UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA

Optimal vehicle suspension characteristics for increased structural fatigue life

by

Hendrik Gerhardus Abraham Breytenbach

Submitted in partial fulfilment of the requirements for the degree

Master of Engineering
(Mechanical Engineering)

in the Faculty of

Engineering, Built Environment and Information Technology (EBIT)

at the

University of Pretoria,
Pretoria

November 2009

Optimal vehicle suspension characteristics for increased structural fatigue life

Author: Hendrik Gerhardus Abraham Breytenbach
Supervisor: Prof. P.S. Els
Department: Mechanical and Aeronautical Engineering
Degree: Master of Engineering

Abstract

The designers of heavy, off-road vehicle suspension systems face unique challenges. The ride comfort versus handling compromise in these vehicles has been frequently investigated using mathematical optimisation. Further challenges exist due to the large variations in vehicle sprung mass. The suspension system must provide adequate isolation from road load inputs throughout its payload operating range. This is imperative if good vehicle structural life is to be ensured. A passive suspension system can only provide optimal isolation at a single payload. The designer of such a suspension system must therefore make a compromise between designing for a fully-laden or unladen payload state. This work deals with suspension optimisation for vehicle structural life.

The work mainly addresses two questions: 1) What are the suspension characteristics required to ensure optimal isolation of the vehicle structure from road loads? and 2) If such optimal suspension characteristics can be found, how sensitive are they to changes in vehicle payload? The study aims to answer these questions by examining a Land Rover Defender 110 as case study.

An experimentally validated mathematical model of the test vehicle is constructed for the use in sensitivity studies. Mathematical optimisation is performed using the model in order to find the suspension characteristics for optimal structural life of the vehicle under consideration. Sensitivity studies are conducted to determine the robustness of the optimal characteristics and their sensitivity to vehicle payload variation. Recommendations are made for suspension characteristic selection for optimal structural life.

Optimale voertuig suspensie karakteristieke vir verlengde struktuurvermoeidheidsleeftyd

Outeur: Hendrik Gerhardus Abraham Breytenbach
Studieleier: Prof. P.S. Els
Departement: Meganiese en Lugvaartkundige Ingenieurswese
Graad: Magister in Ingenieurswese

Opsomming

Ontwerpers van swaar, veldvoertuig suspensie stelsels staar unieke uitdagings in die gesig. Die ritgemak teenoor hantering kompromie in hierdie voertuie is reeds telkemale ondersoek, ook met wiskundige optimering. Verdere uitdagings bestaan as gevolg van die groot veranderinge in geveerde massa by hierdie voertuie. Die suspensiestelsel moet gepaste isolasie van pad insette oor `n wye reeks van bedryfstoestande lewer. Dit is veral belangrik indien daar verseker wil word dat die voertuig goeie struktuurleeftyd het. `n Passiewe suspensiestelsel kan egter slegs optimale isolasie by `n enkele vragtoestand lewer. Die ontwerper van `n passiewe suspensie stelsel moet dus `n kompromie aangaan tussen ontwerp vir `n vol of leë vragtoestand. Hierdie studie handel oor suspensie optimering vir struktuur leeftyd.

Die werk spreek hoofsaaklik twee vraagstukke aan: 1) Watter suspensie karakteristieke word benodig om die voertuig struktuur optimaal van padinsette te isoleer? en 2) Indien sulke optimale karakteristieke gevind kan word, wat is hulle sensitiwiteit vir veranderinge in voertuig vrag? Die studie mik om hierdie vraagstukke aan te spreek deur ondersoeke op `n Land Rover Defender 110 toetsvoertuig.

`n Eksperimenteel gevalideerde, wiskundige model van die toetsvoertuig word saamgestel met die oog op sensitiwiteitstudies. Wiskundige optimering word met die model uitgevoer om sodoende die suspensie karakteristieke vir optimale struktuurleeftyd vir die betrokke toetsvoertuig te bepaal. Sensitiwiteitsanalises word gedoen om die robuustheid van die optimale karakteristieke, met betrekking tot veranderinge in voertuig vrag, vas te stel. Aanbevelings word gemaak oor die keuse van suspensie karakteristieke vir optimale struktuur leeftyd.

Acknowledgments

Ek is met voltooiing van hierdie verhandeling aan verskeie mense dank verskuldig:

- Heinrich Bauermeister by MMS tegnologie vir die twee jaar verlof op my beursverpligtinge, wat my nagraadse studie moontlik gemaak het.

- My ouers, wat met niks begin het en vir my en my broer elke moontlike geleentheid geskep het. Julle liefde en ondersteuning is onmeetlik en het hierdie mylpaal in my lewe moontlik gemaak.

- Prof. Els vir vyf jaar se mentorskap deur Baja, finalejaar ontwerp, finalejaar skripsie en my nagraadse studies. Prof het `n geweldige bydra aan my beroepsmondering gemaak.

- My mede-studente vir die insette, kritiek, aanmoediging en vele tegniese en ander besprekings gedurende die afgelope twee jaar. Die saam werk, leer en kuier oor die twee jaar sal ek vir niks verruil nie!

- Melissa vir jou vriendskap, liefde, ondersteuning en ook opoffering oor bietjie meer as vyf jaar. Jou aandeel in my menswees is onmisbaar.

Table of contents

List of figures

List of tables

List of abbreviations

ABS	Anti-Lock Braking System
ADAMS	Automatic Dynamic Analysis of Mechanical Systems
CDAS	Command and Data Acquisition System
CG	Centre of Gravity
COV	Coefficient of Variation
DADS	Dynamic Analysis and Design System
DLC	Dynamic Load Coefficient
dml	Dynamic Move Limit
DOF	Degree(s) of Freedom
DSD	Displacement Spectral Density
dsf	Damper Scale Factor
EoM	Equation(s) of Motion
FBD	Free-body Diagram
FFT	Fast Fourier Transform
GPS	Global Position System
IRI	International Roughness Index.
KD	Kinetic Diagram
LDT	Linear Displacement Transducer
LF	Left Front
LR	Left Rear
NAASRA	National Association of Australian State Road Authorities
ODE	Ordinary Differential Equation
ov	Optimisation Variable
PSD	Power Spectral Density
RF	Right Front
RR	Right Rear
RMS	Root Mean Square
Std. Dev.	Standard Deviation
sgv	Static Gas Volume
S-N	Stress-Life / Stress at Failure–Number of Cycles to Failure
4S4	4 State Semi-Active Suspension System

List of symbols

Roman symbols:

a	=	Distance from rear wheels to CG [m] (Chapter 3, Appendix A)
b	=	Distance from front wheel to CG [m] (Chapter 3, Appendix A)
c	=	Half of vehicle track width [m] (Chapter 3, Appendix A)
d	=	Half of suspension attachment width [m] (Chapter 3, Appendix A)
c_{tlf}	=	Tyre damping rate/coefficient, left front [Ns/m] (Chapter 3, Appendix A)
c_{tlr}	=	Tyre damping rate/coefficient, left rear [Ns/m] (Chapter 3, Appendix A)
c_{trf}	=	Tyre damping rate/coefficient, right front [Ns/m] (Chapter 3, Appendix A)
c_{trr}	=	Tyre damping rate/coefficient, right rear [Ns/m] (Chapter 3, Appendix A)
$f(t)$	=	Random data set, cross-correlation function, (Chapter 3)
$f(\bar{x})$	=	Cost function (Chapter 5)
f_{tol}	=	Function value convergence tolerance for Dynamic-Q (Chapter 5)
g	=	Gravitational acceleration constant [9.81m/s^2] (Chapter 3, Appendix A)
$g_{1,2,3}$	=	Inequality constraint function (Chapter 5)
$g(\bar{x})$	=	Inequality constraint function (Chapter 5)
$g_{new\ \ j}$	=	Robust constraint function (Chapter 5)
h	=	Centre of mass vertical position [m] (Chapter 3)
$h(\bar{x})$	=	Equality constraint function (Chapter 5)
k	=	Stiffness or spring rate [N/m] (Appendix B)
k_{ef}	=	Effective stiffness [N/(m/m)] (Chapter 3)
k_j	=	Penalty parameter for robust constraints (Chapter 5)
k_{loop}	=	Maximum number of iterations in Dynamic-Q (Chapter 5)
k_{stat}	=	Stiffness about the static position [N/m] (Appendix B)
k_{tlf}	=	Tyre stiffness, left front [N/m] (Chapter 3, Appendix A)
k_{tlr}	=	Tyre stiffness, left rear [N/m] (Chapter 3, Appendix A)
k_{trf}	=	Tyre stiffness, right front [N/m] (Chapter 3, Appendix A)
k_{trr}	=	Tyre stiffness, right rear [N/m] (Chapter 3, Appendix A)
k_{β_f}	=	Stiffness due to bulk modulus of a fluid column, [N/m] (Chapter 3, Appendix A)
m	=	Mass [kg] (Appendix A, B)
m_b	=	Vehicle sprung mass [kg] (Chapter 3, Appendix A, B)
m_{uf}	=	Unsprung mass, front [kg] (Chapter 3, Appendix A)
m_{ur}	=	Unsprung mass, rear [kg] (Chapter 3, Appendix A)
n_p	=	Polytropic constant (Chapter 3, Appendix B)
n	=	Number of variables in the optimisation problem (Chapter 5)
p	=	Pressure [Pa] (Chapter 2, 3, Appendix B)
p_{stat}	=	Static pressure [Pa] (Chapter 3, Appendix B)
t	=	Time [s] (Chapter 5)
t_{total}	=	Total simulation time [s] (Chapter 5)

$\overline{u}$ = Unit vector (Chapter 5)
v = Velocity [m/s] (Chapter 5)
x = Displacement [m] (Chapter 3, Appendix A, B)
x_{air} = Air volume displacement [m] (Chapter 3, Appendix B)
x_{nR} = Range of design variable n (Chapter 5)
x_{oil} = Oil volume displacement [m] (Chapter 3, Appendix B)
x_{stat} = Static displacement [m] (Chapter 3, Appendix A, B)
x_{susp} = Suspension displacement [m] (Chapter 3, Appendix B)
x_{tol} = Variable size convergence tolerance for Dynamic-Q (Chapter 5)
$x(t)$ = Random data set, cross-correlation function (Chapter 3)
$\dot{x}$ = Relative sliding velocity [m/s] (Chapter 3)
$\dot{x}_c$ = Critical velocity, modified Stribeck model [m/s] (Chapter 3)
$\dot{x}_s$ = Velocity based transient response parameter, Stribeck model [m/s] (Chapter 3)
$\overline{x}$ = Vector of cost function design variables (Chapter 5)
$\overline{x}^*$ = Vector of optimal design variables (Chapter 5)
$\overline{x}^*_{robust}$ = Vector of robust, optimal design variables (Chapter 5)
z_b = Sprung mass (body) vertical displacement [m] (Chapter 3, Appendix A)
z_{rlf} = Road displacement, left front [m] (Chapter 3, Appendix A)
z_{rlr} = Road displacement, left rear [m] (Chapter 3, Appendix A)
z_{rrf} = Road displacement, right front [m] (Chapter 3, Appendix A)
z_{rrr} = Road displacement, right front [m] (Chapter 3, Appendix A)
z_{uf} = Unsprung mass displacement, front [m] (Chapter 3, Appendix A)
z_{ur} = Unsprung mass displacement, rear [m] (Chapter 3, Appendix A)
$\dot{z}_b$ = Sprung mass (body) vertical velocity [m/s] (Chapter 3, Appendix A)
$\dot{z}_{rlf}$ = Road input velocity, left front [m/s] (Chapter 3, Appendix A)
$\dot{z}_{rlr}$ = Road input velocity, left rear [m/s] (Chapter 3, Appendix A)
$\dot{z}_{rrf}$ = Road input velocity, right front [m/s] (Chapter 3, Appendix A)
$\dot{z}_{rrr}$ = Road input velocity, right rear [m/s] (Chapter 3, Appendix A)
$\dot{z}_{uf}$ = Unsprung mass velocity, front [m/s] (Chapter 3, Appendix A)
$\dot{z}_{ur}$ = Unsprung mass velocity, rear [m/s] (Chapter 3, Appendix A)
$\ddot{z}$ = Acceleration in vertical direction [m/s^2] (Appendix A)
$\ddot{z}_b$ = Sprung mass (body) vertical acceleration [m/s^2] (Chapter 3, Appendix A)
$\ddot{z}_{uf}$ = Unsprung mass vertical acceleration, front [m/s^2] (Chapter 3, Appendix A)
$\ddot{z}_{ur}$ = Unsprung mass vertical acceleration, rear [m/s^2] (Chapter 3, Appendix A)

Roman capitals:

A = Area [m^2] (Appendix B), constant in stress-life power law [MPa] (Chapter 3)
B = Constant in stress-life power law (Chapter 3)
E = Young's modulus of elasticity [GPa] (Chapter 3)
F = Force [N] (Chapter 2, 3, Appendix B)
$F_{air}(x_{air})$ = Air spring force as a function of air volume [N] (Chapter 3, Appendix B)

F_{bump_lf} = Force in the left front bump stop [N] (Chapter 5)

F_{bump_lr} = Force in the left rear bump stop [N] (Chapter 5)

F_{bump_rf} = Force in the right front bump stop [N] (Chapter 5)

F_{bump_rr} = Force in the right rear bump stop [N] (Chapter 5)

F_C = Minimum coulomb friction force [N] (Chapter 3)

F_{damper} = Damper force [N] (Chapter 5)

F_f = Friction force [N] (Chapter 3)

F_s = Maximum static friction force [N] (Chapter 3)

F_{slf} = Suspension force, left front [N] (Chapter 3, Appendix A)

F_{slr} = Suspension force, left rear [N] (Chapter 3, Appendix A)

F_{srf} = Suspension force, right front [N] (Chapter 3, Appendix A)

F_{srr} = Suspension force, right rear [N] (Chapter 3, Appendix A)

F_{stat} = Static force [N] (Chapter 3, 5, Appendix A)

F_{tyre} = Tyre contact forces [N] (Chapter 5)

F_{tyre_lf} = Magnitude of the left front tyre contact forces [N] (Chapter 5)

F_{tyre_lr} = Magnitude of the left rear tyre contact forces [N] (Chapter 5)

F_{tyre_rf} = Magnitude of the right front tyre contact forces [N] (Chapter 5)

F_{tyre_rr} = Magnitude of the right rear tyre contact forces [N] (Chapter 5)

F_v = Viscous friction coefficient, Stribeck model [Ns/m] (Chapter 3)

F_z = Force in the Z-direction [N] (Appendix A)

I_{xx} = Sprung mass roll mass moment of inertia [kg.m^2] (Chapter 3, Appendix A)

I_{xx_a} = Unsprung mass roll mass moment of inertia [kg.m^2] (Chapter 3, Appendix A)

I_{yy} = Sprung mass pitch mass moment of inertia [kg.m^2] (Chapter 3, Appendix A)

M_{CG} = Moments about the centre of gravity [Nm] (Appendix A)

N = Number of equations (Chapter 3)

N_f = Number of cycles to failure (Chapter 3)

R = Design space radius of interest (Chapter 5)

$R_{fx}(\tau)$ = Cross-correlation function, (Chapter 3)

$\Re$ = Real domain (Chapter 5)

T = Observation period, cross-correlation function [s] (Chapter 3)

V = Volume [m^3] (Chapter 3, Appendix B)

V_{stat} = Static Volume [m^3] (Chapter 3, Appendix B)

X = Longitudinal direction (Chapter 3)

Y = Lateral direction (Chapter 3)

Z = Vertical direction (Chapter 3)

Greek symbols:

α = Sprung mass roll angle [radians] (Chapter 3, Appendix A)

$\dot{\alpha}$ = Sprung mass roll velocity [radians/s] (Chapter 3, Appendix A)

$\ddot{\alpha}$ = Sprung mass roll acceleration [radians/s^2] (Chapter 3, Appendix A)

β = Sprung mass pitch angle [radians] (Chapter 3, Appendix A)

β_f = Fluid bulk modulus [GPa] (Chapter 3, Appendix B)

$\dot{\beta}$ = Sprung mass pitch velocity [radians/s] (Chapter 3, Appendix A)

$\ddot{\beta}$ = Sprung mass pitch acceleration [radians/s^2] (Chapter 3, Appendix A)

δ = Transient response parameter, Stribeck model. (Chapter 3)

δ_i = Variable perturbation magnitude for gradient estimation (Chapter 5)

δ_{stlf} = Tyre static deflection, left front [m] (Chapter 3, Appendix A)

δ_{stlr} = Tyre static deflection, left rear [m] (Chapter 3, Appendix A)

δ_{strf} = Tyre static deflection, right front [m] (Chapter 3, Appendix A)

δ_{strr} = Tyre static deflection, right rear [m] (Chapter 3, Appendix A)

$\overline{\delta}$ = Variable perturbation for gradient estimation (Chapter 5)

ε = Strain in the suspension mounting [m/m] (Chapter 3)

ϕ = Rear unsprung mass roll angle [radians] (Chapter 3, Appendix A)

$\dot{\phi}$ = Rear unsprung mass roll velocity [radians/s] (Chapter 3, Appendix A)

$\ddot{\phi}$ = Rear unsprung mass roll acceleration [radians/s^2] (Chapter 3, Appendix A)

θ = Front unsprung mass roll angle [radians] (Chapter 3, Appendix A)

$\dot{\theta}$ = Front unsprung mass roll velocity [radians/s] (Chapter 3, Appendix A)

$\ddot{\theta}$ = Front unsprung mass roll acceleration [radians/s^2] (Chapter 3, Appendix A)

σ = Stress [MPa] (Chapter 3)

τ = Time delay, cross-correlation function [s] (Chapter 3)

ω = Small cap omega, natural frequency [radians/s] (Appendix B)

Vectors and matrices:

$\{f\}$ = Vector of external forces (Chapter 3, Appendix A)

$\{z\}$ = Coordinate vector (Chapter 3, Appendix A)

$\{\dot{z}\}$ = First time derivative of the coordinate vector (Chapter 3, Appendix A)

$\{\ddot{z}\}$ = Second time derivative of the coordinate vector (Chapter 3, Appendix A)

$\{Z\}$ = State vector (Chapter 3, Appendix A)

$\{\dot{Z}\}$ = First time derivative of the state vector (Chapter 3, Appendix A)

$[C]$ = Diagonal damping matrix (Chapter 3, Appendix A)

$[I]$ = The identity matrix (Chapter 3, Appendix A)

$[K]$ = Diagonal stiffness matrix (Chapter 3, Appendix A)

$[M]$ = Diagonal mass matrix (Chapter 3, Appendix A)

$[0]$ = The zero matrix (Chapter 3, Appendix A)

Chapter 1

1. Introduction and background

Manufacturers in the commercial vehicle sector have great pressure on them to increase the payload capacity of the vehicle without exceeding the legal vehicle mass restrictions. This can only be achieved by reducing the mass of the vehicle structure. The improvements must be accomplished without compromising the reliability of the vehicle's structural and other systems.

The goal of reducing vehicle mass drives designers to explore the use of exotic materials, novel construction techniques and innovative structural geometry. The reduction of input loads to the vehicle structure must also be keenly considered. These loads may be due to payload requirements or road loads, passed to the vehicle structure through the tyres, wheels and suspension system. Reducing the loads on the vehicle structure by selecting appropriate suspension characteristics will be the focus of this study.

1.1 Optimal suspension characteristics for fatigue life

Literature on the selection of suspension characteristics for ride comfort or handling can be found for virtually any type of vehicle. The amount of literature available on the selection of suspension characteristics for fatigue damage to the vehicle structure stands in sharp contrast to this. The focus seems to have been on the prediction of road loads on vehicle structures and subsequently fatigue life of these structures.

Wannenburg (2007) notes that the inputs required for fatigue life prediction of any structure are the structural geometry, material properties and input loading. He states that the structural geometry and material properties are generally well known and the greatest concern therefore involves determining the input loading. Conle and Chu (1997), Li and Sandu (2007), as well as Kat (2009) corroborate Wannenburg's view and discuss some of the challenges in creating useful vehicle simulation models for fatigue life prediction. They are in agreement that tyre models are a major challenge in accurately predicting input loads to the vehicle structure. Li and Sandu (2007) use the polynomial chaos method for dealing with uncertainty in modelling parameters.

Kat (2009) also presents a thorough discussion on the importance of validating simulation model using experimental results. He stresses the need for validating simulation models for the parameters that they are intended to predict. Kat warns that good model correlation for accelerations alone may not imply good correlation for other predicted parameters, such as suspension forces.

Letherwood and Gunter (2001) provide a thorough account of the vehicle dynamic modelling used by the United States Army's Tank-automotive and Armaments Research, Development and Engineering Centre. The vehicle models they discuss are used for investigations and vehicle development throughout the vehicle's life cycle. The studies conducted using their models include investigations into vehicle stability,

safety and vehicle integrity. The models are validated experimentally to ensure that predictions of vehicle dynamic response correlate with reality. Special mention is made of the challenges faced in the modelling of off-road vehicles due to the presence of high non-linear response characteristics caused by large suspension, tyre and soil deflections.

Medepalli and Rao (2000) conducted a study into the prediction of road loads for fatigue design. An ADAMS (Automatic Dynamic Analysis of Mechanical Systems) multi-body dynamics model of the test vehicle was constructed and validated against experimental data for the vehicle dynamics and suspension loads. The model could be used to predict the vehicle dynamics by allowing it to traverse different terrains. The model was employed in a sensitivity analysis on the effects of unsprung mass, suspension geometry and spring rate on the loads on the vehicle body. The sensitivity analysis took the form of a comparison between a rigid axle and fully independent suspension system. One should thus be careful to generalise the results obtained.

Conle and Mousseau (1991) also employed a DADS (Dynamic Analysis and Design System) multi-body dynamics model to predict stresses in vehicle components. The model used measured wheel load histories as inputs, rather than a tyre model traversing a road profile. The model calculated loads on a suspension control arm, which were used as inputs to a finite element model of the control arm. The stress histories calculated in the finite element model were used for damage estimation of the component. Conle and Mousseau stress the importance of experimental correlation of the prediction process (which is lacking in their study) and note the necessity of considering computational expense in the prediction process.

Mrzyglod and Zielinski (2007) used a similar approach for prediction of road loads from multi-body dynamics models and using the loads to compute stress histories. They employed their models in the mathematical optimisation of a car's lower suspension arm. Their optimisation used component mass as objective function and was subject to a fatigue life constraint. The models employed were not validated.

Contrary to most authors, who consider the prediction of vehicle fatigue damage, Zeiler and Barkey (2001) conducted an analytical study into the effect of suspension characteristic on fatigue damage. They used a linear pitch bounce vehicle model (not validated) combined with the Forman crack growth law to calculate the sensitivities of fatigue damage to spring and damper characteristics. Their results indicate a decreasing fatigue life when spring and damper rates are increased. No suggestion is made as to what the optimal characteristics may be. Zeiler (2002) extended the work to include a simple, formal mathematical optimisation with the model developed by Zeiler and Barkey (2001).

A wider search of the literature showed that some work has been done in other fields on the topic of minimising a system's fatigue damage by selecting appropriate stiffness and damping characteristics. Tuned mass dampers are an example of such systems. Tuned mass dampers are used to dampen vibrating structures within a narrow frequency band. They are often employed in civil engineering applications to reduce the motion of high rise buildings under the influence of wind or earthquake loads. Li and Hu (2002) explored the design of a tuned mass damper for minimizing fatigue damage to a primary structure. They used a simple two degree-of-freedom

model in their mathematical optimisation. The cost function employed was fatigue damage to the primary structure due to random excitation representing wind loading. The design variables selected were the stiffness and damping rates of the tuned mass damper. They were able to achieve a 75% improvement in the fatigue damage to the structure by introducing their tuned mass damper. The reduction in loads on the primary structure was however only effective in a very narrow frequency band.

Literature on the selection of suspension characteristics for reducing damage to the vehicle structure could not be found. The current study aims to address this shortfall in the literature.

1.2 The laden versus unladen compromise

The ideal vehicle suspension system would provide near optimal isolation at different speeds over varying road profiles for widely varying payloads. A passive system cannot function optimally under these varying conditions and will inevitably result in a compromised solution. The high changes in the unsprung mass of heavy vehicles present even greater challenges in selecting suitable suspension characteristics. A laden vs. unladen compromise exists, which is in many ways analogous to the ride comfort vs. handling compromise so often investigated in vehicle dynamics research.

The design of a passive suspension system will inevitably be a compromised solution for greatly varying payload states. Large variations in payload include the further complication of changing vehicle centre of mass position. Von Glasner et al. (1993) discuss these problems and state that a single state suspension system will necessarily have compromised performance when operating at payload conditions other than the design state. The compromise may involve performance factors as diverse as ride comfort, handling and braking capability (which implies safety).

1.3 Possible solutions to the laden versus unladen compromise

The promise of ideal behaviour over a wide range of operating conditions makes active suspension systems attractive for solving the challenges in vehicle suspension design due to varying operational conditions. Vaughan et al. (2003) conducted a theoretical study, using a linear four degree of freedom (DOF) pitch-plane model, into the active control of a vehicle suspension system in order to counter the effects of change in vehicle payload. They also consider the effects of the change in the vehicle centre of mass position. It was concluded that, even though the addition of payload and the change in vehicle centre of mass positions greatly affect the dynamic behavior of the vehicle, the use of active suspension control can moderate these changes. It is mentioned that the high power consumption of active suspensions may make this solution infeasible.

Semi-active or adaptive suspension systems offer many of the advantages of active systems at a fraction of the energy requirements and complexity (Von Glasner et al, 1993). These systems have consequently been suggested to remedy the laden vs. unladen compromise in heavy vehicles. Von Glasner et. al. performed a simulation study on the effects of variable damping on heavy vehicle ride comfort and stopping distances. A model of a light truck was constructed. The mass of the truck was

7500kg when fully laden and 4200kg unladen. The model included the Anti-Lock Braking System (ABS) algorithm of the real vehicle. The vehicle stopping distances from 80km/h for both the laden and unladen case were very similar using the standard vehicle suspension system. An improvement of the vehicle ride comfort of 19% was achieved in the unladen case by lowering the spring and damping rates proportionally to the vehicle mass. This was however achieved at a 5.6% increase in stopping distance. The stopping distance could however be decreased by 4.2% of the baseline stopping distance by switching the dampers to twice the baseline rate, during a braking maneuver.

Bode et al. (1997), conducted tests on a rather elaborate quarter vehicle experimental setup in order to determine the influence of payload on the vertical body acceleration and vertical wheel load fluctuation. The experimental setup consisted of a frame (representing the quarter vehicle body) supported in the centre of a truck-trailer by low friction vertical guide-ways. The mass of the frame could be varied between 1.1 and 4.0 tons and was suspended on a double wishbone front suspension system from a commercial truck. The suspension system consisted of a pneumatic spring and variable damper, which could be switched manually between various passive characteristics or actively controlled. The truck-trailer could then be towed over different obstacles or road profiles in order to measure the response of the quarter vehicle model to different road inputs. The study showed that for both the fully laden and unladen cases large improvements to both the vertical acceleration and vertical wheel load fluctuation could be obtained by employing appropriate settings of the variable damper. The influence of adapting the spring rate to load conditions was not examined.

1.4 The development of the $4S_4$ suspension system

The Four State Semi-Active Suspension System or $4S_4$ was originally developed by Els (2006) as a possible solution to the ride comfort vs. handling compromise. The $4S_4$ system is a hydro-pneumatic suspension system, which allows switching of both the spring and damping rates between different passive states. The system can be toggled between two damping and two spring states, thus leading to four passive states. A schematic of one of the $4S_4$ units is shown in Figure 1-1.

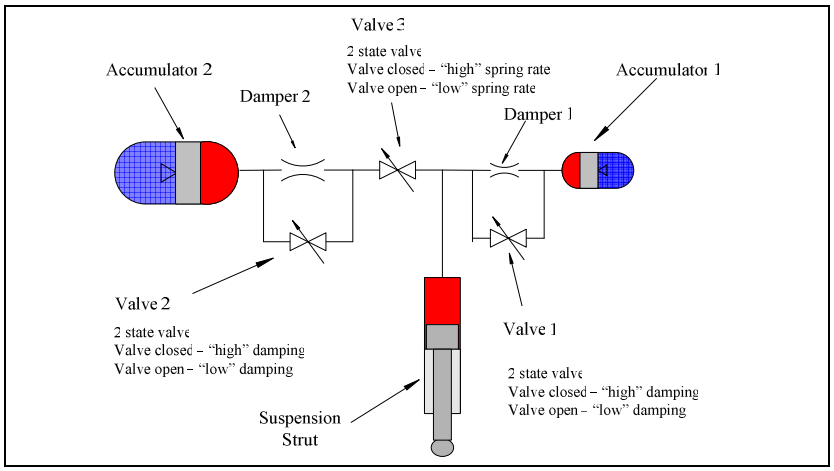

Figure 1-1: $4S_4$ schematic.

The $4S_4$ consists of a suspension strut, two gas accumulators, two damper packs and three solenoid valves. The different spring and damper rates are achieved by opening or closing solenoid valves to regulate the flow paths of the working fluid in the system. The low spring rate is achieved by opening valve 3 in Figure 1-1 and thereby activating both accumulators. The high spring rate is achieved by closing valve 3 and thereby limiting the gas compression volume to accumulator 1. The low damping rate is achieved by opening valves 1 and 2 and allowing the oil flow to bypass the damper packs. Finally, high damping rates are achieved by closing valves 1 and 2 and forcing all the oil flow through the two damper packs.

The ability of the $4S_4$ to switch between vastly varying spring and damper states is its right of existence. The system in its current implementation solves the ride comfort handling compromise for a Land Rover Defender 110. The vehicle's mass properties were characterised in a laboratory environment (Uys, 2006) and subsequently an ADAMS model of the vehicle was constructed. The model was validated against field test results (Els, 2006). Thoresson (2003) used the vehicle model in mathematical optimisation to find the optimal suspension characteristics for ride comfort and handling for the vehicle.

Thoresson's results confirmed the commonly accepted notion that the spring and damper rates for ride comfort and handling lie at opposite sides of the design space. The optimum characteristics determined by Thoresson were subsequently implemented on the Land Rover in the form of the $4S_4$. The high damping, high spring rate state of the $4S_4$ has been set up for the handling optimum characteristics, while the low damping, low spring rate state has been configured to the ride comfort optimum. The system can be toggled between the ride and handling at any time and switching is virtually undetectable by vehicle occupants. The change in vehicle suspension characteristics after switching is, however highly apparent to both driver and passengers.

It is suggested that the $4S_4$ system should be investigated as a possible solution to the laden versus unladen compromise which is faced by heavy vehicles.

1.5 Problem statement and scope of work

The problem which has to be solved is that of finding the spring and damper characteristics required for optimal vehicle structural life of a Land Rover Defender 110 sports utility vehicle. The optimal characteristics are to be found for different vehicle payload conditions.

These results will be obtained through a process of mathematical modelling and optimisation. A mathematical model which is simple enough to be computationally efficient, yet complex enough to representatively predict the vertical dynamics of the vehicle with some confidence, is thus required. The mathematical model should be validated against field tests performed on the actual vehicle.

Mathematical optimisation will be conducted using existing mathematical optimization methods where possible, with a cost function which is representative of damage to the vehicle structure due to road loads.

Once the optimal characteristics have been determined, their sensitivity to different operating conditions will be determined. Finally, recommendations will be made on the suitability of the $4S_4$ system for alleviating the laden versus unladen compromise in heavy vehicles from a structural life perspective.

1.6 Dissertation overview

The remainder of this document describes the work performed on finding the optimal vehicle suspension characteristics for increased structural life.

Chapter 2 describes the field test work conducted on the test vehicle used in the study. The data obtained from the experimental work in chapter 2 was used in the validation of the mathematical model developed in chapter 3.

Chapter 3, in addition to describing the mathematical modelling, devotes some pages to the experimental characterization of some of the suspension system components.

Chapter 4 presents the results of the experimental work as well as the simulation results from the mathematical model. The results are compared for correlation purposes and the validity and limitations of the constructed mathematical model is discussed.

The optimisation problem, which this dissertation set out to investigate, is finally formulated in chapter 5. The choice of cost function, constraints and optimisation are laid out in detail in this chapter.

Chapter 6 summarises the conclusions drawn from the study and makes recommendations for future investigations.

Chapter 2

2. Experimental work

The purpose of the experimental work discussed in this chapter was to gather data, which could be used for the validation of the mathematical model of the test vehicle. The vehicle used for this study was a Land Rover Defender 110. It has been used extensively in the research of the Dynamic Systems Group at the University of Pretoria. A great deal of instrumentation is thus permanently fitted to vehicle. This section discusses the parameters required for model validation as well as the instrumentation on (or added to) the experimental vehicle in order to measure these parameters. The field tests conducted are also discussed.

2.1 Instrumentation of the experimental vehicle

The model is required to predict the dynamics of the test vehicle with some confidence. The measurements for the validation of the model should thus include the vehicle body motions as well as the motions of the vehicle wheels and suspension. The parameters measured for validation of the vehicle dynamics are given in Table 2-1.

Table 2-1: Parameters measured for validation of vehicle dynamics.

Parameter	Transducer
Vehicle speed	e-DAQ GPS
Vehicle speed	VBOX GPS
Vehicle speed	Proximity probe measuring drive shaft speed
Left rear (LR) vertical acceleration	Accelerometer
Right rear (RR) vertical acceleration	Accelerometer
Right front (RF) vertical acceleration	Accelerometer
Roll velocity	Solid state gyroscope
Pitch velocity	Solid state gyroscope
Left front (LF) suspension displacement	Linear Displacement Transducer (LDT)
RF suspension displacement	LDT
LR suspension displacement	LDT
RR suspension displacement	LDT

Vehicle speed was measured by three different methods to ensure redundancy on this critical parameter. The body motions were captured by measuring vertical acceleration in three positions as well as pitch and roll velocities. The vehicle suspension and wheel motions are measured by Linear Displacement Transducers (LDT) integrated into each $4S_4$ suspension strut. The measurement positions for the transducers described are given in Figure 2-1.

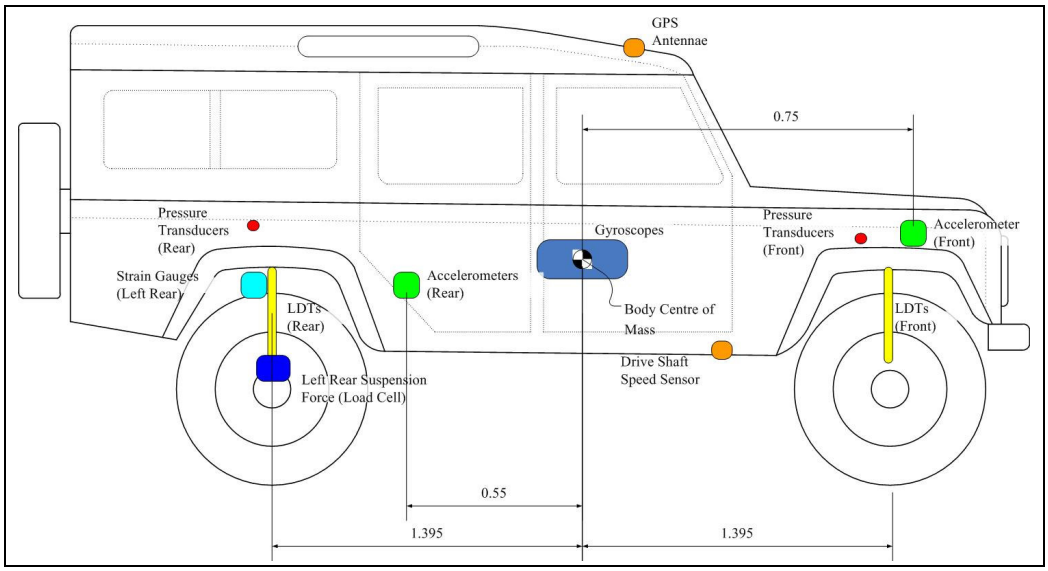

Figure 2-1: Schematic of measurement positions.

It is further required that the model be capable of predicting "structural damage" and this implies the measurement of suspension forces, which are the inputs to the vehicle structure and thus determine the structural damage. Strains on the vehicle chassis (near the left rear suspension attachment points) were measured for the purpose of validating the "structural damage" prediction. The parameters which were recorded for this purpose are listed in Table 2-2.

Table 2-2: Parameters measured for the validation of structural damage.

Parameter	Transducer
LR suspension pressure	Pressure transducer
RR suspension pressure	Pressure transducer
RF suspension pressure	Pressure transducer
LF suspension pressure	Pressure transducer
LR suspension force	Tension-compression load cell
6 x Strain	Foil strain gauge

The $4S_4$ struts are each fitted with a pressure transducer that measures the oil pressure within the strut. The measured pressures were recorded and later converted to forces using their area of application ($F = pA$). The forces neglect the friction and inertial forces and thus only represent the forces due to the spring and damper characteristic of the $4S_4$ strut. The inertial forces are considered negligible in this case. This is illustrated in the free-body diagram of the $4S_4$ suspension strut given in Figure 2-2. The illustration shows that the friction force is unobservable using the pressure transducer.

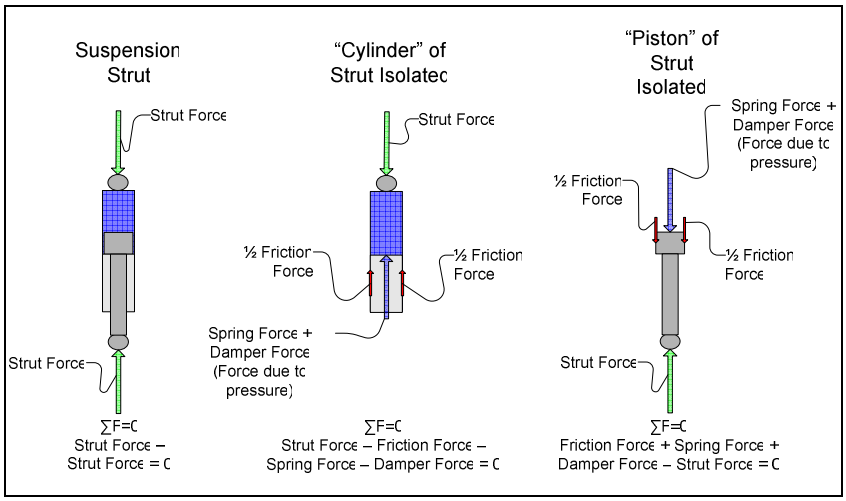

Figure 2-2: Free-body diagram of suspension strut.

The fact that the friction force is unobservable by the pressure transducer, combined with uncertainty concerning the frequency response characteristics of the pressure transducers, lead to the integration of an additional tension-compression load cell in line with the $4S_4$ strut. The load cell was installed in the left rear position where the strain gauges were applied to the suspension mounting. The additional load cell measures the total force transferred through the strut directly. The load cell and the strain gauges on the left rear suspension mounting are shown in Figure 2-3.

Preliminary test runs with the load cell installed revealed the presence of high friction in the $4S_4$ struts. In an attempt to alleviate the high friction, all four suspension units were removed and partially reconditioned. This included cleaning and re-lubricating the spherical joints used for mounting the $4S_4$ struts, as well as replacing any worn seals. There was however little improvement in the friction after all this effort. It was concluded that the friction was inherent to the system due to the hydraulic seal configuration employed. The friction would thus have to be included in the vehicle model.

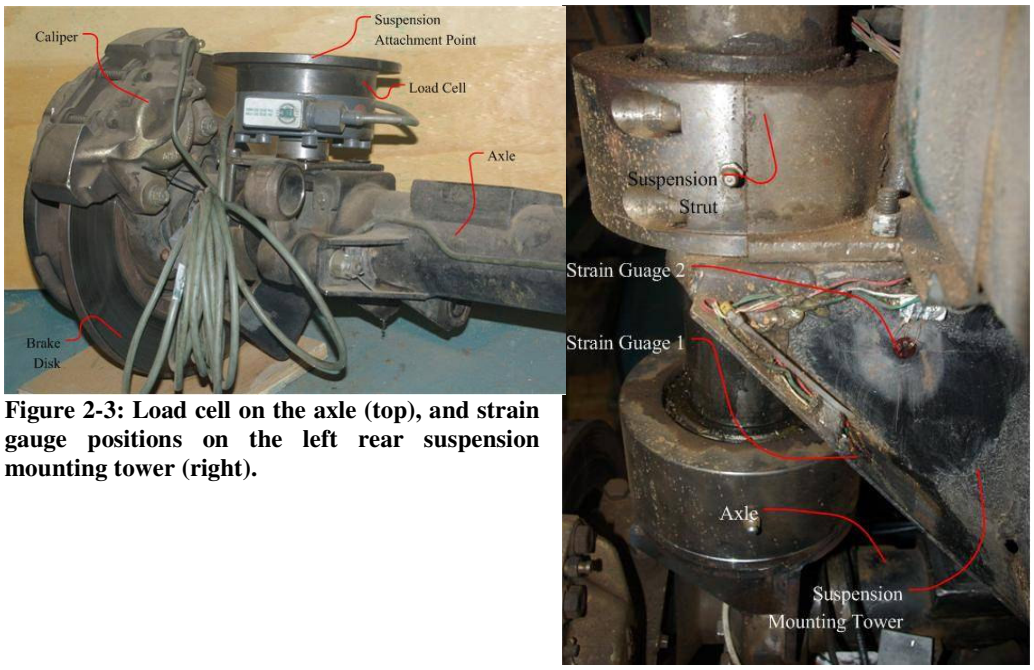

Figure 2-3: Load cell on the axle (top), and strain gauge positions on the left rear suspension mounting tower (right).

Strains were measured using two tri-axial rosette strain gauges. The strain measurements were necessary in order to obtain load-time histories for structural damage estimation. The strain gauges were applied in two different positions on the suspension mounting tower on the vehicle chassis. The positions of the strain gauges are shown in Figure 2-3.

The suspension mounting tower was chosen for the strain measurements, since it is in the direct path of force transmission between vehicle suspension and chassis. Additionally the mounting is a stiff structure which would have limited structural response at the low frequencies dominant in the vehicle dynamics. The suspension system would provide some additional isolation from high frequencies. Neglecting structural response would significantly simplify the mathematical modelling.

All data channels were logged on an e-DAQ Lite data acquisition computer. A sampling frequency of 1000 Hz was employed on all runs. The e-DAQ's built in linear roll-off anti-aliasing filter was employed. A roll-off frequency of 333 Hz was used. The data was filtered with an additional 25 Hz ideal low-pass filter in the post processing environment. This additional filtering was done as the frequencies of interest in vehicle dynamics application are generally below 20 Hz.

2.2 Vertical dynamics model validation tests

The evaluation of the vertical dynamics of a vehicle is generally conducted over discrete obstacles or what Letherwood and Gunter (2001) call "RMS cross country courses". The latter is generally referred to as a random road surface. Discrete obstacles reveal the transient response of the vehicle system. Random road surfaces excite a wide band of frequencies simultaneously and give a good representation of the vehicle vertical response. Complete and thorough vertical dynamics model validation thus requires test data for both discrete obstacles and random road surfaces. The vehicle validation tests were conducted at the Gerotek Testing Facilities near Pretoria, South Africa (Gerotek, 2009).

2.2.1 Discrete obstacle tests

The type of discrete obstacle chosen for the model validation testing is a trapezoidal speed bump. This type of obstacle is recommended by Bode et al. (1997) and Letherwood and Gunter (2001). An example of a pair of trapezoidal bumps is shown in Figure 2-4. Obstacles of this type are well suited to vertical vehicle model validation as they predominantly excite the vehicle suspension dynamics without excessive participation of the tyre dynamics. The relatively low frequency content of a trapezoidal bump can be credited for this behaviour. Data from these tests is best compared in the time domain.

Figure 2-4: Trapezoidal bumps.

The discrete obstacle testing was conducted on the Tank Track at Gerotek. Three test layouts were set up, using two different heights of trapezoidal bumps, on which to conduct the testing. The dimensions (in millimeters) of the trapezoidal bumps used are given in Figure 2-5.

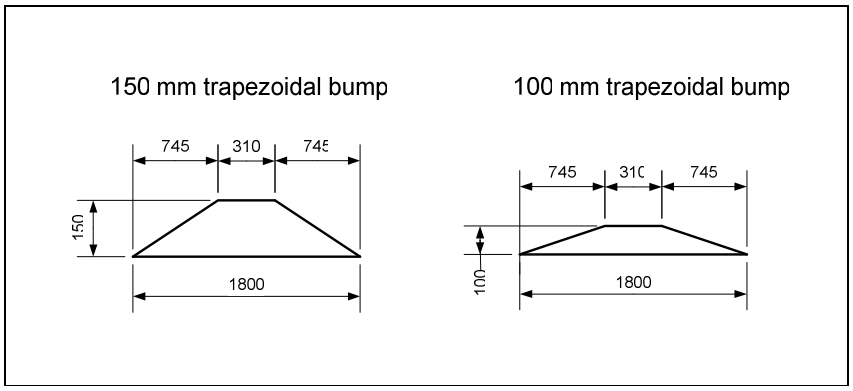

Figure 2-5: Dimensions of the trapezoidal bumps employed.

The first layout consisted of two sets of symmetrically laid out trapezoidal bumps, first a 150 mm and then a lower 100 mm bump. This first layout was used for the characterisation of the vehicle dynamics with the $4S_4$ set to the "ride comfort" or "soft" setting. The track layout is shown schematically in Figure 2-6.

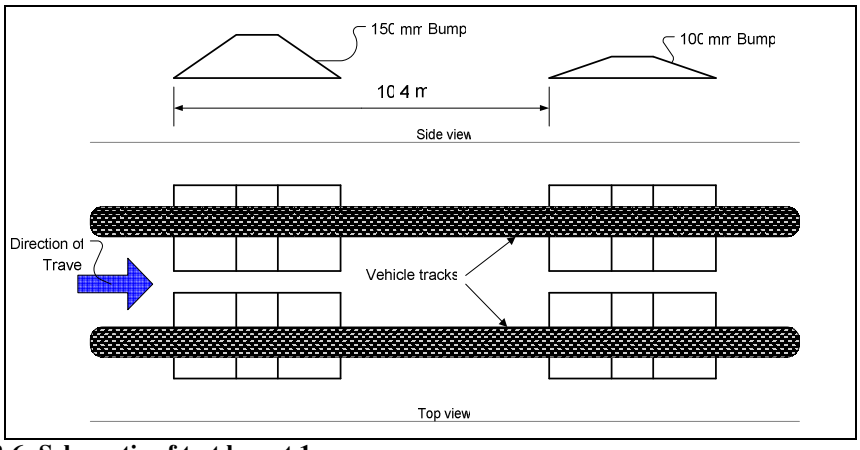

Figure 2-6: Schematic of test layout 1.

The harshness of the "handling" or "hard" setting of the $4S_4$, necessitates the use of a less severe test layout. The 100 mm bump would therefore be excluded from the layout for the test with the suspension set to hard. The second layout thus consisted of the single pair of the 100 mm trapezoidal bumps as indicated in Figure 2-7.

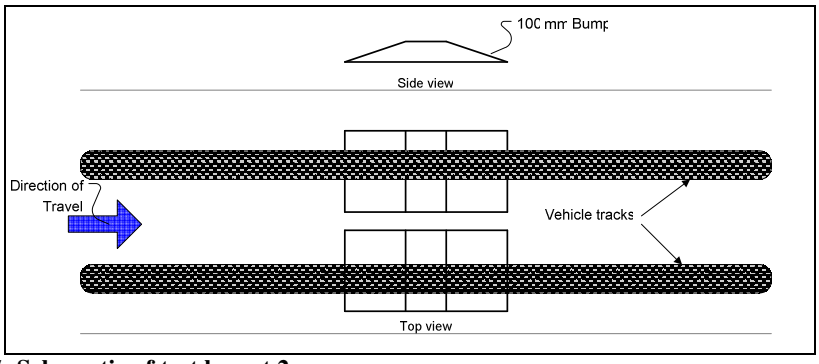

Figure 2-7: Schematic of test layout 2.

The first two layouts would mainly excite the pitch plane motions of the vehicle's vertical dynamics. The layout was thus aimed at exciting the roll plain vertical dynamics of the vehicle. The large and smaller bumps were again positioned as in the first layout, but in this case one of each of the pairs of bumps was removed. The bumps were staggered to excite the roll dynamics in both directions. This third trapezoidal bump configuration is represented in Figure 2-8.

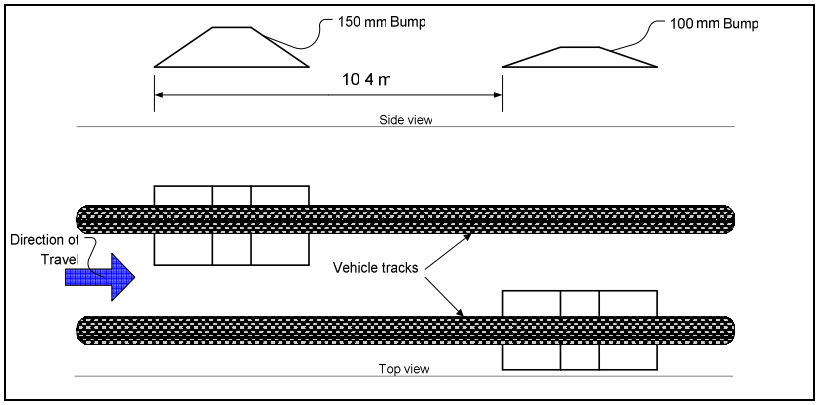

Figure 2-8: Schematic of test layout 3.

2.2.2 Random road surface tests

The Belgian paving at Gerotek was used for evaluating the vertical dynamics of the vehicle over random rough terrain. Data recorded on this type of terrain is best compared in the frequency domain. The Belgian paving track is pictured in Figure 2-9. It should be made clear, that although the track is referred to as "the Belgian paving", it is of much higher roughness, than the Belgian paving surfaces usually referred to in vehicle dynamics literature.

Figure 2-9: The Belgian paving at Gerotek.

The Belgian paving at Gerotek was profiled by Bekker (2008) and its roughness profile is presented in Figure 2-10 in the form of a Displacement Spectral Density (DSD) plot. The vertical lines in the figure indicate the lower and upper reporting limits as recommended by the ISO 8608 standard (International Organisation for Standards, 1995). These frequency limits indicate the minimum and maximum spatial frequencies which influence an off-road vehicle's general vertical dynamics. The DSD plot in the figure follows an even linear trend over most of the frequency range of interest, indicating a predominantly random profile. There is however a sharp peak at a spatial frequency of 6 cycles/meter. This equates to a wavelength of 167 mm which corresponds to the size of the cobbles on the Belgian paving track. This should not be cause for concern, but may add some high frequency energy to the system.

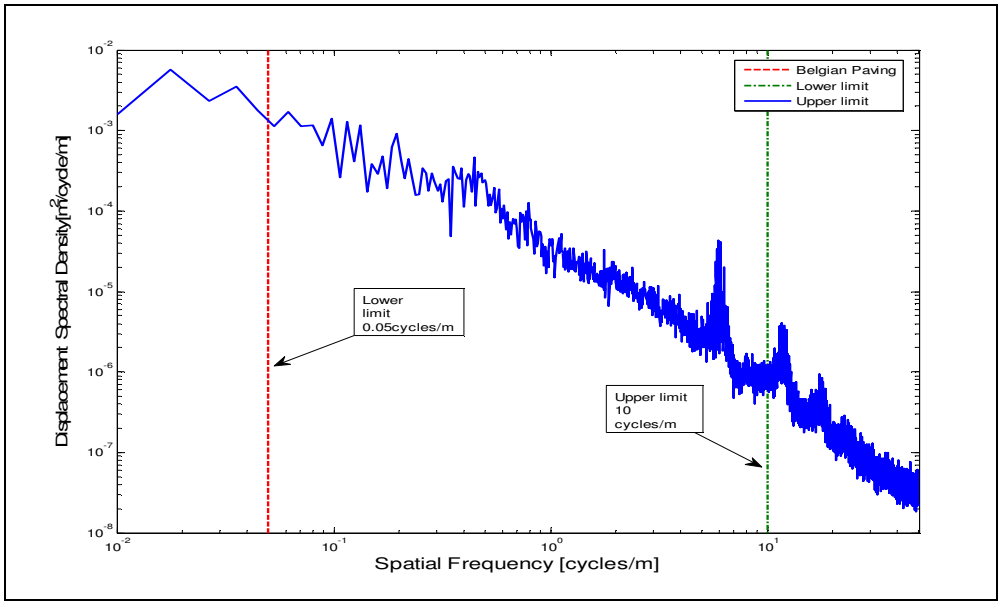

Figure 2-10: DSD of Belgian paving. Modified from Bekker (2008).

2.3 Experimental procedure

The discrete obstacle and Belgian paving tracks were all traversed at different speeds and with different suspension settings. The speed was kept as constant as possible by driving the diesel powered Land Rover against the engine governor in gear. Several runs were performed on each track with each of the chosen speed and suspension setting combinations. The multiple runs performed allow the repeatability of the tests to be gauged.

It was furthermore ensured that the tests were carried out in a fairly random order. This procedure was aimed at reducing the effects of any systematic error in the measurements. The test runs performed are listed in Table 2-3 for each layout. The speed and suspension settings as well as the run numbers used are also listed.

Table 2-3: List of tests performed.

Run no.	Test description	Suspension setting	Speed [km/h]
01, 05, 10	Belgian paving	ride comfort	±14.5 (low range, 1st gear)
02, 06, 11	Belgian paving	ride comfort	±54 (low range, 4th gear)
03, 08, 12	Belgian paving	handling	±14.5 (low range, 1st gear)
04, 09, 13	Belgian paving	handling	±54 (low range, 4th gear)
16, 19, 20, 21, 22	Bump course, layout 1	ride comfort	±14.5 (low range, 1st gear)
23, 24, 25, 26, 27	Bump course, layout 2	handling	±14.5 (low range, 1st gear)
28, 30, 32, 33	Bump course, layout 3	ride comfort	±14.5 (low range, 1st gear)
29, 31, 34	Bump course, layout 3	handling	±14.5 (low range, 1st gear)

The results of these tests will be presented as required in Chapter 4 of this document. Selected results will however be presented in the current chapter in order to show the extent of the repeatability which was obtained during testing. Data will be presented for runs over the large trapezoidal bump with the $4S_4$ set to ride comfort, runs over the small trapezoidal bump with the $4S_4$ set to handling and finally runs over the Belgian paving.

The data for three runs over the large trapezoidal bump with the $4S_4$ set to the ride comfort setting has been overlaid for comparison in terms of the repeatability. Figure 2-11 shows the left rear and right front suspension displacements for the three runs considered. The data was filtered using a 25 Hz ideal low-pass filter in order to remove excessive noise. Additionally, the mean value of each data set was subtracted in order to compensate for small zero offset changes and improve the ease of comparison. The data sets are almost indistinguishable, showing excellent repeatability. The quality of the results (by inspection) is similar to that obtained by Bode et al. (1997).

The measured pitch velocity for the three runs, presented in Figure 2-12, indicates equally good repeatability. The suspension forces, which are so critical for this study, also show good repeatability (Figure 2-13). The calculated force was determined by multiplying the pressure with the area of application (p. 2-2). The reader should note that the suspension force data has not been corrected for zero offset.

