

The Use of Simulation to Expedite Experimental Investigations of the Effect of High-Performance Shock Absorbers

by

Christopher M. Boggs

Dissertation submitted to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in
Mechanical Engineering

Dr. Steve C. Southward, Co-Chairman

Dr. Mehdi Ahmadian, Co-Chairman

Dr. John B. Ferris

Dr. Saied Taheri

Dr. Corina Sandu

Dr. Leigh McCue

January 19, 2009

Danville, Virginia

keywords: vehicle dynamics, system identification, dynamic substructuring,
modeling, simulation, 8-post rig, shock absorber, damper

Copyright 2009, Christopher M. Boggs

The Use of Simulation to Expedite Experimental Investigations of the Effect of High-Performance Shock Absorbers

Christopher M. Boggs

Abstract

Successful race teams rely heavily on track testing to search for the ideal suspension setup. As more restrictions are placed on the amount of on-track testing by major racing sanctioning bodies, such as NASCAR, teams have increased their attention to alternate testing methods to augment their track data and better understand the dynamics of their racecars. One popular alternate to track testing is 8-post dynamic shaker rig testing. Eight-post rig testing gives the team a better understanding of the vehicle's dynamics before they arrive at the race track, allowing them to use their limited track testing time more efficiently.

While 8-post rig testing certainly is an attractive option, an extensive test matrix is often required to find the best suspension setups. To take full advantage of 8-post rig tests, more efficient experimental methods are needed. Since investigating shock absorber selection is often the most time-consuming task, this study focuses on developing more efficient methods to select the best shock absorber setups.

This study develops a novel method that applies dynamic substructuring and system identification to generate a mathematical model that predicts the results of future tests as both command inputs and components are changed. This method is used to predict the results of 8-post rig tests as actuator commands and shock absorber forces are varied. The resulting model can then be coupled with shock absorber models to simulate how the vehicle response changes with shock absorber selection. This model can then be applied to experimental design.

First, a physically-motivated nonlinear dynamic shock absorber model is developed, suitable for quickly fitting experimental data and implementing in simulation studies. Next, a system identification method to identify a vehicle model using experimental data is developed. The vehicle model is then used to predict response trends as shock absorber selection is varied. Comparison of simulation and experimental results show that this model can be used to predict the response levels for 8-post rig tests and aid in streamlining 8-post rig testing experimental designs.

Acknowledgements

I would first like to thank my advisors, Dr. Steve Southward and Dr. Mehdi Ahmadian, for their guidance and support during this project and throughout my graduate studies in the Mechanical Engineering Department at Virginia Tech. I would also like to thank Dr. John Ferris, Dr. Saied Taheri, Dr. Corina Sandu, and Dr. Leigh McCue for their advice and for serving on my graduate committee.

I would also like to thank the Virginia Institute for Performance Engineering and Research (VIPER) for their support in providing excellent facilities and equipment for my research along with ample laboratory time required to test this work. Special thanks are due to Dr. Steve Southward and Bryan Pittman for serving as an 8-post rig operator for all my rig testing. Thanks are also due to Mike Alex and Chandler Reubush in providing me background on current trends in 7-post rig testing and analysis.

Finally, I would like to thank my family – my mother and father, Carolyn and Doug Boggs, and my brother Gregory for their support during my studies at Virginia Tech.

Contents

Abstract.....	ii
Acknowledgements.....	ii
Acknowledgements.....	iii
Contents	iv
List of Figures.....	vii
List of Tables	xii
Nomenclature.....	xiii
CHAPTER 1 INTRODUCTION.....	1
1.1 Motivation.....	1
1.2 Objectives	4
1.3 Approach.....	4
1.4 Contribution.....	6
1.5 Outline	7
CHAPTER 2 BACKGROUND	8
2.1 Analytical Framework	8
2.2 Suspension Design and the Role of the Shock Absorber.....	18
2.3 Shock Absorber Construction.....	20
2.4 Shock Absorber Modeling.....	24

2.5 8-Post Testing	25
2.6 The Role of Experimental Data in Vehicle Modeling	31
CHAPTER 3 SHOCK ABSORBER MODELING	34
3.1 Physical Motivation	34
3.2 Experimental Setup.....	48
3.3 Shock Absorber Modeling.....	54
3.4 Model Validation	59
3.5 Summary.....	69
CHAPTER 4 QUARTER-CAR MODEL DEVELOPMENT	71
4.1 Quarter-Car Model, Decoupled Analysis	72
4.2 Quarter-Car Model, Simulated Data.....	79
4.3 System Identification Methods	81
4.4 System Identification on Simulated Data	88
4.5 System Identification for Quarter-Car Rig Data.....	92
4.6 Issues with Extending to Full-Vehicle Modeling	101
CHAPTER 5 FULL-VEHICLE MODELING AND IDENTIFICATION	103
5.1 Full-Vehicle Model.....	103
5.2 System Identification Methods	118
5.3 System Identification Results	128
5.4 Summary.....	137

CHAPTER 6 8-POST RIG TESTING AND IDENTIFICATION	139
6.1 Experimental Setup.....	139
6.2 Linear Model Validity	152
6.3 8-Post Rig System Identification	162
6.4 Model Validation.....	170
6.5 Summary.....	178
CHAPTER 7 PREDICTING THE INFLUENCE OF SHOCK ABSORBER SETUP.....	179
7.1 Shock Build Database.....	179
7.2 Simulated Shock Trends	182
7.3 Experimental Results	188
7.4 Summary	190
CHAPTER 8 CONCLUSIONS	191
8.1 Summary and Conclusions	191
8.2 Recommendations for Future Research.....	194
REFERENCES.....	200

List of Figures

Figure 1.1. Synergy Racing car on 8-post rig (photo by author, 2009)	2
Figure 1.2. Dynamic substructuring example	4
Figure 2.1. Dynamic substructuring example	9
Figure 2.2. Substructuring for car on track	11
Figure 2.3. (a) Vehicle dynamics present on 8-post rig, (b) Dynamics excited by 8-post rig	11
Figure 2.4. Vehicle on 8-post rig and shock model substructuring	12
Figure 2.5. System identification flowchart	14
Figure 2.6. Penske 7300 shock absorber: (a) External view (photo by author, 2009), (b) Section view (adapted from [9], used with permission of Randy Lawrence, President, Penske Racing Shocks, 2009), (c) Diagram	21
Figure 2.7. Construction of the main piston (adapted from [9] , used with permission of Randy Lawrence, President, Penske Racing Shocks, 2009)	22
Figure 2.8. Flows through the main piston during rebound: (a) Low speed bleed flow, (b) High speed valve flow (adapted from [9] , used with permission of Randy Lawrence, President, Penske Racing Shocks, 2009)	23
Figure 3.1. Basic shock absorber construction	35
Figure 3.2. Equivalent fluid system for shock absorber	36
Figure 3.3. Mechanical equivalents: (a) Damper in series with a spring, (b) Damper in series and in parallel with springs	41
Figure 3.4. Sample shock dynamometer data and curve fit	42
Figure 3.5. (a) Error in extrapolating a high-order polynomial fit, (b) Regions for polynomials in Polylinear model	43
Figure 3.6. LPF1 shock absorber model	44
Figure 3.7. The effect of a stiffening fluid on time constant: (a) Fluid stiffness, (b) Time constant	45
Figure 3.8. Output weighting for LPNL2 model	46
Figure 3.9. Penske 7300 shock absorber on a Roehrig EMA (photo by author, 2009)	48
Figure 3.10. Force-velocity for the three shock builds tested: (a) Build 1, (b) Build 2, (c) Build 4	50
Figure 3.11. Segment of the random drivefile	52

Figure 3.12. Static test drivefile.....	53
Figure 3.13. Bump test drivefile.....	53
Figure 3.14. Polylinear model, force-velocity plot.....	55
Figure 3.15. Polylinear model, error vs. time	56
Figure 3.16. LPF1 model, force-velocity plot	57
Figure 3.17. LPF1 model, error vs. time.....	57
Figure 3.18. Random validation: (a) Build 1, (b) Build 2, (c) Build 4	61
Figure 3.19. Force-velocity for build 2, 1 click: (a) 50 psi, (b) 150 psi.....	62
Figure 3.20. Comparison of sine validation error versus model type for build 1, 14 clicks, 50 psi	64
Figure 3.21. Force-velocity for sine tests: (a) low velocity, (b) high velocity.....	65
Figure 3.22. Static test time plots: (a) Drive profile, (b) Force	66
Figure 3.23. Static test force-displacement.....	67
Figure 3.24. Force vs. time for build 4 bump test: (a) Compression, (b) Rebound	68
Figure 3.25. Force-velocity plot for build 4 bump test.....	69
Figure 4.1. The quarter-car model.....	73
Figure 4.2. Effect of damping on tire force	76
Figure 4.3. Effect of damping on ride height.....	77
Figure 4.4. Effect of suspension stiffness on tire force.....	78
Figure 4.5. Effect of suspension stiffness on ride height.....	78
Figure 4.6. Coupling of vehicle and shock absorber models.....	79
Figure 4.7. Simulink layout for quarter-car model	81
Figure 4.8. Diagram of complete system and the subsystem to identify	82
Figure 4.9. System ID experiments: (a) Linear shock for road and aero inputs, (b) Second shock for shock input.....	85
Figure 4.10. ID results for tire force	89
Figure 4.11. ID results ride height.....	90
Figure 4.12. Simulink layout for simulation using the ID model	91
Figure 4.13. Relative amplitude spectra for drivefile	92

Figure 4.14. Comparison of tire force standard deviation	92
Figure 4.15. Quarter-car rig at CVeSS (photo by author, 2009).....	94
Figure 4.16. Sensor mounting: (a) Accelerometer, (b) Shock potentiometer, (c) String potentiometer with integrated tachometer (photos by author, 2009).....	95
Figure 4.17. Other components: (a) dSPACE AutoBox, (b) dSPACE Controldesk software, (c) MTS 407 controller (photos by author, 2009).....	96
Figure 4.18. Quarter-car rig FRF estimates.....	99
Figure 4.19. Coherence for FRF estimates	99
Figure 4.20. Comparison of simulated and measured time signals	100
Figure 4.21. Comparison of simulated versus measured response measures	101
Figure 5.1. Chassis geometry and forces	105
Figure 5.2. Forces applied due to anti-rollbar.....	108
Figure 5.3. Shock velocity frequency response.....	113
Figure 5.4. Tire force frequency response	114
Figure 5.5. Ride height frequency response	114
Figure 5.6. Coupling of vehicle and shock absorber models.....	116
Figure 5.7. Full-vehicle model implemented in Simulink	117
Figure 5.8. Diagram of system to identify.....	118
Figure 5.9. Shock velocity FRF estimate: (a) FRF estimate, (b) Coherence	129
Figure 5.10. Tire force FRF estimate: (a) FRF estimate, (b) Coherence	130
Figure 5.11. Ride height FRF estimate: (a) FRF estimate, (b) Coherence.....	131
Figure 5.12. Comparison of ideal FRF, FRF estimate, and modified FRF estimate	133
Figure 5.13. Result of parametric identification.....	134
Figure 5.14. Simulink model.....	135
Figure 5.15. Simulation time trace comparison.....	136
Figure 5.16. Simulation PSD comparison	137
Figure 5.17. Metric comparison	137
Figure 6.1. 8-post rig: (a) NASCAR Cup car on rig, (b) Rig with pit cover removed, (c) Wheelloader and aeroloaders mounted under pit cover (photos by author, 2009)	140

Figure 6.2. Rig attachments to vehicle: (a) Wheel platen and wheel restraint, (b) Wheel platen load cells, (c) Aeroloader mounting and load cell (photos by author, 2009).....	141
Figure 6.3. Test control: (a) SmarTest controller, (b) FasTest PC interface (photos by author, 2009)	142
Figure 6.4. Test vehicle suspension: (a) Independent SLA front suspension, (b) Solid-axle trailing-arm rear suspension (photos by author, 2009).....	143
Figure 6.5. Shock instrumentation: (a) Shock potentiometer, (b) Shock load cell (photos by author, 2009)	144
Figure 6.6. Accelerometers: (a) Wheel and chassis accelerometers, (b) Cockpit accelerometers (photos by author, 2009)	145
Figure 6.7. Chassis spring rating tests	150
Figure 6.8. Shock builds: (a) Baseline, (b) Alternate	151
Figure 6.9. Shock velocity due to wheelloader input: (a) Time response, (b) FRF magnitude, (c) Coherence.....	154
Figure 6.10. Shock velocity due to aeroloader input: (a) Time response, (b) FRF magnitude, (c) Coherence	155
Figure 6.11. Tire force due to wheelloader input: (a) Time response, (b) FRF magnitude, (c) Coherence	156
Figure 6.12. Tire force due to aeroloader input: (a) Time response, (b) FRF magnitude, (c) Coherence.	157
Figure 6.13. (a) Aeroloader and rear spring locations, (b) Simplified rear axle model	157
Figure 6.14. Ride height due to wheelloader input: (a) Time response, (b) FRF magnitude, (c) Coherence	158
Figure 6.15. Ride height due to aeroloader input: (a) Time response, (b) FRF magnitude, (c) Coherence	159
Figure 6.16. Shock velocity due to shock force input: (a) Change in time response, (b) FRF magnitude, (c) Coherence.....	160
Figure 6.17. Tire force due to shock force input: (a) Change in time response, (b) FRF magnitude, (c) Coherence.....	161
Figure 6.18. Ride height due to shock force input: (a) Changer in time response, (b) FRF magnitude, (c) Coherence.....	161
Figure 6.19. Total coherence for select outputs.....	166
Figure 6.20. Shock velocity FRF estimate.....	168
Figure 6.21. Tire force FRF estimate.....	168
Figure 6.22. Ride height FRF estimate.....	169
Figure 6.23. Hub acceleration FRF estimate	169

Figure 6.24. Transfer function fits for shock velocity	170
Figure 6.25. Simulink model	171
Figure 6.26. Baseline response for Richmond drivefile	172
Figure 6.27. Shock setup RMS trending for Richmond drivefile	173
Figure 6.28. Shock setup RMS correlation for Richmond drivefile	174
Figure 6.29. Sine heave time results.....	175
Figure 6.30. Sine heave testing RMS results.....	175
Figure 6.31. Bump testing results.....	176
Figure 6.32. Static testing results.....	177
Figure 7.1. Front ride height simulation trending.....	183
Figure 7.2. Front ride height simulation tradeoffs.....	184
Figure 7.3. Total tire force simulation trending.....	185
Figure 7.4. Total tire force simulation tradeoffs.....	186
Figure 7.5. Right front hub acceleration simulation trending.....	187
Figure 7.6. Right front hub acceleration simulation tradeoffs.....	187
Figure 7.7. Comparison of simulated and measured shock trending.....	189

List of Tables

Table 2.1. Comparison of vehicle excitation and behavior on track and on 8-post rig.....	28
Table 3.1. Shock configurations tested.....	50
Table 3.2. Definition of relative amplitude spectrum.....	51
Table 3.3. Summary of sine wave tests.....	53
Table 3.4. Summary of model fitting for build 1, 50 psi, 1 click.....	59
Table 3.5. RMS error analysis, lb.....	60
Table 4.1. Baseline parameter values.....	76
Table 4.2. Comparison of natural frequency and damping ratios.....	90
Table 4.3. Relative amplitude spectrum for random drivefile.....	96
Table 4.4. Sine wave tests.....	97
Table 4.5. Shock configurations tested.....	97
Table 5.1. Model parameters.....	112
Table 6.1. Summary of measured and calculated channels.....	145
Table 6.2. Sine test frequencies and amplitudes.....	147
Table 6.3. Random signal relative amplitude spectrum.....	148
Table 6.4. Importance of coherence and output energy for linear modeling.....	153
Table 6.5. Zeroed input-output pairs.....	166
Table 6.6. Number of active feedback paths.....	167
Table 7.1. Number of shock builds and simulation times.....	180
Table 7.2. Shim sizes (adapted from [9], used with permission of Randy Lawrence, President, Penske Racing Shocks, 2009).....	181
Table 7.3. Shock build database.....	182
Table 7.4. RMS level extremes.....	183
Table 7.5. Shock setups for 8-post rig testing.....	188

Nomenclature

This section defines all major variables used in this document. While the nomenclature section for each chapter lists all the variables used in that chapter for convenience, consistent nomenclature across all chapters was used in most cases.

Chapter 2

Δt	Time increment
$u[k], y[k]$	Input and output at time index k
U^k, Y^k	Input and output time history up to time index k
$M(\theta_{LF}, \theta_{RF}, \theta_{LR}, \theta_{RR})$	Input-output mapping for 8-post rig system parameterized by shock parameters
$\theta_{LF}, \theta_{RF}, \theta_{LR}, \theta_{RR}$	Shock parameterization
Θ	Shock selection space
$J(U^K, Y^K)$	Objective function
$G(U^K, Y^K)$	Response constraints
$M_{shock}(\theta_{shock})$	Shock substructure mapping
$M_{car}(\theta_{car})$	Car substructure mapping

Chapter 3

Q_c, Q_r	Rate of change of volume for compression and rebound chambers
Q_g	Flow that displaces gas piston
Q_p	Flow through main piston
\bar{Q}_c, \bar{Q}_r	Effective compressive flows for compression and rebound chambers

p_c, p_r	Compression and rebound chamber pressure
p_{co}	Initial chamber pressure
k_c, k_r	Stiffness of compression and rebound chamber compressibility
k_g	Gas chamber stiffness
k_g'	Stiffness of gas spring and compression chamber in series
C	Flow resistance for main piston
A_c, A_r	Piston area on compression and rebound sides
A_{rod}	Piston rod area
x, v	Piston rod displacement and velocity
F	Shock absorber force
C_{inc}, K_{inc}, F_o	Incompressible shock absorber damping, stiffness, and preload force
s	Laplace variable
τ	Shock model time constant
k_p, k_s, c	Equivalent mechanical model parallel stiffness, series stiffness, and damping
$F_s, f(v)$	Algebraic shock force
F_{lag}	LPF1 force
$\alpha(F_s)$	LPNL2 weighting function
$x_{unscaled}(t)$	Unscaled time signal
A_i, f_i, ϕ_i	Amplitude, frequency, and phase at i^{th} frequency

Chapter 4

m_1, m_2	Sprung and unsprung masses
k, k_t	Suspension and tire stiffness
\bar{c}	Damping

w_1, w_2	Sprung and unsprung mass displacement
y_R, F_A, F_S	Inputs from road, aerodynamics, and shock
w	Quarter-car degrees of freedom
M, C, K, L	State-space mass, damping, stiffness, and input matrices
x	State matrix
u	Input matrix
y	Output matrix
A, B, C, D	Continuous-time state-space matrices
A_d, B_d, C_d, D_d	Discrete-time state-space matrices
x_s	Suspension travel
x_{RH}	Ride height
F_{spring}, F_t	Spring and tire forces
g	Gravitational constant
H	Frequency response function (FRF)
s	Laplace operator
T	Sample time
z	Unit-time advance operator
j	Imaginary unit
ω	Frequency, radians/second
f	Frequency, Hz
$F_{polylinear}$	Polylinear force
F_{LPNL2}	LPNL2 force
τ	Time constant
H_1, H_2, H_3	FRF matrices
u_1, u_2, u_3	Road, aerodynamic, and shock force inputs
ε	Modeling error in frequency domain
Y, U	Outputs and inputs in frequency domain for all data blocks

