The determination of the coefficients applied for the steering driver is now discussed. The baseline vehicle's response as indicated in Figure 3.8 is used for the fitting of the parameters.

3.4.3 Determination of Factors

The peak factor D is determined by plotting the maximum yaw acceleration value $\ddot{\psi}$ against the vehicle speed $\dot{x}$. For this the graphs have to be interpolated. Quadratic curves were fitted through the vehicle's response curves, and the estimated peak values were used. The peak factor is defined as:

$$D = a_1\dot{x}^2 + a_2\dot{x} \tag{3.7}$$

The peak factor curve was fitted through the estimated peak values, with emphasis on accurately capturing the data for vehicle speeds of 50 to 90 km/h. The 90 km/h peak was taken as the point where the graph changed due to the maximum yaw acceleration just prior to roll-over. The resulting quadratic curve fit to the predicted peak values of the yaw acceleration is shown in Figure 3.10. It is observed that the fit for the Magic Formula is poor for 30 km/h. This is attributed to the almost linear curve fit through the yaw acceleration vs. steering rate for speeds of 10 and 30 km/h Figure 3.6, resulting in an unrealistically high prediction of the peak values.

In the original paper (Bakker et al. 1989), BCD is defined as the cornering stiffness, here it will be termed the 'yaw acceleration gain'. For the yaw acceleration gain the gradient at zero steering rate is plotted against vehicle speed as illustrated in Figure 3.11. The camber term γ of the original paper will be ignored so that coefficient a_5 becomes zero. The yaw acceleration gain is fitted with the following function:

$$BCD = a_3 sin(2arctan(\dot{x}/a_4))(1 - a_5\gamma) \tag{3.8}$$

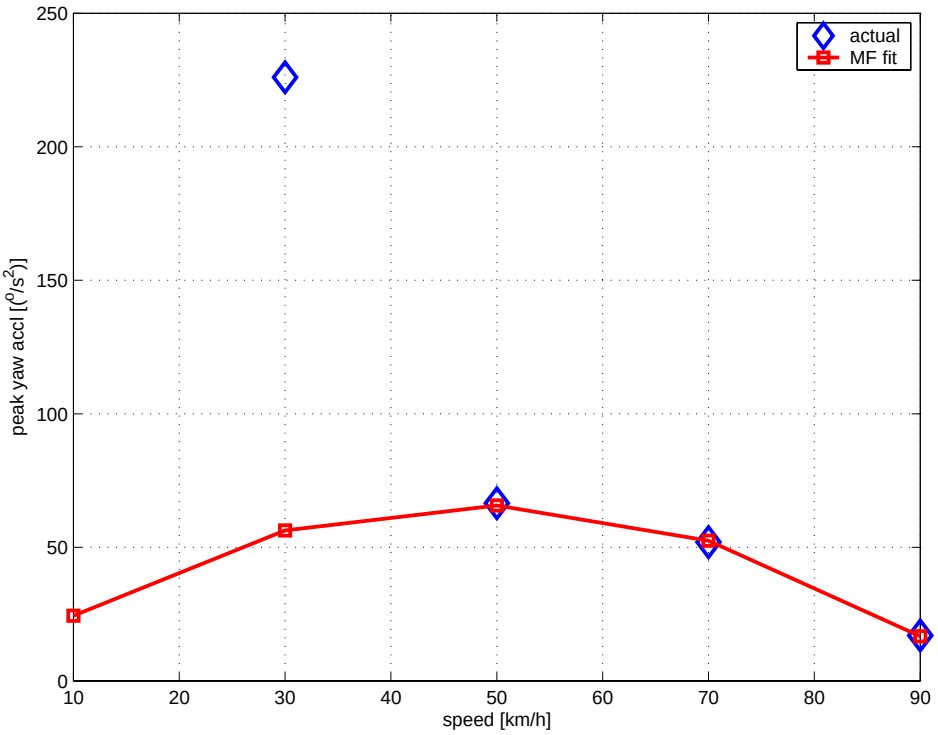

Figure 3.10: Magic Formula coefficient quadratic fit through equivalent peak values

For the determination of the curvature E, quadratic curves were fitted through each of the curves in Figure 3.8. These approximations could then be differentiated twice to obtain the curvature for each. This curvature is plotted against vehicle speed $\dot{x}$, and the straight line approximation:

$$E = a_6\dot{x} + a_7 \tag{3.9}$$

is then fitted through the data points, in order to determine the coefficient a_6 and a_7. The straight line approximation fitted through the points is shown in Figure 3.12.

The shape factor C, is determined by optimising the resulting Magic Formula fits to the measured data. This parameter is the only parameter that has to be adjusted in order to achieve better Magic formula fits to the original

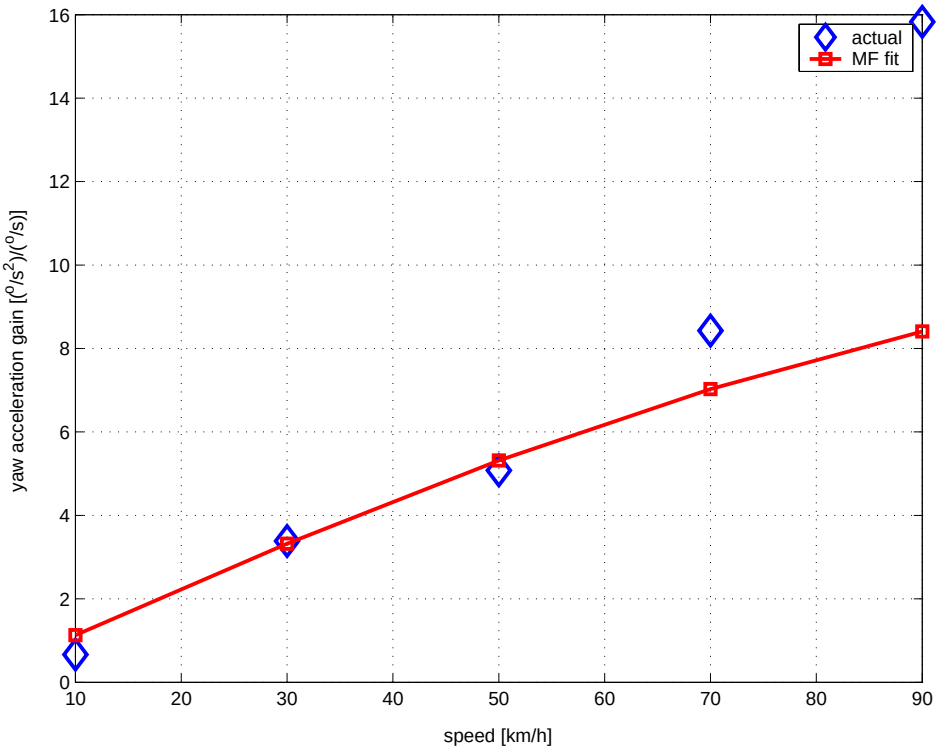

Figure 3.11: Magic Formula fit of yaw acceleration gain through the actual data

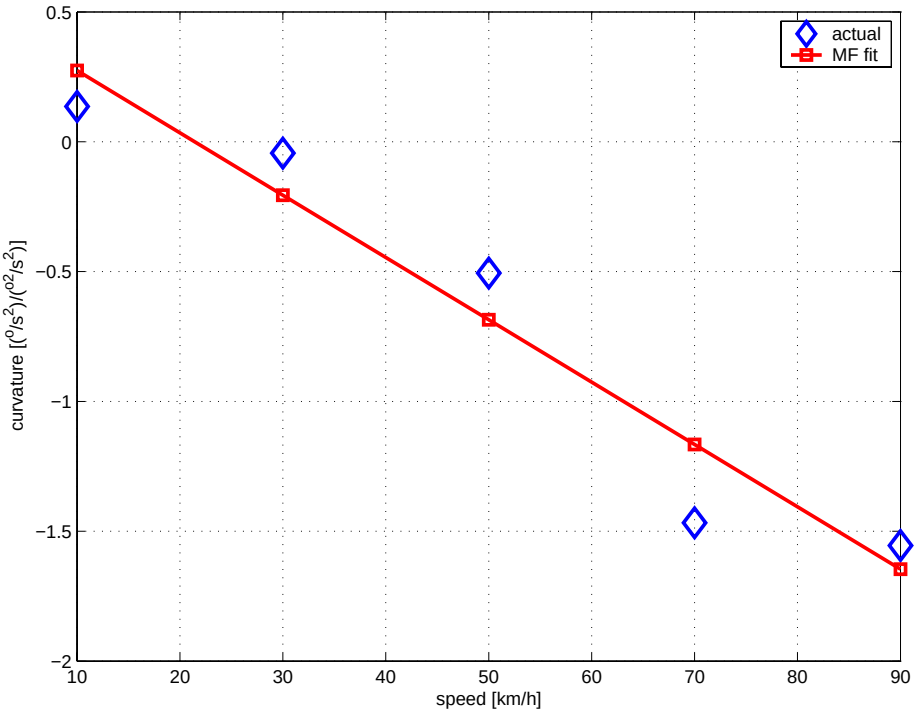

Figure 3.12: Determination of curvature coefficients

data. It is defined in terms of the Magic Formula coefficient a_0 as follows:

$$C = a_0 \tag{3.10}$$

The stiffness factor B is determined by dividing BCD by C and D:

$$B = BCD/CD \tag{3.11}$$

In the current research the horizontal and vertical shift of the curves were ignored allowing coefficients a_8 to a_{13} to be assumed zero. The Magic Formula fits to the original data are presented in Figure 3.13. It can be seen that most of the fits except for 90 km/h are very good. The vehicle simulation failed for most of the steering rate inputs before reaching a steady state yaw acceleration at 90 km/h, thus this can be viewed as an unstable regime. With the Magic formula coefficients being determined, the manipulation of the Magic formula for the driver application is discussed.

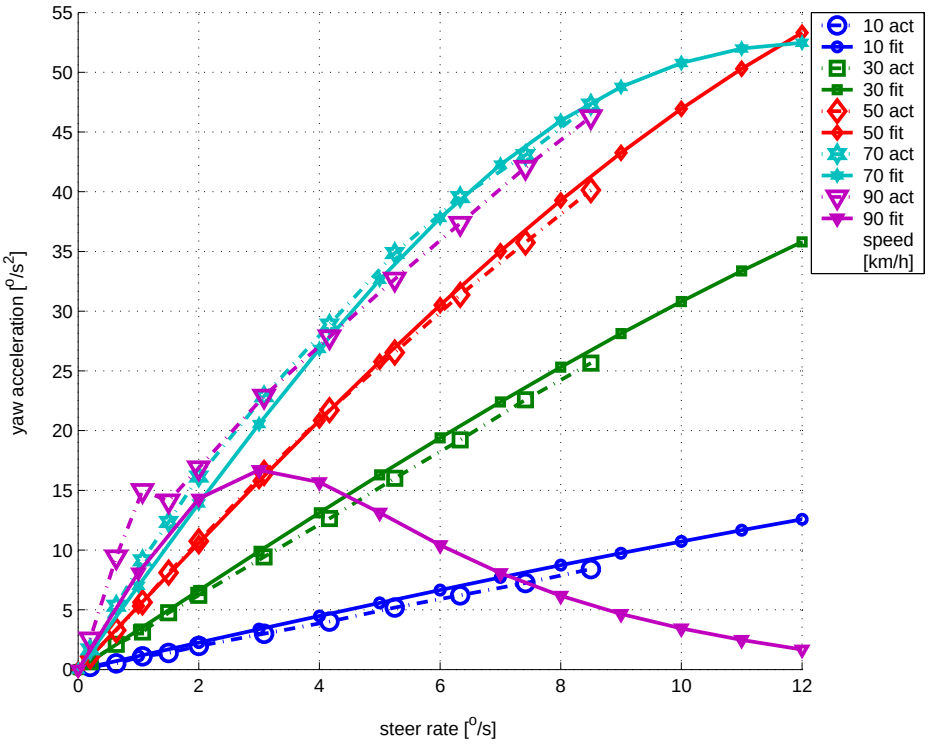

Figure 3.13: Magic Formula fits to original vehicle steering behaviour

3.4.4 Reformulated Magic Formula

The steering driver requires, as output, the steering rate $\dot{\delta}$. For this reason the Magic Formula's subject of the formula must be reformulated, to make

it possible to have as inputs, vehicle velocity $\dot{x}$ and required vehicle yaw acceleration $\ddot{\psi}$, and as output required steering rate $\dot{\delta}$. However, due to the nature of the Magic Formula it is not possible to change the subject of the formula, so the *arctan* function is described by the *pseudo arctan* function as suggested by Pacejka (2002) as:

$$psatan(x) = \frac{x(1 + a\,|x|)}{1 + 2(b\,|x| + ax^2)/\pi} \tag{3.12}$$

where $a = 1.1$ and $b = 1.6$. The Magic Formula can now be written as:

$$F = Bx - E\left(Bx - \frac{Bx(1+a|Bx|)}{1+2(b|Bx|+a(Bx)^2)/\pi}\right)$$
$$F = tan\left(\frac{arcsin(\frac{y}{D})}{C}\right) \tag{3.13}$$

This equation was solved symbolically for x using MATLAB's Symbolic Toolbox, and returns an exceptionally long equation, of three terms, not presented here due to its shear size. This resulting equation is coded into the Simulink model consisting of the MSC.ADAMS model and the steering controller. It should be noted that the solution to equation (3.13) will return multiple answers as the shape of the graphs in Figure 3.13 suggest. Only the first part of the graphs, up to the peak/maximum point, was used, with the peak point used as a limit for higher steering rates.

3.4.5 Implementation of Results

In order to validate the performance of the proposed methodology, the Magic Formula driver was implemented in the vehicle simulation model. The vehicle was simulated performing the ISO3888-1 (1999) double lane change manœuvre. The excellent comparison to measured results is presented in Figure 3.14, for kingpin steering angle, yaw velocity, left rear (lr) spring displacement and left front (lf) lateral acceleration. It is important to note that the double lane change is simulated at a constant speed (see results in Figure 3.5) while the measured results show that the driver decreased speed during the manœuvre, explaining the slight discrepancies towards the end of the double lane change.

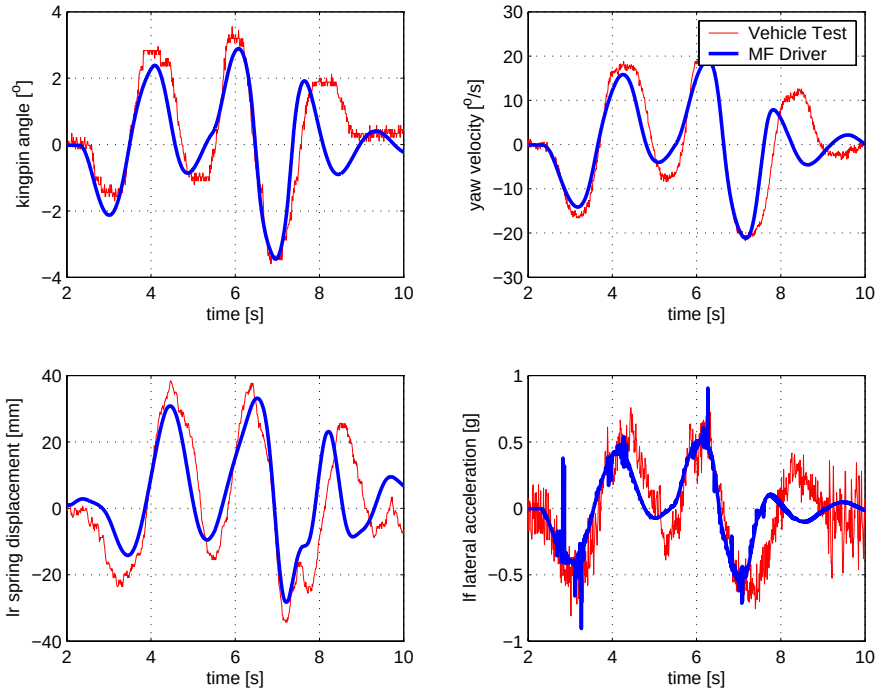

Figure 3.14: Correlation of Magic Formula driver model to vehicle test at an entry speed of 63 km/h

The driver model was then analysed for changing the vehicle's suspension system from stiff to soft, for various speeds. Presented in Figure 3.15 is the driver model's ability to keep the vehicle at the desired yaw angle (Genta 1997) over time. From the results it can be seen that a varying preview time with vehicle speed, would be beneficial, however, it is felt that for this preliminary research the constant 0.5 seconds preview time is sufficient. Also it is evident that the softer suspension system, and 70 km/h vehicle speed, are slightly unstable, as seen by the oscillatory nature at the end of the double lane change manœuvre.

The results show that the driver model provides a well controlled steering input. Also there is a lack of high frequency oscillation typically associated with single point preview driver models, when applied to highly non-linear vehicle models like SUV's, that are being operated close to their limits in the

double lane change manœuvre.

3.5 Conclusions

It has been shown that the Magic Formula, traditionally used for describing tyre characteristics, can be fitted to the vehicle's steering response, in the form of yaw acceleration vs. steering rate, for different vehicle speeds.

A single point steering driver model has been successfully implemented on a highly non-linear vehicle model. The success of the driver model, is attributed to the modelling of the vehicle's response with the Magic Formula. The success of the single point steering driver can be related to the non-linear gain factor, that changes in value with vehicle speed and required yaw acceleration.

Future work should entail an investigation into determining the parameters of the vehicle that modify the tyre characteristic Magic Formula coefficients to arrive at the steering rate and yaw acceleration parameters. Ideally the tyre Magic Formula coefficients should be multiplied by some modifying factor, based on vehicle characteristics, to be used directly for the control of the vehicle steering. This would eliminate the need for the computationally expensive characterisation currently required. A further aspect that could be considered is determining the value of varying preview time with vehicle speed. The driver model is, however, sufficiently robust to be used in the optimisation of the vehicle's suspension characteristics for handling.

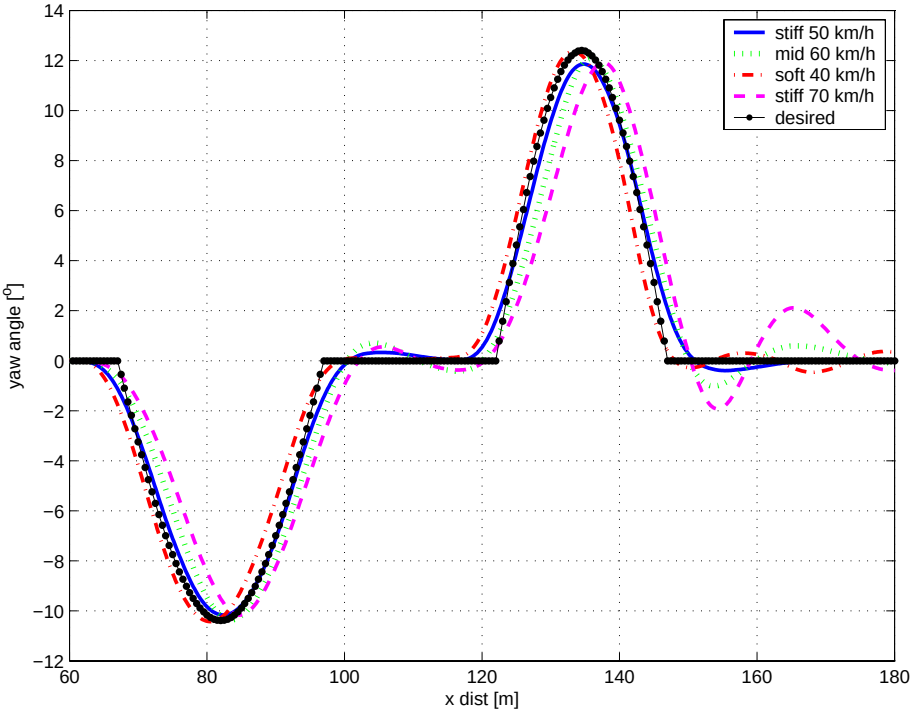

Figure 3.15: Comparison of different suspension settings and vehicle speeds, for the double lane change manœuvre, where the desired is as proposed by Genta (1997)

Chapter 4

Finite Difference Gradient Information

In the masters thesis Thoresson (2003), the feasibility of using gradient-based approximation methods for the optimisation of the spring and damper characteristics of an off-road vehicle, for both ride comfort and handling, was investigated. The Sequential Quadratic Programming (SQP) algorithm and the locally developed Dynamic-Q method were the two successive approximation methods used for the optimisation. The determination of the objective function value is performed using computationally expensive numerical simulations that exhibit severe inherent numerical noise. The use of forward finite differences and central finite differences for the determination of the gradients of the objective function within Dynamic-Q is also investigated. The results of this study, presented here, proved that the use of central finite differencing for gradient information improved the optimisation convergence history, and helped to reduce the difficulties associated with noise in the objective and constraint functions.

This chapter presents the feasibility investigation of using gradient-based successive approximation methods to overcome the problems of poor gradient information due to severe numerical noise. The two approximation methods applied here are the locally developed Dynamic-Q optimisation algorithm of Snyman and Hay (2002) and the well known Sequential Quadratic

Programming (SQP) algorithm, the MATLAB implementation being used for this research (Mathworks 2000b). This chapter aims to provide the reader with more information regarding the Dynamic-Q successive approximation algorithm, that may be used as an alternative to the more established SQP method. Both algorithms are evaluated for effectiveness and efficiency in determining optimal spring and damper characteristics for both ride comfort and handling of a vehicle.

The initial vehicle model was used in this part of the optimisation investigation. This vehicle is fitted with the $4S_4$ suspension system, and up to four design variables are considered in the optimisation of the suspension characteristics. The details of which were discussed in detail in Chapter 3, Section 3.1.

It is found that both optimisation algorithms perform well in optimising handling. However, difficulties are encountered in obtaining improvements in the design process when ride comfort is considered. This is attributed to the very noisy nature of the ride comfort objective function, which incorporates computed vertical accelerations. Nevertheless, meaningful design configurations are still achievable through the proposed optimisation process, at a relatively low cost in terms of the number of simulations that have to be performed.

4.1 Optimisation Algorithms

The following optimisation algorithms are evaluated in this chapter:

- The Dynamic-Q method (see Section 2.3) which constructs a sequence of simple spherical quadratic approximations to the original problem, and successively solves these sub-problems via the LFOPC (leapfrog) algorithm (Snyman 2000). The gradients used by the algorithm are evaluated by forward finite differences, or by central finite differences, with sufficiently large steps to smooth out the numerical noise.

- The Sequential Quadratic Programming (SQP) method (see Section 2.2) with Broyden- Fletcher-Goldfarb-Shanno approximations to the Hessian matrix and one-dimensional minimization by quadratic polynomial interpolation in the search directions. This method is invoked by the MATLAB m-file fmincon, which also uses forward finite difference approximations for the gradients.

4.2 Gradient Approximation Methods

Most gradient-based optimisation algorithms require the determination of the first and/or second order gradient information of the objective and constraint functions with respect to the design variables. In most engineering optimisation problems this gradient information is not analytically available. The only information available to the designer is the values of objective and constraint functions obtained via expensive simulations. This research investigated the use of forward and central finite differences in the Dynamic-Q optimisation algorithm, for the determination of the first order gradient information. The formulation of forward finite difference gradient information was presented in Section 2.4.1, and central finite difference gradient determination described in Section 2.4.2.

4.3 Optimisation

The vehicle model used is the initial vehicle model that was modelled in ADAMS View 12 and described in Section 3.1. This model made use of the 521 tyre model, which was later found to be not sufficiently suited for the desired correlation with measured data. In this initial study the optimisation parameters were defined as in the following sections.

4.3.1 Design Variables

In choosing the design variables for optimisation, the assumption is made that the left hand and right hand suspension settings will be the same, but that front and rear settings may differ. The design variables chosen for

optimisation are the static gas volume (Figure 4.1), and damper force scale factor (Figure 4.2), on both the front and rear axles. Thus there are two variables per axle.

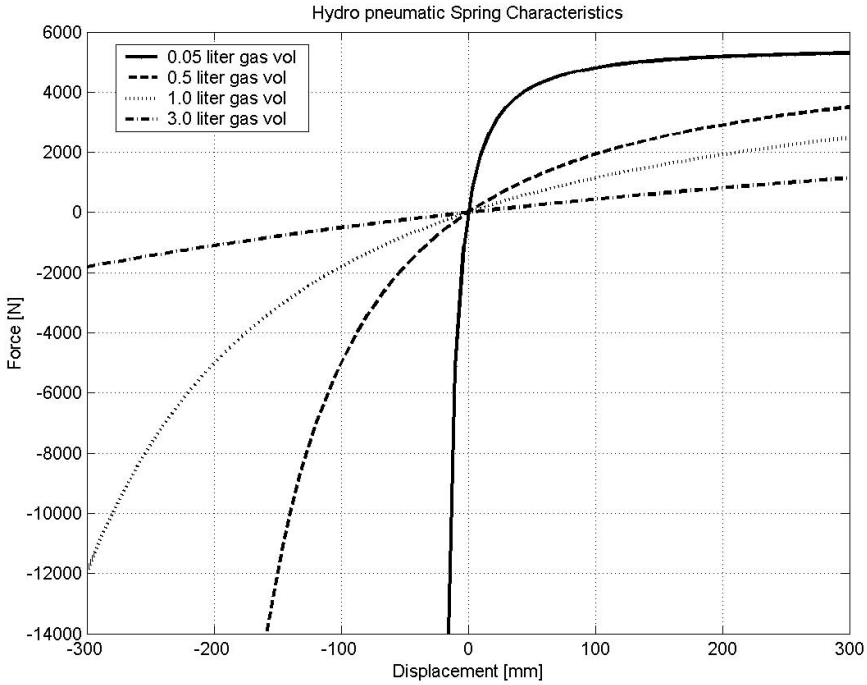

Figure 4.1: Definition of spring characteristics for various gas volumes

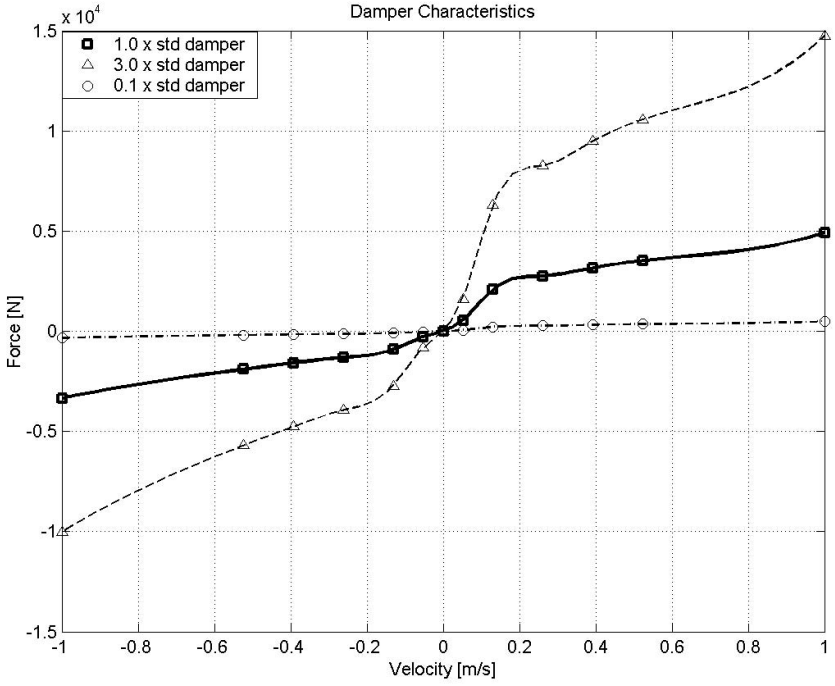

Figure 4.2: Definition of damper characteristics for various damper scale factors

For this initial study the standard damper force characteristic is multiplied by a factor which constitutes the damping design variable (Figure 4.2). The general shape and switch velocities of the damper are thus kept the same. This chapter only considers the cases of two and four 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).

4.3.2 Two Variable Case

The two design variable study is an important starting point in the optimisation procedure as it gives the necessary insight into the problem. For this two design variable study, it was decided to use the same design variables as those considered by Els and Uys (2003) in their preliminary study, namely the static gas volume and the damper force scale factor. Figure 4.1 illustrates the spring characteristics for various static gas volumes. Figure 4.2 illustrates the damper characteristics for various damper scale factors.