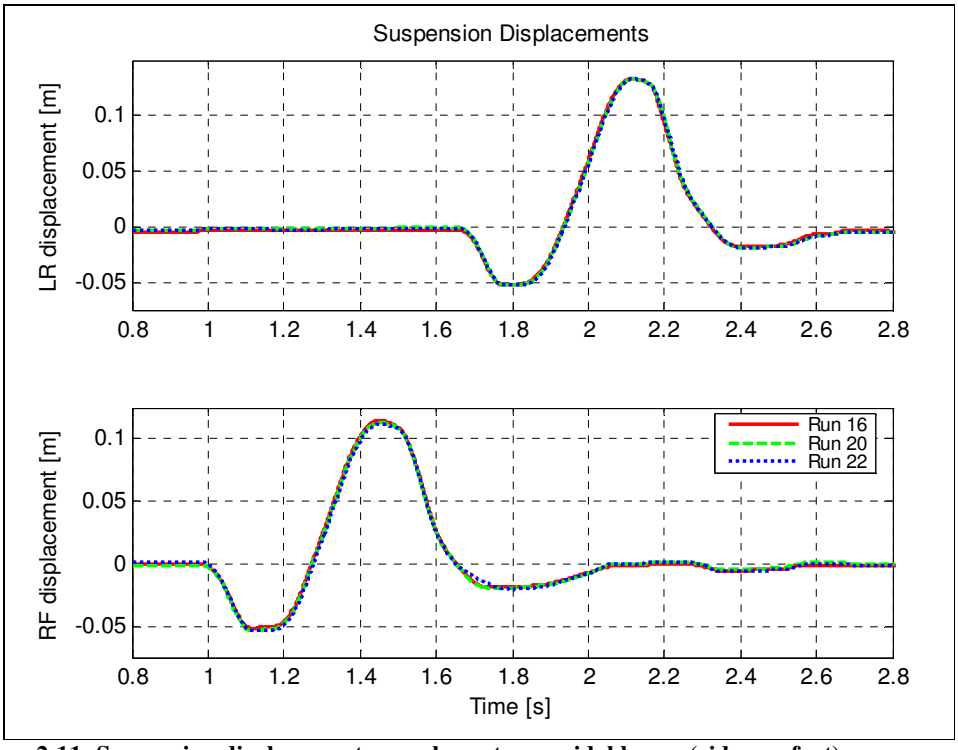

Figure 2-11: Suspension displacements over large trapezoidal bump (ride comfort).

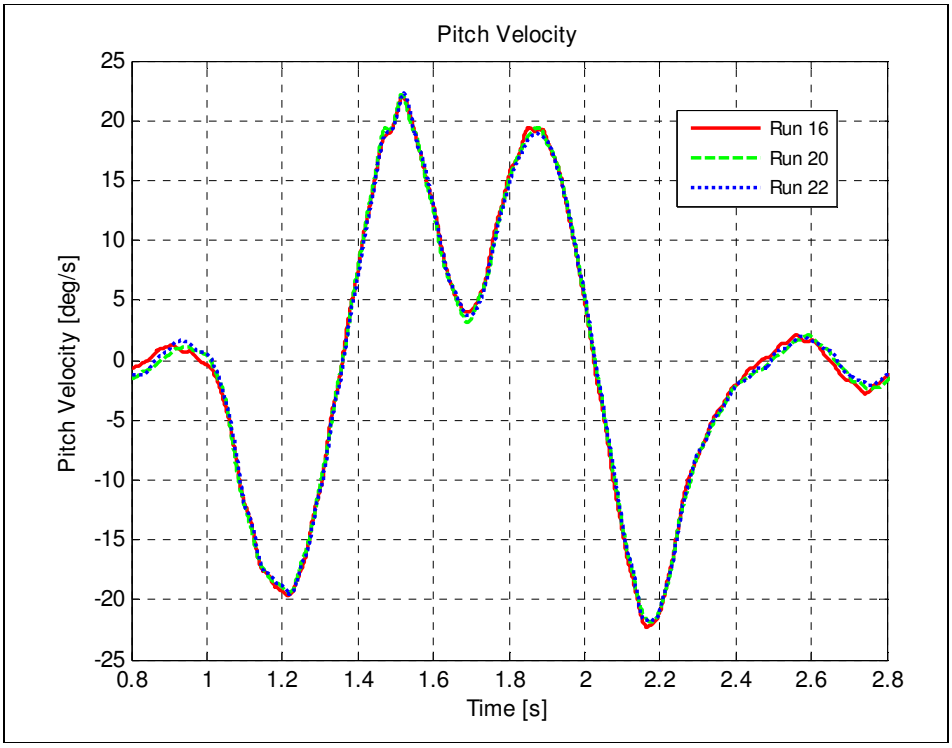

Figure 2-12: Body pitch velocity over large trapezoidal bump (ride comfort).

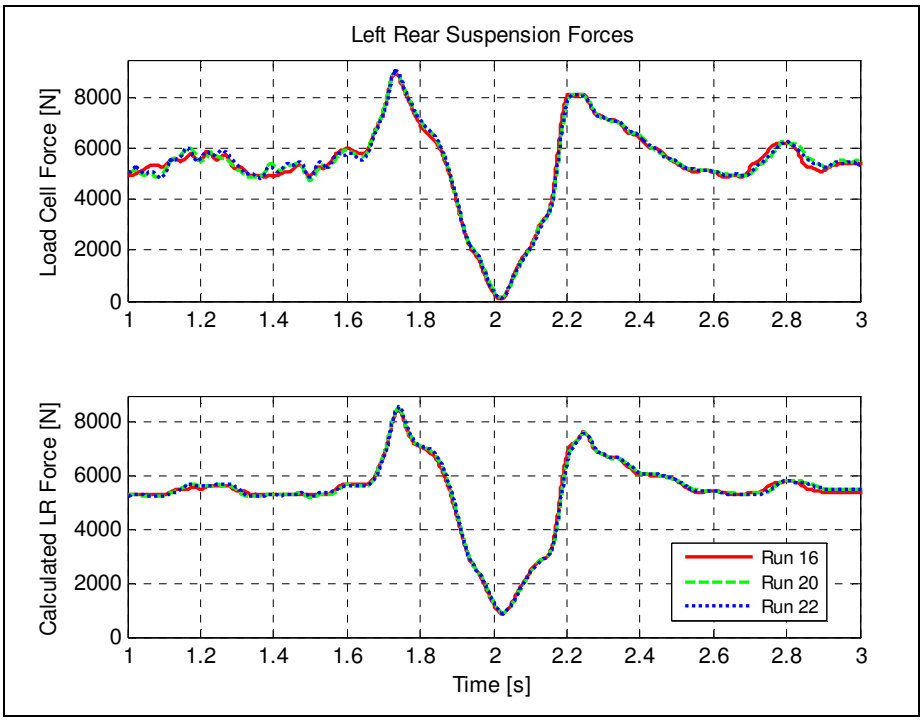

Figure 2-13: LR suspension forces over large trapezoidal bump (ride comfort).

The vehicle body vertical accelerations, which are the parameter often used for vehicle model correlation, are shown in Figure 2-14. The data shows excellent repeatability.

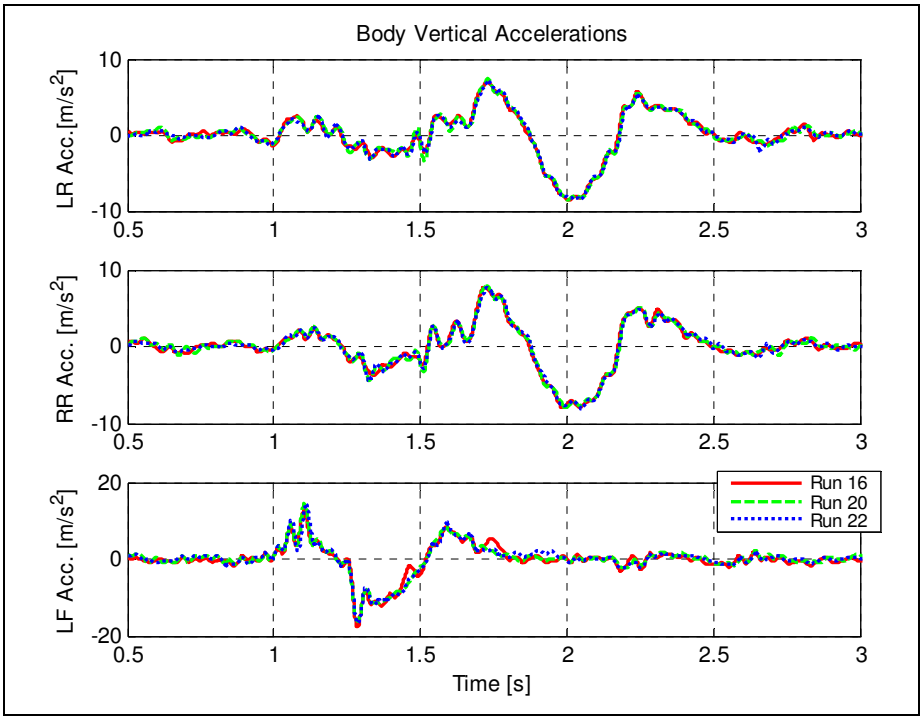

Figure 2-14: Body vertical accelerations over large trapezoidal bump (ride comfort).

The data for three runs over the small trapezoidal bump with the $4S_4$ set to the handling mode is now presented. The suspension displacements recorded for the three runs considered are plotted in Figure 2-15. The data again shows good repeatability, although not to the extent observed with the softer suspension setting. It should be noted that the magnitude of the suspension displacements for these runs is much smaller and the variation on the peak values, thus appear much larger. The variation is in the order of 2-3mm.

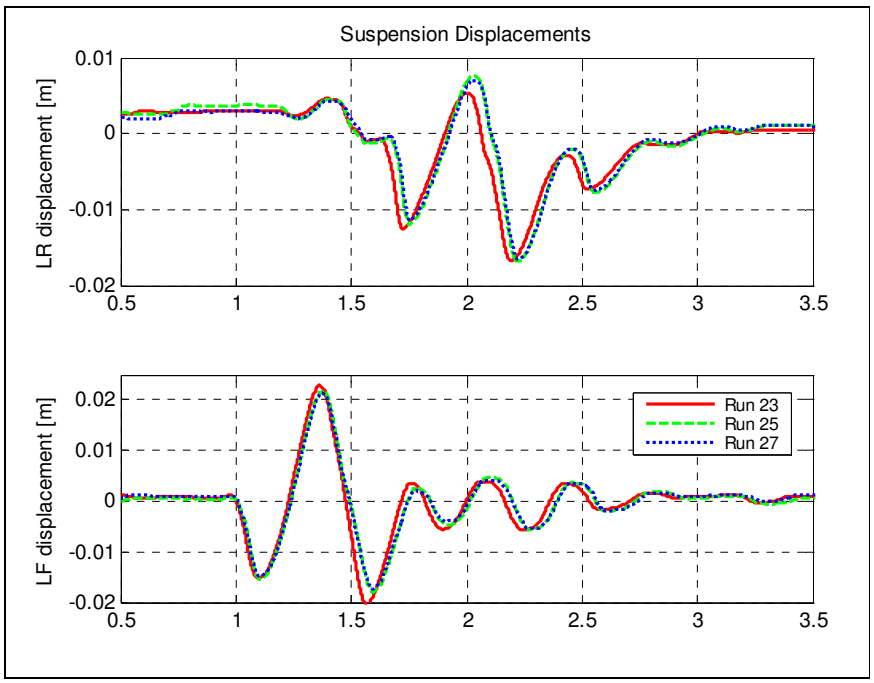

Figure 2-15: Suspension displacements over small trapezoidal bump (handling).

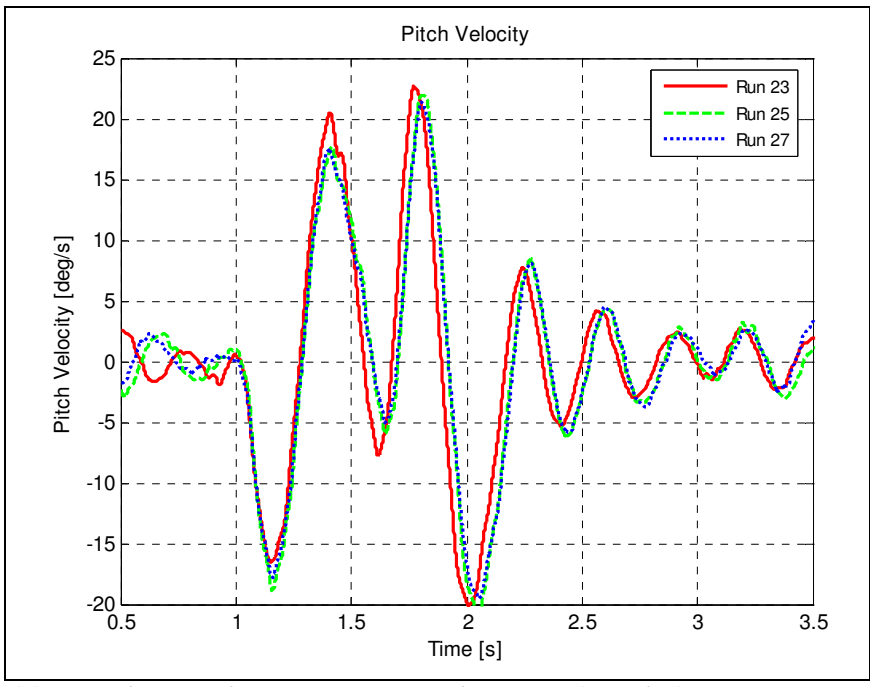

Figure 2-16: Body pitch velocity over small trapezoidal bump (handling).

The vehicle body pitch velocity, suspension forces and body vertical accelerations are presented in Figure 2-16, Figure 2-17 and Figure 2-18 respectively. These datasets all show good repeatability between different runs.

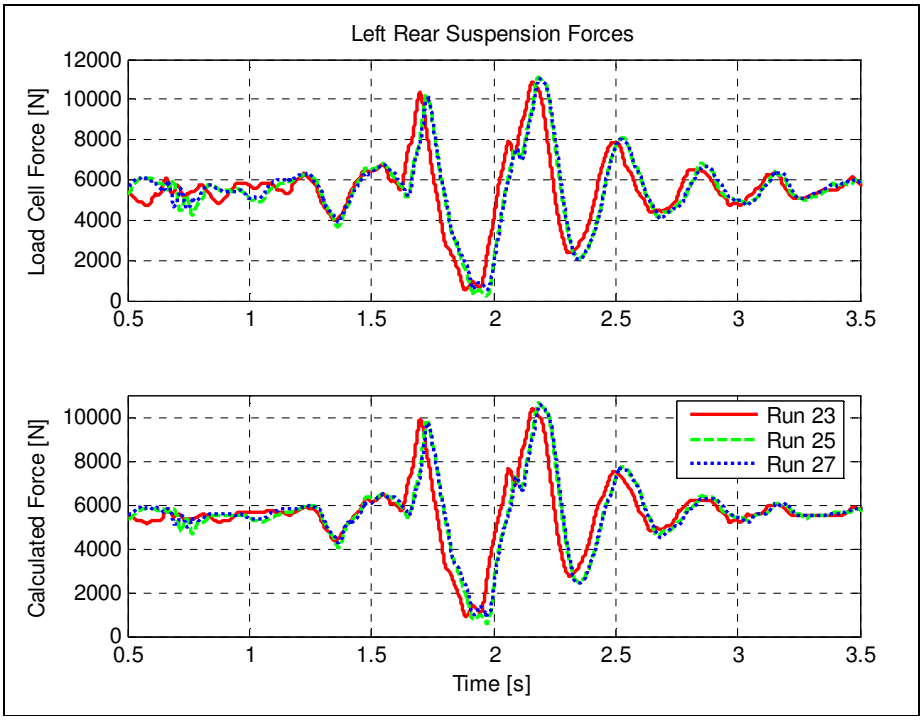

Figure 2-17: LR Suspension forces over small trapezoidal bump (handling).

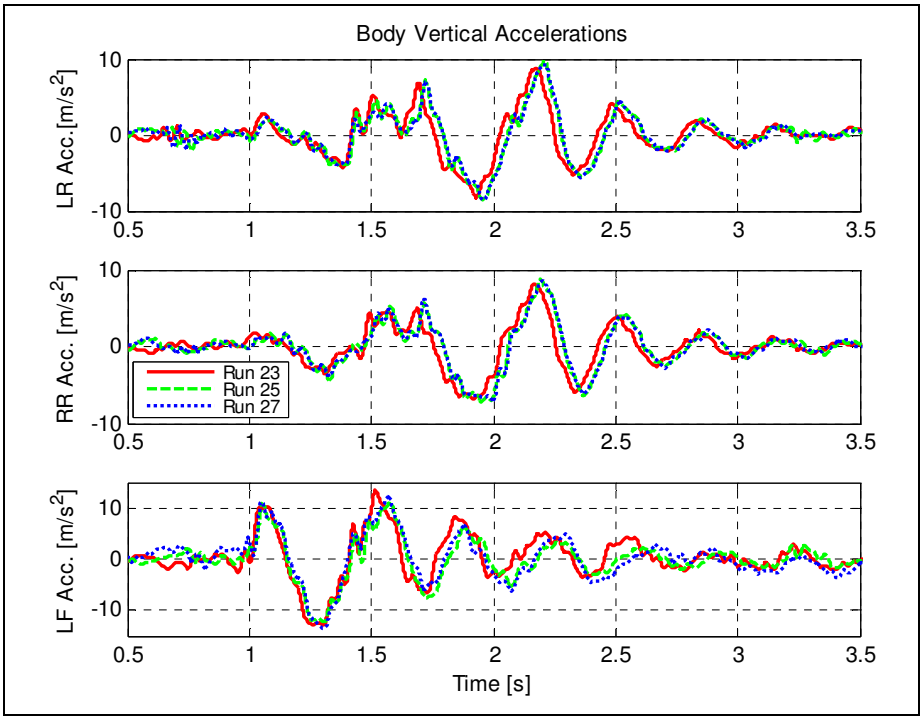

Figure 2-18: Body vertical accelerations over small trapezoidal bump (handling).

The data recorded over the Belgian paving cannot be compared in the time domain and will thus be compared in tabular form in terms of Root Mean Square (RMS) values. The RMS suspension force, pitch velocity and body accelerations for runs 01, 05 and 10 are compared in Table 2-4. The fatigue damage calculated from measured strain 3 has also been included (See Chapter 3 for the details of this calculation). These runs were conducted at low speed with the suspension set to ride comfort.

Table 2-4: Comparison for three runs on Belgian paving (ride comfort, ±14.5 km/h).

	Run no.	RMS LR load cell force [N]	RMS LR calculated force [N] (Pressure)	RMS LR vertical acceleration [m/s^2]	RMS RR vertical acceleration [m/s^2]	Damage calculated from strain 3.
	01	5532	5454	1.333	1.319	5.1717e-7
	05	5522	5457	1.379	1.355	5.1229e-7
	10	5539	5539	1.391	1.399	5.4410e-7
Mean		**5531**	**5483**	**1.368**	**1.358**	**5.2452e-7**
Std. Dev.		**8.54**	**48.32**	**0.031**	**0.040**	**1.7131e-8**
COV:		**0.15%**	**0.88%**	**2.24%**	**2.95%**	**3.27%**

The Coefficient of Variation (COV) on the load cell forces for this set of runs is smaller than 1% indicating very good consistency. The COV, on both the calculated damage and vertical accelerations, is smaller than 4% which is considered to be within the measurement accuracy. The runs at higher speed with the suspension set to ride comfort show similar results, Table 2-5.

Table 2-5: Comparison for three runs on Belgian paving (ride comfort, ±54 km/h).

	Run no.	RMS LR load cell force [N]	RMS LR calculated force [N] (Pressure)	RMS LR vertical acceleration [m/s^2]	RMS RR vertical acceleration [m/s^2]	Damage calculated from strain 3.
	02	5589	5462	2.715	2.604	2.3033e-6
	06	5617	5511	2.822	2.523	2.2346e-6
	11	5596	5515	2.756	2.565	2.1174e-6
Mean		**5601**	**5496**	**2.764**	**2.564**	**2.2184e-6**
Std. Dev.		**14.57**	**29.51**	**0.054**	**0.041**	**9.3999e-8**
COV:		**0.26%**	**0.54%**	**1.95%**	**1.58%**	**4.24%**

The runs with the suspension set to handling at low speed are compared in Table 2-6. The runs show slightly increased COV on suspension forces when compared to the runs with the softer suspension. The vertical acceleration COVs at the slower speed are slightly reduced from the runs on ride comfort mode. The variation on fatigue damage is somewhat higher than in previous data. Again the variation is considered to be acceptable.

Table 2-6: Comparison for three runs on Belgian paving (handling, ±14.5 km/h).

	Run no.	RMS LR load cell force [N]	RMS LR calculated force [N] (Pressure)	RMS LR vertical acceleration [m/s^2]	RMS RR vertical acceleration [m/s^2]	Damage calculated from strain 3.
	03	5965	5888	3.011	3.195	5.6150e-6
	08	5909	5903	3.048	3.175	5.3625e-6
	12	5786	5781	3.030	3.175	5.5139e-6
Mean		**5887**	**5857**	**3.030**	**3.182**	**5.4971e-6**
Std. Dev.		**91.57**	**66.53**	**0.019**	**0.012**	**1.2708e-7**
COV:		**1.56%**	**1.14%**	**0.61%**	**0.36%**	**2.31%**

The data for the higher speed runs with the suspension set to handling mode are shown in Table 2-7. The COV for the force data is below 2% which is excellent. The variation in the calculated damage and the RMS vertical acceleration is however quite high. The higher speed and severe road conditions lead to higher excitation frequencies, which are thought to excite some structural modes in the vehicle. The vibration of the panels to which the transducers were mounted could cause this slight inconsistency. The variation in the data could also be caused by the fact that the human driver may not have followed exactly the same path across the Belgian Paving for each of the tests. The variation in the data, although larger than preferred, is not large enough to warrant changes to the test procedure or transducer positions. Discretion will however be required in comparing the measured data to simulation results.

Table 2-7: Comparison for three runs on Belgian paving (handling, ±54 km/h).

	Run no.	**RMS LR load cell force [N]**	**RMS LR calculated force [N] (Pressure)**	**RMS LR vertical acceleration [m/s^2]**	**RMS RR vertical acceleration [m/s^2]**	**Damage calculated from strain 3.**
	04	6264	6167	5.053	4.476	1.2887e-5
	09	6278	6199	5.450	5.232	1.3178e-5
	13	6078	6035	5.435	5.051	1.1678e-5
Mean		**6207**	**6134**	**5.313**	**4.920**	**1.2581e-5**
Std. Dev.		**111.65**	**86.93**	**0.225**	**0.395**	**7.9544e-7**
COV:		**1.80%**	**1.42%**	**4.24%**	**8.02%**	**6.32%**

2.4 Summary and conclusions to experimental work

The test vehicle was instrumented for testing in order to capture its dynamics. Discrete obstacle tests were used in order to excite the vehicle's vertical dynamics, while random rough terrain was used in order to excite the vehicle's response. Selected data sets were presented in order to show the soundness and repeatability of the test procedure employed. The test procedure is concluded to be highly repeatable. The data recorded can thus be used to validate the vehicle dynamics model which is developed in the following chapter. The results of the experimental work are presented in Chapter 4 as required to validate the mathematical model. We can now progress to the construction of the mathematical model of the vehicle.

Chapter 3

3. Mathematical modeling

This chapter considers the construction of the mathematical model of the vehicle's dynamics. The results generated from the model are to be used to calculate the cost function value for the mathematical optimisation, which is the ultimate goal of the study. The model should therefore be as simple as possible in order to be computationally efficient, and yet maintain enough complexity to be representative of the actual vehicle's dynamics.

At the core, vertical dynamics models consist of a system of 2^{nd} order Ordinary Differential Equations (ODE). A 2^{nd} order Equation of Motion (EoM) exists for each DOF of the system being considered. The solution of these equations requires rewriting them as a system of 1^{st} order Ordinary Differential Equations and applying numerical integration schemes, such as Runge Kutta integration. The number of computations required for this operation is of order N^3, with N the number of equations involved. The solution time of the vehicle model is thus a cubic function of the degrees of freedom included in the model and as such the degrees of freedom should preferably be kept to a minimum.

An ADAMS multi-body dynamics model of the Land Rover has been developed and used in several studies by Thoresson (2003, 2007), Cronje (2008) and several others. The model contains 23 bodies each with six degrees of freedom, leading to 138 degrees of freedom in total of which 16 are unconstrained. The model is therefore very computationally expensive and was considered unsuitable for the current study. The Land Rover's rigid axles and simple suspension kinematics could be exploited to develop a more computationally efficient model of the vehicle.

A linear pitch-plane four degree of freedom model, with non-linear spring and damper characteristics was constructed as a first attempt. The model was implemented in the Matlab environment. The results from this model showed some resemblance to reality, but were of insufficient quality to be used in the optimisation. The model was improved by introducing non-linear suspension characteristics, paying special attention to spring and damper characteristics, suspension friction, suspension bushings and bump stops. The model in this form gave excellent correlation over the discrete obstacles, but correlation over the random terrain was still lacking, due to the absence of the roll degree of freedom.

The model was thus expanded to a seven DOF model, which is able to observe the roll degrees of freedom of the vehicle body and axles. This model gave much improved results over the random terrain, while maintaining the excellent correlation over the discrete obstacles. The computational expense of the more complex model was kept the same as the four DOF model by exploiting the sparsity of the mass, stiffness and damping matrices. The development of the seven DOF model is presented in this chapter.

3.1 Seven degree of freedom model

The seven degree of freedom model is shown in the diagram of Figure 3-1. The model consists of three masses, representing the vehicle body (sprung mass) and the front and rear unsprung masses. The sprung mass has a translational DOF in the vertical direction and two rotational DOF in the pitch and roll directions. The sprung mass represents the vehicle body and chassis. The unsprung masses each have a translational DOF in the vertical direction as well as a rotational degree of freedom in the roll direction. The unsprung masses represent the vehicle axles.

The sprung mass is connected to each of the unsprung masses by actuators representing the $4S_4$ suspension struts. The vehicle's tyres are each represented by a spring and damping element in parallel.

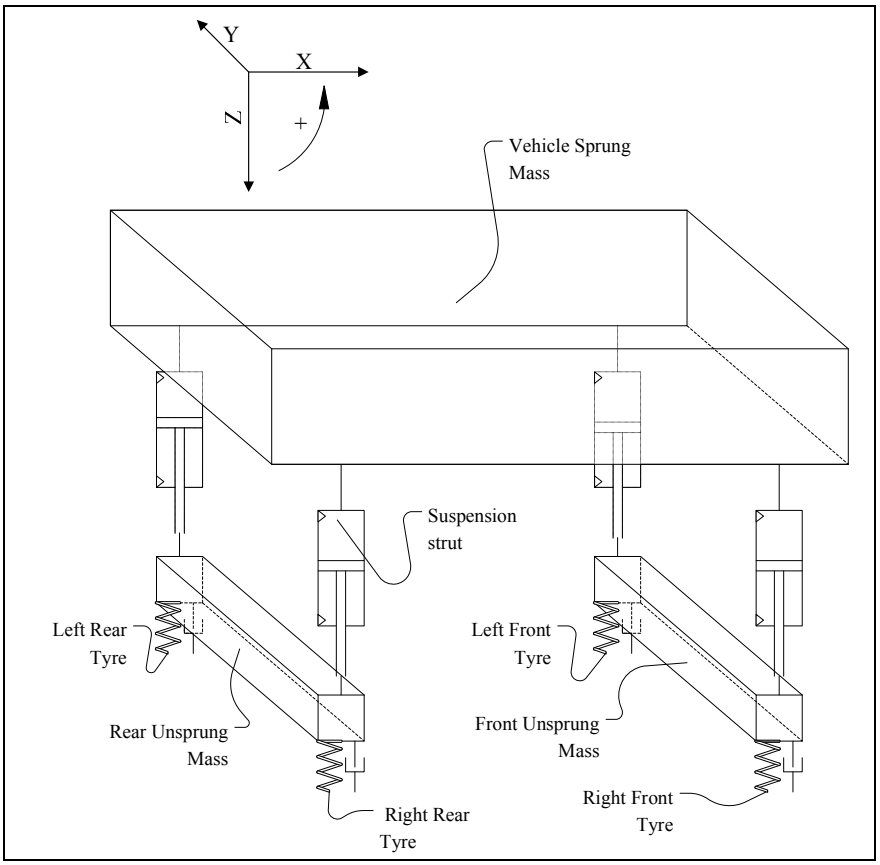

Figure 3-1: Seven DOF vehicle model.

The equations of motion for the system are derived by setting up the force and moment equations for each mass considered. The complete derivation, for the equations of motion for the model, is given in Appendix A.

The seven 2nd order differential equations, which are the equations of motion for the system, are presented below:

$$m_b \ddot{z}_b = m_b g - F_{srr} - F_{slr} - F_{slf} - F_{srf}$$

$$m_{uf} \ddot{z}_{uf} + k_{trf} z_{rrf} + k_{trf} c\theta + c_{trf} \dot{z}_{uf} + c_{trf} c\dot{\theta} + k_{tlf} z_{uf} - k_{tlf} c\theta + c_{tlf} \dot{z}_{uf} - c_{tlf} c\dot{\theta} \ldots$$
$$= k_{trf} z_{rrf} + c_{trf} \dot{z}_{rrf} - k_{trf} \delta_{stf} + k_{tlf} z_{rlf} + c_{tlf} \dot{z}_{rlf} - k_{tlf} \delta_{stf} + m_{uf} g + F_{srf} + F_{slf}$$

$$m_{ur} \ddot{z}_{ur} + k_{trr} z_{rrr} + k_{trr} c\phi + c_{trr} \dot{z}_{ur} + c_{trr} c\dot{\phi} + k_{tlr} z_{ur} - k_{tlr} c\phi + c_{tlr} \dot{z}_{ur} - c_{tlr} c\dot{\phi} \ldots$$
$$= k_{trr} z_{rrr} + c_{trr} \dot{z}_{rrr} - k_{trr} \delta_{str} + k_{tlr} z_{rlr} + c_{tlr} \dot{z}_{rlr} - k_{tlr} \delta_{str} + m_{ur} g + F_{srr} + F_{slr}$$

$$I_{xx} \ddot{\alpha} = (F_{slr} + F_{slf}) d - (F_{srr} + F_{srf}) d$$

$$I_{yy} \ddot{\beta} = (F_{slf} + F_{srf}) b - (F_{srr} + F_{slr}) a$$

$$I_{xx\_a} \ddot{\theta} - ck_{tlf} z_{uf} + c^2 k_{tlf} \theta - cc_{tlf} \dot{z}_{uf} + c^2 c_{tlf} \dot{\theta} + ck_{trf} z_{uf} + c^2 k_{trf} \theta + cc_{trf} \dot{z}_{uf} + c^2 c_{trf} \dot{\theta} \ldots$$
$$= -ck_{tlf} z_{rlf} - cc_{tlf} \dot{z}_{rlf} + ck_{tlf} \delta_{stf} + ck_{trf} z_{rrf} + cc_{trf} \dot{z}_{rrf} - ck_{trf} \delta_{stf} + d(F_{srf} - F_{slf})$$

$$I_{xx\_a} \ddot{\phi} - ck_{tlr} z_{ur} + c^2 k_{tlr} \phi - cc_{tlr} \dot{z}_{ur} + c^2 c_{tlr} \dot{\phi} + ck_{trr} z_{ur} + c^2 k_{trr} \phi + cc_{trr} \dot{z}_{ur} + c^2 c_{trr} \dot{\phi} \ldots$$
$$= -ck_{tlr} z_{rlr} - cc_{tlr} \dot{z}_{rlr} + ck_{tlr} \delta_{str} + ck_{trr} z_{rrr} + cc_{trr} \dot{z}_{rrr} - ck_{trr} \delta_{str} + d(F_{srr} - F_{slr})$$

In matrix form this becomes:

$$[M]\{\ddot{z}\} + [C]\{\dot{z}\} + [K]\{z\} = \{f\}, \text{ where}$$

$$[M] = \begin{bmatrix} m_b & & & & & & \\ & m_{uf} & & & & [0] & \\ & & m_{ur} & & & & \\ & & & I_{xx} & & & \\ & & & & I_{yy} & & \\ & [0] & & & & I_{xx\_a} & \\ & & & & & & I_{xx\_a} \end{bmatrix},$$

$$[C]=\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & c_{trf}+c_{tlf} & 0 & 0 & 0 & c(c_{trf}-c_{tlf}) & 0 \\ 0 & 0 & c_{trr}+c_{tlr} & 0 & 0 & 0 & c(c_{trr}-c_{tlr}) \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & c(c_{trf}-c_{tlf}) & 0 & 0 & 0 & c^2(c_{tlf}+c_{trf}) & 0 \\ 0 & 0 & c(c_{trr}+c_{tlr}) & 0 & 0 & 0 & c^2(c_{tlr}+c_{trr}) \end{bmatrix},$$

$$[K]=\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & k_{trf}+k_{tlf} & 0 & 0 & 0 & c(k_{trf}-k_{tlf}) & 0 \\ 0 & 0 & k_{trr}+k_{tlr} & 0 & 0 & 0 & c(k_{trr}-k_{tlr}) \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & c(k_{trf}-k_{tlf}) & 0 & 0 & 0 & c^2(k_{tlf}+k_{trf}) & 0 \\ 0 & 0 & c(k_{trr}+k_{tlr}) & 0 & 0 & 0 & c^2(k_{tlr}+k_{trr}) \end{bmatrix},$$

$$\{f\}=\begin{Bmatrix} m_b g - F_{ssr} - F_{slr} - F_{slf} - F_{srf} \\ k_{trf}(z_{rrf}-\delta_{stf})+c_{trf}\dot{z}_{rrf}+k_{tlf}(z_{rlf}-\delta_{stf})+c_{tlf}\dot{z}_{rlf}+m_{uf}g+F_{srf}+F_{slf} \\ k_{trr}(z_{rrr}-\delta_{str})+c_{trr}\dot{z}_{rrr}+k_{tlr}(z_{rlr}-\delta_{str})+c_{tlr}\dot{z}_{rlr}+m_{ur}g+F_{srr}+F_{slr} \\ (F_{slr}+F_{slf}-F_{srr}-F_{srf})d \\ (F_{slf}+F_{srf})b-(F_{srr}+F_{slr})a \\ c[k_{tlf}(\delta_{stf}-z_{rlf})-c_{tlf}\dot{z}_{rlf}-k_{trf}(\delta_{stf}-z_{rrf})+c_{trf}\dot{z}_{rrf}]+d(F_{srf}-F_{slf}) \\ c[k_{tlr}(\delta_{str}-z_{rlr})-c_{tlr}\dot{z}_{rlr}-k_{trr}(\delta_{str}-z_{rrr})+c_{trr}\dot{z}_{rrr}]+d(F_{srr}-F_{slr}) \end{Bmatrix}$$

$$\{z\}=\begin{Bmatrix} z_b \\ z_{uf} \\ z_{ur} \\ \alpha \\ \beta \\ \theta \\ \phi \end{Bmatrix},\ \{\dot{z}\}=\begin{Bmatrix} \dot{z}_b \\ \dot{z}_{uf} \\ \dot{z}_{ur} \\ \dot{\alpha} \\ \dot{\beta} \\ \dot{\theta} \\ \dot{\phi} \end{Bmatrix} \text{ and } \{\ddot{z}\}=\begin{Bmatrix} \ddot{z}_b \\ \ddot{z}_{uf} \\ \ddot{z}_{ur} \\ \ddot{\alpha} \\ \ddot{\beta} \\ \ddot{\theta} \\ \ddot{\phi} \end{Bmatrix}.$$

This system of seven second order differential equations can be decomposed to form the following system of 14 first order differential equations.

$$\{\dot{Z}\} = \begin{Bmatrix} \{\dot{z}\} \\ \{\ddot{z}\} \end{Bmatrix} = \begin{bmatrix} [0] & [I] \\ -[M]^{-1}[K] & -[M]^{-1}[C] \end{bmatrix} \begin{Bmatrix} \{z\} \\ \{\dot{z}\} \end{Bmatrix} + \begin{Bmatrix} \{0\} \\ [M]^{-1}\{f\} \end{Bmatrix}, \text{ with}$$

$$\{Z\} = \begin{Bmatrix} \{z\} \\ \{\dot{z}\} \end{Bmatrix} \text{ and } \{\dot{Z}\} = \begin{Bmatrix} \{\dot{z}\} \\ \{\ddot{z}\} \end{Bmatrix}$$

The system of first order differential equations of motion is now in a form amicable to solution by a numerical integration scheme, such as Runge Kutta integration. The Runge Kutta algorithm for systems, as presented in Burden and Faires (2005), was programmed in Matlab in order to solve the equations of motion. Coding the Runge Kutta algorithm was necessary due to convergence problems encountered with Matlab's built in differential equation solvers. The Matlab solvers failed to converge when friction models were introduced. The coded Runge Kutta algorithm gave the author better control over the convergence behaviour of the solver, which lead to more robust convergence.

3.2 Modelling parameters for the vehicle model

The mass properties and force element characteristics for the Land Rover Defender 110 are fairly well known, thanks to the various studies already conducted on this experimental vehicle at the University of Pretoria. The force element characteristics are also well documented. The mass properties and force element characteristics employed in the initial seven DOF vehicle model will now be discussed.

3.2.1 Mass properties

The mass properties of the unladen vehicle (i.e. masses, centre of mass position, and mass moments of inertia) were determined experimentally and published by Uys et al., (2006). The masses and mass moments of inertia employed in the model for the unladen vehicle are listed in Table 3-1. The masses and moments of inertia employed are representative of the condition of the test vehicle during testing. The dimensions quoted in the table are shown in a front and side view of the vehicle in Figure 3-2 and Figure 3-3, respectively.

Table 3-1: Model mass properties for the unladen vehicle.

Mass property	Value [Units]
Sprung mass, m_b	1734 [kg]
Sprung mass pitch moment of inertia, I_{yy}	2440 [kg.m2]
Sprung mass roll of inertia, I_{xx}	688 [kg.m2]
Front unsprung mass, m_{uf}	229 [kg]
Front unsprung mass roll moment of inertia, I_{xx_a}	33.1 [kg.m2]
Rear unsprung mass, m_{ur}	229 [kg]
Rear unsprung mass roll moment of inertia, I_{xx_a}	33.1 [kg.m2]
Centre of mass longitudinal position (a,b)	1.395, 1.395 [m]
Centre of mass vertical position (h)	1.190 [m]
Wheel track width (2c)	1.5, (c=0.75) [m]
Suspension track width (2d)	1.1, (d=0.55) [m]

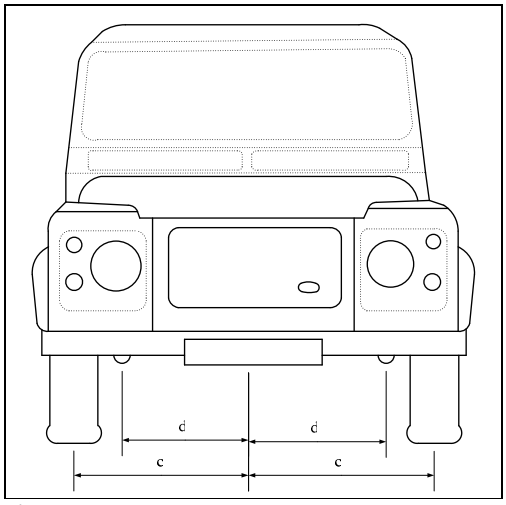

Figure 3-2: Vehicle front view.

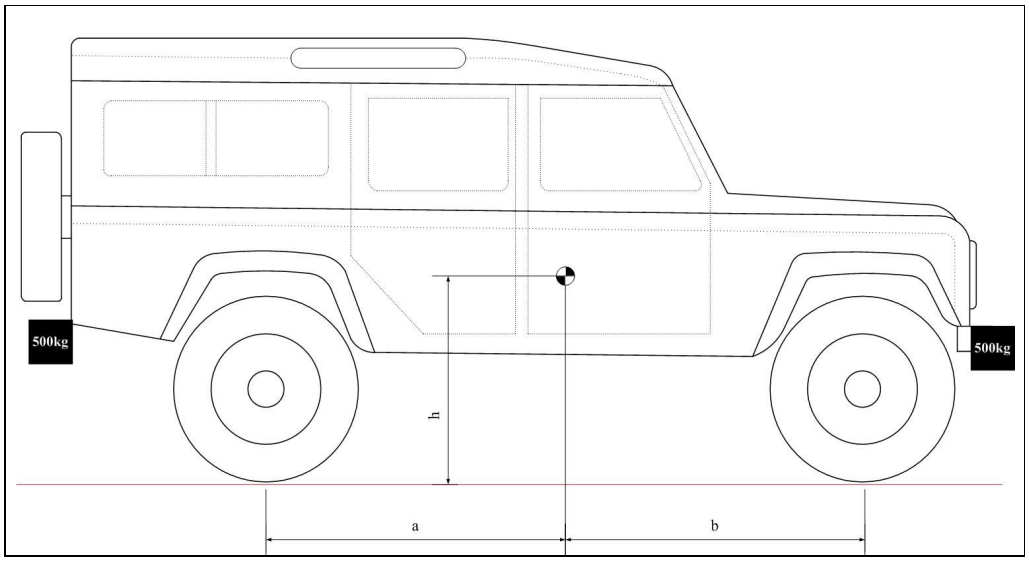

Figure 3-3: Vehicle in the fully laden configuration.

It is of interest to investigate the vehicle suspension characteristics for other vehicle payload cases. Three additional load cases have been chosen to this goal. The first is a fully laden case which is representative of the test vehicles maximum load capability. This load case can be achieved by attaching 500 kg of steel weights to the front and rear bumpers of the vehicle as pictured in Figure 3-3. This arrangement was chosen as it can be conveniently reproduced on the real vehicle for testing, if required. The vehicle mass properties for this configuration are given in Table 3-2.

Table 3-2: Model mass properties for the fully laden vehicle.

Mass property	**Value [Units]**
Sprung mass, m_b	2734 [kg]
Sprung mass pitch moment of inertia, I_{yy}	5080 [kg.m2]
Sprung mass roll of inertia, I_{xx}	1298 [kg.m2]
Front unsprung mass, m_{uf}	229 [kg]
Front unsprung mass roll moment of inertia, I_{xx_a}	33.1 [kg.m2]
Rear unsprung mass, m_{ur}	229 [kg]
Rear unsprung mass roll moment of inertia, I_{xx_a}	33.1 [kg.m2]
Centre of mass longitudinal position (a,b)	1.150, 1.640 [m]
Centre of mass vertical position (h)	0.975 [m]
Wheel track width (2c)	1.5, (c=0.75) [m]
Suspension track width (2d)	1.1, (d=0.55) [m]

The third case has been chosen to represent the laden vs. unladen load ratio usually experienced in commercial truck applications. Cole (2001) states that the ratio of unsprung to sprung mass in heavy road vehicles can range between 0.1 and 1, for different load cases. In the case of the Land Rover being considered here an unsprung mass to sprung mass ratio of 0.1 would imply a sprung mass value of 4580 kg. The inertia of the sprung mass will be scaled proportionally to the mass increase. The proportional scaling of the inertia to the mass is valid assuming that the extra mass added to the vehicle is distributed uniformly throughout the vehicle. This payload condition is similar to that of armoured versions of the Land Rover. The mass properties for the vehicle in the heavy load condition are presented in Table 3-3.

Table 3-3: Model mass properties for the heavy load condition.

Mass property	**Value [Units]**
Sprung mass, m_b	4580 [kg]
Sprung mass pitch moment of inertia, I_{yy}	6897 [kg.m2
Sprung mass roll of inertia, I_{xx}	1802 [kg.m2]
Front unsprung mass, m_{uf}	229 [kg]
Front unsprung mass roll moment of inertia, I_{xx_a}	33.1 [kg.m2]
Rear unsprung mass, m_{ur}	229 [kg]
Rear unsprung mass roll moment of inertia, I_{xx_a}	33.1 [kg.m2]
Centre of mass longitudinal position (a,b)	1.395, 1.395 [m]
Centre of mass vertical position (h)	1.190 [m]
Wheel track width (2c)	1.5, (c=0.75) [m]
Suspension track width (2d)	1.1, (d=0.55) [m]

Commercially available semi-trailers have fully laden to laden mass ratios of between five and seven, implying a sprung mass of between 8500 and 12000 kg for the Land Rover. The extreme payload case sees the vehicle sprung mass being increased to 10000 kg. The high payload variations observed in semi-trailer applications is represented by a purely theoretical load case. This load case will be referred to as the extreme load condition. The extreme load condition has very little real analogy in the Land Rover application, but was nonetheless considered due to its importance for the heavy vehicle industry. The properties for this load case are shown in Table 3-4.

Table 3-4: Model mass properties for the extreme load condition.

Mass property	Value [Units]
Sprung mass, m_b	10000 [kg]
Sprung mass pitch moment of inertia, I_{yy}	15412 [kg.m2]
Sprung mass roll of inertia, I_{xx}	3120 [kg.m2]
Front unsprung mass, m_{uf}	229 [kg]
Front unsprung mass roll moment of inertia, I_{xx_a}	33.1 [kg.m2]
Rear unsprung mass, m_{ur}	229 [kg]
Rear unsprung mass roll moment of inertia, I_{xx_a}	33.1 [kg.m2]
Centre of mass longitudinal position (a,b)	1.395, 1.395 [m]
Centre of mass vertical position (h)	1.190 [m]
Wheel track width (2c)	1.5, (c=0.75) [m]
Suspension track width (2d)	1.1, (d=0.55) [m]

3.2.2 Force element characteristics

The suspension forces are made up of five contributing components namely: the hydro-pneumatic spring force, the hydraulic damping force, the bushing force in the vehicle trailing arms, the bump stop forces and the friction force in the $4S_4$ suspension struts. The four $4S_4$ suspension struts have been designed to have the same characteristics for each of the four corners of the vehicle. The front and rear struts differ only by the maximum suspension travel. No anti-roll bars were installed on the vehicle during testing.

3.2.2.1 Hydro-pneumatic spring force

The hydro-pneumatic spring is modelled as a polytropic gas compression process. A schematic of the hydro-pneumatic spring is shown in Figure 3-4.

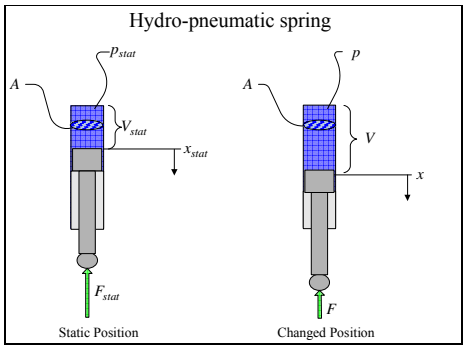

Figure 3-4: Schematic of the hydro-pneumatic spring.

The derivation of the force-displacement characteristic for a hydro-pneumatic spring with constant area is given in Appendix B. The force displacement characteristic is given by:

$$F = p_{stat} A \left(\frac{x_{stat}}{x} \right)^{n_p},$$

with:

F	=	force in the pneumatic spring,
p_{stat}	=	static pressure (constant),
A	=	area (constant),
x_{stat}	=	static displacement (constant),
x	=	hydro-pneumatic spring displacement,
n_p	=	polytropic gas constant.

The static pressure, static volume (product of area and static displacement) and polytropic constant are the quantities which determine the pneumatic spring's force-displacement characteristic. The values used for these quantities for the unladen payload case are listed in Table 3-5. Load levelling functionality was used to compensate for the higher payload cases.

Table 3-5: Constants determining the spring characteristic.

Quantity	Value [Units]
Static pressure, p_{stat}	4253 [N]
Static volume, V_{stat}	0.0001 [m3] for hard, 0.0005 [m3] for soft suspension settings
Area, A	1.963×10^{-3} [m2], (50mm diameter)
Total static displacement, x_{stat}	0.0509 [m] for hard, 0.2546 [m] for soft suspension settings
Polytropic constant, n_p	1

The polytropic constant is a function of the compression rate of the spring and varies between unity and the universal gas constant for a real process. A polytropic constant of unity models an isothermal process, while a polytropic constant equal to the universal gas constant models an adiabatic process. Different polytropic constants were attempted in the model, but the isothermal model adequately predicted the vehicle dynamics in most cases. Consequently a polytropic constant equal to unity was used.

The spring characteristic is also influenced, especially in the handling mode, by the compressibility of the oil in the suspension unit. This finding was stated by Els (2006) and the bulk modulus of the oil was subsequently measured in the $4S_4$, by Els. The bulk modulus of the AeroShell fluid 41 used in the $4S_4$ was measured as 1.368 GPa. The bulk modulus of the oil, β_f, can be included in the calculation of the spring force in order to give more accurate results.

The stiffness of the oil is modelled as a spring in series with the air spring characteristic. The system representing the spring characteristic thus consists of a linear spring (oil stiffness) and non-linear spring (air spring stiffness) in series. The system is solved for the displacement in the air spring and the displacement of the

bulk oil, given an initial estimate of the displacements. Once the displacements are known the force in the suspension can be calculated, from either the air spring or the oil stiffness characteristic. The system to be solved for the displacements is:

$$[F] = \begin{bmatrix} k_{\beta_f} x_{oil} - F_{air}(x_{air}) \\ x_{susp} - (x_{air} + x_{oil}) \end{bmatrix} = \bar{0},$$

where:

k_{β_f}	=	the stiffness of the bulk oil,
x_{oil}	=	the displacement due to the oil volume,
x_{susp}	=	the total suspension displacement, and
$F_{air}(x_{air})$	=	air spring force as a function of air volume.

The stiffness due to the spring was obtained in the derivation given in Appendix B. The stiffness due to the oil was found to be $k_{\beta_f} = -3.296$ MN/m. The hydro-pneumatic spring characteristic remains the same as that given earlier in this section.

The system described in the preceding paragraphs was solved using Matlab's non-linear equation solver fsolve.m. The solver employs the Newton-Raphson method for non-linear systems. The initial estimate for the air spring displacement was taken equal to the total suspension displacement, while the oil displacement was taken as zero.

Solution of the system introduced severe additional computational expense into the spring model. The solution time of the spring model more than doubled when the effects of oil bulk stiffness were included. The gains in suspension force prediction accuracy were also negligible and the additional computational expense could not be justified. The effects of the bulk oil stiffness were therefore excluded from the final model.

3.2.2.2 Hydraulic damper force and friction force

The damper characteristics of the $4S_4$, although not previously published, are well documented. This is due to the numerous studies which have been conducted on the experimental vehicle in question. The damper characteristics for the four struts were intended to be the same. Small variations within the manufacturing tolerances have however led to varying characteristics. The four damper characteristics have previously been individually characterised. The experimental characteristics found have been superimposed in Figure 3-5 and Figure 3-6, representing the handling and ride comfort characteristics respectively. The average characteristic is superimposed on the measured data, while the coefficient of variation of the experimental characteristics across the measured range is plotted in the subplots of both figures.

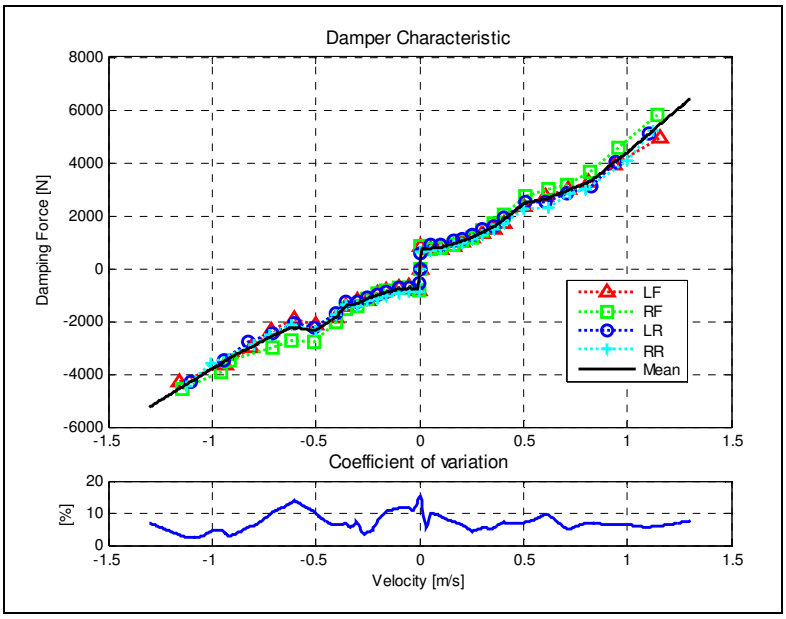

Figure 3-5: Damper characteristic for $4S_4$ (ride comfort).

The soft characteristic in the Figure 3-5 shows a pronounced friction "dead-band" between +-0.05 m/s, as well as blow-off knees at +-0.5 m/s. The COV varies between 3 and 17% across the range of speeds considered. The variation is larger than desirable. The average characteristic shown in black will be used in the model.

The hard characteristic shows a similar friction dead-band between +-0.05 m/s, to that observed with the soft characteristic. The data for the hard characteristic shows large variations, with COV values of up to 30%. The large variations are most certainly undesirable, but are due to manufacturing tolerances in the hardware.

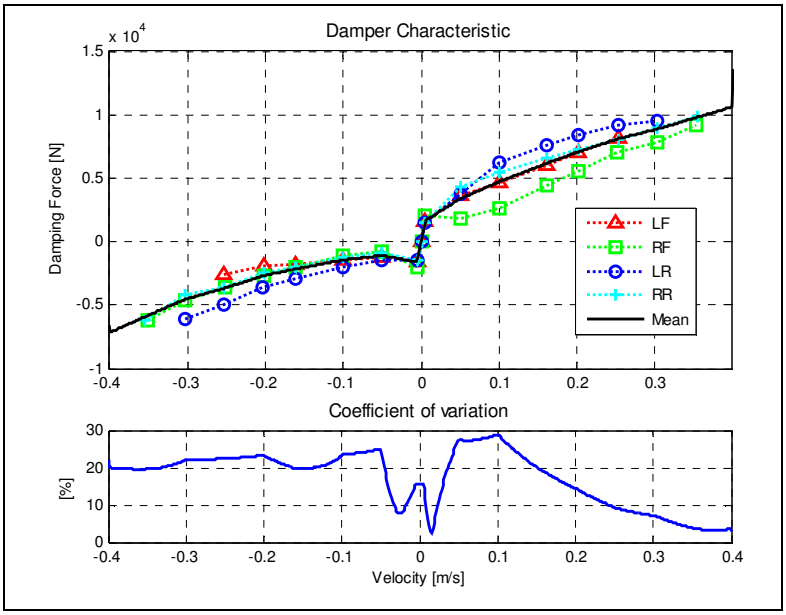

Figure 3-6: Damper characteristic for $4S_4$ (handling).

During these characterisation tests, the soft characteristic was determined for a velocity range of -1.2 to 1.2 m/s, while the handling characteristic was only determined for velocities between -0.35 and 0.35 m/s. The measured field test data contains suspension velocities of up to 1m/s for both the handling and ride comfort suspension settings. The additional data would thus have to be considered to define the handling characteristic over sufficiently wide range. One of the struts was however characterised over a greater velocity range (up to +-0.63 m/s) on the hard setting. This measured characteristic is shown in Figure 3-7. The additional data reveals a blow-off at 0.35 m/s.

The characteristic used will be constructed from the mean of the four individual characteristics between -0.35 and 0.35 m/s, the additional data for absolute velocities up to 0.63 m/s and then linearly extrapolated for larger absolute velocities. The characteristic to be used is shown in black in Figure 3-7.

The force characteristics for both the ride comfort and handling suspension settings are seen to have a friction component and a viscous component. The friction component is caused by friction between the hydraulic seals, wear rings and the strut cylinder walls. The friction is thus inherent to the system and cannot be changed. The damper characteristic, which in fact implies the viscous component, will however serve as a design variable in the optimisation and as such will be subject to modification. It is thus desirable to model the friction and viscous damping separately, in order to obtain realistic results from the optimisation.

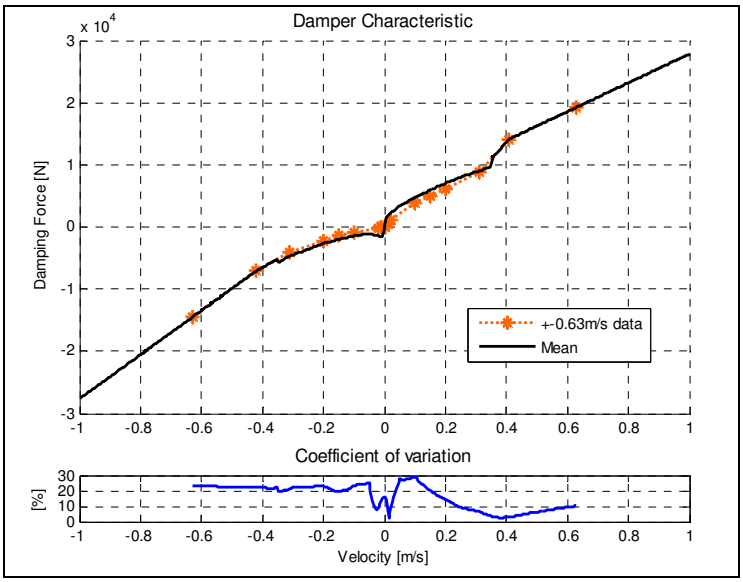

Figure 3-7: Extrapolated damper characteristic for $4S_4$ (handling).