J	Objective function
S_{uu}, S_{uy}	Input spectrum and cross spectrum
F_{shock}	Shock force
c	Linear damping
v	Shock velocity
S	Matrix to select shock velocity from output vector
\tilde{H}_1, \tilde{H}_2	FRF matrices with linear shock
\bar{H}_1, \bar{H}_2	Shock force FRF with linear shock
$y_{\text{experiment 1}}$	Output from baseline experiment
$B(s), A(s), b_n, a_m$	Transfer function numerator and denominator polynomials and their associated polynomial coefficients
$w_i(k)$	Weighting function

Chapter 5

Vehicle Parameters

t, a, b	Track width and distance from CG to front and rear axles
M, J_θ, J_ϕ	Chassis inertial properties
$m_t^{LF}, m_t^{RF}, m_t^{LR}, m_t^{RR}$	Wheel masses

Vehicle Position

$z_1^{LF}, z_1^{RF}, z_1^{LR}, z_1^{RR}$	Chassis vertical displacements at four corners
$z_2^{LF}, z_2^{RF}, z_2^{LR}, z_2^{RR}$	Wheel vertical displacements at four corners
z, θ, ϕ	Chassis heave, roll, and pitch
z_1	Chassis vertical displacements vector
z_2	Wheel vertical displacements vector
Φ	Chassis heave, roll, and pitch vector

T Transformation from chassis mode displacements to vertical displacements

Inputs

F_z, M_θ, M_ϕ External forces applied to chassis

$z_R^{LF}, z_R^{RF}, z_R^{LR}, z_R^{RR}$ Road input at four corners

z_R, F_{ext}, F_D Road inputs, external chassis forces, and arbitrary suspension forces

7DOF Model Development

M_c, M_t Diagonal mass matrices for chassis and tires

$F_{LF}, F_{RF}, F_{LR}, F_{RR}$ Suspension forces at four corners

$F_t^{LF}, F_t^{RF}, F_t^{LR}, F_t^{RR}$ Tire forces at four corners

F_c, F_t Suspension and tire force vectors

$k^{LF}, k^{RF}, k^{LR}, k^{RR}$ Suspension stiffness

$k_t^{LF}, k_t^{RF}, k_t^{LR}, k_t^{RR}$ Tire stiffnesses

K_s, K_t, C_p 4x4 diagonal suspension stiffness, tire stiffness, and suspension damping matrices

K_p 7x7 primary suspension and tire stiffness matrix

k_R^F, k_R^R Front and rear anti-rollbar stiffnesses

K_1, K_{roll} 4x7 and 7x7 anti-rollbar stiffness matrix

Structural and State-Space Equations

Z Degrees of freedom

\bar{M}, C_s, K, L Mass, damping, stiffness, and input matrices

u, y, x Input, output, and state vectors

A, B, C, D State-space matrices

Output Equations

v_{shock} Shock absorber velocity

FRH, RRH	Front and rear ride heights
l_1, l_2, l	Ride height geometry
<u>FRF Estimation</u>	
u_{rig}, u_{shocks}	Input vectors corresponding to actuators and shock forces
B_{rig}, B_{shocks}	Input matrices corresponding to actuators and shock forces
D_{rig}, D_{shocks}	Output matrices corresponding to actuators and shock forces
s	Laplace operator
H_{rig}, H_{shocks}	Frequency response due to actuator and shock force inputs
S	Matrix to select shock velocity from output vector
Y, U	Outputs and inputs in frequency domain for all data blocks
E	Modeling error in the frequency domain
H	Frequency response function (FRF)
J	Objective function
S_{uu}, S_{uy}	Input spectrum and cross spectrum
\bar{C}	Linear damping
H_A	FRF matrix with linear shocks
ΔF	Arbitrary shock force in addition to linear shock force
$\Delta f_1, \Delta f_2, \Delta f_3, \Delta f_4$	Arbitrary shock force in addition to linear shock force
$u_{shocks}^1, u_{shocks}^2, u_{shocks}^3, u_{shocks}^4$	Shock force inputs
$\bar{c}_1, \bar{c}_2, \bar{c}_3, \bar{c}_4$	Linear damping coefficients
$v_s^1, v_s^2, v_s^3, v_s^4$	Shock velocities
$\tilde{H}_{rig}, \tilde{H}_{shocks}$	FRF matrices with linear shocks and an additional arbitrary shock force
U_{rig}	Inputs for all blocks applied to baseline setup
$\bar{c}'_1, \Delta c_1$	Alternate left front shock damping coefficient and change in damping coefficient from baseline

$y_0[k], \varepsilon_0[k]$	Output and error for k^{th} baseline data block
$y_i, u_i^T, \varepsilon_i$	Output, input, and error for i^{th} alternate experiment
E_{rig}	Error for all data blocks for baseline setup
$J_i _{\tilde{H}_{rig}}$	Objective function for i^{th} alternate experiment given the FRF estimate for the baseline inputs
$y_i[k], u_i[k]$	Output and input for i^{th} alternate experiment and k^{th} data block
$\Delta y_i[k], \Delta f_i[k], \varepsilon_i[k]$	Change in output, change in shock force, and modeling error for i^{th} alternate experiment and k^{th} data block
$\Delta Y_i, \Delta F_i, E_i$	Change in output, change in shock force, and modeling error for i^{th} alternate experiment and all data blocks
\tilde{H}_{shock_i}	Shock FRF for i^{th} shock
f	Frequency, Hz
e	Euler's number
j	Imaginary unit
H_{high}, H_{low}	High and low frequency residual fit form

Transfer Function Fitting

$B(s), A(s), b_n, a_m$	Transfer function numerator and denominator polynomials and their associated polynomial coefficients
$w_i(k)$	Weighting function

Chapter 6

X, X_m, T_x	Aeroloader displacements, aeroloader mode displacements, and the associated transformation
F, F_m, T_f	Aeroloader forces, aeroloader mode forces, and the associated transformation
b', L', a', t_f, t_r	Aeroloader geometry about reference point along vehicle centerline

y	Output vector
u_{rig}, u_{shocks}	Input vectors corresponding to actuators and shock forces
H_{rig}, H_{shocks}	Frequency response due to actuator and shock force inputs
\bar{C}	Linear damping
v_{shocks}	Shock absorber velocity
ΔF	Arbitrary shock force in addition to linear shock force
S	Matrix to select shock velocity from output vector
$\tilde{H}_{rig}, \tilde{H}_{shocks}$	FRF matrices with linear shocks and an additional arbitrary shock force
$y^{(i)}[k], u_{rig}[k]$	Output and input for i^{th} alternate experiment and k^{th} data block
$\Delta F^{(i)}[k], \varepsilon^{(i)}[k]$	Change in shock force and modeling error for i^{th} alternate experiment and k^{th} data block
Y, U	Outputs and inputs in frequency domain for all data blocks
E	Modeling error in the frequency domain
H	Frequency response function (FRF)
$[h_{ij}]$	Logical matrix with zeroes if corresponding entry in \tilde{H}_{shocks} was zeroed
f	Frequency, Hz
$w_{aero}, w_{wheels}, w_{shocks}$	Weighting function for aeroloader, wheelloader, and shock force transfer function fitting
$F_{aero}, X_{aero}, F_{zero}$	Aeroloader force, displacement, and force at zero displacement
K	Spring matrix, expressed in terms of aeroloader forces and displacements
K_m	Spring matrix, expressed in terms of aeroloader mode forces and displacements

Chapter 1

Introduction

The following sections describe the motivation for the proposed research, research objectives, the approach, and the expected contributions.

1.1 Motivation

Successful race teams search for every possible advantage over the competition, no matter how small, to reduce lap times and win the race. While track testing is ultimately the most valid testing method available, the amount of time and resources track testing takes is prohibitive when searching for the perfect setup. As more restrictions are placed on the amount of on-track testing by major racing sanctioning bodies, such as NASCAR, teams have increased their attention to alternate testing methods to augment their track data and better understand the dynamics of their racecars.

One popular alternative to track testing that can be used to better understand the vehicle's dynamics is shaker rig testing, such as 7-post or 8-post rig testing, as shown in Figure 1.1. Seven-post rig testing is widely used in Formula-One racing, but has become popular in NASCAR only in the last few years [1, 2]. During a 7-post or 8-post rig test, the tires are supported by 4 actuated wheel platforms, or wheelloaders, that simulate inputs from the track surface. Three or four actuators, or aeroloaders, attach to the vehicle's chassis to provide forces which simulate the effect of inertial and aerodynamic forces present during a track test. The vehicle is shaken on the rig to simulate how the car would respond at a particular track or to characterize the vehicle response to more general waveforms such as sine waves or broadband signals. In this study all tests will be conducted on an 8-post rig, but all work is applicable to a 7-post rig, so we will refer to 8-post testing unless specifically referring to 7-post testing.



Figure 1.1. Synergy Racing car on 8-post rig (photo by author, 2009)

Eight-post testing is an attractive option to race teams, and has become very popular in the past several years. Unlike track testing, rig testing can be conducted as much as the team wants. Also, the costs due to facility rental, travel, number of personnel, gas, tires, etc. is much lower for an 8-post rig test. The test conditions are more repeatable during an 8-post rig test than during a track test, which makes it easier to distinguish between the effect of changing the setup and test variability.

While 8-post rig testing certainly has advantages, there are also downsides. Some of the major challenges will be discussed in further detail in Chapter 2. One issue with 8-post rig testing is that while it is more efficient than track testing, an extensive test matrix can still take a long time. To take full advantage of 8-post rig tests, more efficient experimental design methods are needed.

During an 8-post rig test, different vehicle adjustments are made to determine the effect of each adjustment and to find ideal setups for different track conditions. Some of the parameters that can be considered in a typical 8-post rig test are shown in Table 1.1.

Table 1.1. Potential 8-post rig testing adjustments

12 shock settings per corner	x4	=	48
1 spring rate and 1 preload per corner	x4	=	8
1 rollbar rate per axle	x2	=	2
1 trackbar height	x1	=	1
1 tire pressure per tire	x4	=	4
			<hr/>
			63 parameters

Of the 63 possible adjustments, investigating the 48 potential shock absorber adjustments can easily be the most time-consuming. In this study, we will use the terms shock and shock absorber interchangeably. This burden is further confounded by the fact that most of those shock adjustments are internal to the shock, so there needs to be a shock built for each adjustment. This motivates us to determine efficient methods to vary the shock parameters, since we have the most to gain by reducing this effort.

Clearly, the shock adjustments can add the most complexity to any 8-post rig test experimental design. If we could find a way to characterize the effect of shock adjustments, and use that information to run a shaker rig test more efficiently, then we could significantly reduce the complexity of our tests.

While running a large test matrix will give the best results, it is often impractical. Most experimental designs in 8-post testing are done in a guess-and-check manner. This consists of the engineers guessing a list of potential setups, testing them, and checking the test data. After checking the data, they may generate further lists. This method clearly has problems. First, only information on those particular setups is gained, not more general sensitivity information that could be used to tune shocks. Second, potential gains can be missed if the setup is not part of the guessed list. Third, if the guesses are poor, there is often no indication of what to try next. Clearly, this method is highly dependent on the engineer's ability to guess setups.

Most teams also test their shocks on a shock dynamometer. Since shocks can be the most challenging components to tune during a rig test, knowing how the shock performs at a component level before running shaker rig tests could be very useful. While finding a practical method to apply shock data to 8-post rig testing remains a major challenge, the potential improvement in testing efficiency could be significant.

1.2 Objectives

The overall objective of this work is to develop methods to use race shock data from shock dynamometer tests to improve shaker rig testing. The detailed objectives of this study are to

1. Characterize the performance of race shocks,
2. Develop a validated mathematical model that uses prior 8-post rig test data to predict the influence of shock absorber selection in future 8-post rig tests, and
3. Apply the model to improve the efficiency of 8-post tests.

1.3 Approach

This study develops a novel method that applies system identification for dynamic substructuring to generate a mathematical model that predicts the results of future tests as both command inputs and components are changed. Dynamic substructuring represents the behavior of a complex system as an assembly of simpler substructures and their interactions, as shown in Figure 1.2. This modeling technique allows us to investigate the effect of changing various system components by swapping substructure models that represent the system components. This study will apply system identification techniques to identify mathematical models for all substructures from experimental data.

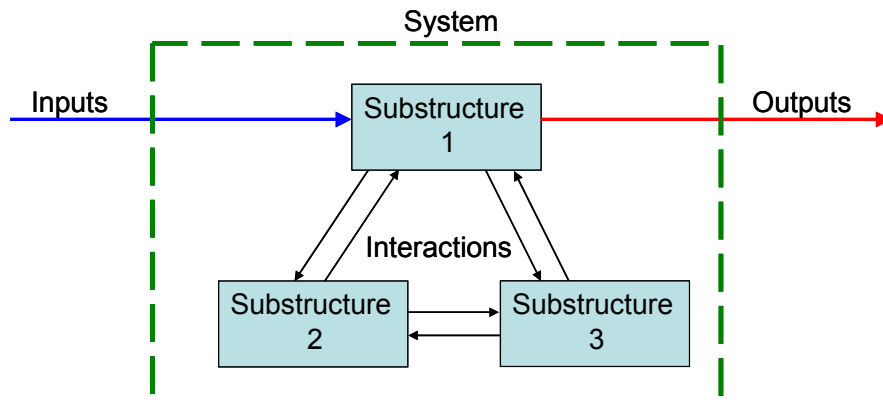


Figure 1.2. Dynamic substructuring example

This method is used to predict the results of 8-post rig tests as actuator commands and shock absorber forces are varied. The resulting model can then be coupled with shock absorber models to simulate how the vehicle response changes with shock absorber selection. This model can then be applied to experimental design.

The detailed steps that will be taken to achieve the objectives are:

Shock Modeling

1. Develop characterizations of race shock performance using a shock dynamometer,

Quarter-Car Study

2. Develop a method to identify a vehicle model using simulated quarter-car rig data,
3. Run experiment on quarter-car rig and apply identification method,

8-Post Rig Study

4. Expand method to full scale using simulated 8-post rig data,
5. Apply method to 8-post rig data and validate model accuracy, and

Apply Method

6. Use model to predict vehicle performance for different setups and help select setups, and
7. Conduct rig tests for different shock setups to validate the model's ability to predict vehicle performance and select setups.

Step 1 involves developing shock absorber models using shock dynamometer data. First, a series of shock dynamometer tests were run on different shock builds, which are used for both model fitting and validation purposes. Candidate shock models are developed, which are then fitted with the experimental data. Additional dynamometer data is used to validate the model performance and see where the model performs well.

To make it easier to develop the vehicle system identification method, it was done first at a quarter-car scale in Steps 2 and 3. To make sure the identification method can extract the vehicle model accurately, the method is first applied to simulated data in Step 2 generated by a known quarter-car model. Once the system identification method has been proven on simulated data, it is then applied to test data collected from a quarter-car rig in Step 3 to verify that the method will work on real data.

The method is expanded to full 8-post rig scale in Steps 4 and 5. Similar to the quarter-car study, the method is first developed using simulated 8-post rig data to work out the bugs before it is applied to actual rig data.

Once the 8-post rig vehicle model is created, Step 6 applies the model to help select shock setups and identify sensitivity information. The results from Step 6 are validated in Step 7 by running a series of rig tests for some of the same setups considered in Step 6.

1.4 Contribution

As discussed above, the contribution of this research to both the racing industry and the larger engineering community is significant. The expected contributions of this work to the racing industry include

1. Providing more accurate empirical shock absorber models that can be used for simulation studies,
2. Providing a method to identify an empirical model of the influence of component selection during 8-post rig testing, and
3. Improving the efficiency of 8-post test sessions by using model to identify trends and help shock selection.

This work may also have a larger impact on the engineering community since

1. The method may be extended to include additional components,

2. The models can be used to suggest targets for shock design or be used as a part of a larger vehicle simulation, and
3. The models and methods may be valid for different classes of vehicles or other dynamic systems.

1.5 Outline

This study begins by developing shock absorber models, proceeds by developing the vehicle model identification method at quarter-car and full-car scales, and concludes by applying the identified model to predict trends in the vehicle response as shock setups are changed. Chapter 2 covers relevant background information. Chapter 3 covers the shock absorber modeling. Chapter 4 covers the quarter-car study. Chapter 5 develops the method at full-vehicle scale using simulated data from a 7 degree of freedom model. Chapter 6 applies the method to actual 8-post rig data and validates the model with additional data sets. Chapter 7 applies the model identified in Chapter 6 to determine which setups will have a high or low RMS response level, which is then verified for several setups on the 8-post rig. Chapter 8 provides summary, conclusions, and recommendations for future work.

Chapter 2

Background

This chapter presents background information that the reader may find useful while reading this document. The chapter begins by casting this research in a larger analytical framework suitable for a larger class of engineering problems. It continues by highlighting several design requirements of a racing suspension and the role of the shock absorber. Next, the construction of the shock absorber used in this study is described. Then, prior work in shock absorber modeling is covered. Next, a summary of the development of 8-post technology and the published work in that area is presented. Finally, the role of experimental data in developing quality vehicle models is discussed.

2.1 Analytical Framework

While this research focuses on the application of shock selection during an 8-post rig test, the approach and methods used in this research are applicable for a larger class of engineering problems. To highlight how this research is applicable for a larger class of problems, this section presents this research within a more general analytical framework. First, the concept of dynamic substructuring is presented as a convenient approach for manipulating complex models. After presenting the concept of dynamic substructuring, the concept will then be applied to the problem of shock selection during 8-post rig testing. The 8-post rig test will be shown to be a test of a substructure of the more complex vehicle system. The vehicle during the 8-post rig test may be further substructured into the components of interest. In this study our focus is shock absorbers, so the vehicle on the 8-post rig will be substructured into 4 substructures for the shock absorbers and one main substructure representing the vehicle and any remaining vehicle components. Once the vehicle substructuring has been established, this framework can then be used to state the shock absorber optimization problem and the 8-post rig

substructure system identification problem. At this point, we will formally state the problem addressed in this research.

2.1.1 Dynamic Substructuring

Dynamic substructuring is a strategy to model a complex system as an assembly of substructures and interactions between the substructures. An example of substructuring a system into three substructures is shown in Figure 2.1. Substructuring involves dividing a larger structure into substructures, which are chosen to best suit the application. Each substructure is treated as a separate system, with inputs and outputs that are used to couple the substructures, accept the system inputs, and generate the system outputs. In mechanical systems, these interactions typically consist of a compatibility constraint on displacement and an equilibrium constraint on force.

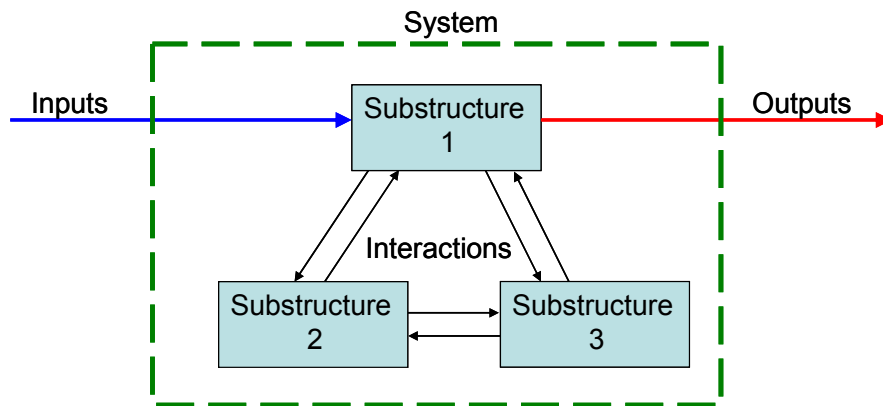


Figure 2.1. Dynamic substructuring example

Since substructuring allows the user to change model components as a substructure, the technique can be thought of as a compromise between black-box modeling and physical modeling.