The static gas volume is denoted by $gvol$, and the damper force scale factor by $dpsf$. These variables are allowed to range from 0.05 to 3 in magnitude, which are accordingly chosen as upper and lower bounds. The design variables are explicitly defined as follows:

$$x_1 = gvol, \qquad x_2 = dpsf \tag{4.1}$$

with bounds

$$0.05 \leq x_i \leq 3, \qquad i = 1, 2 \tag{4.2}$$

4.3.3 Four Variable Case

For the four design variable problem the front and rear settings are uncoupled. This means 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 variables are also allowed to range from 0.05 to 3 in magnitude. Thus the design variables are defined explicitly as follows:

$$\begin{aligned} x_1 &= dpsff, \qquad x_2 = gvolf, \\ x_3 &= dpsfr, \qquad x_4 = gvolr \end{aligned} \tag{4.3}$$

with bounds

$$0.05 \leq x_i \leq 3, \qquad i = 1, ..., 4 \tag{4.4}$$

4.3.4 Definition of Objective Functions

For ride comfort, the motion of the vehicle is simulated for travelling in a straight line over the Belgian paving (Gerotek 2006) and the sum of driver and rear passengers British Standard (BS6841 1987) weighted root mean square (RMS) vertical accelerations are used for the objective function. 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 $1 \times 10^{-4}\ m^2/(cycles/m)$, and a terrain index ω of 4 (Thoresson 2003). In a study performed by Els (2005), it was found that the BS6841 weighed RMS vertical acceleration corresponds well with subjective responses of ride comfort in off-road vehicles. For this reason the weighted RMS vertical accelerations will be used for the objective function, when considering ride comfort optimisation. The motion sickness component was ignored as it requires long run times and the Belgian paving test track is not long enough to evaluate motion sickness. No additional measures were used for the ride comfort objective function, despite numerous studies (Alleyne and Hedrick 1995, Kim and Ro 1998, Pilbeam and Sharp 1996, Miller 1998) where the tyre deflection or force is used as a measure of road holding, when considering a quarter car model. This was ignored as the suspension system has the ability to switch to the handling setting should a handling condition be detected. Also handling

lateral and roll degrees of freedom cannot be simulated with a quarter car model. Kim and Ro (2001) also suggest that tyre deflection is insufficient for evaluating handling parameters. This afforded the optimisation for ride comfort and handling to be done separately.

For handling, the vehicle performs a ISO3888-1 (1999) double lane change manœuvre at 80 km/h and the maximum body roll angle at the first peak (Els and Uys 2003) is used as the objective function. For this initial investigation only roll angle was used as a measure of handling as suggested by Uys et al. (2006a).

4.3.5 Design Space

For the two design variable optimisation, surface plots of the objective function over the complete design space were generated. However, with an increasing number of variables added this is not possible. These objective function surfaces were generated for the optimisation of handling (Figure 4.3) and ride comfort (Figure 4.4) separately. From the figures it can be seen that for excellent handling capability we require high damping and high spring stiffness, however, the damping does not really contribute to the improvement if the spring stiffness is high (a small gas volume in Figure 4.3).

However, for ride comfort (Figure 4.4) we find that the opposite holds. The lowest spring stiffness and low damping is required, but a medium spring stiffness and low damping results in a minimal decrease in ride comfort compared to the optimum. The damper scale factor has a more noticeable effect on the ride comfort, as established previously by Els and Uys (2003) for the heavier version of this vehicle.

4.3.6 Handling Results

No significant problems were encountered in applying the algorithms to the optimisation of handling. For handling optimisation with two design variables both algorithms converged to an optimum without difficulty. Figure

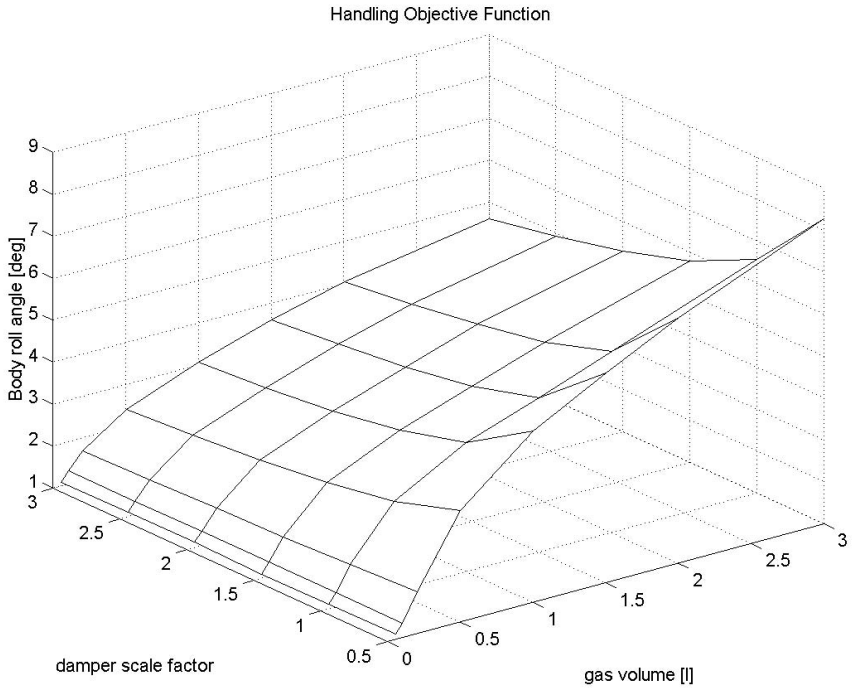

Figure 4.3: Vehicle roll angle, double lane change at 80 km/h for the two variable design space

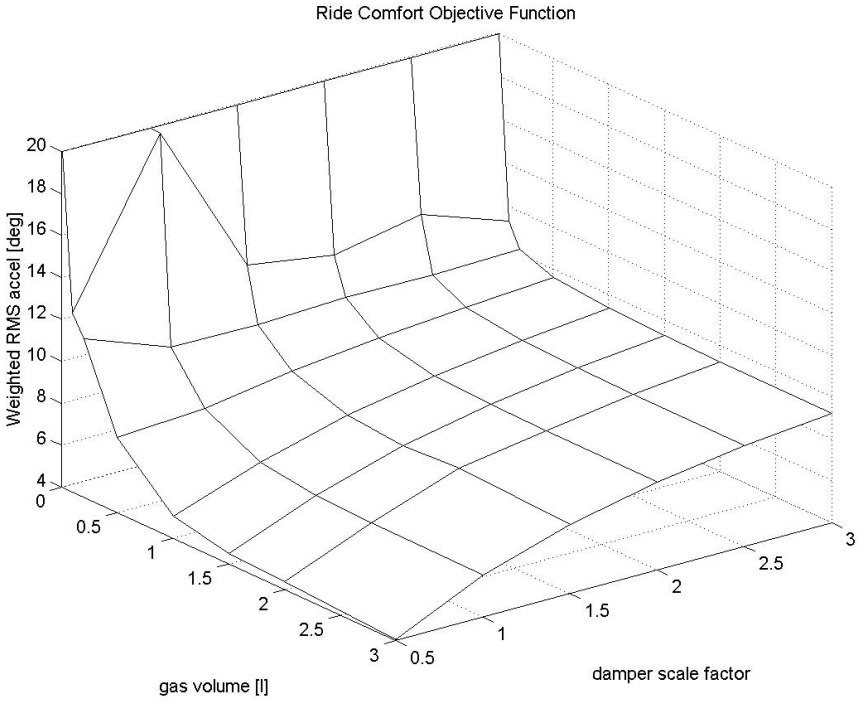

Figure 4.4: Vehicle Ride comfort, Belgian paving at 60 km/h for the two variable design space

4.5, depicting the convergence histories for each algorithm, plots the objective function value against number of function evaluations. The number of function evaluations is used for the x-axis, instead of iterations, as the cost of the optimisation depends on the number of function evaluations performed, and not the number of iterations, due to the computationally expensive nature of the simulation model. Note that each successive marker on the graphs denotes a new iteration, and that each iteration does not necessarily require the same number of function evaluations. The SQP convergence history for handling optimisation (Figure 4.5) indicates two local minimum solution sets with the same objective function value. This is observed when comparing the two distinct values of design variable $x(2)$ (at 15 and 30 function evaluations for example) that result in little or no change in the objective function value. Because of the cost of the function evaluations, the objective function values are plotted against cumulative number of function evaluations at the iteration point. The use of Dynamic-Q with 10 % move limit (Figure 4.5) re-iterates the fact that design variable two (damper multiplication factor) has a limited effect on the objective function value as has already been established in Figure 4.3. Using a 20 % move limit (Figure 4.5) Dynamic-Q progresses faster to a minimum. Because of the excellent performance of the forward finite difference method the use of central finite differences at additional cost was not necessary.

The handling optimisation results for four design variables (Figure 4.7) were not really different to that for two variables. This can be expected as the dynamics of the system has not changed substantially. It is interesting to note that a move limit of as big as 30 % of the variable's range may be used in Dynamic-Q using forward finite differences. It can also be seen from Figure 4.6 that the optimisation histories are very well behaved. Figure 4.6 again indicates the definite existence of more than one local minimum with the same objective function value, but significantly different design variable values. This is attributed to the 'flat' region in the design space, where the objective function is relatively insensitive to the design variables. The SQP algorithm performed almost similarly to Dynamic-Q, and also found two different local minima, with the same objective function value. SQP converged in 9

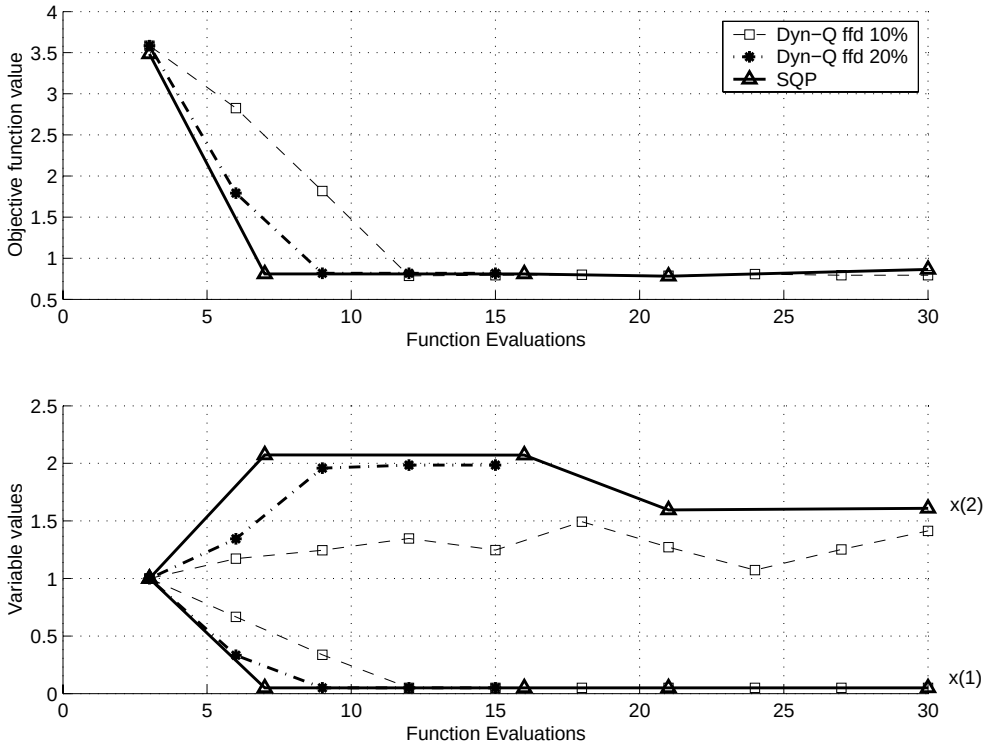

Figure 4.5: Handling optimisation, 2 design variables

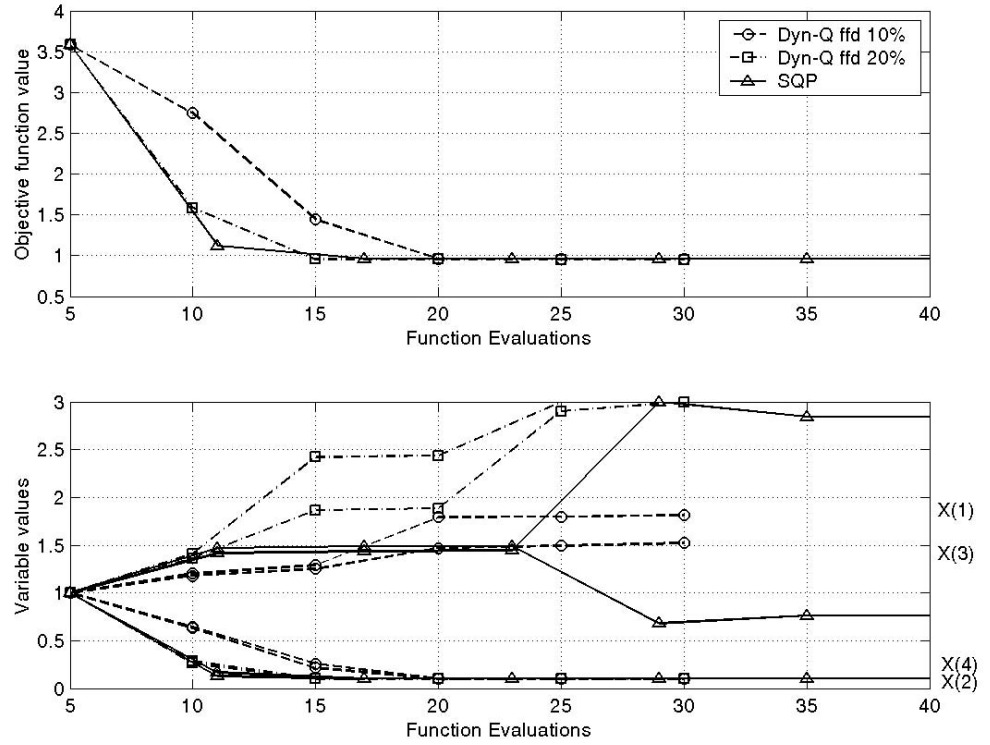

Figure 4.6: Handling optimisation, 4 design variables

iterations (49 function evaluations), and Dynamic-Q in 5 iterations (25 function evaluations). The two different optima correspond to the cases where the damper scale factors are respectively the same and different at the front and rear. Large differences in performance of the optimisation algorithms are only expected when the number of design variables increases. This reinforces our initial conclusion from the results for two variables: that the damper scale factor has negligible effect on the vehicle's handling performance (body roll angle) through the double lane change manœuvre at the optimum (stiff) spring rate. It can also be observed from the results that the spring gas volume (two design variable optimisation, variable x_1, four design variable optimisation, variables x_2 and x_4) runs to the lowest bound, corresponding to the maximum possible spring stiffness (smallest possible gas volume).

4.3.7 Ride Comfort Results

A summary of the results of the optimisation for ride comfort is shown in Table 4.1. Two design variable ride comfort optimisation encountered the problems associated with a noisy objective function. It is postulated that the severe noise present in the ride comfort objective function, as opposed to the smooth nature of the handling objective function, is related to the fact that for ride comfort, accelerations are used, while for handling, angular displacement is used for the objective function measure, however, the filtering and use of the RMS value is a traditional smoothing effect, which is not evident from the results though. Although not immediately apparent at this stage, especially when considering the course mesh of Figure 4.4 used to get a feel of the form of the optimisation problem, it will be shown that the objective function exhibits severe noise.

The SQP method (Figure 4.7) took 8 iterations (33 function evaluations) to stabilise on a minimum, corresponding to the lowest possible damping and stiffness, as expected from Figure 4.4. The Dynamic-Q method experienced greater difficulties in obtaining a stable minimum. For this reason, the

central finite difference method for determining the gradient, was introduced to obtain stability in the optimisation process. The Dynamic-Q method with central finite differences, with a 10 % move limit (Figure 4.7) took 9 iterations (50 function evaluations) to find a minimum, with inspection showing that this minimum is effectively reached after only 4 iterations (25 function evaluations). The vertical acceleration at this point is, however, significantly higher than that found with SQP indicating the existence of a separate interior local minimum. A 20 % move limit (Figure 4.7) took 6 iterations (30 function evaluations), finding a local minimum not far off the SQP minimum. The Dynamic-Q minimum design variable values are not at the extrema found by the SQP method, reinforcing the fact that the ride comfort design space has a flat plateau of local minima.

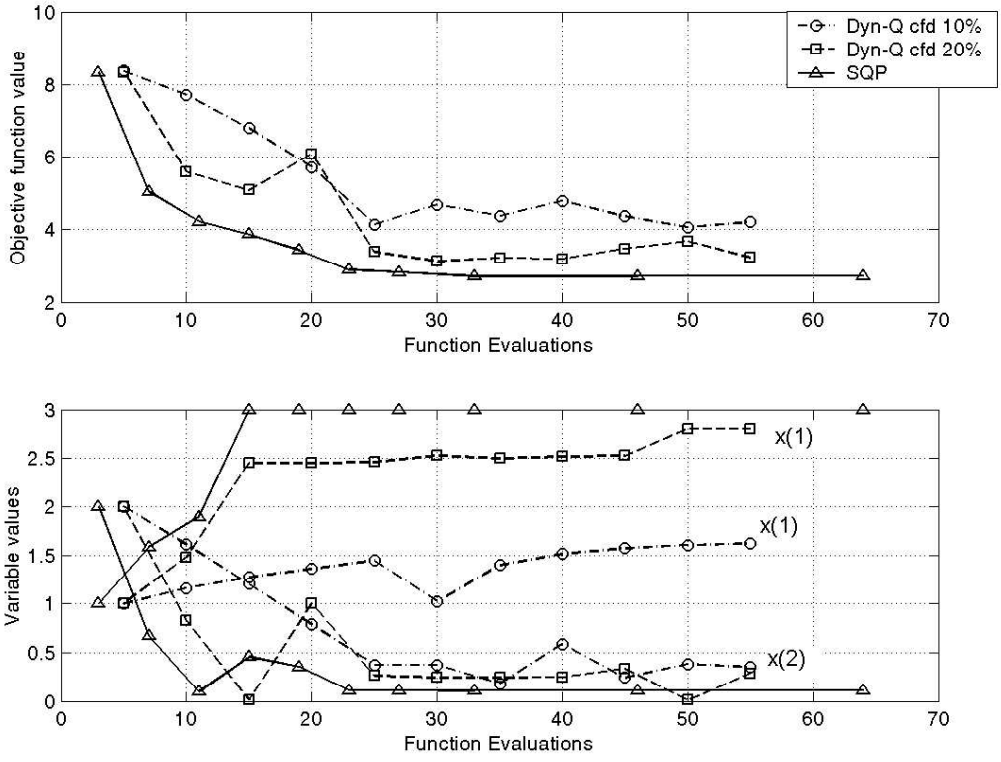

Figure 4.7: Ride comfort optimisation, 2 design variables

For the four design variable optimisation Dynamic-Q was modified so that the move limit for each iteration is 90 % of the move limit of the previous iteration. This was done so as to stabilise the convergence behaviour of the

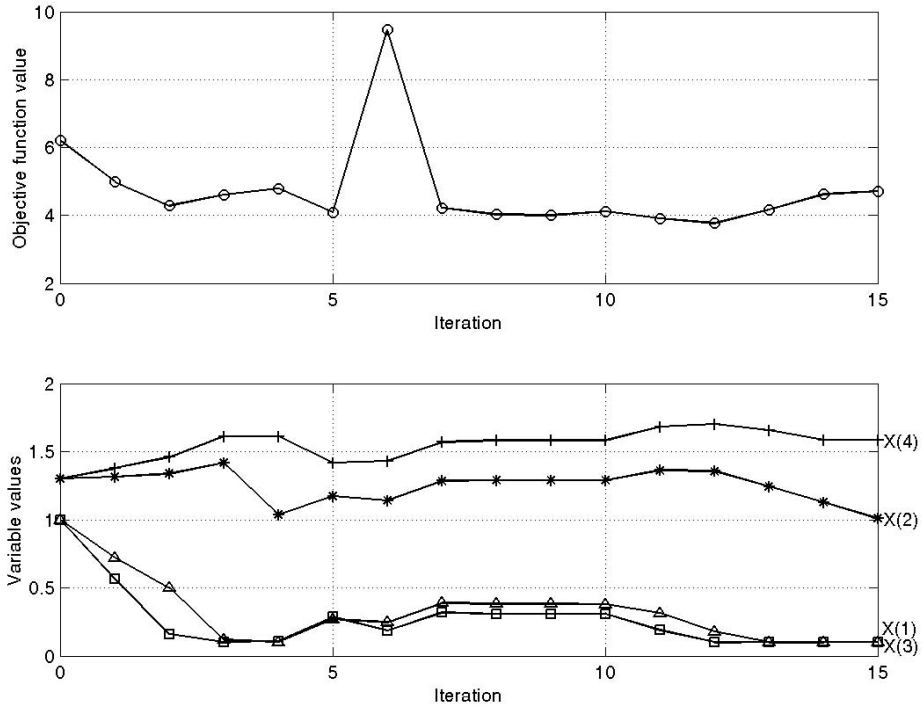

Figure 4.8: Dynamic-Q ffd ride comfort, 4 design variables, 10 % move limit

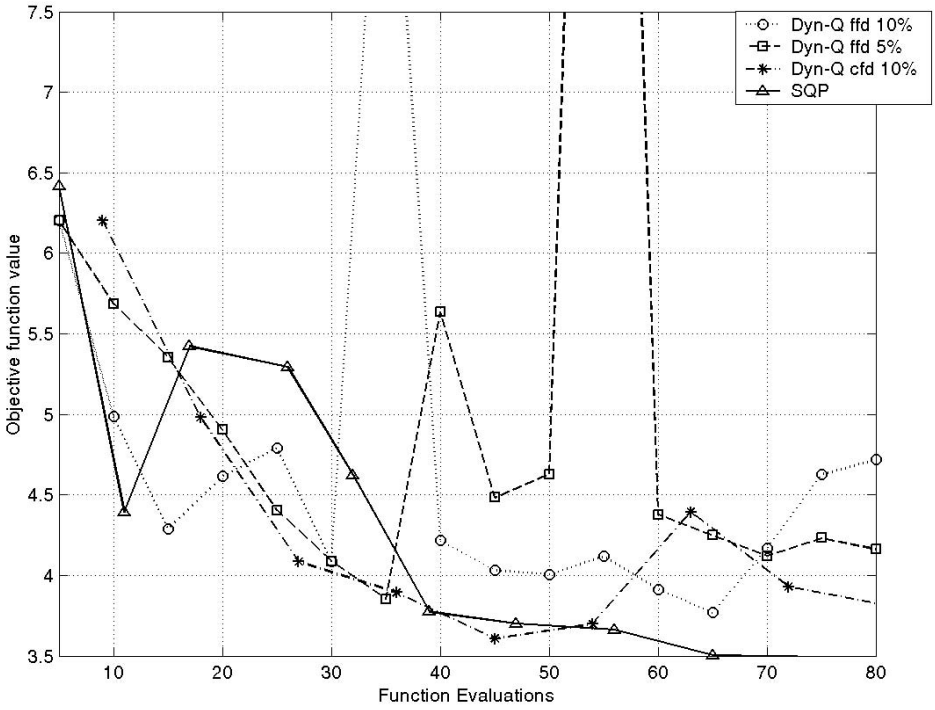

Figure 4.9: Ride comfort optimisation, 4 design variables

Table 4.1: Ride Comfort Optimisation Results

Algorithm	Move Limit	Figure Number	Iterations	Function Evaluations	Optimum value
Two Design Variables					
SQP	-	4.7	8	33	2.7
Dynamic-Q central	10 %	4.7	9 (4)	50 (25)	4.1
finite differences	20 %	4.7	6	30	3.1
Four Design Variables					
Dynamic-Q forward	10 %	4.8	12	65	3.8
finite differences	5 %	4.9	6	35	3.9
Dynamic-Q cfd	10 %	4.9	4	45	3.6
SQP	-	4.9	8	65	3.5

algorithm and to try and prevent high spikes in the optimisation process. These spikes are caused by a poor approximation to the objective function close to the minimum, resulting in the LFOPC algorithm finding a minimum of the approximated problem on the slope of the steep valley close to the actual minimum. However, Dynamic-Q quickly recovers within a single iteration (5 function evaluations) as can be seen in Figure 4.8 iteration 7. The results of the optimisation are presented for both central finite differences and forward finite differences used for the gradient approximations in Figure 4.9.

An alternate explanation for the spiky nature of the optimisation convergence histories is the numerical noise of the objective function. This becomes apparent when evaluating Figure 4.7 where it can be observed that for a relatively small change in the design variable values (Dyn-Q cfd 10 %) there is a relatively large change in the objective function value. This is also observed in Figure 4.8 iterations 5 and 6.

From Figures 4.8 and 4.9 for the forward finite difference Dynamic-Q implementation, it can be seen that the smaller move limit of 5 % is more stable reaching a minimum within 6 iterations (35 function evaluations),

while 10 % move limit takes 12 iterations (65 function evaluations). The algorithm however does not converge due to the noisy objective function with steep valley. The convergence behaviour for central finite differences coupled to Dynamic-Q is shown in Figure 4.9 requiring 4 iterations (45 function evaluations). Again it has been determined that the smaller move limit is beneficial to finding the minimum. The central finite difference gradient evaluation builds into the system a level of robustness. From the results it can be seen that around 1.5 liter gas volume and limited damping returns the best results. The central finite difference results show that by increasing the rear gas volume with minimal damping, a better overall ride can be achieved (Figure 4.9).

SQP also found similar good results within 8 iterations (65 function evaluations) (Figure 4.9). From Table 1 it is concluded that Dynamic-Q with forward finite differences does not reach the same minimum as Dynamic-Q with central finite differences. Dynamic-Q with central finite differences is also comparatively economical to SQP, finding an minimum within 5 % of the SQP minimum objective function value.

The levels of tyre vertical acceleration associated with the obtained optimum design conditions were evaluated. It was found that compared to the baseline vehicle (Figure 4.10), the tyre does experience high levels of acceleration, which is associated with tyre hop, when driving in a straight line over the Belgian paving. However, these accelerations are not transmitted to the vehicle body, so that the objective function value calculated is indeed an optimal value. Due to the presence of tyre hop, it is suggested that even when considering ride comfort and handling separately, a measure of the tyre hop (vertical force, deflection or acceleration) should also be considered when optimising the vehicle's suspension for ride comfort, this is later included in the work performed in Chapter 5.

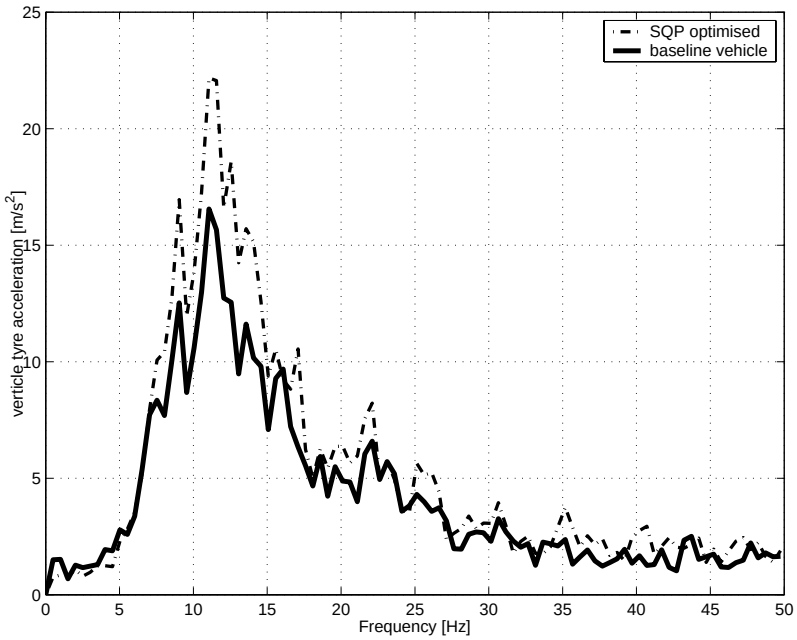

Figure 4.10: Tyre hop investigation. Vertical tyre accelerations for SQP optimised suspension compared to baseline vehicle.