The force-velocity characteristics resembles a Stribeck curve at low velocities. Since Stribeck first introduced this concept of velocity dependent friction in 1902, several attempts have been made to describe it mathematically (Armstrong-Hélouvry et al., 1994). The exponential model given below is proposed by Armstrong-Hélouvry et al.

$$F_f(\dot{x}) = F_C + (F_s - F_C)e^{-\left(\dot{x}/\dot{x}_s\right)^{\delta}} + F_v\dot{x}$$

In the model, $\dot{x}$ is the relative velocity between the sliding surfaces. The constants F_C and F_s describe the minimum level of Coulomb friction and the level of static friction. The constant F_v is used to include simple viscous friction, while the constants δ and $\dot{x}_s$ are empirical parameters, which determine the behaviour in transitional region between static and Coulomb friction. These parameters mostly have good physical significance, which makes system identification somewhat more intuitive. Our aim is to model the viscous friction separately, thus it will be neglected in the friction model, which then becomes:

$$F_f(\dot{x}) = F_C + (F_s - F_C)e^{-\left(\dot{x}/\dot{x}_s\right)^{\delta}}.$$

This form of the model possesses a sharp discontinuity at zero velocity. Such a discontinuity is difficult to deal with numerically. Numerical solvers tend to oscillate around the point of discontinuity and struggle to converge. The model was therefore adapted to remove the discontinuity. The friction force is interpolated linearly between zero velocity and a critical velocity, $\dot{x}_c$. The modification is as follows:

$$F_f(\dot{x}) = \begin{cases} \left(\dfrac{F_C + (F_s - F_C)e^{-\left(\dot{x}_c/\dot{x}_s\right)^{\delta}}}{\dot{x}_c} \right) \dot{x} & ,|\dot{x}| \le \dot{x}_c \\ F_C + (F_s + F_C)e^{-\left(\dot{x}_c/\dot{x}_s\right)^{\delta}} & ,|\dot{x}| > \dot{x}_c \end{cases}$$

The friction characteristic employed was found by fitting the Stribeck model to both the hard and soft damping curves and averaging the two Stribeck models obtained. The model was fitted to the damping curves between -0.05 and 0.05 m/s. This was done to ensure minimal participation of viscous damping effects in the friction model. The viscous damping curves used as well as the friction damping curve for both the hard and soft suspension settings are given in Figure 3-8 and Figure 3-9, respectively.

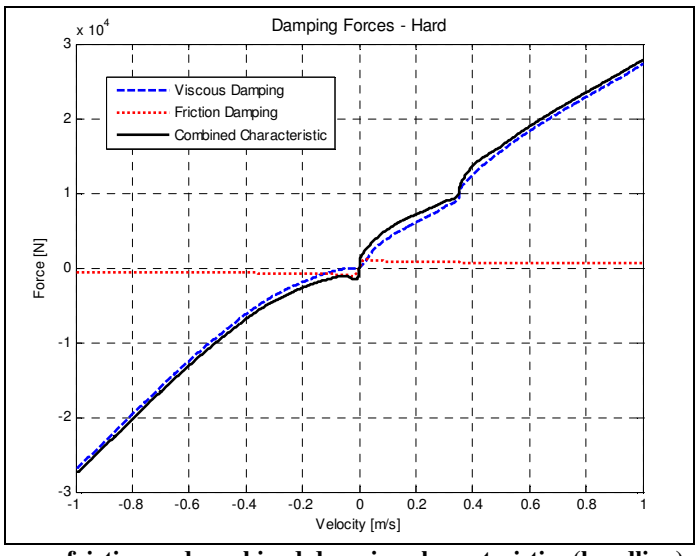

Figure 3-8: Viscous, friction and combined damping characteristics (handling).

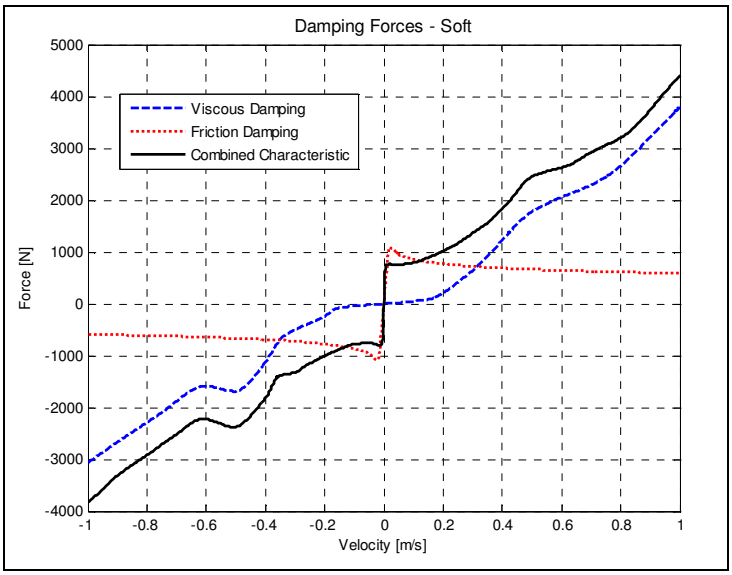

Figure 3-9: Viscous, friction and combined damping characteristics (ride comfort).

It should be clear from the figures that the softer suspension setting will be most greatly affected by the high friction in the system. The correlation of the model would thus be affected. Further investigation of the friction in the $4S_4$ was deemed important.

The limitation of this model, for modelling the $4S_4$ friction, is that it is only dependent on velocity. The model does not take the pressure in the system or loading conditions in account. Since the pressure in the system provides preloading for the hydraulic lip seals employed, it may have a significant effect on the friction in the system. Bonchis et al. (1999) developed a pressure-based model for a double acting hydraulic cylinder. The model is attractive due to its ability to capture both the influence of pressure and velocity on the system. It does however suffer from the fact that it uses five empirical parameters, which have vague physical analogies. The model has the further limitation that, as with the Stribeck model above, it is a steady state model and does not take dynamic friction effects into account.

A popular dynamic friction model is the LuGre model. The model takes the deflection of surface asperities into account by modeling them as elastic bristles. Yanada and Sekikawa (2008) applied the model to a hydraulic cylinder application. They found that the model gave good correlation for velocity variations in one direction without stoppages. The LuGre model gave poor results for stoppages and velocity reversals. Yanada and Sekikawa attribute this poor correlation to the lack of consideration of lubricant film dynamics in the LuGre model. They proposed modifications to the LuGre model in order to take the film dynamics into account. The modifications to the model greatly improved the model's response to stoppages and velocity reversals.

The hydraulic cylinder application considered by the above mentioned authors have many similarities to the hydro-pneumatic $4S_4$ system. The friction in the $4S_4$ system may thus show similar behaviour to this system. An experimental characterisation of the friction in the $4S_4$ system was thus conducted in order to better characterise the friction in the $4S_4$ system. The experimental setup is pictured in Figure 3-10.

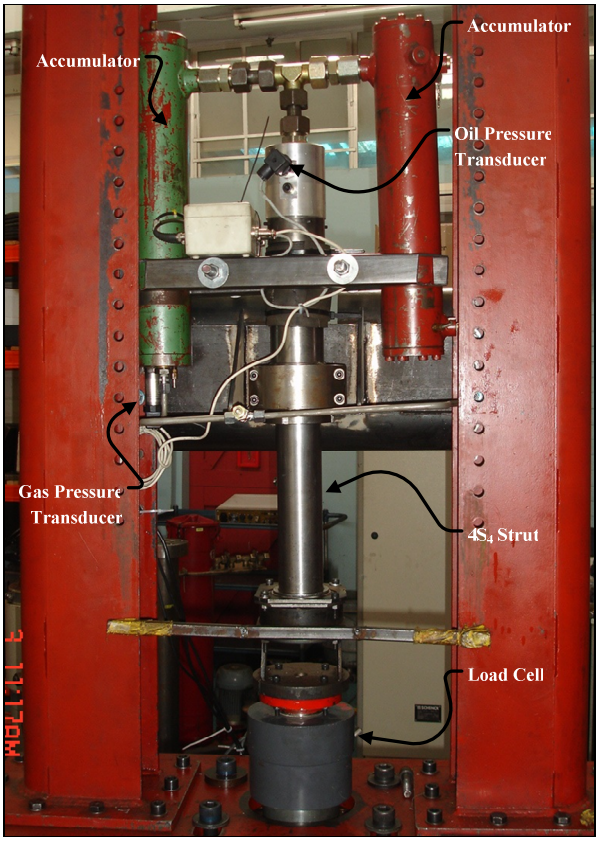

Figure 3-10: Friction characterisation setup.

The setup consisted of a $4S_4$ strut, which had the valve block and accumulators replaced, with two large volume accumulators (10 ℓ in total). The large accumulators would ensure that the pressure in the system would vary as little as possible throughout the strut's stroke length. The two large accumulators were connected by a tube to ensure equal pressure. This would allow the friction to be characterised at different, almost constant pressures. The gas and hydraulic working fluids were separated by floating pistons in the accumulators. The valve block used in the test contained no damper packs and was designed with large flow passages to ensure that viscous damping in the system was negligible.

The strut assembly for the test was mounted to a rigid frame for characterisation using a servo-hydraulic actuator. The strut assembly was supported by spherical bearings to ensure that no moment was applied to the strut. The pressure of the gas and oil were measured using pressure transducers located as indicated in Figure 3-10. The force exerted by the strut was measured using a load cell between the strut and the actuator. During the tests, prescribed displacements were applied to the strut and the force and pressure responses were recorded using a CDAS data acquisition system. Data was sampled at 1000 Hz and a 300 Hz cut-off anti-aliasing filter was used.

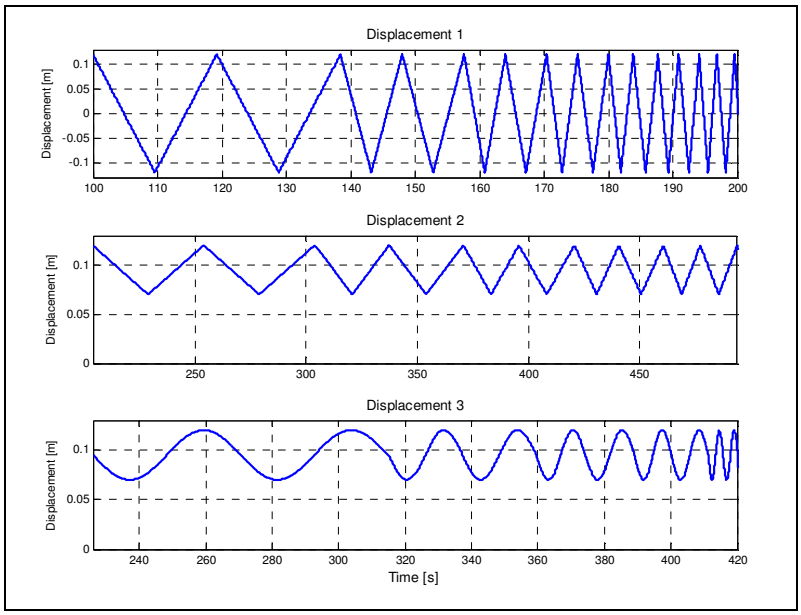

Figure 3-11: Prescribed displacements for friction characterisation.

The tests were conducted at static gas pressures of 500, 1500, 2000, 2500 and 3000 kPa. These pressures correspond to 100, 300, 400, 500 and 600 kg load on the strut, with the strut at maximum extension of 130 mm. The static load in the struts on the Land Rover is approximately 450 kg. The accumulators were charged to the desired pressure with the strut at maximum extension.

Three different prescribed displacements signals, shown in Figure 3-11, were used during the tests. The first was a large amplitude triangular wave with increasing frequency. The triangular displacement implies constant velocities, which allow steady state to be achieved. The amplitude of the displacement remained 0.12 m, while the velocity was varied between 0.01 and 0.25 m/s.

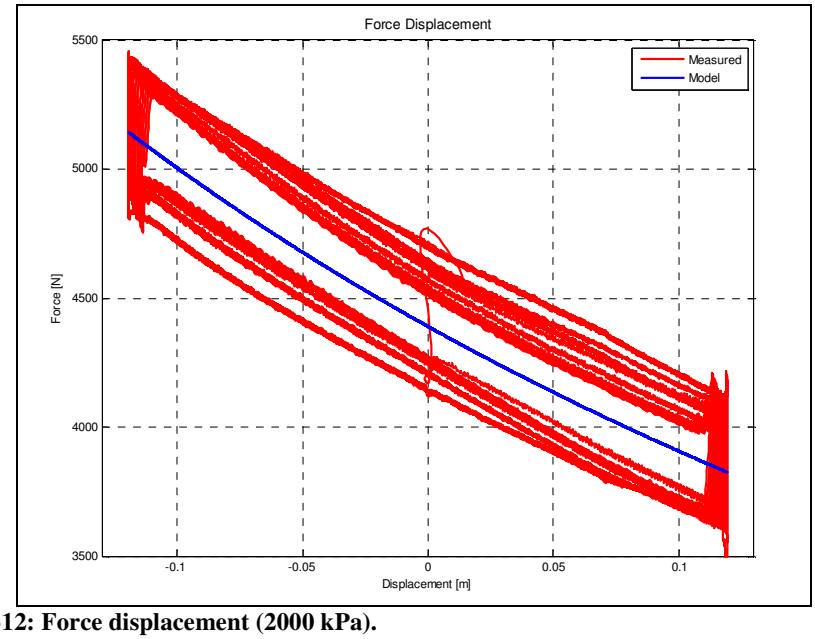

Figure 3-12: Force displacement (2000 kPa).

The data from a test with the first displacement signal is plotted in Figure 3-12. The data shown is for a gas pressure of 2000 kPa. The hysteresis in the measured characteristic is due to the friction present. The change in the size of the hysteretic loop is caused by changes in the friction force as the velocity changes. A hydro-pneumatic spring model of the same form as given in section 3.2.2.1 was fitted to the data. The force characteristic of the model has been included in the plot. The hysteretic loop is notably absent from the simulated characteristic.

The second prescribed displacement was also a triangular wave of increasing frequency. The amplitude was kept small (0.025 m) to ensure small changes in pressure due to the compression of the gas volume. The frequency was varied such that the velocity varied between 0.001 and 0.2 m/s. The measured force data is the sum of the friction and the force due to the pressure in the strut. The friction force was isolated by subtracting the pneumatic spring force predicted by the model from the measured force data. The static force-velocity friction characteristics obtained by this method are overlaid in Figure 3-13. The measured characteristics indicate much lower friction than the characteristic previously estimated from the damper characteristics.

The data given in Figure 3-13 shows that the friction characteristic is biased towards the compression direction. There is some dependence on pressure in the friction characteristic. The effect is attributed to an increased preload on the hydraulic seals as pressure increases. The pressure dependence is considered small enough to neglect for our purpose. The characteristic measured for 2000 kPa was used to model the friction behaviour.

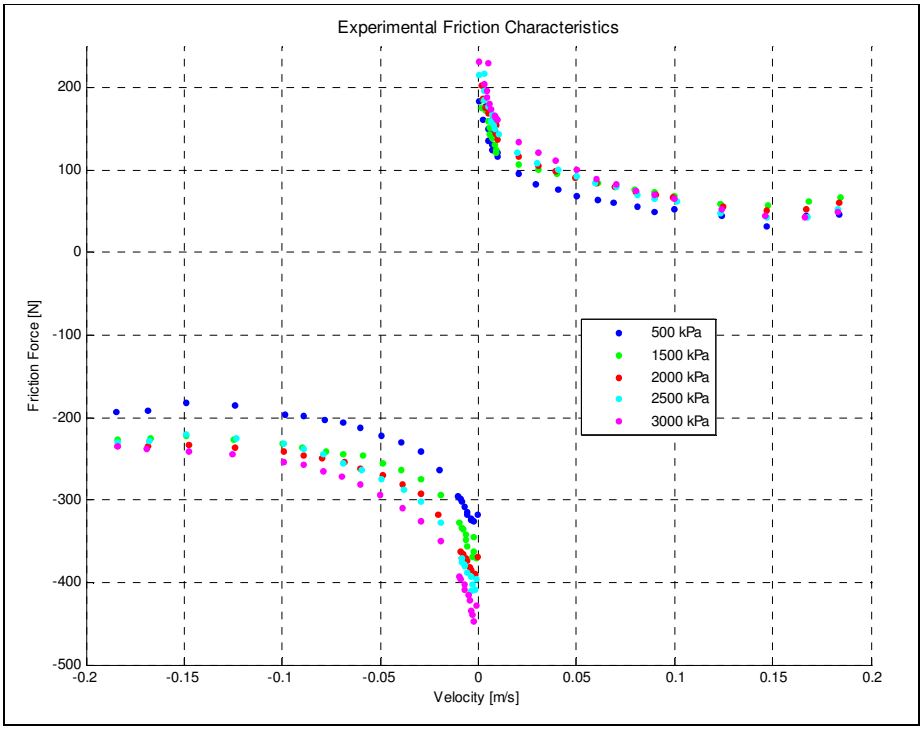

Figure 3-13: Static friction characteristic for different pressures.

The third prescribed displacement (0.025 m amplitude sine wave of increasing frequency) was used to evaluate the friction model fitted to the data. The maximum velocity of the sine wave varied between 0.003 and 0.15 m/s. The measured and predicted force data for the test at 2000 kPa with displacement 3 is shown in Figure 3-14. The correlation of the amplitude of the friction force is good for both directions. The correlation for other pressures was of similar quality.

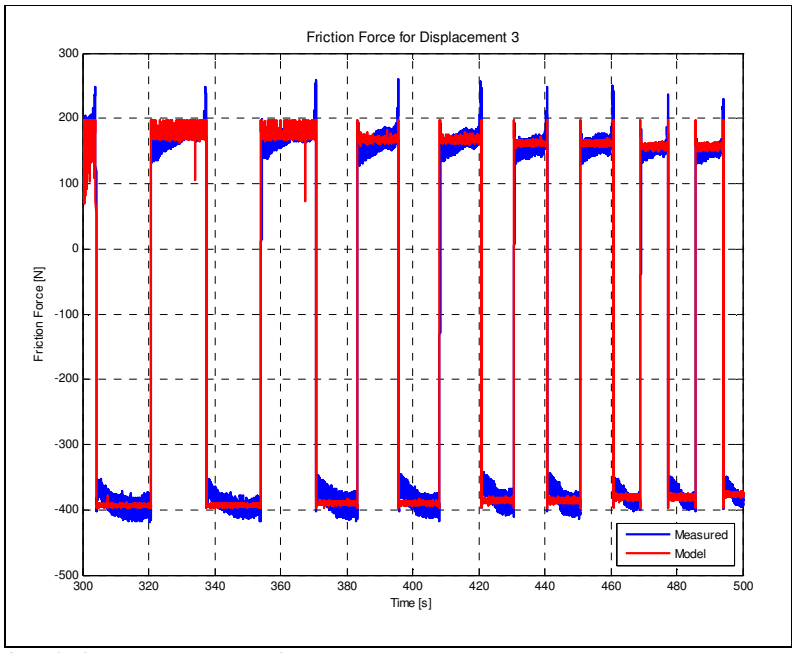

Figure 3-14: Friction model correlation.

The peaks at the velocity reversals from positive to negative velocity do not correlate exactly. Also, the measured friction force continues to develop throughout each uni-directional motion interval. These results are not captured by the model as they are unsteady. A dynamic friction model will be required to accurately model these effects. The differences between measured and simulated data are however considered to be small enough to ensure good representation of the vehicles dynamics. The implementation of a dynamic friction model is thus left to future studies.

3.2.2.3 Suspension bushing force

The stiffness of the bushings in the trailing arms of the Land Rover suspension influences the force characteristics of the vehicle suspension. These characteristics were previously determined experimentally. The characteristic used is given in Figure 3-15.

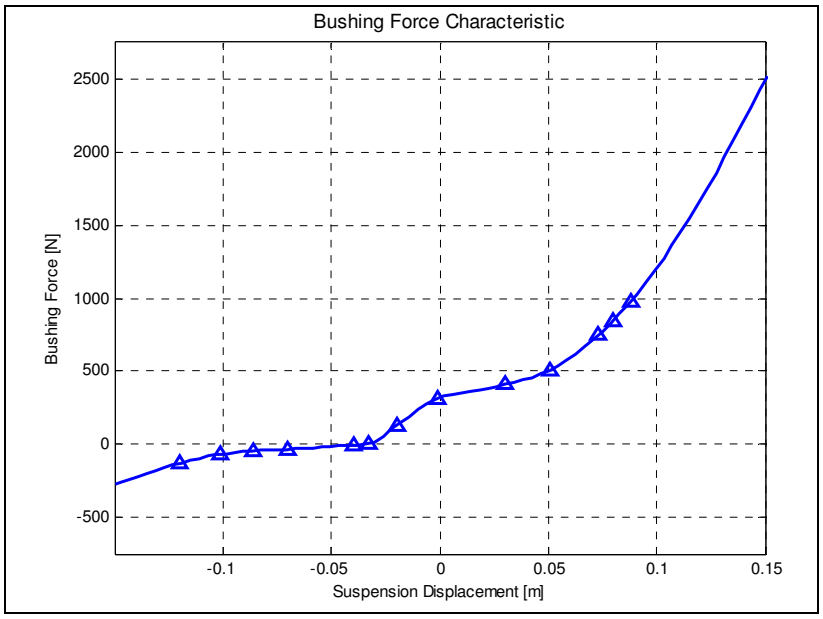

Figure 3-15: Suspension bushing force characteristic.

3.2.2.4 Bump stop force

The Land Rover suspension is fitted with rubber bump stops. The bump stops act as an additional stiffness in the suspension in the case of excessive suspension displacement. The bump stops have been previously characterised and the characteristics used here are presented in Figure 3-16. As can be deduced from Figure 3-16, the rear suspension has a greater available working space than the front suspension.

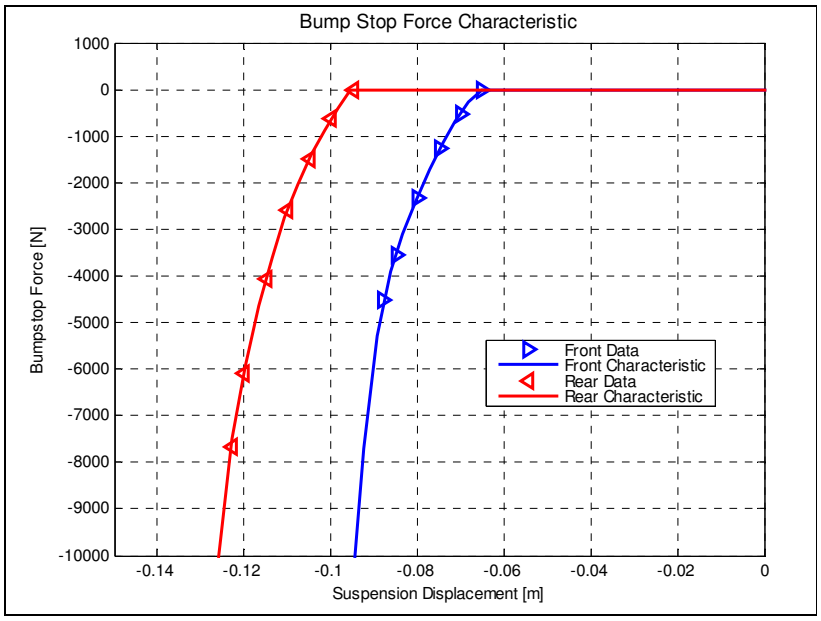

Figure 3-16: Bump stop force characteristic.

3.2.2.5 Modelling the tyre-road interface

The tyre is modelled as a point follower model with a parallel linear spring and linear damper. The non-rolling dynamic tyre stiffness and damping characteristics were previously determined experimentally using a pendulum type tyre tester and are presented in Figure 3-17. A linearised tyre stiffness of 250000 N/m, and a linearised damping coefficient of 1200 Ns/m were used. A logical argument in the implementation of the tyre model allows it to loose ground contact, i.e. the tyre force becomes zero when the tyre deflection exceeds the free radius.

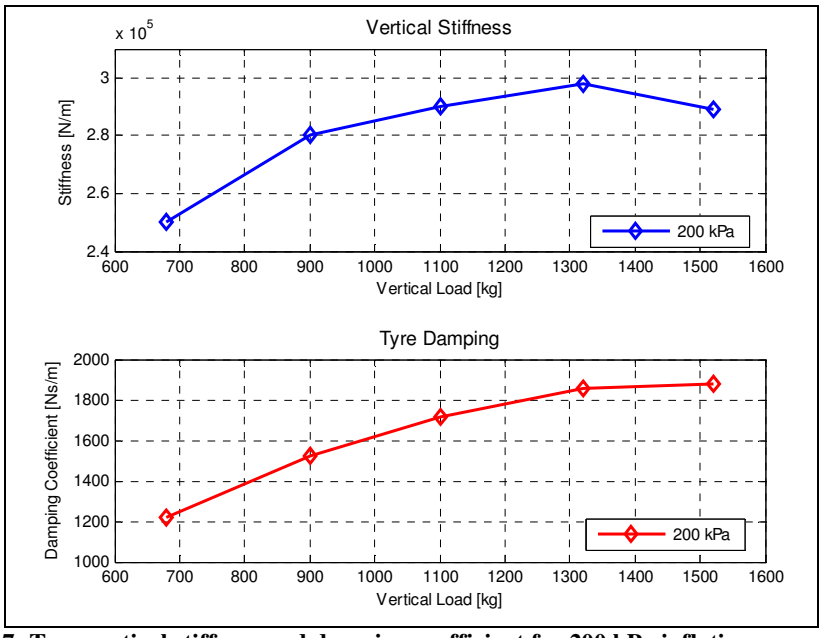

Figure 3-17: Tyre vertical stiffness and damping coefficient for 200 kPa inflation.

The road was modelled as a displacement and velocity input to the tyre spring-damper system. As mentioned in Chapter 2, the road profile of the Belgian paving was characterised by Bekker (2008). The exact vertical displacement versus longitudinal displacement of 24 tracks along the Belgian paving are known from Bekker's study. Additionally, the dimensions of the trapezoidal bump tracks used are also known. The measured road profile could thus be used as input to the tyre model.

The circular profile and local compliance or hysteresis of the real tyre rolling over rough terrain filters some of the higher frequency content of the road profile. The filtering due to the rolling wheel is however asymmetric in the sense that positive vertical displacements in the road profile (bumps) are filtered with less severity than negative vertical displacements (dips or potholes). This is referred to as rigid wheel filtering.

A second generally accepted filtering methodology is to low-pass filter the road profile using a spatial cut-off frequency corresponding to half the tyre contact patch length. This method is thought to be very effective for smooth roads, but is thought to break down over rough terrains such as the Belgian paving.

A portion of the measured Belgian paving profile is plotted in Figure 3-18. The same road profile, filtered by the two techniques described above, are superimposed on the measured profile for comparison. A spatial cut-off frequency of 20 cycles/m, corresponding to a contact patch length of 0.1 m, was used for the low-pass filtering. A wheel radius of 0.4 m was employed for the rigid wheel filtering. It should be clear from the figure that a vast difference exists between the results obtained from the two filtering techniques. The rigid wheel filtering is suggested to be more representative of the actual wheel filtering and warrants some further investigation. The use of more complex tyre contact models, such as footprint or sector models, was not considered due to their high computational expense.

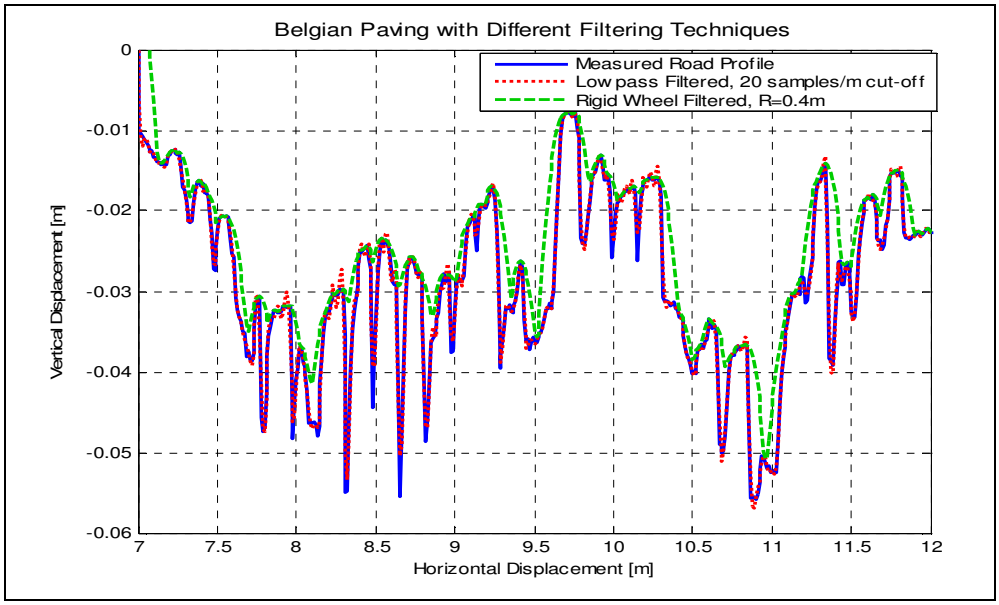

Figure 3-18: Belgian paving with different filtering techniques.

The FFTs and DSDs of the data shown in Figure 3-18 are presented in Figure 3-19. The plots have only been generated for spatial frequencies between 0.05 and 10 cycles per meter as recommended by ISO 8608 (International Organisation for Standards, 1995). The reader should note that the 20 cycles per meter cut-off for the low-pass filtered data is beyond the quoted frequency range. Comparison of the rigid wheel filtered data to the measured road profile data shows a smooth roll-off starting from 2.6 cycles/m. Interestingly this corresponds to a wavelength of 0.385 m which is close to the filtering radius. It should also be noted that the peak corresponding to the gap between the cobbles on the Belgian paving has been completely smoothed in the rigid wheel filtered data.

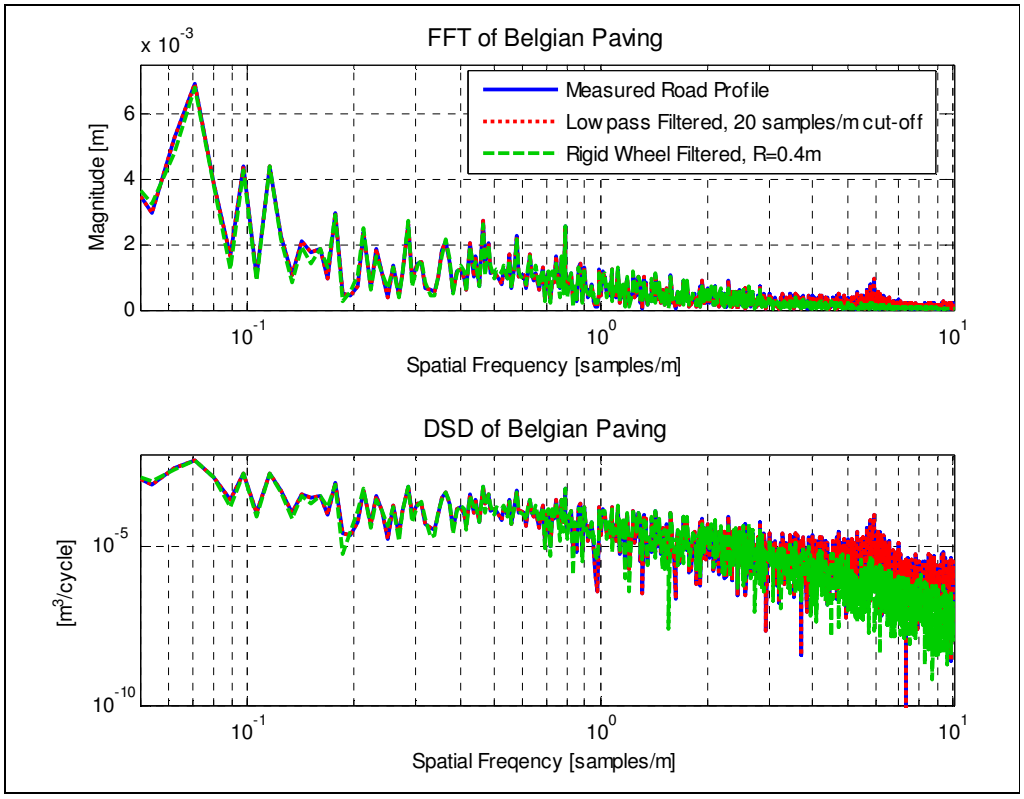

Figure 3-19: FFT and DSD of Belgian paving with different filtering techniques.

The effect of the filtering on the velocity input to the tyre is also of interest, since this will influence the tyre damping forces. The velocity input to the tyre will be proportional to the gradient of the road displacement data and the vehicle speed. The displacement data was thus differentiated in order to find the gradient of the road displacement gradients. The displacement gradients as well as their FFTs and DSDs have been plotted in Figure 3-20 and Figure 3-21. The gradients clearly show the effects of the filtering. The magnitudes of the peak gradients for the measured data go up to 2, while the magnitudes for the peak gradients of the low-pass filtered data, do not exceed 1. The rigid wheel filtered data leads to gradients which do not exceed 0.4. The extent of the effects of the filtering is echoed by the frequency domain plots in Figure 3-21. It can be seen from the velocity results that the filtering will have an even more profound effect on tyre damping forces than on tyre stiffness forces.

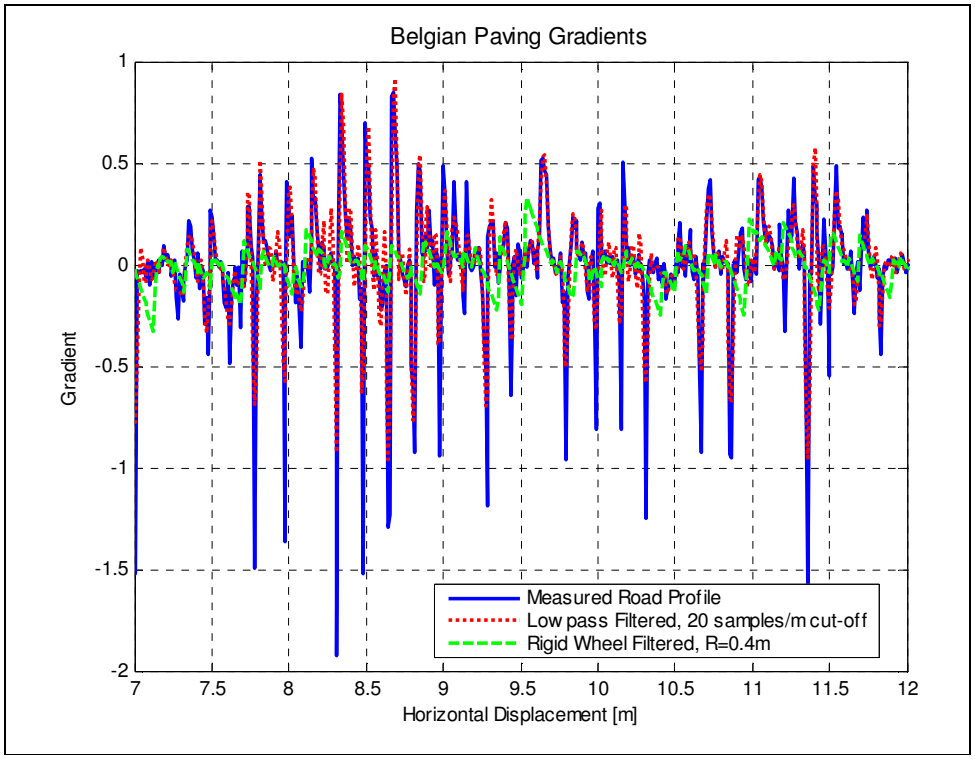

Figure 3-20: Belgian paving gradients.

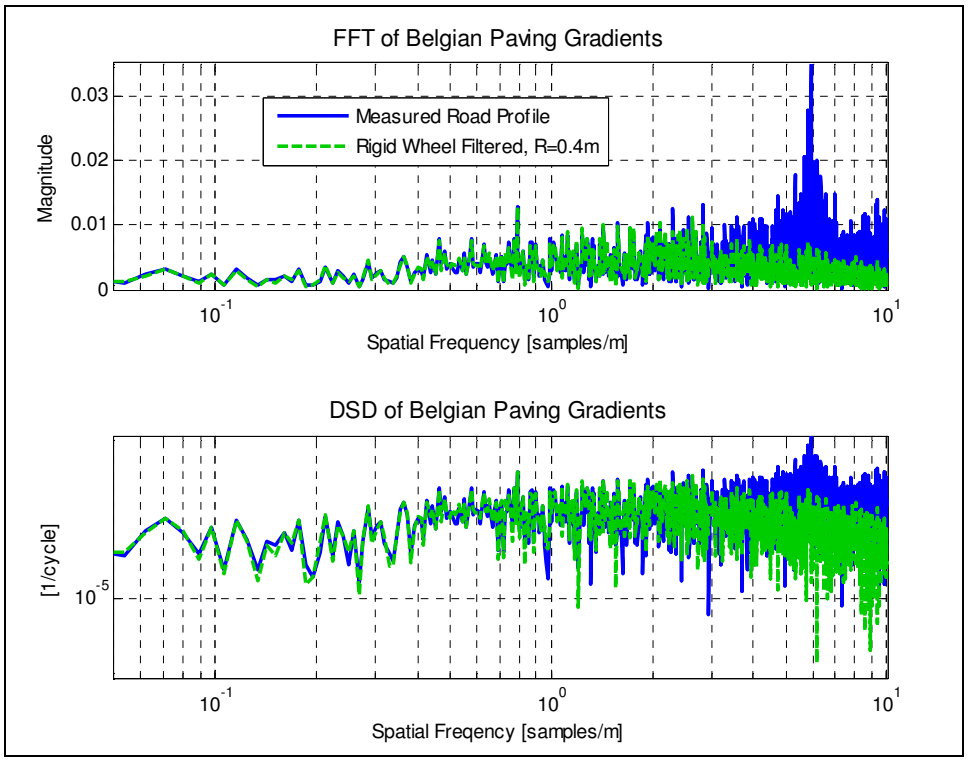

Figure 3-21: FFT and DSD of Belgian paving gradient.

3.3 Damage prediction

The aim of the study is to optimise vehicle suspension characteristics for increased structural fatigue life. The requirement thus exists, that the mathematical model must be able to predict structural life or structural damage in some manner. The assumption is made that failure in the vehicle chassis will be caused by high cycle fatigue. The damage estimation process used will therefore be stress based.

The process of estimating life cycle prediction in our case consists of two parts. The first is the prediction of the load-time history. The second is a cycle counting, in which the severity of the load history is determined in terms of equivalent constant amplitude cycles. The final component of fatigue life estimation consists of a damage summation, where the counted cycles are translated to damage to the structure.

The focus first moves to the prediction of the load-time history of the suspension strut mounting. During the experimental work, the load-time history of the left rear suspension mounting strut was measured using six strain gauges. The strains measured by these strain gauges would have to be predicted by the mathematical vehicle model. The intention is not to accurately predict the complex three-dimensional stress state or expected life in the left rear suspension mounting tower. It is rather intended to find the measured strain channel which is most representative of the stresses in the suspension strut mounting due to the vertical forces applied to this mounting.

The assumption was made that the structure will not be loaded beyond its linear elastic range. Loading beyond the linear region of the structure would constitute failure. A further assumption is made, namely that the first natural frequency of the suspension mounting is well above the low range of frequencies of interest that can typically be influenced by suspension spring and damper characteristics (<25 Hz). This implies a quasi-static analysis in which the structures dynamic response is negligible, can be applied. These assumptions allow us to model the strains in the structure by a direct proportionality to the forces applied to the structure as given below. A physical analogy of this may be found in a linear spring, where the elongation (strain) is proportional to the applied force by a constant (stiffness).

$$F_{sr} = k_{ef}\,\varepsilon$$

$$\therefore \varepsilon = \frac{F_{sr}}{k_{ef}}$$

where,

ε	=	strain,
F_{sr}	=	rear suspension force, and
k_{ef}	=	the effective stiffness of the suspension mounting.

It was decided to identify the measured strain channel which most accurately represents the strains in the suspension mounting, caused by the vertical components of the suspension forces. Cross-correlation functions were used as the metric for distinguishing the suitability of the strain channels for our purpose.

Cross-correlation functions are used to describe the general dependence between two sets of random data. The cross-correlation coefficient between two time dependent sets of data, $f(t)$ and $x(t)$, for an observation period T are defined as:

$$R_{fx} = \lim_{T\to\infty} \frac{1}{T} \int_0^T f(t)x(t+\tau)dt$$

The integral given above is evaluated for different values of the time delay, τ, in order to find the cross-correlation function, $R_{fx}(\tau)$. The interdependence of the two sets of data is highest where the cross-correlation function approaches unity. Matlab's built in cross-correlation function, crosscorr.m, was used to calculate the cross-correlation for all six strains with the left rear suspension force. This was done for the experimental data for runs 01 to 04 over the Belgian paving. The results of the calculation are presented in Figure 3-22 to Figure 3-25.

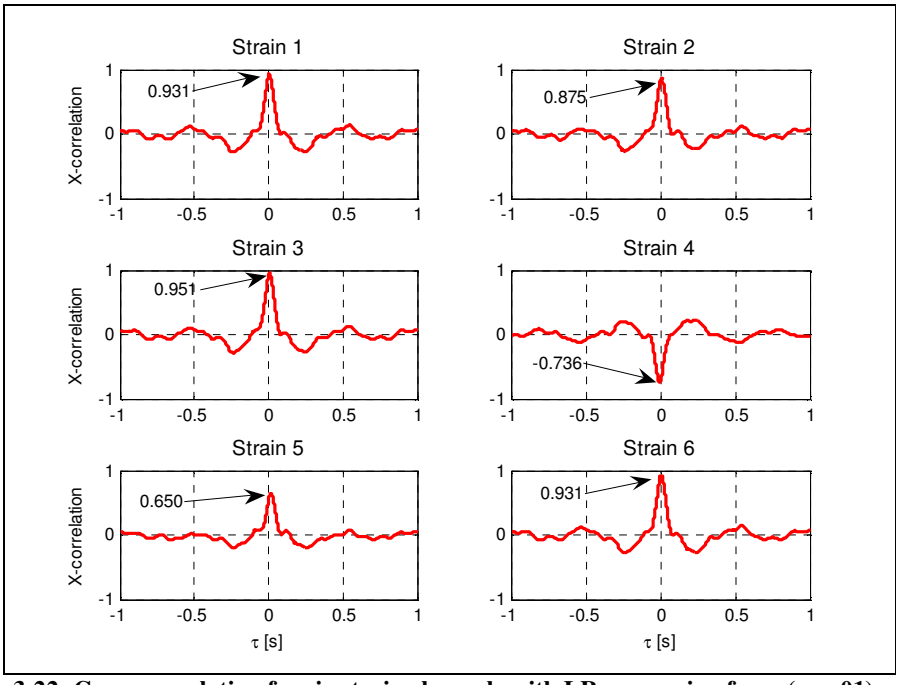

Figure 3-22: Cross-correlation for six strain channels with LR suspension force (run 01).

The cross-correlation coefficients for run 01 (Figure 3-22) indicates that strain channels 1, 3 and 6 are good candidates for determining the equivalent stiffness of the suspension mounting. Strain channels 2, 4 and 5 have peak correlation coefficients which are much lower than unity, indicating that these strain channels may be sensitive to either lateral or longitudinal suspension forces.

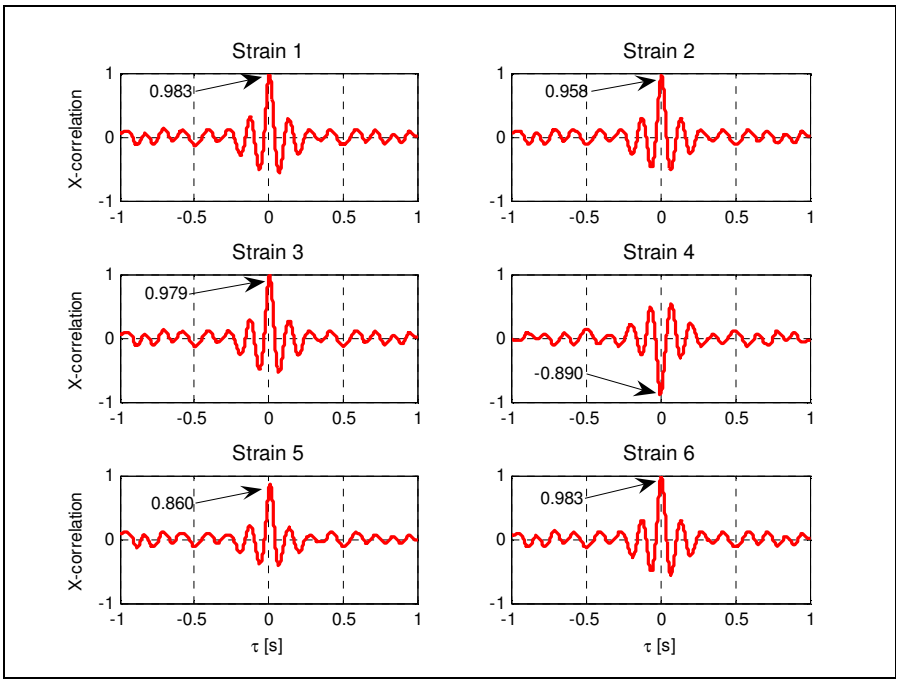

Figure 3-23: Cross-correlation for six strain channels with LR suspension force (run 02).

A similar result is seen with the data from run 02 (Figure 3-23), with channels 1, 3 and 6 again proving good candidates. This performance is repeated for run 03 and run 04 as indicated in Figure 3-24 and Figure 3-25. Strain channel 3 was chosen for determining the effective stiffness of the suspension mounting. This channel performed well in all four test runs and gave the best correlation for three runs (run 01, run 03 and run 04).

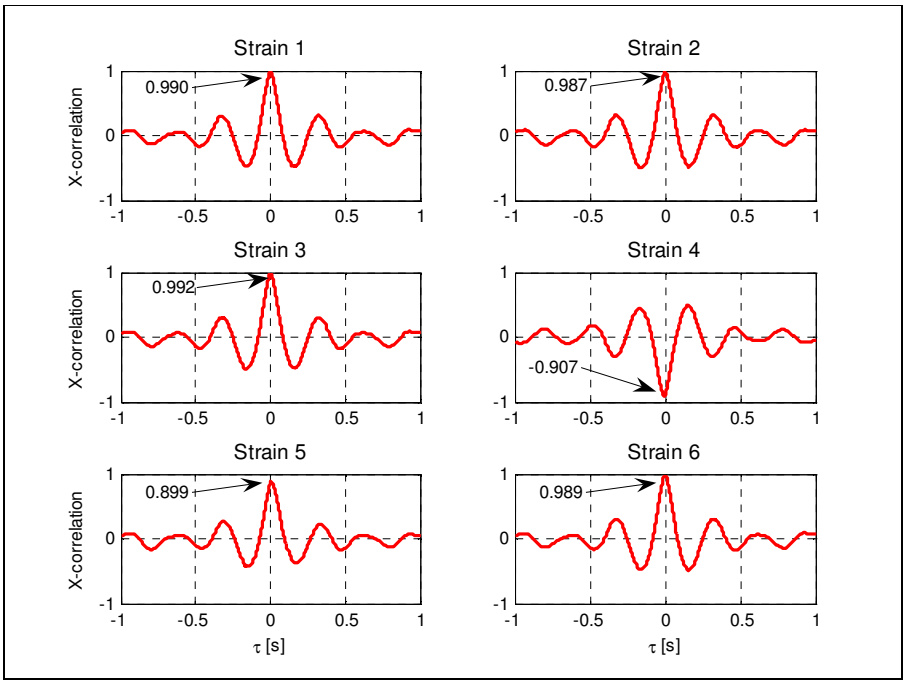

Figure 3-24: Cross-correlation for six strain channels with LR suspension force (run 03).

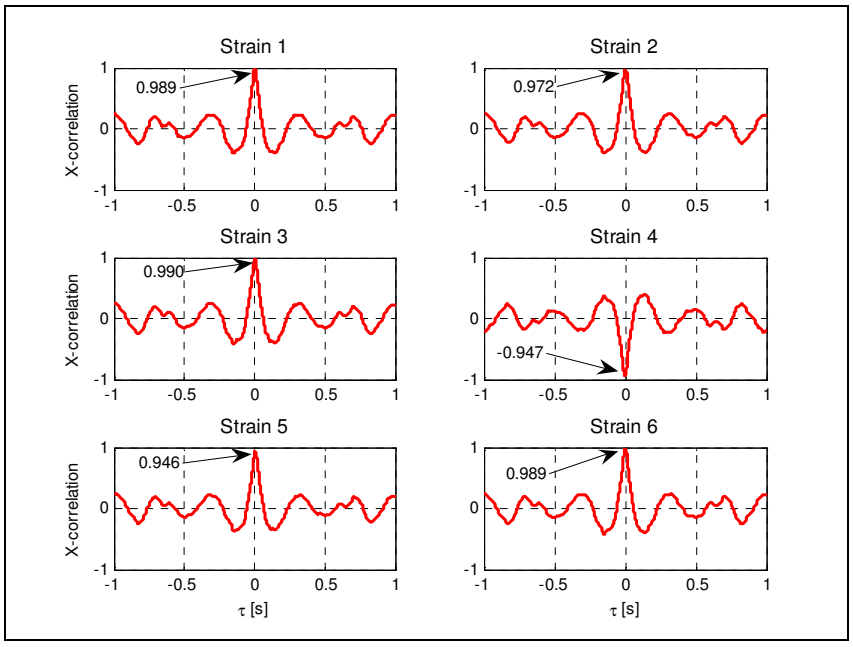

Figure 3-25: Cross-correlation for six strain channels with LR suspension force (run 04).

The most suitable strain channel for determining the effective stiffness has been determined, but the task of finding the effective stiffness remains. The strain versus suspension force data for runs 01 to 04 is plotted in Figure 3-26. The data sets have been fitted with linear relationships for which the gradients are given in the plots. The reciprocal of the average of these gradients will be taken to represent the effective stiffness of the suspension mounting. The mean of effective stiffness is $k_{ef} = \frac{1}{2.228\times10^{-8}} = 4.488\times10^{7}$ N/(m/m) with a COV of less than 6% .

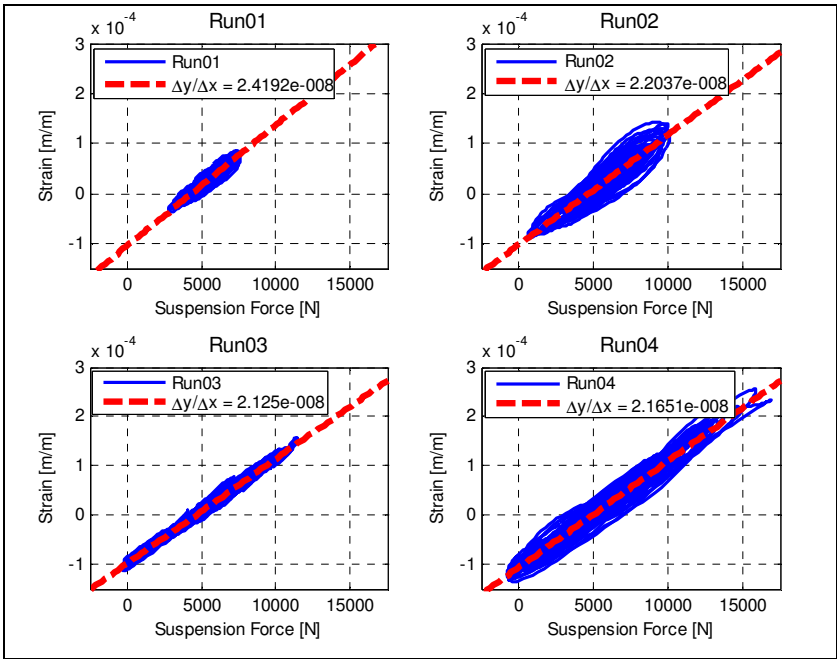

Figure 3-26: Strain 3 versus LR suspension force for four runs.

The effective stiffness allows us to determine strain load-time histories. The conversion of strain to stress histories again leads us to invoke the linearity assumption made earlier. The stress is found from Hooke's law for uni-axial stress, $\sigma = E\varepsilon$ (Gere, 2004). The relation relates the stress, σ, to strain, ε, by Young's Modulus of Elasticity, E. The modulus of elasticity of the suspension mounting is to be determined if we are to use Hooke's law in the conversion of strain histories to stress histories.

The Land Rover's rear suspension mounting is fabricated from steel sheet metal. The exact grade of material is however unknown. This presents a problem in finding the materials Modulus of Elasticity. Fortunately, the Modulus of Elasticity of a material is dependent on the molecular structure of the material. The vast majority of steels therefore have elasticities which range between 195 and 210 GPa. A typical value for the Modulus of Elasticity of plain carbon, low alloy steels will thus be used for the unknown suspension mounting material. According to Callister (2003), the modulus of elasticity of a typical plain carbon, low alloy steel is 207 GPa.

Now that we have developed a method for predicting the stress load-time history, we can progress to the second part of the damage estimation process, which is cycle counting. Cycle counting is generally done in the time domain. Rain-flow counting or Range-Pair-Range counting are methods of time domain cycle counting. Frequency domain methods have also been developed. The frequency domain methods make deductions from the power spectral density (PSD) curves of stress or strain histories, in order to estimate the expected number of cycles at specified stress ranges. The frequency domain methods are used infrequently as they are very sensitive to the frequency ranges and frequency increments of the PSD they are applied to (Haiba et. al., 2002). Time domain cycle counting will therefore be used in this study.

The Palmgren-Miner rule, developed by A. Palmgren and made popular by M.A. Miner, is a frequently used time-domain method of estimating damage due to variable amplitude loading (Dowling, 2007). The method uses a stress-life cycle or S-N curve in order to predict the damage which would be caused to a structure by a load cycle. The Palmgren-Miner rule is simple to implement and reliable and will be used for the damage estimation.

S-N curves can be represented by a power law of the form $\sigma = AN_f^{-B}$. The relation relates the number of cycles to failure, N_f, to the constant amplitude stress, σ, which would cause the failure at the number of cycles in question. The constant, B, represents the gradient or slope of the S-N (stress versus number of cycles) curve on a log-log plot. The constant, A, then scales the S-N curves magnitude. The constants, and thus the shape of the S-N curve, are functions of material properties or manufacturing techniques.

The problem of the unknown material properties of the suspension mounting now re-emerges. The problem of determining the fatigue life constants, A and B, cannot be circum-navigated as easily and robustly as was done for the modulus of elasticity. Experimental methods of determining the constants would be far too expensive and time consuming for the purposes of this study. The reader should recall that the study

does not aim to give exact predictions of the component life, but rather intends to find a good representation of the vehicle structural life behaviour as a function of suspension characteristics. A different solution was thus required.

The Standards South Africa (The South African Buro for Standards) has published a standard on the structural use of steel. The SANS 10162-1:2005 (Standards South Africa, 2005) code presents a method of dealing with the fatigue of welded joints. The method assumes weld fatigue behaviour described by the power law relation of the form described in the previous paragraphs. The code suggests that the slopes of the stress-life curves of most types of welded joints (and thus the constant, B) can be assumed to be equal. The code advocates that using a value of $1/3$ for the constant, B. The difference in the fatigue behaviour of different welded joints is then solely accounted for by the constant, A. Values for A, are prescribed according to the classification of the weld. The stress ranges of the load-time histories are used in this method and the effects of mean stress are neglected.

A class E weld was chosen arbitrarily for the fatigue life prediction in the study. The existence of a fatigue limit was neglected. The S-N curve used for the class E weld, together with the values of the constants A and B, is presented in Figure 3-27.