Some of the potential advantages of substructuring include

1. Using a combination of empirical and analytical models;
2. Allowing work to be done on a simpler system, which can be useful for
 - a. Model Derivation,

- b. Experiments,
 - c. Analysis,
 - d. Simulation, and
 - e. System Identification;
3. Exploring the effect of changing components and geometry by changing substructures and topology;
 4. Increasing computational efficiency; and
 5. Predicting the internal response, such as forces applied to a component.

2.1.2 Vehicle Substructuring

A car on the road may be substructured by drawing a control surface around the chassis, cutting through the shocks, springs, and control arms. This control surface will also intersect the brakes, throttle, and steering, which are controlled by the driver. These components act to couple the chassis to the wheels. At each cut, the interaction is represented by geometric compatibility and force equilibrium. The tire is mounted on the wheel, and the tire interacts with the road surface. Also, aerodynamic forces are generated due to the chassis velocity, height from the road surface, and orientation. This substructuring is shown in Figure 2.2. The vehicle on the 8-post rig is consistent with the subsystems enclosed by the dashed line.

On the 8-post rig, the forces applied are primarily in the vertical direction. While the lateral and longitudinal tire grip on the low friction wheelloader platens will apply small lateral and longitudinal forces to the vehicle, these are often not large enough to create significant motion in the lateral, longitudinal, and yaw degrees of freedom. These degrees of freedom, which may be considered the yaw-plane dynamics, are present on the rig but are not significantly excited by the rig. On the track, these dynamics are excited by the lateral and longitudinal tire grip and the inertial forces due to the vehicle's trajectory on the track. When on the track, the yaw-plane dynamics couple with the remaining dynamics, which we will call the vertical dynamics. This coupling can be represented as shown in Figure 2.3.

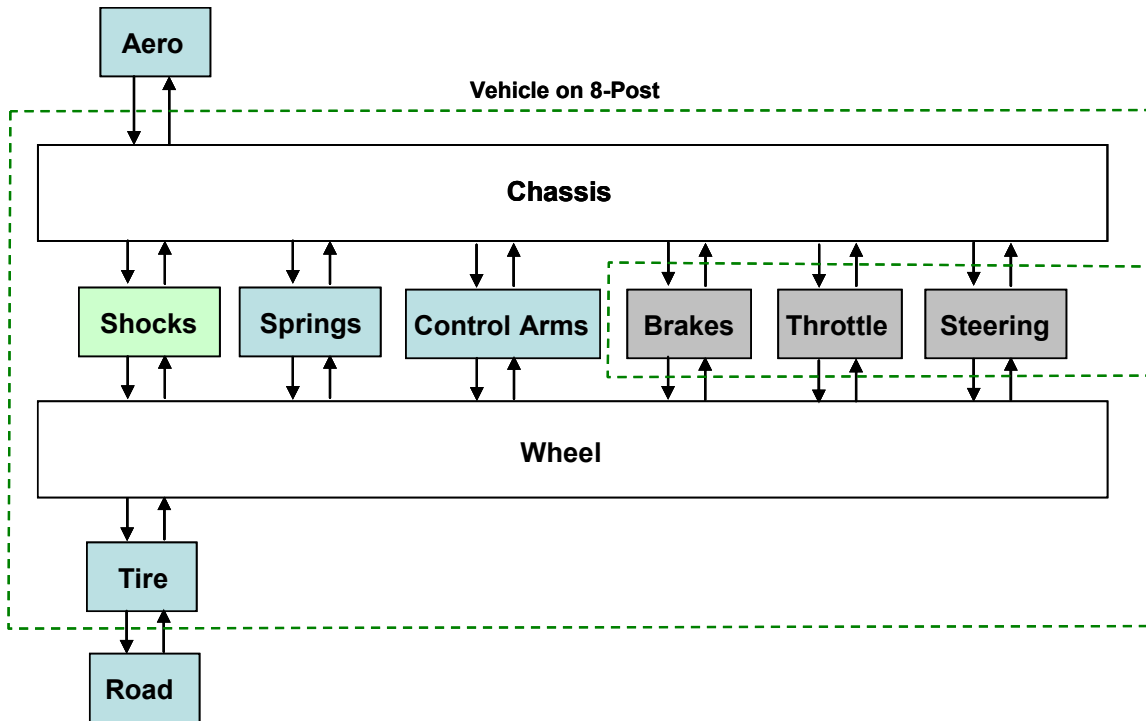


Figure 2.2. Substructuring for car on track




Figure 2.3. (a) Vehicle dynamics present on 8-post rig,
 (b) Dynamics excited by 8-post rig

Since the 8-post rig test only excites a substructure of the complete vehicle system that we wish to optimize, optimization of the vehicle performance during an 8-post rig test becomes a challenging task. There are several methods that can be applied to apply 8-post rig data to improve performance on the race track, which include

1. **Correlation:** Neglecting the coupling and tuning for metrics on the rig that have shown correlation with good track performance in previous studies,

2. **Hardware-in-the-Loop:** Including a model of the dynamics not excited on the 8-post rig and modifying the actuator commands to reflect how the coupling would affect the vertical dynamics, and
3. **Identification and Synthesis:** Using the data collected from an 8-post rig test to identify a model of the vertical dynamics to be synthesized with other components for vehicle simulations.

In this study, we will apply method 1. The model that this work produces could also be integrated into a larger vehicle model as part of method 3. Method 2 requires actually modifying the actuator commands during the 8-post rig test and is not used here.

2.1.3 8-Post Rig Substructuring

Since this study focuses on shock absorber selection on the 8-post rig, we substructure the 8-post rig model into a car model and 4 shock absorber models, as shown in Figure 2.4. Each model block consists of a model structure $M(\bullet)$, which is parameterized by the parameter vector $\theta \in \Theta$.

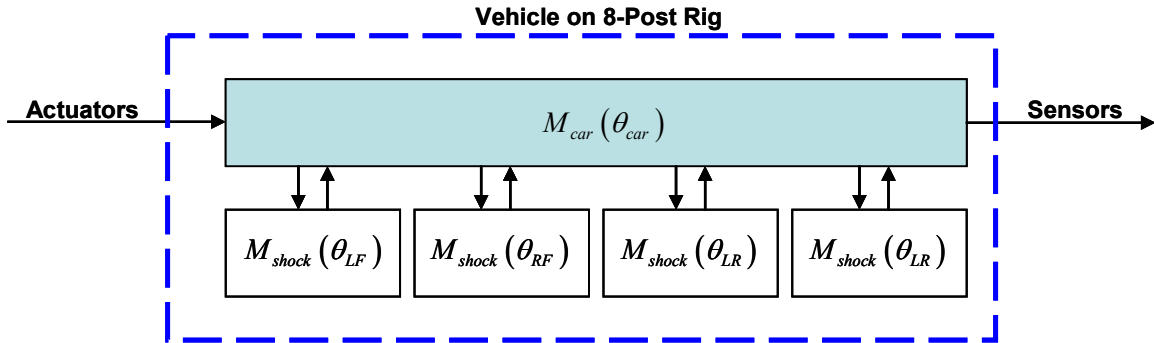


Figure 2.4. Vehicle on 8-post rig and shock model substructuring

Using this framework, we can define the problems of shock optimization and system identification that will be addressed in this research.

2.1.4 Optimization Problem

For a given car, we wish to select the best shock absorbers out of a set of candidate shock absorbers. The problem of shock absorber selection may be defined more formally as an optimization problem. The quality of a particular shock selection will be determined by a series of tests that will be conducted and an objective measure that will be calculated from the recorded data. There may also be secondary measures that must not exceed certain levels for an acceptable setup. This can be stated formally as an optimization problem.

Define k to be the sample index corresponding to time $k\Delta t$, where Δt is the sampling time. We then define $u[k]$ and $y[k]$ to be the vector inputs and outputs recorded at sample index k . Also define

$$\begin{aligned} U^k &= [u[1] \quad u[2] \quad \cdots \quad u[k]] \\ Y^k &= [y[1] \quad y[2] \quad \cdots \quad y[k]] \end{aligned} \quad (2.1)$$

which is all the sampled signals up to time index k . Defining the test to be K samples long, the matrix U^K defines the drivefiles used for rig testing. The dynamics of the 8-post rig can then be described by the nonlinear mapping

$$M(\theta_{LF}, \theta_{RF}, \theta_{LR}, \theta_{RR}): U^K \rightarrow Y^K \quad (2.2)$$

where the parameters θ_i define the shock selection, which may be selected from the shock selection space Θ . The optimization problem of selecting the best shock absorbers can then be stated

$$\begin{aligned} &\text{Minimize } J(U^K, Y^K) \\ &\text{subject to} \\ &M(\theta_{LF}, \theta_{RF}, \theta_{LR}, \theta_{RR}): U^K \rightarrow Y^K \\ &\theta_{LF}, \theta_{RF}, \theta_{LR}, \theta_{RR} \in \Theta \\ &G(U^K, Y^K) \leq 0 \end{aligned} \quad (2.3)$$

where J represents the objective function and G represents constraints on secondary response measures. This problem can be solved using gradient-based nonlinear programming methods if the candidate shock model parameter space is continuous and $J(\bullet)$ and $M(\bullet)$ are differentiable. If the candidate shock model parameters are discrete, integer programming methods must be used.

2.1.5 System Identification Problem

To perform the optimization problem shown in Equation (2.3) in simulation or analytically, the dynamics of the vehicle on the 8-post rig and the influence of shock selection must be known. A convenient method for determining the dynamics of the system is to perform system identification on the substructures shown in Figure 2.4.

The generic system identification process is shown in Figure 2.5. The process is initiated by defining a target system, which we wish to describe using a mathematical model. A series of tests are performed on the target system, where the target system input and output data is recorded. The experimental data is used both to perform system identification and to assess the model quality. System identification applies the experimental data to generate a mathematical model of the system. Once the model is identified, the response predicted by the model is compared to the response measured from the target system. If this comparison shows that the identified model is not sufficient for the particular application, the testing or system identification processes must be adjusted until an acceptable model is produced.

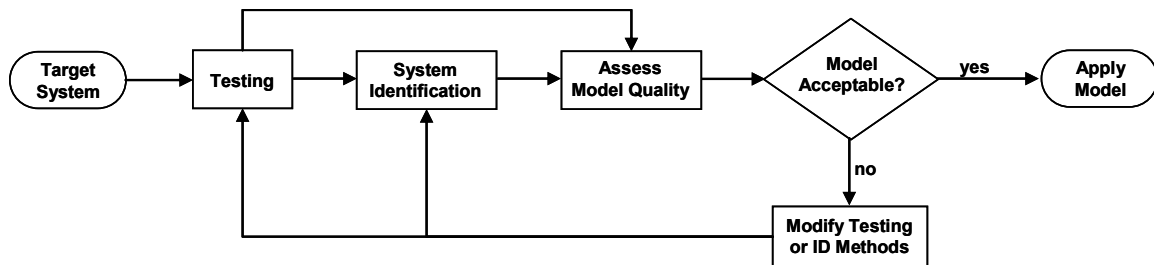


Figure 2.5. System identification flowchart

This study will first develop the testing and system identification methods in simulation before performing laboratory testing. To develop the methods in simulation, a known

mathematical model is used for the target system. Simulation results are used as simulated data sets, which can be used for system identification and assessing model quality. By eliminating modeling error due to signal noise, system nonlinearity, and repeatability, the simulation study will highlight weaknesses in the testing and system identification processes.

To perform system identification, three things must be defined

1. A model structure $M(\bullet)$,
2. A measure of model quality, and
3. A method of selecting model parameters given sufficient experimental data.

Another requirement is that the excitation signals must be sufficiently rich (sufficient richness) and the model structure must be such that the true model parameters are a unique optimal solution in terms of model quality (identifiability). This restriction requires linearly independent excitation for each input and that the model structure is not overly complex.

In this study, we will consider models for the vehicle and the shock absorbers. The model for the shock will be

$$M_{shock}(\theta_{shock}): U_{shock}^K \rightarrow Y_{shock}^K \quad (2.4)$$

where U_{shock}^K and Y_{shock}^K are the input and output histories for the shock substructure. Multiple shock absorbers will be tested at the component level on a shock dynamometer to provide the data required for shock system identification. Since multiple shocks will be tested and modeled, there will be a parameter vector θ_{shock} for each shock tested. This means that the space of candidate shock models, Θ , will consist of a finite number of points. Since Θ is not continuous, integer programming methods will be used to solve the optimization problem in Equation (2.3).

The model for the vehicle will be

$$M_{car}(\theta_{car}): U_{car}^K \rightarrow Y_{car}^K \quad (2.5)$$

Where U_{car}^K and Y_{car}^K are the input and output histories for the car substructure. Since we will only test with one vehicle in this study, there will be only one parameter vector θ_{car} . One conventional testing method for collecting the data needed for system identification would involve applying independent external forces to each input of the car substructure. This research will focus on identifying the car substructure model in Equation (2.5) with only data from complete 8-post structure testing and shock substructure testing.

2.1.6 Problem Statement

The main question that we attempt to answer in this research is:

Can the time required to find desirable shock absorber setups on the 8-post rig be reduced using dynamic substructuring for system identification and optimization?

To perform this, we will need

1. Shock absorber models,
2. A model of the vehicle on the 8-post rig and how it interacts with the shock absorbers, and
3. An objective function to rate the quality of different shock setups.

For the shock absorber substructure models, we will choose a model structure which is motivated by the dominant velocity dependence of the device and the secondary influence of fluid compressibility to add phase lag.

For the vehicle substructure model, we will use a linear model structure. The assumption of linearity may prove to be insufficient to capture some of the dynamics, but it should predict general trends that will help in shock selection. These issues will be discussed in more detail later in this chapter.

The choice of a successful objective function remains a challenge. As discussed above, the 8-post rig test can be considered a test of a substructure of the complete vehicle on the

racetrack. Rigorous analytical or simulation methods would require modeling of the entire vehicle, including the tire/track interface, aerodynamics, braking, powertrain, and steering. These models would either be used in a hardware-in-the-loop sense or to be synthesized with a model of the vertical dynamics identified from the rig. Developing and validating this level of modeling requires a large database of experimental data, which is not available at this time and is beyond the scope of this research.

The most common approach in the 8-post testing industry is to use statistics, or response measures, which have shown to be consistent with favorable track performance. The selection of the proper response measures to analyze is typically based on the experience of the engineers, not a formal correlation analysis. These response measures are typically weighted combinations of the RMS or MS response of the tire force, ride height, suspension deflection, and other measured or calculated signals in the form

$$J = \sum_{i=1}^I \rho_i \text{RMS}(y_i) \quad (2.6)$$

where ρ_i is the weighting factor for the RMS of signal y_i . The signal may be bandpass filtered to the frequency range of interest, and may include an entire test run or focus on a critical segment of the test.

The models should be chosen by the user to sufficiently characterize the desired dynamics for the application and to be evaluated quickly for efficient simulations. To be considered successful, the resulting model must be fast enough to be run many times in a day to provide suggestions for further 8-post testing.

In practice, the user would also need to select their objective function based on their best knowledge of what is most appropriate for their application. Since we do not have a validated objective function to suggest to users and this method could be used for a range of applications, our metric of success will be the ability to predict how the RMS levels shown in Equation (2.6) vary when changing shock setups. If we can successfully identify trends in these RMS measures, we will consider this method a success.

2.1.7 Dynamic Substructuring Research

Most research in system identification in dynamic substructuring focuses on linear second-order modal models [3]. In these studies, sensors must be placed on the structure to identify the modes of interest and viscous or structural damping assumptions were often made.

More recent work has applied more generic input/output model structures to substructuring [4, 5]. The presented method identifies models of each substructure independently, after which the substructure models are synthesized to create the structural model. This method requires testing of all structures independently and also requires excitation and sensors at all interface degrees of freedom. One method to provide excitation at a boundary degree of freedom is to perform tests with multiple point masses [4].

In 2008, a study was published that performs identification using a mixture of structure and substructure data [6]. This is particularly useful when one of the structures is difficult to test independently or behaves significantly different as part of the structure. This method requires inputs to both substructures and the interface degrees of freedom.

2.2 Suspension Design and the Role of the Shock Absorber

The ability of the tire to maintain contact with the race track over bumps and around corners is critical for achieving the best performance of a race car [7]. The suspension determines the coupling between the chassis and the wheel, which determines how the work done from external forces from the road and air propagate through the vehicle as kinetic energy, potential energy, and dissipated energy. The proper design and tuning of the suspension can mean the difference between the winning team and one that does not even qualify for the race.

Some objectives for high-performance suspension design may include:

1. Providing the desired tire normal force distribution between the four tires,

2. Minimizing tire normal force variations due to disturbances,
3. Providing the desired body ride height and attitude,
4. Minimizing body ride height and attitude changes due to disturbances,
5. Controlling dynamic load transfer distribution and load transfer rate during transient maneuvers,
6. Providing sufficient sensory feedback to the driver to allow for driving decisions, and
7. Maintaining acceptable vehicle stability to allow for disturbances.

Unlike passenger car design, driver comfort is not typically a primary concern.

Tire traction or grip analysis is often simplified by breaking the total grip into components of mechanical grip and aero grip [8]. Mechanical grip results from the application of inertial forces and road disturbances to the vehicle, while aero grip is caused by high-speed aerodynamic downforce. The best setup for a particular track will be a tradeoff between the best mechanical grip and the best aero grip, which will depend on the importance of aerodynamics on that particular race track. For example, at a half-mile short track race, mechanical grip is often critical, while at a 2.5 mile superspeedway track, it is often possible to sacrifice mechanical grip for improved aero grip.

To optimize mechanical grip, the tire normal forces must be distributed ideally for the current track maneuver (braking, cornering, accelerating, etc.) and the normal force variation must be minimal. The ideal normal force distribution will allow the four tires to achieve the best lateral and longitudinal force balance for the given maneuver.

Minimal normal force variation will allow the driver to drive at the absolute limit of the vehicle's performance envelope. If there is more normal force variation, it will disturb the available grip on the four tires. This will then disturb the vehicle's lateral acceleration, longitudinal acceleration, and yaw rate on the racetrack and the driver will need to adjust by steering, changing the vehicle's line, or slowing the vehicle. This alteration from the ideal racing line can make lap times slower. Often disturbances occur

too fast for the driver to correct, so they will either need to adjust their driving style to anticipate the disturbances or they may lose control. This anticipation often requires sub-limit driving, which will also make lap times slower.

To achieve optimal aero grip, the chassis must maintain an ideal ride height and attitude to provide the best total aerodynamic downforce for total grip and the desired downforce distribution to the front and rear axles for balance. Any disturbance from this ideal position will result in a change in downforce, which will cause a variation in the tire normal force.

The suspension also controls dynamic load transfer during straight-line braking, turn entry, turn exit, and straight line acceleration. A proper dynamic load transfer distribution and load transfer rate are also important for good handling.

Two of the main components that can be tuned in a suspension are the springs and the shocks. Quite often, however, the selection of spring rates, anti-rollbar stiffnesses, and roll center heights is constrained by the desired steady-state load transfer distribution and the desired DC chassis position at top speeds [7]. This means that the shock absorbers often have the most freedom to be tuned for mechanical disturbances, aero disturbances, and dynamic load transfer.

The shock absorber converts energy stored in the springs and the work done by external forces into heat instead of kinetic energy of the chassis and wheel. This is done by forcing a fluid through a restriction with a flow rate that is a function of the velocity across the suspension. Since the shock absorber is a dissipative element, it is particularly well-suited for controlling resonance, dynamic balance, and load transfer rate.

2.3 Shock Absorber Construction

This study focuses on the Penske 7300 shock absorber, shown in Figure 2.6. This shock is commonly used in NASCAR Nextel Cup and other similar racing series. This shock,

like many other racing shocks, is designed to be highly adjustable, and consists of several components, described below.