4.4 Conclusion

The feasibility of using gradient-based approximation methods for the optimal design of a vehicle's suspension was investigated. An industry-standard version of the SQP method, and the in-house Dynamic-Q method, were evaluated. The determination of the objective function was performed using a full multi-body vehicle simulation model that was both computationally expensive, and exhibited severe inherent numerical noise when considering, in particular, ride comfort. The goal was to determine the vehicle's optimal spring and damper characteristics for both ride comfort and handling.

It is concluded that both optimisation algorithms work exceptionally well when optimising for handling. However, it is felt that the damper should play a role in the optimum suspension settings and thus a damper dependent measure should also be considered in the definition of the handling objective function.

Although difficulties were experienced in the ride comfort optimisation due to the severe inherent noise in the objective function, both algorithms sufficiently overcame this problem by yielding locally optimal feasible solutions. It was found that for the Dynamic-Q algorithm the use of central finite differences for the gradient approximations, at the cost of $2n + 1$ function evaluations per iteration, achieved meaningful optima at a lower cost in terms of total number of function evaluations (simulations), than the SQP method. This can be attributed to the inherent stability that the central finite differencing technique introduces by considering information ahead and behind the current iteration point. Dynamic-Q can thus be strongly recommended for applications in vehicle suspension optimisation. The increase in tyre hop over the baseline vehicle suggests that tyre hop should be added as a constraint when optimising ride comfort, and cannot be neglected, as previously postulated.

The gradient-based approximation methods considered here prove to be feasible optimisation methods when noisy objective functions are to be optimised. These methods have the distinct advantage of requiring relatively few function evaluations, each of which corresponds to an expensive numerical simulation, before reaching an optimal design. It is concluded from this work that other means of eliminating the negative effects of the numerical noise should be investigated. Also a means of decreasing the total number of expensive numerical simulations even further, should be investigated as this method will still be prohibitively expensive when more design variables are considered. It is also concluded that a move limit of 10 % is a good general value to be used in Dynamic-Q. The Dynamic-Q algorithm will now be used further as it is an in-house code making access to the source code easy, for the implementation of the ideas of the rest of this research.

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

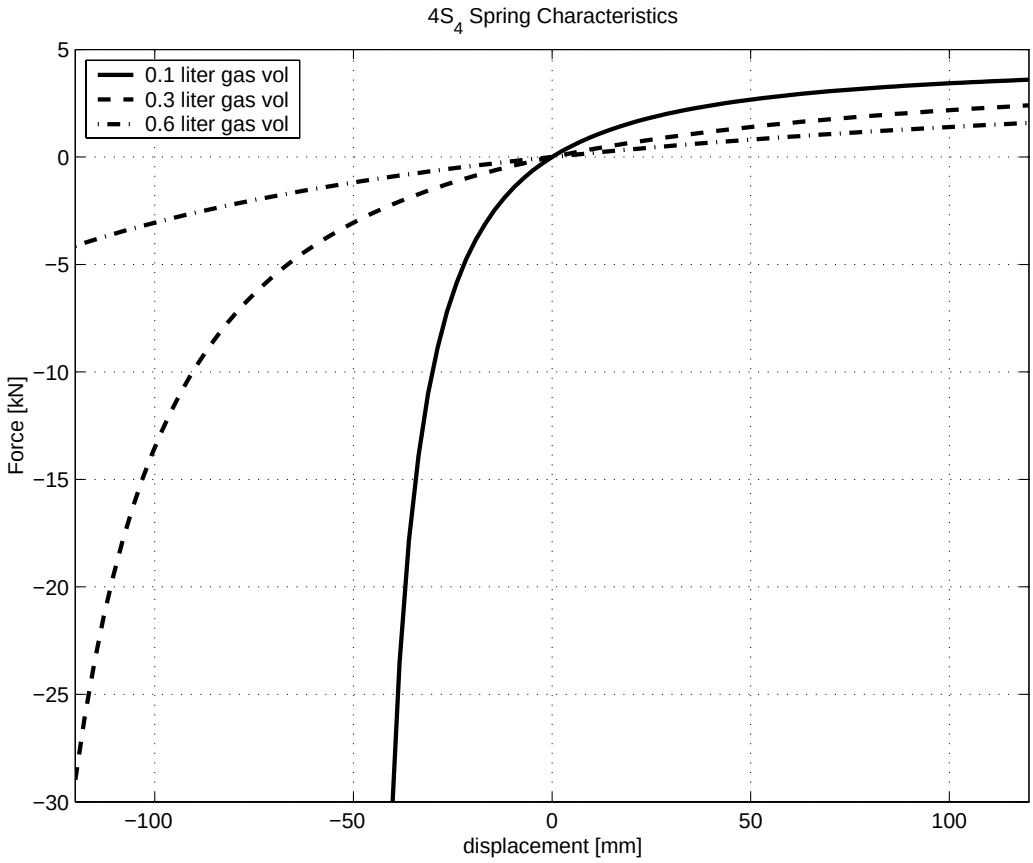

Figure 5.1: Definition of $4S_4$ spring characteristics for various gas volumes

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

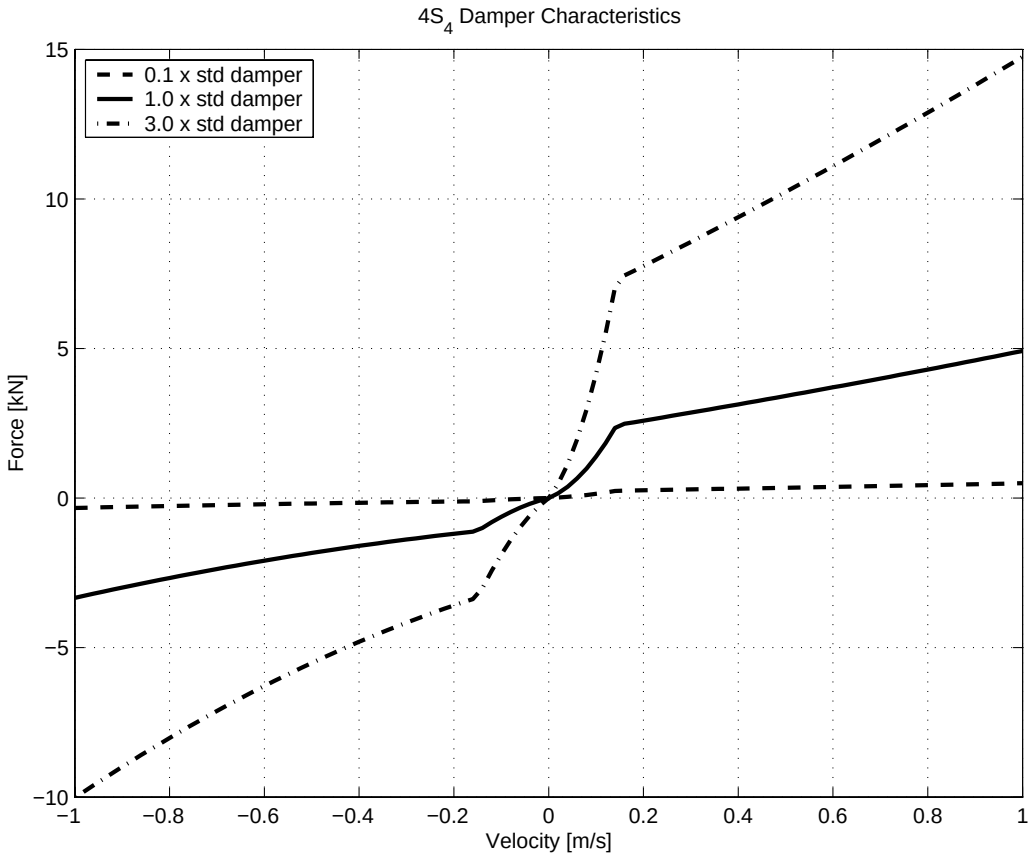

Figure 5.2: Definition of $4S_4$ damper characteristics for various damper scale factors

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

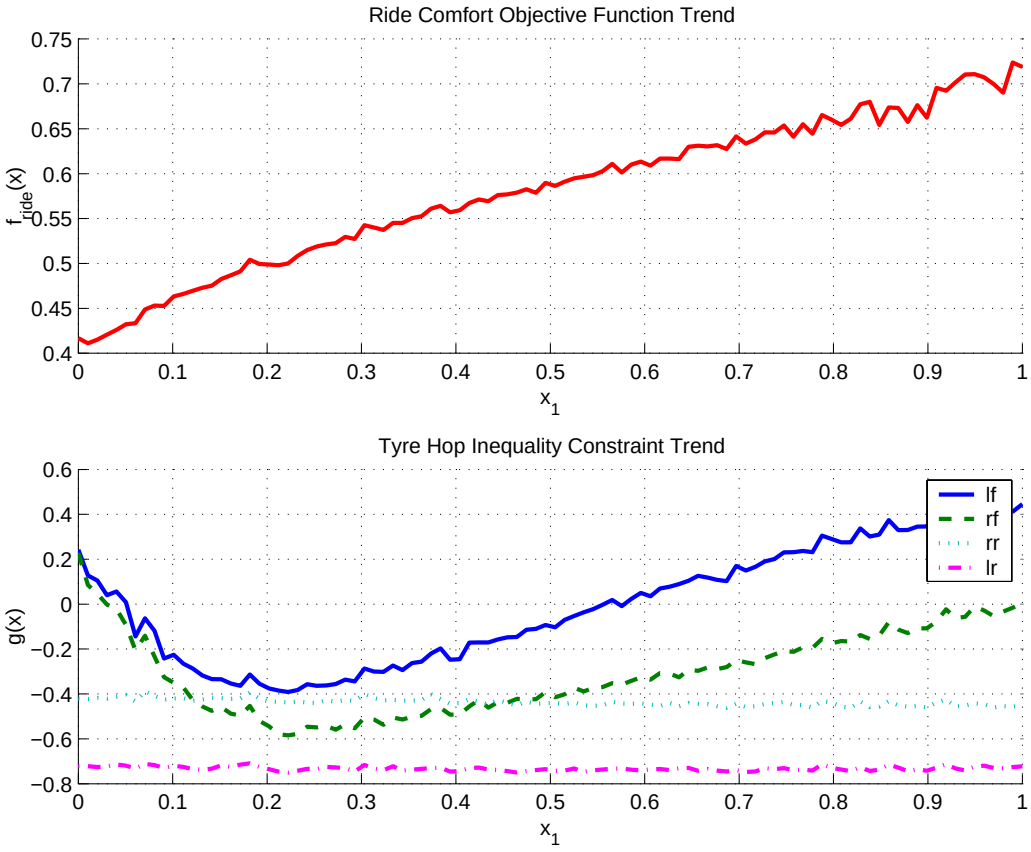

Figure 5.3: Level of inherent numerical noise in objective function and inequality constraints, for change in front damper design variable x_1, for full vehicle MSC.ADAMS model

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:

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

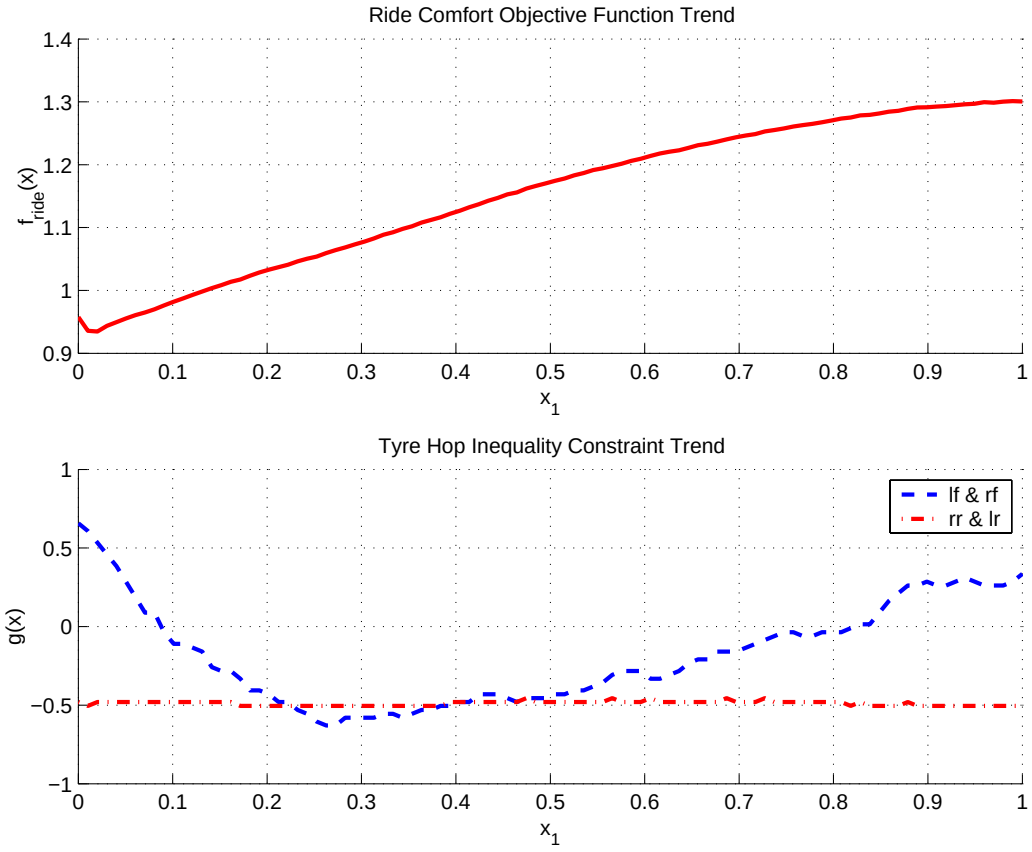

Figure 5.4: Level of inherent numerical noise in objective function and inequality constraints, for change in front damper design variable x_1, when considering the simplified MATLAB model

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

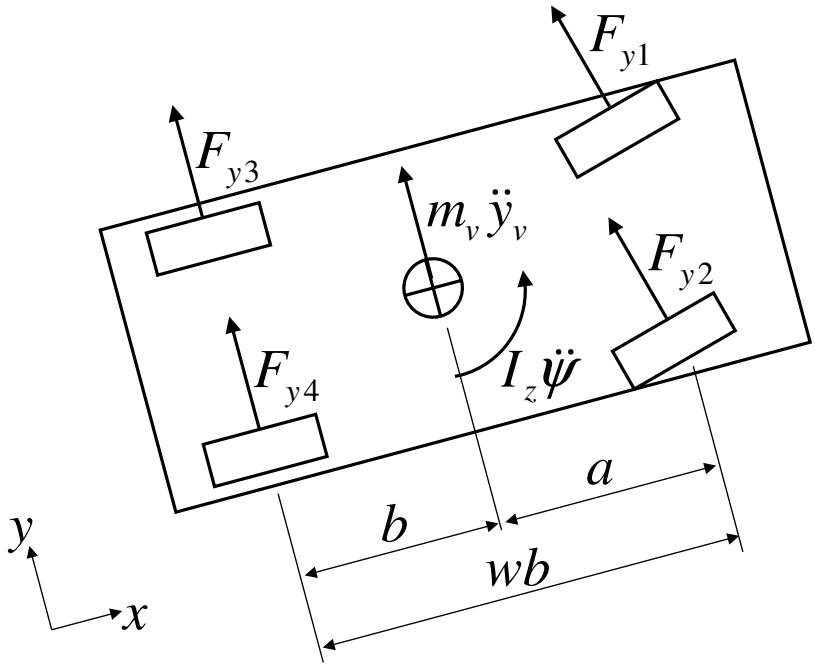

Figure 5.5: Top view of vehicle in handling manœuvre

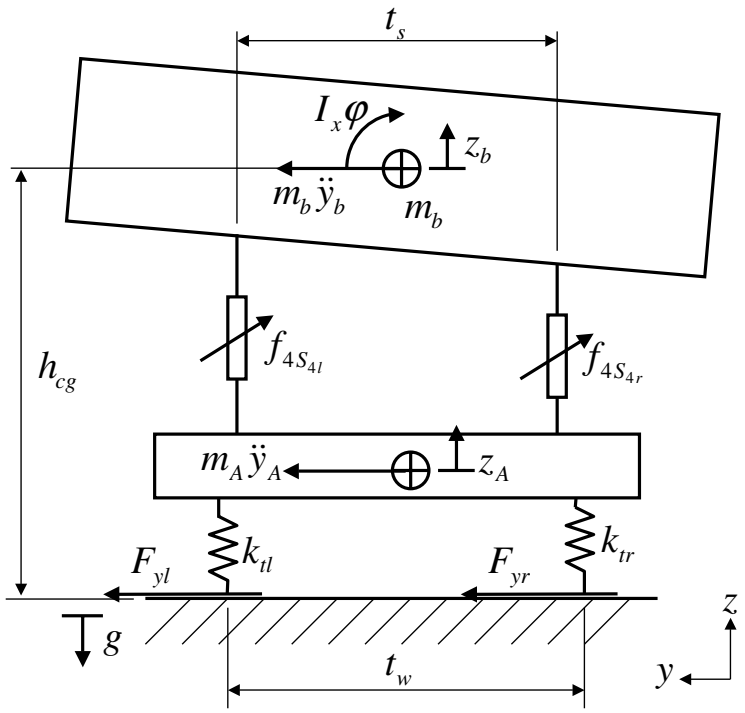

Figure 5.6: Rear view of vehicle indicating body roll

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:

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

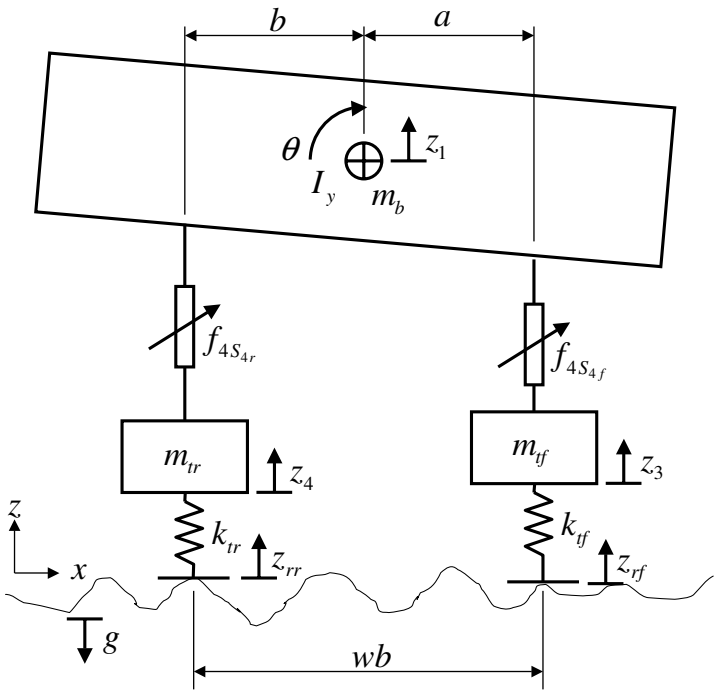

Figure 5.7: Simple pitch-plane vehicle model

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.

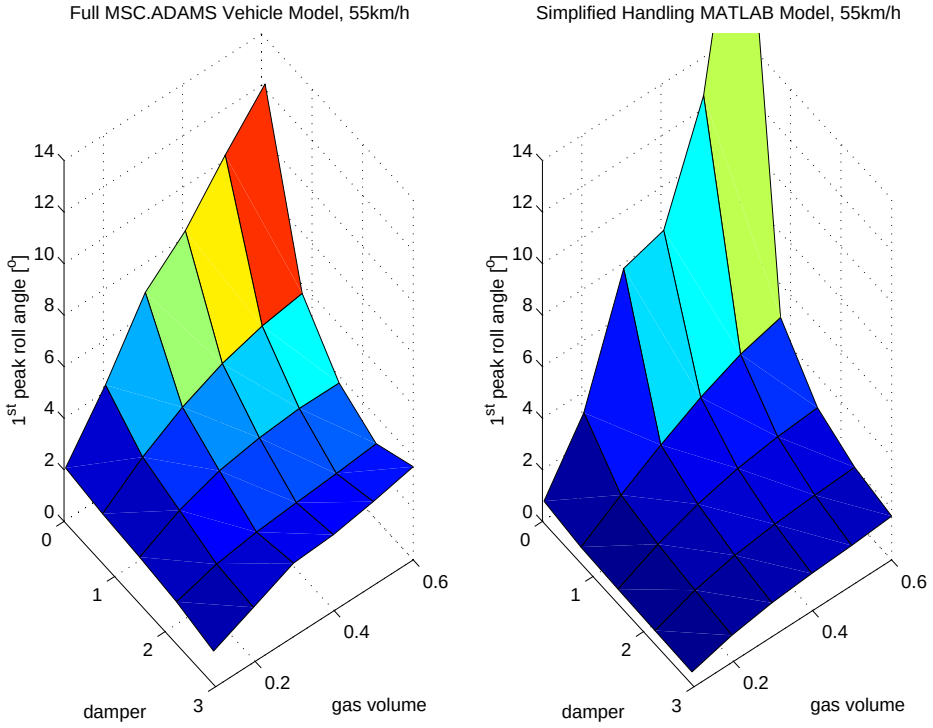

Figure 5.8: Validation of 1st peak roll angle over design space, for double lane change.

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.

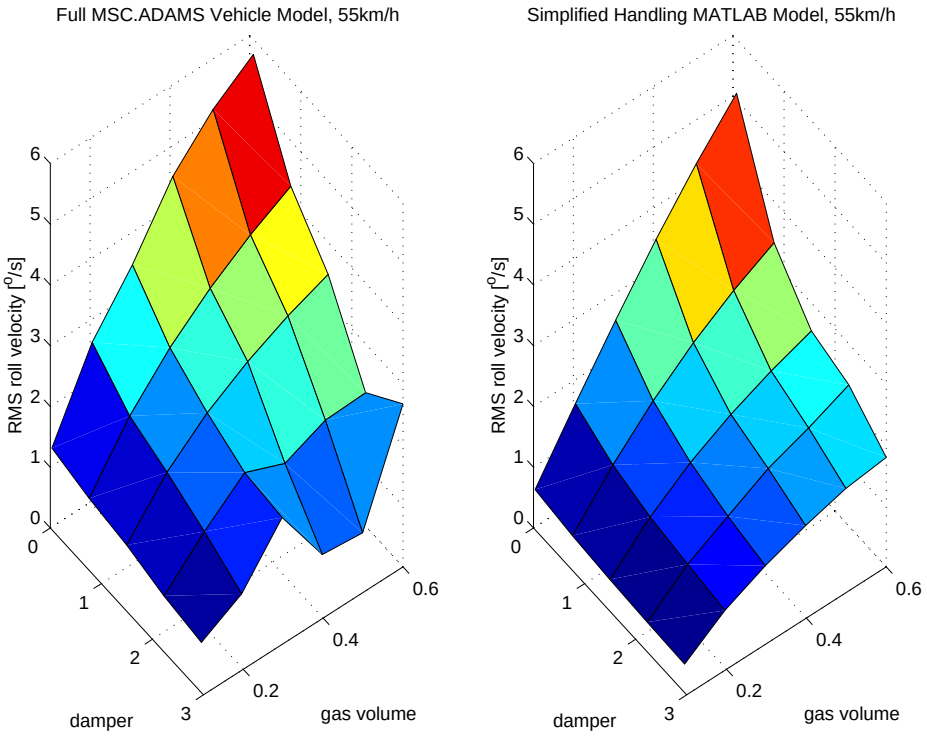

Figure 5.9: Validation of RMS roll velocity over design space, for double lane change.

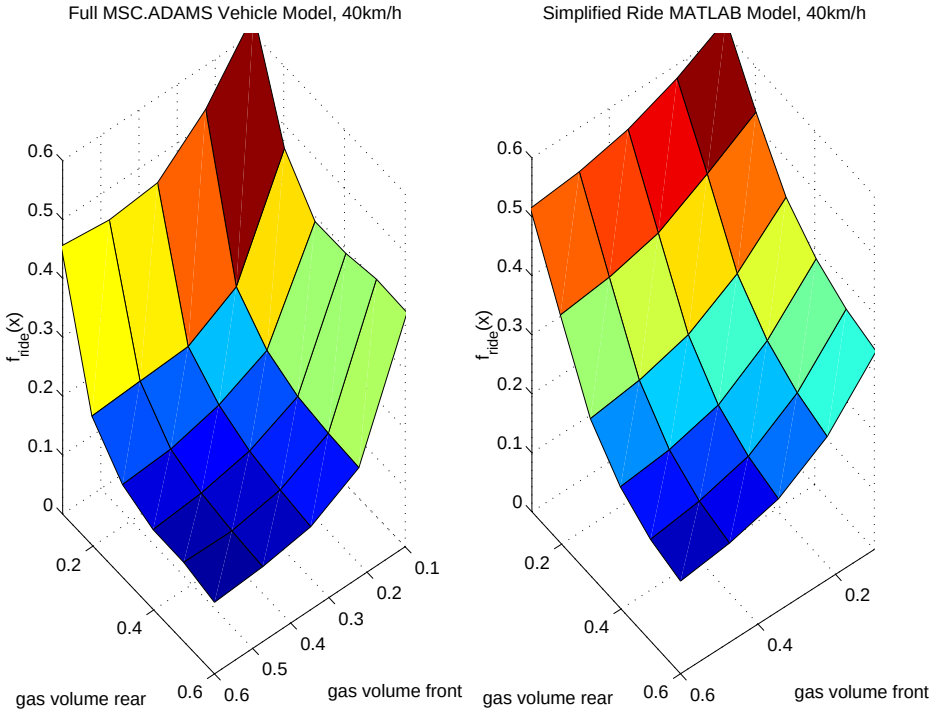

Figure 5.10: Model validation of ride comfort for differing front and rear gas volumes

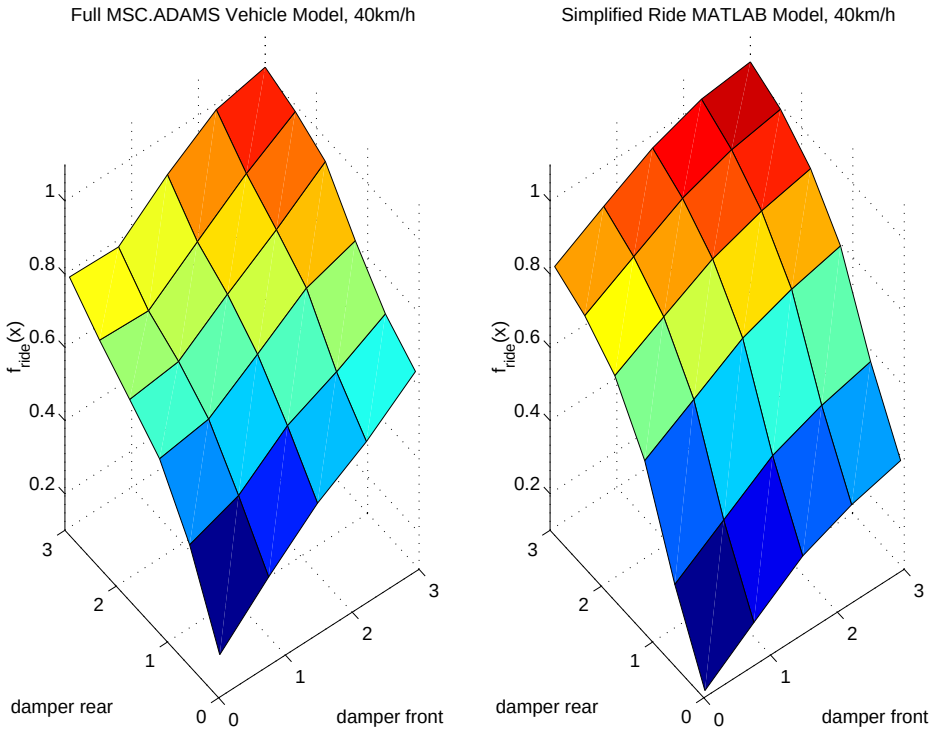

Figure 5.11: Model validation of ride comfort for differing front and rear damper scale factors

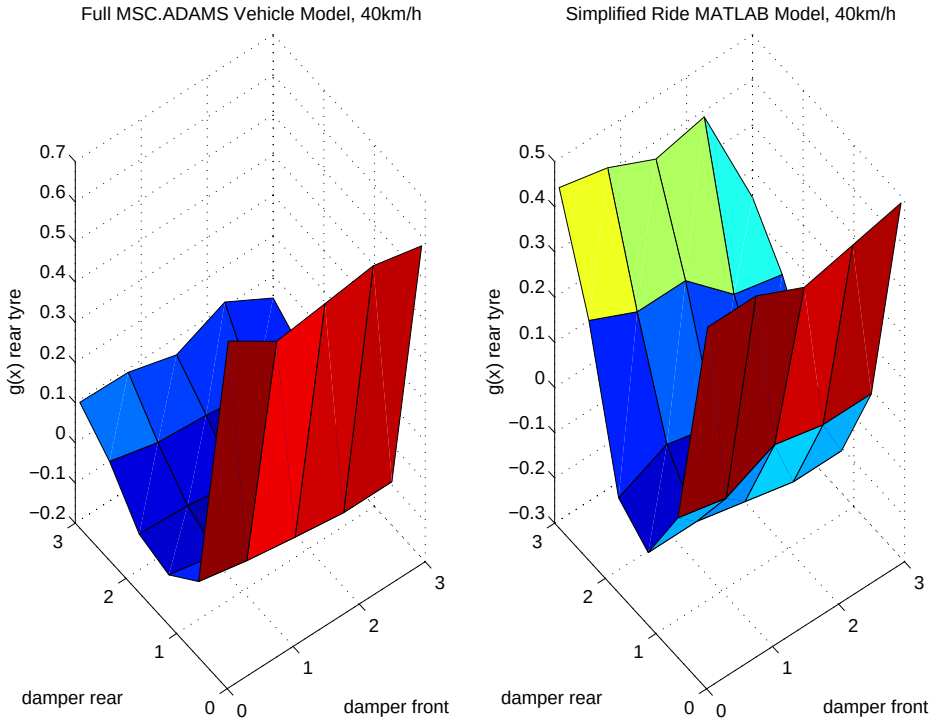

Figure 5.12: Model validation of ride comfort for differing front and rear damper scale factors: effect on rear tyre hop

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.