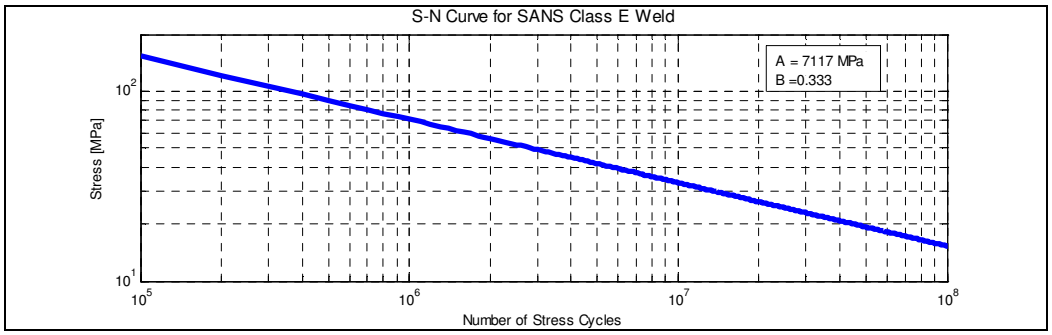

Figure 3-27: S-N curve for SANS class E weld.

3.4 Summary and conclusions to mathematical modelling

The construction of the mathematical model of the experimental vehicle was discussed in this chapter. The model was developed for use in mathematical optimisation. The modelling philosophy followed, was to keep the model as simple as possible, in order to be computationally efficient, yet to include as much complexity as was necessary to sufficiently represent reality. The modelling assumptions will be submitted to the rigorous scrutiny by validation with experimental results in Chapter 4, before applying it to its intended goal of mathematical optimisation in Chapter 5.

Chapter 4

4. Experimental results and model validation

Several field tests were conducted in order to capture and characterise the dynamics of the test vehicle, as was discussed in Chapter 2. Great effort was expended into ensuring the quality and repeatability of the experimental data. A mathematical model was then constructed in order to predict the experimental vehicles dynamic response (Chapter 3). This current chapter presents the experimental results together with simulation results in order to validate the mathematical model.

The results are presented in two parts. The first deals with the model validation over the discrete obstacles. Comparison for this data mainly occurs in the time domain. The second part of the model validation compares the model simulation results with the experimental data from the tests conducted over the Belgian paving. This data is best compared in the frequency domain.

During the field tests, suspension displacements, suspension forces, body pitch velocity and body accelerations were measured. The model correlation for all measured parameters was considered. The correlation in this chapter is presented for selected parameters only. The limited presentation was necessitated due to the sheer volume of data available. Complete results are given in Appendix C. The correlation of body vertical acceleration is given in this chapter since this quantity is often used in the correlation of vehicle models. The suspension forces and damage calculated from the measured strains will also be compared with predicted damage obtained from the vehicle mathematical model, due to its importance to the current study. Other data will also be presented here as required. Since the repeatability has been proven to be of exceptional quality in Chapter 2, only one data set for each obstacle or road surface has been considered for correlation.

4.1 Correlation over discrete obstacles

The experimental data over the discrete obstacles is compared with the model for both symmetric and asymmetric obstacle configurations. The first set of results which is considered is for the first trapezoidal bump configuration (see Figure 2-6). The test layout saw the Land Rover first traversing a pair of large and subsequently a pair of small trapezoidal bumps. The trapezoidal bumps were of symmetric configuration. The measured and simulated results for run 16 have been superimposed in Figure 4-1 and Figure 4-2. The run was conducted at relatively low speed (1st gear, low range, ±14.5 km/h) with the 4S_4 set to the ride comfort setting. The average of the measured drive shaft speed and the VBOX speed was used as input to the model.

The measured and simulated vertical body accelerations are shown in Figure 4-1. Visually, there is excellent correlation between the trends of the measured and simulation results. The magnitudes of the peaks and valleys also seem to correspond well. The phase difference in the acceleration data for the second bump is caused by a slight velocity mismatch.

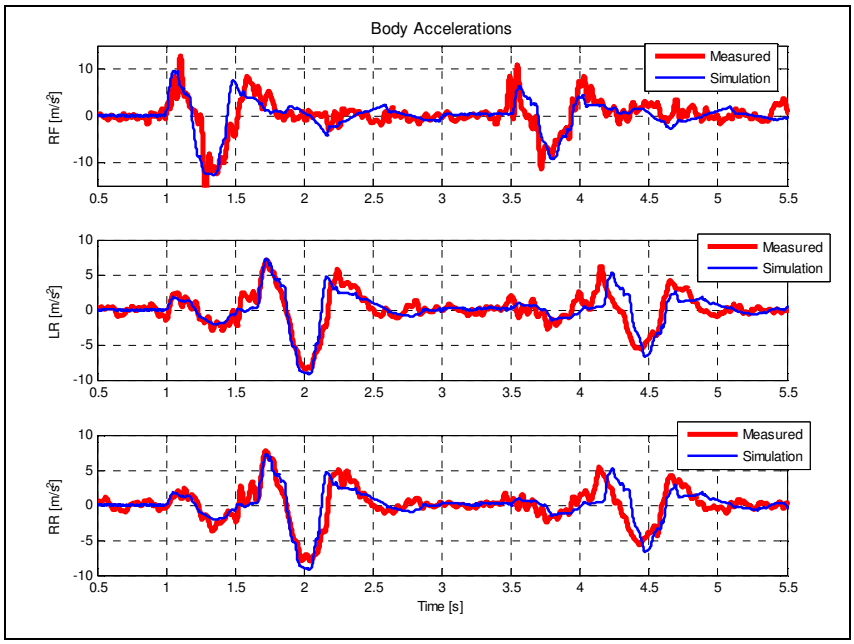

Figure 4-1: Body accelerations (run 16).

The correlation for body pitch and roll velocities as well as suspension displacements and suspension forces is of similar quality as the acceleration data. The predicted stresses in the left rear suspension mounting correlated very well with the measured values, as can be seen in Figure 4-2. This is especially true, when the simple stress prediction technique is considered. The trends in the simulated results and the peak values in the measured data match quite closely. The stresses for the run varied between -20 and 15 MPa, implying a maximum stress range of 35 MPa. The model's correlation over symmetric discrete obstacles, with the suspension set to ride comfort mode is very good.

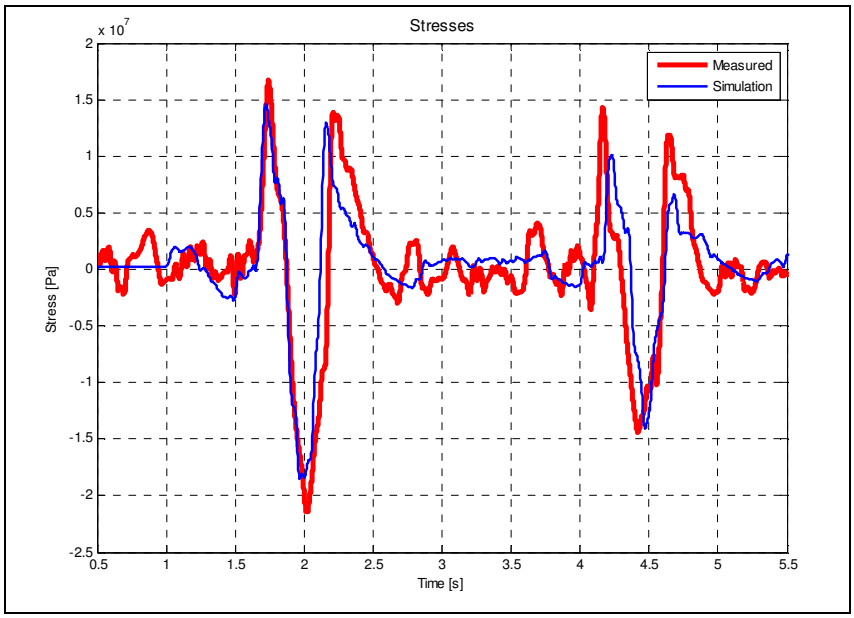

Figure 4-2: Measured and predicted stresses (run 16).

The second set of data which will be considered is for the second trapezoidal bump course configuration. In this set of tests, the test vehicle was driven over a single pair of smaller symmetrically arranged trapezoidal bumps (see Figure 2-7). The suspension $4S_4$ suspension system was set to the handling or hard mode. The tests were conducted at the same speed (±14,5 km/h) as over the previous discrete obstacle course configuration. The test data for run 24 was taken to represent this set of tests and is presented, together with the simulation results in Figure 4-3 through to Figure 4-6.

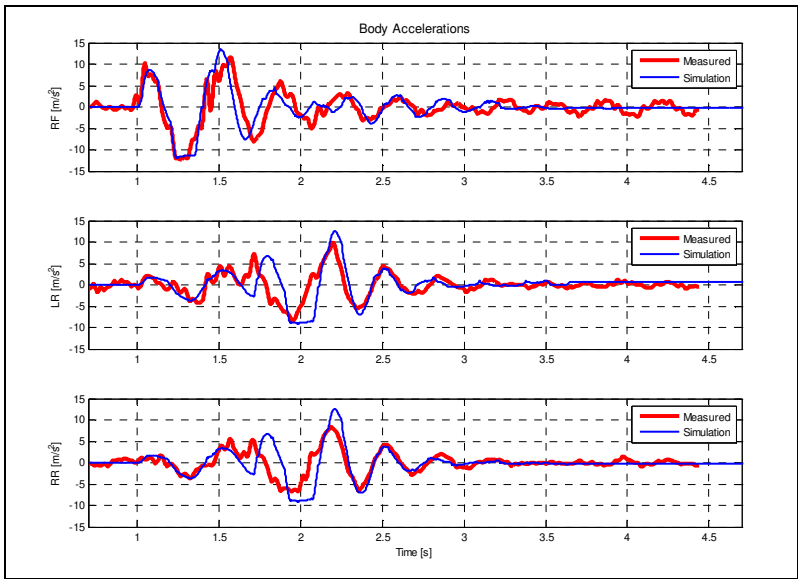

Figure 4-3: Body accelerations (run 24).

The vehicle body accelerations are plotted in Figure 4-3. The correlation for the rear vertical accelerations for run 24 is not of the same quality as was observed for run 16. The trends in the measured and simulated data are similar, but the peak accelerations are over-predicted. The situation is somewhat better in the case of the front vertical accelerations. The discrepancy is due to a loss of wheel-ground contact in the model (at about 2 seconds) which does not occur to such a severe extent in the actual vehicle. The point contact of the tyre model is thought to worsen the situation. The vehicle model becomes much more sensitive to tyre-road modelling parameters as the suspension stiffness is increased. The increased sensitivity is further aggravated by the problem's sensitivity to vehicle longitudinal velocity. The correlation for body vertical accelerations for run 24 is however considered satisfactory.

The measured and simulated suspension displacements for run 24 are superimposed in Figure 4-4. The simulated results show similar trends to the measured data. There is however some discrepancy in the magnitudes. The discrepancies seem quite large relative to the displacements observed in the measured results. The errors range from between 5 to 10 mm in magnitude, which is well within the accuracy bounds to which the exact terrain profile is known. Furthermore, the initial displacements could have been different due to the hysteretic effects of the friction in the suspension. The errors are thus considered acceptable. An error in the measured road profiles could therefore also result in the poor correlation, body pitch velocity and body accelerations. Again, the vehicle dynamics with the $4S_4$ set to hard seems very sensitive to uncertainty in the road profile.

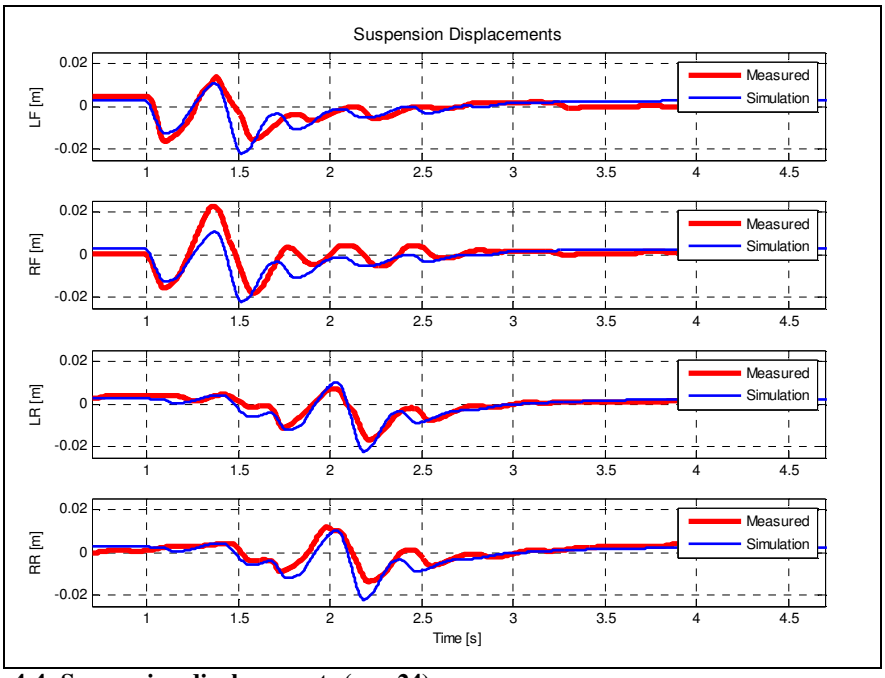

Figure 4-4: Suspension displacements (run 24).

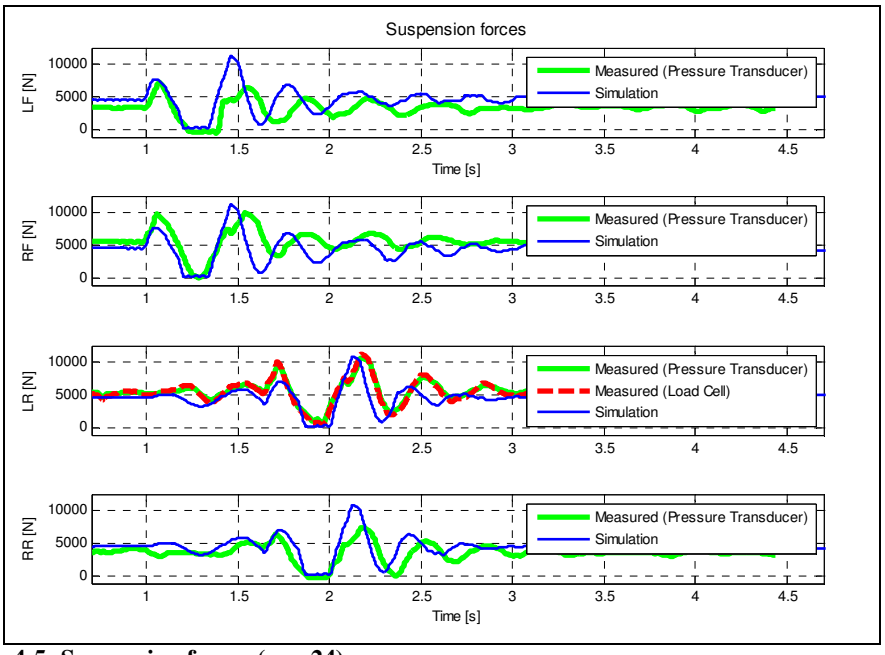

Figure 4-5: Suspension forces (run 24).

The predicted suspension forces are given in Figure 4-5, while stresses are given in Figure 4-6. The correlation for suspension forces is good, except for overshoot on the front suspension forces. The good correlation of the rear suspension forces is reflected in the stress prediction in Figure 4-6. The stresses range from -20 MPa to 25 MPa or a maximum stress range of 45 MPa. The accurate stress prediction gives some confidence in the model's likely ability to predict damage. The correlation of the model with the suspension set to handling mode over the discrete obstacles is thus considered satisfactory.

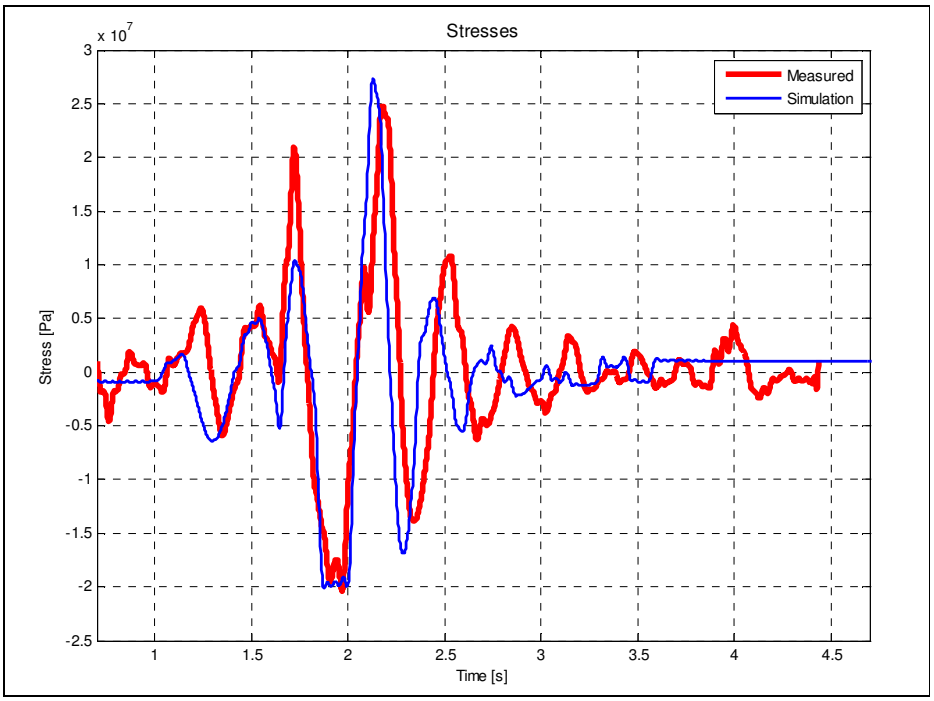

Figure 4-6: Measured and predicted stress (run 24).

The correlation of the model over asymmetrically arranged discrete obstacles is considered next. Run 30 was selected to represent the Land Rovers response over trapezoidal bump layout 3 (Figure 2-8), with the suspension set to ride comfort mode. The run was again conducted at ±14.5 km/h. The vertical accelerations for the run are presented in Figure 4-7. The correlation for vertical accelerations is reasonably good, as is the correlation for the vehicle pitch velocities shown in Figure 4-8. The same can not be said for the roll velocities, also in Figure 4-8. Although the trends are arguably similar, the model vastly over-predicts the roll velocity.

It was initially imagined that the small angle assumption, which was made in modelling the roll degree of freedom, was to blame. The body roll angle in the model does however not exceed 6 degrees, and is thus still small enough to warrant the small angle assumption. A sensitivity analysis was consequently conducted to determine the modelling parameters which could most likely cause the high roll velocities. The vehicle speed, masses, inertias, dimensions, tyre parameters, suspension damper and spring characteristics, suspension friction, suspension bushings and road profile were considered. The road profile gave the highest sensitivity followed by the vehicle sprung mass magnitude, roll inertia, speed and suspension static gas volume. Only limited improvement could be obtained by varying these parameters. Roll stiffness in the suspension bushings, which is not captured by the model, may be to blame for the high roll velocities predicted by the model. Frame flexibility, which is not modelled, could be a contributor, but this is not considered the main cause.

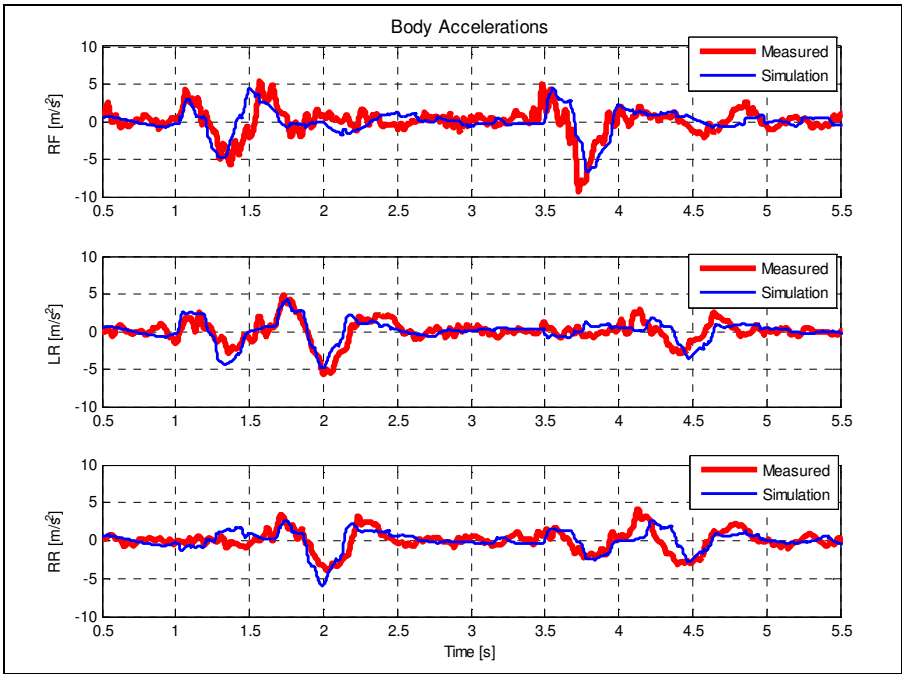

Figure 4-7: Body vertical accelerations (run 30).

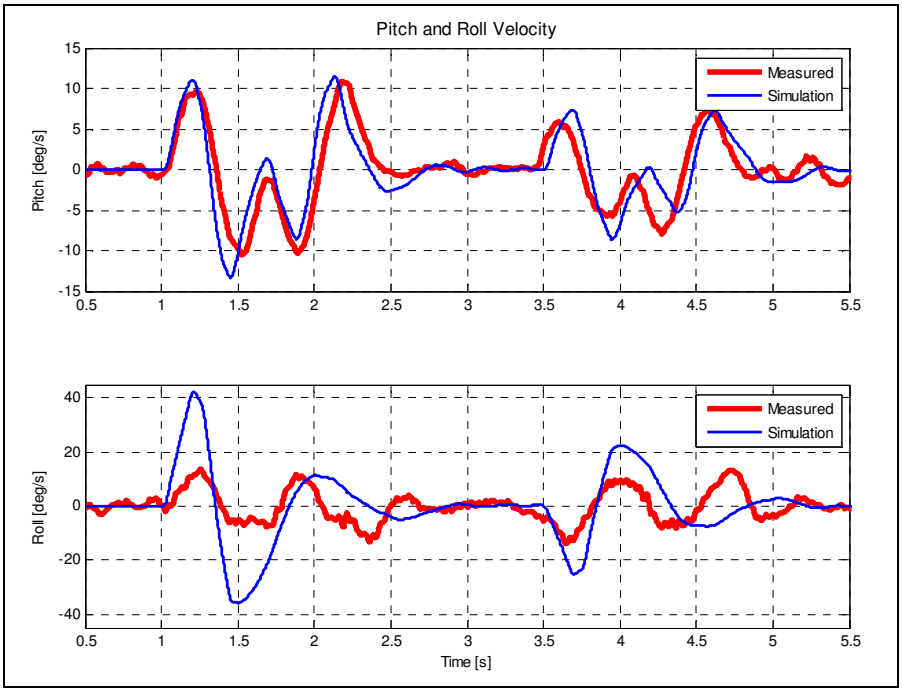

Figure 4-8: Body pitch- and roll-velocities (run 30).

The predicted stress for run 30 is given in Figure 4-9. The correlation for stress is adequate, with the simulation following the trends in the measured data closely, but failing to achieve the same peak values as seen in the measured data.

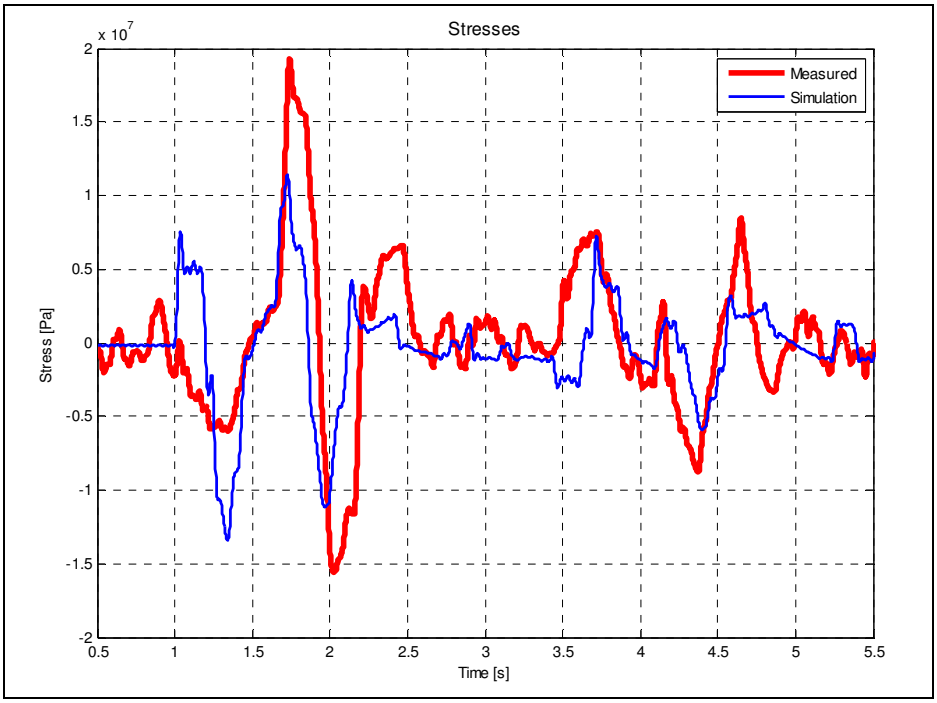

Figure 4-9: Measured and predicted stress (run 30).

The final set of data which is considered for discrete obstacles is that of run 31. The run was also conducted over the third trapezoidal bump layout at a speed of ±14,5 km/h. The $4S_4$ was set to handling mode. The measured vertical accelerations are plotted in Figure 4-10. The correlation is poor. The trends in measured and simulated data are seemingly consistent, but the peak values cannot be compared. The factors detracting from the correlation of the runs in handling mode and the runs over the asymmetric obstacles are compounded here. Over-prediction of the roll velocity is again to blame.

The sensitivity analysis performed for run 30 was repeated for run 31. The correlation was again found to be most sensitive for the road profile. Vehicle speed plays a significantly more dominant roll when the suspension stiffness is increased. The exact width between the tyre contact patches, the vehicle sprung mass and the static gas volume complete the list of the five most critical parameters from the sensitivity analysis. One can see that the tyre-road interaction plays an important role in modeling the vehicle dynamics.

The predicted stress is superimposed on the measured stress in Figure 4-11. Considering the poor correlation observed for the vehicle dynamics on run 31, the predicted stress corresponds quite well to the measured data. The stress prediction is considered adequate for predicting the trends in fatigue damage to the vehicle structure.

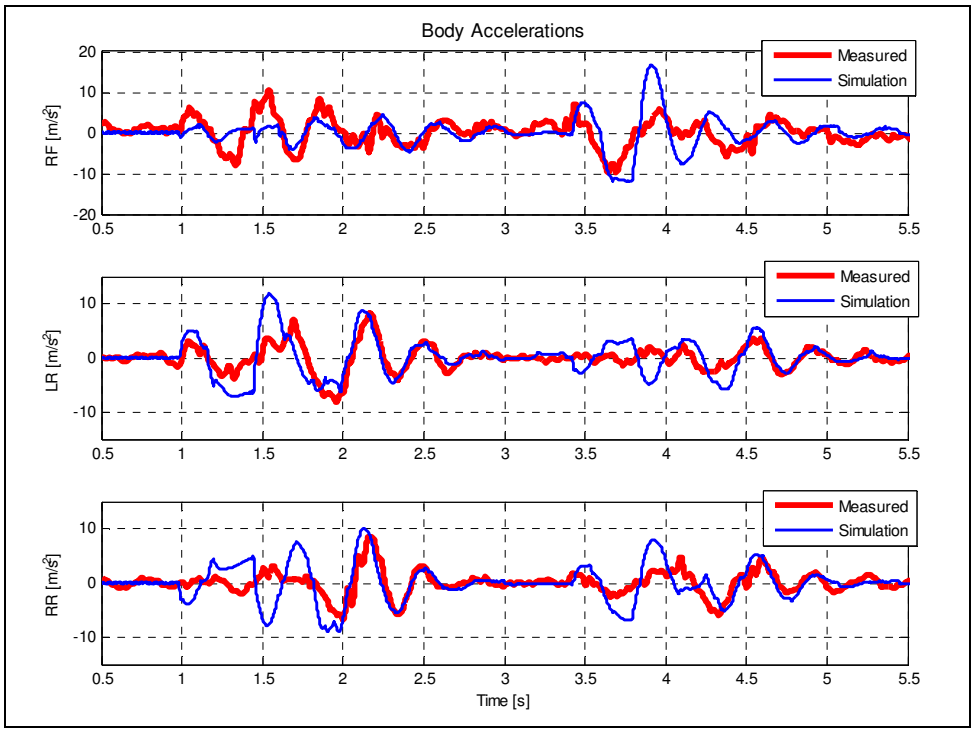

Figure 4-10: Vertical accelerations (run 31).

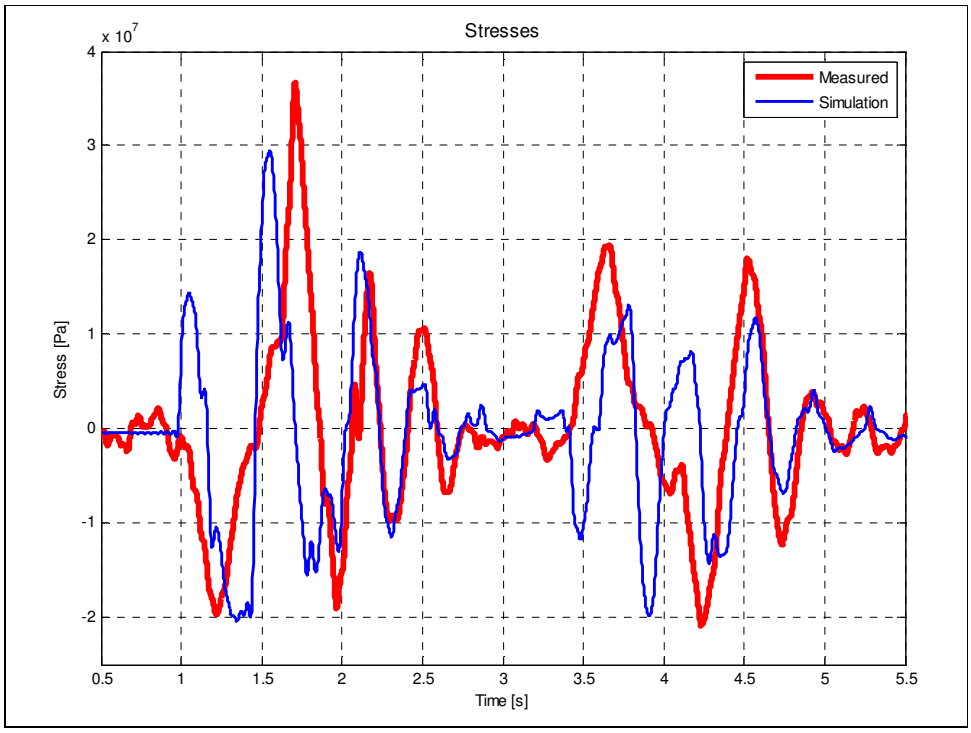

Figure 4-11: Measured and predicted stress (run 31).

4.2 Correlation over rough terrain

The comparison of data recorded over rough terrain is exceedingly difficult, if considered in the time domain. The model validation for rough terrain was thus conducted, solely using frequency domain information or metrics such as root mean squared values (RMS). As discussed in Chapter 2, the rough terrain field testing of the Land Rover was conducted for combinations of two speed conditions and the two suspension modes. The high repeatability of the data allows the consideration of a single data set for each combination of speed and suspension mode. The data for runs 01, 02, 03 and 04 will be considered in the correlation. The correlation results for vertical accelerations, suspension forces and stress are given here, but the complete correlation results have been included in Appendix C. Additional data is presented as required.

Run 01 was conducted at low speed (1st gear, low range, ±14.5 km/h) with the 4S$_4$ set to the ride comfort or soft mode. The speed was kept constant by driving against the diesel engine's governor. The FFTs of the measured and simulated vertical accelerations for run 01 are shown in Figure 4-12. The simulation results show good correlation up to 5 Hz. The body natural frequency at the rear is captured almost perfectly at 1.8 Hz. The peak between 1.8 and 2.2 Hz on the front accelerometer is attributed to a combination of the body pitch and bounce natural frequencies. The model contains higher noise between 5 and 10 Hz than is present in the actual vehicle. The high noise at frequencies over 5 Hz is thought to be due to a combination of the static friction model employed, the point follower tyre model employed and inaccuracy in predicting the body roll.

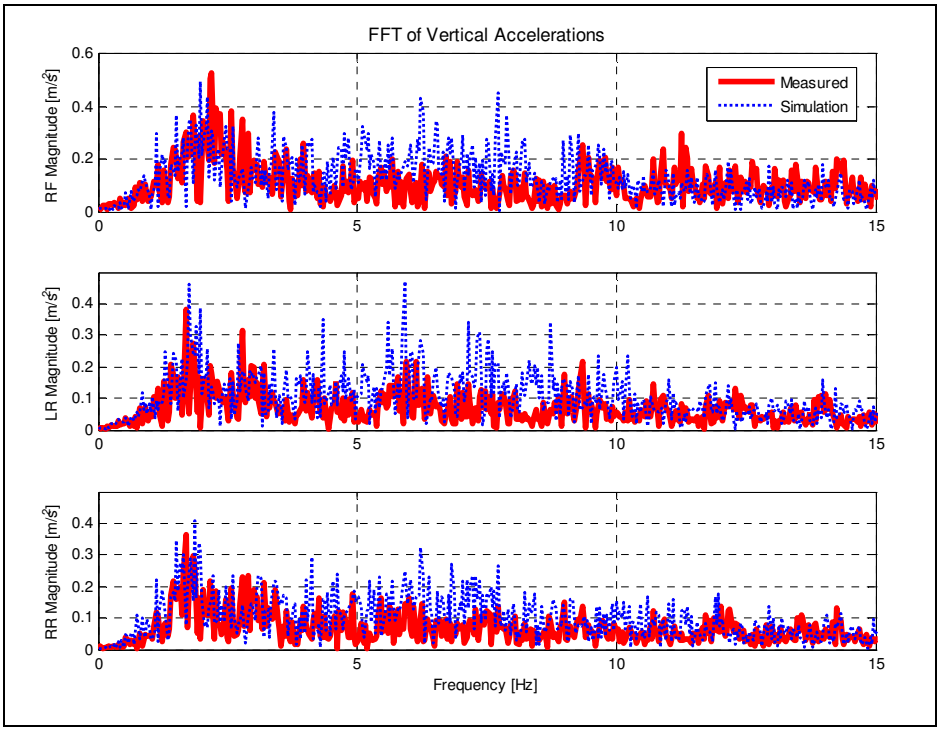

Figure 4-12: FFT of vertical accelerations (run 01).

The FFTs of the measured and simulated pitch- and roll velocity for run 01 are plotted in Figure 4-13. The FFT of pitch velocity reveals several decaying peaks, which have a similar appearance to sidebands which are usually observed in vibrating systems with decaying response. The frequency locations of the first five peaks which are observed in the FFT of the measured data are mirrored closely in the simulation results. The amplitudes and shapes of the first too peaks, at 0.8 Hz and 2.2 Hz, in the simulation results show good correlation with the measured data. The larger of these peaks between 2 Hz and 2.5 Hz corresponds to the pitch natural frequency. The noise in the simulation results at frequencies above 5 Hz can again be observed here.

The main peak on the roll velocity FFTs correlate well in terms of frequency location, but the roll velocity of the simulation is severely over-estimated. The over-estimation occurs mainly at frequencies below 5 Hz. The roll degree of freedom is thus only partially to blame for the increased noise in the data at frequencies above 5 Hz.

The FFTs as well as the histograms of the measured and predicted stresses for run 01 are given in Figure 4-14. The stress prediction again shows the correct trend, but with over prediction at frequencies higher than 5 Hz. The over prediction of stress is also observable in the stress histogram.

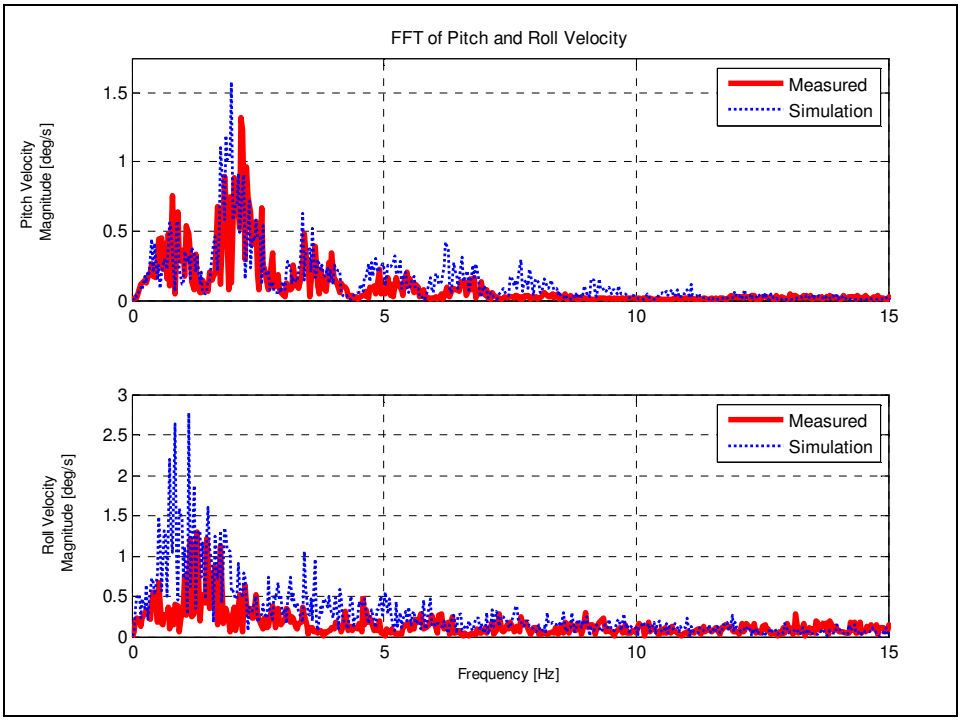

Figure 4-13: FFTs of pitch- and roll-velocity (run 01).

The correlation of the simulation data for run 01 generally seems reasonable for frequencies lower than 5 Hz, but is plagued by the presence of noise at higher frequencies. The static friction model is considered to contribute to the noise in the simulation above 5 Hz.

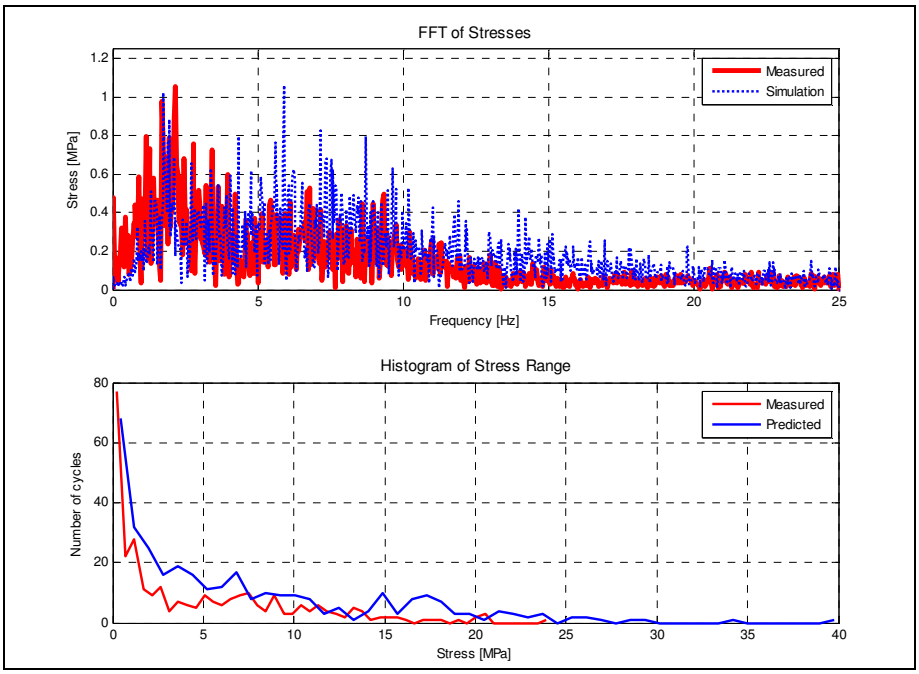

Figure 4-14: FFT and histogram of measured and predicted stress (run 01).

Run 02 was conducted at higher speed than run 01 (4^{th} gear, low range, ±54 km/h) with the $4S_4$ set to the ride comfort or soft mode. As in the case of run 01 the speed was kept constant by driving against the diesel engine's governor. The FFTs of the measured and simulated body vertical acceleration for run 02 are shown in Figure 4-15. The reader should note that the right front accelerometer was faulty during the remaining runs over the Belgian paving and the data will not be presented.

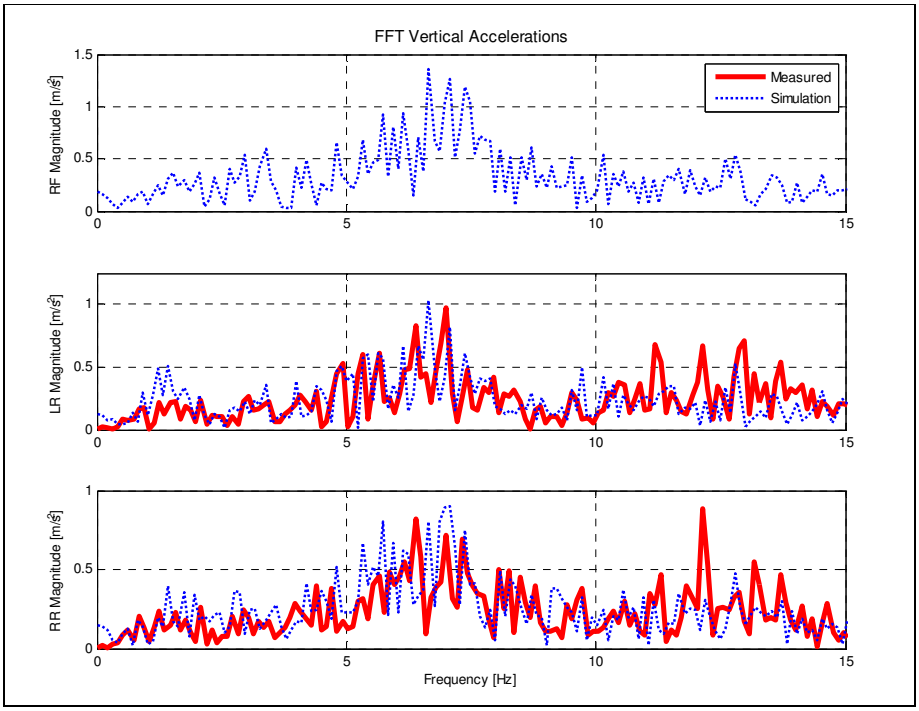

Figure 4-15: FFTs of vertical accelerations (run 02).

The body natural frequency is again visible in both measured and simulated data at 1.8 Hz. The wheel hop frequency can be distinguished from the measured right rear acceleration data at 12 Hz. The simulation data shows the same trend as observed in the simulation data up to about 10 Hz. The simulation again over-estimates the vertical acceleration in this region. At higher frequencies, the simulation response decreases below that predicted by the model and no longer follows the trends in the measured acceleration data. It is unclear whether the accelerations above 15 Hz in the measured data are due to the rigid body dynamics of the vehicle or due to the structural dynamics in the vicinity of the transducer location. The fact that this energy is not reflected in the suspension force FFT (Figure 4-16) seems to confirm the suspicion that there may be participation of the structural dynamics.

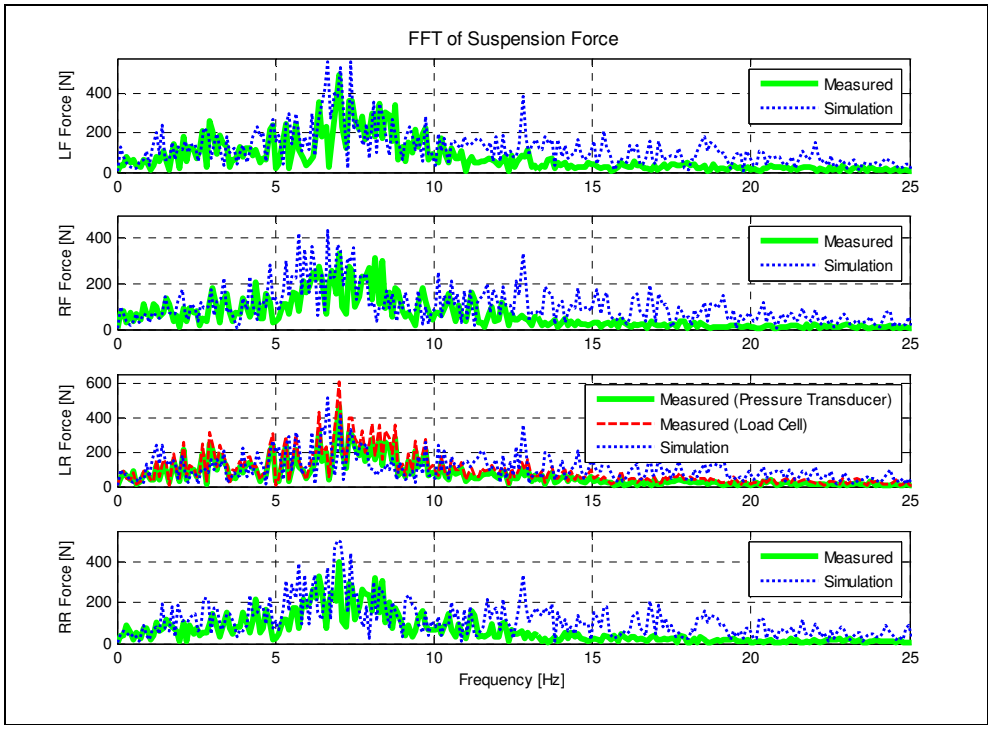

Figure 4-16: FFTs of suspension force (run 02)

The FFT of measured and predicted pitch- and roll-velocities for run 02 are shown in Figure 4-17. The pitch natural frequency between 2 Hz and 2.5 Hz can again be observed in the measured data. The correlation for the simulated pitch velocity with the measured data is very good. The model over-predicts the pitch velocity between 5 and 10 Hz. The roll-velocities from the simulation show poor correlation with measured results.

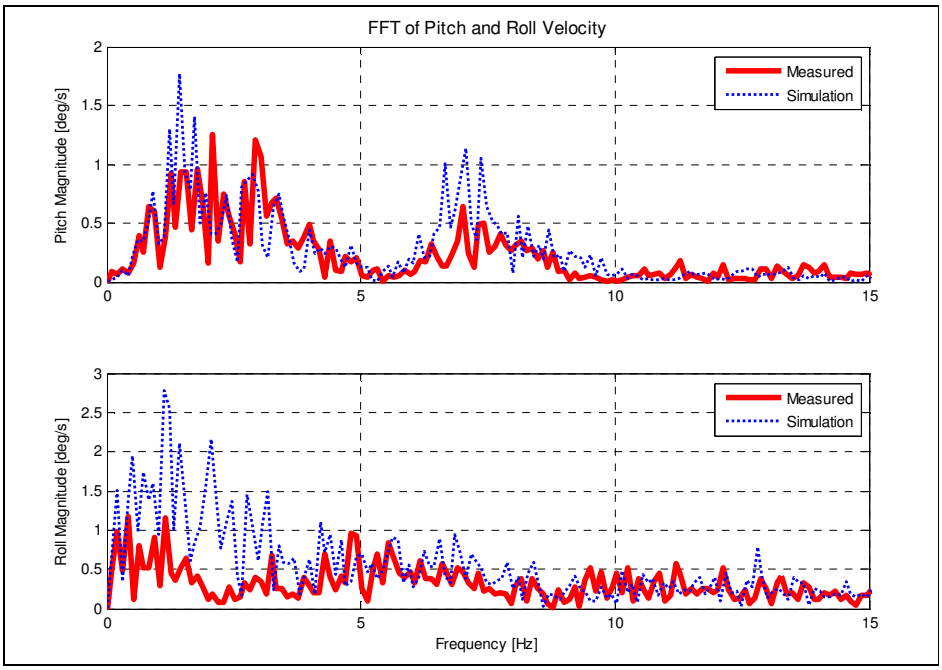

Figure 4-17: FFT of pitch- and roll-velocity (run 02).

The predicted FFT of stress also shows fair correlation with the measured results as seen in Figure 4-18. The stress is however under-predicted at low frequencies. The trends for the histogram of stress range correspond well, but the simulated results are once again over-predicted.

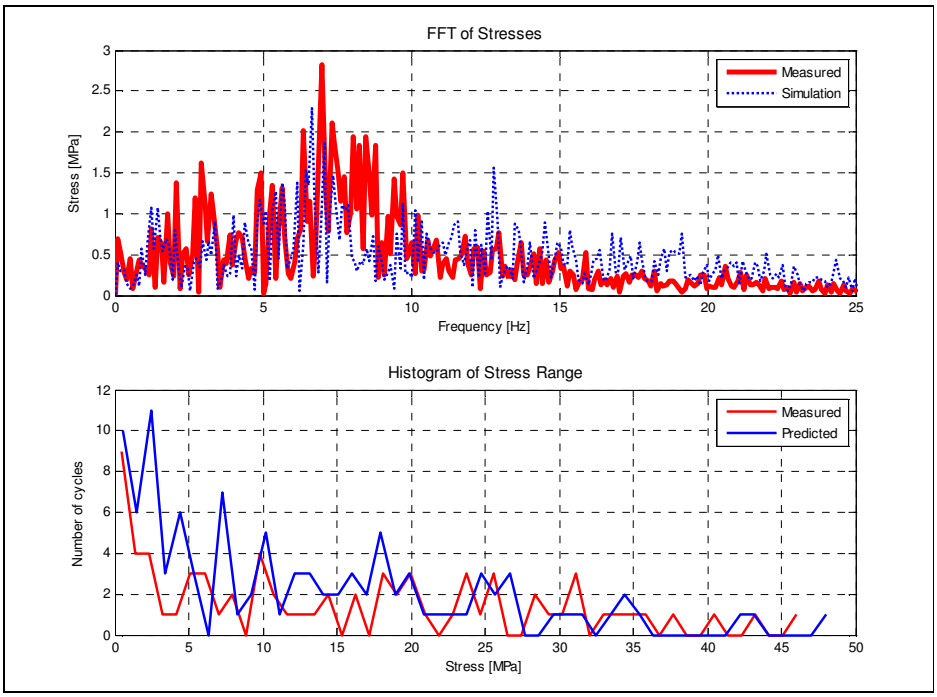

Figure 4-18: FFT and histogram of measured and predicted stress (run 02).

Run 03 was conducted at the same lower speed of ±14.5 km/h as in run 01. This time however, the $4S_4$ was set to the handling mode. The reader is made aware of the shift in the suspension natural frequency due to the change from ride comfort to handling mode. The new suspension natural frequency at 3.2 Hz shows up as a well defined peak in the FFT of rear accelerations given in Figure 4-19. The FFT of vertical accelerations obtained from the simulation data follows the trend of the measured results very well. The over-prediction between 3 and 10 Hz as observed in run 01 is however still present in this data set.

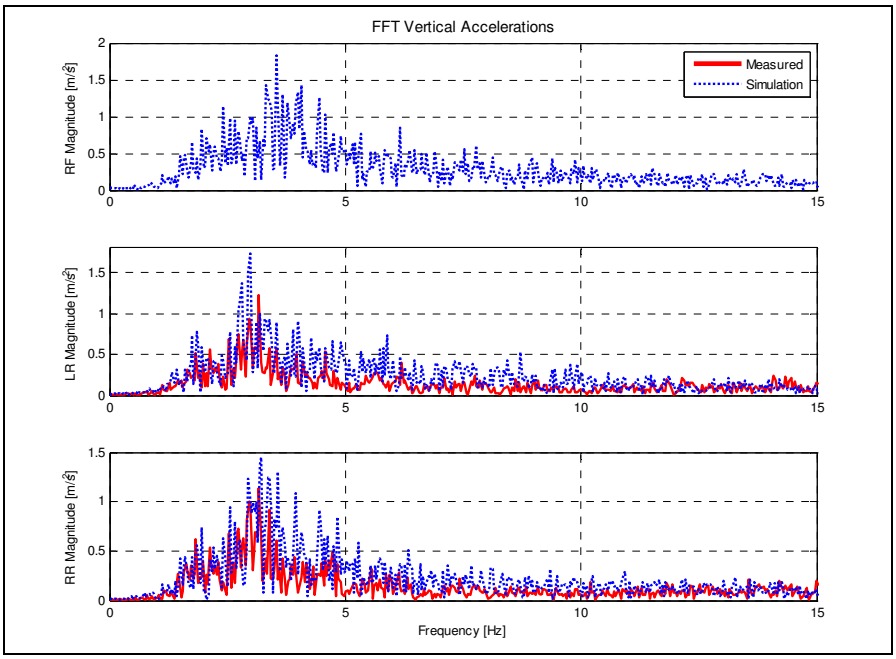

Figure 4-19: FFT of vertical accelerations (run 03).

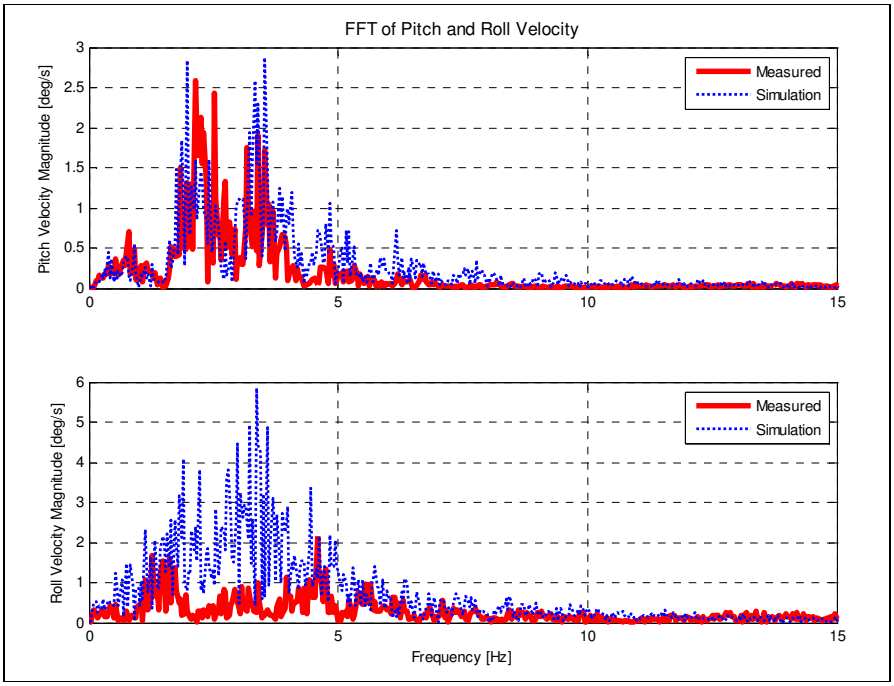

Figure 4-20: FFTs of pitch- and roll-velocity (run 03).

The FFTs of the simulated pitch velocities in Figure 4-20 show excellent correlation with the measured FFT. The correlation of the roll velocities is very poor, with the model over-predicting the roll response below 5 Hz.

The predicted stress FFTs and stress histograms in Figure 4-21 both show the same trends as the measured data. The model shows higher energy than the measurements above 5 Hz. The higher energy content is again attributed to a combination of tyre and friction modelling as well as roll effects.