Figure 2.6. Penske 7300 shock absorber: (a) External view (photo by author, 2009), (b) Section view (adapted from [9], used with permission of Randy Lawrence, President, Penske Racing Shocks, 2009), (c) Diagram

The shock is mounted on the car using two spherical joints, which allow rotational motion and prevent the application of side loads to the shock, which would degrade performance. On the top of the shock is a gas chamber, which is separated from the hydraulic fluid by a floating gas piston. The main purpose of the gas chamber is to compensate for the volume of the piston rod as it enters the shock body. The gas chamber is pressurized with nitrogen gas using a Schrader valve on the top of the shock. This pressurization preloads the hydraulic fluid, reducing the effects of fluid compressibility and the potential for fluid cavitation due to pressure drops across the piston.

Below the gas chamber is the head valve, which separates the head valve chamber from the compression chamber. The head valve is threaded into the shock body, and can be removed when rebuilding the shock. The head valve is designed to accept different bleed jets and shim stack configurations, which can be used to fine-tune low speed and high

speed shock behavior, respectively. The main purpose of the head valve is to provide a higher pressure in the compression chamber during the compression stroke, forcing fluid into the rebound chamber to avoid cavitation.

Below the head valve is the compression chamber, which is separated from the rebound chamber by the main piston, shown in Figure 2.7. The main piston provides most of the energy dissipation in the shock, since it sees more flow than the head valve. This makes it the main factor that determines the shock's force-velocity characteristic. The components of the main piston assembly include the piston, compression shim stack, rebound shim stack, bleed jet, and needle valve. Large-scale shock curve adjustments can be made by selecting different piston types, which changes the valving. The bleed jet threads into the shaft, and can be exchanged to adjust low speed damping. The needle valve allows for the bleed area to be adjusted by changing the needle's position, which is controlled by an external knob on the bottom of the shaft. This knob is indexed by clicking the adjuster. The click setting can range from fully open or 1 click, to fully closed or 30 clicks. The compression and rebound shim stack stiffness and preload can be adjusted to control high speed damping.




Figure 2.7. Construction of the main piston (adapted from [9] , used with permission of Randy Lawrence, President, Penske Racing Shocks, 2009)

Now that the shock components have been introduced, shock function will be described. When the shock is extending or in rebound, the piston-shaft assembly moves downward relative to the shock body. This reduces the volume of the rebound chamber, forcing

fluid from the rebound chamber to flow into the compression chamber through the main piston. The fluid has two main paths through the main piston, as shown in Figure 2.8. At low speeds, the fluid will flow through the bleed. As speed increases, the pressure drops across the piston. When the pressure drop exceeds the preload in the rebound shims, the rebound shim stack will open, causing valve flow.




Figure 2.8. Flows through the main piston during rebound:

(a) Low speed bleed flow, (b) High speed valve flow (adapted from [9] , used with permission of Randy Lawrence, President, Penske Racing Shocks, 2009)

Also as the piston-shaft assembly moves downward, the combined volume of the compression and rebound chambers increases due to the removal of the rod. This total volume change must be replaced by fluid from the head valve chamber. Like the main piston, there are also bleed and valve flow for the head valve. The loss of shock oil volume in the head valve chamber is then compensated for by a downward motion of the gas piston and expansion of the gas charge in the gas chamber.

The flow across the main piston bleed and valves causes a pressure drop in the fluid. This pressure drop creates a net force on the piston, which must then be reacted by the rod. This is the main source of the shock force, with smaller contributions due to seal drag, bearing drag, and piston-shaft assembly mass.

2.4 Shock Absorber Modeling

Shocks are complex nonlinear devices whose internal dynamics have been modeled for a variety of applications. Lang developed a comprehensive physical shock model in 1977 [10]. This model included approximately 80 parameters, and is not practical for multiple parametric studies or simulations. Reybrouck created a model with only 20 parameters in 1994 that captured many important shock dynamics [11]. This model relied on semi-empirical coefficients, and did not consider the effect of internal modifications on shock performance. Talbott developed a model of an Öhlins NASCAR Cup shock in 2002 that considered the effect of shock tuning parameters, such as shim stack stiffness, shim stack preload, and bleed orifice area on shock performance [12]. Emmons extended Talbott's model to include the head valve used in current Penske NASCAR Cup shocks [13].

The physical models listed above are useful for predicting the effect of changing shock internals prior to assembling and testing a particular build. However, once a particular shock has been run through a series of tests on the dynamometer, a more efficient empirical model can be made by neglecting internal dynamics and only fitting the input-output characteristics. Compared to physical models, empirical models can be more accurate and can be evaluated faster, which is better for simulation purposes.

Empirical dynamic models for shock absorbers can often be loosely classified into linear, shape function, and universal approximators.

Linear models use linear dynamic systems, such as spring-damper combinations and transfer functions, which are useful for analysis and fast for simulation, but are often very inaccurate [14-17].

Shape function models use prior knowledge of the shock's force-velocity characteristic to specify a shape function, whose parameters reflect the shape of different parts of the curve. These algebraic shape function methods provide decent approximations and some physical insight, but often rely on transcendental functions which take longer to fit and evaluate, which is less suitable for simulation [18, 19].

Universal approximator models apply a family of functions which have the property of being able to approximate any well-behaved function to an arbitrary degree of accuracy. Universal approximators include neural networks, radial basis functions, polynomials, splines, and wavelets. While the universal approximation models theoretically offer great potential in dynamic modeling, the generic model structure often requires a much higher order model compared to an ad-hoc method that takes advantage of knowledge about the system [19-22]. These high-order models cause slower computation speed and longer time to fit the model. Another issue with these models is they have limited ability to predict performance outside the range of the fit data, which may introduce large errors [20].

One important trend that should be used in an efficient shock model is that shocks are quite often dominantly a velocity-dependent device. A simple algebraic shock model is a polynomial curve fit of the force-velocity characteristic. While this model neglects any hysteresis that may be present, it adequately represents the general trend of the force-velocity profile. The current study modifies this algebraic polynomial model slightly, and uses this algebraic function as an input to a nonlinear low-pass filter to better capture the lags present in the shock absorber at higher frequency.

2.5 8-Post Testing

Proper suspension tuning is critical for good performance of a race car. The most valid method to tune suspensions is track testing; however, the amount of available track testing is extremely limited due to regulations and cost. This has created a need for alternate methods for suspension tuning. One of the most popular alternate methods is 8-post testing [2].

During an 8-post test, the tires are supported by 4 actuated wheel platforms, or wheelloaders, that support the vehicle weight and simulate inputs from the track surface. Four actuators, or aeroloaders, are attached to the vehicle's chassis to provide heave force, roll moment, pitch moment, and warp moment which simulate the effect of inertial

and aerodynamic forces present during a track test. The vehicle is shaken on the rig with drivefiles that simulate track conditions or that are used to characterize the vehicle response to more generic inputs such as sine waves or broadband signals. Wheelloaders are typically controlled in position control, while aeroloaders are controlled in force control. The drivefile is repeated for several suspension setups as sensor measurements are recorded and the sensor responses are compared between setups. This information is often used to find or eliminate potential setups prior to track testing or racing. It can also be used to provide sensitivity information to allow the race engineer to make adjustments on the track.

Common sensors during a rig test can be broken into rig sensors and vehicle-specific sensors. The rig sensors available for every test include actuator force, position, and acceleration which can be used to determine control error, tire force, and chassis position. Additional instrumentation may be placed on the vehicle as desired. These measurements typically focus around the suspension, tires, and chassis positions. Common vehicle instrumentation includes suspension displacement; suspension forces such as spring forces, shock forces, and bump stop forces; wheel accelerometers; and chassis positions at critical locations on the car.

Track-based drivefiles are typically generated using an iterative process called drivefile identification or drivefile iteration, which is described by Kelly in [1]. Drivefile iteration is a process that applies sensor data from a track test session to generate an 8-post rig drivefile that represents a particular track. First, the vehicle is tested on the desired track, while desired sensor measurements are recorded. The vehicle is then installed on the 8-post rig with the same sensors and a transfer function H is developed between the rig actuators and the vehicle sensors. The vehicle is then shaken with an initial guess at the drivefile X , and the error is calculated between the rig sensors Y_a and the track sensors Y_d . The initial drivefile is then updated in the frequency domain using

$$X_{update} = X + KH^{-1}(Y_a - Y_d) \quad (2.7)$$

Where K is a step size for the update. The new drivefile is run on the rig, and the drivefile is updated again. This process is repeated until the error reaches an acceptable

level. The identified drivefile is then used to “replay” the track for all future shock setups.

The vehicle response on the rig is analyzed in different ways. The most basic method is to simply look at the sensor time traces, power spectra, or frequency responses and compare them between runs. This will give a detailed picture of the response, and is best used to investigate a specific detail of the vehicle response. Since there are often many sensors and runs, it can be difficult to understand the general trends of the test by looking at time traces.

To get a summary of the trends in the test, the sensor signals are often used to calculate a series of RMS-based response measures. Each response measure is a scalar statistic that is calculated from the data vectors. Two common RMS-based response measures are the “Grip Number”, which is based on the RMS of the total tire force, and the “Aero Number”, which is based on RMS of the ride height variation. The response measures are often broken down by lap markers and can be bandpass filtered to highlight a desired frequency range. Once the response measures are calculated, they can be compared between runs or as a function of the varied suspension parameters.

While this analysis only considers a small portion of the dynamic behavior of the vehicle on the 8-post rig and does not address differences between the vehicle dynamics on the rig and on the track, this method has been used with positive results. Due to these limitations, the interpretation of results from 8-post rig tests has largely been experienced-based, making it prone to error when conditions occur outside the test engineer’s experience base.

While rig testing has become a popular tool in the racing industry in the past few years, even the best teams struggle with how to successfully use rig data to improve lap times [2]. Three of the major issues in 8-post testing today include

1. **Drivefile Development:** How to take track data and use it to quickly generate a drivefile that can accurately simulate the track,

2. **Data Analysis:** How to relate data from a laboratory testing rig with non-rolling tires, sitting on teflon pads, with how the vehicle will respond on a race track, and
3. **Experimental Design:** How to run the experiments to efficiently determine the best setups and the sensitivity of the response to different setup parameters.

This study focuses on issue 3, which has been discussed in Chapter 1. Another critical area of research in 8-post rig testing is dealing with inconsistencies between the vehicle excitation and behavior between the track and laboratory environments which may lead to false trends observed on the 8-post rig. A comparison of some of these differences is summarized in Table 2.1.

Table 2.1. Comparison of vehicle excitation and behavior on track and on 8-post rig

	On Track	On 8-Post
Road Input at time t	A distribution of heights that interacts with the tire contact patch	A single "equivalent" wheelloader displacement
Aerodynamic Forces	A function of the chassis position, which changes with setup	Determined once during drivefile iteration
Driving Line - Road under tires - Acceleration	Driver will seek ideal driving line, which depends on setup	Constant
Tire Dynamics	Rolling, stressed, heated On asphalt or concrete	Stationary on Teflon pad
Anti-Features / Jacking Forces	Braking, acceleration, and cornering interact with vertical motion due to suspension geometry	Lateral/longitudinal acceleration minimal Brakes and driving torque not applied

Due to the highly competitive nature of racing, very little of the development work done in 7-post testing has been published. There have been 3 technical papers with data from a 7-post rig test that are applicable to this study, each giving a small piece of information about how 7-post rigs are used for suspension development. Since there is so little published in this area, the key details of these three papers are shown below.

Kelly et al. discusses the process they apply to perform drivefile iteration on the 7-post rig at the Automotive Research Center (ARC) in Indianapolis [1]. This process is described above. This paper performed drivefile iteration on a Champ car using track data from Sebring raceway. The wheel hub accelerometers were bandpass filtered from 1-25 Hz to iterate the wheelloader drivefile, and the lateral acceleration, longitudinal

acceleration, and vehicle speed were used to determine the aerolader drivefile. The drivefile error converged from 40% RMS error to 20% RMS error in 8 iterations.

The paper suggests analysis of suspension setups using both frequency response functions and drivefile playback. Swept sine testing can be used to make sure the damping levels are acceptable for the stiffness levels. Track simulation can be used for further investigation of how the car will respond to specific track events.

Miller reported on a 7-post test session with Clemson University's Formula SAE car on Öhlins' 7-post rig [23]. This test session conducted two studies, first to tune the anti-roll bars (ARBs), and second to tune the shocks. Initial testing indicated a large amount of friction in the pushrod, bellcrank, and A-arm bearings. To tune the ARBs, a drivefile was used which slowly rolled the chassis to the right using the aeroloaders. The front and rear ARBs were varied in a 3x3 matrix, and the load transfer distribution was used to determine the best ARB settings. Once the best ARB settings were determined, they were used for all future tests. Next, another drivefile was used to tune the shock compression adjustment, by cycling the wheelloaders in heave at 10 in/s peak velocity, while the aeroloaders were used to simulate 1.4 g cornering acceleration and 0.2 g longitudinal acceleration. The shock compression adjusters for the front and rear were also varied in a 3x3 matrix in low, medium, and high settings. The data was used by the Öhlins engineers to calculate response measures for each test. The post-processing steps used to calculate these measures from rig data were not described. These measures included mean left grip disturbance, mean right grip disturbance, mean overall grip disturbance, roll plot, LF grip disturbance, RF grip disturbance, LR grip disturbance, RR grip disturbance, body heave, body pitch, and body roll. These response measures were used to create contour plots for the ARB matrix and the shock matrix, which provides a summary of the sensitivity to the adjustments.

Kowalczyk uses a linear 7-degree-of-freedom vehicle model to simulate the vehicle response on the rig for the purpose of helping focus rig data analysis, in particular to identify tradeoffs and determine realistic limitations of tuning [24]. A chirp drivefile was used for the wheelloaders in heave and pitch, with the amplitude envelope adjusted to

match track data power spectral density (PSD). Contour plots were generated for the peak of the power spectrum for heave and pitch acceleration for different front and rear shock settings. A tradeoff between optimizing heave acceleration and pitch acceleration was shown, which forces the race engineer to decide if the potential benefit in improving one metric outweighs the other. They chose to minimize pitch variations, which would reduce the variations in the distribution of normal force variation front to rear. The power spectrum from rig data was then presented which showed a significant improvement in pitch response, with only a slight degradation in heave response.

Seven-post testing for motorsports applications originated from 4-post testing which was primarily used for comfort and noise vibration harshness (NVH) applications. One of the earliest developments for motorsports testing was to put the car on a 4-post rig and attach soft springs to the chassis. These springs were preloaded to simulate a DC aerodynamic downforce. In 1994, Kasprzak performed testing on one of these modified 4-post rigs and compared it to simulation data [7]. The simulations first used a 4th-order linear half-car model, and later substituted the linear shocks with a lookup table using shock dynamometer data from a sine test. Kasprzak states that the model trends show good correlation with experimental data. Once promising setups were found using the simulations, they could be tested on the rig and the setup could be fine-tuned. They used this process for a race team, and the team reported the setup was better.

The last two papers by Kowalczyk and Kasprzak have suggested that linear models have been successfully applied to predict trends and accelerate setup tuning on 7-post rig tests. Linear models are also used in an iterative process to develop drivefiles. These trends can be used to select potential setups to be tested on the rig, identify parameter sensitivity for tuning, and show tradeoffs and realistic limitations to suspension tuning. Unfortunately, both of these papers presented limited rig data and did not quantify the accuracy of their linear models.

Another thing that was not described in these papers is how they developed their vehicle models. Having an accurate vehicle model and knowing the limitations of the models can be critical to successfully applying simulation to help rig testing.

Chan introduced a 12 degree-of-freedom nonlinear vehicle model for virtual 7-post testing, which includes tire separation and large rotations, for use in future 7-post research, including identification studies [25]. While this model accounts for things that are not possible with a linear model, it does not account for linear chassis compliance that is present in most oval and road racing cars, and it does not account for the coupling present in a solid-axle suspension.

2.6 The Role of Experimental Data in Vehicle Modeling

To create an accurate model with known performance limitations, it helpful to use experimental data to validate the model accuracy. If desired, the model can also be updated to better match the experimental data. This process of model validation and updating can be performed on the component level, the system level, or both depending on the application. Some relevant examples showing trends in applying experimental data to model validation and updating are shown below.

In the automotive industry, vehicle models are quite commonly used early in the development cycle when vehicle test data is not available. This makes model updating using vehicle system data difficult, so the focus is often on integrating validated component models. Vilela correlated proving ground data collected for a selection of different proving ground tests with a detailed ADAMS model assembled using component models [26]. The correlation was determined to be acceptable, indicating that the model could have been used before the vehicle was manufactured.

Often, system level data for an existing vehicle is used to refine an existing vehicle model, which will be used in future studies. The experimental data is often assumed to be the ideal response for the model to try to achieve, although experimental errors can make this a poor assumption. One basic method to use test data to update models is to first perform a parametric study to determine the sensitivity of error to model parameters. The model parameters can then be adjusted using the sensitivity information to reduce modeling error [27].

A more systematic method for model updating is to apply optimization methods to find parameter values that minimize the model error. Hu performed model validation and updating on a half-car ride model using proving ground data [28]. The initial guess at the parameters was defined using component tests, which were then updated to better match the proving ground data using optimization of modeling error.

While many studies use proving ground data for model development, there has been an increasing trend to use laboratory test data since the experiments are run in a more controlled environment with higher repeatability [29]. For example, Cheli and Sabbioni performed 4-post testing on a van to identify the unknown stiffness, damping, and inertial parameters [30]. The parameters were optimized to best fit the experimental data using nonlinear optimization. Further data was used to show that there was acceptable correlation between the model and data. The ability of the model to predict the effect of changing parameters was not discussed. Andersen performed parameter identification for a nonlinear multibody dynamics model of a McPherson strut suspension using laboratory data from a quarter-car rig [31].

Most model updating studies attempt to match the displacements and/or accelerations of the sprung and unsprung masses. One additional sensor measurement that is important for an 8-post model to predict is the tire normal force. Ziegenmeyer optimized the parameters of a linear quarter-car model to minimize the prediction error for a quarter-car rig test on a quarter-car rig which included real suspension components and a tire sitting on a wheel platen [32]. The model was then used to predict the accelerations of the sprung and unsprung masses and the tire normal force with reasonable accuracy.

The above studies use optimization methods to update model parameters that are highly dependent on gradient estimates of the error as a function of model parameters. Depending on the model structure, this error function can have local minima, and the optimization method may converge to these local minima. This can be further complicated with experimental data with noise and nonlinearities. The success of these optimization methods can also be dependent on a good initial guess, which requires accurate component models. Also, gradient methods are not applicable when the

parameters can only be chosen from a discrete space. Due to these issues, alternative to gradient-based optimization is often desired.

For example, Alasty generated target data using an ADAMS simulation and applied genetic algorithms to find a global minimum to the error function of a nonlinear seven degree of freedom model [33]. He also discusses identifiability conditions for nonlinear models about an operating point.

The above studies all use structural models (both linear and nonlinear), where the models have been constructed using components that have a specific physical meaning. Many applications do not require accurate estimation of the structural parameters, only a model of the input-output behavior. If this is the case, a more generic linear or nonlinear dynamic model structure can be used, which has more modeling freedom and can be identified using generic system identification tools. By using a generic input-output model structure, there is no need for prior knowledge of the geometry of how components in the physical system are arranged or the boundary conditions for their connections.

For example, if the system operates within a range in which the dynamics are linear, the system may be represented by many different linear dynamic models. General linear dynamic models can be characterized using frequency responses or impulse response, which can be identified easily using least-squares regression techniques. If a transfer function or state-space model is required, a variety of system identification methods can be applied to fit these types of models [34].