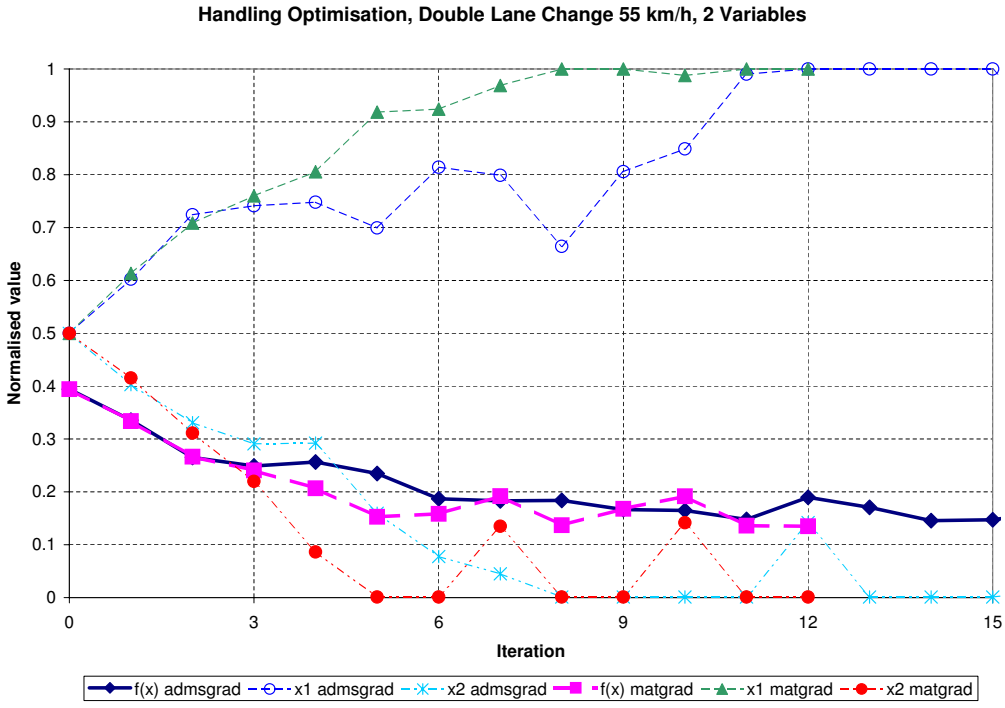

Figure 6.1: Handling optimisation convergence histories for full MSC.ADAMS model, and using the simplified MATLAB model for gradient information, 2 design variables

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).

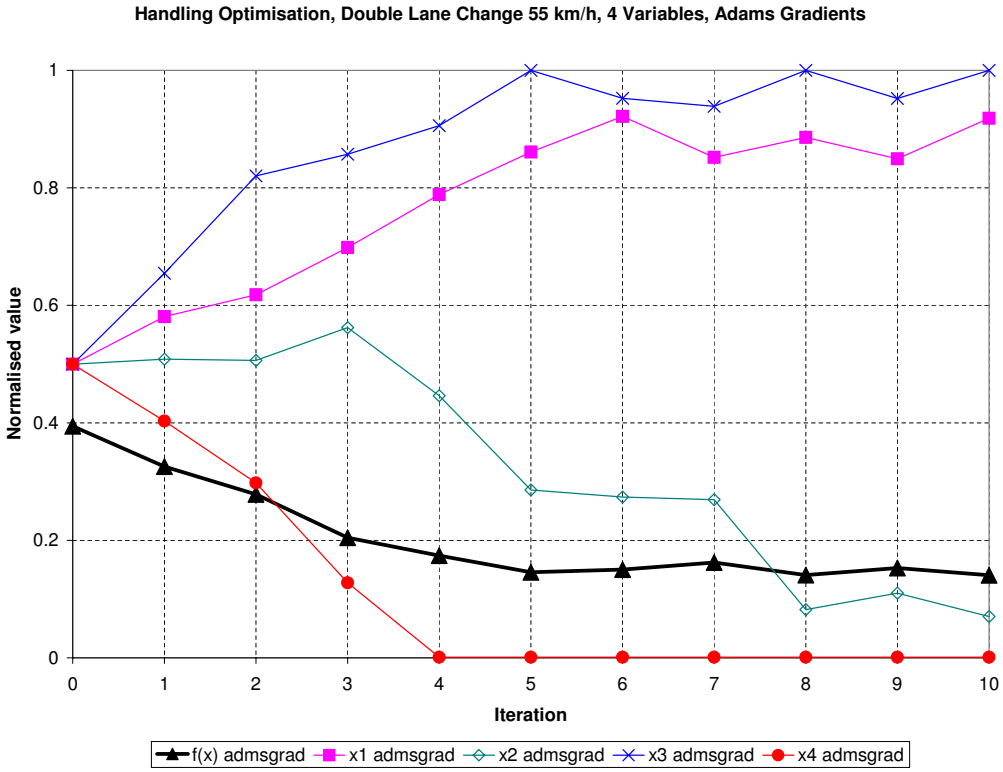

Figure 6.2: Handling optimisation convergence history using the full MSC.ADAMS model for gradient information, 4 design variables

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.

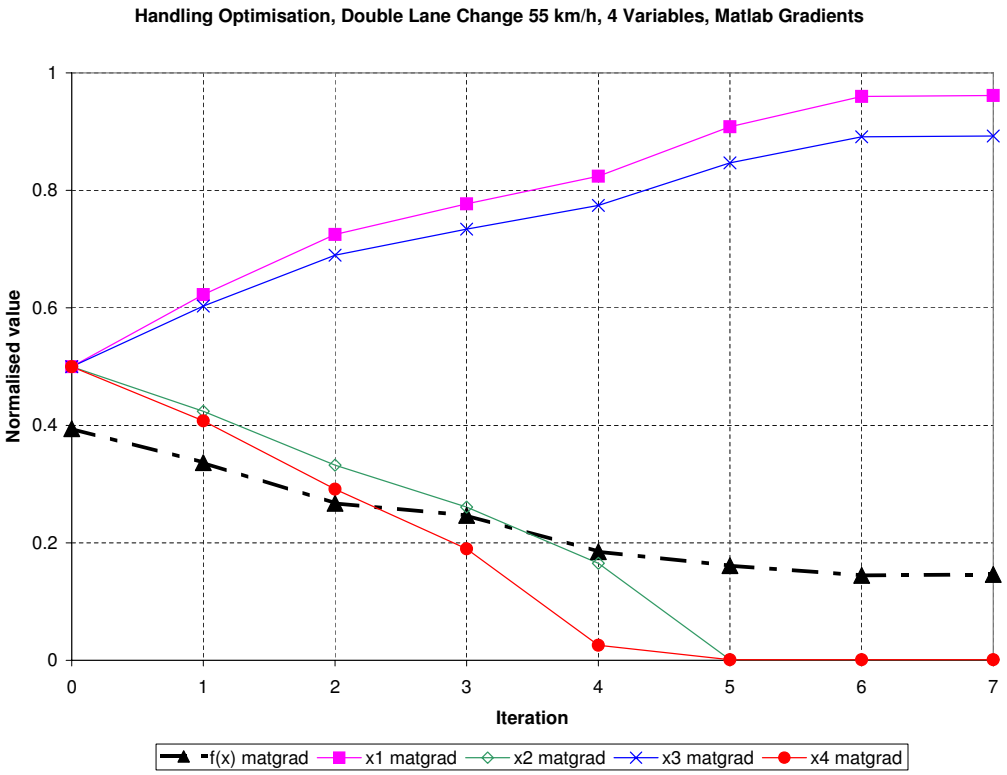

Figure 6.3: Handling optimisation convergence histories using the simplified MATLAB model for gradient information, 4 design variables

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.

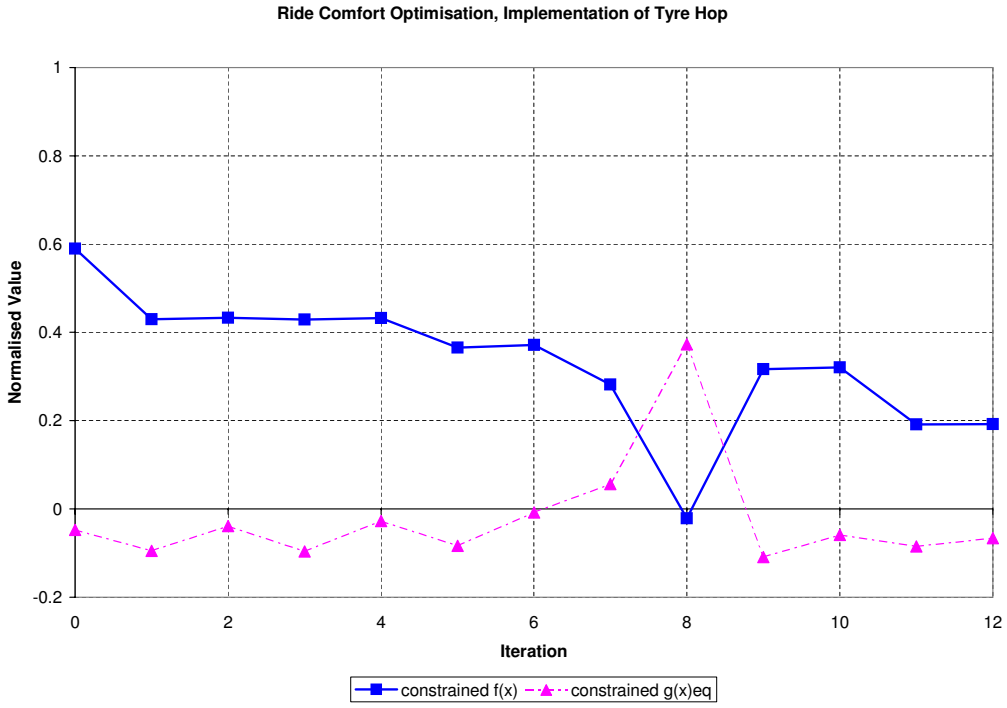

Figure 6.4: Implementing tyre hop as a constraint in ride comfort optimisation

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

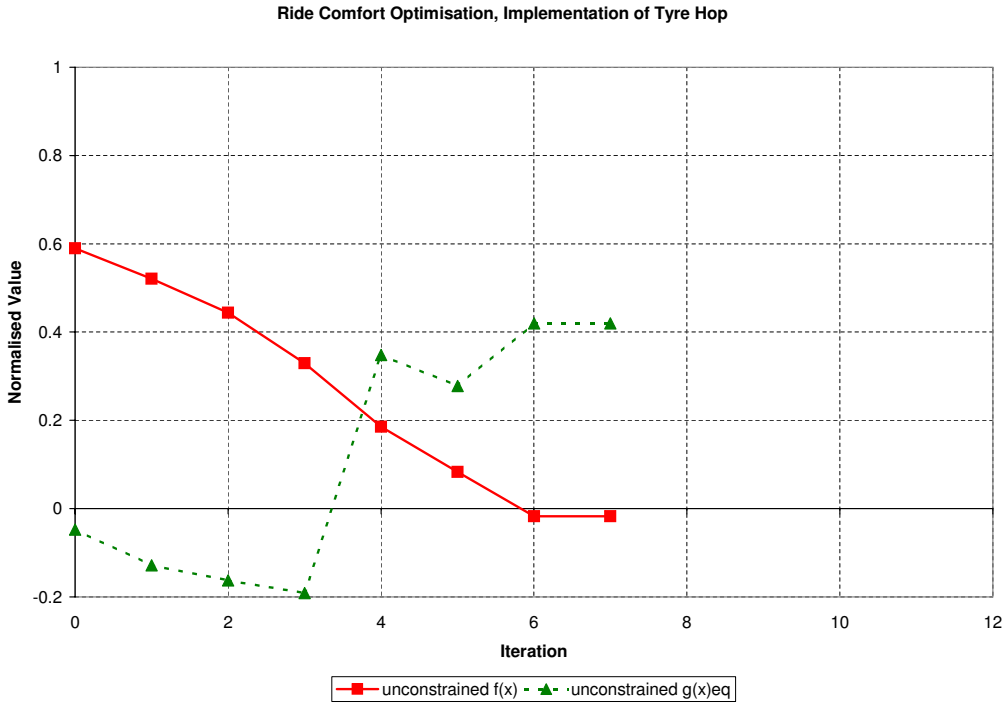

Figure 6.5: Observing tyre hop value while performing ride comfort optimisation

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.

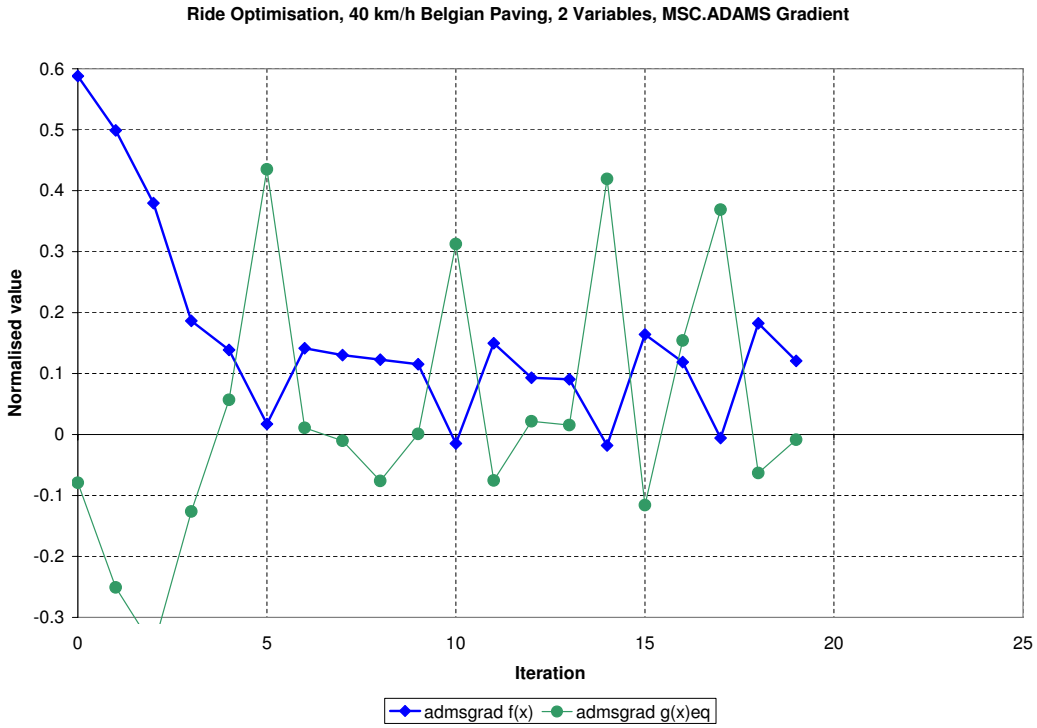

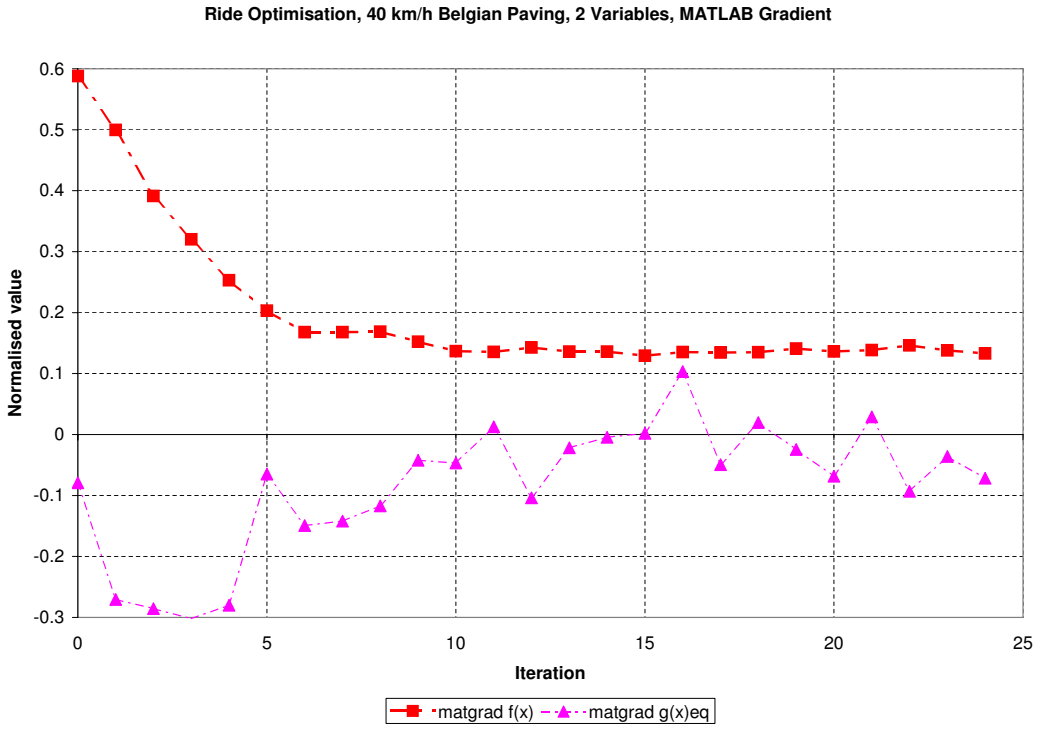

Figure 6.6: Comparison of the optimisation histories for the MSC.ADAMS gradient and simple MATLAB model gradient methods for 2 design variable ride comfort optimisation

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.

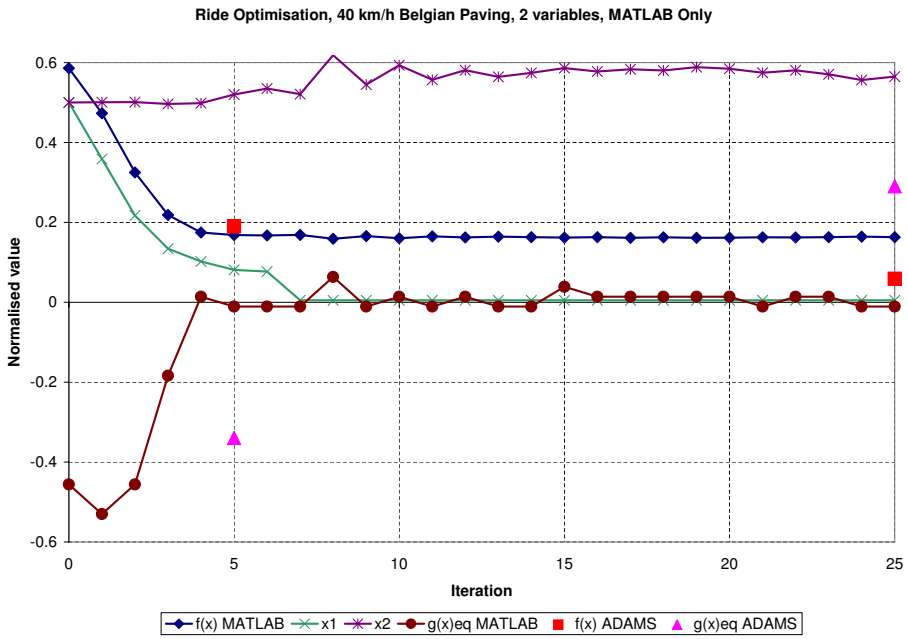

Figure 6.7: Ride comfort optimisation convergence history for using only the simple MATLAB based model, for objective function value, gradients, and tyre hop information.

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.

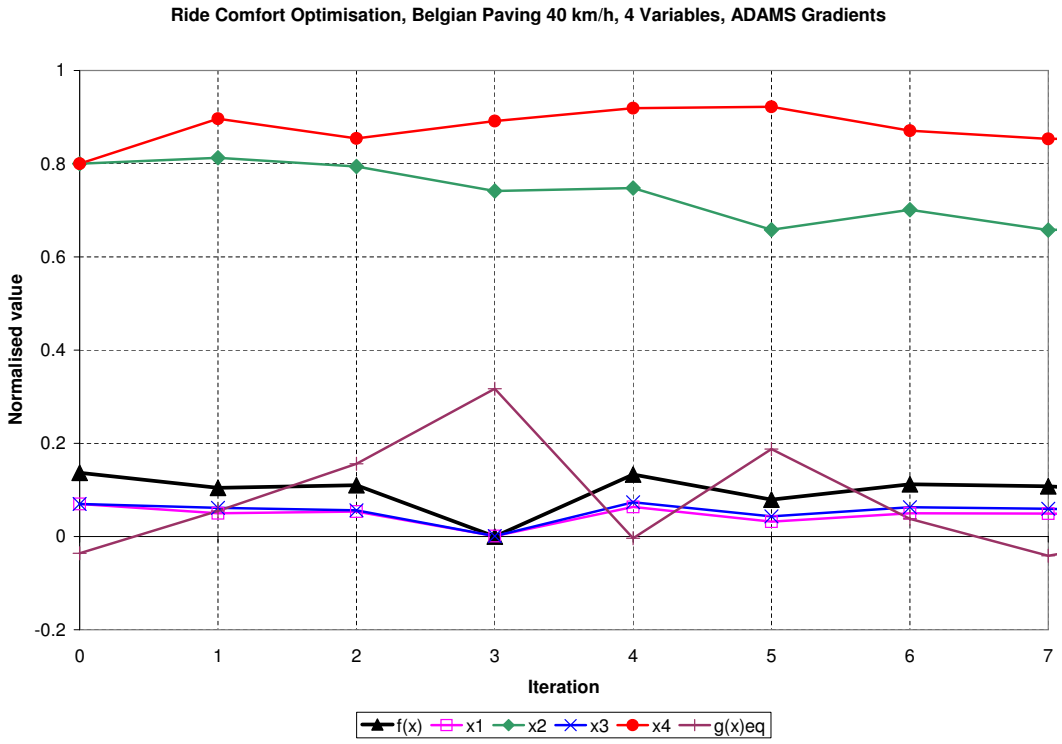

Figure 6.8: Ride Comfort optimisation convergence history for 4 design variables using the full MSC.ADAMS model for gradient information, starting at the optimum from two design variables

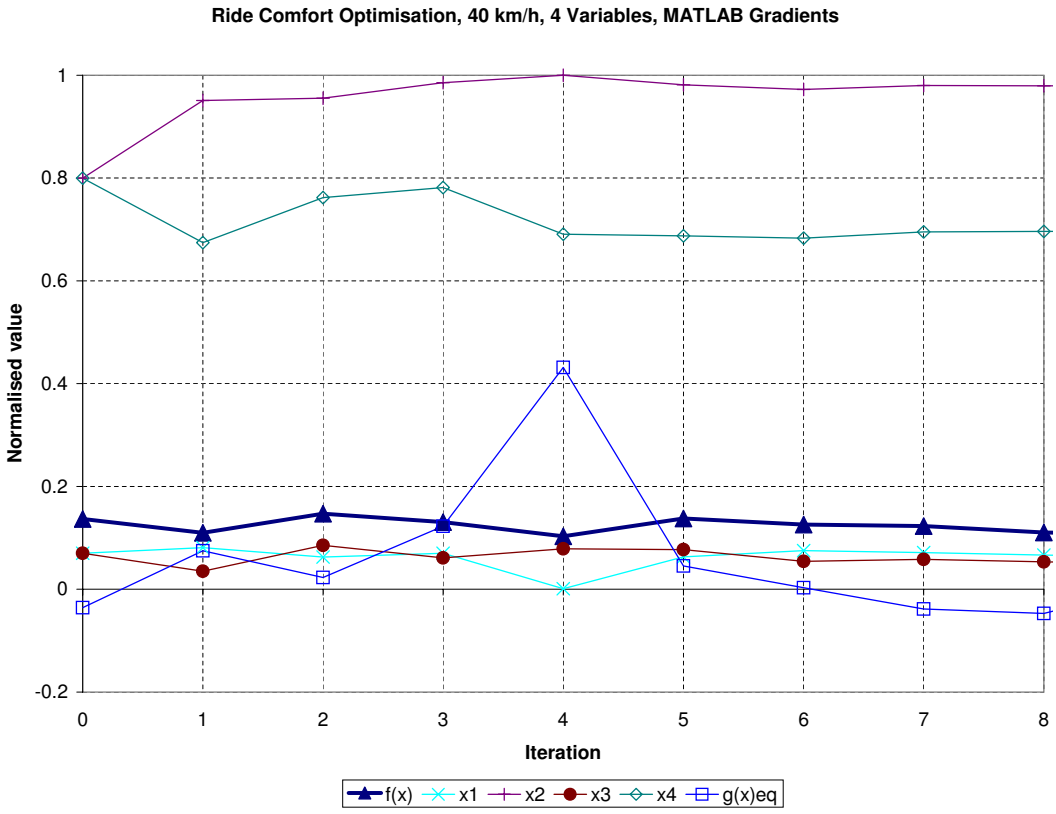

Figure 6.9: Ride Comfort optimisation convergence history for 4 design variables using the simple matlab model for gradient information, starting at the optimum from two design variables

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.

Chapter 7

Numerous Design Variables

In previous chapters, the use of simplified models and central finite differencing for the determination of gradient information, when optimising the off-road vehicle's suspension characteristics for ride comfort and handling, was shown to be beneficial. The problems considered, however, looked at only a few multiplication factors to define the suspension characteristics. In this chapter the suspension characteristics are defined by up to 14 design variables, dramatically increasing the complexity of the optimisation problem. The design variables are used to define the non-linear spring and damper characteristics, with these characteristics being optimised for the vehicle's handling and ride comfort. This chapter highlights the complications involved with the higher number of design variables. Poor scaling and sensitivity effects are illustrated in typical optimisation convergence histories, and solutions highlighted. The improved scaling discussed, dramatically helps to improve the convergence history with respect to noise. This chapter thus aims to give the optimisation engineer techniques for identifying and correcting complications associated with gradient-based vehicle suspension optimisation. It is normally these complications that lead to the adoption of less efficient stochastic based optimisation methods. While not all the complications are solved, reasons for the complications are investigated.