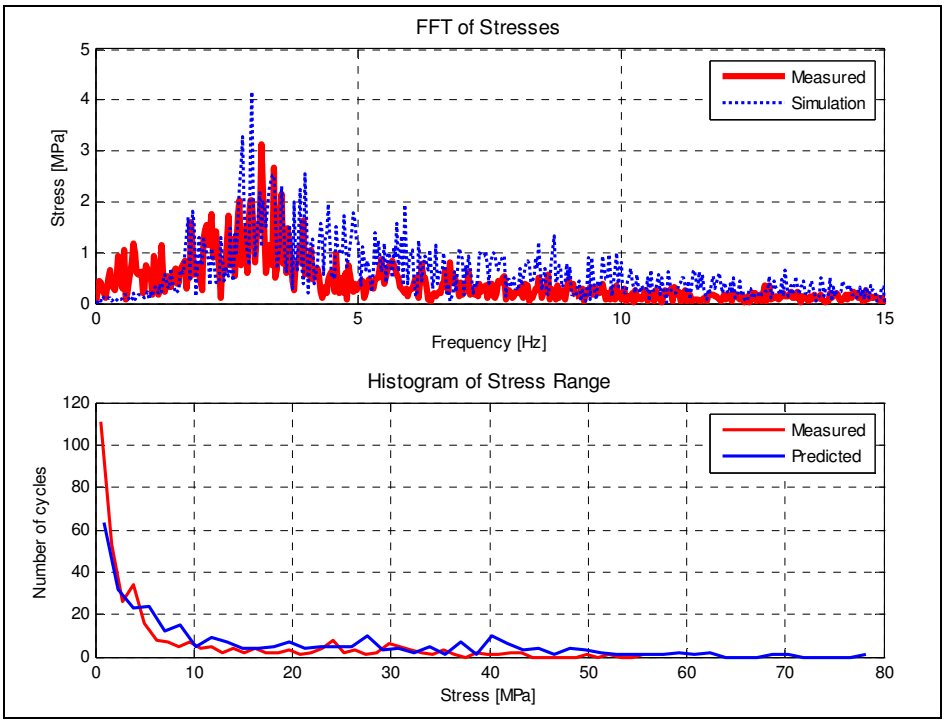

Figure 4-21: FFT and histogram of stress (run 03)

The final dataset considered is for run 04. The run was conducted at the higher speed of ±54 km/h as in run 02. The simulated FFT of left rear vertical accelerations follows the trend and magnitude of the measured FFT reasonably well. The suspension natural frequency at around 3 Hz is again observed in both simulated and measured FFTs.

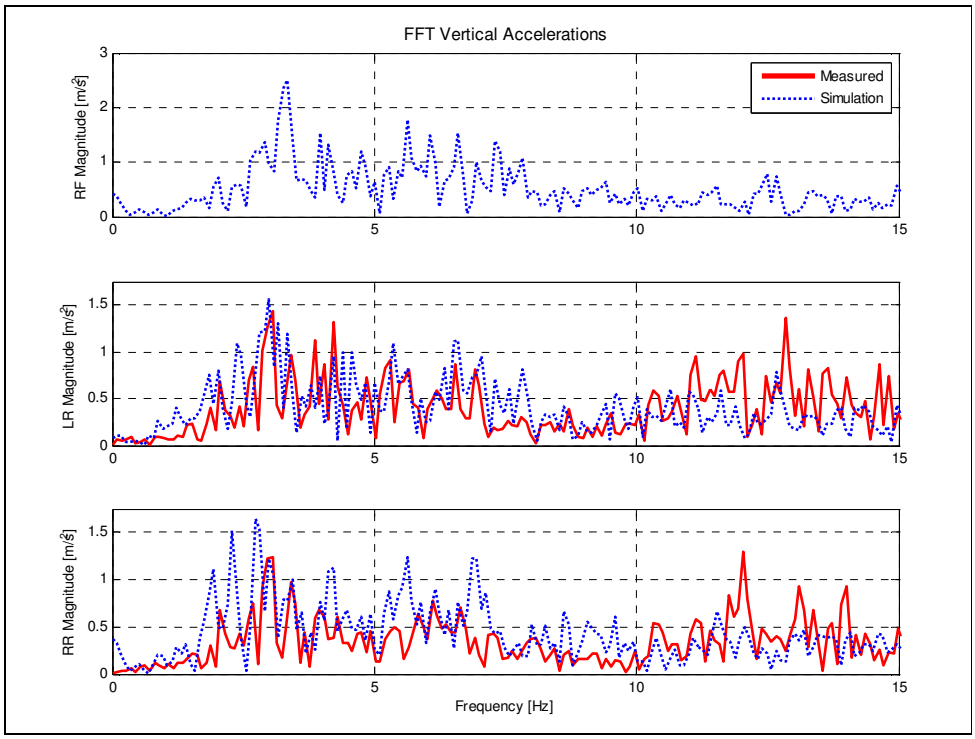

Figure 4-22: FFT of rear vertical accelerations (run 04)

The simulated pitch velocity data (Figure 4-23) shows good correlation with the measured data. The measured pitch natural frequency peak at 3 Hz is almost double that predicted by the simulation, but the frequency locations of the peaks correspond exactly. The roll velocity is over-predicted for frequencies below 3 Hz, but shows good correlation at higher frequencies.

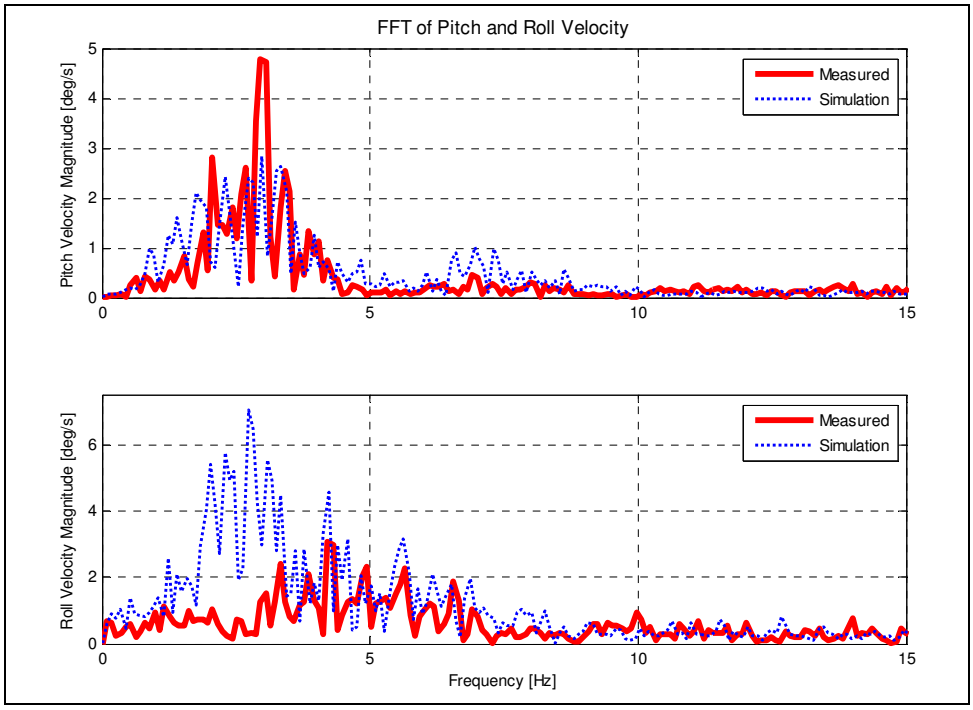

Figure 4-23: FFT of pitch- and roll-velocities (run 04)

The FFT of the measured and predicted stress, as well as the histogram of the stress ranges counted stress ranges is shown in Figure 4-24. The FFT and histogram of the simulated stresses are both true to the trends in the measured data.

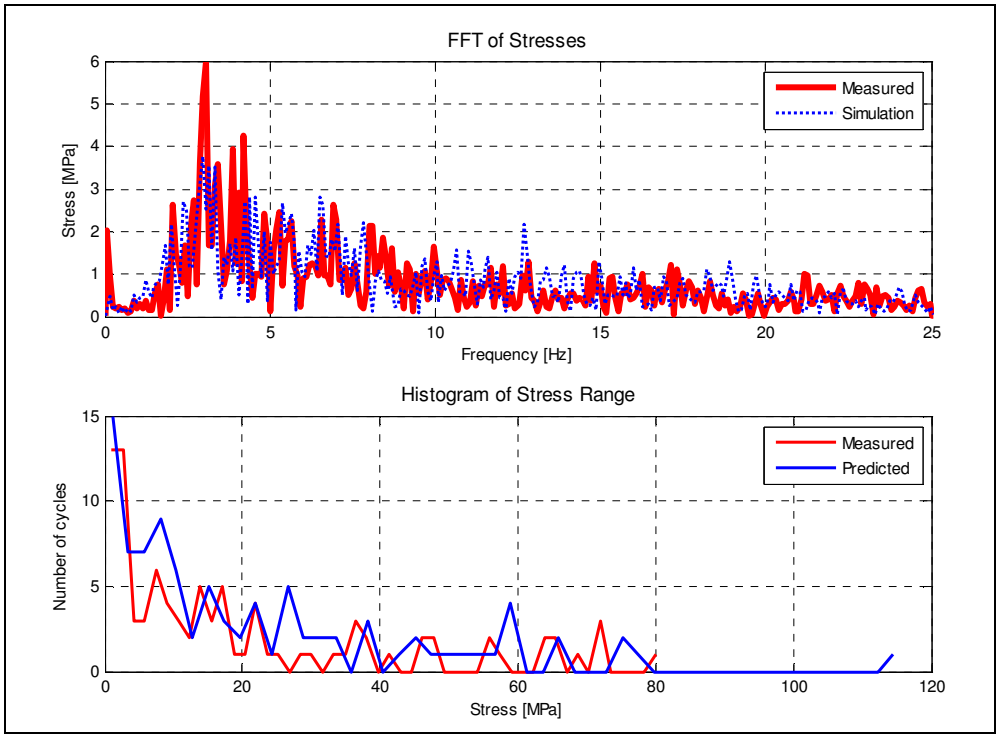

Figure 4-24: FFT and histogram of stress (run 04).

The data for all four runs considered is compared in terms of RMS values and fatigue damage in Table 4-1. The RMS and fatigue damage values allow for objective comparison. The RMS values are compared for rear vertical accelerations, pitch velocities and suspension forces. The mean values have been subtracted from the data before the RMS values were calculated (see Chapter 3 for a discussion on the fatigue damage calculation method).

Table 4-1: Comparison of measured and predicted RMS and fatigue damage values

Run number		Rear vertical acceleration RMS [m/s^2]	Rear suspension forces RMS [N]	Fatigue damage
run 01 ±14.5 km/h ride comfort	**Measured**	1.401	752	0.553e-6
	Predicted	1.993	1103	1.880e-6
	% Error	**42.27%**	**46.65%**	**239.13%**
run 02 ±54 km/h ride comfort	**Measured**	3.149	1715	2.408e-6
	Predicted	3.276	1687	2.280e-6
	% Error	**-4.04%**	**-1.61%**	**-5.14%**
run 03 ±14.5 km/h handling	**Measured**	3.273	2007	5.994e-6
	Predicted	5.061	2868	20.60e-6
	% Error	**-54.61%**	**42.91%**	**243.27%**
run 04 ±54 km/h handling	**Measured**	6.106	3776	12.73e-6
	Predicted	6.616	3676	16.10e-6
	% Error	**-8.37%**	**-2.64%**	**26.7%**

The correlation for the low speed runs is poor, with all quantities considered being greatly over predicted. The data for the higher speed runs shows better correlation in terms of both RMS and fatigue damage values. The difference in the measured and predicted damage prediction is still undesirably high.

The reader is however urged to consider the trends in the measured damage data. There is a factor five difference in measured damage between the ride comfort and handling settings at higher speed. At low speed, this difference is even more pronounced, with the difference in damage between the ride comfort and handling being an order of magnitude. The magnitude of the conceivable gains from the mathematical optimisation with the model thus eclipse the possible error being made.

The reader's attention is drawn to the fact that although the correlation for vertical accelerations and suspension forces is generally good, the correlation for damage is not of such high quality. This serves to confirm the statements made in the literature survey (Chapter 1) that a model should be validated for the parameters, which it is required to predict. It cannot be assumed that a model, which has only been validated for acceleration, will be able to predict fatigue damage.

4.3 Summary and conclusions of the model validation

The simulation results were compared with measured data for both discrete obstacles and random terrain. The comparison was done for all the tested speed and suspension setting configurations. The correlation over symmetric discrete obstacles is excellent for all quantities considered. The correlation over asymmetric discrete obstacles is tainted by the over-prediction of the roll velocity. The cause of the over-prediction is unclear, but the friction characteristic is considered to be the most likely cause. This suspicion is shared by other authors (Cronje, 2008). In a study on improving the handling of the Land Rover, Cronje found that friction in the suspension had a significant effect on the correlation of the lateral and roll dynamics of the vehicle. Cronje also found that the torsional flexibility of the vehicle chassis did not play a significant roll dynamics correlation.

The simulation results for runs over the Belgian paving showed very good correlation at high speed. The correlation at low speed is very poor. The friction and tyre model, uncertainty concerning the exact road profile traversed as well as a small error in the velocity measurement most likely carry the blame for the poor correlation at low speed. The poor correlation leads to the conclusion that the model cannot be used to predict the vehicle behaviour at low speed over the Belgian paving. The measured damage at low speed is, however much lower than that at high speed and high speed correlation is thus more important from a fatigue life perspective. The quality of the simulation results at high speed is considered adequate for the mathematical optimisation in this work. The model's correlation for high speed runs over the Belgian paving are of a similar quality as was obtained by Becker (2008) with the full ADAMS model of the Land Rover.

The model generally gives better results with the suspension set to ride comfort mode. This may be partly explained by recalling that the uncertainty in the hard damper characteristic is much greater than that of the soft characteristic. The greater

participation of the tyre dynamics, when the hard handling setting is activated, is also considered a contributing factor. The current tyre model cannot capture the complex hysteretic behaviour of a real tyre traversing very rough terrain.

Chapter 5

5. Mathematical optimisation

The work conducted in previous chapters has laid the foundation for the goal of the study. The aim is to investigate the required suspension characteristics to extend a vehicle's structural life. The implied mathematical optimisation problem in its most basic form could be posed as follows: Maximize the vehicle structural life with respect to the suspension spring and damper characteristics. Constraints may be added to this problem to govern the broader feasibility of the design. Mathematically the problem could be formulated as a standard optimisation problem of the form:

$$\min_{w.r.t.\ \bar{x}} f(\bar{x}), \qquad \bar{x} = [x_1, x_2, ..., x_n]^T \in \Re^n,$$

subject to the constraints:

$$g_j(\bar{x}) \leq 0, \qquad j = 1,2,...,m$$
$$h_j(\bar{x}) = 0, \qquad j = 1,2,...,r$$

The problem may be solved by any number of optimization algorithms. The choice of algorithm varies from direct search methods (particle swarm or evolutionary algorithms) to gradient based methods. Work done by Thoresson (2007), at the University of Pretoria, showed that gradient based methods are well suited to suspension optimisation problems. Thoresson used the Dynamic-Q algorithm developed by Snyman and Hay (2002) in his optimisation for suspension ride comfort and handling characteristics. The Dynamic-Q algorithm makes use of successive spherical quadratic approximations to the cost function, which are then solved by the LFOPC dynamic trajectory method. The method is especially well suited to engineering problems, where the cost function evaluations take the form of simulations, which may have high inherent numerical noise content and are inherently expensive.

The similarities between the work done by Thoresson (2007) and the current study as well as the favourable results obtained by Thoresson lead to use of Dynamic-Q for the mathematical optimisation in this work. The current study attempted to build on the foundation laid by Thoresson and others in the gradient based mathematical optimisation of vehicle suspension spring and damper characteristics at the University of Pretoria.

5.1 Objective function definition

The challenge in defining a suitable metric to represent damage in the vehicle structure due to road loads is that changes in load levels are not a direct measure of changes in fatigue damage to the structure (Zeiler and Barkey, 2001). This is due to the fact that fatigue is not only dependent on load levels, but also on the number of load cycles applied. A simple parameter such as standard deviation of suspension forces will therefore not provide unambiguous insight into the influence of suspension characteristics on fatigue damage in the vehicles structure. The discussion of input loads to the vehicle structure naturally leads to the suggestion of cumulative fatigue damage as suitable metric (see Chapter 3 for a full discussion). The objective function is therefore defined as the cumulative damage incurred by the vehicle when driving over a road consisting of 100 m of ISO class A road and 100 m of the Gerotek Belgian paving at a constant speed of 54 km/h. The design variables are the spring and damper characteristics of the vehicle as is defined in paragraph 5.2.

5.2 Definition of design variables

The force displacement characteristic of the hydro-pneumatic spring (for constant static gas pressure) is uniquely determined by the static gas volume in the system. The reader is referred to Appendix B for a complete discussion. The static gas volume in the $4S_4$ strut may theoretically be varied between 0.85 ℓ and 0 ℓ. The volume of 0 ℓ would lead to a suspension stiffness which is solely due to the compressibility of the hydraulic oil. The likeliness of an optimum existing at extreme highs of suspension stiffness is considered negligible. Additionally the computational expense of including the oil stiffness in the results is prohibitive and the lowest static gas volume which will be considered is thus 0.05 ℓ. The larger volume of 0.85 ℓ is the maximum gas volume, which the $4S_4$ accumulators can currently accommodate. The static gas volume, sgv, is thus defined as the first design variable. The baseline static gas volume was chosen to be 0.3 ℓ. The stiffness of the hydro-pneumatic spring about the static position for a static gas volume of 0.3 ℓ is the same as the linear baseline Land Rover spring stiffness. The range of possible hydro-pneumatic spring characteristics for the unladen case is shown in Figure 5-1.

The damper characteristic is not so easily defined by a single variable. The damper characteristic is a complex curve in both the bump and rebound directions, which cannot be captured by a single variable as was done with the spring characteristic. Thoresson (2007) followed an approach were the standard Land Rover damper characteristic was modelled using piecewise quadratic fits. This approach has a solid physical foundation as the general description of the forces generated by fluid flowing through an orifice is of the form:

$$F = kv^2$$

The force, F, is thus a quadratic function of the velocity of the flow, v. The factor k is a correlation factor, which may be dependent on the size of the orifice, the entry

and exit conditions and the surface roughness. In different applications it may either be described as a "drag coefficient" or a "loss factor".

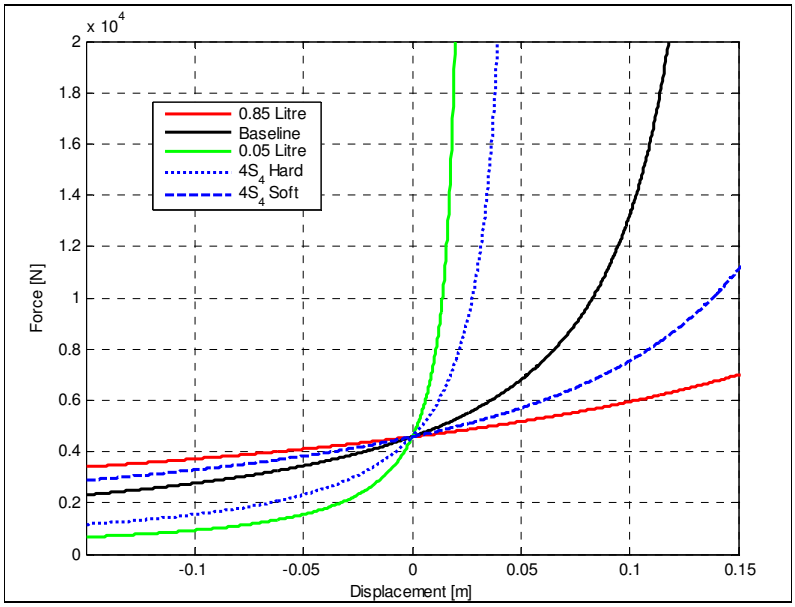

Figure 5-1: Range of possible spring characteristics (unladen).

Thoresson (2007) showed the effectiveness of the technique and the approach will be repeated here. The damper characteristic is defined using the 4 piecewise quadratic fits:

$$fit_1 = dsf_1\left(-1084.5v^2 + 1378.2v\right) - dsf_2\left(878.2\right)$$
$$fit_2 = dsf_3\left(-2220.2v^2 + 4076.3v\right)$$
$$fit_3 = dsf_4\left(7316.8v^2 + 6483.7v\right)$$
$$fit_4 = dsf_5\left(324.88v^2 + 2522.7v\right) + dsf_6\left(2067.9\right)$$

The damper force is then determined from the piecewise fits by the logical argument:

$$\begin{aligned} &if \quad v \le 0 \\ &F_{damper} = \max\left[fit_1(v), fit_2(v)\right] \\ &else \\ &F_{damper} = \max\left[fit_3(v), fit_4(v)\right] \\ &end \end{aligned}$$

An infinite number of damper characteristics can be achieved by varying the six damper scale factors $dsf_1 \rightarrow dsf_6$. These scale factors are defined as design variables 2 through to 7. The damper characteristics obtained in this manner all have a single blow-off in both the bump and rebound direction. Each of the scale factors is allowed to vary between 0.1 and 4. The range of possible damper characteristics is plotted in Figure 5-2. The baseline characteristic (all scale factors set to unity), as well as the two measured $4S_4$ damper characteristics have been included in the plot.

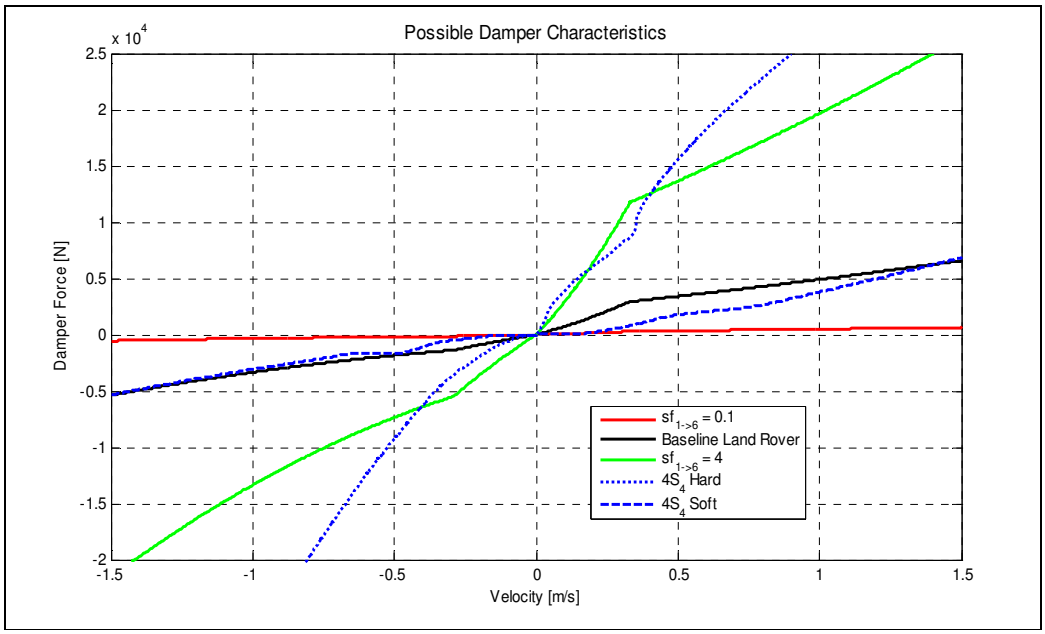

Figure 5-2: Range of possible damper characteristics

5.3 Constraint functions

The performance of a vehicle suspension system cannot be judged on only one criterium as has been done thus far. Some constraints are necessary in the optimisation to ensure convergence to realistic and feasible optima. Vehicle suspensions have to perform adequately in criteria as diverse as ride comfort, suspension working space, vehicle traction and damage to road surfaces. Adequate performance in these criteria may be introduced into the design process in the form of constraint functions to the optimisation.

The first constraint which will be considered relates to the suspension working space. This available working space for the suspension is fixed by the vehicle design and suspension geometry. The suspension layout in the Land Rover allows for 250 mm of total suspension displacement in the front and 300 mm total suspension displacement in the rear. The working space is constrained by the vehicle bump stops. The bump stops represent hard stops for the suspension and it is thus undesirable for the vehicle suspension to encounter the bump stops during normal operation. This is especially true from a ride comfort and fatigue damage perspective. The vehicle spring and damper characteristics must therefore be such that enough of the road input is absorbed in the available working space that the bump stops are not encountered during normal operation. A bump stop force of 100 N was chosen as the limit for the bump stop forces.

The working space constraint function is thus defined as follows:

$$g_1 = \frac{mean\left[t\left(F_{bump\_lf}\right) > 100, t\left(F_{bump\_rf}\right) > 100, t\left(F_{bump\_lr}\right) > 100, t\left(F_{bump\_rr}\right) > 100\right]}{t_{total}} \leq 0,$$

where,

$F_{bump\_lf,rf,lr,rr}$ = bump stop forces,

t_{total} = the total simulation time.

This gives the fraction of total simulation time for which bump stop contact has exceeded 100 N.

Wheel hop has often been employed as a measure of vehicle handling and safety in sensitivity studies (Holdmann and Holle 1999; Ivers and Miller, 1989; cited in Uys et. al., 2004). The philosophy behind this statement, is that when the wheel is not in contact with the road surface the ability of the tyre to generate lateral or longitudinal force is lost.

The sensitivity of wheel hop to suspension damping, observed by Uys et. al. (2007), prompted Thoresson (2007) to include a wheel hop constraint in his optimisation for ride comfort. The constraint was triggered when the tyre lost ground contact for more than 10% of the total simulation time. The constraint was implemented in the current study as:

$$g_2 = \frac{mean\left[t\left(F_{tyre\_lf}\right) \leq 0, t\left(F_{tyre\_rf}\right) \leq 0, t\left(F_{tyre\_lr}\right) \leq 0, t\left(F_{tyre\_rr}\right) \leq 0\right]}{t_{total}} - 0.1 \leq 0,$$

where,

$F_{tyre\_lf,rf,lr,rr}$ = magnitude of the tyre contact forces.

Cole (2001) mentions the importance of considering infrastructure damage due to heavy vehicle suspension systems. Designing road friendly suspension systems will in the long run reduce vehicle structural damage, thanks to reduced road input loads which are caused by deteriorating road conditions. According to Cole, the primary failure modes of continuous flexible pavements are rutting and fatigue cracking. Gross vehicle weight is considered to be the dominant factor in rutting and is thus largely out of control of the suspension designer. Fatigue cracking, however, is caused by the dynamic components of the vehicle tyre forces.

The dynamic tyre forces are influenced by suspension characteristics, vehicle speed and road surface roughness (Cole and Cebon, 1996; cited by Cole, 2001). Cebon (1989) noted in his review of the literature, that there was agreement that soft suspension and tyre characteristics were required to minimise dynamic tyre loads. There does, however appear to be an optimal level of damping which minimises road damage.

A parameter frequently used, in the vehicle-terrain interaction environment, to characterise the relative magnitude of the tyre dynamic forces is the Dynamic Load Coefficient (DLC). The literature does not seem entirely united over the definition of the DLC. Steyn (2001) quotes the DLC as the "Coefficient of Variance" of the applied load, while Cebon (1989) defines the DLC as the ratio of the RMS dynamic tyre force to the static tyre force. These quantities although similar are not equivalent. In this study Cebon's definition will be used.

$$\therefore DLC = \frac{RMS\left(F_{tyre}\right)}{F_{stat}}$$

The road damage due to dynamic loads is usually modelled as functions of the coefficient of variation of dynamic tyre force or the DLC. Eisenmann's Road Stress Factor, which is a fourth power function of the DLC, is very popular (Cebon, 1999). Cebon states that the DLC in heavy vehicles under normal operating conditions ranges between 0.1 and 0.3, depending on the suspension type employed (Sweatman, 1983; cited by Cebon, 1989). DLC values of up to 0.4 were measured by Sweatman on what Cebon calls "particularly poor" suspension design. Cebon notes that the DLC is very sensitive to road roughness and vehicle speed.

The work done by Sweatman (Sweatman, 1983) was studied thoroughly in an attempt to apply his findings concerning DLC in heavy vehicles as a constraint in this work. Comparison with Sweatman's work proved to be very difficult. Although Sweatman carefully noted the speed and roughness conditions for every test, the road profile roughness was recorded in NAASRA roughness. The National Association of Australian State Road Authorities (NAASRA) roughness indicator is a single parameter road roughness indicator. It is analogous to the International Roughness Index (IRI), (Sayers and Karamihas, 1998).

The single parameter IRI and NAASRA roughness indicators are measures of the relative motion between the sprung and unsprung masses of the vehicle travelling over a specific terrain. The indicators inevitably contain no direct frequency information. In the case of the IRI, the vehicle used for the roughness computation is standardised in the "golden car" model. The addition of the adjective, golden, in this case refers to the standardisation. The standardisation is done to remove effects from variation in vehicle vertical response. A specification for the vehicle used in calculation or measurement of the NAASRA roughness indicator seems to be absent.

Additional criticism of the NAASRA roughness indicator can be expressed on the units of measurement. IRI is measured in meters of relative vertical motion, per kilometer of longitudinal vehicle travel. The NAASRA roughness is measured in cycles of relative vertical motion per kilometer of longitudinal vehicle travel. Any relative motion between the vehicle sprung and unsprung masses, in one direction, which is larger than 15.2 mm is counted as a cycle (Hunt and Bunker, 2003). The user thus runs the risk of measuring a very similar roughness index for vastly different road profiles or vice versa.

The lack of comparable road profile data made it very difficult to define a suitable road damage constraint. The inclusion of the road damage constraint also caused a breakdown of the search algorithm's convergence. These problems eventually lead to

the exclusion of the road damage constraint from the optimisation. Several trial optimisation runs were however conducted using road damage (DLC) as objective function. The road damage optimisation converged to a similar region of the design space as the vehicle damage optimisation. The requirements for road damage optimal and vehicle damage optimal suspension can therefore not be entirely dissimilar. The relationship between vehicle damage optimal and road damage optimal suspension characteristics should be investigated in future work.

5.4 Dynamic-Q parameter selection

The behaviour and performance of the Dynamic-Q algorithm is governed by several parameters. These include the dynamic move limit, dml; the convergence tolerance on step size, x_{tol}; the convergence tolerance on the function value, f_{tol}; the maximum number of function evaluations, k_{loop}; and the gradient estimation method and gradient perturbation step size. The scaling of the cost function and design variables are also important. This section briefly summarises the selection of these parameters.

5.4.1 Cost function and variable scaling

The scaling of the cost function for use with the Dynamic-Q algorithm should be such that the design space is as "spherical" as possible. This ensures that the minima of the quadratic approximations of the cost function will correspond closely to the minima of the cost function. The cost function was scaled such that the baseline function value is equal to unity. This is achieved by multiplying the calculated damage by 5×10^5.

The design variables were scaled such that their ranges varied between 0 and 1. The scaled design variables will be referred to as optimisation variables (ov). The scaling of the design variables was accomplished as presented below.

$$ov_1 = \frac{(sgv - 0.05)}{0.8}$$

$$ov_{2\to7} = \frac{(dsf_{1\to6} - 0.1)}{3.9}$$

5.4.2 Gradient estimation and gradient sensitivities

The Dynamic-Q algorithm, which is to be used for the mathematical optimisation in this study, is a gradient based method. As such, the accurate estimation of the gradient of the cost function is imperative to the method's efficiency. There are essentially two methods which are used for gradient estimation in cases where the cost function evaluation is by simulation. Forward finite difference approximation is often employed and has a computational cost of $n+1$ function evaluations, where n represents the number of variables on which the cost function depends.

The forward finite difference approximation to the gradient is given by:

$$\frac{\partial f(\bar{x})}{\partial x_i} \cong \frac{\Delta f(\bar{x})}{\delta_i} = \frac{f(\bar{x} + \bar{\delta}_i) - f(\bar{x})}{\delta_i},$$

where $\bar{x}$ is the vector of design variables and the perturbation, $\bar{\delta} = [0 \ldots \delta_i \ldots 0]^T$.

Central finite difference approximation is often employed for gradient estimation, where the cost function contains severe numerical noise. The extra robustness does however come with increased computational cost. Central finite difference approximation has a computational cost of $2n$ function evaluations. The effectiveness of this technique was clearly illustrated by Thoresson (2007). The central finite difference approximation to the gradient is given by:

$$\frac{\partial f(\bar{x})}{\partial x_i} \cong \frac{\Delta f(\bar{x})}{\delta_i} = \frac{f\left(\bar{x} + \bar{\delta}_i/2\right) - f\left(\bar{x} - \bar{\delta}_i/2\right)}{\delta_i}$$

The size of the perturbation for the gradient estimation, δ_i, is critical to the quality and robustness of the gradient estimation. Choosing the δ_i too large may cause inaccurate estimation. This does however not mean that δ_i can be chosen to be arbitrarily small. At small values of δ_i, rounding or other numerical effects may cause the gradient estimation to become unreliable. Snyman (2004) suggests that gradient sensitivity study should be conducted on the size of δ_i in order to determine a suitable size. Conducting the sensitivity study with both forward and central finite difference approximation, would also aid in deciding between the two gradient approximation methods. Such a sensitivity study, although computationally expensive and time consuming, will pay dividends in terms of quick and robust convergence in the eventual mathematical optimisation.

The gradients were evaluated at the baseline spring and damper characteristics in order to gauge their sensitivity. The size of δ_i was varied logarithmically between 10^{-15} and 10^{-1}. The gradient was calculated at 45 points between the boundaries specified. This was conducted for both the forward and central finite difference gradient approximation techniques.

The sensitivity study showed that the gradient estimation techniques were both stable from 10^{-13} up to 10^{-4}. Beyond these limits the gradient estimates seem to become unstable. A large range of stable gradient approximations bodes well for the optimisation. The results were verified for other points in the design space.

The results of the gradient sensitivity analysis initially lead to the choice of forward finite difference gradient approximation for the gradient estimation in the Dynamic-Q algorithm. Reliable convergence was achieved with unconstrained problems, but the convergence deteriorated significantly with the introduction of constraint functions. The use of central finite difference approximation improved the situation

significantly. The perturbation step size for the gradient estimation, δ, was taken as 10^{-6}. Central finite differences will be employed for the optimisation in the study.

5.4.3 The dynamic move limit

The dynamic move limit in the dynamic Q algorithm limits the maximum step which can be taken by the algorithm. It is recommended that the dynamic move limit should be selected such that it is in the same order of magnitude as the radius of the region of interest. Let the radius of interest then be determined by taking square root of the sum of the squares of half the variable ranges of interest as shown below:

$$R = \sqrt{\left(\frac{x_{1R}}{2}\right)^2 + \left(\frac{x_{2R}}{2}\right)^2 + \ldots + \left(\frac{x_{nR}}{2}\right)^2} = \frac{1}{2}\sqrt{x_{1R}^2 + x_{2R}^2 + \ldots + x_{nR}^2}$$

The relation given can be simplified further, by assuming that the variables have been correctly scaled in order to have a range between 0 and 1. The radius of interest is then simply $R = \sqrt{n}/2$. The dynamic move limit, dml, was thus initially taken as half of the square root of the number of variables being considered. A dynamic move limit of between one third and one fifth the radius of interest was however found to given more stable convergence.

5.4.4 Convergence tolerances

There are no guidelines given for the selection of the convergence tolerances. The relative variable and function convergence tolerances, x_{tol} and f_{tol}, were thus chosen as 10^{-6}. These represent minute changes, considering the possible variation in both design variable and cost function values. The maximum function evaluation limit, k_{loop}, which simply avoids excessive iteration in the case of failed convergence, was selected to be 100 function evaluations. This represents two days of simulation time in the case of the seven design variable problem. The Dynamic-Q parameters are summarised in Table 5-1.

Table 5-1: Dynamic-Q parameter settings for the optimisation.

Parameter	Value
dml	$\sqrt{n}/10 \rightarrow \sqrt{n}/6$
x_{tol}	10^{-6}
f_{tol}	10^{-6}
k_{loop}	100
Gradient estimation technique	Central finite differences
δ	10^{-6}

5.5 The two-variable problem

The optimisation problem, considered in its most basic form, is a function of two variables; the spring characteristic and the damper characteristic. The hydro-pneumatic spring static gas volume was kept as variable, but the six damper scale factors were consolidated into a single damper scale factor. Considering this simplest form of the problem allows visualisation of the cost function as a surface in three dimensional space. This attribute is lost when more than two variables are considered. The two-variable formulation thus allows one to develop an intuition for the problem, which may extend into the multi-variable problem. The optimisation methodology was consequently developed on the two-variable problem before attempting multi-variable problems.

5.5.1 Cost function visualisations

The cost functions for the unladen, fully laden, heavy load and extreme load cases were each evaluated at 100 equi-spaced points across the design space. Surface plots of each of the cost functions are shown in Figure 5-3. The plots do not give detailed information about the presence of noise or local minima in the cost functions, but the general trend of the cost functions are revealed. The cost functions in all four cases indicate a largely uni-modal design space. The minima can be found in the low damping, soft spring region of the design space. The cost functions exhibit a large region of near minimum values. This may be indicative of the existence of highly robust optima.

The cost functions seem to be more sensitive to the damper characteristics than the static gas volume, which determines the pneumatic spring characteristic. The sensitivity to the static gas volume increases as the vehicle payload is increased. The stiffness of the pneumatic spring is a hyperbolic function of the static gas volume, as is shown in Appendix B. This is at least the partial cause of the cost functions' low sensitivity to static gas volume at high static gas volumes. The load-levelling effect of the hydro-pneumatic suspension is thought to be the main cause. The reader should take note that the constraints are in no way taken into account in the visualisations of the design space. One should therefore be careful not to make inferences about the locations or existence of feasible minima.

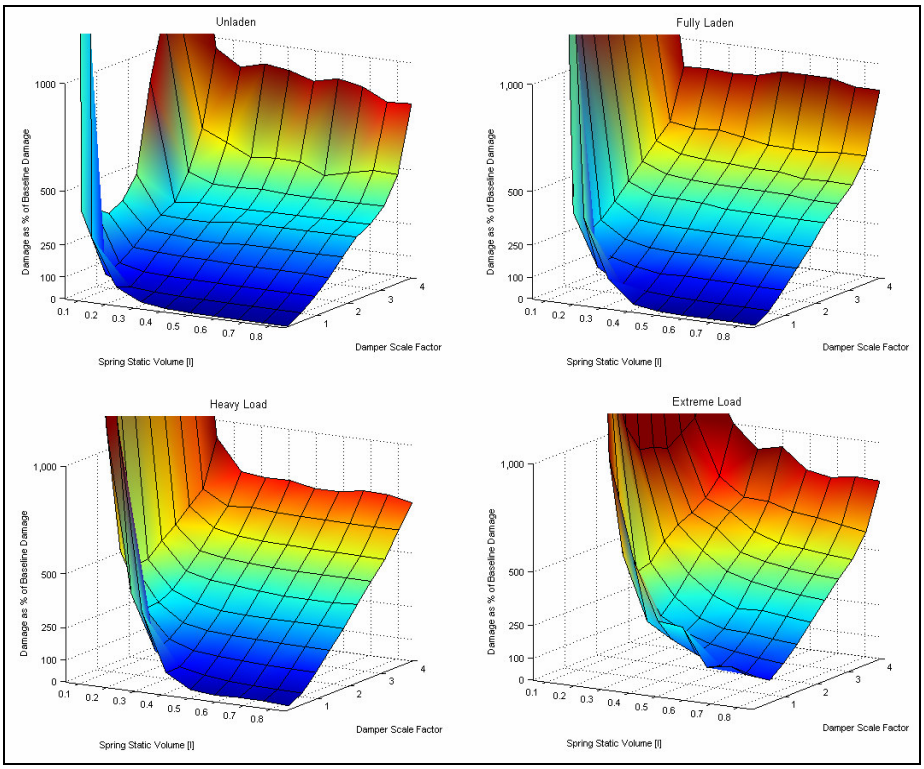

Figure 5-3: Cost function plots for the four load cases.

5.5.2 Two-variable optimisation

The Dynamic-Q algorithm was applied to find the feasible minima of each of the cost functions. The optimisation was performed from several different starting points distributed across the design space to ensure that the converged optimum corresponds to the global optimum. The optimisation was performed from each starting point with the move limit set to both $\sqrt{n}/6$ and $\sqrt{n}/10$. The starting points used are listed in Table 5-2.

Table 5-2: Two-variable optimisation starting points.

Unladen		**Fully Laden**		**Heavy Load**		**Extreme Load**	
Static Gas Vol. [ℓ]	Damper Scale Factor	Static Gas Vol. [ℓ]	Damper Scale Factor	Static Gas Vol. [ℓ]	Damper Scale Factor	Static Gas Vol. [ℓ]	Damper Scale Factor
0.3	1	0.3	1	0.3	1	0.3	1
-	-	0.3	1.5	0.3	2.5	0.3	4
0.2	0.9	0.2	0.9	0.2	0.9	0.2	0.9
0.2	3.2	0.2	3.2	0.2	3.2	0.2	3.2
0.7	0.9	0.7	0.9	0.7	0.9	0.7	0.9
0.7	3.2	0.7	3.2	0.7	3.2	0.7	3.2
0.45	2	0.45	2	0.45	2	0.45	2

The optimisation converged fairly rapidly from most starting points. The convergence histories for selected results have been superimposed on the cost function plots in Figure 5-4. The convergence rate from all runs was reliable for most of the design space. This deteriorated significantly when the algorithm entered the flat dish-like region of near optimum values. The optimisation did not converge to distinct minima, but settled instead in a wide region of near optimal values.

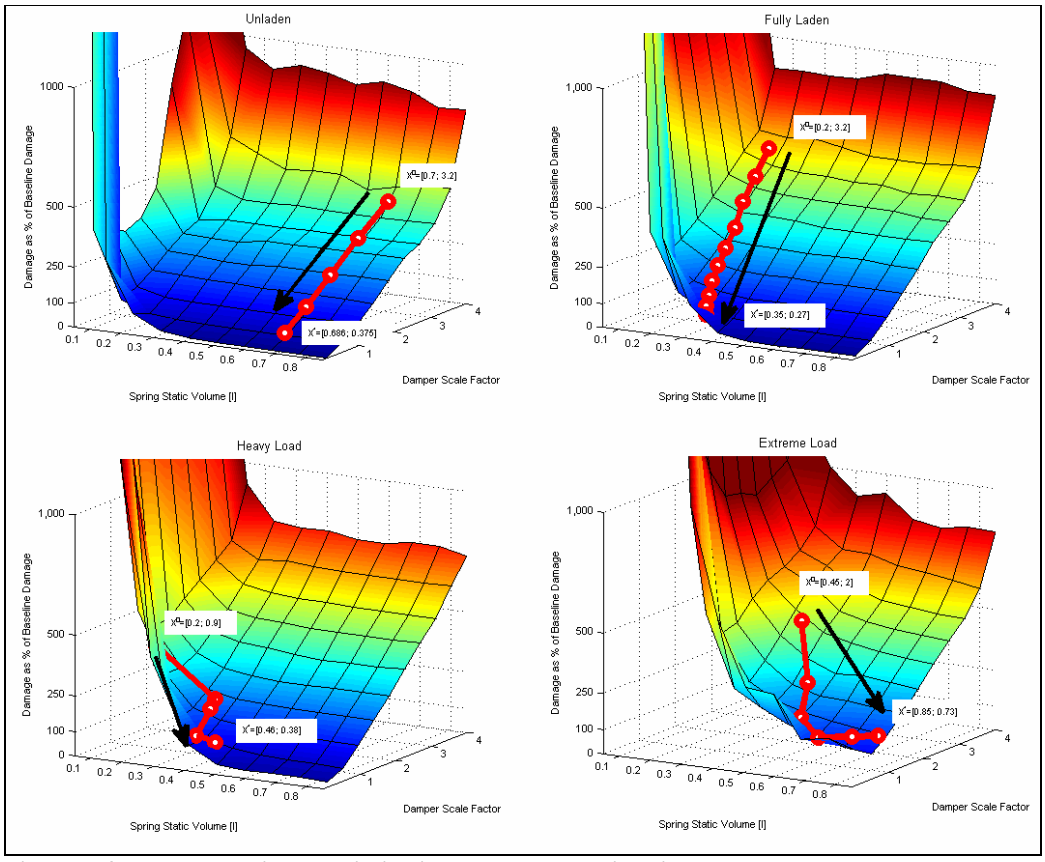

Figure 5-4: The two-variable optimisation convergence histories.

Attempts to improve the convergence in this region by adopting different scaling of the cost function, and refining the convergence tolerances proved ineffective. The converged points were also used as starting points for optimisation with small move limits and tight convergence criteria, but this also failed to deliver a distinct optimum. The poor convergence near the optimum is attributed to the presence of high numerical noise in the both the cost and constraint functions as well as low gradients. As the magnitude of the gradient decreases near the optimum the numerical noise is thought to dominate the gradient approximation, which leads to the deterioration of convergence. The results of the Dynamic-Q optimisation are presented in Table 5-3. The optima are quoted in terms of ranges of the best five converged solutions. Statistical data for the converged points has been included.

Table 5-3: Two-variable optimisation results.

	Unladen			**Fully Laden**		
	Static Gas Vol. [ℓ]	Damper Scale Factor	Cost function value [% of baseline]	Static Gas Vol. [ℓ]	Damper Scale Factor	Cost function value [% of baseline]
Range	0.3-0.7	0.2-0.39	24-28%	0.34-0.5	0.21-0.52	50-65%
Mean	0.49	0.34	25.5%	0.39	0.38	57%
Std. Dev.	0.16	0.11	1.0%	0.09	0.21	6.7%
Optimum	0.69	0.38	24%	0.35	0.27	50%
	Heavy Load			**Extreme Load**		
	Static Gas Vol. [ℓ]	Damper Scale Factor	Cost function value [% of baseline]	Static Gas Vol. [ℓ]	Damper Scale Factor	Cost function value [% of baseline]
Range	0.4-0.85	0.33-0.66	65-76%	0.6-0.85	0.55-0.73	125-148%
Mean	0.53	0.45	72%	0.69	0.6	141%
Std. Dev.	0.2	0.17	4.1%	0.10	0.09	8.2%
Optimum	0.46	0.38	65%	0.85	0.73	126%

The widely distributed range of near optimum points may be due to the inadequate convergence of the optimisation or it may be indicative of very robust optima. In order to clarify the situation and confirm the optima, a Monte Carlo type simulation was performed for each cost function in the region to which convergence occurred. To this purpose, the static gas volume and the damper scale factor were modelled as normally distributed, independent random variables. The mean and standard deviations of the variables were taken as listed in Table 5-3. The cost function was evaluated at 500 random points for each simulation. The results of the Monte Carlo simulation are shown on contour plots of the cost functions in Figure 5-5 to Figure 5-8. The optima determined from the Monte Carlo simulation are also given in the plots.

The optima determined in the optimisation with Dynamic-Q were confirmed by the Monte Carlo simulations. The contour lines in the figures again confirm the insensitivity of the cost functions with respect to the static gas volume. The sensitivity increases as the unsprung mass is increased. The optimum damper scale factors are in all four cases constrained by the bump stop constraint. The boundary of this constraint can clearly be distinguished between the feasible and infeasible cost function evaluations. The constraint boundary appears to be fairly smooth and limited to the low damping regions of the design space.

The fatigue damage optima, for the unladen case, are compared to the ride comfort and handling optima determined by Thoresson (2003) in Table 5-4. The spring and damper characteristics required for fatigue damage and ride comfort are very similar. The $4S_4$ system may conceivably be configured to fatigue damage and handling optima, instead of the ride comfort and handling optima currently implemented.

Table 5-4: Fatigue damage, ride comfort and handling optima.

Unladen vehicle		
	Static Gas Volume [ℓ]	Damper Scale Factor
Fatigue damage	0.3-0.7	0.2-0.39
Ride comfort	0.5	0.3
Handling	0.1	3

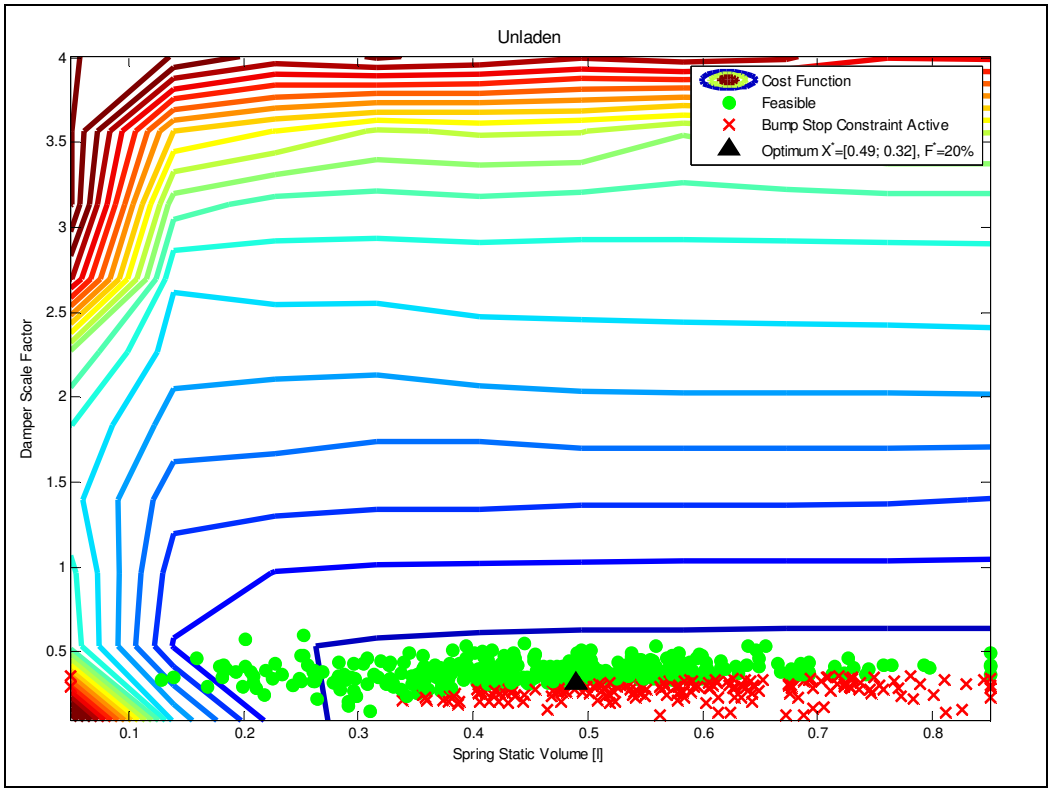

Figure 5-5: Contour plot of the unladen two-variable cost function.

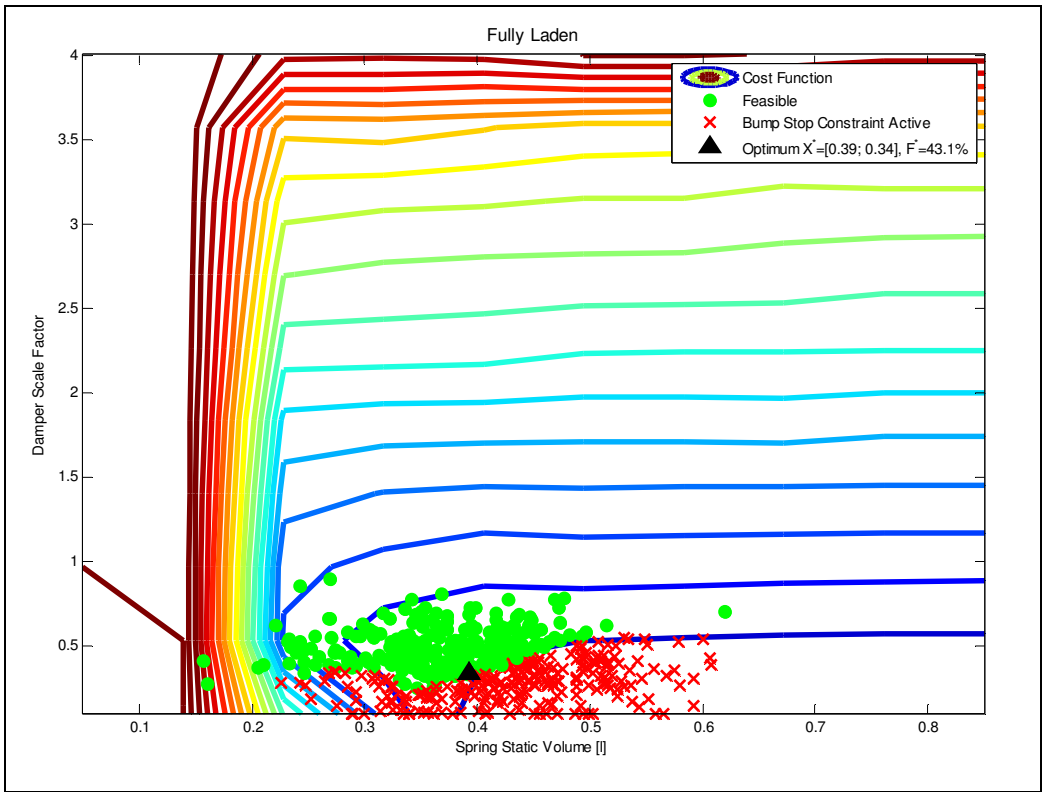

Figure 5-6: Contour plot of the fully laden two-variable cost function.

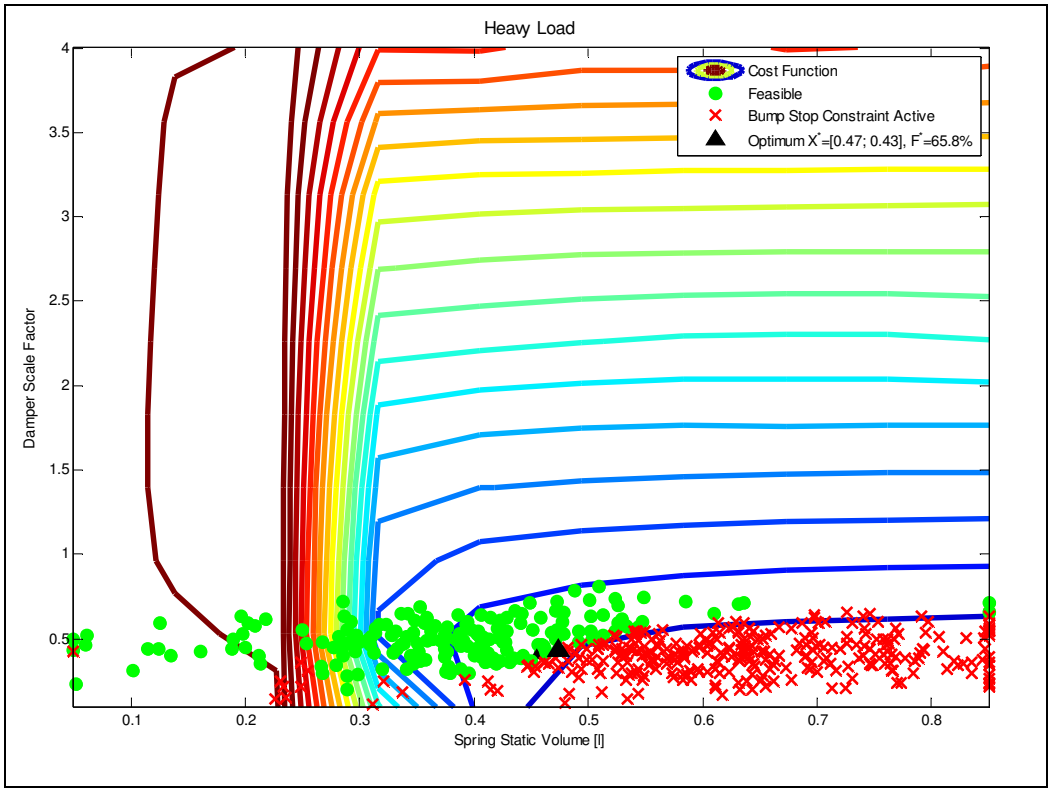

Figure 5-7: Contour plot of the heavy laden two-variable cost function.