Chapter 3

Shock Absorber Modeling

This chapter develops a physically-motivated nonlinear dynamic shock absorber model that can be quickly fit to experimental data and implemented in simulation studies. This model is based on the understanding that the shock is dominantly a velocity-dependent device, with lag due to compressibility effects.

The model consists of an algebraic backbone, which is a function of velocity alone, and a nonlinear low-pass filter, which has been designed based on the observation that shocks often exhibit less hysteresis at higher velocities. Due to the simplicity of the model, it can be fitted with data and evaluated quickly.

The next section presents a simple physical shock absorber representation that motivates several shock models. This is followed by a description of the experimental procedure for shock dynamometer testing. These models are then fitted using the experimental data for 12 different shock configurations. After the models have been fitted to the experimental data, they are then validated using both broadband and sine wave inputs. This analysis allows us to specify the conditions under which the model can accurately predict force and where the model may need some improvement.

3.1 Physical Motivation

This section describes the physical motivation for the shock models presented later in this chapter. First, a simple linear model for the internal flows inside a shock absorber is presented. Using this model, both the incompressible and compressible equations are derived. The fluid dynamic equations derived for the incompressible and compressible cases are then represented as mechanical system equivalents, which is a more intuitive form for a component model in a mechanical system.

The equivalent mechanical system is found to be a damper in series with one spring, representing compressibility effects, and in parallel with another spring, which represents the resistance from the gas charge to the rod entering the rebound chamber.

3.1.1 Physical Model of the Shock Absorber

A simplified construction of a shock absorber is shown in Figure 3.1. The construction includes a shock body, a piston-rod assembly, and a floating piston. The piston separates the compression and rebound chambers which are both filled with hydraulic fluid, while the floating piston separates the fluid in the compression chamber from the nitrogen charge in the nitrogen chamber. When compressed, the piston-rod assembly moves up into the shock body, which changes the volume of the compression and rebound chamber. Since the rod must be inserted into the shock body, the fluid displaced by the rod insertion must be compensated for by motion of the floating piston. The change in volume in the rebound chamber causes flow through the bleed holes and shims stacks in the main piston. This flow creates a resistance to motion that is a function of velocity.




Figure 3.1. Basic shock absorber construction

Another method of visualizing the fluid circuit inside the shock absorber is the analogous fluid circuit shown in Figure 3.2. This explicitly shows each effect as an absolute fluid flow, which allows the engineer to visualize the shock flows and gain insight on the

equations governing shock function. If inertial effects are neglected, only relative flow rates matter. To simplify the fluid system, first we fix the main piston in the shock body. We then represent the effect of moving the piston in the shock body by changing the volumes of the compression and rebound chambers the same amount that they are changed by moving the piston-rod assembly in the shock. The effect of the nitrogen reservoir can be modeled as a spring-loaded accumulator.

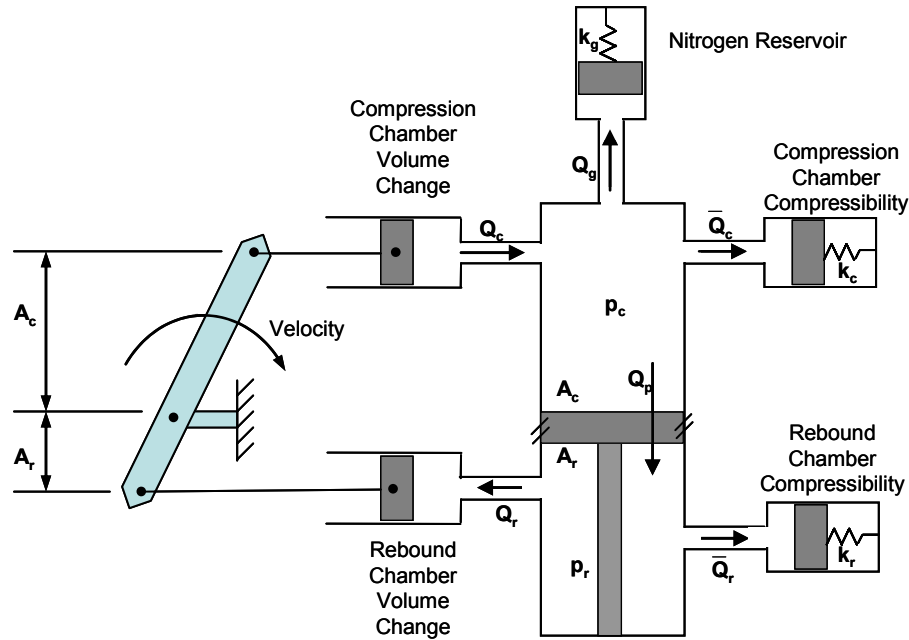


Figure 3.2. Equivalent fluid system for shock absorber

When the fluid is incompressible, conservation of mass reduces to conservation of volumetric flow rate. Conservation of volumetric flow rate is easy to visualize in a fluid circuit, as the net flow for any control volume must be zero. When the fluid is compressible, conservation of volumetric flow rate no longer applies. To approximate the incompressible system using a simplified incompressible system, the error between conservation of volume and conservation of mass can be approximated using an equivalent flow to an accumulator. This model is consistent with a compressible fluid with constant bulk modulus. The equivalent flow due to compression can be thought of as fluid volume and energy storage due to fluid compression. This allows us to treat a more complex compressible system like a simple incompressible system and view fluid compression as a flow compressing a spring.

The rate of change of volume for the compression and rebound chambers is

$$Q_c = A_c v \quad (3.1)$$

$$Q_r = A_r v \quad (3.2)$$

where v is the rod velocity, Q_c and Q_r are the rate of volume change in the compression and rebound chambers, and A_c and A_r are the areas on the main piston on the compression and rebound side. In Figure 3.2, the changes in volumes are caused by controlling two pistons, which displace the correct amount of fluid. When the piston on the compression side causes a flow to the compression chamber, the flow must displace the floating piston, pass through the main piston, or go to the compression chamber compressibility. Similarly, as the rebound chamber expands, flow must come from the main piston or from fluid expansion. This gives continuity equations for the compression and rebound chambers of

$$Q_c = Q_g + Q_p + \bar{Q}_c \quad (3.3)$$

$$Q_p = Q_r + \bar{Q}_r \quad (3.4)$$

Where \bar{Q}_c and \bar{Q}_r are the effective compressive flows for the compression and rebound chambers, while Q_p is the flow through the main piston and Q_g is the flow that displaces the floating piston.

For the purpose of analysis, we will model the accumulators as linear springs and the piston flow as a linear flow resistance. The gas charge will behave like a linear spring for small deflections or low gas pressure, while the fluid compressibility will act like a linear spring if the bulk modulus is constant. The piston flow resistance for a constant area orifice can be modeled as linear for laminar flow or small velocities. Linear flow characteristics can also be achieved in other conditions by the opening of the shim stacks, which effectively increases the orifice area with velocity. We will extend the results to motivate the development of nonlinear shock models in a later section. The constitutive equations for the accumulators and main piston are

$$\text{Compression Chamber: } \bar{Q}_c = \frac{1}{k_c} \dot{p}_c \quad (3.5)$$

$$\text{Rebound Chamber: } \bar{Q}_r = \frac{1}{k_r} \dot{p}_r \quad (3.6)$$

$$\text{Gas Chamber: } Q_g = \frac{1}{k_g} \dot{p}_c \quad (3.7)$$

$$\text{Main Piston: } p_c - p_r = CQ_p \quad (3.8)$$

where p_c and p_r are the pressures in the compression and rebound chambers, k_c and k_r are the stiffnesses for the compression and rebound chamber compressibility, k_g is the stiffness of the gas charge, and C is the flow resistance for the main piston.

The shock force can be found by drawing a free-body diagram of the piston-rod assembly. This yields the shock compression force as

$$F = p_c A_c - p_r A_r \quad (3.9)$$

This model neglects losses due to piston friction and also kinetic energy of the fluid and shock components. These equations are further developed for the incompressible and compressible cases in the following sections.

3.1.2 Incompressible Model

When the fluid is incompressible, the effective compression and rebound flows due to compression are zero. Combining Equations (3.1) – (3.8) yields

$$A_c v = \frac{1}{k_g} \dot{p}_c + \frac{1}{C} (p_c - p_r) \quad (3.10)$$

$$\frac{1}{C} (p_c - p_r) = A_r v \quad (3.11)$$

Substituting (3.11) into (3.10) and integrating allows us to solve for the compression pressure

$$p_c = k_g A_{rod} x + p_{co} \quad (3.12)$$

where p_{co} is the initial fluid pressure, x is the piston rod displacement, and A_{rod} is the area of the rod. Substituting the compression pressure into Equation (3.11) yields the rebound chamber pressure

$$p_r = p_c - CA_r v \quad (3.13)$$

Now that we have the two chamber pressures, we can use them to find the shock force

$$\begin{aligned} F &= (CA_r^2)v + (k_g A_{rod}^2)x + p_{co} A_{rod} \\ &= C_{inc} v + K_{inc} x + F_o \end{aligned} \quad (3.14)$$

where C_{inc} , K_{inc} , F_o are constants that define the shock model's damping, stiffness, and nominal force properties at the system level. This shows that the incompressible model force can be represented as a linear damper in parallel with a linear spring and an offset force.

3.1.3 Compressible Model

Combining Equations (3.1) – (3.8) yields

$$\begin{aligned} \left(\frac{1}{k_g'}\right) \dot{p}_c &= -\frac{1}{C}(p_c - p_r) + A_c v \\ \left(\frac{1}{k_r}\right) \dot{p}_r &= \frac{1}{C}(p_c - p_r) - A_r v \end{aligned} \quad (3.15)$$

Where k_g' is stiffness of the gas spring and the compression chamber spring mechanically in series. While k_g' will usually be finite due to the gas charge, k_r may be infinite if the rebound chamber is incompressible. Assuming that both k_g' and k_r are finite, we may solve for the pressure change rates and add Equation (3.9) to form the state-space model

$$\begin{bmatrix} \dot{p}_c \\ \dot{p}_r \end{bmatrix} = \begin{bmatrix} -\frac{k_g'}{C} & \frac{k_g'}{C} \\ \frac{k_r}{C} & -\frac{k_r}{C} \end{bmatrix} \begin{bmatrix} p_c \\ p_r \end{bmatrix} + \begin{bmatrix} k_g' A_c \\ -k_r A_r \end{bmatrix} v \quad (3.16)$$

$$F = [A_c \quad -A_r] \begin{bmatrix} p_c \\ p_r \end{bmatrix}$$

Rewriting the model as a transfer function yields

$$\frac{F}{V}(s) = \frac{A_{rod}^2 \left(\frac{k_g' k_r}{k_g' + k_r} \right)}{s} + \frac{\left(\frac{C}{k_g' + k_r} \right) \left[A_c^2 k_g' + A_r^2 k_r + A_{rod}^2 \left(\frac{k_g' k_r}{k_g' + k_r} \right) \right]}{\left(\frac{C}{k_g' + k_r} \right) s + 1} \quad (3.17)$$

This equation can be further simplified by assuming the gas stiffness is much softer than the rebound and compression stiffnesses and that the compression, rebound, and rod areas are similar in magnitude. This assumption simplifies Equation (3.17) to

$$\begin{aligned} F &= \left[\frac{1}{\left(\frac{C}{k_r} \right) s + 1} \right] (C A_r^2) v + (k_g A_{rod}^2) x + F_o \\ &= \left(\frac{1}{\tau s + 1} \right) C_{inc} v + K_{inc} x + F_o \end{aligned} \quad (3.18)$$

Where C_{inc} , K_{inc} , and F_o are the damping rate, spring rate, and initial force found for the incompressible case. This shows that the only difference between the incompressible and the compressible models is a first-order low-pass filter on the shock force. The lag, τ , is caused by the compressibility in the rebound chamber. This provides the motivation for the shock models that are developed in Section 3.15.

3.1.4 Equivalence to Mechanical Systems

Shocks are often modeled as an ideal dashpot in series and/or in parallel with springs. Modeling a shock absorber using mechanical elements is a natural fit when the shock is a component in a mechanical system. In this section we show that the incompressible model is equivalent to a damper in parallel with a spring and in the compressible case, the model is equivalent to a damper in series and in parallel with springs.

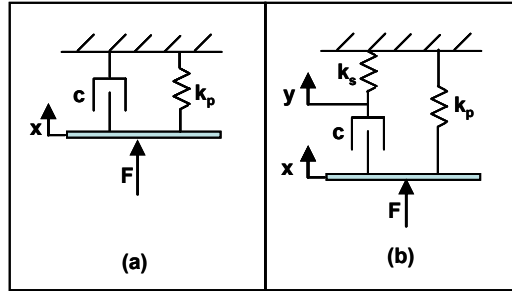


Figure 3.3. Mechanical equivalents: (a) Damper in series with a spring, (b) Damper in series and in parallel with springs

The equation for a damper in parallel with a spring as shown in Figure 3.3a is

$$F = c\dot{x} + k_p x \quad (3.19)$$

This is in direct agreement with the incompressible model of Equation (3.14), with

$$\begin{aligned} c &= CA_r^2 \\ k_p &= k_g A_{rod}^2 \end{aligned} \quad (3.20)$$

The equation for a damper in series and in parallel with springs as shown in Figure 3.3b is

$$F = \left(\frac{c}{\left(\frac{c}{k_s}\right)s + 1} \right) \dot{x} + k_p x \quad (3.21)$$

This is in direct agreement with the compressible model of Equation (3.18), where

$$\begin{aligned} k_p &= k_g A_{rod}^2 \\ c &= CA_r^2 \\ k_s &= k_r A_r^2 \end{aligned} \quad (3.22)$$

The previous sections have derived a linear compressible shock absorber model. This work will next be applied to develop models for real shocks.

3.1.5 Extension of Physical Motivation to Shock Models

While the above models are linear, shock absorbers are inherently nonlinear devices. This section extends the above linear discussion to modeling real shocks. A common nonlinear shock model is a simple curve fit of the force as a function of velocity

$$F_s = f(v) \quad (3.23)$$

The curve fit is often a polynomial, but some studies have used transcendental shape functions. An example force-velocity plot and a curve fit model are shown in Figure 3.4. While algebraic curve fits do not capture frequency-dependency, they capture the general shape of the force-velocity response and are an excellent starting point for shock modeling.




Figure 3.4. Sample shock dynamometer data and curve fit

The first approach that was tried for an algebraic model was a polynomial curve fit. Since the force-velocity plot often has a complex shape, a high-order polynomial was used before increasing the polynomial order produced no noticeable reduction in modeling error. Because a high-order polynomial was required to capture the complex shape, significant errors occur if the model is used to predict a force outside of the fit data. For one model, a velocity less than 5% outside of the range of the fit data caused more than 10,000 lb of error, as shown in Figure 3.5a.

While a high-order polynomial is useful to drive down the error, it cannot be used for extrapolation. Further examination of multiple shock curves yielded the observation that the shape of the shock curve at high velocity is often close to linear and the high-order polynomial is only needed at low velocity. This observation prompted the modification to fit a high order polynomial at low velocities and fit a low order polynomial at high velocity compression and another low-order polynomial at high velocity rebound, as shown in Figure 3.5b. In the transition regions between low and high velocities, the two polynomials are linearly combined to provide a continuous transition. We label this algebraic shock model as “Polylinear”.




Figure 3.5. (a) Error in extrapolating a high-order polynomial fit, (b) Regions for polynomials in Polylinear model

While the algebraic curve fit model typically captures the general trend of the response, it cannot capture the dynamic behavior. This is particularly important in shocks where the hysteresis is significant. For example, the shock data in Figure 3.4 shows almost 300 lb hysteresis, which is 20% of the 1500 lb total force range. If we want an accurate simulation, we need to capture the shock’s dynamics.

In the previous sections, we have established that the behavior of the linear compressible shock absorber model shown in Figure 3.2 was equivalent to a linear dashpot with springs both in series and in parallel. Since race shocks often run low gas spring rates, we will drop the gas stiffness term for now. The model in differential equation form is

$$\dot{F} = \frac{1}{\tau}(-F + F_s) \quad (3.24)$$

where F_s is the algebraic force the shock would achieve if there was no compliance in series or at low frequency. The term $\tau = \frac{c}{k_s}$ acts as a system time constant. We can replace the proportional shock force $F_s = c\dot{x}$ with the nonlinear algebraic shock force calculated by the curve fit model in Equation (3.23). The resulting dynamic shock model takes on the structure shown in Figure 3.6. This model includes a linear first-order differential equation, and a nonlinearity on the velocity input. We label this shock model as “LPF1”.




Figure 3.6. LPF1 shock absorber model

The simple first-order lag model may work reasonably well for some shocks, but many shocks only exhibit hysteresis at lower velocities when the flow is changing directions, pressures in the compression and rebound chambers are just starting to develop, and the valves are starting to open. At higher velocities, the valves are fully open and the fluid has been compressed, so the force-velocity plots often exhibit very little hysteresis and can be accurately modeled using an algebraic model. When using the LPF1 model, the lag causes hysteresis at low velocity where it is needed, but it also causes lag at higher velocities where zero lag is best.

This phenomenon may be explained by extending our physical shock model to include fluid stiffness that changes as a function of velocity. We then find that the time constant for the linear model is $\tau = \frac{c}{k_s} = \frac{c}{k_r}$. High lag at low velocities and no lag at high velocity suggest that the fluid stiffens at higher velocity and behaves more like an incompressible fluid, as shown in Figure 3.7. At low velocities, the fluid has not changed volume much and the fluid has a nominal compressibility, giving the shock a nominal time constant. At

higher velocities, the fluid has already been compressed and is almost incompressible, making the time constant almost zero. This suggests that the fluid compressibility is more accurately represented as a stiffening spring, not a linear spring.




Figure 3.7. The effect of a stiffening fluid on time constant:

(a) Fluid stiffness, (b) Time constant

It would seem that all that would be needed to model this behavior is to make the time constant in Equation (3.24) a function of velocity, with a curve similar to the one shown in Figure 3.7b. While this is possible if there is significant lag at high velocity, many shocks show very little hysteresis at high velocity. A small time constant at high velocities would require small time steps for accurate numerical integration. This will cause slow simulations, so it is not an acceptable solution.

A more efficient way to get high lag at low velocity and no lag at high velocity is to use the lagged force at low velocity and use direct feedthrough of the algebraic force at high velocity. The model equations are

$$\begin{aligned} \dot{F}_{lag} &= \frac{1}{\tau} (F_s - F_{lag}) \\ F &= \alpha(F_s) F_{lag} + (1 - \alpha(F_s)) F_s \end{aligned} \quad (3.25)$$

where the weighting function α is unity at low velocities, zero at high velocities, and a linear combination of the two between low and high velocities, as shown in Figure 3.8.

This model consists of a linear first-order differential equation and nonlinearities at both the input and output. We label this model “LPNL2”.

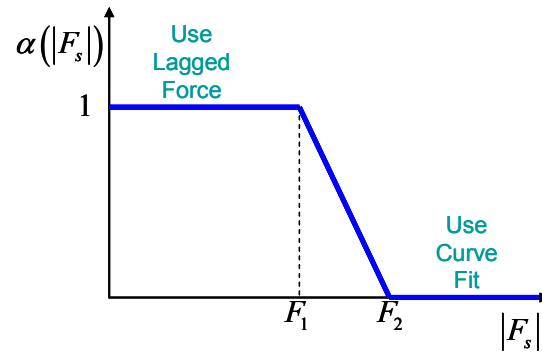


Figure 3.8. Output weighting for LPNL2 model

3.1.6 Summary

The purpose of this section was to motivate the development of physically-based shock absorber models. We started by developing a linear compressible shock absorber model. The equations were solved for both the incompressible and compressible cases. Next, mechanical equivalents were found for both the incompressible and compressible models. In the case of the incompressible model, the mechanical equivalent is a damper in parallel with a spring that represents the gas charge. In the case of the compressible model, the mechanical equivalent is a damper in series with a spring that represents the rebound chamber compressibility and in parallel with a spring that represents the gas charge.