Based on the success of the non-linear simplified models describing vehicle handling and ride comfort, to obtain the gradient information for optimisation

problems with two and four design variables, the problem is expanded to 14 design variables. These 14 design variables better describe the shape of the damper characteristics front and rear, and the static gas volume front and rear.

The concept of using the simplified models for the calculation of the gradient information, as proposed by Balabanov and Venter (2004) for finite element structural problems, and Chapter 5 for vehicle suspension optimisation, is now assumed to be sufficiently representative of the full simulation model. This chapter will thus only define the design variables and discuss the optimisation results, with emphasis on the adjustments needed when considering many design variables. The vehicle model used is the same as in Chapter 6, except for the design variables that define the damper characteristics.

7.1 Definition of Design Variables

The front and rear static gas volumes are kept as design variables, defining the non-linear spring stiffness. The design variables that define the damper characteristics are redefined in order to achieve a more accurate description of the required damper characteristics. The standard rear damper characteristic is used, and redefined in terms of piecewise quadratic approximations, as illustrated in Figure 7.1. This gives a very accurate approximation to the measured damper characteristics. The damping force is primarily generated as a result of oil flow through an orifice, and for this reason the quadratic approximation is used to describe the characteristics. The general description of the force generated by oil flow through an orifice can be described by the quadratic relation:

$$F = kv^2 \tag{7.1}$$

where F is the damper force, v the velocity of the relative displacement of the piston (in this case of the suspension strut between the axle and the body), and k a correlation coefficient, dependent on the area and drag factor (C_d) of the orifice.

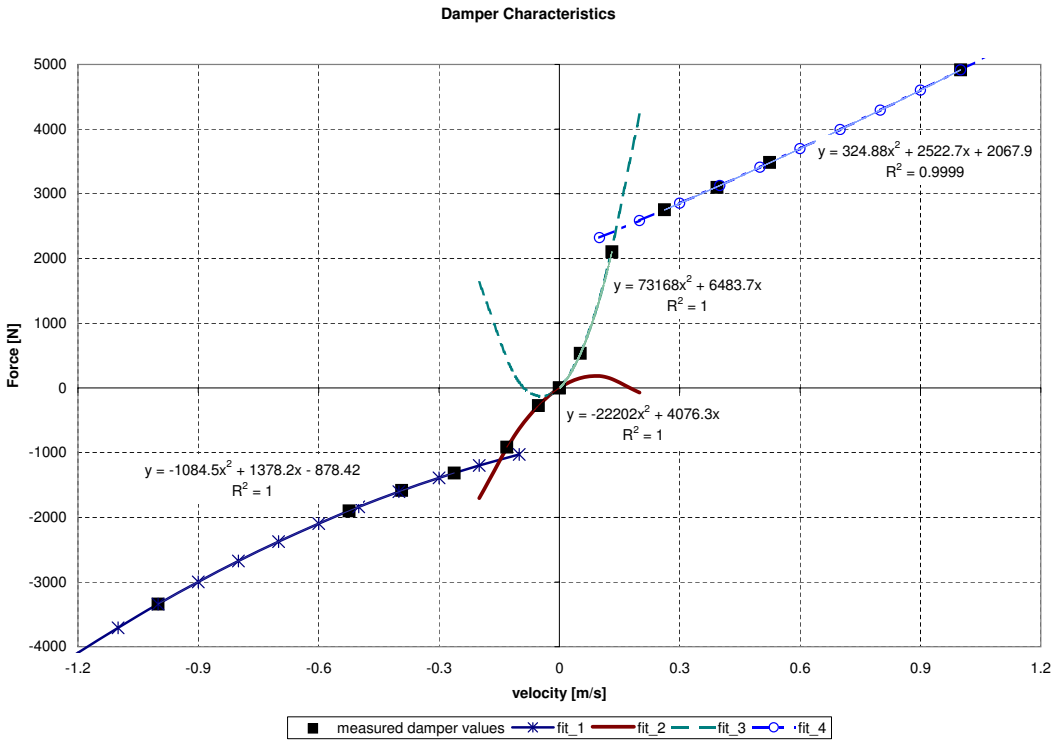

Figure 7.1: Definition of damper characteristics with quadratic approximation to the baseline rear Land-Rover damper

The damper fits with scale factors $sf_{1\rightarrow 6}$ are defined as follows:

$$\begin{aligned} fit_1 &= sf_1(-1084.5v^2 + 1378.2v) - sf_2(878.2) \\ fit_2 &= sf_3(-2220.2v^2 + 4076.3v) \\ fit_3 &= sf_4(7316.8v^2 + 6483.7v) \\ fit_4 &= sf_5(324.88v^2 + 2522.7v) + sf_6(2067.9) \end{aligned} \tag{7.2}$$

The damper force F_{dmp}, using the above piecewise fits to the measured damper characteristic, is defined as follows:

$$\begin{aligned} &if \quad v \leq 0 \\ &F_{dmp} = max(fit_1(v), fit_2(v)) \\ &else \\ &F_{dmp} = min(fit_3(v), fit_4(v)) \\ &end \end{aligned} \tag{7.3}$$

The full damper force velocity characteristic can now be defined in terms of the six scale factors $sf_{1\rightarrow 6}$. These damper scale factors are allowed to range

between 0.1 and 3. The design variables can then be stated as follows:

$$x_{1\to 6} = \frac{sff_{1\to 6}-0.1}{3-0.1}, \qquad x_7 = \frac{gvolf-0.1}{0.6-0.1}, \tag{7.4}$$
$$x_{8\to 13} = \frac{sfr_{1\to 6}-0.1}{3-0.1}, \qquad x_{14} = \frac{gvolr-0.1}{0.6-0.1},$$

with bounds:

$$0.001 \leq x_i \leq 1, \qquad i = 1, ..., 14 \tag{7.5}$$

where sff and sfr denotes the front and rear damper scale factors, $gvolf$ and $gvolr$ the front and rear static gas volumes, of the $4S_4$ suspension system. As before the static gas volumes range between 0.1 and 0.6 liter. All the design variables are then scaled to range from zero and one as suggested by Snyman (2005b).

The normalised objective and constraint functions defined in Chapter 5 are again used for the optimisation of ride comfort and handling.

7.2 Handling Optimisation

The handling optimisation was performed using the middle of the design space as a starting point, the opposite of the 4 design variable optimum (i.e. the infeasible point), and a random point in the design space. The results indicated that design variables 3, 4, 10 and 11 (the scale factors of fit_2 and fit_3, front and rear) all converged to the maximum boundary value of one, while design variables 7 and 14 (the gas volumes) converged to the minimum boundary value of almost zero. However, the other design variables did not change from their initial starting value. Yet when using different values starting values for these variables, different minima $f^*(\mathbf{x})$, less than the above result were obtained. This indicates that design variables 1, 2, 5, 6, 8, 9, 12 and 13 do have an effect on the local minimum found, yet not as strong as 3, 4, 7, 10, 11 and 14. Figure 7.2 indicates the objective function convergence history and the relative summed change in the design variables, and objective function, from one iteration to the next. It could be argued that the convergence/termination criteria are not strict enough allowing

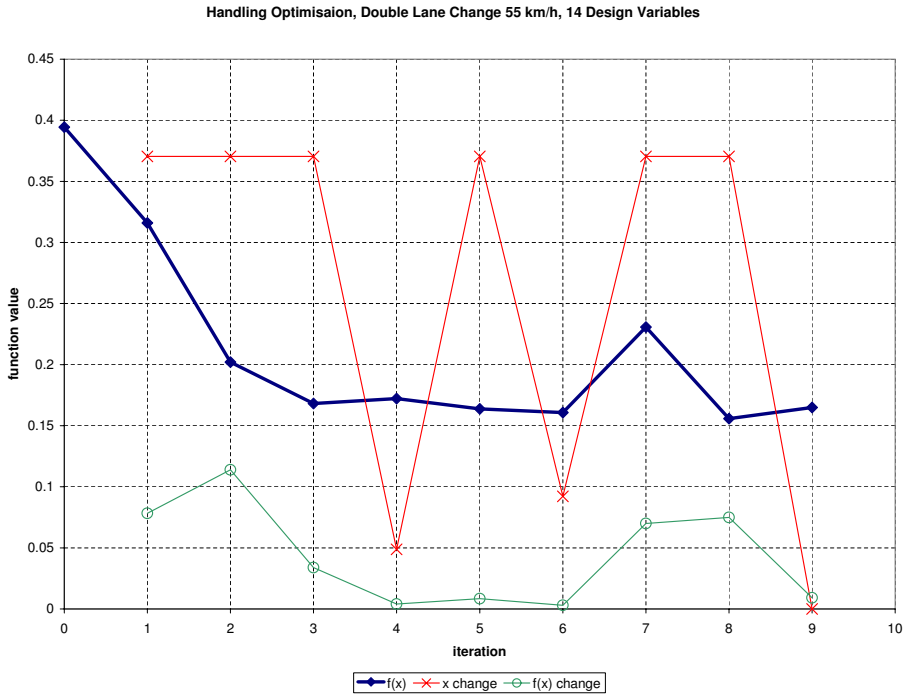

Figure 7.2: Handling 14 design variable optimisation convergence history.

premature termination of the optimisation algorithm on a non-optimum point. The termination criteria were then made 10% of previous, yet the optimisation still converged to the same points. This indicates that the design variable's current scaling flattens out their effect, resulting in an almost zero gradient, or low sensitivity.

This low sensitivity could be overcome by one of two methods. Firstly by rescaling the particular variables that they have the same magnitudes but over a much smaller range, effectively increasing their sensitivity. Alternatively by using a much larger perturbation of the design variables when calculating the gradient by central finite differences. However, Figure 7.3 indicates that the relative change in the design variable is so small, that, a change in the perturbation when calculating the gradient will not work. When observing the change in the normalised objective function values with respect to, for example design variable x_2 (Figure 7.4), it can be seen that there is a definite minimum of the objective function with respect to x_2. The objective function with respect to these design variables, or the design

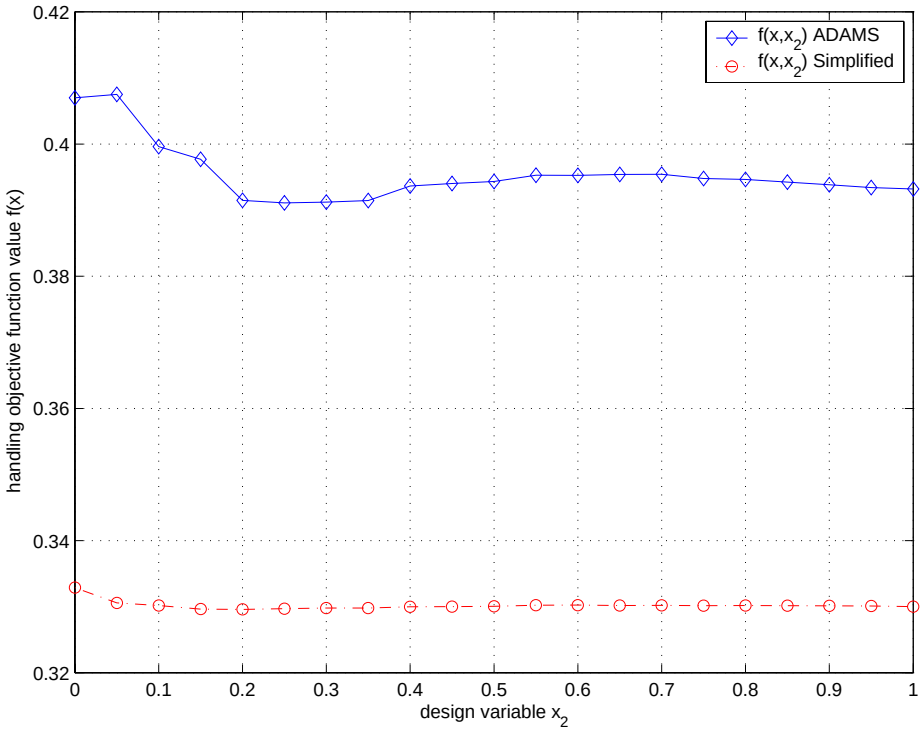

Figure 7.3: Change in objective function value with respect to design variable x_2, when the other design variables are in the middle of the design space

variables themselves has to be rescaled. This highlights that the scaling of the design variables between zero and one, as suggested by Snyman (2005b) does not necessarily guarantee good convergence to a minimum.

The design variables should thus be scaled to have almost equivalent sensitivity, without deviating too far away from similar ranges and magnitudes. This implies that it is more desirable to have an objective function of a spherical nature rather than an elliptic nature. Figure 7.5 illustrates this point, by showing the more direct and faster convergence to the two design variable optimum, when the objective function is scaled to be more spherical, as opposed to the elliptic objective function that has the design variables scaled between zero and one. In the elliptic objective function graph, it can be seen that due to slight errors in the approximate quadratic approximations, the design variables 'jump' around the optimal

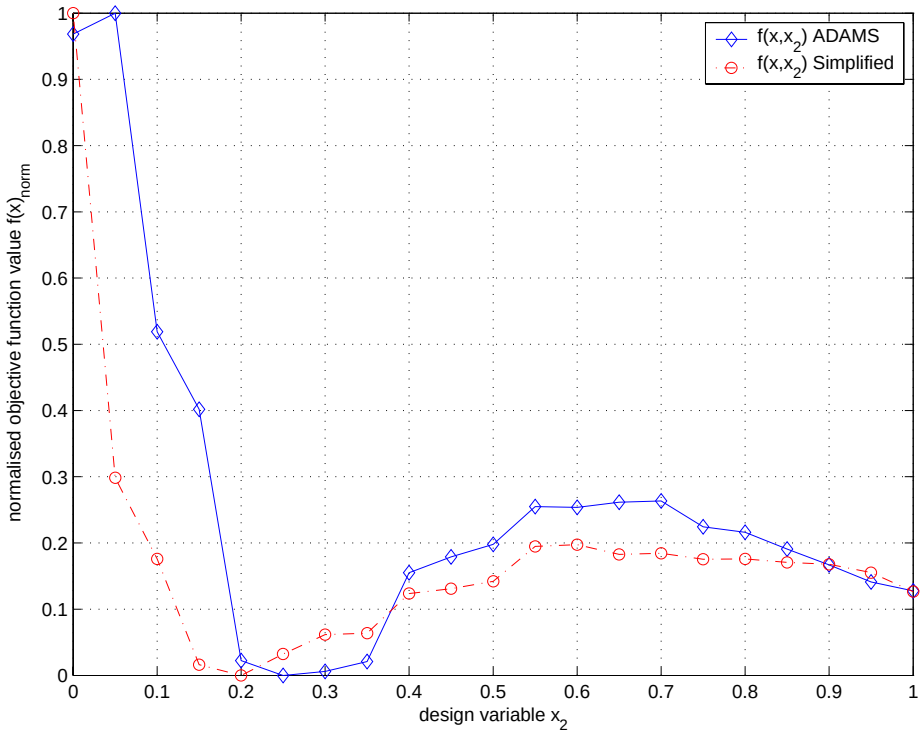

Figure 7.4: Normalised change in objective function value with respect to design variable x_2, when the other design variables are in the middle of the design space

line to the optimum. This jumping is another reason for the spikes seen in the objective function's optimisation convergence history when close to the optimum. Thus, although the objective function's history may appear erratic, the design variables are moving closer to the optimum point in the design space. This effect is amplified when the objective function is severely elliptic, with respect to the design variables, i.e. combination of a steep valley with respect to the one design variable and a shallow valley with respect to the other design variable. Also to consider is that Dynamic-Q constructs successive spherical quadratic approximations to the optimisation problem, thus if the optimisation problem exhibits a more spherical nature, the successive approximations will be a more accurate approximation, resulting in faster convergence to the actual optimum.

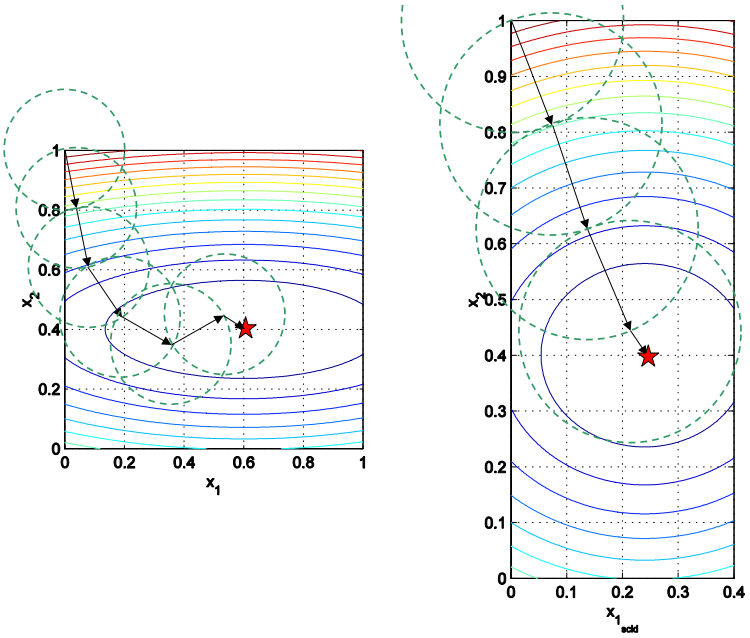

Figure 7.5: Illustration of the effect of ellipticity and sphericality on the convergence to the optimum

7.3 Ride Comfort Optimisation

The optimisation was performed with 14 design variables for ride comfort. Contrary to the handling optimisation case, it was found that design variables 1, 2, 5, 6, 8, 9, 12 and 13 (damper scale factors of fit_1 and fit_4 front and rear) had the largest effect on the erratic behaviour of the objective function value. Figure 7.6 illustrates the convergence history for the first 20 iterations, where it can be seen that changes in design variables 1 and 12, correspond to spikes in the objective function value, while design variable 14 is well behaved. Design variables 1, 2, 5, 6, 8, 9, 12 and 13 all exhibited similar trends that appeared erratic. On closer inspection it was noted that for a small change in the design variable relative to it's allowable range, there is a relatively dramatic change in the objective function value, as observed in Figure 7.6. It is thus proposed that these design variables should be rescaled so that the original normalised range of 0 to 0.4 (resulting in a scale factor range of 0.1 to 1.26) becomes their new 0 to 1 range. This effectively decreases their

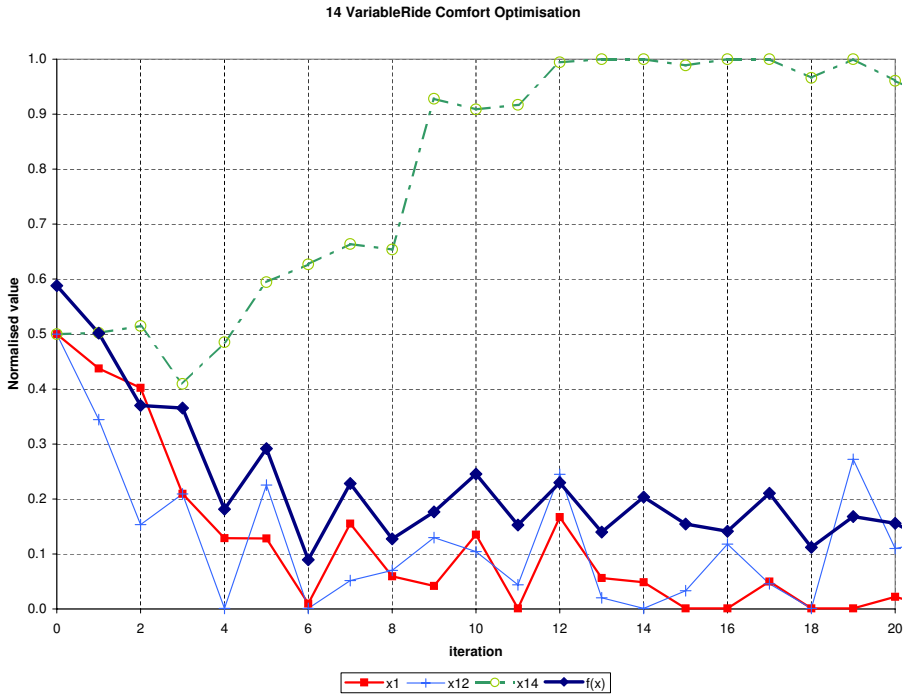

Figure 7.6: Ride comfort optimisation convergence history illustrating first 20 iterations

sensitivity, and flattens out the change in objective function with respect to a change in design variable value. This is opposite to what is needed for the handling optimisation described above. The design variables can then be stated as follows:

$$\begin{aligned} x_{1\to2} &= \tfrac{sff_{1\to2}-0.1}{1.27-0.1}, & x_{3\to4} &= \tfrac{sff_{3\to4}-0.1}{3-0.1}, \\ x_{5\to6} &= \tfrac{sff_{5\to6}-0.1}{1.27-0.1}, & x_{7} &= \tfrac{gvolf-0.1}{0.6-0.1}, \\ x_{8\to9} &= \tfrac{sfr_{1\to2}-0.1}{3-0.1}, & x_{10\to11} &= \tfrac{sff_{3\to4}-0.1}{1.27-0.1}, \\ x_{12\to13} &= \tfrac{sff_{5\to6}-0.1}{1.27-0.1}, & x_{14} &= \tfrac{gvolr-0.1}{0.6-0.1}, \end{aligned} \tag{7.6}$$

with bounds:

$$0.001 \leq x_i \leq 1, \qquad i = 1, ..., 14 \tag{7.7}$$

where sff and sfr denotes the front and rear damper scale factors, and $gvolf$ and $gvolr$ the front and rear static gas volumes of the $4S_4$ suspension system, as defined in Section 7.1. The results obtained in Figure 7.7 for the

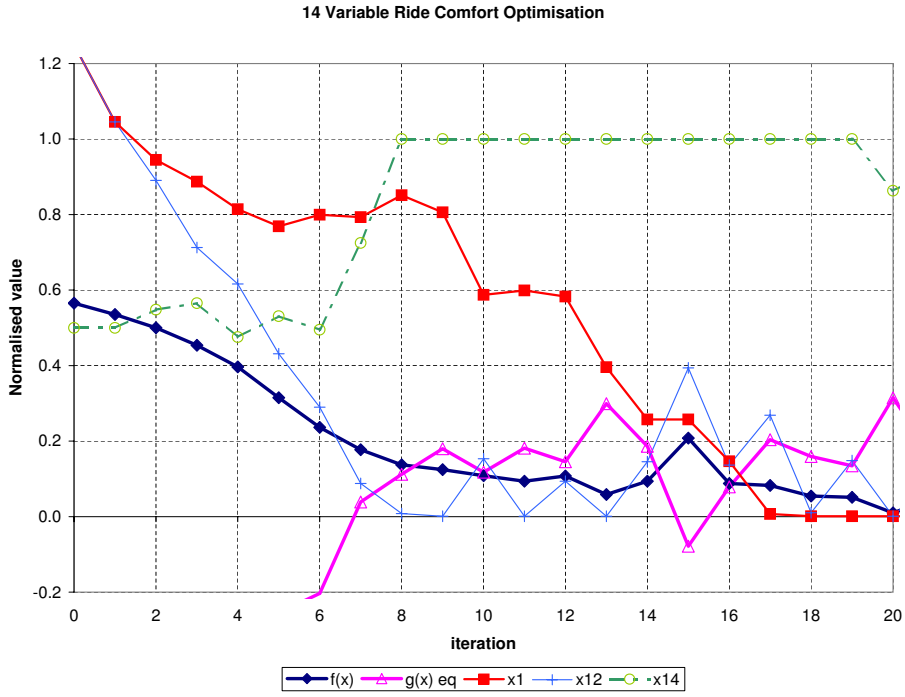

Figure 7.7: Ride comfort optimisation convergence history with rescaled design variables

rescaled problem show a dramatic improvement in the objective function's optimisation convergence history compared to Figure 7.6. However, the inequality constraint is poorly satisfied for most of the optimisation iterations. Experimentation with the penalty function parameters within the LFOPC solver of Dynamic-Q did not have sufficiently noticeable effects on the inequality constraint's convergence history. This is because the LFOPC solver finds a feasible optimum of the optimisation approximate sub-problem for every iteration, regardless of the changes in the penalty function values. This can be attributed to the smooth nature of the spherically quadratic approximate objective and constraint functions. It is postulated that the complication arises from a poor approximation of the objective and constraint functions due to the high levels of numerical noise in the full simulation model, and or gradient information.

Closer inspection of the tyre deflection at very low damping values, indicated unrealistically high levels of tyre deflection. The high tyre deflection is

attributed to the linear vertical tyre stiffness used in the ADAMS model. The tyre damping was increased to help overcome this effect with no suitable improvement at double the measured tyre damping. The logical step is to implement a non-linear vertical tyre stiffness, to overcome the problem. This was not implemented due to time constraints as the current tyre model used in the MSC.ADAMS model cannot accommodate a non-linear vertical tyre stiffness. It is thus suggested that a non-linear vertical tyre stiffness should be implemented in the tyre model used, before a decrease in the high levels of noise associated with the tyre hop inequality constraints can be achieved.

Another suggested method of overcoming the noise levels present in the objective and constraint functions as a result of the simulation model, is by re-formulating the multi-body dynamics solver's convergence criteria. The implementation of the proposed method, however, requires access to the code, and for this reason was not implemented for this research.

7.4 Conclusions

This chapter investigated the optimisation of an off-road vehicle's suspension characteristics for ride comfort and handling, where the suspension characteristics are defined by numerous design variables.

The handling optimisation highlighted the design variables that have a predominant effect on the handling performance, but also that the other design variables do contribute to the improvement of the optimum objective function achievable. However, some variables showed poor sensitivity and needed to be rescaled to improve the sphericality of the optimisation problem. This highlighted, that ensuring that the design variables vary over the same range and have equal orders of magnitude, does not necessarily guarantee good convergence to the optimum.

The ride comfort optimisation, illustrated that erratic optimisation convergence histories can be a result of over sensitive design variables in

comparison to the rest of the design variables. These over sensitive design variables were identified and rescaled, resulting in greatly improved optimisation convergence history of the objective function. However, difficulty was encountered with satisfying the tyre hop inequality constraints. This difficulty could be as a result of the tyre model's use of a linear vertical tyre stiffness, leading to unrealistically high tyre deflections in the presence of low suspension damping. Increasing the vertical tyre damping did not result in a sufficient improvement. A non-linear vertical tyre stiffness should be implemented in the model in future.

This chapter provides the optimisation engineer with some valuable methods for identifying scaling problems in the definition of the optimisation problem. While all the complications associated with noise in the optimisation process have not been addressed, feasible suggestions for future work have been proposed.

Chapter 8

Automatic Scaling of Design Variables

With the difficulties encountered in Chapter 7, it is proposed that an automatic scaling methodology be implemented within Dynamic-Q. This should limit the number of investigations and time the optimisation engineer spends on the formulation of the optimisation problem, performing optimisation runs with poor convergence, and repeating the process. This automatic scaling methodology is proposed for unconstrained optimisation problems with only design variable upper and lower bounds. This scaling aims to improve the sphericality of the optimisation problem. Figure 8.1 illustrates the typical problem with an elliptic problem, whereby the optimisation is only very sensitive to one variable, and thus takes longer to reach the optimum point. The primary assumption with this methodology is that the design variables are uncoupled. While this is an oversimplification of the optimisation problem, from the surfaces generated in Chapter 5 this assumption is not far from the physical problem.