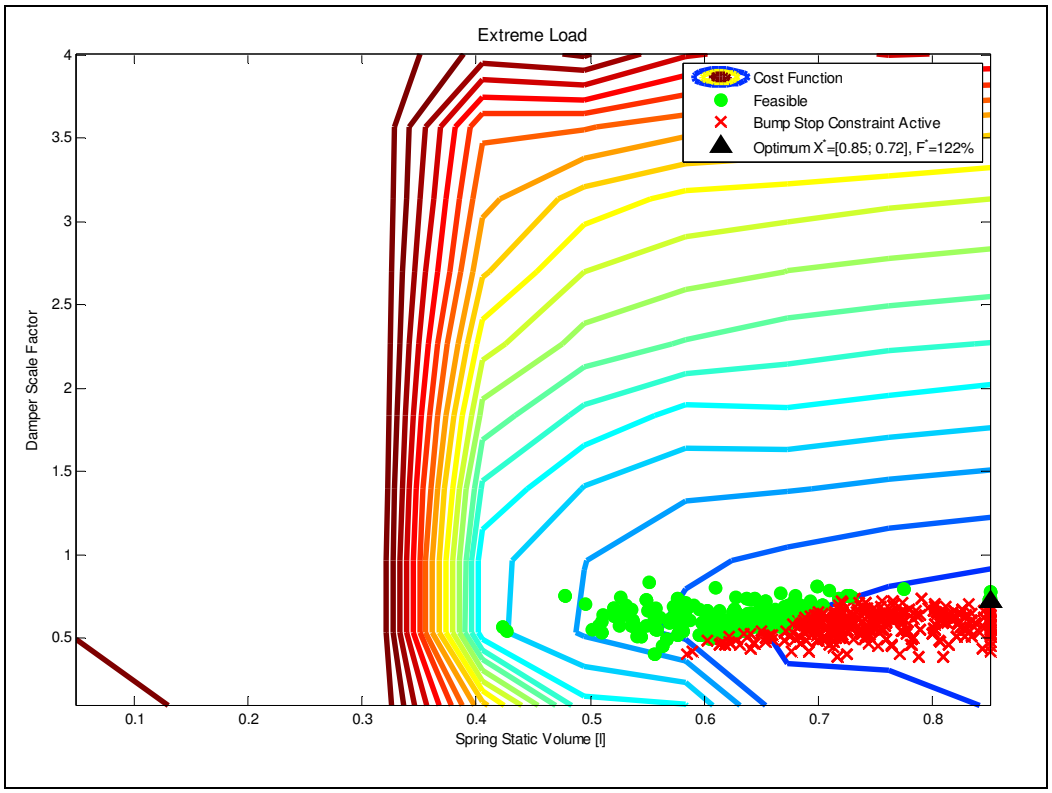

Figure 5-8: Contour plot of the extreme load two-variable cost function.

5.5.3 Determining robust optima

The shallow dish-like shape of the objective function in the optimal region makes an optimum in this region highly robust in terms of cost function value. The optima determined thus far are in all cases found on the constraint boundary. Small variations in the design variables in the implementation of the optimum solution could thus cause violations of the constraints. The optima are thus not robust in terms of possible constraint violations. A method is required for selecting optima deep enough in the feasible design space, so as to avoid constraint violations in case of small variations on the design variables.

A conservative choice for the damper scale factor could be made by observing that the bump stop constraint seems to be bounded by a critical damping level for the design space. The insensitivity of the static gas volume then allows a choice of a wide range of static gas volumes without large variations in the cost function value. The optimum variable ranges chosen by this method for each cost function are listed in Table 5-5. It should be noticed that the increase in robustness comes at a cost in objective function values.

Table 5-5: Conservative robust optima

	Static Gas Vol. [ℓ]	Damper Scale Factor	Cost function value [% of baseline]
Unladen	0.3-0.85	0.4	26-35%
Fully Laden	0.35-0.85	0.6	60-80%
Heavy Load	0.5-0.85	0.75	87-107%
Extreme Load	0.7-0.85	0.85	143-167%

Beyer and Sendhoff (2007) provide a thorough, though somewhat fragmented review of the literature on robust optimisation. Although much work has been done in the field, the developments of methods which can be applied to realistic engineering problems relying on expensive simulations, seems to be in its infancy. The deterministic methods which have been developed essentially all rely on constructing a robust version of the cost function and its constraint functions. The robust version of the cost function may be found either by regularisation of the cost and constraint functions or similar penalty based methods.

These methods may require knowledge of the mean and variance of the cost and constraint functions across the design space. The determination of the mean and variance of an unknown, simulation based cost or constraint function is however problematic. The determination of the mean and variance using a design of experiments approach i.e. probabilistic methods is prohibitive.

Lee and Park (2000) suggest a method for dealing with constraints in finding a robust optimum. Their method involves defining new, more conservative, constraint functions for the problem which are of the form:

$$g_{new\ \ j} = g_j + k_j \sum_{i=1}^{n} \left| \frac{\partial g_j}{\partial x_i} \right| \Delta x_i .$$

The new constraints are thus composed of the old constraints and a penalty parameter determined by the gradient of the old constraint and the expected maximum variation of the design variables. The optimisation is then performed with the new constraints. The approach does however require that the constraint functions and their derivatives are defined well into the feasible region of the design space.

The bump stop constraint employed here is formulated such that neither the constraint function values nor constraint function derivatives are continuous into the feasible region. It is also undesirable to redefine the constraints as it would require rerunning all of the optimisation which has thus far been conducted. A different method for dealing with constraint violations due to design variable variations is thus desired. The method implemented should also be such that it can be extended to the solution of the multivariable problem.

The proposition is made that a suitable robust optimum may be determined by taking a step from the deterministic optimum in the direction of the feasible design space. The direction of the step is determined by examining the gradient of the constraint functions. As mentioned, the gradient is not available at the feasible optimum. The Monte Carlo simulation results do however provide several points which are very near the optimum, but in the infeasible region of the design space. The constraint function gradient information at the optimum can thus be approximated by evaluating the gradient at the infeasible point, closest to the optimum. The direction of the step into the feasible design space, is taken as the normalised steepest descent direction of the constraint function gradients, $\bar{u} = \dfrac{\bar{\nabla} g(\bar{x})}{\left\| \bar{\nabla} g(\bar{x}) \right\|}$.

The size of the step is determined by the expected maximum variation on the variables as a percentage of their magnitude, Δx. The magnitude of the step is scaled by:

$$\frac{\Delta x}{1-\Delta x}\bar{x}^{*}.$$

The position of the robust optimum is thus defined as:

$$\bar{x}^{*}_{robust} = \bar{x}^{*} + \left(\frac{\Delta x}{1-\Delta x}\right) \times \left[\bar{x}^{*} \cdot \bar{u}\right]$$

The maximum variable variation was selected to be 10% of the mean value, i.e. $\Delta x = 0.1$. The variables are assumed to have a standard normal distribution. The probability of the variables varying by more than 10% of the mean value is less than 5%, when a standard deviation of 5% is assumed (Johnson, 2005).

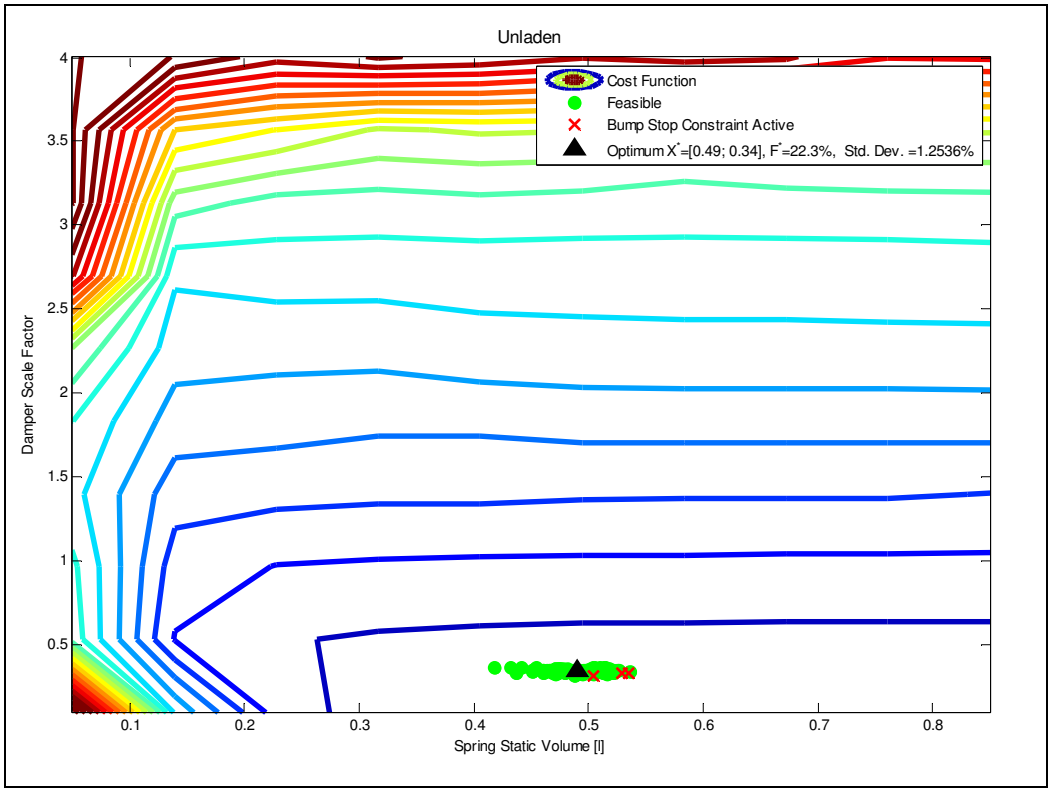

Figure 5-9: Contour plot of the unladen cost function including the robust optimum.

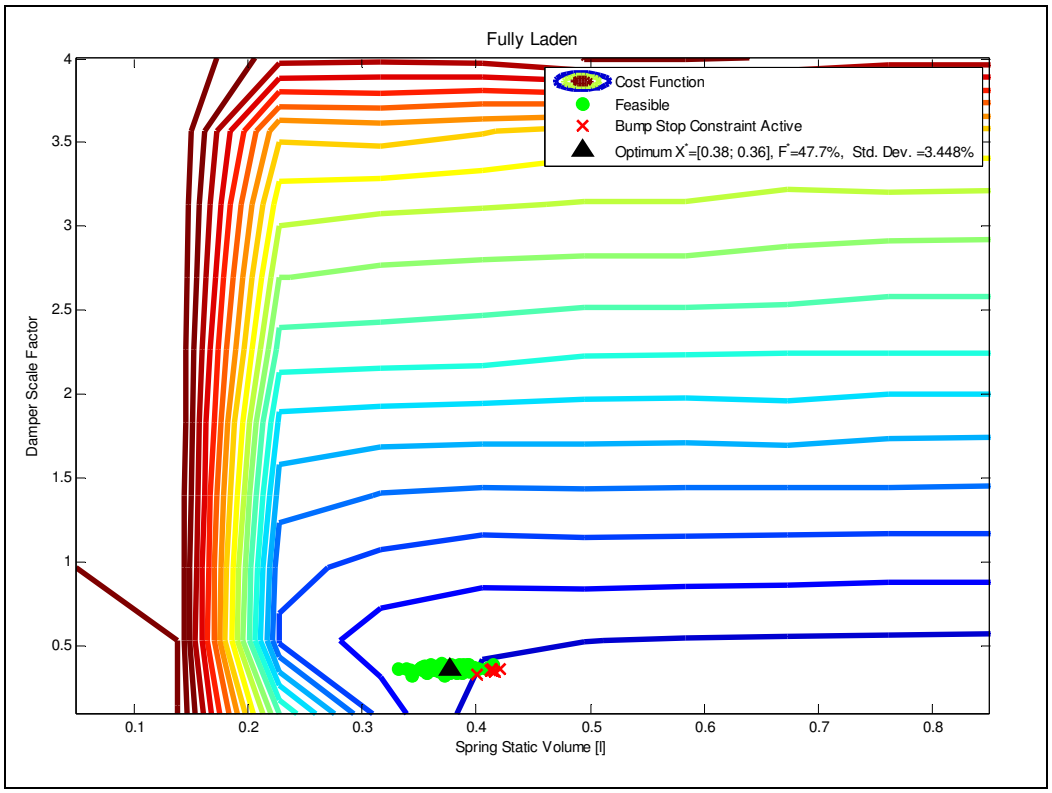

Figure 5-10: Contour plot of the fully laden cost function including the robust optimum.

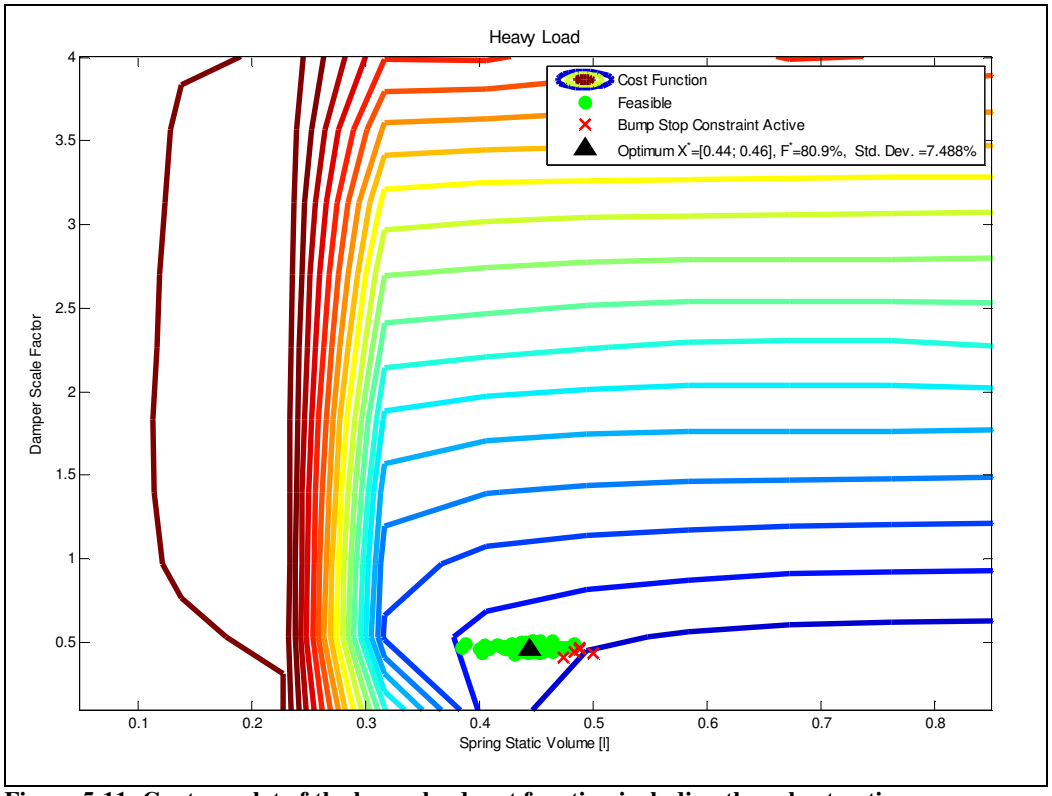

Figure 5-11: Contour plot of the heavy load cost function including the robust optimum.

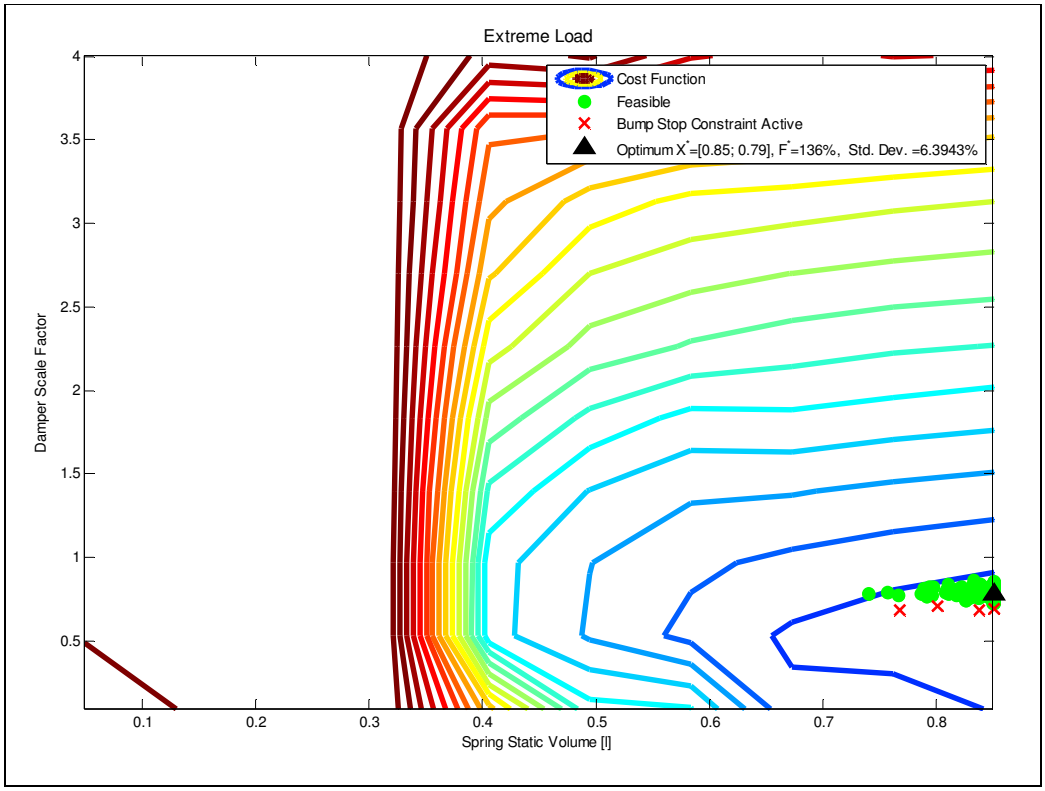

Figure 5-12: Contour plot of the extreme load cost function including the robust optimum.

The robust optima determined by this method for each of the cost functions have been plotted in Figure 5-9 to Figure 5-12. The robustness in terms of constraint violation was confirmed by conducting a second Monte Carlo simulation using 100 points around the robust optimum. The means and standard deviation for the cost function around the robust optimum are also shown.

The robust optima are compared to the deterministic optima in Table 5-6. The damper scale factors required for the robust optima are all slightly larger than those of the deterministic optima. This avoids bump stop contact and causes a slight increase in cost function value.

Table 5-6: Comparison of robust and deterministic optima.

	Deterministic			**Robust**		
	Static Gas Vol. [ℓ]	Damper Scale Factor	Cost Function Value [% of Baseline]	Static Gas Vol. [ℓ]	Damper Scale Factor	Cost Function Value [% of Baseline]
	Unladen					
Mean (Std. Dev.)	0.49	0.32	20%	0.49 (0.025)	0.34 (0.017)	22.3% (1.25%)
	Fully Laden					
Mean (Std. Dev.)	0.39	0.34	34.1%	0.38 (0.019)	0.36 (0.018)	47.7% (3.45%)
	Heavy Load					
Mean (Std. Dev.)	0.47	0.43	65.8%	0.44 (0.022)	0.46 (0.023)	80.9% (7.49%)
	Extreme Load					
Mean (Std. Dev.)	0.85	0.72	122%	0.85 (0.043)	0.79 (0.040)	136% (6.39%)

The number of infeasible points in the Monte Carlo simulations was 3 for the unladen case, 7 for the fully laden case, 5 for the heavy load case and 4 for the extreme load case. This implies 95.25% of the points evaluated lie within the feasible region. The strategy for finding a robust optimum thus seems to be successful.

The success of the technique for determining robust optima, developed in this section, provides some confidence that it could prove equally successful if extended to multiple variables. The technique can now be applied to find the robust optima from the deterministic optima in more variables.

5.6 Multi-variable optimisation

The optimisation methodology which has been developed and refined using the two-variable problem is now applied to problems with more variables in this section. The seven-variable problem originally formulated in 5.2 is attempted. A three-variable formulation of the optimisation problem is also constructed in order to ease the interpretation of the results of the seven-variable optimisation.

5.6.1 Seven-variable optimisation

The Dynamic-Q algorithm was once again applied to find the feasible minima of each of the cost functions. Several different starting points distributed across the design space were used to ensure that the converged optima correspond to the global optimum. The move limit was set to $\sqrt{n}/6$ and the starting points used are listed in Table 5-7. The starting points are similar to those used for the two-variable optimisation. The optima from the two-variable optimisation have been added as starting points (last row of data in Table 5-7).

Table 5-7: Multi-variable optimisation starting points.

Unladen		**Fully Laden**		**Heavy Load**		**Extreme Load**	
Static Gas Vol. [ℓ]	Damper Scale Factors	Static Gas Vol. [ℓ]	Damper Scale Factors	Static Gas Vol. [ℓ]	Damper Scale Factors	Static Gas Vol. [ℓ]	Damper Scale Factors
0.3	0.1	0.3	1	0.3	1	0.3	1
-	-	0.3	1.5	0.3	2.5	0.3	4
0.2	0.9	0.2	0.9	0.2	0.9	0.2	0.9
0.2	3.2	0.2	3.2	0.2	3.2	0.2	3.2
0.7	0.9	0.7	0.9	0.7	0.9	0.7	0.9
0.7	3.2	0.7	3.2	0.7	3.2	0.7	3.2
0.45	2	0.45	2	0.45	2	0.45	2
0.49	0.32	0.39	0.34	0.47	0.43	0.85	0.72

The convergence rate was again good from most starting points. The optimisation could not converge to a better optimum when started from the optima for the two-variable optimisation. This indicates that these starting points are at least local minima. The convergence history to the optima for the unladen payload case has been included in Figure 5-13 (Starting point: $sgv = 0.2\ell$, $dsf_{1\to 6} = 3.2$). The convergence histories to the optima for the remaining three load cases have been included in Appendix D.

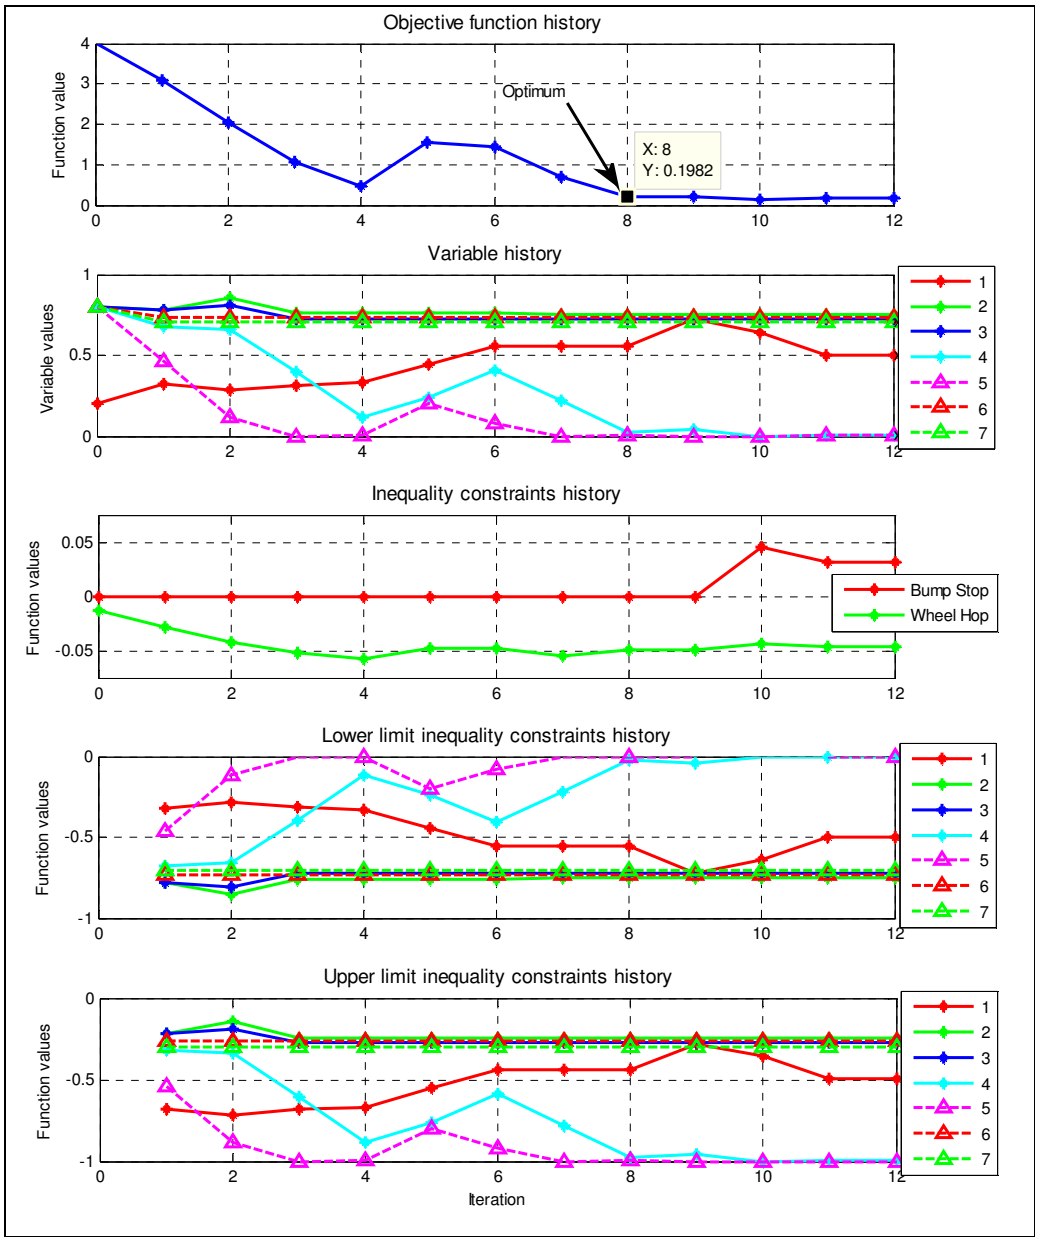

Figure 5-13: Convergence history to optimum for unladen payload case.

The design variable governing the static gas volume showed a definite trend towards large static gas volumes in all cases. The damping scale factors did not reveal such definite trends. Careful examination eventually revealed that design variables 4 and 5, which govern the damper characteristic below the blow-off velocity, tended towards low damping values. It is suspected that the other damping variables were rendered insensitive due to the location of the blow-off knee. This occurs due to the blow-off knee being located at higher velocities than are present in the simulation data. This was true for at least some cases. The optimum damper characteristics for the different load cases have been superimposed in Figure 5-14. It is difficult to distinguish any strong trends from the results. The reader should note that the positive velocity damper characteristics for the unladen, heavy load and extreme load cases are exactly the same.

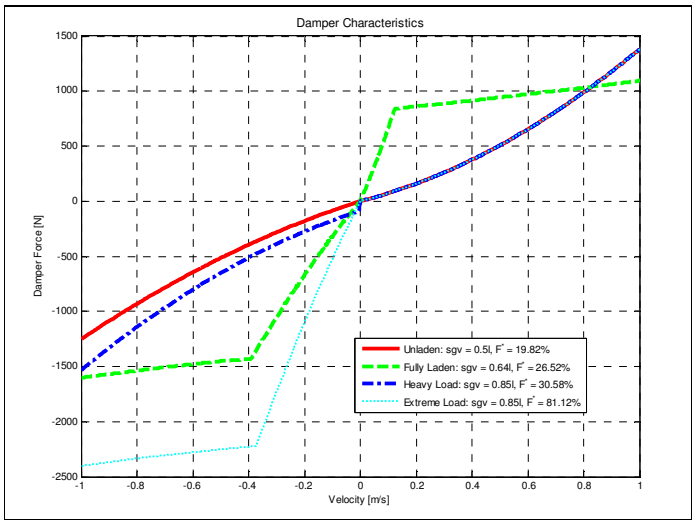

Figure 5-14: Optimum damper characteristics for multi-variable optimisation.

5.6.2 Three-variable optimisation

In an attempt to clarify the results and possibly reveal more definite trends, it was decided to simplify the multi-variable optimisation somewhat. The design variables were redefined such that the damper characteristic was defined by two damper scale factors. The first damper scale factor would govern the negative velocity region of the damper characteristic, while the second scaled the positive velocity region. Negative velocity in this case corresponds to the bump direction, while positive velocity indicates rebound. The definition of the optimisation variables, for the three-variable optimisation, was as follows:

$$ov_1 = \frac{(sgv - 0.05)}{0.8}$$

$$ov_2 = \frac{(dsf_{1\to3} - 0.1)}{3.9}$$

$$ov_3 = \frac{(dsf_{4\to6} - 0.1)}{3.9}$$

The starting points for the three-variable optimisation are the same as listed in Table 5-7. The optimisation converged well from almost all starting points. As with the two-variable optimisation, the algorithm converged to a wide range of near optimum points. The convergence trends were much more consistent than observed with the seven-variable optimisation. The first design variable once again tended towards large static gas volumes. An interesting observation is that the optima all show the same trend of higher damping in the bump direction, while the damping in the rebound direction is minimised. The cost function values at the optima show some improvements over the two-variable optimisation optima. The reduction in cost function values grow more pronounced with increase in payload. The optima obtained are presented in Table 5-8. The three-variable optimisation has thus succeeded in clarifying trends in the data. The optimum damper characteristics have been plotted in Figure 5-15.

Table 5-8: Three-variable optimisation results.

	Unladen			
	Static Gas Vol. [ℓ]	Damper Scale Factor (Bump)	Damper Scale Factor (Rebound)	Cost function value [% of baseline]
Range	0.35-0.85	0.3-0.8	0.1-0.3	18-29%
Mean	0.6	0.51	0.15	19%
Std. Dev.	0.25	0.32	0.2	4%
Optimum	0.74	0.6	0.13	18%
	Fully Laden			
	Static Gas Vol. [ℓ]	Damper Scale Factor (Bump)	Damper Scale Factor (Rebound)	Cost function value [% of baseline]
Range	0.40-0.85	0.32-0.56	0.1	24-28%
Mean	0.45	0.44	0.1	26%
Std. Dev.	0.15	0.2	0	2%
Optimum	0.43	0.365	0.1	24%
	Heavy Load			
	Static Gas Vol. [ℓ]	Damper Scale Factor (Bump)	Damper Scale Factor (Rebound)	Cost function value [% of baseline]
Range	0.50-0.85	0.39-0.69	0.1	30-40%
Mean	0.70	0.52	0.1	35%
Std. Dev.	0.22	0.23	0	3%
Optimum	0.74	0.58	0.1	30%
	Extreme Load			
	Static Gas Vol. [ℓ]	Damper Scale Factor (Bump)	Damper Scale Factor (Rebound)	Cost function value [% of baseline]
Range	0.85	0.81-1.11	0.1	76-86%
Mean	0.85	0.91	0.1	79%
Std. Dev.	0	0.22	0	4%
Optimum	0.85	0.81	0.1	76%

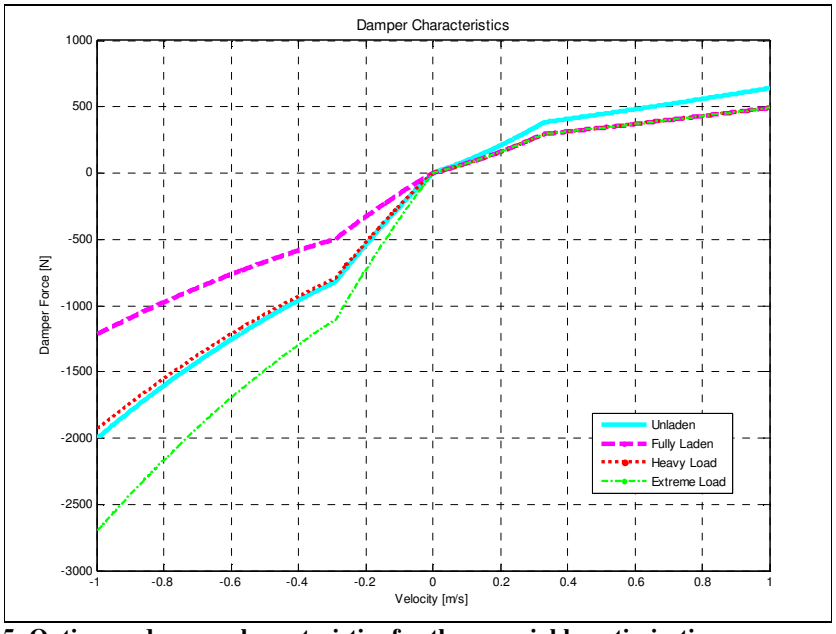

Figure 5-15: Optimum damper characteristics for three-variable optimisation.

5.6.3 Feasibility of the three-variable optimisation results

The cause of the reduction in cost function values with the change to an asymmetric damper characteristic was unclear. Determining the cause before the robust optima are determined is desirable, since the outcome of the inquiry could conceivably render the optima infeasible.

It was initially postulated that the asymmetric damper characteristic caused a steady increase in the vehicle ride height as the vehicle traversed the rough terrain. The increase in ride height would lead to a greater static gas volume, thereby reducing the stiffness of the non-linear hydro-pneumatic suspension. Evaluation of the simulation results with the optimum characteristics revealed that this was not the case. The ride height increased no more than 30 mm across the rough terrain. This was confirmed for all four payload cases considered. The terrain elevation rises and falls by more than 100 mm across the 100 m of the Belgian paving. A 30 mm increase in ride height is thus not considered significant. The reader is furthermore urged to recall that the optimisation problem being considered is relatively insensitive to changes in static gas volume. The increase in ride height is therefore not considered to be the cause of the reduced fatigue damage.

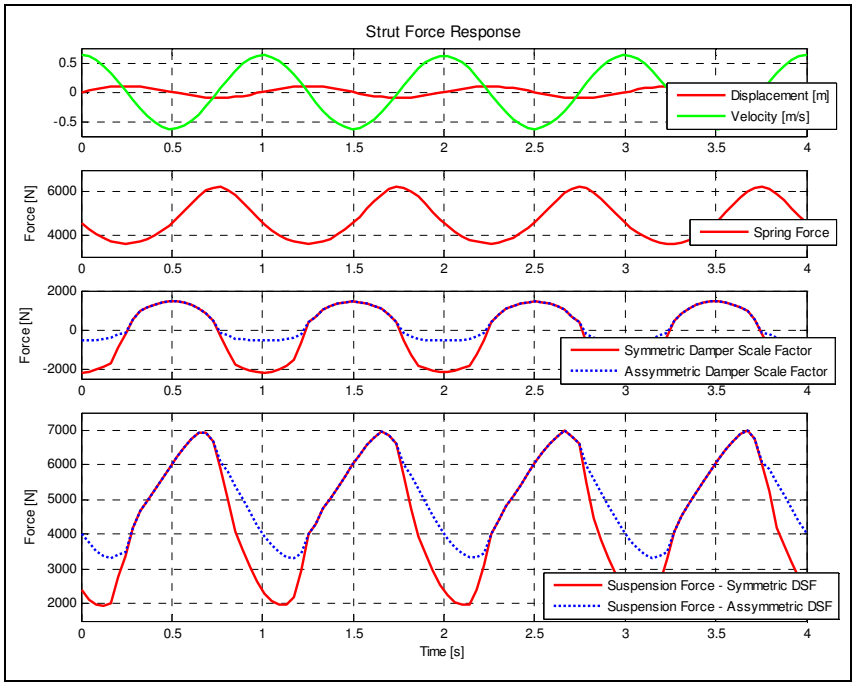

Figure 5-16: Strut force response to sinusoidal input.

The force response of the suspension strut to sinusoidal displacement input was evaluated in an attempt to explain the reduced damage due to the asymmetric characteristic. The input sinusoid had an amplitude of 0.1 m at 1 Hz frequency. The maximum velocity of the prescribed displacement was 0.62 m/s. The prescribed displacement and corresponding velocity are plotted in the first subplot of Figure 5-16. The hydro-pneumatic spring force was calculated from the model and is also presented in the figure. The damping force was calculated with equal damper scale factors in bump and rebound, as well as with asymmetric damper characteristic. The

friction force was added to the damper forces. The difference in the damper rebound force response should be clear. The total suspension force, which is equal to the sum of the spring, damper and friction forces, is presented in the last subplot of Figure 5-16. The force cycles created by the symmetric damper scale factors are much larger than those created by the asymmetric damper characteristic. It should be clear that the larger cycles would be detrimental to the damage.

5.6.4 Robust optima for multi-variable optimisation

The cause of the reduced function values which accompany the asymmetric damper characteristic has been determined and does not seem to pose a restriction on the feasibility of the optima determined in the three-variable optimisation. The process of determining the robust optima was thus initiated.

The method developed in paragraph 5.5.3. was again employed here. A Monte Carlo simulation was performed for each payload configuration. The design variables were modelled as normally distributed random variables with mean and standard deviations as presented in Table 5-8. Each Monte Carlo simulation was conducted using 1000 function evaluations. The minima achieved from these simulations are given in Table 5-9.

Table 5-9: Minima from the Monte Carlo simulations for the three-variable problem.

	Static Gas Vol. [ℓ]	Damper Scale Factor (bump)	Damper Scale Factor (rebound)	Cost function value [% of baseline]
Unladen	0.58	0.49	0.17	17.6%
Fully Laden	0.45	0.43	0.10	23%
Heavy Load	0.79	0.55	0.10	29%
Extreme Load	0.85	0.76	0.1	75.4%

The optima for all four payload cases tend to minimize the damper scale factor in the rebound direction. The damping in the bump direction is increased above that required by the symmetric damper scale factor case. The optima determined from the Monte Carlo simulations were then used to determine the robust optima. The robust optima were selected such that a maximum variable variation of 10% would not cause any constraint violations.

The robust optima are compared to the optima determined with the Monte Carlo simulations in Table 5-10. A final Monte Carlo simulation was performed with 1000 function evaluations to confirm the robustness of the optima. The mean values of the variables were set equal to the robust optima, while a standard deviation of 5% was assumed.

Table 5-10: Comparison of robust and deterministic optima for three variables.

	Deterministic				Robust			
	Static Gas Vol. [ℓ]	Damper Scale Factor (bump)	Damper Scale Factor (rebound)	Cost function value [%]	Static Gas Vol. [ℓ]	Damper Scale Factor (bump)	Damper Scale Factor (rebound)	Cost function value [%]
Unladen								
Mean (std.dev.)	0.58 -	0.49 -	0.17 -	17.6% -	0.58 (0.029)	0.52 (0.026)	0.17 (0.009)	25.3% (6%)
Fully Laden								
Mean (std.dev.)	0.45 -	0.43 -	0.1 -	23% -	0.44 (0.022)	0.46 (0.023)	0.1 (0.005)	25.6% (2.2%)
Heavy Load								
Mean (std.dev.)	0.79 -	0.55 -	0.1 -	29% -	0.82 (0.004)	0.60 (0.030)	0.1 (0.005)	28.4% (1.6%)
Extreme Load								
Mean (std.dev.)	0.85 -	0.76 -	0.1 -	75.4% -	0.85 (0.043)	0.82 (0.041)	0.1 (0.005)	82.6% (10%)

The resulting robust optima show a similar trend as observed with the two-variable optimisation. The damping required for the robust optima is slightly higher than for the absolute optima. The gain in robustness is achieved at a small cost in objective function value.

5.7 Load-levelling suspension effects

The optimal suspension characteristics for fatigue damage have now been determined for different payloads and different definitions of the damper characteristic. It was concluded that the cost functions were more sensitive to the damper characteristic than the static gas volume for all cases considered. The observation was stated to result from a combination of the hyperbolic increase in stiffness with decreasing static gas volume, and the load levelling operation of the hydro-pneumatic suspension system. The effects of replacing the load-levelling hydro-pneumatic suspension with a linear, non-load-levelling system, is critically examined in this section. The evaluation was achieved by comparing the cost functions of the current model with hydro-pneumatic suspension, to that of a non-load-levelling linear suspension system. The comparison was done for both the laden and extreme payload cases.

The vehicle model was modified for the purposes of the comparison to have a non-load levelling linear suspension system. The suspension stiffness was varied linearly between 10000 N/m and 180000 N/m. The range of stiffness is comparable to the variation of the linearised stiffness in the static position of the hydro-pneumatic suspension system employed during the optimisation. The static ride height was allowed to vary as the suspension stiffness was varied in both payload cases. Visualisations of the cost functions, similar to those presented in 5.5.2, were constructed and are presented in Figure 5-17.

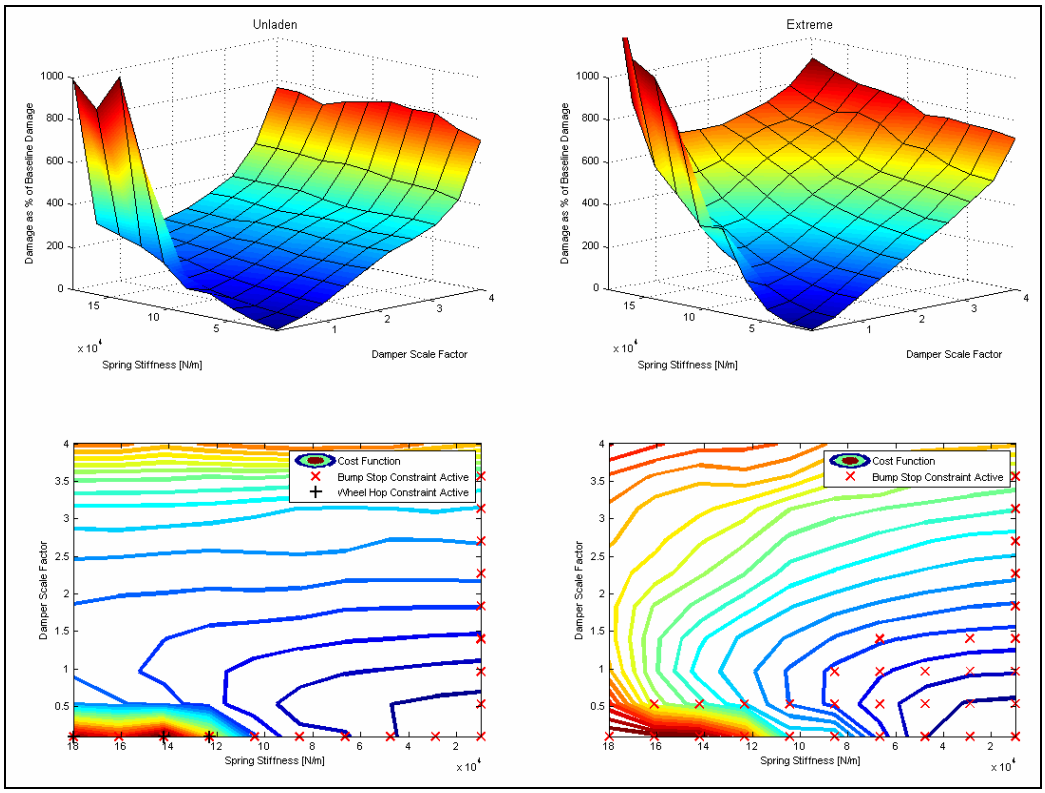

Figure 5-17: Cost function visualisation for linear, non-load-levelling suspension.

The unladen cost function for the linear, non-load-levelling suspension shows a higher sensitivity to spring stiffness, when compared to the cost function for load-levelling air suspension (Figure 5-3). Contrary to the observation made for load-levelling suspension, the cost function sensitivity to spring stiffness does not increase as significantly with increasing load. The cost function sensitivity to the damper scale factor seems to have remained the same as for load-levelling pneumatic suspension.

The infeasible regions of the design space are represented in the contour plots of Figure 5-17. The red crosses and black plusses, respectively indicate points where the bump stop and wheel hop constraints were active. The feasible region for the unladen payload case seems to be of similar size and location to that observed for the load-levelling suspension. The extreme payload case shows a greatly increased infeasible region, due to the bump stop constraint. This is to be expected as the suspension workspace decreases as ride height lowers with increasing payload, thereby reducing the available workspace.

The results indicate that a linear suspension system can achieve very similar fatigue performance to a hydro-pneumatic suspension system at the design payload state. The load-levelling hydro-pneumatic suspension system is however much more robust, both in terms of objective function value and bump stop constraint violation. The implication is that the load-levelling hydro-pneumatic suspension system is much more forgiving toward a vehicle designer, when variations in operational conditions are considered.

5.8 Conclusions to the mathematical optimisation and sensitivity analyses

Mathematical optimisation has been performed to find the optimal spring and damper characteristics for fatigue damage. The fatigue damage to the vehicle structure (calculated as in paragraph 3.3) was defined as the objective function for the mathematical optimisation. The suspension spring and damper characteristics were defined as design variables. A suspension workspace or bump stop constraint as well as a wheel contact constraint were defined in order to ensure that the optimal characteristics were realistic. The implementation of a third road damage constraint was unsuccessful.

The nature of the objective function and constraint functions was first studied using a simple two-variable formulation of the optimisation problem. The cost function scaling, search method parameters and general optimisation methodology were refined using this simpler form of the optimisation problem. The two-variable problem allowed visualisation of the cost functions which would not be possible for multi-variable formulations.

The two-variable problem revealed that the objective functions are uni-modal, smooth and well behaved throughout most of the design space. Significant numerical noise was however present, which influenced the quality of the gradient approximations and thus the convergence behaviour of the Dynamic-Q algorithm. This was especially true when constraint functions were introduced. The use of central finite difference gradient approximation alleviated the problem to some extent.

The objective functions for all four payload conditions considered showed minima in the low damping, low stiffness regions of the design space. A threefold improvement in cost function value was achieved for the unladen payload case. The optimisation converged to this region in all four load cases. The cost functions were found to be significantly more sensitive to the damper characteristic than to the spring static gas volume. The optima were shown to be highly robust with respect to cost function value, through the use of Monte Carlo simulation. The Monte Carlo simulations also revealed the shape of the constraint function boundaries

The optima are, however not robust in terms of constraint function violation. This is due to the fact that the optima are constrained by the dominant suspension workspace constraint. The suspension workspace constraint became active in the low damping regions of the design space. A method for determining constraint violation robust optima from the deterministic optima found by Dynamic-Q was developed. The method delivered suitably robust optima.

The optimisation was repeated with multi-variable formulations of the optimisation problem. The additional variables were used to define a complex damper characteristic. The seven-variable formulation failed to reveal clear trends for the damper characteristic. The complex formulation may have created a multi-modal design space with several local minima or regions of low sensitivity in some of the variables. A three-variable problem was formulated in an attempt to show clearer trends.

The three-variable formulation maintained the static gas volume as design variable. The damper characteristic was split into a bump and rebound characteristic, each described by a separate damper scale factor. A strong trend emerged in which the rebound damper scale factor was minimised, while the bump direction damper scale factor was increased from the two-variable optimum. The optimum for the unladen case showed that a cost function value of 20% of the baseline was attainable, i.e. the fatigue life of the vehicle could be improved by a factor of five by using the optimised spring and damper characteristics.

The effect of replacing the load-levelling pneumatic suspension with a linear non-load-levelling suspension was investigated. It was concluded that although similar optimum performance could be attained for both systems, the load-levelling suspension would provide a more robust solution.

Chapter 6

6. Conclusions and recommendations

The study set out to investigate the suspension characteristics required to extend vehicle structural life. The sensitivity of these characteristics to changes in payload was also required. The process followed was one of mathematical modelling and optimisation.

The instrumented Land Rover Defender 110 of the University of Pretoria's vehicle dynamics research group was chosen as test vehicle. A load cell was added to the vehicle in order to measure the suspension force in the left rear suspension. The left rear suspension mounting was also instrumented with six strain gauges to enable the calculation of damage to the structure. The vehicle dynamics was characterised over both discrete obstacles and random terrain. The dynamics was characterised for different suspension settings and vehicle speeds. The repeatability of the tests conducted was shown to be excellent. The results from the tests were used to validate a mathematical model of the test vehicle.

A mathematical model of the test vehicle was developed. The model complexity was kept as low as possible in order to ensure low computational expense. The model complexity was systematically increased as was required to yield adequate representation of the vehicle dynamics.

The model gave excellent correlation over symmetric discrete obstacles for all measured quantities considered. The model could not be validated over asymmetric discrete obstacles. The cause for the lack of correlation over asymmetric discrete obstacles is unclear, but the static friction model is thought to be a major contributor.

The model correlation over random rough terrain was excellent at high speed. The fatigue damage predicted by the model correlated well with experimental data. At low speed the model gave poor correlation. The poor low speed correlation is mainly attributed to the point follower tyre model and static friction model employed. The correlation in all cases was generally better with the suspension set to ride comfort mode. The greater participation of tyre dynamics when the suspension is set to handling mode is thought to be the cause of this occurrence.

The recommendation is made that tyre models for use on rough undeformable terrain be investigated to improve correlation. Detailed investigations into the dynamic behaviour of the friction on the $4S_4$ suspension system should also be conducted as this is especially critical for fatigue damage correlation. The friction investigations should include a study into the effects of working fluid contamination and degradation on the exposed sliding surfaces of the $4S_4$.

The mathematical model correlation, although imperfect, was considered adequate for mathematical optimisation and sensitivity studies. The model was considered validated for high speed driving over rough terrain.

Four different payload cases were considered for the purpose of the optimisation. The empty state in which the vehicle was tested was defined as the unladen payload case. The remaining payload cases considered were a fully laden case, representing the vehicle with 1 ton payload; a heavy load case, representing an armoured version of the Land Rover; and a fictitious extreme load case representing the unladen to fully laden payload variation in commercial truck-trailers. The suspension optimisation was conducted for all four payload cases.

The Dynamic-Q algorithm, which has been used successfully by Thoresson (2007) in the optimisation of vehicle suspension characteristics, was chosen for the optimisation. Fatigue damage to the vehicle structure, as predicted by the vehicle model was chosen as cost function. The fatigue damage was calculated by applying Miner's damage summation rule to the stress-time histories predicted by the vehicle model. The static gas volume which determines the characteristic spring characteristic of the hydro-pneumatic suspension system was chosen as the first design variable. A complex damper characteristic was defined using six additional design variables. Constraint functions were defined in order to restrict loss of wheel ground contact and bump stop contact. An attempt was made to include a third, road damage constraint, but little success was achieved with this. The road damage constraint requires further investigation.

A simple two-variable formulation of the optimisation problem was defined to aid in the development and refinement of the optimisation methodology. The two-variable formulation proved invaluable as it allowed visualisation of the cost function which was not possible with multi-variable formulations of the problem. The visualisations of the objective functions for the four payload cases revealed that the cost functions are generally smooth and uni-modal across the design space. The minima could be found in the low spring stiffness, low damping region of the design space. The cost functions exhibited a much greater sensitivity to the damper scale factors, than to the static gas volume. The relative insensitivity to static gas volume was attributed to the load-levelling functionality of the hydro-pneumatic suspension system and the hyperbolic nature of the suspension system's stiffness as a function of static gas volume. The sensitivity to static gas volume increased with increasing payload.

The optimisation converged to a wide, shallow, dish-like region surrounding the optima. The optimum for the unladen case delivered a threefold improvement in cost function value over the baseline. The convergence was generally good, though the quality of the convergence deteriorated significantly near the optima. It was concluded that the deteriorating convergence rate was due to the presence of high numerical noise in the cost function and a very shallow design space.

The large region of near optimal values was however indicative of highly robust optima in terms of cost function value. Monte Carlo simulations were conducted in the region of the optima to evaluate the robustness of the optima. The investigation proved the robustness of the optima in terms of cost function value and revealed that the optima were all constrained by the dominant bump stop contact constraint. A clear constraint boundary could be distinguished at low damping regions on contour plots of the objective functions.

Although the optima were highly robust in terms of cost function value, their location on the constraint function boundary rendered them non-robust in terms of constraint violations. A method for selecting constraint robust optima from optima determined by the search algorithm was developed. The method was shown to deliver constraint robust optima. The method proved promising, but should be further refined for future work.

The optimisation was repeated with the seven-variable formulations of the optimisation problem. The results were unsatisfactory as it was difficult to determine clear trends from the data. The problem was once again reformulated to a simpler form in an attempt to reveal clearer trends. A three-variable formulation was selected, where the static gas volume was kept as variable, while the damper characteristic was described by a single variable for bump and rebound directions respectively. The three-variable formulation proved successful in revealing clearer trends.

The optima for the three-variable optimisation were found using Dynamic-Q and their robust counterparts were determined using the methodology developed. The damper characteristic for the rebound direction was minimised, while the bump characteristic was increased slightly above that of the two-variable optimum. A further improvement in the objective function values was achieved. The improvement in cost function value grew more pronounced as payload increased. The attainable improvements in fatigue damage warrant the possible implementation of fatigue damage optimised suspension characteristics.

The effect of using linear spring characteristics instead of the load-levelling hydro-pneumatic suspension was investigated. It was concluded that the linear characteristic could possibly perform as well as the load-levelling hydro-pneumatic at a single design state. The robustness, both in terms of cost and constraint functions, imparted by the load-levelling effects of the hydro-pneumatic suspension system leads to the recommendation of the hydro-pneumatic suspension as the preferred system.

The use of the $4S_4$ suspension system cannot be recommended solely based on a fatigue damage argument. This is due to the fact that the optima required similar static gas volumes and the optimisation problems considered generally exhibited a greater sensitivity for the damper characteristic. If fatigue damage is the only discriminator a variable damping system would be sufficient to ensure optimal suspension characteristics at varying payloads. It was however shown that the fatigue damage optimal suspension characteristics are biased towards the ride comfort optimal suspension characteristics. The fatigue damage optimal characteristics may therefore lead to compromised handling. The $4S_4$ suspension system may prove invaluable if additional performance requirements, such as vehicle handling are considered.

Bibliography

Armstrong-Hélouvry, B., DuPont, P., De Wit, C.C., 1994, A Survey of Models, Analysis Tools and Compensation Methods for the Control of Machines and Friction, *Automatica*, 30 (7), pp. 1083-1138.

Bekker, C.M., 2008, *Profiling of rough terrain*, Unpublished Masters Dissertation, University of Pretoria, South Africa.

Beyer, H.-G., Sendhoff, B., 2007, Robust optimization – A comprehensive survey, *Computer methods in the applied mechanics and engineering*, 196, pp. 3190-3218.

Bode, O., von Glasner, E.-C., Holnkaus, H., Pflug, H.-C., Povel, R., 1997, Einfluss von Fahrzustand und Fahrwerk auf die Vertikaldynamik von Nutzfahrzeugen, *Automobiltechnische Zeitschrift*, 99 (7/8), pp. 396-402.

Bonchis, A., Corke, P.I., Rye, D.C., 1999, A Pressure-Based, Velocity Independent, Friction Model for Asymmetric Hydraulic Cylinders, *Proceedings of the 1999 IEEE International Conference on Robotics & Automation*, pp. 1746-1751.

Burden, R.L., Faires, J.D., 2005, *Numerical Analysis*. 8th Ed. Thomson - Brooks/Cole.

Callister, D.J., 2003, *Materials Science and Engineering: An Introduction*. 6th Ed. John Wiley and Sons.

Cebon, D., 1989, Vehicle-Generated Road Damage: A Review, *Vehicle System Dynamics*, 89, pp. 107-150.

Cebon, D., 1999, *Handbook of Vehicle-Road Interaction*, Swets & Zeitlinger.

Cole, D.J., 2001, Fundamental Issues in Suspension Design for Heavy Road Vehicles, *Vehicle System Dynamics*, 35 (4-5), pp. 319-360.

Cole, D.J., Cebon, D., 1996, Truck suspension design to minimize road damage, *Proc. Institution of Mechanical Engineers*, 210:2 Part D, pp. 95-107. (Indirect reference: cited by Cole, 2001)

Conle, F.A., Chu, C.-C, 1997, Fatigue analysis and local stress-strain approach in complex vehicular structures, *International Journal of Fatigue*, Vol. 19 Supp. No. 1, pp. S317-S323.

Conle, F.A., Mousseau, C.W., 1991, Using vehicle dynamics simulations and finite-element results to generate fatigue life contours for chassis components, *International Journal of Fatigue*, 13 (3), pp. 195-205.

Cronje, P.H., 2008, *Improving off-road vehicle handling using an active anti-roll bar*, Unpublished Masters Dissertation, University of Pretoria, South Africa.

Dowling, N.E., 2007, *Mechanical Behaviour of Materials: Engineering Methods for Deformation, Fracture, and Fatigue*, 3rd Ed. Pearson - Prentice Hall.

Els, P.S., 2006, *The ride comfort vs. handling compromise for off-road vehicles*, Unpublished PhD Thesis, University of Pretoria, South Africa.

Els, P.S., Uys, P.E., Snyman, J.A., 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*, 43, pp. 787-801.

Gere, J.M., 2004, *Mechanics of Materials*. 6th Ed. Thomson – Brooks/Cole.

Gerotek Vehicle Testing Facility,
http://www.armscorbusiness.com/Subsites/Gerotek/Gerotek01_landing.asp (accessed on 2009/09/18)

Haiba, M., Barton, D.C., Brooks, P.C., Levesley, M.C., 2002, Review of life assessment techniques applied to dynamically loaded automotive components, *Computers and Structures*, 80, pp. 481-494.