The physical insight gained by developing the linear models was then extended to define physically-motivated shock absorber models. While the incompressible linear shock model was a linear function of velocity, real shocks are nonlinear. To capture the shock force as a nonlinear function of velocity, we defined a Polylinear model, which curve fits a high-order polynomial at low velocities and a low-order polynomial at high velocities. This was motivated by the fact that while the shock curve may have a complex shape at lower velocity, the shock often exhibits near-linear behavior at higher velocities. Having a low-order curve fit at high velocities also has the advantage of being well-suited for extrapolation to velocities beyond the original curve fit data without significant error.

While a curve fit of the force-velocity plot captures much of the shock response, shocks often exhibit significant hysteresis which must be captured in any accurate shock model. Motivated by our linear compressible shock model, our first dynamic shock model (LPF1) was just the curve fit filtered by a first-order lowpass filter.

The observation was then made that shocks typically only exhibit hysteresis at low velocities, while at higher velocities there is almost no hysteresis to the force-velocity curve. Unfortunately, the LPF1 model adds thickness at all velocities. This motivated a dynamic shock absorber model (LPNL2) which used the lagged force of LPF1 at low velocities and the Polylinear curve fit at high velocities.

While this model was developed for a shock absorber commonly used in motorsports applications, the minimal assumptions that led us to arrive at this model structure allow it to be applied to a wider class of dynamic systems. First, we assume the force output is dominantly an algebraic function of the velocity input. Next, the high velocity behavior can be described by a low-order polynomial. The dynamic behavior can be approximated by a nonlinear stiffening spring in series with a damper, producing lag and attenuation at low velocity, but not at high velocity, where the spring is effectively rigid. In our case, the stiffening effect is a result of a fluid that has a stiffening bulk modulus, but it could be used to describe other stiffening compliances in series with a damper, such as compliant rubber bushings. It may also be possible to accommodate other dynamic behavior by changing the weighting function.

To fit the three model structures developed in this section to the desired shocks, the shocks are first tested on a shock dynamometer and then the data is used to fit the models and to validate that the model accurately represents the shock performance under different conditions than the fit data to avoid overfitting. These topics are covered in the following sections.

3.2 Experimental Setup

This section describes the shock testing conducted to collect the data necessary to fit the shock models described in the previous section and to validate the model's ability to predict forces. We start by describing the Roehrig EMA shock dynamometer used for testing. We continue by describing the 12 different shock configurations we tested. We conclude by describing the drivefiles that were used for model fitting and validation of the shock models.

3.2.1 Roehrig EMA Dynamometer

To develop an empirical shock absorber model, a Penske 7300 NASCAR Nextel Cup shock absorber was tested on a Roehrig Electromagnetic Actuator (EMA) shock dynamometer, as shown in Figure 3.9. This dynamometer uses an electromagnetic actuator to actuate the shock. This allows a wide variety of inputs, including sine waves, triangle waves, random inputs, track data, or any user-defined input.



Figure 3.9. Penske 7300 shock absorber on a Roehrig EMA (photo by author, 2009)

The Roehrig EMA consists of a load cell located at the top of the dynamometer, an electromagnetic actuator located at the base of the machine, and an infrared thermocouple

to measure shock body temperature. The system is driven by two electromagnetic actuators that act on a ram connected to the lower shock mount. The Roehrig EMA is capable of forces up to 2000 lb, 8 inches of stroke, excellent frequency response up to 80 Hertz, and speeds of 100 in/s [35].

Tests on the EMA are controlled using Roehrig's Shock software. This software allows test definition from both standard waveforms and user-defined inputs. Test signals are defined about a user-defined DC offset, which was set to 3 inches in this study to put the shock absorber near the middle of its range of travel. In addition, Shock allows a warm-up sequence to be run prior to each test session, where the shock is run until it reaches a specified temperature. This allows for more uniform testing conditions. For this study, the shock was warmed up to 110°F prior to each test. During each test, measurements of shock position, velocity, force, and temperature are recorded at 2,000 Hz. At the end of each experiment, this data is saved for further analysis.

3.2.2 Shock Configurations

As discussed in Chapter 2, each Penske 7300 shock is designed to be highly adjustable and can be rebuilt with different internal components, giving the shock engineer flexibility to rebuild the shock and tune the shock characteristics to a desirable setting. The main internal adjustments are the compression shims, rebound shims, and bleed jet on the main piston and the head valve. The external adjustments are the bleed adjuster and gas pressure.

To investigate the effectiveness of the shock models for a range of shock settings, we made 3 different shock builds. The first build was used as a baseline. The second build has a softer compression shim stack on the head valve. The third build has a softer compression shim stack on the main piston. Each build was tested at 50 psi and 150 psi gas pressure and at a low setting (1 click) and a medium setting (14 clicks) for the bleed adjuster. A force-velocity plot for the three builds at 50 psi and low bleed setting are shown in Figure 3.10. A summary of the 12 different configurations is shown in Table 3.1.




Figure 3.10. Force-velocity for the three shock builds tested:

(a) Build 1, (b) Build 2, (c) Build 4

Table 3.1. Shock configurations tested

Build Name	Description	Bleed Adjuster	Gas Pressure
Build 1	Baseline	1 click	50 psi
Build 1	Baseline	14 clicks	50 psi
Build 1	Baseline	1 click	150 psi
Build 1	Baseline	14 clicks	150 psi
Build 2	Softer HV	1 click	50 psi
Build 2	Softer HV	14 clicks	50 psi
Build 2	Softer HV	1 click	150 psi
Build 2	Softer HV	14 clicks	150 psi
Build 4	Softer HV	1 click	50 psi
Build 4	Softer HV	14 clicks	50 psi
Build 4	Softer HV	1 click	150 psi
Build 4	Softer HV	14 clicks	150 psi

3.2.3 Drivefiles

Each of the shock configurations described above was tested using a variety of drivefiles on the shock dynamometer. For this study, we used random signals for model fitting and validation purposes. A random drivefile was generated using a multisine waveform by first defining the shape of the relative amplitude spectrum of the position command at 10 different points between 0.1 and 20 Hz, as shown in Table 3.2. This shape was chosen to excite the frequencies of interest while not introducing unrealistic velocities into the shock absorber.

Table 3.2. Definition of relative amplitude spectrum

Frequency Hz	Amplitude Relative
0.1	1
1	1
2	0.8
3	0.7
4	0.5
5	0.3
8	0.2
10	0.1
20	0.05
30	0.01

To create a more evenly distributed amplitude spectrum, the spectrum defined in Table 3.2 was interpolated at every 0.1 Hz to define the amplitude spectrum at 300 uniformly spaced points. The unscaled time signal was then created by summing the amplitude spectrum

$$x_{unscaled}(t) = \sum_i A_i \sin(2\pi f_i t + \phi_i) \quad (3.26)$$

where A_i is the value of the unscaled amplitude spectrum at frequency f_i . The phase ϕ_i at each frequency is assigned by MATLAB's uniformly distributed random number generator, *rand*, over the range 0 to 360 degrees. This was used to create a 30 second random signal. Multiple realizations of the random signal were calculated and the one with the lowest crest factor was chosen to avoid extreme peaks in the drivefile. The signal is then rescaled to have the desired peak amplitude on the shock dynamometer. For our drivefile, the peak amplitude was selected to be 1 inch. This drivefile is then saved for import to the Roehrig shock dynamometer. A 5-second segment of the drivefile is shown in Figure 3.11.




Figure 3.11. Segment of the random drivefile

In addition to the random drivefile described above, we also tested each of the 12 shock settings with sine waves over a wide range of amplitudes and frequencies, as shown in Table 3.3. The purpose of the sine wave tests in this study is to serve as a validation data set for the model. This method will allow us to see at what amplitudes and frequencies the model performs well, and where the model could use some improvement [36]. This information can be quite useful when deciding if a shock model is suitable for a particular simulation application.

In addition to the random and sine tests that were performed for each shock configuration, a static test was performed for each combination of build and gas pressure. The static profile is shown in Figure 3.12. It starts at an initial test offset of 2 inches and stays there for 20 seconds. It then moves up one inch at 0.2 in/s and then maintains the new position for 20 seconds. This process is repeated until the shock is moved up 3 inches, at which point the process is repeated in the downward direction until the shock reaches the initial test offset of 2 inches. The static test gives a spring rating of the shock at three positions and in two directions, which gives us both the force-deflection curve and also the amount of hysteresis in that curve caused by directional effects such as seal drag.

Each of the three builds were also tested using a 0.5 inch bump test, with the drivefile shown in Figure 3.13. The transition speed of the bump was 30 in/s.

Table 3.3. Summary of sine wave tests

Amplitude (in)	Velocity (in/s)	Frequency (Hz)
1	1	0.16
1	3	0.48
1	6	0.95
1	10	1.59
1	20	3.18
1	30	4.77
0.5	5	1.59
0.5	10	3.18
0.5	20	6.37
0.25	10	6.37
0.25	20	12.73
0.1	2.5	3.98
0.1	5	7.96
0.1	10	15.92




Figure 3.12. Static test drivefile




Figure 3.13. Bump test drivefile

3.2.4 Summary

This section has presented the experimental procedures used to collect the data that will be used to fit the shock models described in Section 3.1 and to validate their ability to predict force under different operating conditions.

Twelve different shock configurations are tested to investigate the effect of head valve shims, compression shims, bleed adjuster, and gas pressure on the model accuracy. Each configuration is tested with a random drivefile for fitting and validation. Each shock configuration is also tested using multiple sine waves for validation. Each combination of build and gas pressure is tested using a static test to determine the shock's dependence on position and how that changes with the direction that the position was approached. Each build was tested using a bump input with transition velocity of 30 in/s to test the model accuracy during a bump event.

3.3 Shock Absorber Modeling

The previous section described three different shock models and our experimental procedure for collecting data to fit and validate the models. This section describes our efforts to fit the models, while the next section covers the model validation. While each of the 12 shock configurations is fit to all three models, we only present the details for fitting shock build 1 with 50 psi gas and 1 click rebound in this section for brevity. We will summarize the key results for the other fits where informative.

The data used for fitting the models was from seconds 1-10 from the random test. All shock dynamometer data is filtered using an 8th order Butterworth lowpass filter with cutoff of 50 Hz using MATLAB's *filtfilt* command, which provides zero phase shift and twice the attenuation by filtering the data forwards and then backwards.

3.3.1 Algebraic Modeling

The Polylinear model presented in Section 3.1.5 was fit using seconds 1-10 of the 30-second random test. The resulting force-velocity plot is shown in Figure 3.14 and a time

plot of the forces and error is shown in Figure 3.15. The model produced good fits with a 15th order polynomial at low velocity and a 3rd order polynomial at higher velocities. To ensure that this polynomial model was not overfitting the fit data, it was also verified that the error in the validation data was also being reduced as model order increased. The transition from low to high velocity was set at 15 in/s and the width of the transition was set to 0.5 in/s. The polynomial model is fast – the 9-second segment was fit in 0.12 seconds and evaluated in 0.02 seconds. The Polylinear model fit was shown previously in Figure 3.4. Due to significant hysteresis in the force-velocity plot, there was significant error in the Polylinear model by as much as 29 lb RMS error. This corresponds to 8% of the signal standard deviation of 366.5 lb.




Figure 3.14. Polylinear model, force-velocity plot




Figure 3.15. Polylinear model, error vs. time

3.3.2 LPF1 Model

The first-order LPF1 filter was fit by trying different time constants ranging from 0.1 to 10 milliseconds in 0.1 millisecond increments and selecting the one with the least error. The resulting force-velocity plot produced by the LPF1 model is shown in Figure 3.16 and a time plot of the forces and modeling error is shown in Figure 3.17. The best time constant was found to be 1.3 milliseconds, and the resulting LPF1 model produced 8.1 lb of RMS error. This is a 72% reduction in error from the Polylinear curve fit, which produced 29 lb of error. This corresponds to 2.2% of the signal standard deviation of 366.5 lb. This clearly indicates that the LPF1 provides a significant improvement in modeling error over a simple force-velocity curve fit. Further inspection of the validation data for the LPF1 model will confirm this result in Section 3.4.




Figure 3.16. LPF1 model, force-velocity plot




Figure 3.17. LPF1 model, error vs. time

3.3.3 LPNL2 Model

The LPNL2 model blends the force predicted by the Polylinear model with the force from the LPF1 model. At low velocities, the force is determined by the LPF1 model. At higher velocities, the force is determined by the Polylinear model. Between the low and high velocity regions, the force is a linear combination of the Polylinear and LPF1 models.

The time constant and the high and low velocity transition forces that define the LPNL2 model were found using MATLAB's constrained optimization function *fmincon*. The time constant was bounded to be from 0.1 to 10 milliseconds. The high and low velocity transition forces were bounded between 10 and 1000 lb. They were also constrained so that the low velocity transition force occurs at a lower force than the high velocity transition force with a minimum difference of 1 lb to avoid a large slope in the weighting function.

The fit resulted in a time constant of 1.3 seconds, which agrees with the time constant in the LPF1 model. The transition forces were 560 lb and 1000 lb. The fit error was 8.0 lb RMS, which is only slightly smaller than the 8.1 lb RMS error for the LPF1 model. The reason for this small improvement in error is that the data for this particular shock build does not exhibit the near-zero hysteresis at high velocities that previous studies exhibit. While this does not indicate a significant benefit of LPNL2 versus LPF1, further analysis of the validation data in Section 3.4 will show that LPNL2 does not do worse than LPF1 and performs noticeably better for some shock builds and drivefiles.

3.3.4 Summary

This section has described the fitting of three different types of shock absorber models for shock build 1 at 50 psi and 1 click on the rebound adjuster. A summary of the results of the fits is shown in Table 3.4. The fit time and evaluation times are the times required to fit the model and to evaluate the model for a 9 second velocity input sampled at 2000 Hz on a 1.8 GHz Pentium M laptop with 512M RAM, and are given for comparison purposes. While Polylinear must be evaluated as an input to LPF1 and LPNL2, the evaluation times for LPF1 and LPNL2 do not include the evaluation time for Polylinear.

Table 3.4. Summary of model fitting for build 1, 50 psi, 1 click

	Fit Error lb RMS	% Reduction from Static	% of Signal STD	Fit Time s	Evaluation Time s
Polylinear	29.38	-	8.0%	0.12	0.02
LPF1	8.1	72.3%	2.2%	32	0.03
LPNL2	8.0	72.7%	2.2%	181	0.04

The Polylinear curve fit model does a good job modeling the shock, with a RMS error of 29 lb, or 8% of the measured force standard deviation. The fit and evaluation times are very small.

The dynamic LPF1 model does significantly better than Polylinear, with a 72% reduction in RMS error. The fit time of 32 seconds could be reduced further by using an optimization routine with a good initial guess or an exhaustive search over fewer lag values. Since the differential equations are linear, evaluation can be done quickly using MATLAB's *lsim* function.

The LPNL2 only does slightly better than LPF1. This is because the particular build tested does not exhibit much of the near-zero lag at high velocities that LPNL2 was designed to accommodate. Further validation analysis will show some significant improvements in LPNL2 over LPF1. The fit time of 181 seconds could be greatly reduced by using a better initial guess of the time constant and the transition forces. Further improvement could be made if the transition forces were fixed based on inspection of the force-velocity curves.

The model fitting was repeated for all 12 shock configurations described in Section 3.2. The effect of shock configuration on modeling error will be explored in the next section.

3.4 Model Validation

This section analyzes how well the different shock models are able to predict the shock's performance. The first section will present the results of the random testing. The next section will present the sine testing, which will allow us to determine which amplitudes and velocities the model will do well and where it may have problems. The third section

presents the static testing, which allows us to determine the effect of the gas spring and seal drag. The last section presents some bump testing results, which indicates the model performance for extreme disturbances.

3.4.1 Random Testing

This section presents the validation results of the random testing. The velocity profile used for random testing was used to predict the shock force using the three different model types for each of the 12 different builds. The model force was then compared to the actual force for that particular shock configuration and the RMS error was calculated, as shown in Table 3.5 and Figure 3.18.

Table 3.5. RMS error analysis, lb

	Build 1 1 click 50 psi	Build 1 14 clicks 50 psi	Build 1 1 click 150 psi	Build 1 14 clicks 150 psi
Polylinear	29.05	34.04	27.2	32.7
LPF1	7.5	9.2	7.9	9.8
LPNL2	7.5	8.7	7.6	9.2

	Build 2 1 click 50 psi	Build 2 14 clicks 50 psi	Build 2 1 click 150 psi	Build 2 14 clicks 150 psi
Polylinear	41.5	45.5	22.3	28.1
LPF1	22.7	22.6	6.1	7.4
LPNL2	21.2	21.5	6	7

	Build 4 1 click 50 psi	Build 4 14 clicks 50 psi	Build 4 1 click 150 psi	Build 4 14 clicks 150 psi
Polylinear	17.9	23.2	27.2	24
LPF1	5.3	7.8	10.0	8.3
LPNL2	5.2	7.7	9.8	8.2

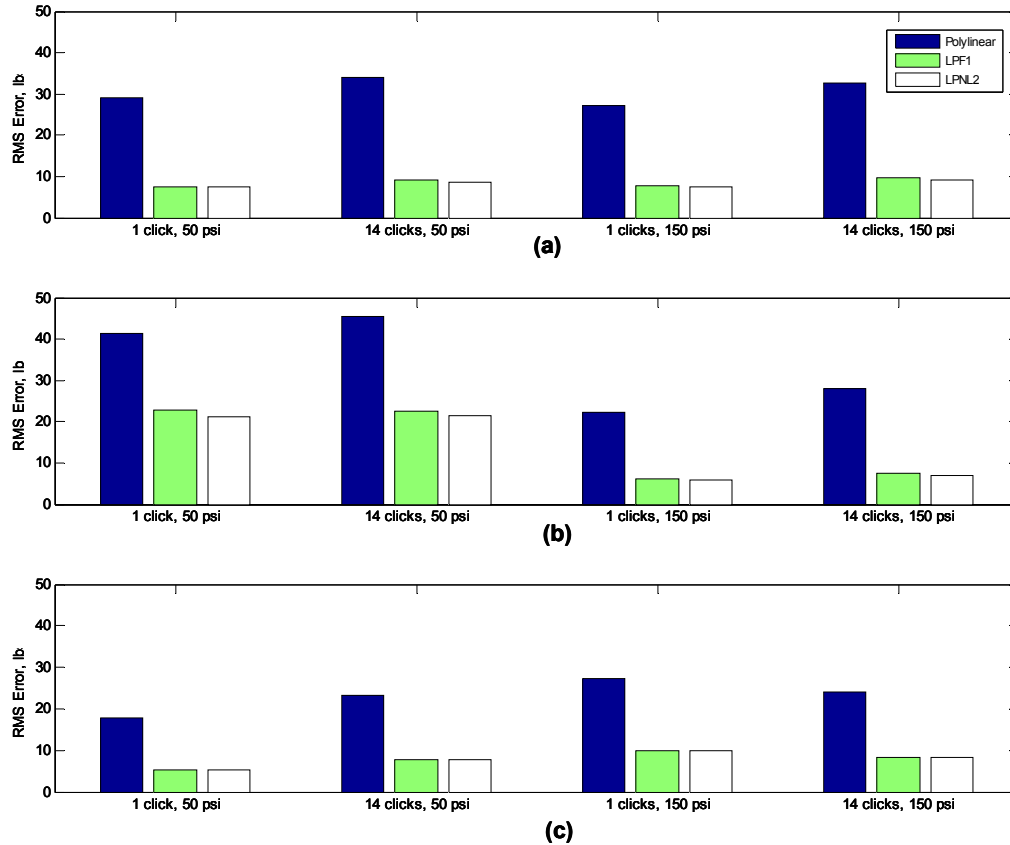


Figure 3.18. Random validation: (a) Build 1, (b) Build 2, (c) Build 4

Several significant observations can be made from these results. First note that in all cases the error for the Polylinear model is much larger than the error for the two dynamic models. The dynamic models ranged from 45-75% less error than the Polylinear model.