The basic proposed methodology can be summarised as follows:

1. Scale design variables using their upper and lower limits to between zero and one.

2. Perform one function evaluation at the middle of the design space

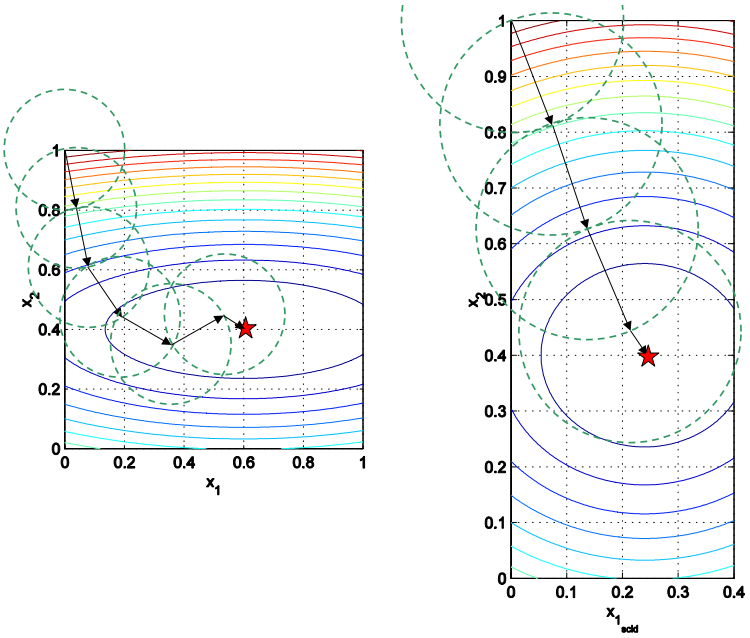

Figure 8.1: Illustration of the effect of ellipticity and sphericality on the convergence to the optimum

3. For each design variable perform 2 additional function evaluations with a perturbation of 50 % about the middle of the design space.

4. Construct a quadratic approximation of the objective function using the above simulation data, with respect to each design variable.

5. Rescale the scaled design variables so that the quadratic coefficient is equivalent to 1.

6. Perform the optimisation with the new rescaled design variables, but report the equivalent unscaled design variables at each iteration point.

This methodology should thus help to eliminate the ellipticity of the problem, if the design variables are approximately uncoupled. In the vehicle dynamics application the simplified models will be used for performing the scaling, however, under normal circumstances the cost of this scaling is equivalent to one iteration, i.e. $2n + 1$ function evaluations where n is the number of design variables, if central finite differencing is used for gradient information. The proposed automatic scaling will now be formally formulated.

8.1 Formulation of Unconstrained Automatic Scaling

Six steps need to be followed for the implementation of the automatic scaling methodology within Dynamic-Q.

Step 1: Scale the design variables x_i between their upper $\check{k}_i$ and lower $\hat{k}_i$ boundary values, so that the design variable ranges from zero to one, as follows:

$$z_i = \frac{x_i - \hat{k}_i}{\check{k}_i - \hat{k}_i} \tag{8.1}$$

where i ranges from 1 to the number of design variables n.

Step 2: Perform $2n + 1$ function evaluations to obtain function values for the construction of the quadratic approximations. This is known as central composite design (CCD). The function evaluations are performed with a 50% perturbation from the middle of the design space. The initial function evaluation $f(\mathbf{x}_{mid})$ is at the middle of the design space (mid-space). The next function evaluations are performed with the design variables at the mid-space value $\mathbf{x}_{mid}$, but a perturbation in only the i^{th} design variable. Thus the function evaluations can be defined as:

$$\hat{f}_i = f(\mathbf{x}_{mid}, \hat{x}_i) \check{f}_i = f(\mathbf{x}_{mid}, \check{x}_i) \tag{8.2}$$

where $\hat{x}_i$ corresponds to the equivalent $\hat{z}_i$ which is defined as:

$$\hat{z}_i = z_{0_i} - 0.5 \tag{8.3}$$

and $\check{x}_i$ corresponds to the equivalent $\check{z}_i$ which is defined as:

$$\check{z}_i = z_{0_i} + 0.5 \tag{8.4}$$

Thus the objective function is evaluated with a 50% perturbation on either side of the mid-space point $\mathbf{x}_{mid}$, resulting in the whole design space being approximated. All the optimisation up until now has been done with a 3% perturbation on either side of the current iteration point for the evaluation

of the central finite difference gradient. It is decided that by considering the whole design range a suitable evaluation of the effective curvature of the optimisation problem is obtained, without being dramatically affected by numerical noise. This then ensures that the sensitivity of the design variables over the whole design space is taken into account.

Step 3: Use the above determined objective function values to construct approximate quadratic approximations of the objective function with respect to the design variable as follows:

$$\tilde{f}(z_i) = a_i {z_i}^2 + b_i z_i + c_i \tag{8.5}$$

Step 4: Rescale the design variables $\mathbf{z}$ so that the corresponding a_i term of the approximated objective function will be 1. The rescaled design variable will thus be defined as:

$$X_i = z_i t_i \tag{8.6}$$

where the design variable scale factor t_i is defined as:

$$t_i = \sqrt{a_i} \tag{8.7}$$

This scaling, however, has the problem that it tends to zero when a_i tends to zero, and X_i tends to infinity when a_i tends to infinity. It is thus proposed that should a_i tend to zero, then the quadratic approximation tends to a straight line, and that this straight line should have a gradient of + or - 1, this means that the b_i term should be used for the scaling. If, on the other hand, a_i becomes very large the scaling will result in a very large design space with respect to that design variable, thus it is proposed that the upper limit of the design space/variable range should be 20, corresponding to an a_i value of 400. The lower limit to the design space is chosen as 0.2, corresponding

to an a_i value of 0.04. The following if loop then applies:

$$\begin{aligned} &if \quad a_i \leq 0.04 \\ &t_i = |b_i| \\ &elseif \quad a_i \geq 400 \\ &t_i = \sqrt{400} = 20 \\ &end \end{aligned} \tag{8.8}$$

Again the problem of tending to infinity when $|b_i|$ tends to infinity, exists. Should $|b_i|$ tend to zero, then t_i also tends to zero. Thus the additional if statement must be inserted:

$$\begin{aligned} &if \quad |b_i| \leq 0.2 \\ &t_i = 0.2 \\ &elseif \quad |b_i| \geq 20 \\ &t_i = 20 \\ &end \end{aligned} \tag{8.9}$$

The rescaled design variables X_i will thus be limited to the following:

$$0.2z_i \leq X_i \leq 20z_i \tag{8.10}$$

Step 5: Change the move limit so that it is still representative for the rescaled problem. The new move limit DM_n is a function of the number of design variables n, the original move limit dml, and the scale factors t_i, and is defined as follows:

$$DM_n = dml\sqrt{\frac{\sum_{i=1}^{n} {t_i}^2}{n}} \tag{8.11}$$

Step 6: Perform the optimisation with the rescaled design variables $\mathbf{X}$, and new move limit DM_n, but report the actual design variable values to the user. The Dynamic-Q design variables $\mathbf{X}$ will thus be converted for printout to the users design variables $\mathbf{x}$ as follows:

$$x_i = \frac{X_i}{t_i}(\check{k}_i - \hat{k}_i) + \hat{k}_i \tag{8.12}$$

8.2 Concept Test

This automatic scaling methodology is first tested on simple severely elliptic, analytical design problems. The first test is using a two design variable analytic problem described by the objective function:

$$f(\mathbf{x}) = {x_1}^2 + 10{x_2}^2 \tag{8.13}$$

Were the objective function $f(\mathbf{x})$ is uncoupled with respect to x_1 and x_2. Upper and lower bounds on the design variables are defined as:

$$-1 \leq x_{1,2} \leq 1 \tag{8.14}$$

and an initial starting point of [1 1]. The gradient information was determined analytically, and the performance of the standard form of Dynamic-Q (Snyman and Hay 2002) was compared to the automatic scaling version of Dynamic-Q to be known as Ascl-Dyn-Q. Figure 8.2 illustrates the comparison of the convergence histories in the design space for the standard form of Dynamic-Q and Ascl-Dyn-Q. It is observed that Ascl-Dyn-Q moves much faster towards the optimum. The function error is determined as defined by Snyman and Hay (2002) as:

$$f(err) = \frac{\|f_{act} - f^*\|}{1 + \|f_{act}\|} \tag{8.15}$$

For the above optimisation problem the results are tabulated in Table 8.1, line *cp* 1.

The second problem *cp* 2 is a skew problem, where the dependence of $f(\mathbf{x})$ with respect to x_1 is coupled to x_2, described by the objective function:

$$f(\mathbf{x}) = {x_1}^2 + 10{x_2}^2 + 3x_1x_2 \tag{8.16}$$

With upper and lower bounds on the design variables defined as:

$$-1 \leq x_{1,2} \leq 1 \tag{8.17}$$

and an initial starting point of [1 1]. It can be seen from Figure 8.3 that the Ascl-Dyn-Q algorithm moves to the optimum in the same manner as the

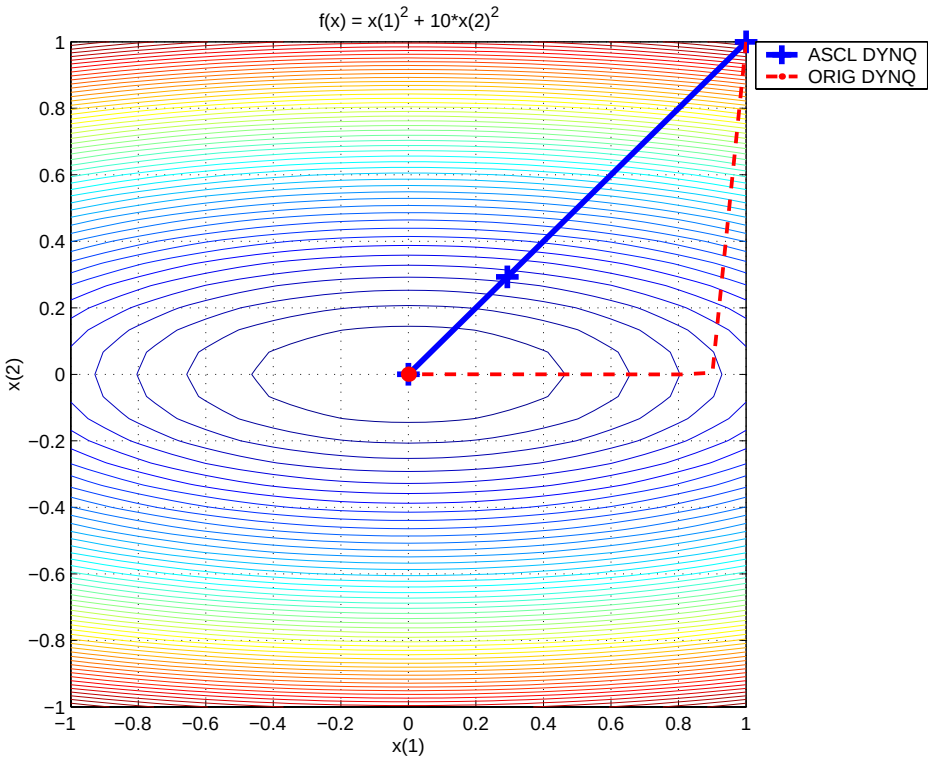

Figure 8.2: Comparison of standard Dynamic-Q convergence to optimum and Dynamic-Q with automatic scaling, for test 1

standard algorithm, but approaches the optimum from the other side. This is the limit of the permissible cross-coupling of design variables, where the algorithms exhibit almost equal performance, *cp* 2 in Table 8.1.

8.3 Modification for Constrained Problems

With the success achieved with the proposed automatic scaling procedure, the methodology was expanded to include constrained optimisation problems. It is proposed that, because Dynamic-Q makes use of LFOPC for the optimisation of the approximate sub-problem at each iteration, the constraints could be included by using the penalty function approach, as in LFOPC. The resulting penalty function should thus be made spherical, as opposed to just the underlying objective function. LFOPC solves the penalty function in a three part approach (Snyman 2000), thus the question must be asked as to what penalty parameter multiplication factor should be used. LFOPC first solves the approximate sub-problem using a low penalty function multiplication

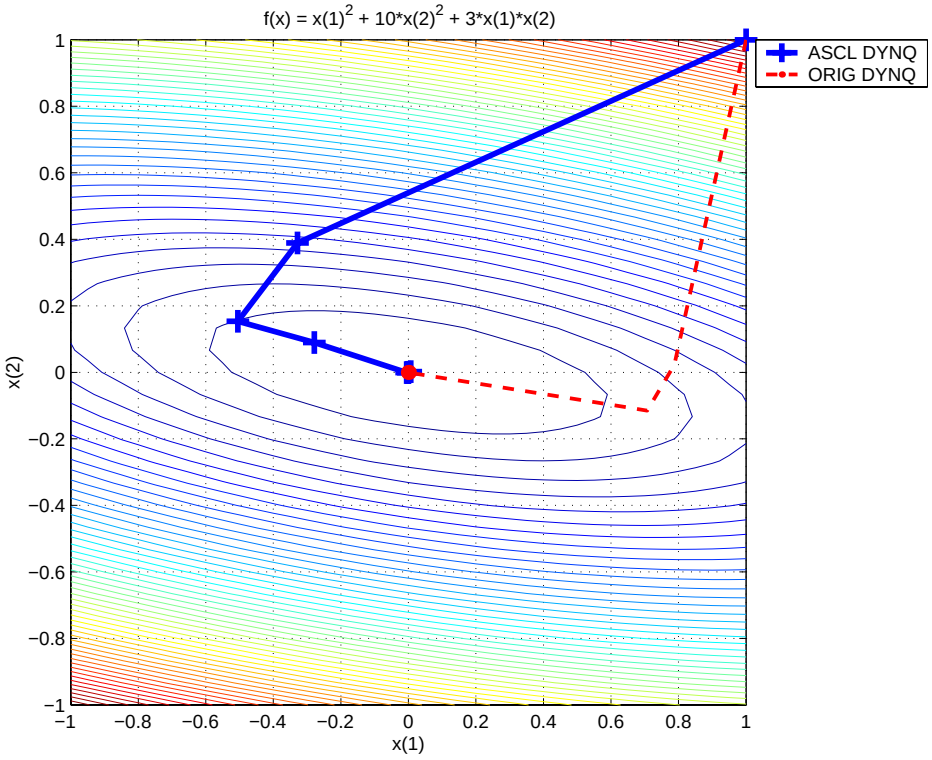

Figure 8.3: Comparison of standard Dynamic-Q convergence to optimum with Dynamic-Q with automatic scaling, for test problem 2

factor, and then increases the penalty function multiplication factor in the next phases. It is thus proposed that the violated constraint should be added to the objective function value using the lowest multiplication factor. The quadratic approximations with respect to each design variable are then fitted to the resulting penalty function. Where the penalty function is defined as follows:

$$\underset{w.r.t.x}{minimize} \quad P(\mathbf{x},\mu) \tag{8.18}$$

where

$$P(\mathbf{x},\mu) = f(\mathbf{x}) + \sum_{j=1}^{m} \mu_j g_j(\mathbf{x}) + \sum_{j=1}^{r} \mu_j h_j(\mathbf{x}) \tag{8.19}$$

where the penalty multiplier μ_j is defined by the if statement for $g_j(\mathbf{x})$ as:

$$\begin{gathered} if \quad g_j(\mathbf{x}) \leq 0 \\ \mu_j = 0 \\ else \quad \mu_j >> 0 \\ end \end{gathered} \tag{8.20}$$

Table 8.1: Results for the standard Dynamic-Q and Auto-Scaling Dynamic-Q methods

Problem #	n	$f(act)$	Dynamic-Q			Ascl Dyn-Q		
			# iter	f^*	$f(err)$	# iter	f^*	$f(err)$
cp 1	2	0.00e+00	7	1.53e-12	1.53e-12	3	8.99e-11	8.99e-11
cp 2	2	0.00e+00	7	1.70e-12	1.70e-12	8	4.24e-11	4.24e-11
Hock 2	2	5.04e-02	7*	4.94e-00	error	11*	4.94e-00	error
Hock 13	2	1.00e+00	6	9.99e-01	1.00e-08	6 nc	1.00e+00	3.00e-07
Hock 15	2	3.07e+02	17	2.13e+02	4.35e-01	14	error	error
Hock 17	2	1	16	1	< 1.00e-08	10 nc	1	< 1.00e-08

* - converged to local minimum

nc - no constraints considered with scaling

and the if statement for $h_j(\mathbf{x})$ as:

$$\begin{aligned} &if \quad h_j(\mathbf{x}) = 0 \\ &\mu_j = 0 \\ &else \quad \mu_j >> 0 \\ &end \end{aligned} \tag{8.21}$$

Some random test problems, of the ones on which Dynamic-Q was tested and presented in Snyman and Hay (2002) are resolved using the auto-scaling methodology. The test problems are from the book of Hock and Schittkowski (1981), and given in Appendix A for the readers convenience.

From the results it can be seen that the auto-scaling improves the optimisation convergence to the optimum, for most test cases, except badly skew elliptic problems like the Rosenbrock problem. From the results the feasibility of using the penalty function for constrained optimisation problems has not been shown, as there are difficulties as to what value to use for the penalty function multiplication factor.

With the success of the automatic scaling a search was done for similar

novel approaches to scaling of the design variables. Most researches suggest the normalisation of the design variables to range between zero and one, as in Snyman (2005b) and Lasdon (2001). Willcox (2006) suggests inspecting the Hessian matrix at the converged optimum design point. The condition number of the Hessian matrix is evaluated. If the condition number is greatly larger than one, the matrix is ill-conditioned, and the design variables are transformed linearly to minimise the condition number of the solution. This is, however, only performed after the optimisation algorithm has converged to a solution, while the scaling proposed in this thesis is done over the whole design space at the begining of the optimisation process. This ensures that the global problem is scaled to be more spherical. The other advantage of the automatic scaling suggested here, is that the Hessian matrix does not need to be constructed, greatly reducing gradient evaluations, that are normally very costly in typical engineering problems.

If the Hessian needs to be calculated in order to better scale the design variables, Danchick (2006) suggests an efficient and accurate method for computation of the Hessian matrix. Danchick makes use of central finite difference quotients and extrapolation-to-the-limit to achieve a h^4 level of accuracy, where h is the finite difference step size. The computational cost is $2n(n+1)+1$ function evaluations per Hessian matrix evaluation, where n is the number of design variables. This is then implemented in an optimisation algorithm that makes use of Hessian decomposition and eigenvalue shifting to follow a ridge in a difficult skew optimisation problem like Rosenbrock's parabolic valley. The determination of this Hessian matrix should be considered over the whole design space before the optimisation process should be considered in the future, but will probably not be necessary for most well defined engineering problems.

8.4 Implementation in the Vehicle Suspension Problem

The automatic scaling method proposed above was implemented for the 14 design variable optimisation, and the results were compared to the those obtained with Dynamic-Q without automatic scaling. Because of the apparent hopping about an optimum point, the convergence histories presented are in the form of the best feasible point at the current iteration point. This approach is borrowed from the genetic algorithm and particle swarm community. If an improved solution is not achieved after a certain number of iterations, the optimisation is terminated. The number of iterations before termination is, however, difficult to select, as this may result in premature termination. It was also determined that the optimisation should not be permitted to terminate within the first 10 iterations.

Presented in Figure 8.4, is the comparison of the optimisation convergence histories for optimisation with automatic scaling (ascl) and without (std). It is observed that the automatic scaling terminates at a better optimum than without automatic scaling, and the design variables do not get stuck in local minima as for the standard optimisation. This local minimum is used to start the optimisation using automatic scaling in order to achieve a better optimum, the results are presented in Figure 8.5. It can thus be concluded that the automatic scaling was successful for the optimisation of 14 design variables for handling.

For ride comfort, the decision of when to terminate the optimisation will impact on the performance of the optimisation methods. Presented in Figure 8.6 is the optimisation convergence histories for the standard form of Dynamic-Q (std) and with automatic scaling implemented (ascl). The use of the penalty function for the automatic scaling as discussed in section 8.3, was not used here, due to difficulties associated with the correct magnitude of the penalty multiplier. It can be seen that the automatic scaling reaches a better

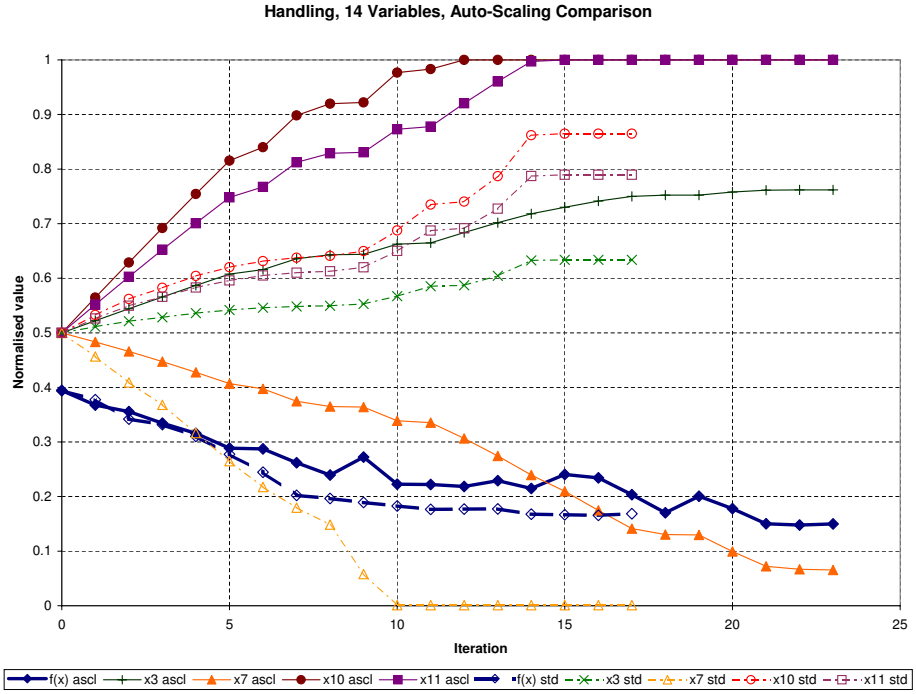

Figure 8.4: Comparison of convergence histories for standard Dynamic-Q and for the implementation of the automatic scaling, for 14 design variable handling

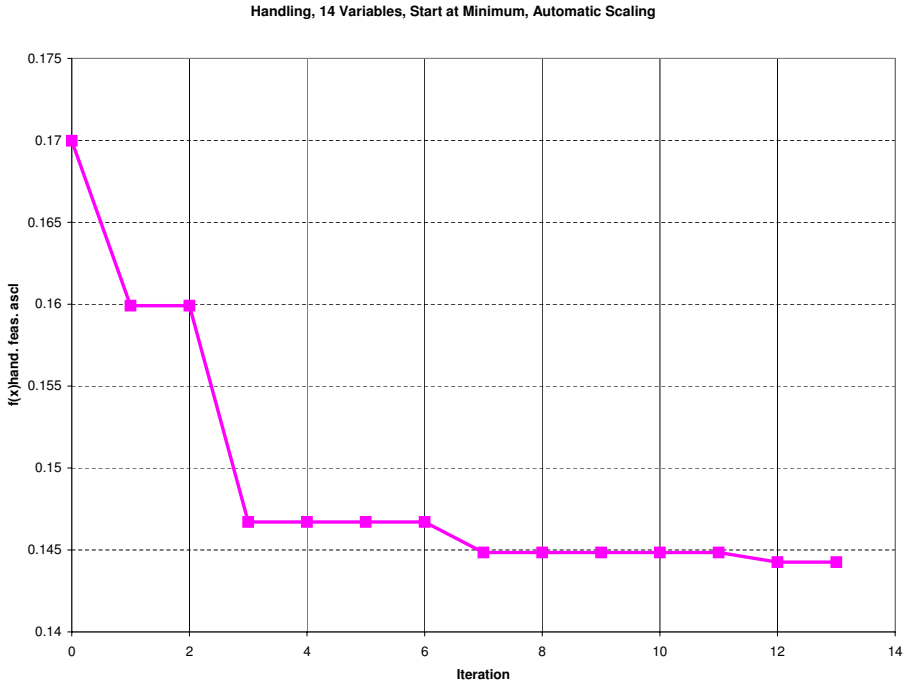

Figure 8.5: Optimisation convergence history for handling, starting at the found minimum

minimum than the standard form, however, this is not the optimum point, as better optima were reached using only 4 design variables (see Chapter 6, Figure 6.9). The minimum point achieved in Figure 8.6, for the automatic scaling (f(x) feas. ascl, iteration 18) was then used as a starting point for the optimisation of the ride comfort, with the design variables subjected to a 30% range about the minimum found in Figure 8.6. The optimisation convergence history is presented in Figure 8.7. It is observed that an equal minimum is reached as for the 4 design variables.

From the optimisation results the optimum damper characteristics for handling (Figure 8.8) and ride comfort (Figure 8.9), are presented. It is also found that for optimum handling the static gas volume must be 0.1 liter, while for optimal ride comfort the static gas volume in front should be 0.39 liter, and at the rear 0.46 liter. The handling gas volume thus ran to the lower boundary, but the ride comfort static gas volume did not. This can be attributed to the tyre hop inequality constraints. The optimal driver vertical RMS acceleration is 1.1 m/s^2 and the rear passenger vertical RMS acceleration is 1.1 m/s^2. The optimal body roll velocity RMS value is 0.52 $^o/s$ and the maximum roll angle is 3.1 ^{o}.

8.5 Conclusions

The automatic scaling methodology was proposed, and implemented on several analytic problems with success. Automatic scaling was then applied to the vehicle dynamics problem of numerous design variables.

The automatic scaling methodology was implemented with success. The optimal handling and ride comfort were determined, where it was found that the handling setting would require a small gas volume, and stiff dampers front and rear, while the ride comfort required soft front and even softer rear damping, and a static gas volume of 0.39 liter in front and 0.46 liter at the rear.