Holdmann, P., Holle, M., 1999, Possibiliies to improve ride and handling performance of delivery trucks by modern mechatronic systems, *JSAE Rev*, 20:505-10. (Indirect reference: cited by Uys et. al., 2004)

Hunt, P.D., Bunker, J.M., 2003, Time Series Analysis of Pavement Roughness Condition Data for Use in Asset Management, *Proceedings: Asia Pacific Conference on Systems Integrity and Maintenance – Cairns, Australia*, pp. 148-154.

International Organization for Standardization, 1995, *International Standard ISO 8606: Mechanical vibration – Road surface profiles – Reporting of measured data*, ISO 8608:1995(E).

Ivers, D.E., Miller, L.R., 1989, Experimental comparison of passive, semi-active on/off, and semi-active continuous suspensions, *SAE Technical paper: Advanved truck suspensions SP-802*, Paper no. 892484 (Indirect reference: cited by Uys et. al., 2004)

Johnson, R.A., 2005, *Probability and Statistics for Engineers*, 7th Ed., Pearson Prentice Hall.

Kat, C., 2009, *Suspension forces on a tri-axle air suspended semi-trailer*, Unpublished Masters Dissertation, University of Pretoria, South Africa.

Letherwood, M.D., Gunter, D.D., 2001, Ground vehicle modelling and simulation of military vehicles using high performance computing, *Parallel Computing*, 27, pp. 109-140.

Lee, K.-H., Park, G.-J, 2000, Robust opimization considering tolerances of design variables, *Computers and Structures*, 79, pp. 77-86.

Li, H., Hu, S.J., 2002, Tuned Mass Damper Design for Optimally Minimizing Fatigue Damage, *Journal of Engineering Mechanics*, June 2002, pp. 703-707.

Li, L., Sandu, C., 2007, On the impact of cargo weight, vehicle parameters, and terrain characteristics on the prediction of traction for off-road vehicles, *Journal of Terramechanics*, 44, p. 221-238.

Medepalli, S., Rao, R., 2000, Prediction of road loads for fatigue design – a sensitivity study, *International Journal of Vehicle Design*, 23 (1/2), p. 161-175.

Mrzyglod, M., Zielinski, A.P., 2007, Multiaxial high-cycle fatigue constraints in structural optimization, *International Journal of Fatigue*, 29, pp. 1920-1926.

Sayers, M.W., Karamihas, S.M., 1998, *The Little Book of Profiling*, University of Michigan.

Sonntag, R.E. , Borgnakke, C., Van Wylen, G.J., 2003. *Fundamentals of Thermodynamics*. 6th Ed. John Wiley and Sons.

Snyman, J.A., Hay, A.M., 2002, The Dynamic-Q Optimization Method: An Alternative to SQP?, *Computers and Math. Applic.*, 44, pp. 1589-1598.

Snyman, J.A., 2004, *Practical Mathematical Optimization*, Springer.

Standards South Africa, 2005, *South African National Standard SANS 10162: The Structural Use of Steel*, SANS 10162-1:2005 2nd Ed.

Steyn, W.J. van der Merwe, 2001, *Considerations of Vehicle-Pavement Interaction for Pavement Design*, Unpublished PhD Thesis, University of Pretoria, South Africa.

Sweatman, P.F., 1983, *A Study of Dynamic Wheel Forces in Axle Group Suspensions of Heavy Vehicles*, Australian Road Research Board, Special Report No. 27.

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

Thoresson, M.J., 2007, *Efficient Gradient-Based Optimisation of Suspension Characteristics for An Off-Road Vehicle.*, Unpublished PhD Thesis, University of Pretoria, South Africa.

Uys, P.E., Els, P.S., Thoresson, M.J., 2004, Criteria for handling measurement, *Journal of Terramechanics*, 43, pp. 43-67.

Uys, P.E., Els, P.S., Thoresson, M.J., 2007, Suspension settings for optimal ride comfort of off-road vehicles travelling on roads with different roughness and speeds, *Journal of Terramechanics*, 44, pp. 163-175.

Uys, P.E., Els, P.S., Thoresson, M.J., Voigt, K.G., Combrinck, W.C., 2006, Experimental determination of moments of inertia for an off-road vehicle in a regular engineering laboratory, *International Journal of Mechanical Engineering Education*, 34 (4), pp. 291-314.

Vaughan, J., Singhose, W., Sadegh, N., 2003, Use of Active Suspension Control to Counter the Effects of Vehicle Payloads, *IEEE*, pp. 285-289.

von Glasner, E.-C., Goehring, E., Povel, R., Schuetzner, P., 1993, Analysis of Intelligent Suspension Systems for Commercial Vehicles, *SAE Transactions*, Paper no. 933008, pp. 35-43.

Wannenburg, J., 2007, *A study of fatigue loading on automotive and transport structures*, Unpublished Phd Thesis, University of Pretoria, South Africa.

White, F.M., 2003, *Fluid Mechanics*. 5th Ed. McGraw-Hill.

Yanada, H., Sekikawa, Y., 2008, Modeling of Dynamic Behaviors of Friction, *Mechatronics*, 18, pp. 330-339.

Zeiler, T.A., Barkey, M.E., 2001, Design sensitivities of fatigue performance and structural dynamic response in an automotive application, *Structural and Multidisciplinary Optimization*, 21, pp. 309-315.

Zeiler, T.A., 2002, Use of strucural dynamic and fatigue sensitivity derivatives in an automotive design optimization, *Structural and Multidisciplinary Optimization*, 23, pp. 390-397.

Appendix

A. Appendix A

Derivation of equations of motion for seven DOF pitch plane model

Consider Figure 3-1, reproduced here for the sake of convenience as Figure A-1.

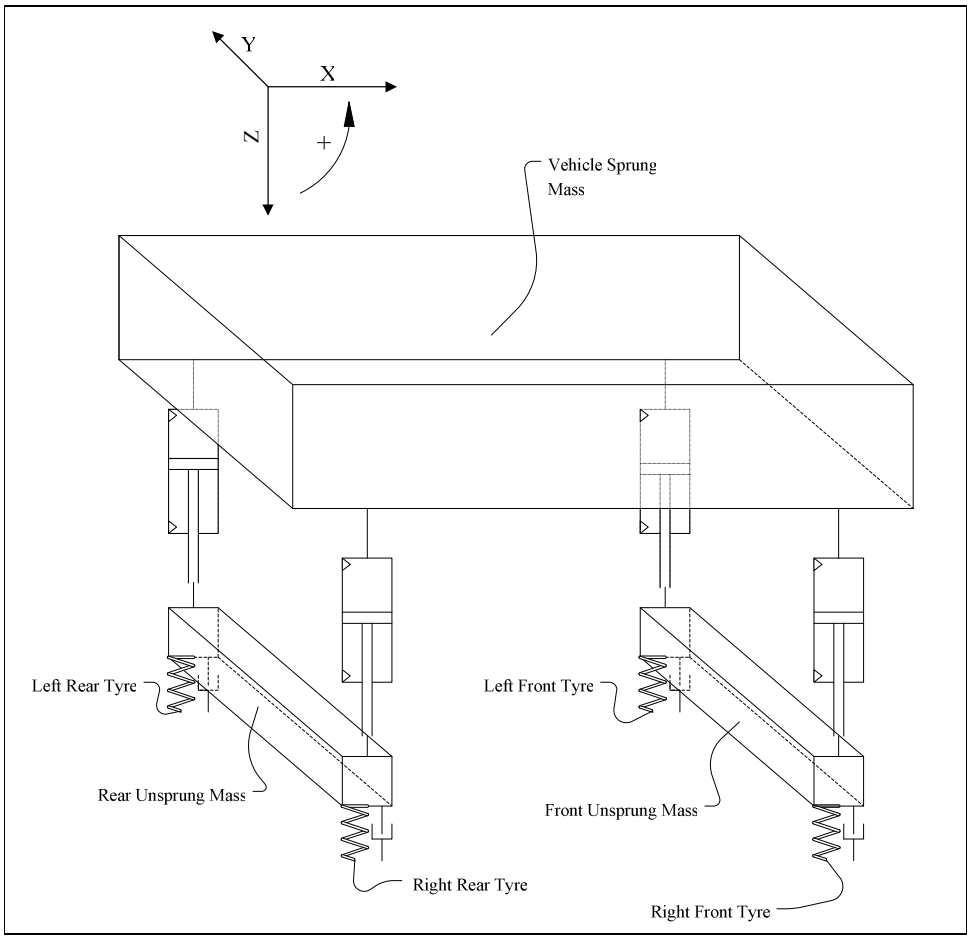

Figure A-1: Seven DOF vehicle model.

The equations of motion are set up for each of the DOF of the system by considering the Free-body Diagram (FBD) and Kinetic Diagram (KD) of each mass in the system. Let us first consider the equations of motion of the vehicle sprung mass in the pitch plane, Figure A-2.

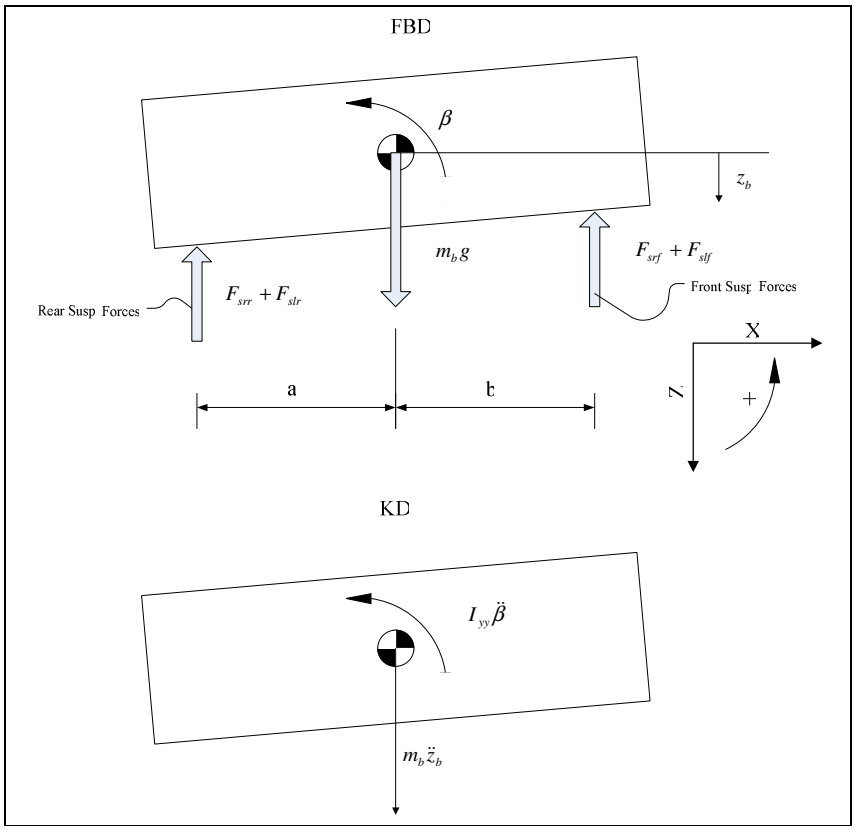

Figure A-2: FBD and KD of sprung mass in the pitch plane.

Taking the sum of the forces in the vertical direction yields:

$$\sum F_z = m\ddot{z}$$

$$m_b \ddot{z}_b = m_b g - F_{srr} - F_{slr} - F_{slf} - F_{srf}$$

Taking the sum of the moments about the centre of mass and assuming that β remains small $(\sin\beta \cong \beta, \cos\beta \cong 1)$ yields:

$$\sum M_{CG} = I_{yy}\ddot{\beta}$$

$$I_{yy}\ddot{\beta} = (F_{slf} + F_{srf})b - (F_{srr} + F_{slr})a$$

This gives two of the three equations of motion associated with the vehicle sprung mass. The final equation of motion for the body is obtained by considering the free-body and kinetic diagram of the body in the pitch plane, Figure A-3.

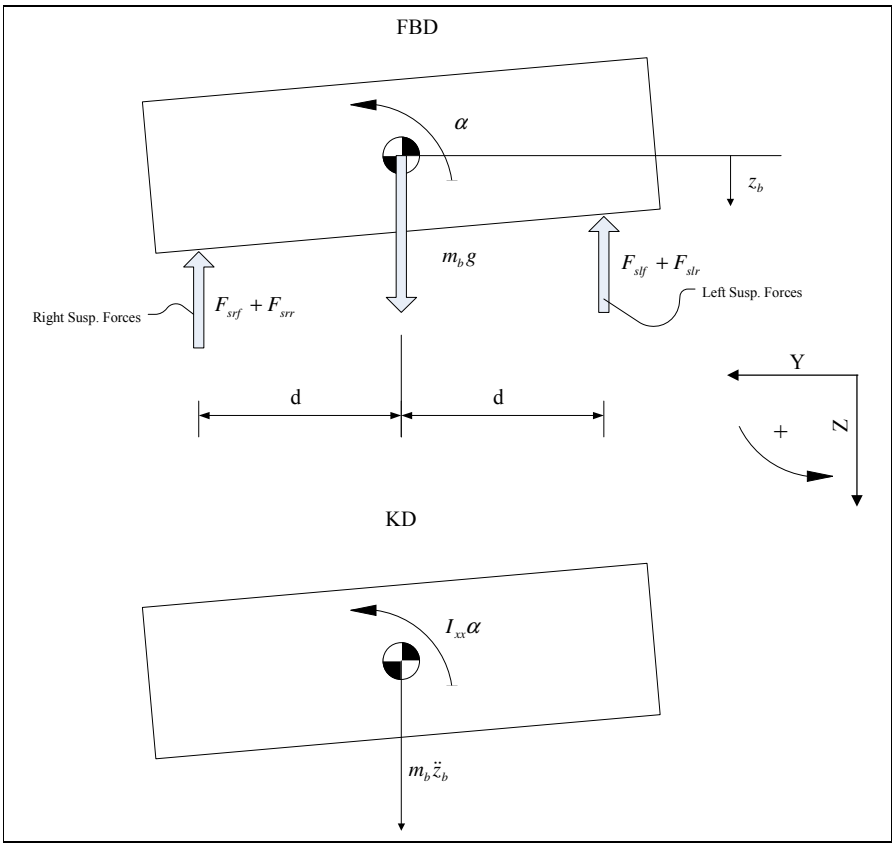

Figure A-3: FBD and KD of sprung mass in the roll plane.

Taking the sum of the moments about the centre of mass and assuming that α remains small $(\sin\alpha \cong \alpha, \cos\alpha \cong 1)$ yields:

$$\sum M_{CG} = I_{xx}\ddot{\alpha}$$

The remaining four equations of motion relate to the front and rear unsprung masses respectively. They are found by observing the free-body and kinetic diagrams of the axles in the pitch plane. The free-body diagrams for the front and rear axles are shown in Figure A-4.

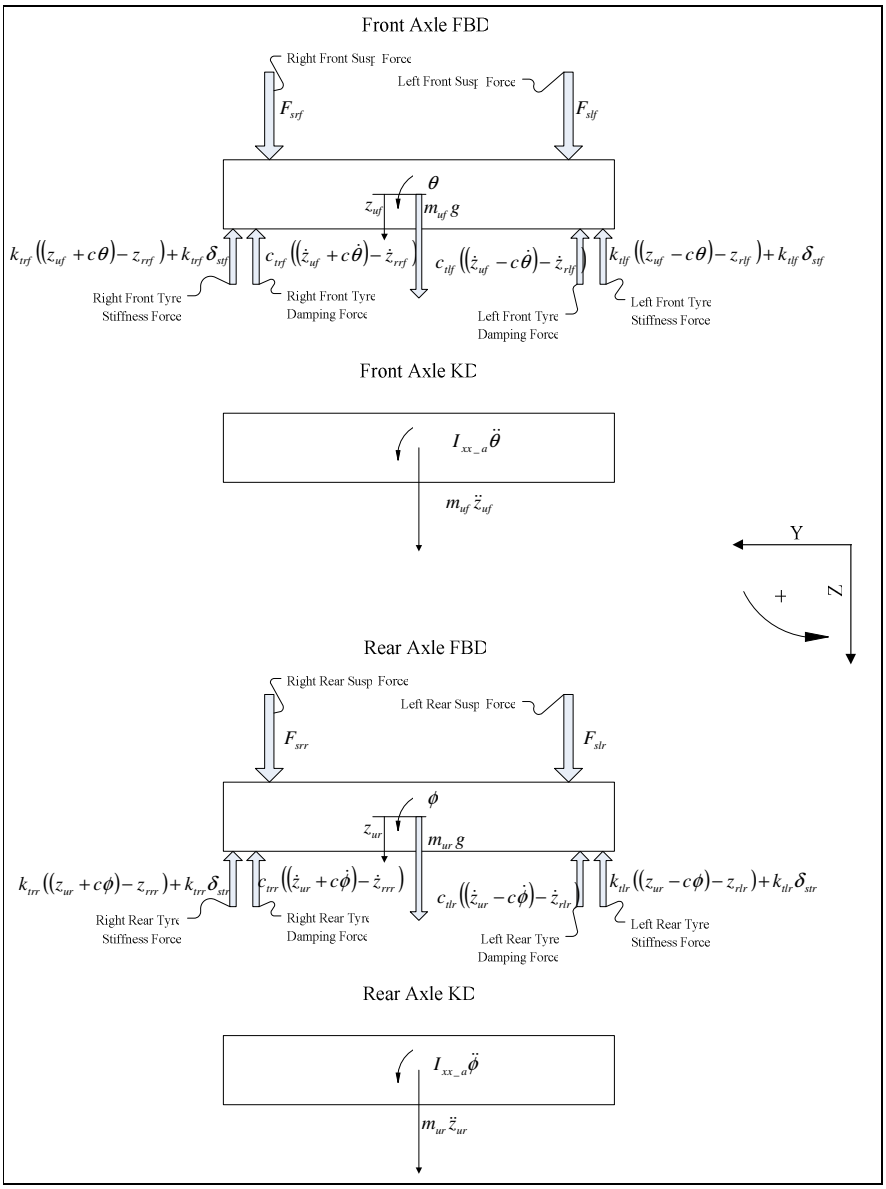

Figure A-4: FBD and KD of front and rear unsprung masses in the roll plane.

The sum of the forces in the vertical direction for the front unsprung mass gives:

$$\sum F_z = m\ddot{z}$$

$$m_{uf}\ddot{z}_{uf} = F_{srf} + F_{slf} - k_{trf}\left(\left(z_{uf}+c\theta\right)-z_{rrf}\right)-c_{trf}\left(\left(\dot{z}_{uf}+c\dot{\theta}\right)-\dot{z}_{rrf}\right)-k_{trf}\delta_{stf},$$

which rearranges to:

$$m_{uf}\ddot{z}_{uf} + k_{trf}z_{rrf} + k_{trf}c\theta + c_{trf}\dot{z}_{uf} + c_{trf}c\dot{\theta} + k_{tlf}z_{uf} - k_{tlf}c\theta + c_{tlf}\dot{z}_{uf} - c_{tlf}c\dot{\theta}\ldots$$
$$= k_{trf}z_{rrf} + c_{trf}\dot{z}_{rrf} - k_{trf}\delta_{stf} + k_{tlf}z_{rlf} + c_{tlf}\dot{z}_{rlf} - k_{tlf}\delta_{stf} + m_{uf}g + F_{srf} + F_{slf}$$

Similarly, taking the sum of the forces in the vertical direction for the rear unsprung mass leads to:

$$m_{ur}\ddot{z}_{ur} + k_{trr}z_{rrr} + k_{trr}c\phi + c_{trr}\dot{z}_{ur} + c_{trr}c\dot{\phi} + k_{tlr}z_{ur} - k_{tlr}c\phi + c_{tlr}\dot{z}_{ur} - c_{tlr}c\dot{\phi}\ldots$$
$$= k_{trr}z_{rrr} + c_{trr}\dot{z}_{rrr} - k_{trr}\delta_{str} + k_{tlr}z_{rlr} + c_{tlr}\dot{z}_{rlr} - k_{tlr}\delta_{str} + m_{ur}g + F_{srr} + F_{slr}$$

Taking the moments about the centre of mass of the front unsprung mass yields:

$$\sum M_{CG} = I_{xx\_a}\ddot{\theta}$$

$$\begin{aligned} I_{xx\_a}\ddot{\theta} = {} & c\left[k_{tlf}\left(\left(z_{uf} - c\theta\right) - z_{rlf}\right) + c_{tlf}\left(\left(\dot{z}_{uf} - c\dot{\theta}\right) - \dot{z}_{rlf}\right) + k_{tlf}\delta_{stf}\right]\ldots \\ & - c\left[k_{trf}\left(\left(z_{uf} + c\theta\right) - z_{rrf}\right) + c_{trf}\left(\left(\dot{z}_{uf} + c\dot{\theta}\right) - \dot{z}_{rrf}\right) + k_{trf}\delta_{stf}\right]\ldots \\ & + d\left(F_{srf} - F_{slf}\right) \end{aligned}$$

After some rearrangement the above expression becomes:

$$I_{xx\_a}\ddot{\theta} - ck_{tlf}z_{uf} + c^2k_{tlf}\theta - cc_{tlf}\dot{z}_{uf} + c^2c_{tlf}\dot{\theta} + ck_{trf}z_{uf} + c^2k_{trf}\theta + cc_{trf}\dot{z}_{uf} + c^2c_{trf}\dot{\theta}\ldots$$
$$= -ck_{tlf}z_{rlf} - cc_{tlf}\dot{z}_{rlf} + ck_{tlf}\delta_{stf} + ck_{trf}z_{rrf} + cc_{trf}\dot{z}_{rrf} - ck_{trf}\delta_{stf} + d\left(F_{srf} - F_{slf}\right)$$

A similar process performed on the rear axle delivers the final equation of motion:

$$I_{xx\_a}\ddot{\phi} - ck_{tlr}z_{ur} + c^2k_{tlr}\phi - cc_{tlr}\dot{z}_{ur} + c^2c_{tlr}\dot{\phi} + ck_{trr}z_{ur} + c^2k_{trr}\phi + cc_{trr}\dot{z}_{ur} + c^2c_{trr}\dot{\phi}\ldots$$
$$= -ck_{tlr}z_{rlr} - cc_{tlr}\dot{z}_{rlr} + ck_{tlr}\delta_{str} + ck_{trr}z_{rrr} + cc_{trr}\dot{z}_{rrr} - ck_{trr}\delta_{str} + d\left(F_{srr} - F_{slr}\right)$$

The seven equations of motion for the system are thus:

$$m_b \ddot{z}_b = m_b g - F_{srr} - F_{slr} - F_{slf} - F_{srf}$$

$$m_{uf} \ddot{z}_{uf} + k_{trf} z_{rrf} + k_{trf} c\theta + c_{trf} \dot{z}_{uf} + c_{trf} c\dot{\theta} + k_{tlf} z_{uf} - k_{tlf} c\theta + c_{tlf} \dot{z}_{uf} - c_{tlf} c\dot{\theta} \ldots$$
$$= k_{trf} z_{rrf} + c_{trf} \dot{z}_{rrf} - k_{trf} \delta_{stf} + k_{tlf} z_{rlf} + c_{tlf} \dot{z}_{rlf} - k_{tlf} \delta_{stf} + m_{uf} g + F_{srf} + F_{slf}$$

$$m_{ur} \ddot{z}_{ur} + k_{trr} z_{rrr} + k_{trr} c\phi + c_{trr} \dot{z}_{ur} + c_{trr} c\dot{\phi} + k_{tlr} z_{ur} - k_{tlr} c\phi + c_{tlr} \dot{z}_{ur} - c_{tlr} c\dot{\phi} \ldots$$
$$= k_{trr} z_{rrr} + c_{trr} \dot{z}_{rrr} - k_{trr} \delta_{str} + k_{tlr} z_{rlr} + c_{tlr} \dot{z}_{rlr} - k_{tlr} \delta_{str} + m_{ur} g + F_{srr} + F_{slr}$$

$$I_{xx} \ddot{\alpha} = (F_{slr} + F_{slf})d - (F_{srr} + F_{srf})d$$

$$I_{yy} \ddot{\beta} = (F_{slf} + F_{srf})b - (F_{srr} + F_{slr})a$$

$$I_{xx\_a} \ddot{\theta} - ck_{tlf} z_{uf} + c^2 k_{tlf} \theta - cc_{tlf} \dot{z}_{uf} + c^2 c_{tlf} \dot{\theta} + ck_{trf} z_{uf} + c^2 k_{trf} \theta + cc_{trf} \dot{z}_{uf} + c^2 c_{trf} \dot{\theta} \ldots$$
$$= -ck_{tlf} z_{rlf} - cc_{tlf} \dot{z}_{rlf} + ck_{tlf} \delta_{stf} + ck_{trf} z_{rrf} + cc_{trf} \dot{z}_{rrf} - ck_{trf} \delta_{stf} + d(F_{srf} - F_{slf})$$

$$I_{xx\_a} \ddot{\phi} - ck_{tlr} z_{ur} + c^2 k_{tlr} \phi - cc_{tlr} \dot{z}_{ur} + c^2 c_{tlr} \dot{\phi} + ck_{trr} z_{ur} + c^2 k_{trr} \phi + cc_{trr} \dot{z}_{ur} + c^2 c_{trr} \dot{\phi} \ldots$$
$$= -ck_{tlr} z_{rlr} - cc_{tlr} \dot{z}_{rlr} + ck_{tlr} \delta_{str} + ck_{trr} z_{rrr} + cc_{trr} \dot{z}_{rrr} - ck_{trr} \delta_{str} + d(F_{srr} - F_{slr})$$

Let us now define $\{z\}=\begin{Bmatrix} z_b \\ z_{uf} \\ z_{ur} \\ \alpha \\ \beta \\ \theta \\ \phi \end{Bmatrix}$, with first and second time derivatives $\{\dot{z}\}=\begin{Bmatrix} \dot{z}_b \\ \dot{z}_{uf} \\ \dot{z}_{ur} \\ \dot{\alpha} \\ \dot{\beta} \\ \dot{\theta} \\ \dot{\phi} \end{Bmatrix}$ and

$\{\ddot{z}\}=\begin{Bmatrix} \ddot{z}_b \\ \ddot{z}_{uf} \\ \ddot{z}_{ur} \\ \ddot{\alpha} \\ \ddot{\beta} \\ \ddot{\theta} \\ \ddot{\phi} \end{Bmatrix}$ The equations of motion can now be written in the matrix form:

$$[M]\{\ddot{z}\}+[C]\{\dot{z}\}+[K]\{z\}=\{f\}$$

The mass, damping, stiffness matrices and forcing vector for this form are given below.

$$[M]=\begin{bmatrix} m_b & & & & & & \\ & m_{uf} & & & & [0] & \\ & & m_{ur} & & & & \\ & & & I_{xx} & & & \\ & & & & I_{yy} & & \\ & [0] & & & & I_{xx\_a} & \\ & & & & & & I_{xx\_a} \end{bmatrix},$$

$$[C]=\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & c_{trf}+c_{tlf} & 0 & 0 & 0 & c(c_{trf}-c_{tlf}) & 0 \\ 0 & 0 & c_{trr}+c_{tlr} & 0 & 0 & 0 & c(c_{trr}-c_{tlr}) \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & c(c_{trf}-c_{tlf}) & 0 & 0 & 0 & c^2(c_{tlf}+c_{trf}) & 0 \\ 0 & 0 & c(c_{trr}+c_{tlr}) & 0 & 0 & 0 & c^2(c_{tlr}+c_{trr}) \end{bmatrix},$$

$$[K]=\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & k_{trf}+k_{tlf} & 0 & 0 & 0 & c(k_{trf}-k_{tlf}) & 0 \\ 0 & 0 & k_{trr}+k_{tlr} & 0 & 0 & 0 & c(k_{trr}-k_{tlr}) \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & c(k_{trf}-k_{tlf}) & 0 & 0 & 0 & c^2(k_{tlf}+k_{trf}) & 0 \\ 0 & 0 & c(k_{trr}+k_{tlr}) & 0 & 0 & 0 & c^2(k_{tlr}+k_{trr}) \end{bmatrix},$$

$$\{f\}=\begin{Bmatrix} m_b g - F_{ssr} - F_{slr} - F_{slf} - F_{srf} \\ k_{trf}(z_{rrf}-\delta_{stf})+c_{trf}\dot{z}_{rrf}+k_{tlf}(z_{rlf}-\delta_{stf})+c_{tlf}\dot{z}_{rlf}+m_{uf}g+F_{srf}+F_{slf} \\ k_{trr}(z_{rrr}-\delta_{str})+c_{trr}\dot{z}_{rrr}+k_{tlr}(z_{rlr}-\delta_{str})+c_{tlr}\dot{z}_{rlr}+m_{ur}g+F_{srr}+F_{slr} \\ (F_{slr}+F_{slf}-F_{srr}-F_{srf})d \\ (F_{slf}+F_{srf})b-(F_{srr}+F_{slr})a \\ c[k_{tlf}(\delta_{stf}-z_{rlf})-c_{tlf}\dot{z}_{rlf}-k_{trf}(\delta_{stf}-z_{rrf})+c_{trf}\dot{z}_{rrf}]+d(F_{srf}-F_{slf}) \\ c[k_{tlr}(\delta_{str}-z_{rlr})-c_{tlr}\dot{z}_{rlr}-k_{trr}(\delta_{str}-z_{rrr})+c_{trr}\dot{z}_{rrr}]+d(F_{srr}-F_{slr}) \end{Bmatrix}$$

The system of seven second order ODEs can be decomposed into a system of fourteen first order ODEs. This is applied by rewriting the matrix equation given above in state variable form.

$$\{\ddot{z}\}=-[M]^{-1}[C]\{\dot{z}\}-[M]^{-1}[K]\{z\}+[M]^{-1}\{f\}$$

Now, define $\{Z\}=\begin{Bmatrix}\{z\}\\ \{\dot{z}\}\end{Bmatrix}$, with time derivative $\{\dot{Z}\}=\begin{Bmatrix}\{\dot{z}\}\\ \{\ddot{z}\}\end{Bmatrix}$ and apply the identity $\{\dot{z}\}=\{\dot{z}\}$. This yields the system of first order ODEs:

$$\{\dot{Z}\}=\begin{Bmatrix}\{\dot{z}\}\\ \{\ddot{z}\}\end{Bmatrix}=\begin{bmatrix}[0] & [I]\\ -[M]^{-1}[K] & -[M]^{-1}[C]\end{bmatrix}\begin{Bmatrix}\{z\}\\ \{\dot{z}\}\end{Bmatrix}+\begin{Bmatrix}\{0\}\\ [M]^{-1}\{f\}\end{Bmatrix}$$

This system of equations may now be solved using a numerical integration scheme.

B. Appendix B

Hydro-pneumatic spring stiffness properties

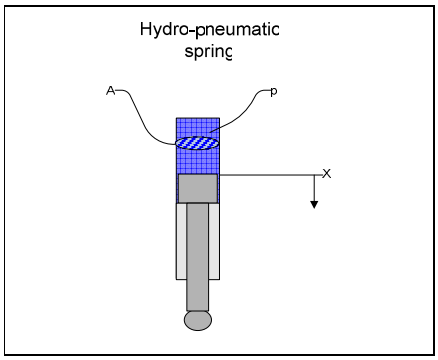

Figure B-1: Sketch of hydro-pneumatic spring

The compression of a hydro-pneumatic spring can be described by a polytropic process (Sonntag et al., 2003):

$$p_1 V_1^{n_p} = p_2 V_2^{n_p}$$

$$\therefore p_2 = p_1 \left(\frac{V_1}{V_2} \right)^{n_p}$$

where:

p	=	internal pressure,
V	=	volume, and
n_p	=	the polytropic gas constant.

The polytropic gas constant given above is highly dependent on the rate of compression or expansion.

The subscript 1 indicates the initial condition, whereas the subscript 2 indicates the final state.

The area is assumed to remain constant (i.e. independent of pressure and displacement). The force exerted by the hydro-pneumatic spring is given by:

$$F = p_1 A \left(\frac{V_1}{V_2} \right)^{n_p}$$

The initial condition can now be taken to correspond to the static position. The initial pressure and volumes thus correspond to the static pressure and volume:

$$F = p_{stat} A \left(\frac{V_{stat}}{V_2} \right)^{n_p}$$

The pneumatic volume of the spring is given by:

$$V = Ax,$$

where x is the displacement of the piston.

The force in the spring thus becomes:

$$F = p_{stat} A \left(\frac{x_{stat}}{x} \right)^{n_p}$$

It can be seen from the equation above that an air spring has a non-linear force characteristic, which is dependent on static pressure, static volume, area, the polytropic constant and the static displacement. We have however already assumed that in the case of our air spring the area and static volume will remain constant. It is further assumed that the polytropic constant will remain unchanged for a given operating condition.

Now, consider that linearised stiffness of a spring element about a point, x_0, is given by: (reference)

$$k = \left. \frac{\partial F}{\partial x} \right|_{x_0}$$

The stiffness of the hydro-pneumatic spring being considered thus becomes:

$$k(x) = -np_{stat} A \left(\frac{x_{stat}}{x} \right)^{n_p - 1} \left(\frac{x_{stat}}{x^2} \right)$$

$$\therefore k(x) = -np_{stat} A \left(\frac{x_{stat}^{n_p}}{x^{n_p + 1}} \right)$$

Since the polytropic constant, the static pressure, the area and the static position have been thus far considered to remain unchanged, the stiffness of the spring is a function of displacement, x, only. The linearised stiffness about the static position reduces to:

$$k_{stat} = -n_p p_{stat} A \left(\frac{1}{x_{stat}} \right)$$

In the case of a load-levelling hydro-pneumatic suspension system, the pressure in the system is increased proportionally to the added payload in order to maintain the same static position.

$$m_b \propto p_{stat}$$

m_b in this case is the vehicle sprung mass.

Let us now examine the change in the stiffness about the static position due to a change in static pressure:

$$\frac{dk_{stat}}{dp_{stat}} = -n_p A\left(\frac{1}{x_{stat}}\right)$$

The result is a constant. This result implies that:

$$k_{stat} \propto p_{stat}$$

Consequently,

$$k_{stat} \propto m_b \text{ and } \therefore \frac{k_{stat}}{m_b} = \text{constant}$$

Take into consideration that the natural frequency of a linear spring-mass system is given by:

$$\omega = \sqrt{\frac{k_{stat}}{m_b}}$$

It is therefore reasonable to assume that the body natural frequency of vehicle with load-levelling suspension as described above is independent of vehicle payload, barring small parameter variations.

The stiffness of a cylindrical column of fluid

The definition of the bulk modulus of a fluid is given by White (2003) as:

$$\beta_f = -V\frac{\partial p}{\partial V}$$

where:

β_f	=	the bulk modulus of the fluid,
V	=	the volume of the fluid column,
∂p	=	an infinitesimal change in the pressure on the fluid,
∂V	=	an infinitesimal change in the volume of the fluid column.

Assuming that the fluid is constrained, such that the cylindrical fluid column's cross sectional area remains constant, the infinitesimal changes in pressure and volume can be expressed as:

$$\partial p = \frac{\partial F}{A}$$
$$\partial V = A\partial x$$

The symbol A here represents the circular cross sectional area of the fluid column, while ∂F and ∂x represent infinitesimal changes in force and displacement respectively. The bulk modulus of the oil can now be expressed as:

$$\beta_f = -V\frac{\partial F}{A^2\partial x}$$

Rearranging this expression gives the stiffness of the fluid column:

$$k_{\beta_f} = \frac{\partial F}{\partial x} = -\frac{\beta_f A^2}{V}$$

The bulk modulus of the oil in the $4S_4$ was measured as $\beta_f = 1.368$ GPa (Els, 2007). The diameter of the oil column in the $4S_4$ is 50 mm. The oil volume in the $4S_4$ during the determination of the bulk modulus of the oil was 1.6 ℓ or 0.0016 m^3. With all these quantities known the stiffness of the oil in the $4S_4$ can be calculated as:

$$k_{\beta_f} = -\frac{\left(1.368\times10^9\right)\left[\left(\pi/4\right)0.05^2\right]^2}{0.0016}$$
$$k_{\beta_f} = -3.296MN/m$$

C. Appendix C

Model validation results

The full model correlation results for both discrete obstacles and random terrain are included here. The list of tests performed from Chapter 2 is repeated here.

Table C-1: List of tests performed.

Run no.	Test description	Suspension setting	Speed [km/h]
01, 05, 10	Belgian paving	ride comfort	±14.5 (low range, 1st gear)
02, 06, 11	Belgian paving	ride comfort	±54 (low range, 4th gear)
03, 08, 12	Belgian paving	handling	±14.5 (low range, 1st gear)
04, 09, 13	Belgian paving	handling	±54 (low range, 4th gear)
16, 19, 20, 21, 22	Bump course, layout 1	ride comfort	±14.5 (low range, 1st gear)
23, 24, 25, 26, 27	Bump course, layout 2	handling	±14.5 (low range, 1st gear)
28, 30, 32, 33	Bump course, layout 3	ride comfort	±14.5 (low range, 1st gear)
29, 31, 34	Bump course, layout 3	handling	±14.5 (low range, 1st gear)

The correlation for tests over the discrete obstacle layout 1 is given below. The results are presented for run 16 in the time domain. The run was conducted at 14.5 km/h with the $4S_4$ set to ride comfort mode. The body vertical accelerations, pitch- and roll-velocities, suspension displacements, suspension forces and stresses are presented in Figures C-1 to Figure C-5.

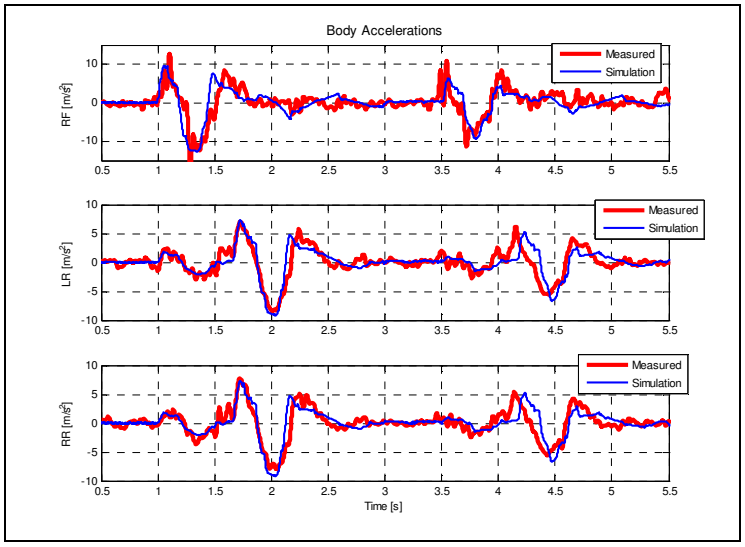

Figure C-1: Body vertical acceleration over trapezoidal bump layout 1 for run 16.

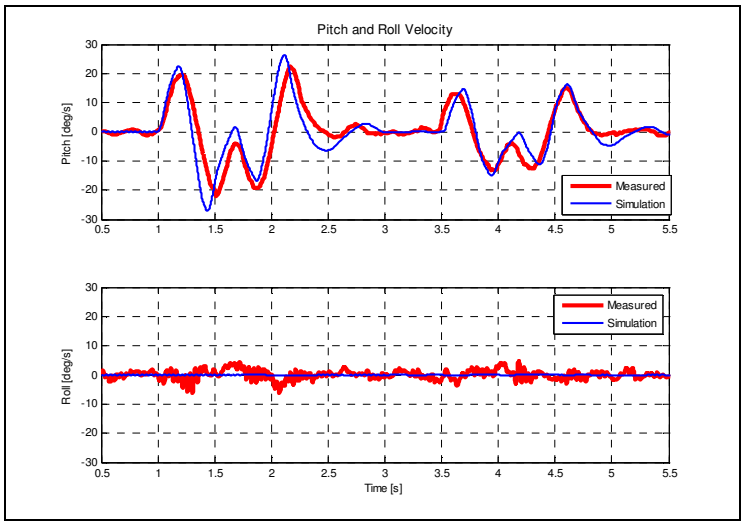

Figure C-2: Body pitch- and roll-velocities over trapezoidal bump layout 1 for run 16.

The suspension displacements for run 16 are given in Figure C-3. The suspension displacements follow the correct trends for both front and rear. The simulated peak rebound displacements of the rear suspension are however about 30 mm lower than that measured results. This is not seen as cause for concern, when the prediction of suspension forces is considered. The non-linear characteristic of the spring reduces in spring rate as the suspension extends and the discrepancy in displacement thus only causes a small error in the force prediction.

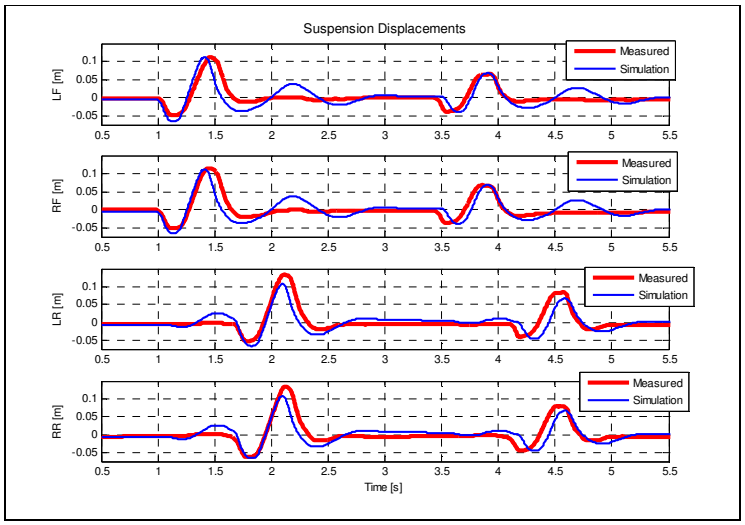

Figure C-3: Suspension displacements over the trapezoidal bump layout 1 for run 16.

The simulation and measured results for the suspension forces as well as the stress in the left rear suspension mounting are plotted in Figure C-4 and Figure C-5. The suspension forces calculated from the measured pressures are plotted in green solid lines, while the left rear suspension force measured by the load cell is plotted in red.

The simulated data should be compared to the measured load cell force. The offset between the simulated and the measured load cell force data can be attributed to the statically indeterminate nature of the vehicle, which causes unequal distribution of the static forces in the four corners of the vehicle suspension.

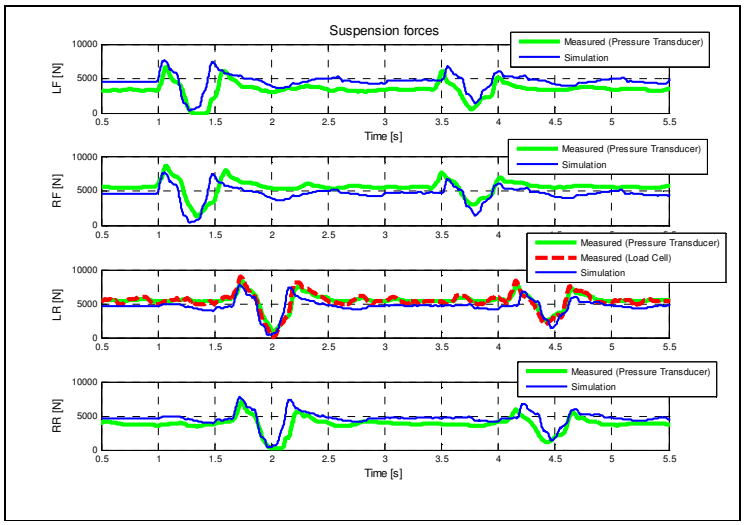

Figure C-4: Suspension forces over the trapezoidal bump layout 1 for run 16.

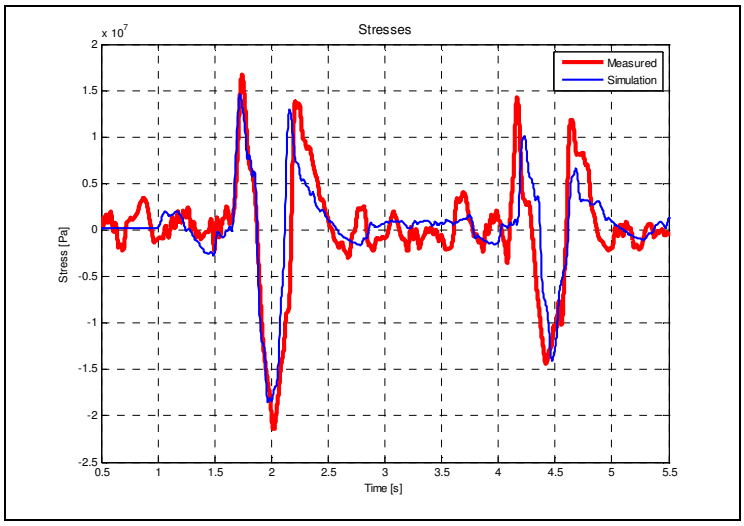

Figure C-5: Stresses over the trapezoidal bump layout 1 for run 16.

The correlation for tests over the discrete obstacle layout 2 is given below. The results are presented for run 24 in the time domain. The run was conducted at 14.5 km/h with the $4S_4$ set to handling mode. The body vertical accelerations, pitch- and roll-velocities, suspension displacements, suspension forces and stresses are presented.

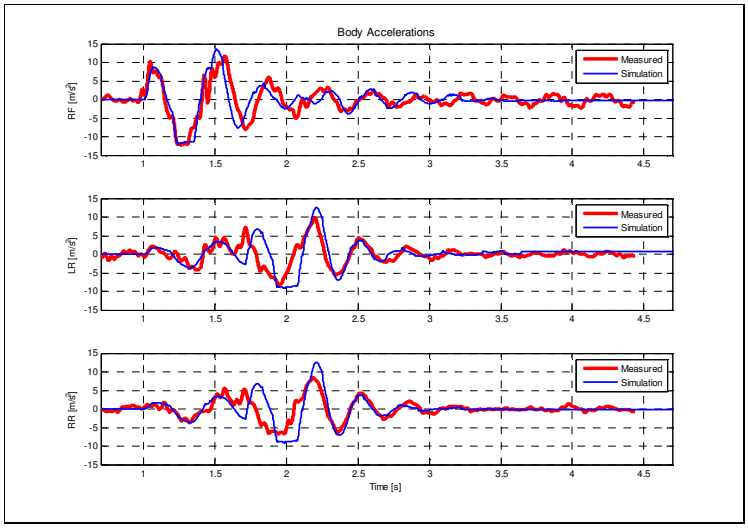

Figure C-6: Body vertical accelerations over trapezoidal bump layout 2 for run 24.

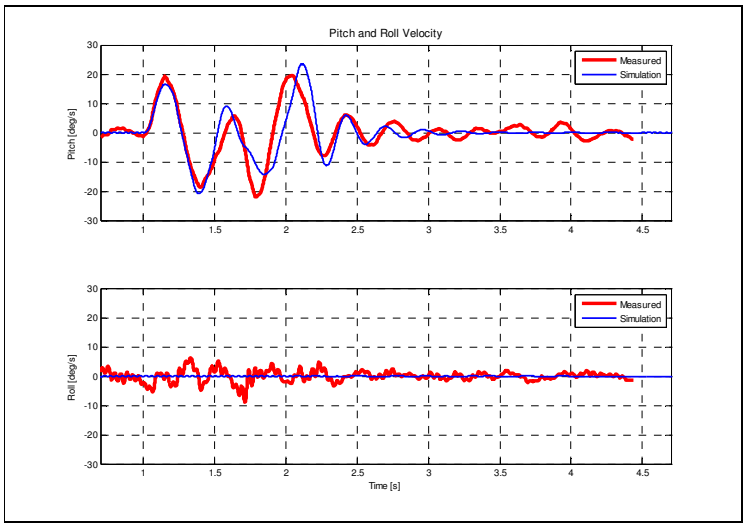

Figure C-7: Body pitch- and roll-velocities over trapezoidal bump layout 2 for run 24

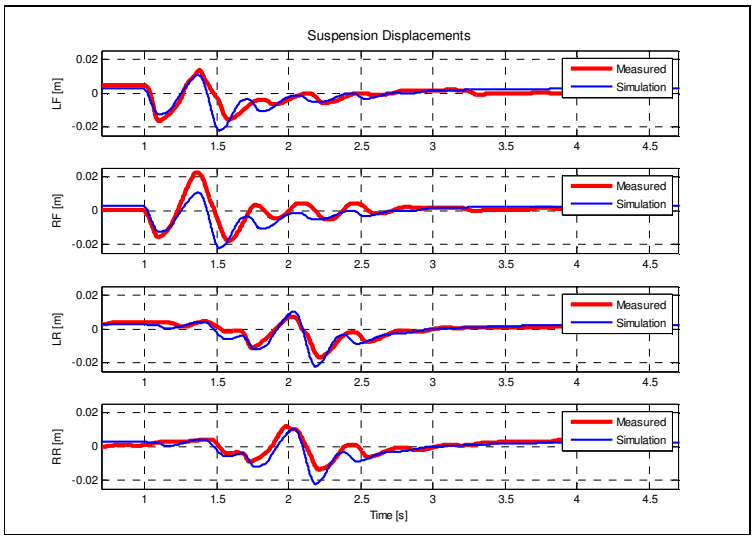

Figure C-8: Suspension displacements over the trapezoidal bump layout 2 for run 24

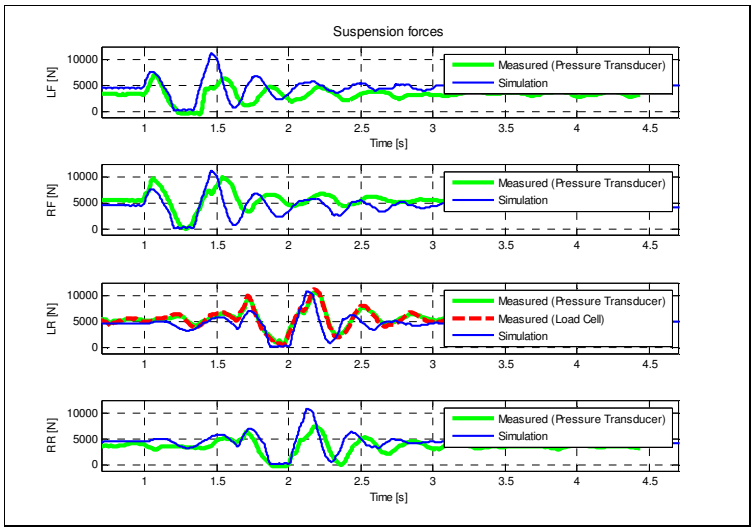

Figure C-9: Suspension forces over the trapezoidal bump layout 2 for run 24

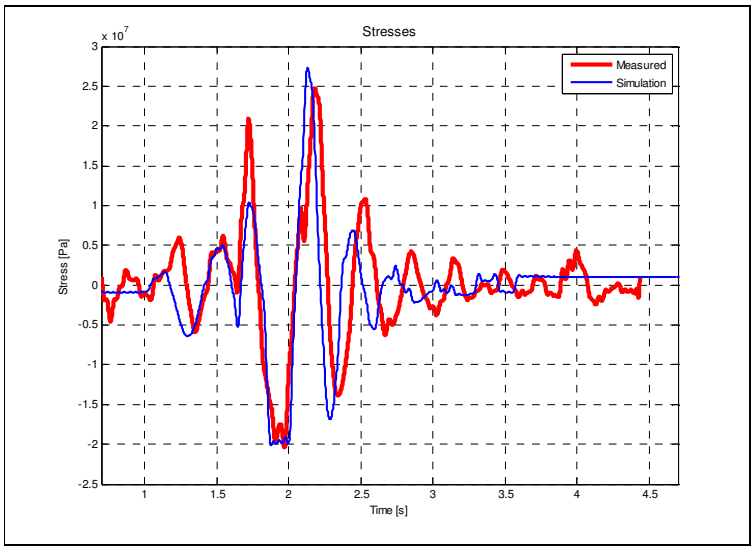

Figure C-10: Stresses over the trapezoidal bump layout 2 for run 24

The correlation for tests over the discrete obstacle layout 3 is given below. The results are presented for run 30 in the time domain. The run was conducted at 14.5 km/h with the $4S_4$ set to ride comfort mode. The body vertical accelerations, pitch- and roll-velocities, suspension displacements, suspension forces and stresses are presented.

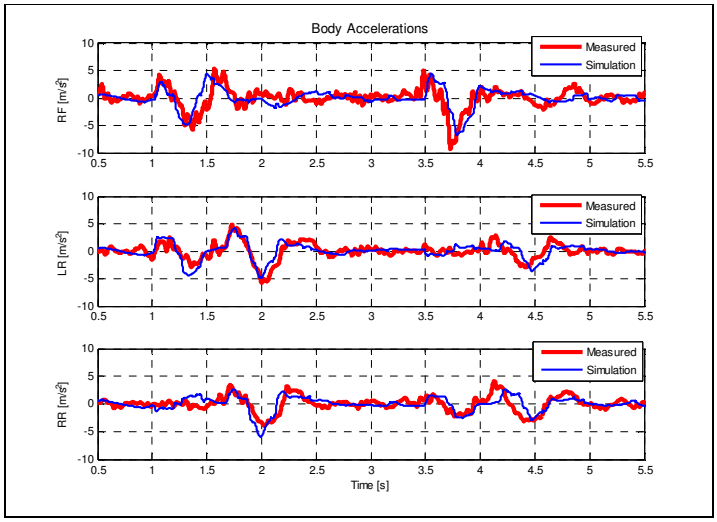

Figure C-11: Body vertical accelerations over trapezoidal bump layout 3 for run 30

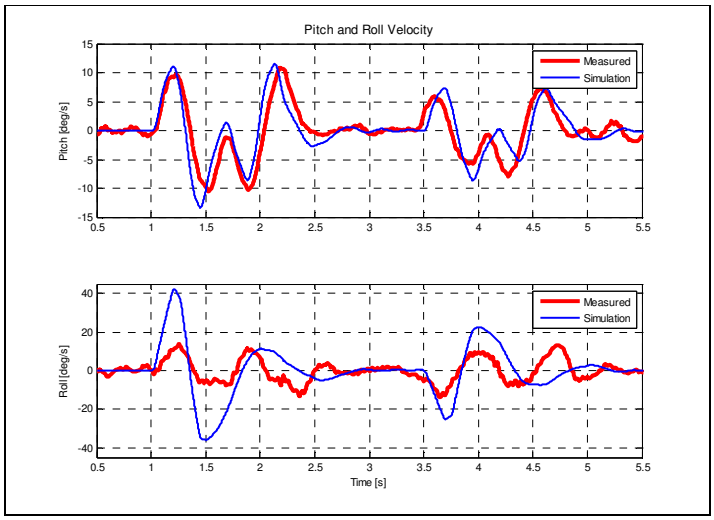

Figure C-12: Body pitch- and roll-velocities over trapezoidal bump layout 3 for run 30

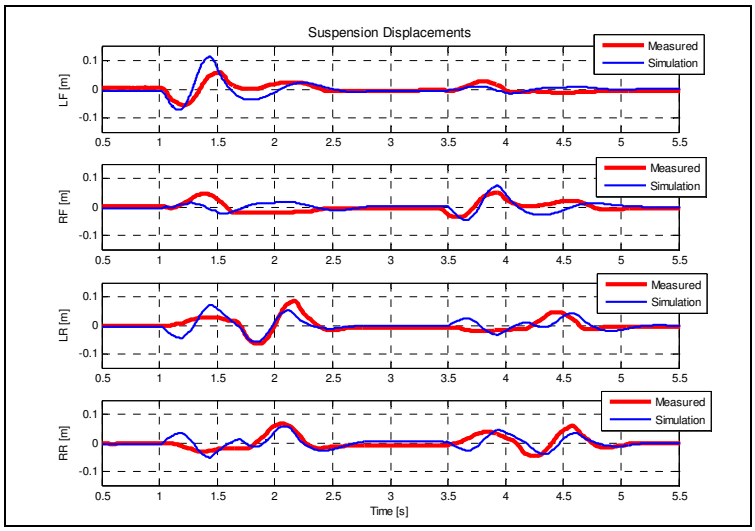

Figure C-13: Suspension displacements over trapezoidal bump layout 3 for run 30.