The errors for the two dynamic models are similar, but the error for the LPNL2 model is less in all cases. The reduction in error from LPF1 to LPNL2 ranged from 1-7%. It is important to note that all of the builds tested had some hysteresis at the highest velocities, so the benefits of the LPNL2 model may not be evident here.

The errors were smallest for build 4, which had softer compression shims. This build had the least amount of hysteresis, with 18 lb RMS error for the Polylinear model at 50 psi and one click, compared to 29 and 42 lb for the other two builds.

The errors were the largest for build 2 at 50 psi, which had softer head valve compression shims. This makes sense, as the reason for using a head valve is that lower gas pressures

can be run without increased hysteresis. This is due to the large amount of hysteresis in the rebound stroke in build 2 at low pressure, as seen in Figure 3.19. During the compression stroke at low pressure, not enough pressure is built up in the compression chamber to create the required flow to the rebound chamber. This causes fluid expansion in the rebound chamber, which must be compressed during the rebound stroke before the force can fully develop. This creates a lag in the force-velocity plot. When we decreased the head valve compression shims, we increased the fluid expansion in the rebound chamber during compression, increasing the lag in the rebound stroke. This effect was eliminated when we increased the gas pressure to 150 psi.

Also notice that the percent improvement for the dynamic models compared with the algebraic model is worst for build 2, low pressure. The improvement was only about 45% for this case, while most other cases showed 60-75% improvement. This indicates that the dynamic models do not do as well with large hysteresis.

An important observation is that the largest modeling errors occur in shocks that do not perform as well as the others. The modeling error is low for shocks with good performance that we may want to put on the car. This indicates that the dynamic models are sufficient for simulations to help choose the best setups. If it is necessary to model a shock with a large pocket of hysteresis like the one in Figure 3.19a, it may be possible to add another lagged state that only influenced the desired velocity range.




Figure 3.19. Force-velocity for build 2, 1 click: (a) 50 psi, (b) 150 psi

For builds 1 and 2, the error increased with a more restrictive bleed setting and lower gas pressure. Both of these factors generally increase hysteresis, causing more modeling error.

Based on the trends with the bleed and head valves, one might expect that softening the head valve compression shims would increase the error. This is true for softening the head valve with low gas pressure, but the exact opposite trend is seen when softening the head valve at high pressure. This can be explained by recalling that the head valve is used to prevent fluid expansion without high gas pressures. At low pressures, where the head valve is working as it was designed, reducing the effectiveness of the head valve increases hysteresis and increases modeling error. At high pressures, the head valve is not needed and acts as another flow restriction, actually increasing compressibility effects. As the effect of the head valve is reduced, the flow restriction is lessened, reducing hysteresis and modeling error. The interaction between the head valve and the gas pressure seem to be a significant factor influencing model accuracy.

3.4.2 Sine Testing

This section presents the sine test validation results, which will give us a better feel for where the model performs well and where it breaks down. The response for the three shock model types and the 12 shock configurations was calculated for the 14 sine waves listed in Table 3.3 and the simulated response was compared to the measured response. The RMS error for each test was calculated and compiled for analysis.

The first thing we consider is how the different model types compare at different amplitudes and frequencies. Figure 3.20 shows how the three model types compare for build 1 with 50 psi gas and 14 clicks rebound. Similar trends exist for the other shock configurations.




Figure 3.20. Comparison of sine validation error versus model type for build 1, 14 clicks, 50 psi

The first thing to notice is that the error for the dynamic models is below 25 lb RMS for all tests except for the low velocity, low frequency. This is also the case with all shock configurations except for the low pressure, soft head valve configuration that was found to have the largest errors discussed in the random validation due to fluid compressibility. This shock configuration had errors exceeding 25 lb RMS at 1 inch at 20 in/s and 30 in/s and at 0.5 inch at 20 in/s. The generally low errors across all shock configurations indicate that the model is reasonable for simulation applications.

For each amplitude, the errors for the three different model types are similar at the lowest velocities, but the dynamic models outperform the algebraic model at higher frequencies. For the 1, 0.5, and 0.1 inch amplitudes, the largest frequencies where the algebraic and dynamic models have almost identical error are 1.6, 1.6, and 4 Hz, respectively. Above these frequencies, the improvement in using the dynamic models is significant. This seems to indicate that the shock dynamics only become significant above roughly 2-4 Hz, at which point there is enough dynamic behavior that a dynamic model is warranted.

Next, notice that while error grows at higher frequencies, the largest error of 32 lb RMS actually occurs at the lowest amplitude and frequency. The reason for this can be seen in Figure 3.21. In Figure 3.21a we see the force-velocity plot at 1 and 6 in/s. Also shown is the Polylinear model. While the data exhibits a large nose, the model is almost linear.

This is because the model was fit over a large velocity range, where the low speed nose is not as noticeable. At higher velocities or with more hysteresis, the effect of the low speed nose is not noticeable, as shown in Figure 3.21b. This issue should be considered before collecting data to make shock models. If a better fit at low velocity and low frequency is desired, this could be achieved by fitting the low velocity polynomial in the Polylinear model using more low velocity data.

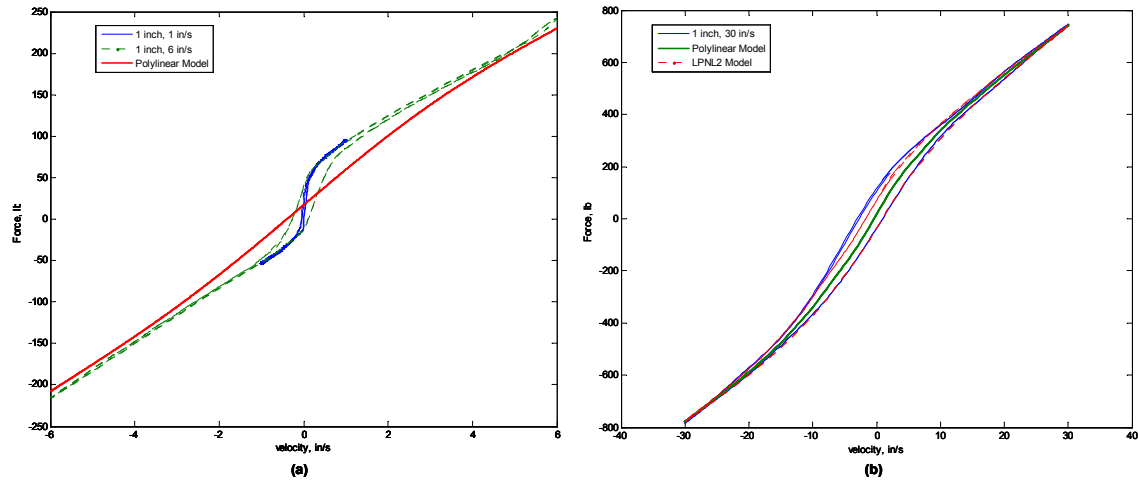


Figure 3.21. Force-velocity for sine tests: (a) low velocity, (b) high velocity

The next thing to notice is that the LPNL2 model is slightly better than the LPF1 at higher frequencies, where hysteresis is more significant. The reason for this can be seen in Figure 3.21b. At low velocities, the data exhibits 200 lb of hysteresis, but at higher velocities, there is very little hysteresis. The transition between LPF1 at low velocity and Polylinear at high velocity allows the model to have lag at low velocity and quickly transition to small lag at high velocity.

3.4.3 Static Testing

Static testing was conducted for the 6 combinations of build and gas pressure for the low bleed setting. The drivefile starts at an initial test offset of 2 inches and stays there for 20 seconds to allow the shock pressures to equalize. It then moves up one inch at 0.2 in/s and maintains the new position for 20 seconds. This process is repeated until the shock is moved up 3 inches from the original test offset, at which point the process is repeated in the downward direction until the shock reaches the initial test offset of 2 inches. The

static test gives a spring rating of the shock at three points, which gives us both the force-deflection curve and also the amount of hysteresis in that curve caused by directional effects such as seal drag. Since this process eliminates the damping effect, this gives us an indication of the stiffness in parallel to the damper.

The time histories of the force for the 6 static tests are shown in Figure 3.22. For reference, the drive profile is also shown in Figure 3.22a. For comparison purposes, each signal was shifted so the force at zero time was zero. The time ranges where the position was held constant are clearly marked on Figure 3.22b. The high force regions between static positions are when it was moving at 0.2 in/s. Once the static position was reached, the shock took about 5 seconds to equalize when compressing and no time in rebound. This is because the shock has a rebound bleed with a check valve that is closed in compression.




Figure 3.22. Static test time plots: (a) Drive profile, (b) Force

The data from Figure 3.22 was reduced to a force-deflection plot by averaging over the time held at each static position, excluding the first 5 seconds and last second of each interval. To allow the results of all 6 force-displacement tests to be considered in one plot, each plot was shifted so that the force at zero position when moving in compression was zero for each test. The resulting force-deflection plot is shown in Figure 3.23.




Figure 3.23. Static test force-displacement

This plot shows essentially two different force-deflection curves, one for 50 psi, as indicated by the thin lines, and a stiffer one for 150 psi, indicated by the thick lines. Notice however that the increase in force for the 150 psi curve is only about 12 lb over 3 inches, at a rate of 4 lb/in. This is much lower than the primary suspension rate, indicating that the parallel spring effect is minimal.

Also note that there is about a 6-10 lb offset in the force-deflection curve between compression and rebound. This is most likely caused by seal drag on the shaft bearing, main piston, and floating piston.

It is possible that this could be integrated into a shock model to further improve the accuracy. However, the above analysis shows that these effects are very small in

comparison with the forces seen in the shock during operation, so any benefit would be minimal.

3.4.4 Bump Testing

Each of the three builds were also tested using a 1 inch bump test, with the drivefile shown in Figure 3.13. The transition speed of the bump was 30 in/s. The purpose of the bump test is to determine how well the model performs for a transient input.

Figure 3.24 shows the force for a compression bump from measured data and predicted by Polylinear and LPNL2. The ripple is caused because the dynamometer was not able to perfectly match our bump profile. As this figure shows, both the Polylinear and LPNL2 models do an excellent job of predicting the bump force in both compression and rebound. Similar responses were observed for the other shock builds.

Both models do well in predicting the general shape of the response, but LPNL2 does a slightly better job at matching the phase of the response, as evident in the force-velocity plot in Figure 3.25. LPNL2 matches the response very well in rebound, but does only slightly better than Polylinear in compression. These results suggest that the models are capable of accurately reproducing transient inputs.




Figure 3.24. Force vs. time for build 4 bump test: (a) Compression, (b) Rebound




Figure 3.25. Force-velocity plot for build 4 bump test

3.5 Summary

This chapter has developed models for a Penske 7300 shock absorber, which is used in NASCAR Cup and similar racing series. These models can be created using simple shock dynamometer tests and then be integrated in a vehicle simulation.

First, we presented a simplified physical model of a shock absorber with compressible flow. The simplified physical model was then used to motivate models that could be fit to real shock data. The first model (Polylinear) was a simple curve fit, based on the fact that most of the shock absorber's response is a function of velocity alone.

While the algebraic model can predict the general trend of the force-velocity curve, it cannot characterize the hysteresis that is significant in many shock curves. Motivated by our mechanical equivalent for our compressible model, our first dynamic model (LPF1) was a simple first-order lag of the force calculated by the curve fit.

Many force-velocity curves exhibit significant hysteresis at low velocity but almost no thickness at higher velocities. To accommodate this behavior, a model (LPNL2) was created to behave similarly to a stiffening spring in series with the damper, by using only the lagged force at low velocity, only the curve fit at high velocity and a linear combination of the two models between low and high velocity regions.

Three different builds in a total of 12 different configurations were tested on a shock dynamometer. The drivefiles included random, sine wave, static, and bump. This testing was designed to allow us to investigate the ability of the models to characterize the shock absorbers for different settings under different operating conditions. The random data was then used to fit models for each of the 12 shock configurations and 3 model types.

After the models were fitted to the experimental data, the models were validated using another random signal. By comparing the RMS modeling error between different shock configurations, several important conclusions were made. First, the reduction in error for the dynamic models compared to the algebraic models was significant, ranging from 45% to 75%. Aside from the builds with cavitation, RMS errors were low, suggesting the model is suitable for simulation applications.

The sine wave validation shows that both the algebraic and dynamic modeling errors were very similar below 2 Hz, suggesting that dynamic models do not provide much value below 2 Hz. Larger errors were also present at the lowest frequencies, because the broadband fit data placed less emphasis on the low velocity nose only significant at lower frequency. This could potentially be improved by fitting the low velocity curve fit using low frequency data. The RMS errors were generally low, being below 25 lb RMS in most cases.

Static validation showed that the gas spring stiffness contributed about 10 lb at 3 inches for the 150 psi case, and showed about 5 lb of friction. The bump validation showed that the shock models could accurately predict transient events.

Now that we have developed validated shock models, we can proceed to integrate them into vehicle simulations. We begin by applying these models in a quarter-car vehicle model in Chapter 4.

Chapter 4

Quarter-Car Model Development

During 8-post rig testing, a vehicle is tested by repeating a particular drivefile as various suspension settings are tuned. The performance of the different runs is compared by calculating a series of performance metrics. These metrics are typically used during 8-post rig tests to help understand how the shock selection affects mechanical grip, aerodynamics, and other factors.

This sensitivity information is one of the main products of the 8-post rig test. The race engineers will use this information to help select or eliminate potential shock packages to try on the race track, and will use the sensitivity information to help make adjustments based on how the vehicle performs on the track.

The main goal of this research is to accurately predict this sensitivity information using simulations, which can augment knowledge gained from 8-post rig testing. To create an accurate model for these simulations, we will use rig test data to perform system identification, which will produce a vehicle model. To make sure our system identification methods are working properly, we will first run the ID on simulated data with a known model. This will allow us to compare the true model with the ID model and tune our ID methods.

Once the ID methods have been developed by performing ID on simulated data, the ID method is then applied to data collected from actual rig data. The ability of the ID model to predict the influence of shock selection on the quarter-car rig's response is then evaluated.

This chapter develops the methods at a quarter-car scale. The first section develops the equations of motion for a linear quarter-car model. This model treats the nonlinear shock model force as an input to a linear dynamic system. This allows us to decouple the nonlinear shock model from the linear vehicle model and consider the frequency response

of the linear vehicle model. A parametric analysis of the linear vehicle model is performed to provide a better understanding of the effect of shock selection on vehicle performance.

Once we have a basic understanding of the behavior of the linear vehicle model, we then couple the linear dynamic model with the nonlinear shock model. This model is used in combination with the shock models presented in Chapter 3 to generate simulated quarter-car rig data.

The methods required to perform the system ID are then discussed and applied to the simulated test data. The results of the ID are compared with the theoretical response to verify that the ID process is working properly.

To verify that the method will work on real-world data, the ID process is then applied to real quarter-car rig data and the ID model is used to predict the response for different shock configurations.

4.1 Quarter-Car Model, Decoupled Analysis

The section develops the equations of motion for a quarter-car vehicle model, where the nonlinear shock model force is treated as an input to a linear dynamic vehicle model. This allows us to decouple the nonlinear shock model from the linear vehicle model and apply linear analysis to better understand the effect of shock selection on vehicle performance.

This section develops the decoupled quarter-car model with a shock force input. Once the model is developed, the behavior of the linear model is analyzed using a parametric study of the frequency response.

4.1.1 State-Space Model

A diagram of the quarter-car model is shown in Figure 4.1. It includes two masses, which represent the inertia of the chassis and the wheel assembly. The two masses are

connected through the suspension, represented by a linear spring and damper. The linear suspension parameters represent wheel rates that take into account the motion ratio between the wheel and the physical spring or damper. The linear damping rate can be used to represent a linear shock absorber or other losses in the suspension.




Figure 4.1. The quarter-car model

The suspension model also includes a generic force input, which will be used in Section 4.2 to integrate our nonlinear shock absorber models. Note however that this is a generic force input that acts on both the sprung and unsprung masses, so it could also be used to add other types of suspension behavior, such as progressive springs, coil-binding, and suspension friction.

The tire is represented as a linear spring. In addition to the shock force input, there are other inputs that represent the road profile and aerolader force. During rig testing, the road inputs are simulated using wheelloaders and the aerodynamic inputs are simulated using aeroloaders.

The equation of motion for the quarter-car model in Figure 4.1 is

$$\begin{bmatrix} m_1 & 0 \\ 0 & m_2 \end{bmatrix} \begin{bmatrix} \ddot{w}_1 \\ \ddot{w}_2 \end{bmatrix} + \begin{bmatrix} \bar{c} & -\bar{c} \\ -\bar{c} & \bar{c} \end{bmatrix} \begin{bmatrix} \dot{w}_1 \\ \dot{w}_2 \end{bmatrix} + \begin{bmatrix} k & -k \\ -k & k+k_t \end{bmatrix} \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} = \begin{bmatrix} 0 & -1 & 1 \\ k_t & 0 & -1 \end{bmatrix} \begin{bmatrix} y_R \\ F_A \\ F_S \end{bmatrix} \quad (4.1)$$

$$M\ddot{w} + C\dot{w} + Kw = Lu$$

where the shock force is positive during compression. Rewriting Equation (4.1) in state-space form yields

$$\begin{aligned} \begin{bmatrix} \dot{w} \\ \ddot{w} \end{bmatrix} &= \begin{bmatrix} \mathbf{0}_{2 \times 2} & I_2 \\ -M^{-1}K & -M^{-1}C \end{bmatrix} \begin{bmatrix} w \\ \dot{w} \end{bmatrix} + \begin{bmatrix} \mathbf{0}_{2 \times 3} \\ M^{-1}L \end{bmatrix} u \\ \dot{x} &= Ax + Bu \end{aligned} \quad (4.2)$$

The output equations are chosen to represent common sensor measurements. From an identification standpoint, it can be important if a sensor is a function of the state alone, a function of only the state and the inputs (direct passthrough), or a function of the states, inputs and system parameters. The sensor equations for the position and velocity of the chassis, wheel, and suspension are a function of the state alone

$$\begin{bmatrix} w_1 \\ w_2 \\ \dot{w}_1 \\ \dot{w}_2 \\ x_s \\ \dot{x}_s \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ -1 & 1 & 0 & 0 \\ 0 & 0 & -1 & 1 \end{bmatrix} x \quad (4.3)$$

The ride height is a function of both the state and the road input

$$x_{RH} = [1 \ 0 \ 0 \ 0]x + [-1 \ 0 \ 0]u \quad (4.4)$$

The spring force, tire force, chassis accelerometer, and wheel accelerometer sensor equations include the states, inputs, and system parameters.

$$\begin{bmatrix} F_{spring} \\ F_t \\ \ddot{w}_1 \\ \ddot{w}_2 \end{bmatrix} = \begin{bmatrix} -k & k & 0 & 0 \\ 0 & -k_t & 0 & 0 \\ A(3,:) \\ A(4,:) \end{bmatrix} x + \begin{bmatrix} 0 & 0 & 0 \\ k_t & 0 & 0 \\ B(3,:) \\ B(4,:) \end{bmatrix} u \quad (4.5)$$

where $A(i,:)$ and $B(i,:)$ refers to the MATLAB notation for the i^{th} row of matrices A and B . Combining sensor Equations (4.3)-(4.5) and rearranging sensor order yields

$$\begin{bmatrix} w_1 \\ w_2 \\ \dot{w}_1 \\ \dot{w}_2 \\ x_s \\ \dot{x}_s \\ \ddot{w}_1 \\ \ddot{w}_2 \\ F_t \\ x_{RH} \\ F_{spring} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ -1 & 1 & 0 & 0 \\ 0 & 0 & -1 & 1 \\ A(3,:) \\ A(4,:) \\ 0 & -k_t & 0 & 0 \\ 1 & 0 & 0 & 0 \\ -k & k & 0 & 0 \end{bmatrix} x + \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ B(3,:) \\ B(4,:) \\ k_t & 0 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} u \tag{4.6}$$

$$y = Cx + Du$$

The decoupled quarter-car model is then found in state-space form by combining Equations (4.2) and (4.6) to get

$$\begin{aligned} \dot{x} &= Ax + Bu \\ y &= Cx + Du \end{aligned} \tag{4.7}$$

where the state-space matrices are defined in Equations (4.1), (4.2), and (4.6).