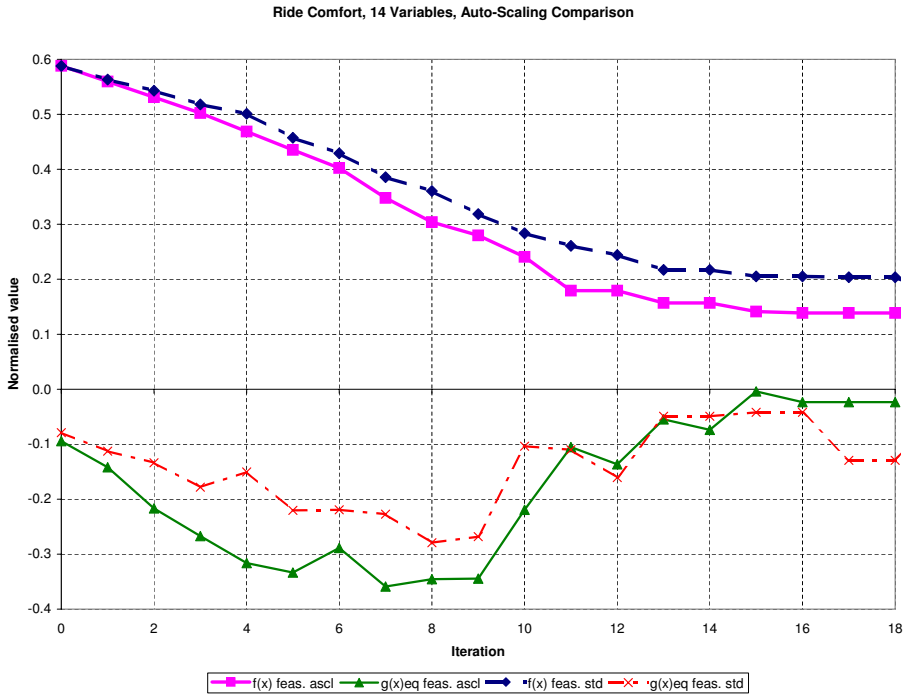

Figure 8.6: Comparison of optimisation convergence histories for standard Dynamic-Q and for the implementation of the automatic scaling (14 design variable ride comfort)

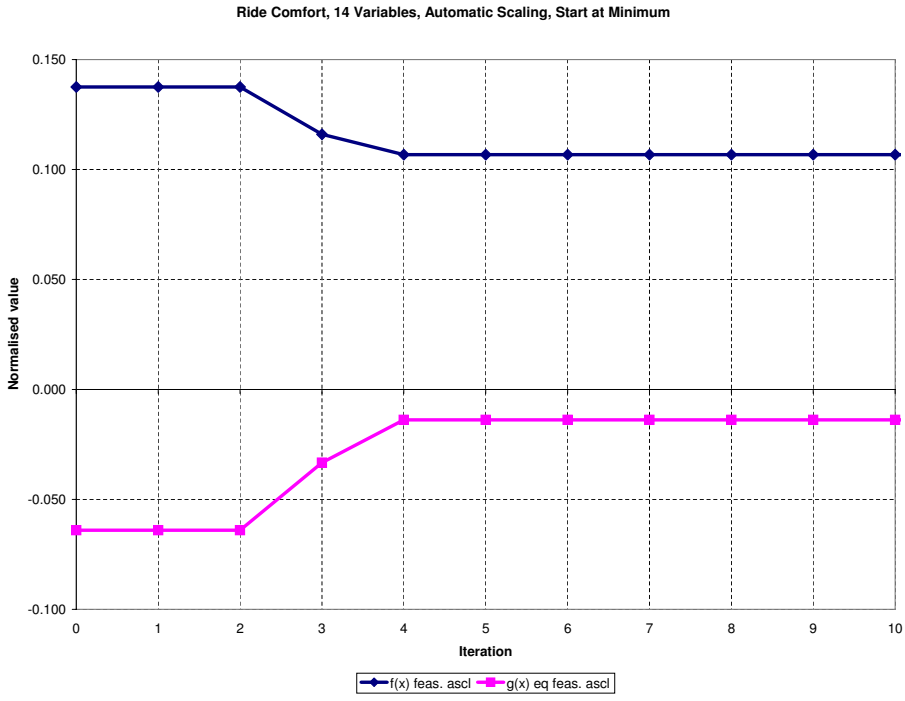

Figure 8.7: Optimisation convergence history for starting at minimum, using automatic scaling

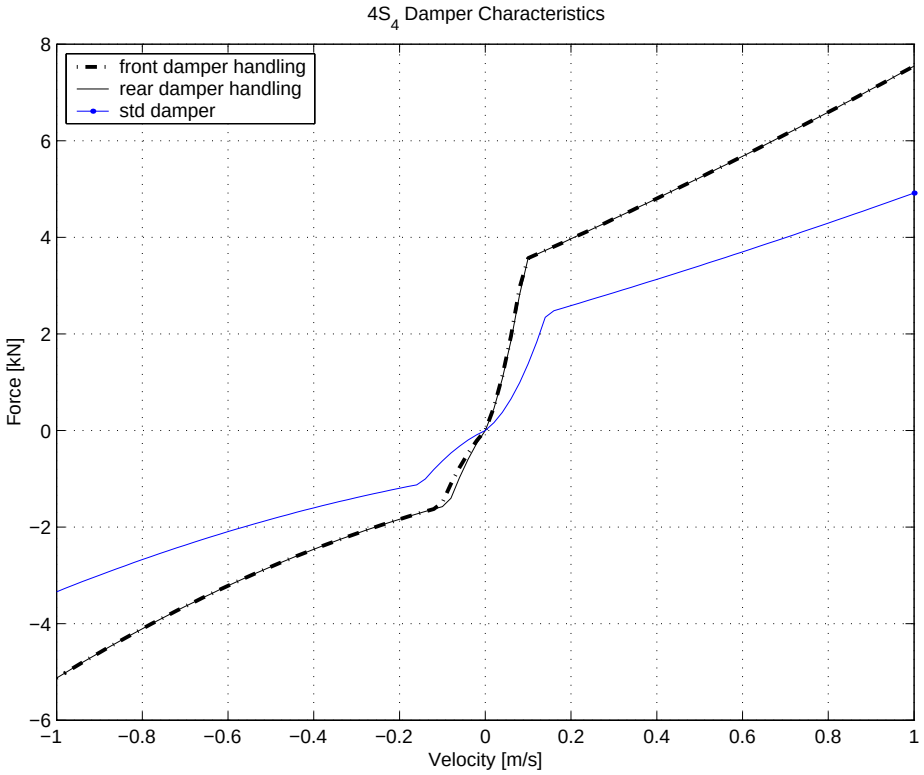

Figure 8.8: Optimum damper characteristics for handling compared to the baseline rear damper

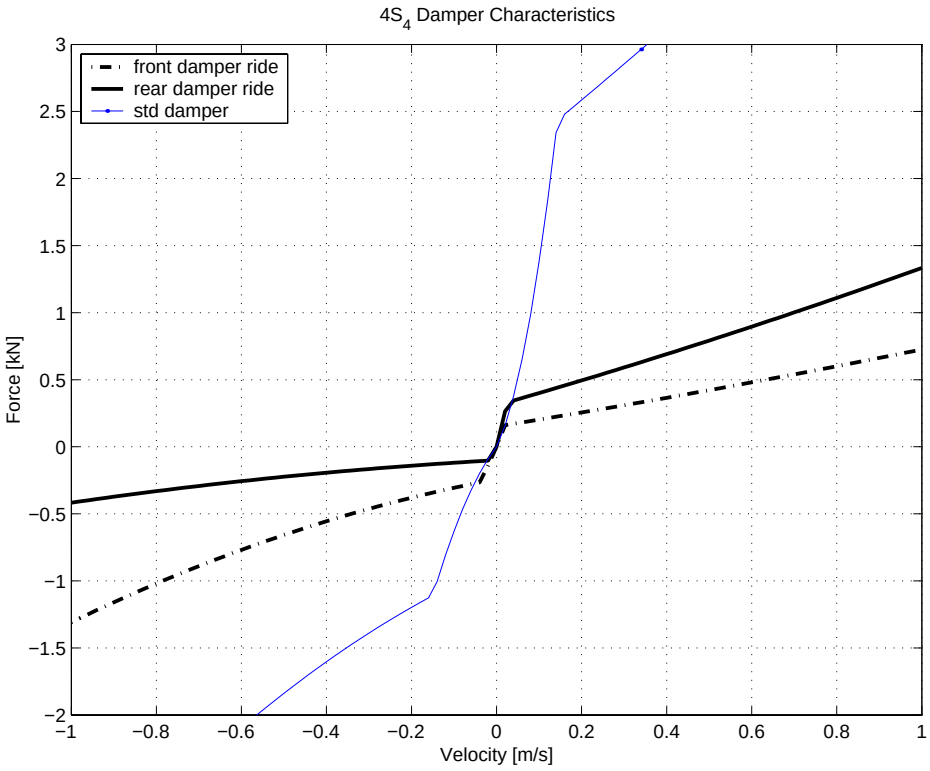

Figure 8.9: Optimum damper characteristics for ride comfort

The automatic scaling methodology can be further improved with the investigation of the Hessian matrix. This Hessian matrix can then be transformed so as to better scale skew problems.

Further investigation into the weight of the penalty function multiplication parameters is needed, so that scaling of constrained problems can be improved.

Chapter 9

Combined Optimisation

With the use of simplified models for gradient information validated, the models are combined to represent the vehicle performing a handling manœuvre on a rough terrain. For the combined ride comfort and handling optimisation, the vehicle performs the double lane change over the Belgian paving. The full simulation model is used, as before, once per iteration for the exact objective function values and constraint values. The Matlab models remain the same. However, the ride model will be used to observe the ride dynamics gradient tendencies, and the handling model for the handling dynamics gradient tendencies. This work was performed before the proposal of the automatic scaling methodology. A study was conducted as to how best to consider the optimisation of the compromise passive suspension. This is done to determine the methodology needed when including the control strategy of the $4S_4$ system for optimisation.

9.1 Handling Followed by Ride Comfort

First the vehicle will be optimised for handling, subject to the tyre hop inequality constraints, and then optimised for ride comfort starting from the point where the handling optimisation converged, for two design variables. The ride comfort is optimised subject to the tyre hop inequality constraints, and an additional inequality constraint that the optimised handling $f^*(\mathbf{x})_{hand}$ may not decrease by more than 20% (compared to the optimised handling

result) as stated below:

$$g(\mathbf{x})_{hand} = 10(f(\mathbf{x})_{hand} - 1.2f^*(\mathbf{x})_{hand}) \leq 0 \tag{9.1}$$

The 20% parameter was selected as it was found that for optimisation runs where the handling constraint was 5 or 10 %, the handling constraint could not be satisfied, if improvements in ride comfort were achieved. The value of 20% was thus found to be a reasonable constraint value. This value would, however, typically depend on the design requirements for the specific vehicle being optimised. The multiplication by 10 was used to better normalise the constraint values between -1 and 1.

The optimisation convergence history for two design variables is presented in Figure 9.1. The equivalent tyre hop constraint is plotted as defined in equation (6.2). The top graph refers to the handling optimisation where the objective function is defined as in equation (5.7), and the bottom graph is for the ride comfort optimisation, where the objective function is defined as in equation (5.6). It can be seen that the optimisation convergence history is well behaved for the handling optimisation, and results in an objective function value of approximately 0.21, which is equivalent to a body roll angle of 3 ^{o}, and a RMS body roll velocity of 1.3 $^o/s$. The ride comfort optimisation, subjected to the handling constraint, has a poorly behaved convergence history, and does not converge to a clear optimum. If iteration 18 is considered as the best minimum, the driver RMS vertical acceleration is approximately 2.2 m/s^2, which is considered as extremely uncomfortable (Els 2005), and needs to be improved. The ride comfort can be greatly improved but at the expense of handling.

9.2 Maximum of Ride Comfort and Handling

The results for handling followed by ride comfort optimisation, prompted the investigation into using the maximum value of the four normalised objective function parameters (roll angle, RMS roll velocity, driver comfort, passenger comfort) as the objective function value. The objective function is thus

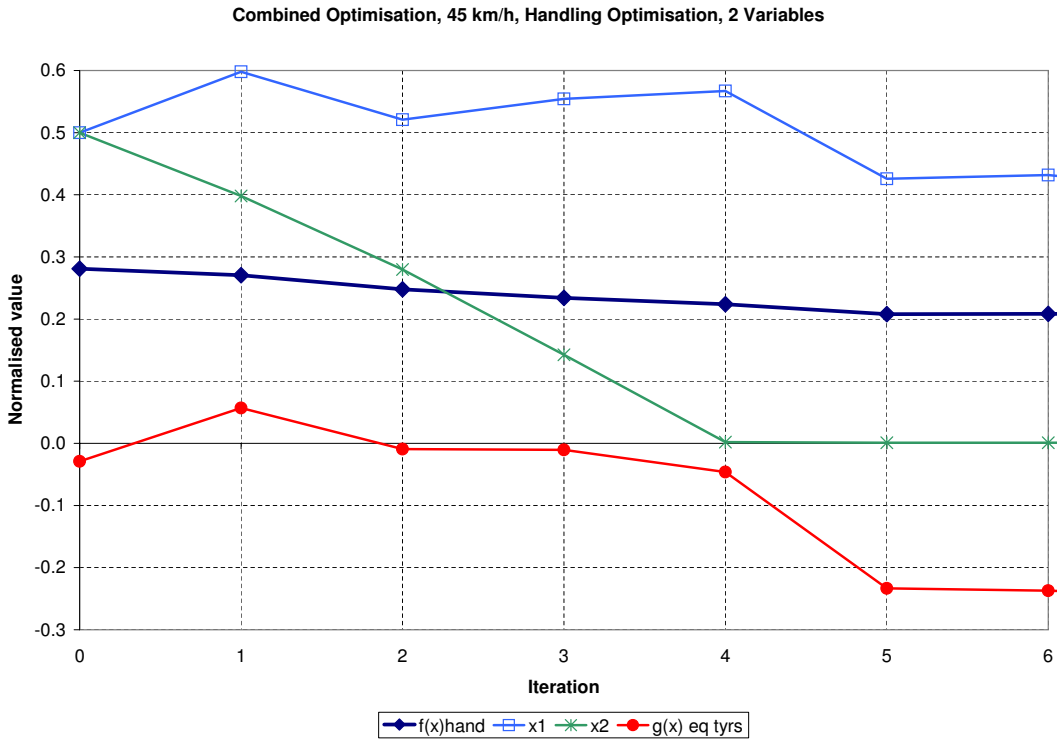

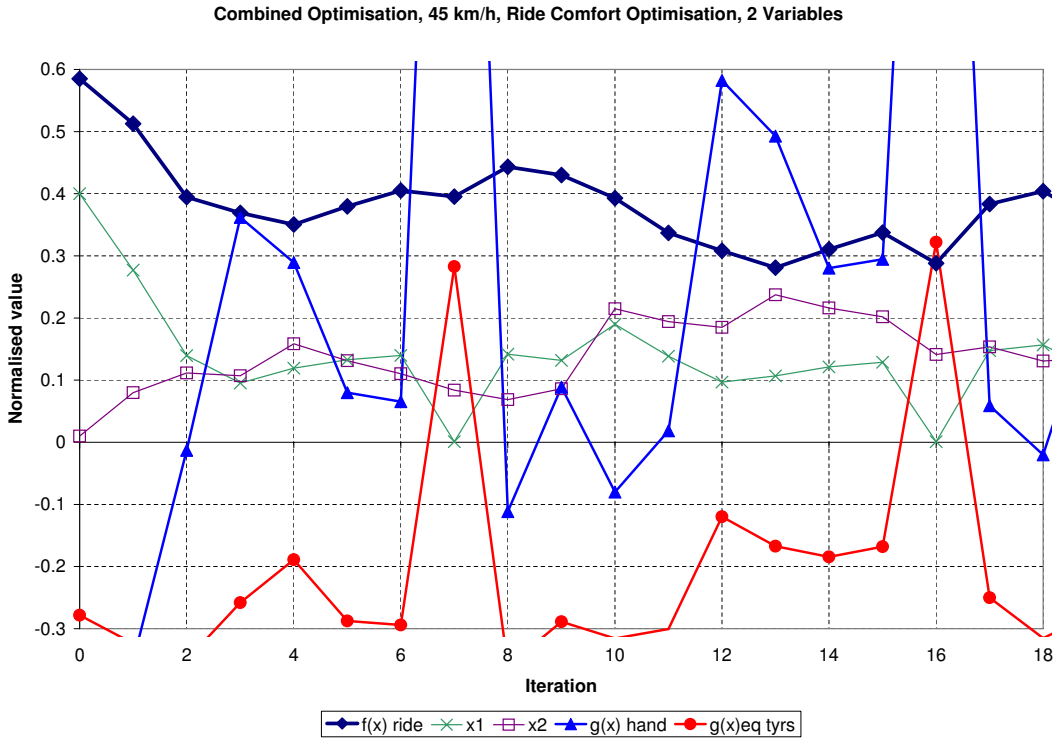

Figure 9.1: Combined convergence history, first handling optimisation, then ride comfort.

defined as follows:

$$f(x) = max(f(x)_{hand}, f(x)_{ride}) \tag{9.2}$$

The disadvantage of the nature of this objective function is the now inherent discontinuities due to the maximum function. However, very reasonable results were achieved as illustrated by Figure 9.2. In the figure, $f(x)hand$ is the handling objective function value as defined by equation (5.7), $f(x)ride$ is the ride comfort objective function as defined by equation (5.6), and the equivalent tyre hop constraint $g(x)eq$ defined by equation (6.2). Additionally it is observed that the overall optimum is the equalization of the two objectives. When considering the final design configuration, iterations 3, 6 and 9, are repeated identical minima, and should be considered for the acceptable band of the design variables, to return objective function values of approximately 0.32. This results in vertical RMS accelerations of approximately 1.8 m/s^2, body roll angle of 4^o, and a RMS roll velocity of 1.9 $^o/s$. The optimisation convergence took fewer iterations than the optimisation of handling followed by ride comfort, even though the objective function is of a discontinuous nature, due to the maximum function.

The use of the maximum function for the objective function was expanded to four design variables, and started in the same place as for two design variables, namely the middle of the design space. The results, presented in Figure 9.3, illustrate the excellent convergence to the optimum, of identical magnitude as for two design variables, but the design variable values differ. Although it is evident that multiple local minima exist, the optimisation converges to identical objective function value minima.

With the difficulty encountered with the definition of ride comfort as a constraint and optimising handling, yet excellent convergence history when using an equal weight of the two objectives, in the form of the maximum function, a pareto front will now be constructed, between the handling and ride comfort objective functions.

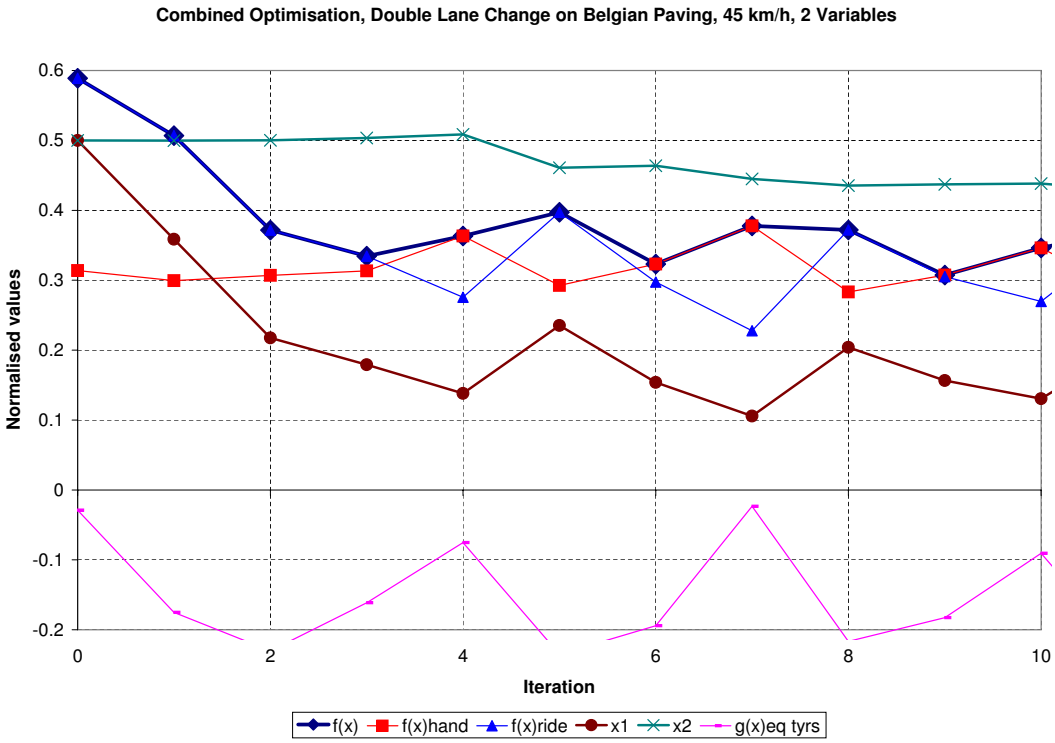

Figure 9.2: Combined optimisation convergence history, maximum of handling and ride comfort objectives, 2 design variables.

9.3 Pareto Optimal Front

With the success of the optimisation results, but the vastly varying optimal design points in the design space, the simplified model was used to investigate the trends in terms of the pareto optimal front of feasible points for ride comfort and handling, subject to the tyre hop constraints. Random points in the four design variable space were generated and their objective and constraint function values evaluated. This would traditionally give the design engineer the necessary insight into which optimal suspension settings to select for a desired combination of ride comfort and handling. However, as shown in Figure 9.4, the random feasible points lie greatly inward of the pareto optimal front. Optimisation runs were performed where the objective function was defined as a weighted sum of the handling and ride comfort objective functions defined in Chapter 5. Figure 9.4, illustrates the optimisation convergence histories of the differing weighted objective

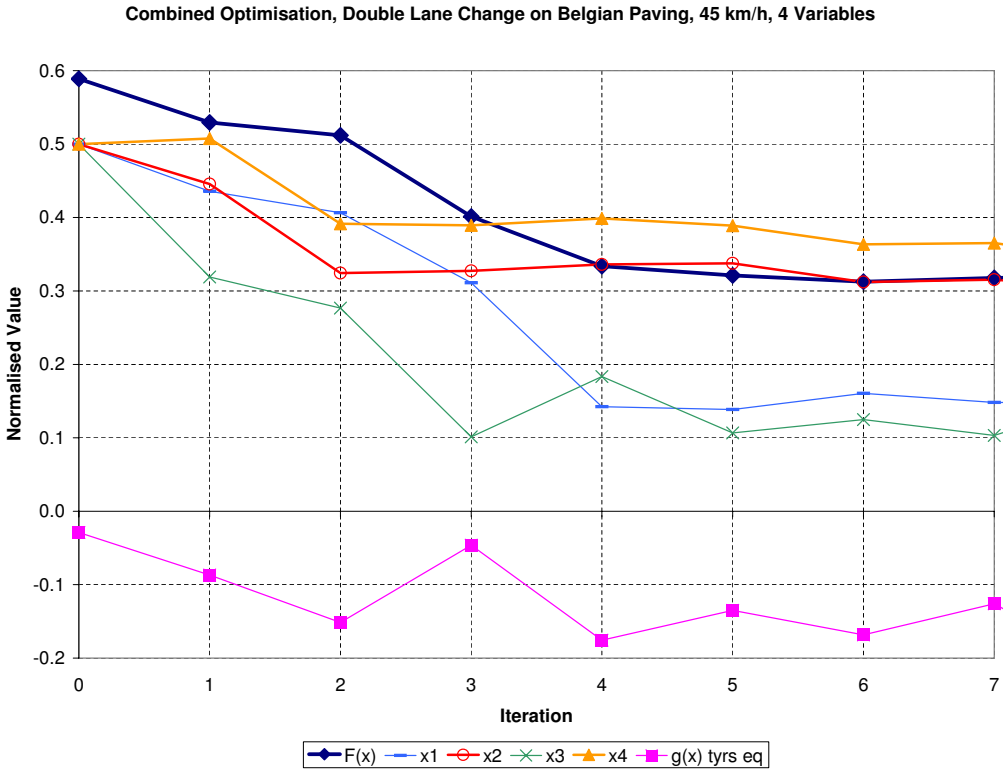

Figure 9.3: Combined optimisation convergence history, maximum of handling and ride comfort objectives, 4 design variables.

functions, to the pareto optimal front. The shortest distance from the pareto front to the zero point is generally accepted as the best compromise, however, this depends on which objective is most important to the vehicle being designed. In the case of the SUV, handling is a safety critical component, as these vehicles at their handling limit roll over before they slide out.

With the converged optimal points of Figure 9.4, the pareto optimal front and design variable values were plotted in Figure 9.5. From Figure 9.5, the change in the design variable with respect to a change in the ride comfort and handling objective function values can be quantified. The design engineer can now use this information to obtain a first order estimate as to the optimal design variable combination in order to achieve a desired point on the optimal pareto front. From Figure 9.5, it is observed that the rear suspension (design variables x_3 and x_4) have the greatest sensitivity on the ride comfort objective function value, when close to the handling optimum (i.e. f($\mathbf{x}$) handling $<$

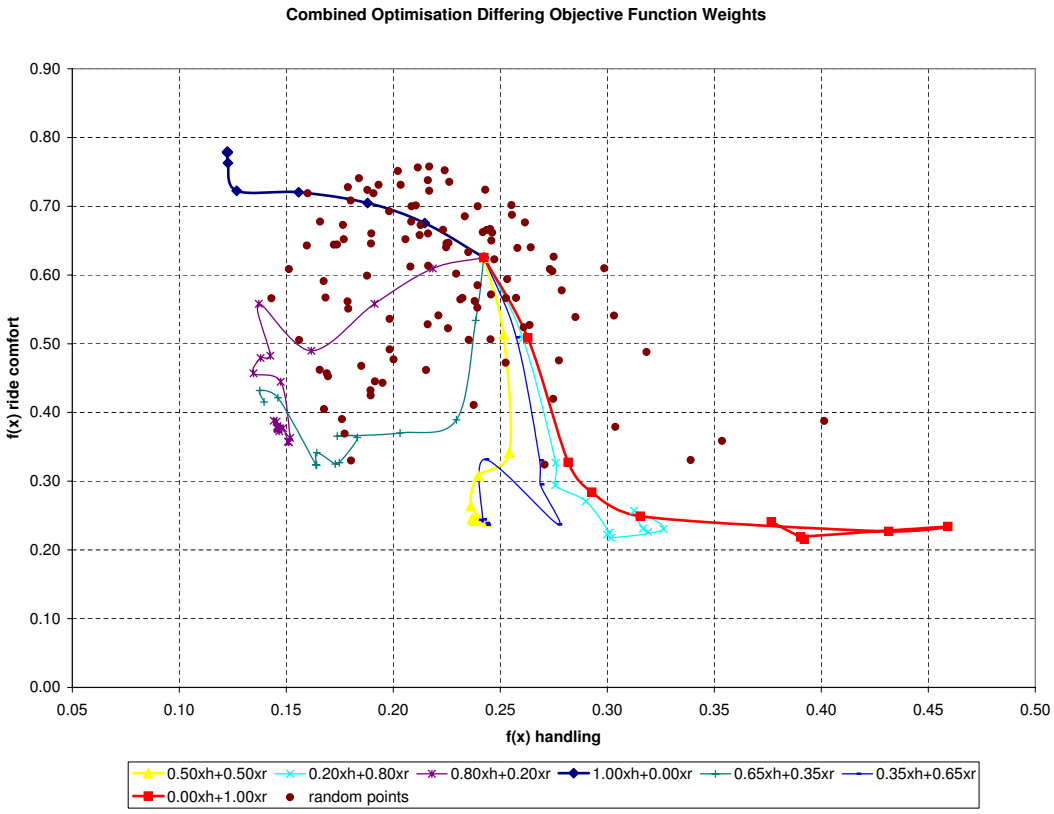

Figure 9.4: Investigation of convergence history to pareto front for different weights of the objective function for handling (h) and ride comfort (r), compared to random feasible points of design space

0.16). A large change in the value of design variables x_1, x_3 and x_4 will result in a dramatic improvement of the ride comfort, but a much smaller decrease in the handling objective function value, from the optimal to worst handling configuration. Also noticeable is the fact that the front gas volume design variable x_2, must be at it's stiffest setting (value of 0) for good handling, yet the rear gas volume design variable x_4 can be as large as 50% (value of 0.3) of the optimal handling gas volume.

From the pareto optimal front results presented in Figure 9.5, it can be concluded that the most feasible compromise point in the design space is for a handling objective function value of 0.16, and a ride comfort objective function value of 0.32. The use of the weighted objective functions to obtain the pareto optimal front is of importance, when designing the vehicle's

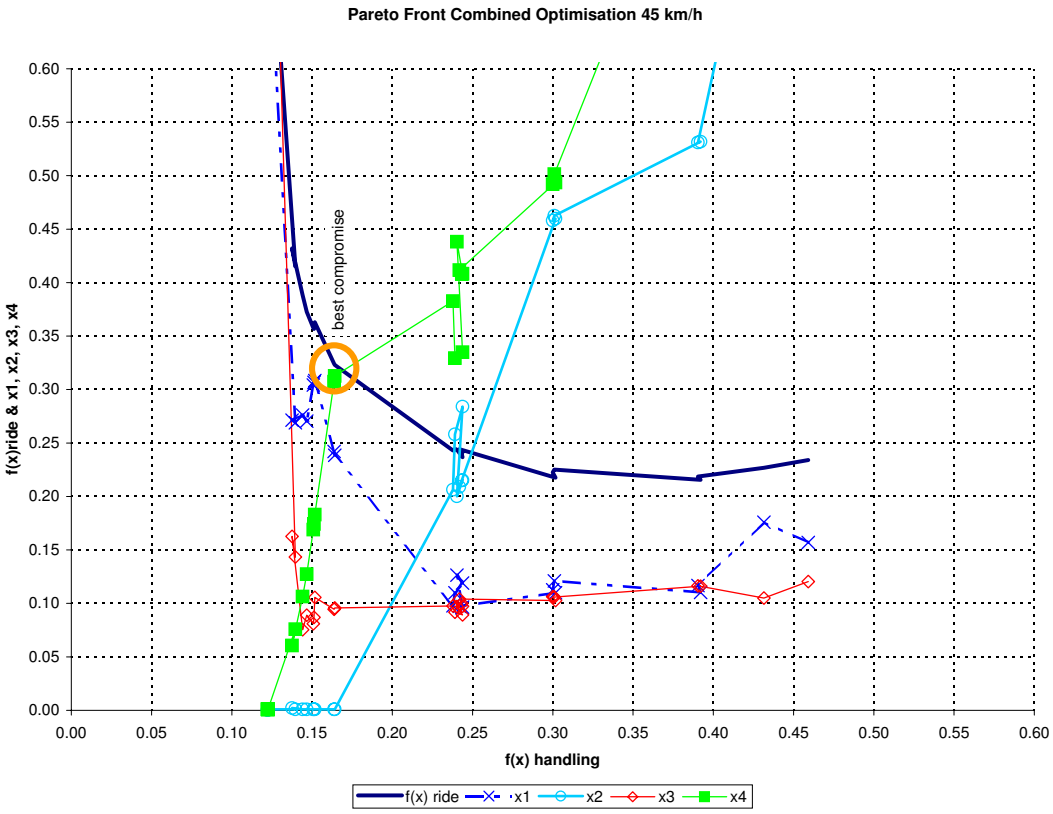

Figure 9.5: Pareto front plot including the change in the design variables along the pareto front

suspension system for the compromise ride comfort vs. handling setup. The necessary insight into the design variables most sensitive to improving the ride comfort with minimal loss of handling ability was obtained.