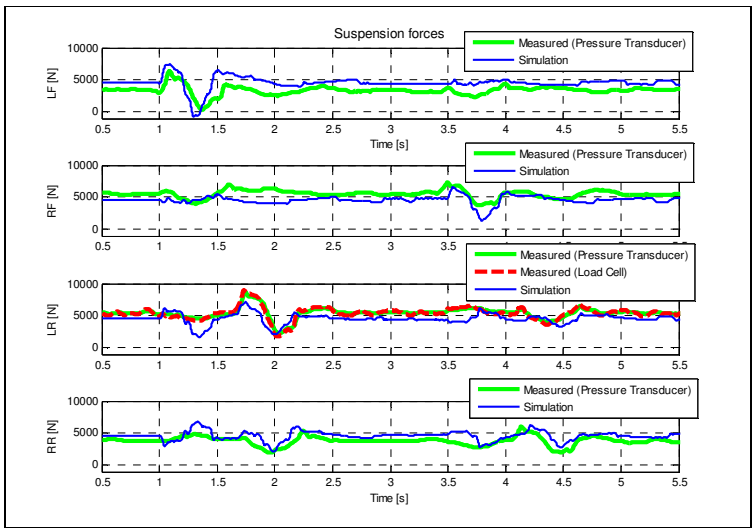

Figure C-14: Suspension forces over trapezoidal bump layout 3 for run 30.

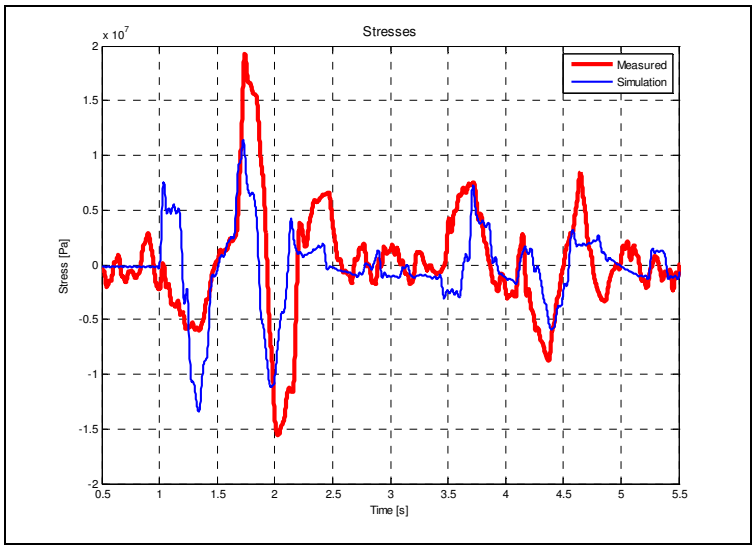

Figure C-15: Stresses over trapezoidal bump layout 3 for run 30.

The correlation for tests over the discrete obstacle layout 3 is given below. The results are presented for run 31 in the time domain. The run was conducted at 14.5 km/h with the $4S_4$ set to handling mode. The body vertical accelerations, pitch- and roll-velocities, suspension displacements, suspension forces and stresses are presented.

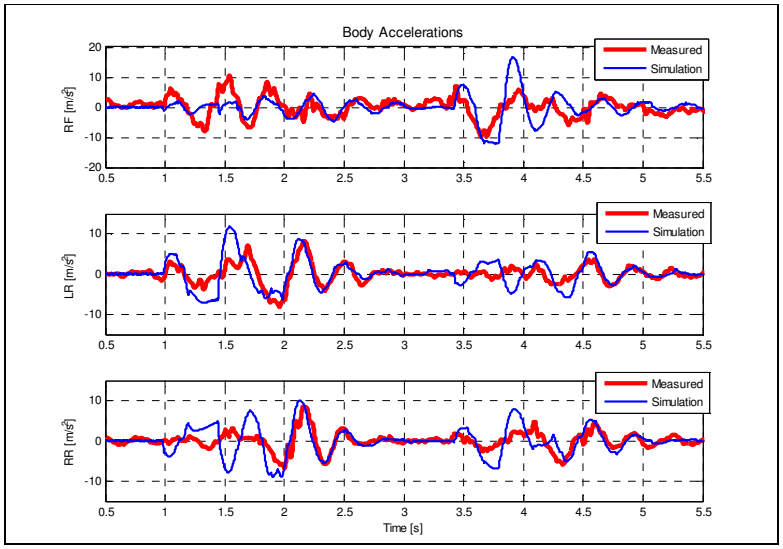

Figure C-16: Body vertical accelerations over trapezoidal bump layout 3 for run 31.

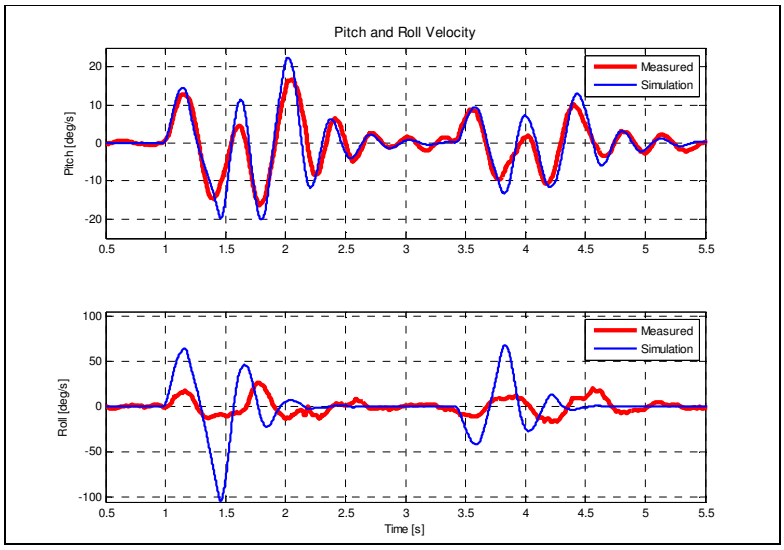

Figure C-17: Body pitch- and roll-velocities over trapezoidal bump layout 3 for run 31.

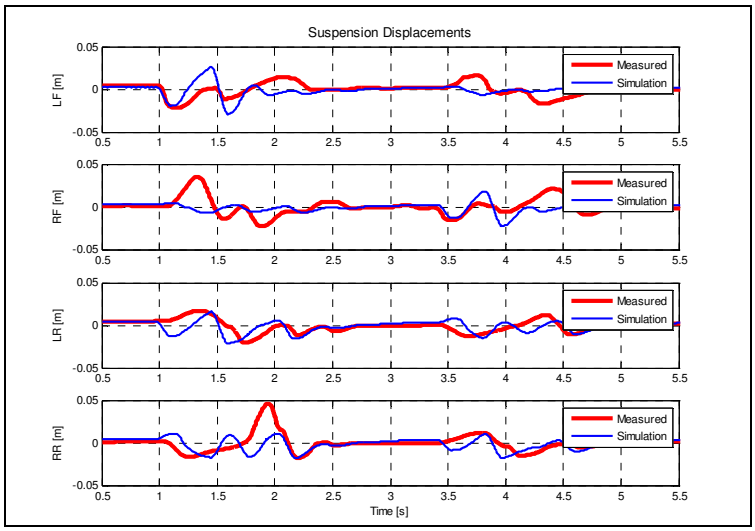

Figure C-18: Suspension displacements over trapezoidal bump layout 3 for run 31.

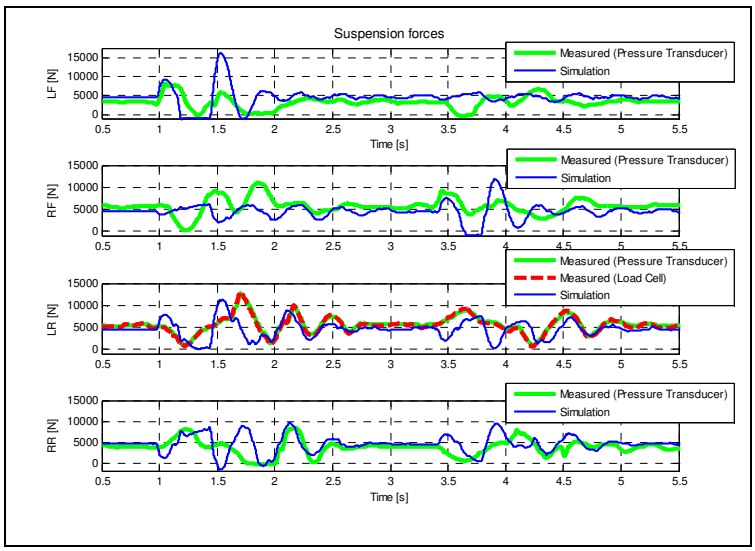

Figure C-19: Suspension forces over trapezoidal bump layout 3 for run 31.

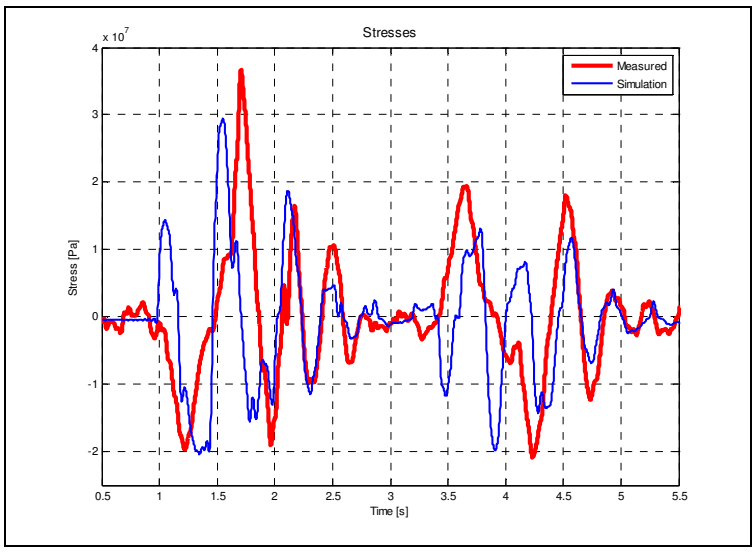

Figure C-20: Stresses over trapezoidal bump layout 3 for run 31.

The correlation for tests over the Belgian paving is given below. The results are presented in the frequency domain for body vertical accelerations, pitch- and roll-velocities, suspension displacements, suspension forces and stresses.

The results for run 01 are presented in the following figures. The run was conducted at 14.5 km/h over the Belgian paving with the $4S_4$ set to the ride comfort mode.

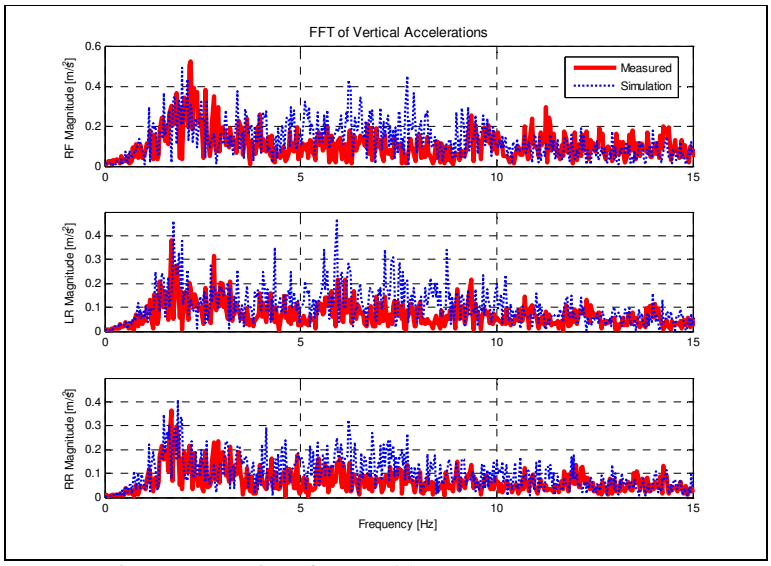

Figure C-21: Body vertical accelerations for run 01.

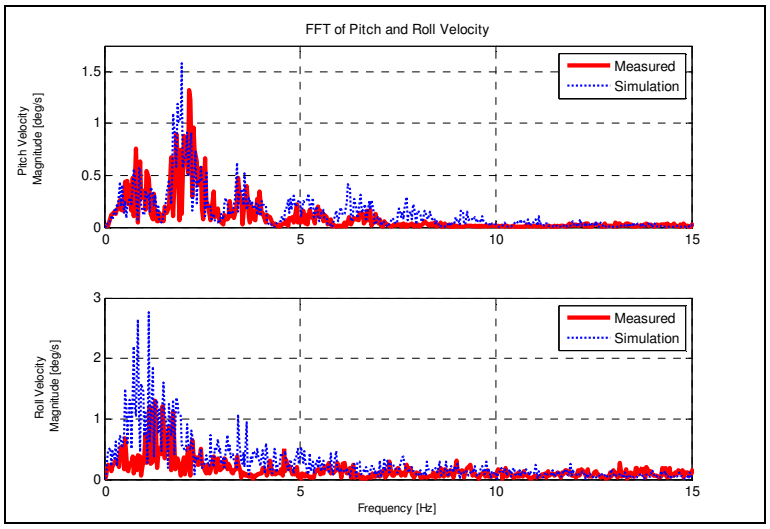

Figure C-22: FFTs of body pitch- and roll-velocities for run 01.

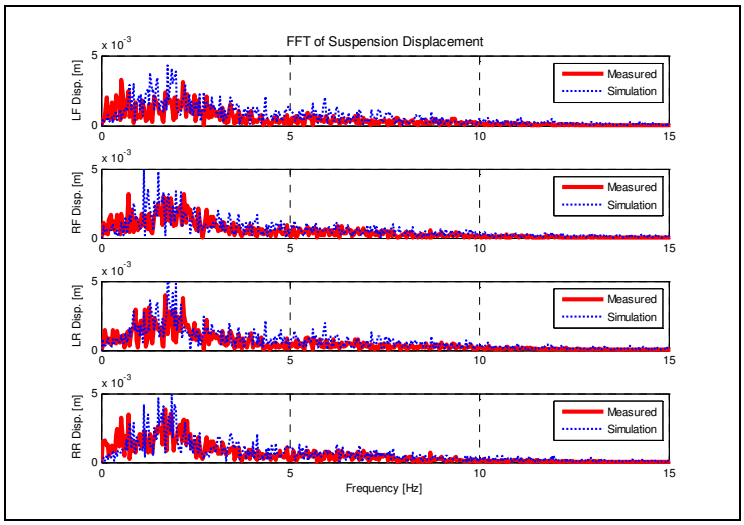

Figure C-23: FFTs of suspension displacements for run 01.

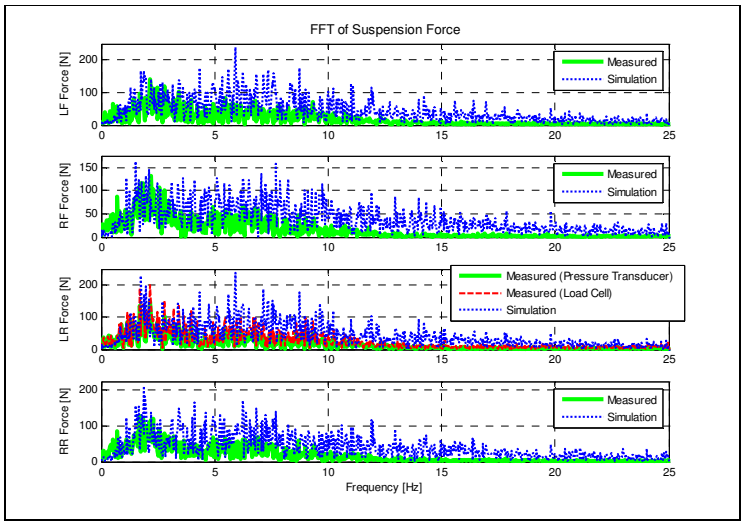

Figure C-24: FFTs of suspension forces for run 01.

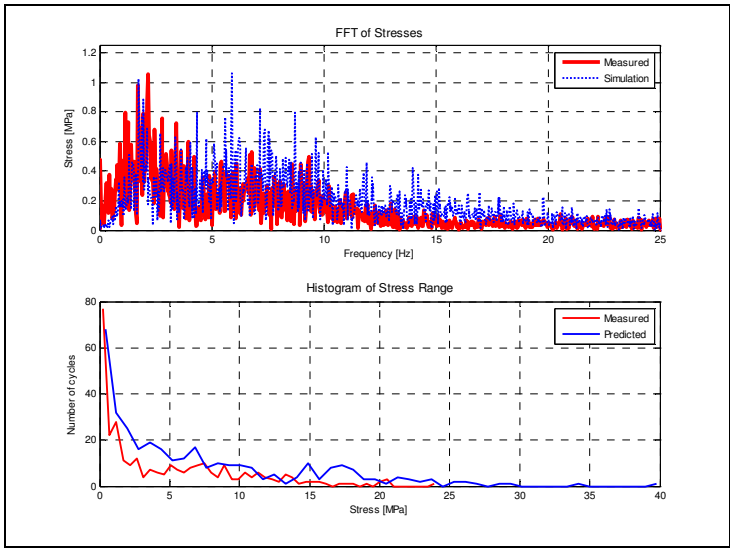

Figure C-25: FFT of stress and histogram of stress ranges for run 01.

The results for run 02 are presented in the following figures. The run was conducted at 55 km/h over the Belgian paving with the $4S_4$ set to the ride comfort mode. Note that the right front accelerations channel was faulty during runs 02, 03 and 04. The data for this accelerometer will therefore not be presented.

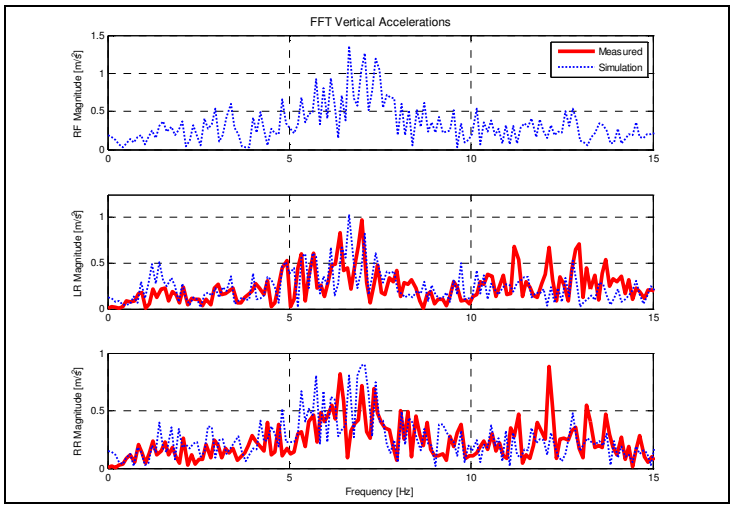

Figure C-26: FFT of vertical accelerations for run 02.

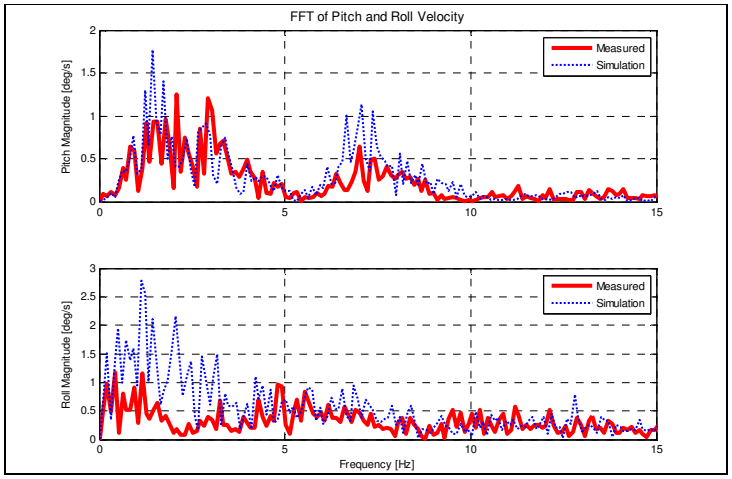

Figure C-27: FFTs of pitch- and roll-velocity for run 02.

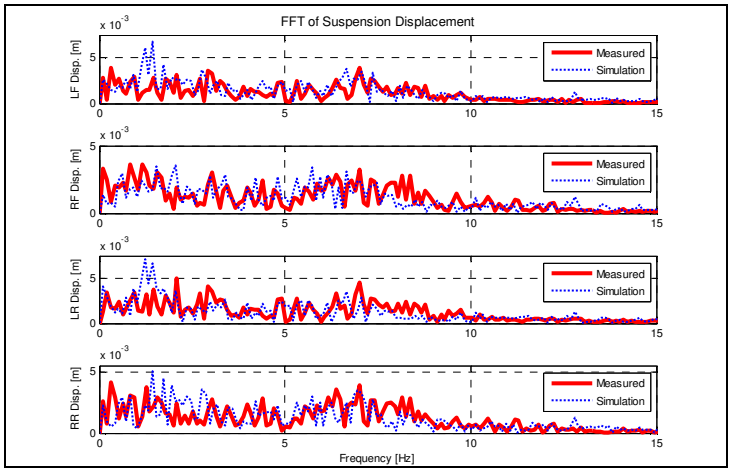

Figure C-28: FFTs of suspension displacement for run 02.

The correlation for the simulation suspension force FFTs with the measured data is excellent as can be seen in the plots of Figure C-29. The FFT of the simulation data follows the trend of the measured FFT exceptionally well and indeed also greatly captures the correct magnitudes. The reader should note that most of the energy is contained between 5 and 10 Hz.

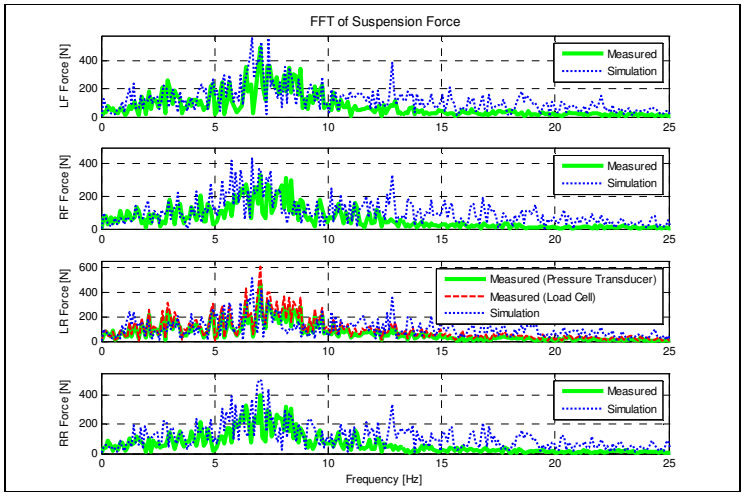

Figure C-29: FFTs of suspension force for run 02.

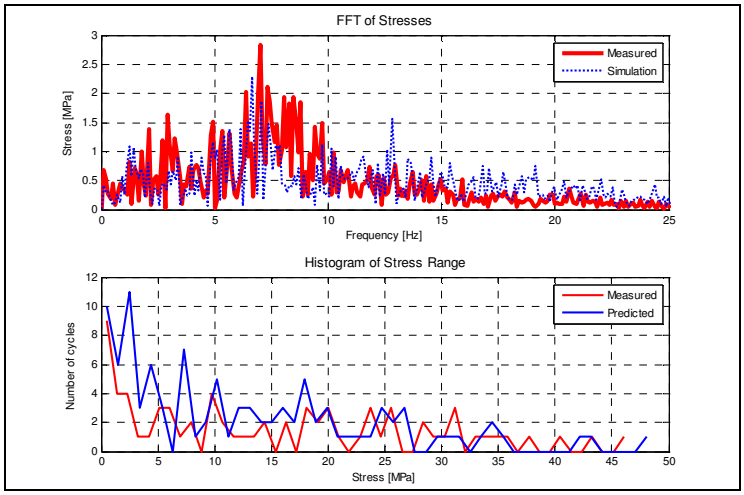

Figure C-30: FFT of stress and histogram of stress ranges run 02.

The results for run 03 are presented in the following figures. The run was conducted at 14.5 km/h over the Belgian paving with the $4S_4$ set to the handling mode.

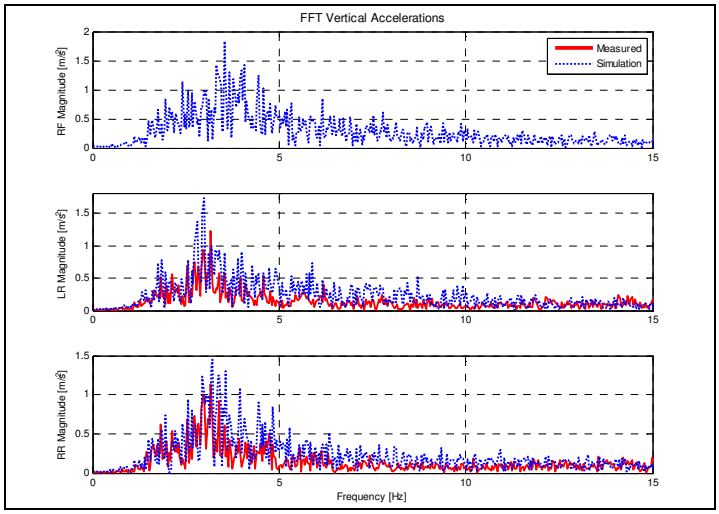

Figure C-31: FFTs of body vertical accelerations for run 03.

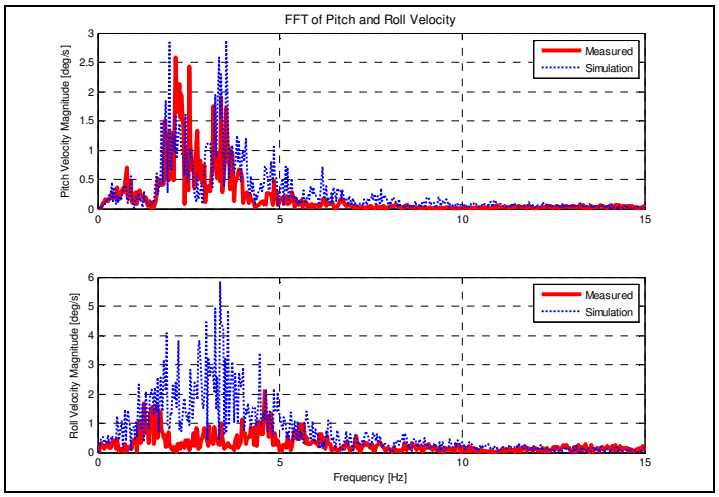

Figure C-32: FFTs of body pitch- and roll-velocity for run 03.

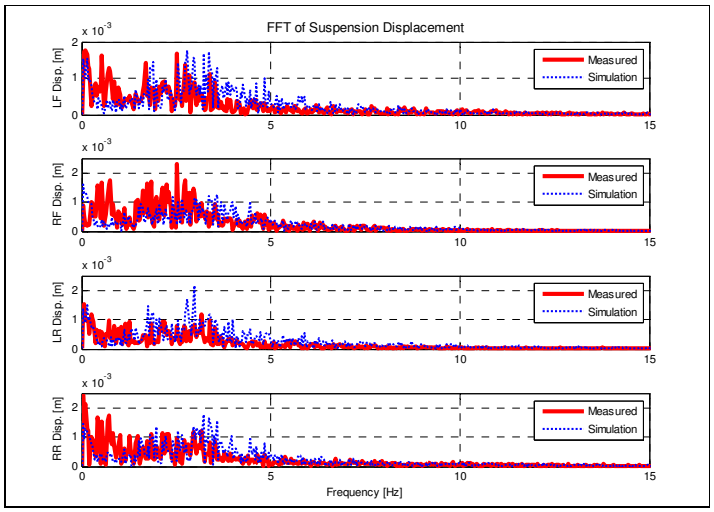

Figure C-33: FFTs of suspension displacements for run 03.

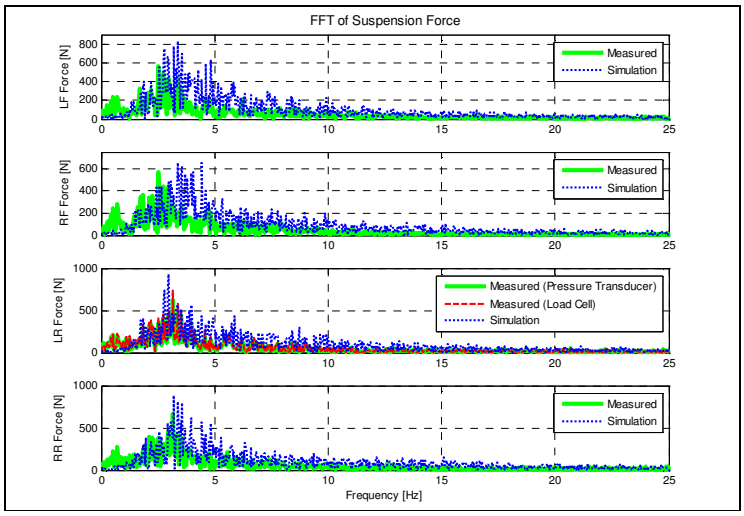

Figure C-34: FFTs of suspension force for run 03.

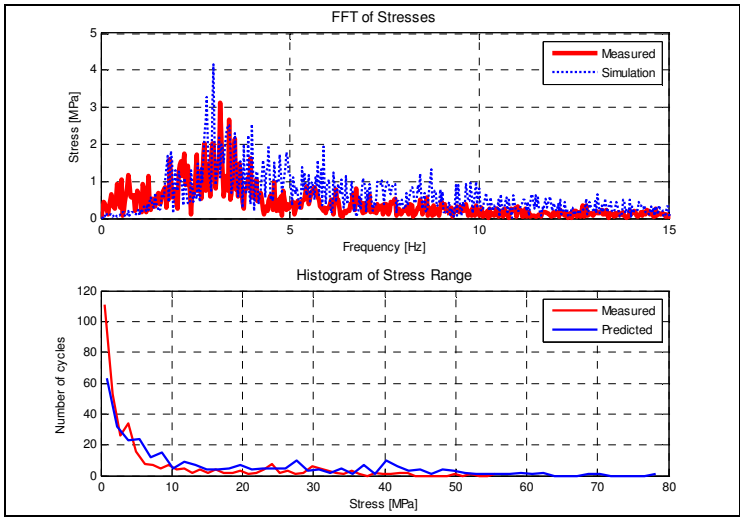

Figure C-35: FFT of stress and histogram of stress ranges for run 03.

The results for run 04 are presented in the following figures. The run was conducted at 54 km/h over the Belgian paving with the $4S_4$ set to the handling mode.

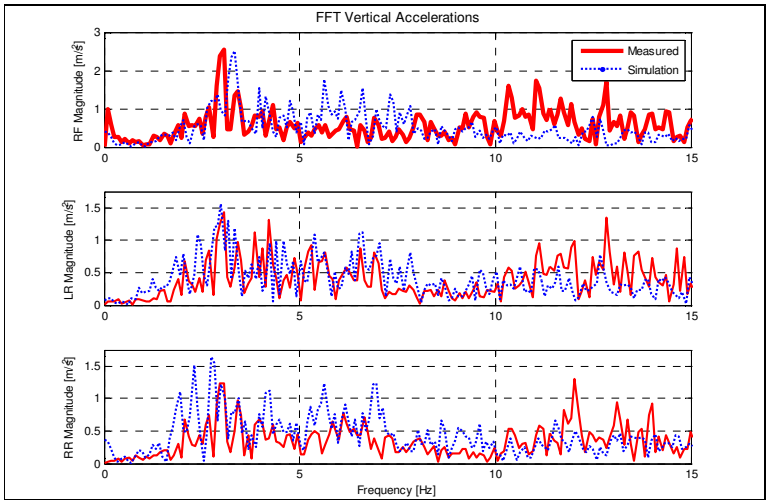

Figure C-36: FFTs of body vertical accelerations for run 04.

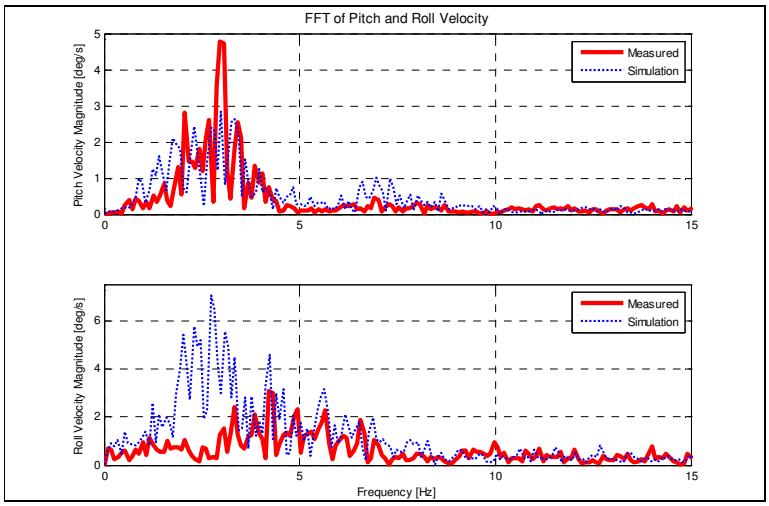

Figure C-37: FFTs of pitch- and roll-velocity for run 04.

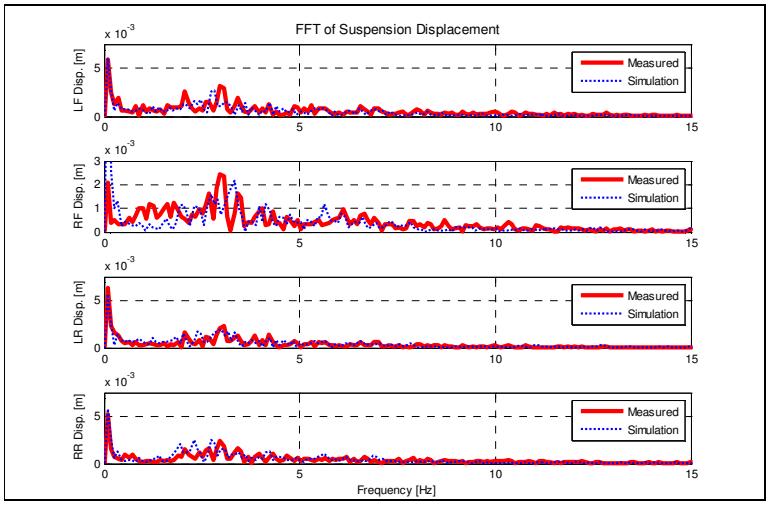

Figure C-38: FFTs of suspension displacement for run 04.

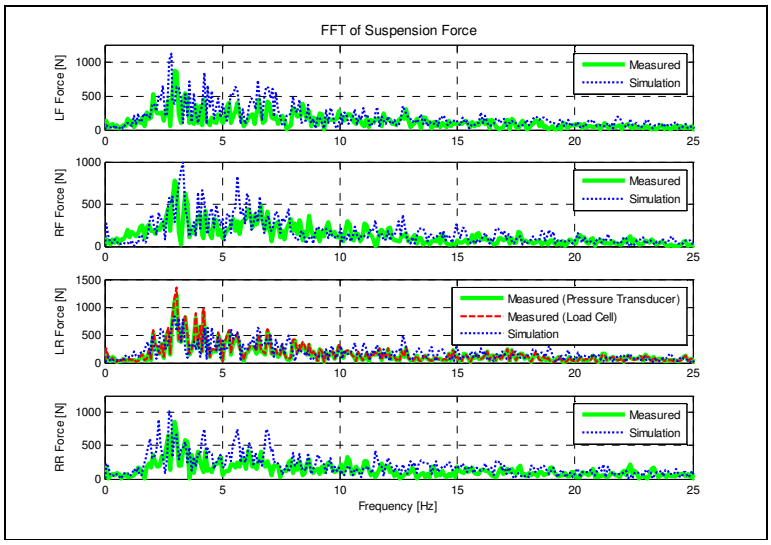

Figure C-39: FFTs of supsension forces for run 04.

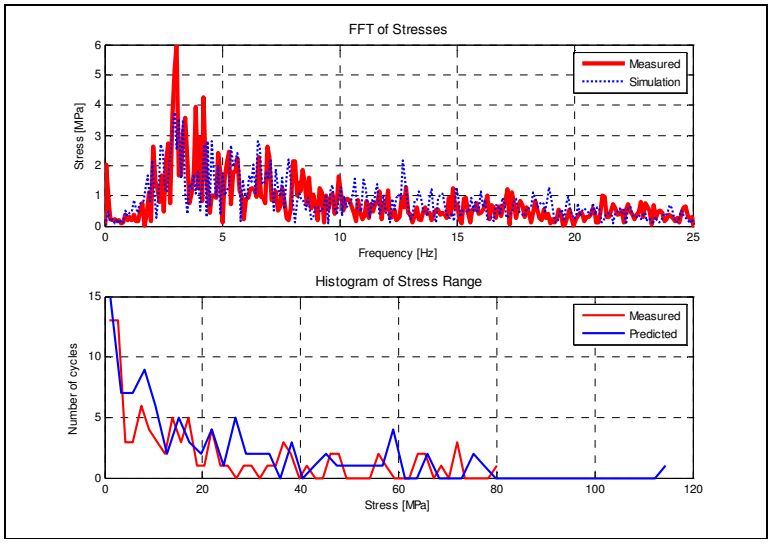

Figure C-40: FFT of stress and histogram of stress ranges for run 04.

D. Appendix D

Convergence histories to the optima for the seven-variable optimisation

The convergence history to the optimum for the unladen case is shown below. The starting point employed was $sgv = 0.2\ell$, $dsf_{1\to6} = 3.2$.

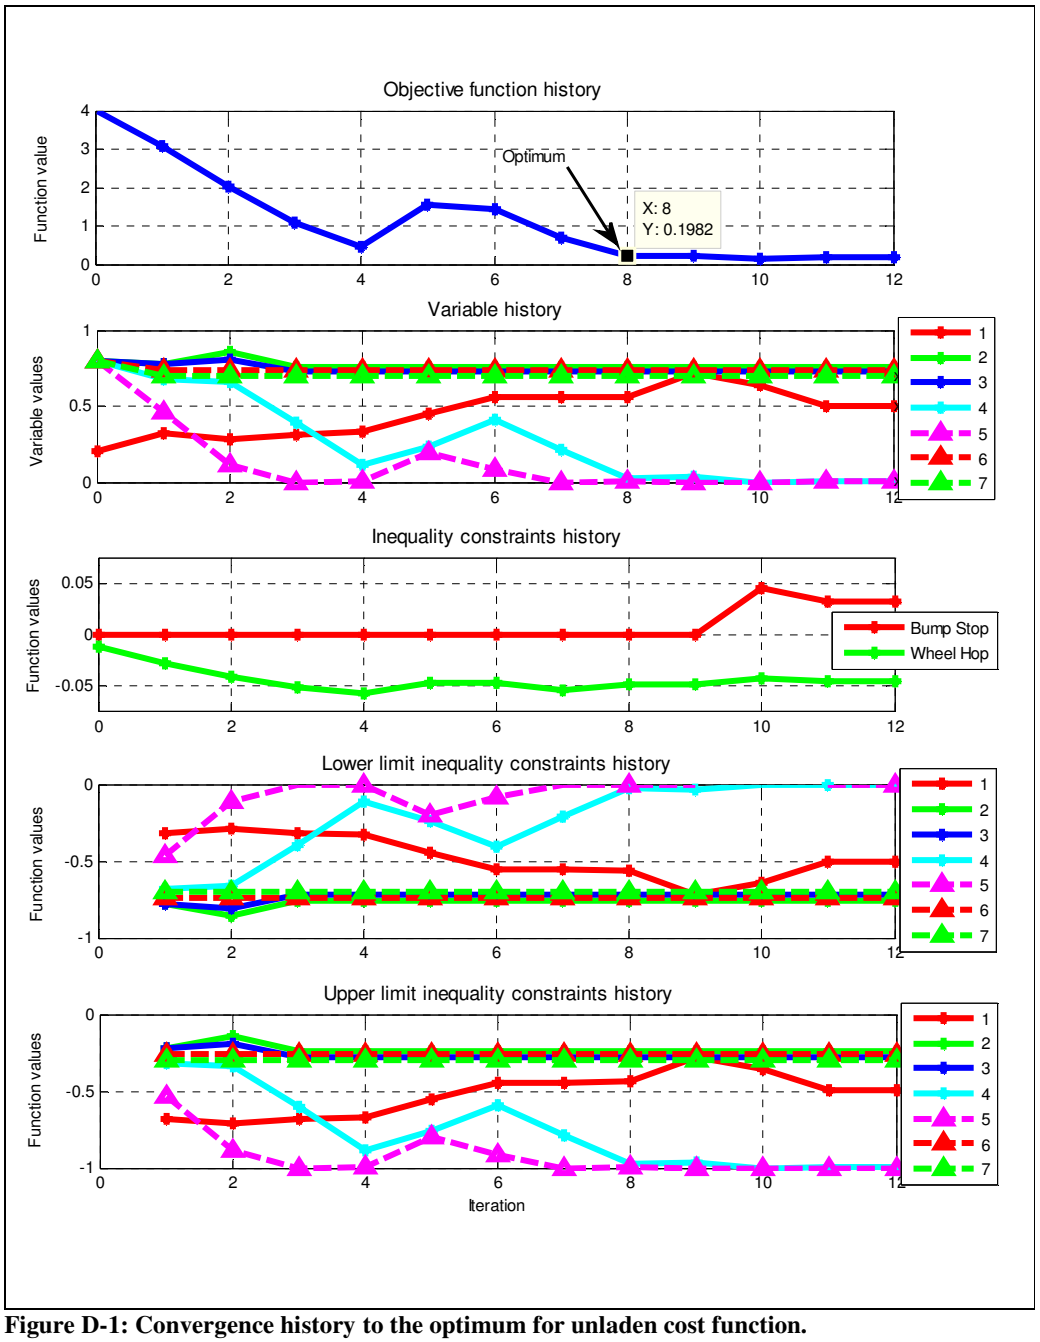

Figure D-1: Convergence history to the optimum for unladen cost function.

The convergence history to the optimum for the fully laden case is shown below. The starting point employed was $sgv = 0.21\ell$, $dsf_{1\to 6} = 0.88$.

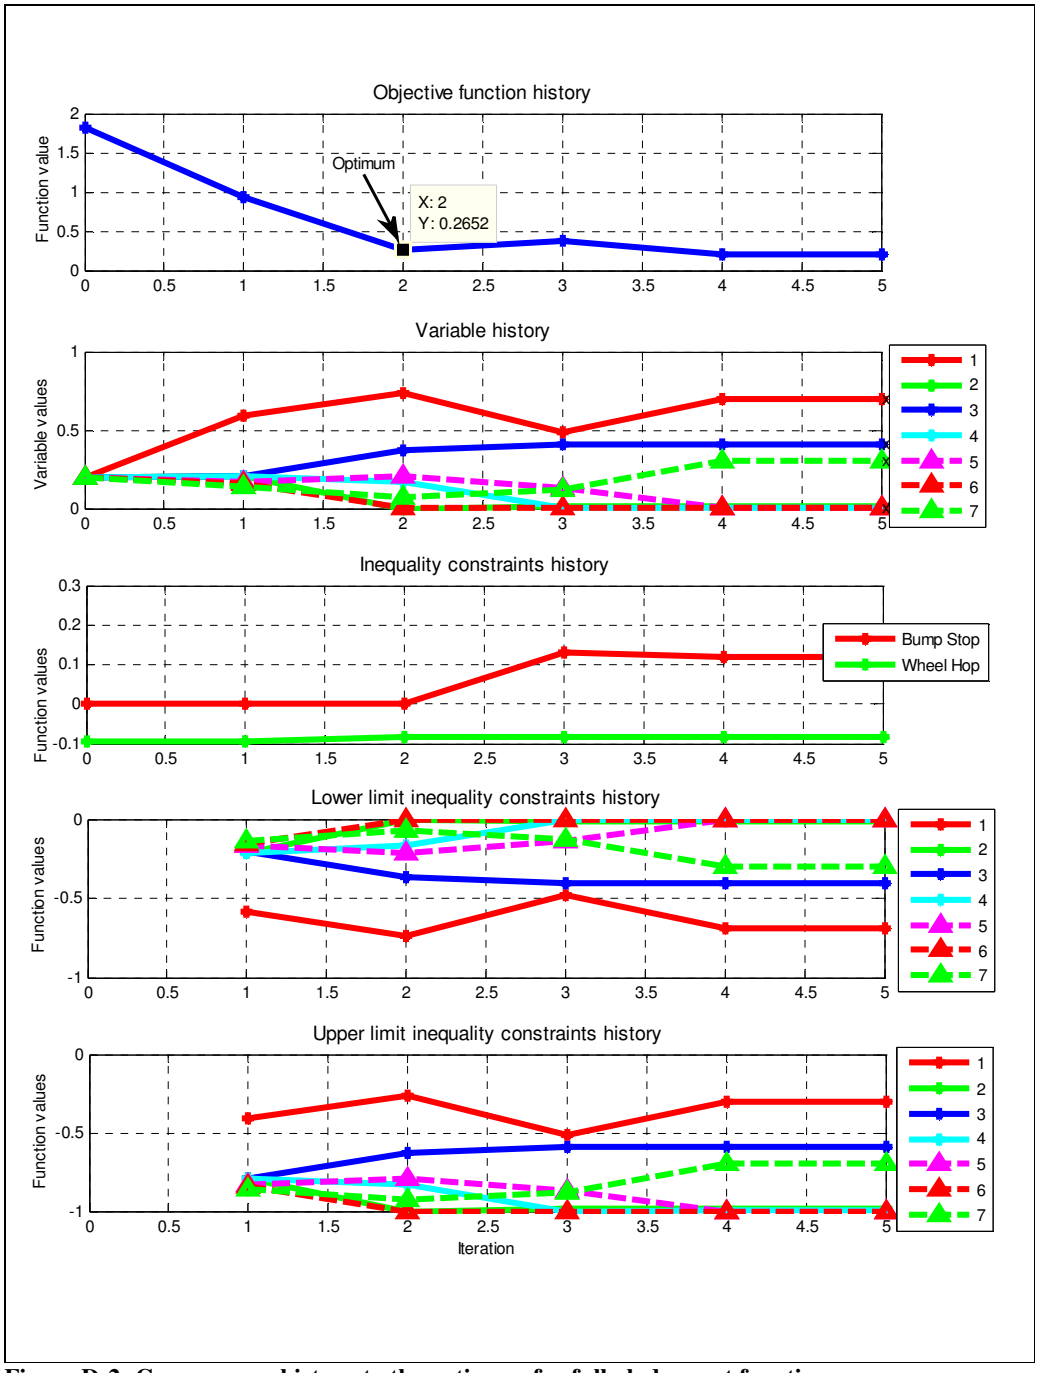

Figure D-2: Convergence history to the optimum for fully laden cost function.

The convergence history to the optimum for the heavy load case is shown below. The starting point employed was $sgv = 0.21\ell$, $dsf_{1\to 6} = 3.2$.

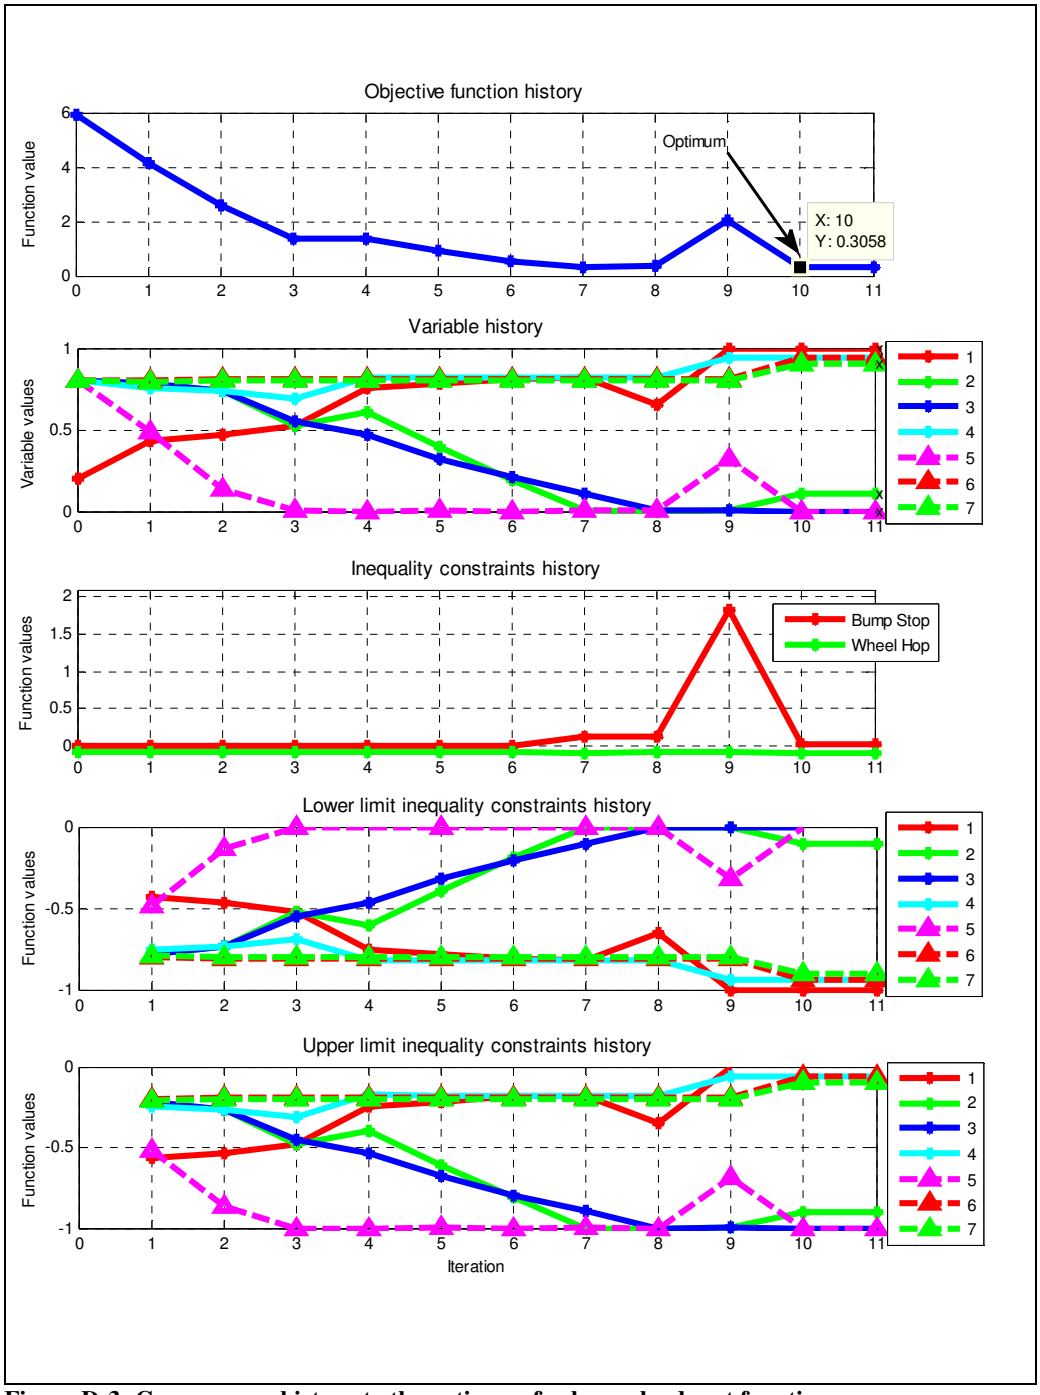

Figure D-3: Convergence history to the optimum for heavy load cost function.

The convergence history to the optimum for the extreme load case is shown below. The starting point employed was $sgv = 0.69\ell$, $dsf_{1\to 6} = 3.2$.

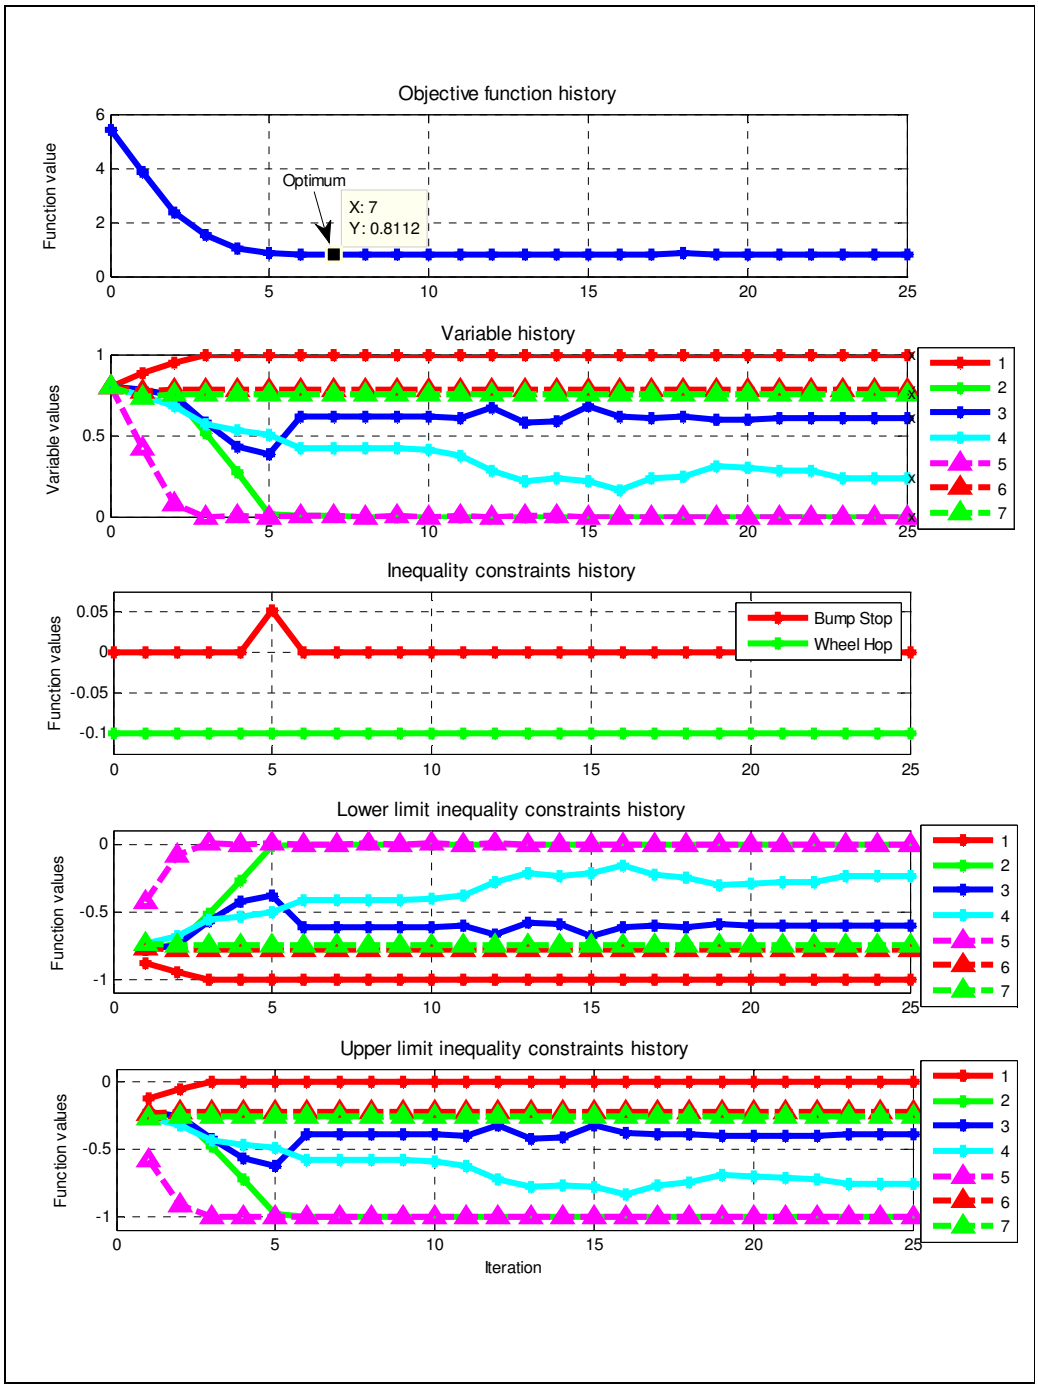

Figure D-4: Convergence history to the optimum for extreme load cost function.