9.4 Summary of Results

Presented in Table 9.1 are the results for the optimisation runs. From the results it can be seen that the combined optimisation is a compromise between handling and ride comfort, especially when considering the use of the maximum function for the objective function. 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 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 9.2. Also noticeable when observing the parameters of the combined optimisation, is that the gas volume lies in the middle of the design space at 0.3 l, but that the damping should be 50% of the current baseline characteristic. This however, severely affects the handling stability of the vehicle as can be observed by the higher RMS roll velocity value. The most feasible compromise suspension setup was, however, achieved when considering the pareto optimal front.

The pareto optimal front provided the necessary insight into the problem in order to select the most feasible compromise. The resulting optimal damping multiplication factors and spring gas volumes were evaluated using the full MSC.ADAMS simulation model. It was found that the pareto optimal front objective function values were optimistic when compared with the actual full simulation model's objective function values. The full simulation pareto front values for four design configurations along the pareto front were evaluated and presented in Table 9.3. From the results it is observed that the pareto front displayed accurate trends, but with optimistic objective function values. From the pareto test points it can be seen that design configuration 3 returns acceptable ride comfort with a average decrease in handling of 31% over test point 1. This pareto methodology can thus in future be used to optimise a controllable suspension with included control system. Only the simplified models are optimised with objective functions being defined by differing weights, to determine the pareto optimal front. Once the pareto front is determined a few test points, along the pareto optimal front, can be used to determine the actual full simulation model objective function values, to scale the pareto optimal front.

It is suggested that a more effective approach for the determination of the pareto optimal front, would be to run the Dynamic-Q optimisation using a weighted objective function consisting of 100% handling and 0% ride comfort. Once the optimisation has converged to this optimum point, the algorithm should then change the objective function weighting to 80% handling and 20% ride comfort. By continually changing the objective function weighting

Table 9.1: 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} $[m/s^2]$
Combined							
2, handling 1^{st}	9.1	6 (9.8)	0.21	1.29	2.9	-	-
2, ride after	9.1	18 (34.2)	0.40	1.52	3.0	2.20	2.18
2, $f_{max}(x)$	9.2	6 (12.6)	0.32	1.86	4.0	1.78	1.78
4, $f_{max}(x)$	9.3	7 (20.8)	0.32	1.83	4.0	1.76	1.62

Table 9.2: Summary of optimum damper factors and gas volumes

opt. run	Fig.	$dpsff$	$gvolf$	$dpsfr$	$gvolr$
Combined					
2, handling 1^{st}	9.1	1.35	0.10	1.35	0.10
2, ride after	9.1	0.55	0.17	0.55	0.17
2, $f_{max}(x)$	9.2	0.51	0.30	0.51	0.30
4, $f_{max}(x)$	9.3	0.53	0.26	0.40	0.28

during the optimisation procedure, the optimisation should progress along the pareto optimal front, from the best handling objective function value to the best ride comfort objective function value.

9.5 Conclusions

This chapter shows that the use of simplified numerical models, originally used for the optimisation of ride comfort and handling separately for gradient information, Chapters 5 and 6, can be successfully used for the combined optimisation of the ride comfort vs. handling compromise suspension configuration.

The advantages and disadvantages of different definitions of the objective

Table 9.3: ADAMS data for 4 test points along pareto optimal front

test point	$dpsff$	$gvolf$	$dpsfr$	$gvolr$	φ_{peak}	$\dot{\varphi}_{RMS}$	a_{RMS_d}	a_{RMS_p}
1	0.91	0.10	0.39	0.15	1.80	2.93	2.55	2.27
2	0.80	0.10	0.39	0.27	2.18	3.24	2.38	1.89
3	0.54	0.18	0.39	0.28	2.48	3.65	1.89	1.71
4	0.39	0.30	0.39	0.38	2.95	4.44	1.46	1.47

and constraint functions for the determination of the optimal suspension characteristics for combined ride comfort and handling are highlighted in terms of achievable optimal ride comfort and handling. It is, however, found that the use of only the simplified models to optimise for the pareto optimal front is most efficient. The design variable values along the pareto front can be used to determine the actual full simulation model's objective function values. The simplified model's pareto front is accurate in terms of design variable values, but the objective function values differ in absolute value.

It is suggested that a future implementation of continuously varying the weighting attached to the different objectives within the objective function definition will more effectively define the pareto optimal front. The optimisation convergence history, will then converge from the best of the one objective gradually towards the best of the other objective, along the pareto optimal front.

The methodology proposed is shown to be an efficient means of optimising a vehicle's suspension system for combined ride comfort and handling. This research illustrated that the use of gradient-based optimisation algorithms are suitable and competitive for determination of the pareto optimal front necessary for optimising vehicle suspension systems when considering combined ride comfort and handling.

Chapter 10

Conclusions

The use of central finite differences with relatively large perturbation sizes has proven to be beneficial in terms of total function evaluations needed to obtain a feasible minimum. This was measured in terms of less noise in the optimisation convergence history, although at an increased number of function evaluations per iteration, but less overall iterations to reach a feasible optimum. This approach holds definite benefit for all gradient-based optimisation algorithms.

A highly nonlinear vehicle model, with large suspension deflection, that returns excellent correlation to measured results was built in MSC.ADAMS. A novel lateral driver model that makes use of the magic formula to define a nonlinear steering gain factor, was proposed and successfully implemented. This nonlinear steering gain factor modelled with the magic formula made it possible to achieve excellent correlation with measured test data, for a single preview point yaw rate steering driver model. This driver model proved to be robust for different suspension setups, when optimising the vehicle's handling for the closed loop double lane change manœuvre.

The necessity of including wheel hop in the ride comfort optimisation problem was investigated. It was found that it is necessary to include wheel hop as an inequality constraint, when optimising the vehicle's suspension for ride comfort, if the vehicle is to remain stable on rough terrain.

Nonlinear simple models that capture the essence of the handling and ride comfort have been developed, and shown to exhibit similar trends to the full computationally expensive numerical simulation model of the off-road vehicle, at approximately 10 % of the simulation cost. The two design variable case of the simplified models had to be scaled to be equivalent to the full simulation model over the design space. These simplified models have been successful in speeding up the optimisation process, by at least 50% of total simulation time needed when using only the full simulation model for gradient, objective and constraint function values, when used for the determination of gradient information by means of central finite differences. This is a novel approach to vehicle suspension design optimisation and has been shown to be accurate and economical when compared to full simulation gradient based optimisation. The contribution in the field of vehicle design is also underlined by the fact that the same principle can be applied to any gradient-based optimisation algorithm.

The optimisation problem was expanded from 4 design variables to 14. Difficulties were encountered with poor scaling of the design variables, and noise associated with infeasible tyre deflections due to the current tyre model only accommodating a linear vertical tyre stiffness. Scaling of the optimisation problem has been investigated, with the result that sphericality of the design space is more important than having equivalent ranges and magnitudes of design variables. Great improvements were achieved in the optimisation convergence histories, when the optimisation problem was better scaled.

The scaling process that was followed, is reformulated into a novel automatic scaling methodology, that can help engineers reduce the time necessary for investigation of design variable scaling. This methodology was tested on analytic functions, and found to improve the optimisation convergence for most tested problems. The methodology was expanded to include constrained optimisation problems in the form of the penalty function, but further

experimentation is required with the penalty function multiplication factor to be used. The automatic scaling methodology was applied to the vehicle suspension optimisation for 14 design variables, with great success. When automatic scaling was applied, better optimum values were reached, than without automatic scaling.

There is at this stage no guarantees that the results achieved are global optima, as concluded from the fact that some of the optimisation results return similar objective function values for different design variable combinations. The aim is however, for an improvement in handling and ride comfort rather than the absolute global optimum suspension setup. There is also no guarantee that if the global optimum is found, the design variables are robust in terms of manufacturing tolerances.

Noise in the inequality constraints, when optimising ride comfort, and combined optimisation of ride comfort and handling, is still problematic. Some avenues were investigated but a more intensive investigation is needed before the problem is fully understood.

The combined optimisation of ride comfort and handling was investigated. Various concepts were investigated for the definition of the objective function. The discontinuous nature of the maximum function in the definition of the objective function was found to pose no difficulties in terms of optimisation convergence. When optimisation was performed using the baseline vehicle's handling as a constraint no improvement was found in ride comfort, and the same applies when the baseline vehicle's ride comfort was used as a constraint and handling optimised no improvement was observed. This was because the baseline vehicle's design point lay outside the feasible design space achievable with the current $4S_4$ suspension system when optimised for the compromise. This will probably be overcome when the control system is included in the optimisation process.

Chapter 11

Discussion of Future Work

The results of the 14 design variable optimisation, postulated that a solver change has the potential to greatly reduce the numerical noise present in the objective functions with respect to small perturbations of the design variables. This would be implemented by interpolation of the equivalent function value, for a constant numerical error, at each time step of the multi-body dynamics solver. It is believed that this alone will greatly contribute to reduced noise in the objective functions obtained from numerical simulations.

The vertical tyre stiffness should be modelled as nonlinear, so as to capture the effect of increasing tyre stiffness with high tyre deflections. This will be beneficial when considering the tyre hop inequality constraint at low suspension damping, which is currently resulting in optimisation convergence difficulties when considering many design variables. This, however, is not easily implementable in the current ADAMS Pacejka '89 tyre model used. Other tyre models would have to be investigated for the implementation of this stiffness characteristic.

A preliminary investigation has been performed in the use of gradient only optimisation algorithms like LFOPC (Snyman 2000) and ETOPC (Snyman 2005a), used with the simplified models, not presented in this thesis. This should be further investigated, as it could prove to be more efficient.

A greater variety of road conditions need to be considered over varying vehicle speeds and loading conditions, before a decision can be made regarding the final overall optimum design. The ultimate test will be the optimisation of the vehicle's performance under severe handling manœuvres on an uneven road. The methodology presented in this research is easily adaptable to these multiple conditions.

The incorporation of the complex model describing the hydro-pneumatic suspension's characteristics as proposed by Theron and Els (2005), should be included with the control of switching proposed by Els (2006), in the final optimisation phase.

The final optimised spring and damper characteristics should be investigated for robustness. This should be done in terms of the effect normal manufacturing tolerances will have on the vehicle's handling and ride comfort.

The proposed automatic scaling methodology should be further researched so as to take Hessian information into account before scaling the design space, in an effort to minimize the negative effects cross-coupled design variables have on the current scaling method.

A variable weighting when performing multi-objective optimisation needs to be investigated, so as to more efficiently plot the pareto optimal front, from the optimum of the one objective to the optimum of the other objective.

Bibliography

Alleyne, A. and Hedrick, J. K.: 1995, Nonlinear adaptive control of active suspensions, *IEEE Transactions on Control Systems Technology* **Vol. 3/1**, pp. 94–101.

Andersson, D. and Eriksson, P.: 2004, Handling and ride comfort optimisation of an inter-city bus, *Vehicle System Dynamics Supplement* **Vol. 41**, pp. 547–556.

Bakker, E., Pacejka, H. B. and Linder, L.: 1989, A new tire model with an application in vehicle dynamics studies, *Society of Automotive Engineers, Technical Paper Series* **Paper no. SAE 890087**.

Balabanov, V. O. and Venter, G.: 2004, Multi-fidelity optimisation of high-fidelity analysis and low-fidelity gradients, 10^{th} *AIAA/ISSMO Symposium on multidisciplinary analysis and optimisation. Albany, NY. 30 August - 1 September* .

Bandler, J. W., Cheng, Q. S., Dakroury, S. A., Mohamed, A. S., Bakr, M. H., Madsen, K. and Søndergaard, J.: 2004, Space mapping: The state of the art, *IEEE Transactions on Microwave Theory and Techniques* **Vol. 52-1**, pp. 337–361.

Bartholomew-Biggs, M., Brown, S., Christianson, B. and Dixon, L.: 2000, Automatic differentiation of algorithms, *Journal of Computational and Applied Mathematics* **Vol. 124**, pp. 171–190.

Baumal, A. E., McPhee, J. and Calamai, P. H.: 1998, Application of genetic algorithms to the optimisation of an active vehicle suspension design,

Computer Methods in Applied Mechanics and Engineering **Vol. 163**, pp. 87–94.

Bischof, C. H., Bücker, H. M., Lang, B., Rasch, A. and Slusanschi, E.: 2005, Efficient and accurate derivatives for a software process chain in airfoil shape optimisation, *Future Generation Computer Systems* **Vol. 21**, pp. 1333–1344.

Boggs, P. T. and Tolle, J. W.: 2000, Sequential quadratic programming for large-scale non-linear optimisation, *Journal of Computational and Applied Mathematics* **Vol. 124**, pp. 123–137.

BS6841: 1987, *Guide to Measurement and Evaluation of Human Exposure to Whole-Body Mechanical Vibration and Repeated Shock*, British Standards Institution.

Dahlberg, T.: 1977, Parametric optimisation of a 1-dof vehicle travelling on a randomly profiled road, *Journal of Sound and Vibration* **Vol. 55-2**, pp. 245–253.

Dahlberg, T.: 1979, An optimised speed controlled suspension of a 2-dof vehicle travelling on a randomly profiled road, *Journal of Sound and Vibration* **Vol. 64-4**, pp. 541–546.

Danchick, R.: 2006, Accurate numerical partials with applications to optimisation, *Applied Mathematics and Computation* **in press, accessed science direct, 18 November 2006**, doi:10.1016/j.amc.2006.05.083.

Eberhard, P., Bestle, D. and Piram, U.: 1995, Optimisation of damping characteristics in non-linear dynamic systems, *First World Congress of Structural and Multidisciplinary Optimisation* pp. 863–870.

Eberhard, P., Schiehlen, W. and Bestle, D.: 1999, Some advantages of stochastic methods in multicriteria optimisation of multibody systems, *Archive of Applied Mechanics* **Vol. 69**, pp. 543–554.

Els, P. S.: 2005, The applicability of ride comfort standards to off-road vehicles, *Journal of Terramechanics* **Vol. 42-1**, pp. 47–64.

Els, P. S.: 2006, *The ride comfort vs. handling compromise for off-road vehicles*, Phd thesis, University of Pretoria, South Africa.

Els, P. S. and Uys, P. E.: 2003, Investigation of the applicability of the Dynamic-Q optimisation algorithm to vehicle suspension design, *Mathematical and Computer Modelling* **Vol. 37**, pp. 1029–1046.

Els, P. S., Uys, P. E., Snyman, J. A. and Thoresson, M. J.: 2006, Gradient-based approximation methods applied to the optimal design of vehicle suspension systems using computational models with severe inherent noise, *Mathematical and Computer Modelling* **Vol. 43-7/8**, pp. 787–801.

Eriksson, P. and Arora, J. S.: 2002, Comparison of global optimisation algorithms applied to a ride comfort optimisation problem, *Structural and Multidisciplinary Optimisation* **Vol. 24**, pp. 157–167.

Eriksson, P. and Friberg, O.: 2000, Ride comfort optimisation of a city bus, *Structural and Multidisciplinary Optimisation* **Vol. 20**, pp. 67–75.

Etman, L. F. P., Vermeulen, R. C. N., van Heck, J. G. A. M., Schoofs, A. J. G. and van Campen, D. H.: 2002, Design of a stroke dependent damper for the front axle suspension of a truck using multibody system dynamics and numerical optimisation, *Vehicle System Dynamics* **Vol. 38-2**, pp. 85–102.

Fiacco, A. V. and McCormick, G. P.: 1968, *Non-linear Programming: Sequential Unconstrained Minimization Techniques*, John Wiley and Sons, Inc, New York.

Genta, G.: 1997, *Motor Vehicle Dynamics, Modelling and Simulation. Series on Advances in Mathematics for Applied Sciences: Volume 43*, World Scientific, New Jersey, USA.

Gerotek: 2006, *Gerotek Testing Facilities, South Africa*, www.gerotek.co.za *accessed 10 April 2006*.

Gobbi, M., Haque, I., Papalambros, P. Y. and Mastinu, G.: 2005, Optimisation and integration of ground vehicle systems, *Vehicle System Dynamics* **43/6-7**, pp. 437–453.

Gobbi, M., Mastinu, G. and Doniselli, C.: 1999, Optimising a car chassis, *Vehicle System Dynamics* **Vol. 32**, pp. 149–170.

Gobbi, M., Mastinu, G., Doniselli, C., Guglielmetto, L. and Pisino, E.: 1999, Optimal and robust design of a road vehicle suspension system, *Vehicle System Dynamics Supplement* **Vol. 33**, pp. 3–22.

Gonsalves, J. P. C. and Ambrósio, J. A. C.: 2005, Road vehicle modelling requirements for optimisation of ride comfort and handling, *Multibody System Dynamics* **13**, pp. 3–23.

Gordon, T. J., Best, M. C. and Dixon, P. J.: 2002, An automated driver based on convergent vector fields, *Proceedings of the Institute of Mechanical Engineers, Part D: Journal of Automotive Engineering* **Vol. 216 Special Issue**, pp. 36–56.

Hock, W. and Schittkowski, K.: 1981, *Test Examples for Nonlinear Programming Codes. Lecture Notes in Economics and Mathematical Systems: Volume 187*, Springer-Verlag, Berlin, Germany.

ISO3888-1: 1999, *Passenger cars - Test track for a severe lane-change manœuvre - Part 1: Double lane-change*, The International Organisation for Standardisation, 06 October 2004.

ISO8608: 1995, *Mechanical Vibration - Road Surface Profiles - Reporting of Measured Data*, The International Organisation for Standardisation, 01 September 1995.

Kim, C. and Ro, P. I.: 1998, A sliding mode controller for vehicle active suspension systems with non-linearities, *Proceedings of International Mechanical Engineers* **Vol. 212 (D)**, pp. 79–92.

Kim, C. and Ro, P. I.: 2001, An accurate model for vehicle handling using reduced-order model techniques, *Society of Automotive Engineering Technical Paper Series* **SAE 2001-01-2520**.

Koziel, S., Bandler, J. W. and Madsen, K.: 2005, Towards a rigorous formulation of the space mapping technique for engineering design, *Circuits and Systems, ISCAS 2005 IEEE International Symposium* **Vol. 23-26 May**, pp. 5605–5608.

Lasdon, L.: 2001, Nonlinear and geometric programming - current status, *Annals of Operations Research* **105**, pp. 99–107.

Mathworks: 2000a, *MATLAB Optimisation Toolbox Users Guide version 2.2*.

Mathworks: 2000b, *MATLAB Release 12 Users guide*.

Miller, L. R.: 1998, Tuning passive, semi-active, and fully active suspension systems, *Proceedings of the 27th Conference on Decision and Control. Austin, Texus. December* .

MSC: 2005, *Getting Started Using MSC.ADAMS View, Version 2005*, MSC-Corporation.

Naudé, A. F.: 2001, *Computer Aided Design Optimisation of Road Vehicle Suspension Systems*, Phd thesis, University of Pretoria, South Africa.

Naudé, A. F. and Snyman, J. A.: 2003a, Optimisation of road vehicle passive suspension systems. Part 1. Optimisation algorithm and vehicle model, *Applied Mathematical Modelling* **Vol. 27**, pp. 249–261.

Naudé, A. F. and Snyman, J. A.: 2003b, Optimisation of road vehicle passive suspension systems. Part 2. Qualification and case study, *Applied Mathematical Modelling* **Vol. 27**, pp. 263–274.

Pacejka, H. B.: 2002, *Tyre and Vehicle Dynamics*, Society of Automotive Engineers, Warrendale, USA.

Papalambros, P. Y.: 2002, The optimisation paradigm in engineering design: promises and challenges, *Computer-Aided Design* **34**, pp. 939–951.

Pilbeam, C. and Sharp, R. S.: 1996, Performance potential and power consumption of slow-active suspension systems with preview, *Vehicle System Dynamics* **Vol. 25**, pp. 169–183.

Proköp, G.: 2001, Modelling human vehicle driving by model predictive online optimisation, *Vehicle System Dynamics* **Vol. 35/1**, pp. 19–53.

Redhe, M. and Nilsson, L.: 2004, Optimisation of the new SAAB 9-3 exposed to impact load using a space mapping technique, *Structural and Multidisciplinary Optimisation* **Vol. 27**, pp. 411–420.

Schuller, J., Haque, I. and Eckel, M.: 2002, An approach for optimisation of vehicle handling behaviour in simulation, *Vehicle System Dynamics Supplement* **Vol. 37**, pp. 24–37.

Sharp, R. S., Casanova, D. and Symonds, P.: 2000, A mathematical model for driver steering control, with design, tuning and performance results, *Vehicle System Dynamics* **Vol. 33**, pp. 289–326.

Snyman, J. A.: 2000, The lfopc leap-frog algorithm for constrained optimisation, *Computers and Mathematics with Applications* **Vol. 40**, pp. 1085–1096.

Snyman, J. A.: 2005a, A gradient-only line search method for the conjugate gradient method applied to constrained optimisation problems with severe noise in the objective function, *International Journal for Numerical Methods in Engineering* **62**, pp. 72–82.

Snyman, J. A.: 2005b, *Practical Mathematical Modelling, An introduction to Basic Optimisation theory and classical and New Gradient-Based Algorithms. Applied Optimisation: Volume 97*, Springer, New York, USA.

Snyman, J. A. and Hay, A. M.: 2002, The Dynamic-Q optimisation method: An alternative to SQP?, *Computers and Mathematics with Applications* **Vol. 44**, pp. 1589–1598.

Theron, N. J. and Els, P. S.: 2005, Modelling of a semi-active hydropneumatic spring-damper unit, *International Journal of Vehicle Design* **accepted for publication**.

Thoresson, M. J.: 2003, *Mathematical optimisation of the suspension system of an off-road vehicle for ride comfort and handling*, Meng thesis, University of Pretoria, South Africa.

Tolsma, J. E. and Barton, P. I.: 1998, On computational differentiation, *Computers in Chemical Engineering* **Vol. 22-4/5**, pp. 475–490.

Uys, P. E., Els, P. S. and Thoresson, M. J.: 2006a, Criteria for handling measurement, *Journal of Terramechanics* **Vol. 43-1**, pp. 43–67.

Uys, P. E., Els, P. S. and Thoresson, M. J.: 2006b, Suspension settings for optimal ride comfort of off-road vehicles travelling on roads with different roughness and speeds, *Journal of Terramechanics* pp. In Press, available on ScienceDirect 11 July 2006.

Uys, P. E., Els, P. S., Thoresson, M. J., Voigt, K. G. and Combrink, W. C.: 2005, Experimental determination of moments of inertia for an off-road vehicle in a regular engineering laboratory, *International Journal of Mechanical Engineering Education* **accepted for publication**.

van Keulen, F. and Toropov, V. V.: 2006, Multipoint approximations for structural optimisation problems with noisy response functions, *Bradford University, webpage: http://0-www.bradford.ac.uk.innopac.up.ac.za/staff/vtoropov/fred/int2a.htm* **accesed 13 September 2006**, pp. 1–8.

Vanderplaats, G.: 1992, *DOT - Design Optimisation Tools. Version 3.0*, VMA Engineering, Colorado Springs, USA.

Vanderplaats, G.: 1999, *DOT - Design Optimisation Tools, Users Manual. Version 5.0*, Vanderplaats Research and Development, Colorado Springs, USA.

Willcox, K.: 2006, Design and optimisation of complex systems, *High Performance Computation for Engineered Systems, Massachusetts institute of technology, webpage: https://0-dspace.mit.edu.innopac.up.ac.za:443/retrieve/4002/HPCES005.pdf* **accesed 18 November 2006**, pp. 1–7.

Appendix A

Auto-Scaling Test Problems

The test problems described here are from the book of Hock and Schittkowski 1981, and used in Chapter 8 to test the automatic scaling theory, before being applied to the vehicle suspension optimisation problem.

A.1 Hock 2

Objective function:

$$f(\mathbf{x}) = 100(x_2 - {x_1}^2)^2 + (1 - x_1)^2 \tag{A.1}$$

Constraints:

$$\begin{aligned} -2 \leq x_1 \leq 2 \\ 1.5 \leq x_2 \leq 3 \end{aligned} \tag{A.2}$$

Starting point:

$$\begin{aligned} \mathbf{x}_0 &= [-2\ 1] \\ f(\mathbf{x}_0) &= 909 \end{aligned} \tag{A.3}$$

Optimum:

$$\begin{aligned} \mathbf{x}^* &= [1.22\ 1.5] \\ f(\mathbf{x}^*) &= 0.05042\ 61879 \end{aligned} \tag{A.4}$$

A.2 Hock 13

Objective function:

$$f(\mathbf{x}) = (x_1 - 2)^2 + {x_2}^2 \tag{A.5}$$

Constraints:

$$\begin{gathered} g(\mathbf{x}) = x_2 + (1 - x_1)^3 \leq 0 \\ -2 \leq x_1 \leq 2 \\ -2 \leq x_2 \leq 2 \end{gathered} \tag{A.6}$$

Starting point:

$$\begin{gathered} \mathbf{x}_0 = [-2 \ \ -2] \\ f(\mathbf{x}_0) = 20 \end{gathered} \tag{A.7}$$

Optimum:

$$\begin{gathered} \mathbf{x}^* = [1 \ 0] \\ f(\mathbf{x}^*) = 1 \end{gathered} \tag{A.8}$$

A.3 Hock 15

Objective function:

$$f(\mathbf{x}) = 100(x_2 - {x_1}^2)^2 + (1 - x_1)^2 \tag{A.9}$$

Constraints:

$$\begin{gathered} g(\mathbf{x}) = 1 - x_1 x_2 \leq 0 \\ g(\mathbf{x}) = -x_1 - {x_2}^2 \leq 0 \\ -2 \leq x_1 \leq 0.5 \\ 1 \leq x_2 \leq 2.5 \end{gathered} \tag{A.10}$$

Starting point:

$$\begin{gathered} \mathbf{x}_0 = [-2 \ 1] \\ f(\mathbf{x}_0) = 909 \end{gathered} \tag{A.11}$$

Optimum:

$$\begin{gathered} \mathbf{x}^* = [0.5 \ 2] \\ f(\mathbf{x}^*) = 306.5 \end{gathered} \tag{A.12}$$

A.4 Hock 17

Objective function:

$$f(\mathbf{x}) = 100(x_2 - {x_1}^2)^2 + (1 - x_1)^2 \tag{A.13}$$

Constraints:

$$\begin{aligned} g(\mathbf{x}) &= x_1 - {x_2}^2 \leq 0 \\ g(\mathbf{x}) &= x_2 - {x_1}^2 \leq 0 \\ -0.5 &\leq x_1 \leq 0.5 \\ -1 &\leq x_2 \leq 1 \end{aligned} \tag{A.14}$$

Starting point:

$$\begin{aligned} \mathbf{x}_0 &= [-2\ 1] \\ f(\mathbf{x}_0) &= 909 \end{aligned} \tag{A.15}$$

Optimum:

$$\begin{aligned} \mathbf{x}^* &= [0\ 0] \\ f(\mathbf{x}^*) &= 1 \end{aligned} \tag{A.16}$$

Appendix B

Vehicle Model Files

Table B.1: Vehicle mass and inertia properties

Body	Mass [kg]	I_{xx}	I_{yy}	I_{zz}
body front	682	909	0	0
body rear	894	952	0	0
tyres	31.5	1.2	1.2	2.0
front axle	166	22.3	0.13	22.3
rear axle	166	22.3	0.13	22.3
steer link	3	0.4	0.4	0

- all other links have 0 mass properties.

List of Tables

List of Figures