Now that we have formed the linear model, we can study the behavior of the system by analyzing the effect of different system parameters on the model's frequency response. For the purpose of this analysis, a set of baseline parameters was chosen to be generally representative of suspensions used in the racing industry, while not representing any specific suspension. The parametric study will consider the parametric sensitivity around this baseline setup.

The baseline parameter values are shown in Table 4.1. The sprung and unsprung masses correspond to a 3000 lb car with 90% sprung mass, and 25% of the mass on one corner. The suspension and tire stiffnesses are typical values. Notice that unlike a passenger car, where the tire may be as much as an order of magnitude stiffer than the soft primary suspension, the stiff suspensions used in racing are often similar to the tire stiffness.

Table 4.1. Baseline parameter values

$m_1 g$	Sprung Weight	675 lbf
$m_2 g$	Unsprung Weight	75 lbf
k	Suspension Stiffness	1280 lbf/in
k_t	Tire Stiffness	1300 lbf/in
c	Suspension Damping	16.67 lbf / (in/s)

4.1.2 Frequency Response

The two commonly considered measurements in rig testing are tire normal force variation and ride height variation, which are indicators of mechanical grip and aero grip. This section considers how the frequency response for the tire force and ride height changes with variations in suspension parameters. Having an understanding of this fundamental behavior can be a useful tool when trying to analyze the more complex nonlinear system. To plot the wide range of FRF magnitude, we will plot magnitude in decibels (dB), which is defined by

$$|H|_{dB} = 20 \log_{10} |H| \quad (4.8)$$

First, the effect of damping is considered by varying damping from 50-300% of the nominal value. Figure 4.2 shows how the frequency responses for the tire force vary with damping. For the road input, the effect of damping is very small until about 2 Hz. Near the resonance at 3 Hz, damping reduces the resonant peak, reducing tire force variation. After the resonant peak, damping worsens the attenuation, increasing tire force variation. The same trend exists for the aeroloader input, but the reduced attenuation at higher frequencies is less pronounced.




Figure 4.2. Effect of damping on tire force

The shock force input shows a reduction at all frequencies for increased damping, which may seem odd at first glance. We can think of the linear spring and damping forces as the nominal linear suspension forces. The “shock force” can include any additional suspension forces in addition to the nominal forces, including shock or spring nonlinearities, friction, or suspension component dynamics. The third subplot in Figure 4.2 can then be thought of as the control authority of the shock force on the tire force. Increased damping reduces the control authority of the shock force near the resonances, while it has little effect elsewhere.

Next, we consider the effect of damping on the ride height, as shown in Figure 4.3. Increasing damping significantly reduces the ride height variation near the resonance for both the wheelloader and aeroloader inputs. The increased response after the resonance is small for the parameter values shown. This suggests using higher damping forces to minimize ride height variations and improve aerodynamic performance. This will also reduce the tire normal force variation near resonance. The tradeoff is that increasing damping will also worsen the attenuation in tire normal force from 4-13 Hz. Depending on the roughness of the track and the aerodynamic sensitivity, a compromise between the two must be found.




Figure 4.3. Effect of damping on ride height

The next parameter we consider is the suspension stiffness. The effect of suspension stiffness on tire force as suspension stiffness is varied from 5-300% of the baseline value is shown in Figure 4.4. The general trend is that increasing stiffness increases the natural frequency and decreases the damping ratio. The significant change in damping ratio with spring stiffness highlights the significance of matching stiffness with damping. This effect reduces normal force variation before the natural frequency and increases it after

the natural frequency. Also worth noting is the effect of making the suspension very soft. This effectively decouples the two-mass system into two single-mass systems. This provides very low tire normal force variations. Also note that softer springs generally increase the control authority of the shock force, except near resonance.




Figure 4.4. Effect of suspension stiffness on tire force

Next, we consider the effect of suspension stiffness on ride height, shown in Figure 4.5. Similar to the tire force, increasing stiffness increases natural frequency and decreases damping ratio. This produces a significant reduction in ride height variation at lower frequencies. Again, there is a tradeoff between mechanical grip and aerodynamic performance that needs to be considered to select the best combination for a particular track.




Figure 4.5. Effect of suspension stiffness on ride height

This section has developed a decoupled linear quarter-car model and explored the influence of changing suspension parameters on the frequency response of the tire normal force and ride height. A tradeoff has been shown, which generally favors more coupling between the chassis and wheel to reduce ride height variation and less coupling to reduce

normal force variation. The importance of selecting the right combination of shock and spring was also highlighted.

Now that we have developed the linear quarter-car model and have a basic understanding of how it performs, we can now integrate the nonlinear shock models created in Chapter 3 into the simulation.

4.2 Quarter-Car Model, Simulated Data

Now that we have defined shock absorber models and a decoupled linear vehicle model, we can combine the two to create a complete quarter-car vehicle model, as shown in Figure 4.6. This model can then be used to create simulated data sets, which we will use in our identification process to determine if our ID was successful in reproducing the known model.

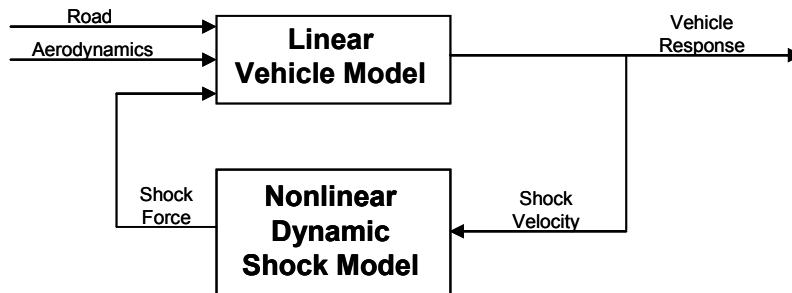


Figure 4.6. Coupling of vehicle and shock absorber models

The quarter-car model equations with inputs for road displacement, aerodynamic force, and shock force were derived in Section 4.1. To couple the shock model with the quarter-car model, the shock velocity from the quarter-car model is treated as an input to the shock model, and the shock force is treated as an input to the quarter-car model.

The shock model presented in Chapter 3 included a curve fit of the force-velocity curve, a first-order lowpass filter of the curve fit, and a weighting function to combine the two. The equations for these three parts are

$$F_{poly} = f(v) \tag{4.9}$$

$$F_{lag} = \frac{1}{\tau s + 1} F_{poly} \quad (4.10)$$

$$F_{LPNL2} = \alpha(F_{poly})F_{lag} + (1 - \alpha(F_{poly}))F_{poly} \quad (4.11)$$

To improve simulation speed for the linear portions of the simulation, the fourth-order state-space model in Equation (4.7) and the first-order lowpass filter in Equation (4.10) are transformed into linear difference equations using the bilinear transformation

$$s = \frac{2}{T} \left(\frac{1 - z^{-1}}{1 + z^{-1}} \right) \quad (4.12)$$

Where s is the continuous-time Laplace operator, z^{-1} is the discrete-time unit delay operator, and T is the sample time. This transformation maps the entire $s = j\omega$ axis to the unit circle, which avoids aliasing and preserves the stability of poles, but causes some warping in the frequency response at higher frequencies [37]. This transformation is done using the MATLAB function `c2d`, with a sample rate of 500 Hz. After transforming Equations (4.7) and (4.10) into discrete time, they may be written as

$$\begin{aligned} x[k+1] &= A_d x[k] + B_d u[k] \\ y[k] &= C_d x[k] + D_d u[k] \end{aligned} \quad (4.13)$$

$$F_{LPNL2}(z) = \left[\frac{1 + z^{-1}}{\left(1 + \frac{2\tau}{T}\right) + \left(1 - \frac{2\tau}{T}\right)z^{-1}} \right] F_{polylinear}(z) \quad (4.14)$$

The final quarter-car model can be implemented in Simulink, as shown in Figure 4.7. The quarter-car block corresponds to Equation (4.13), while the shock model consists of the blocks polylinear, first-order lag, and LPNL2 weighting.

system in Figure 4.8. The three inputs u_1 , u_2 , and u_3 correspond to the wheelloader, aeroloader, and shock force inputs respectively.

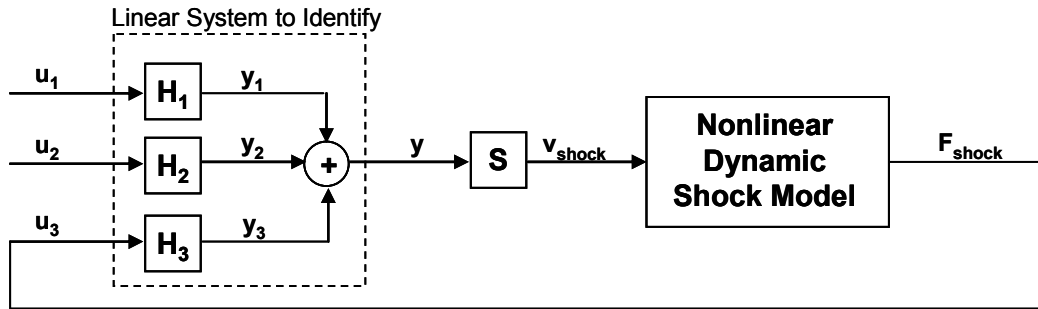


Figure 4.8. Diagram of complete system and the subsystem to identify

First, we will provide a summary of a nonparametric frequency response estimation method that can be applied for independent inputs. Next, the issue of the shock force being dependent to the first two inputs will be addressed. Finally, a method to convert the nonparametric frequency response estimates into a parametric model that is useful for simulation is discussed.

4.3.1 Nonparametric Frequency Response Estimation

If the inputs were independent, a common practice is to create a nonparametric estimate of the transfer function matrix, which has a coefficient for each input, output, and frequency. For simplicity, we show the method for only one frequency and one output. The model that we want to fit is

$$y = H_1 u_1 + H_2 u_2 + H_3 u_3 \quad (4.15)$$

where the inputs and outputs consist of the frequency domain component for one frequency. To fit the model, we will collect several data blocks and choose the model coefficients to minimize model error between the data blocks. Data blocks can be taken from a series of tests or from selecting multiple blocks out of a longer data set. The data for the k^{th} data block will have some deviation from the ideal model. If we assume the error is caused by the outputs, the error model for the k^{th} data block can be represented as

$$y[k] = H_1 u_1[k] + H_2 u_2[k] + H_3 u_3[k] + \varepsilon[k] \quad (4.16)$$

If we collect K data sets, we can arrange Equation (4.15) in matrix form as

$$\begin{bmatrix} y[1] \\ \vdots \\ y[K] \end{bmatrix} = \begin{bmatrix} u_1[1] & u_2[1] & u_3[1] \\ \vdots & \vdots & \vdots \\ u_1[K] & u_2[K] & u_3[K] \end{bmatrix} \begin{bmatrix} H_1 \\ H_2 \\ H_3 \end{bmatrix} + \begin{bmatrix} \varepsilon[1] \\ \vdots \\ \varepsilon[K] \end{bmatrix} \quad (4.17)$$

$$Y = UH + \varepsilon$$

We can choose our nonparametric estimate to minimize the norm of the error vector

$$J = \|\varepsilon\|^2 = \varepsilon^* \varepsilon \quad (4.18)$$

This is a linear least-squares regression problem, which yields the solution

$$\begin{aligned} H &= (U^* U)^{-1} (U^* Y) \\ &= S_{uu}^{-1} S_{uy} \end{aligned} \quad (4.19)$$

where S_{uu} and S_{uy} are the autospectrum and cross-spectrum estimates for the K blocks. Note that the rank of the autospectrum matrix is at most equal to the rank of U , which requires that there be at least three linearly independent data block rows and three independent input columns for a unique solution H to exist.

This process is repeated for all frequencies and outputs to complete the nonparametric frequency response estimation. The model generated by this nonparametric frequency response estimation process is often called a H_l frequency response estimate.

4.3.2 Frequency Response Estimation with Dependent Data

The H_l frequency response estimation method described in the previous section requires three independent inputs for a unique solution of Equation (4.19) to exist. To see why this condition may be violated in our case, consider a linear shock model

$$\begin{aligned} F_{shock} &= cv \\ &= cSy \end{aligned} \quad (4.20)$$

where the vector S selects the shock velocity from the output vector. Substituting into Equation (4.15) and solving for y yields

$$\begin{aligned}
 y &= H_1 u_1 + H_2 u_2 + H_3 c S y \\
 (I - H_3 c S) y &= H_1 u_1 + H_2 u_2 \\
 y &= (I - H_3 c S)^{-1} [H_1 u_1 + H_2 u_2]
 \end{aligned} \tag{4.21}$$

$$y = \tilde{H}_1 u_1 + \tilde{H}_2 u_2$$

This means the shock force will be

$$\begin{aligned}
 F_{shock} &= c S y \\
 &= c S [\tilde{H}_1 u_1 + \tilde{H}_2 u_2]
 \end{aligned} \tag{4.22}$$

$$F_{shock} = \bar{H}_1 u_1 + \bar{H}_2 u_2$$

The shock force is linearly dependent with the other two inputs if the shock force can be written

$$F_{shock} = c_1 u_1 + c_2 u_2 \tag{4.23}$$

for some complex constants c_1 and c_2 . This proves that it is not possible to find a unique H_I frequency response estimate when the shock is linear. If the shock has significant nonlinearity, this will add uncorrelated energy into the shock force input, making a unique solution possible at this frequency.

One method to address the linear dependence issue is to test the car first with a linear shock (or no shock), then with a different shock. The first test is used to identify the transfer functions for the first two inputs only. The second test is used to identify the transfer function for the shock input only. This two-part system identification is shown in Figure 4.9. This strategy is motivated by the fact that some systems lose identifiability if tested while closed-loop feedback is applied, but can be made identifiable by using two or more different control policies [38, 39]. In our case, the control policy corresponds to the shock model.

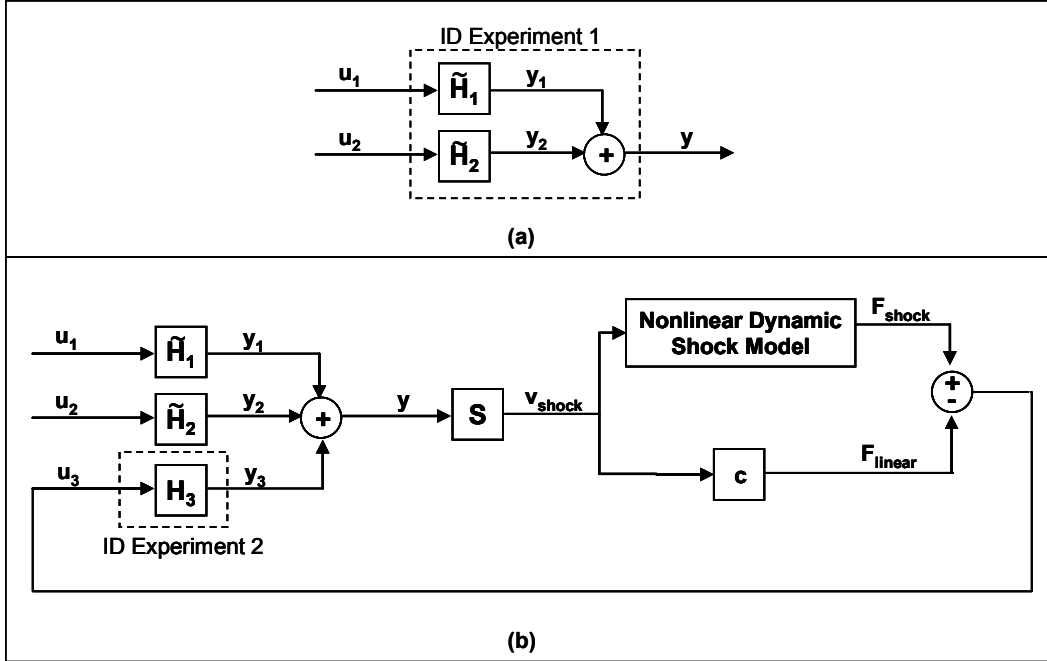


Figure 4.9. System ID experiments: (a) Linear shock for road and aero inputs, (b) Second shock for shock input

Equation (4.22) shows that the 3 input model of Equation (4.15) can be reduced to two inputs in the case of a linear shock. The frequency responses \tilde{H}_1 and \tilde{H}_2 describe the linear system behavior with the shock included. Note that if the shock is removed, the frequency responses reduce to H_1 and H_2 , the original frequency responses in Equation (4.15). If we run test with a linear shock (or no shock), we can perform the H_1 identification procedure to identify \tilde{H}_1 and \tilde{H}_2 .

Next, we run a second ID experiment with an arbitrary shock. Since the shock force that would be created by the linear shock from ID experiment 1 is already included in the frequency response for the first two inputs, we subtract it from the shock force. The difference between these two forces is then used for input 3.

$$u_3 = F_{shock} - cv \quad (4.24)$$

The shock for the first and second experiment must be different, so there will be energy in u_3 . The output due to u_3 , labeled y_3 in Figure 4.8b, can be found since we know \tilde{H}_1 and \tilde{H}_2 . In general,

$$\begin{aligned} y &= y_1 + y_2 + y_3 \\ &= \tilde{H}_1 u_1 + \tilde{H}_2 u_2 + y_3 \end{aligned} \tag{4.25}$$

$$y_3 = y - (\tilde{H}_1 u_1 + \tilde{H}_2 u_2)$$

If the drivefile for experiment 1 and experiment 2 are the same, this reduces to

$$y_3 = y - y_{\text{experiment 1}} \tag{4.26}$$

where $y_{\text{experiment 1}}$ is the output from experiment 1.

Now that we have u_3 and y_3 , we can find the FRF estimate of H_3 . Combining our estimate of \tilde{H}_1 and \tilde{H}_2 from ID experiment 1 and our estimate of H_3 from ID experiment 2 gives us a complete nonparametric frequency response estimate of the linear system in Equation (4.15). Now that we have addressed the issue with the dependence of the shock force on the other two inputs, we can treat the frequency response estimates as if they were obtained using a test with three independent inputs.

4.3.3 Parametric System Identification

The nonparametric frequency response estimate consists of as many coefficients as there are data points in a single block of data. While this could be used for simulation, it is more practical to fit the nonparametric model with a parametric model with fewer parameters. The parametric model has the advantage of smoothing the frequency response function and improving the model in small frequency ranges where the nonparametric model was poor. It also allows us to see the poles and zeros for the different input/output pairs.

There are many different methods of performing parametric system identification, which can be chosen depending on the particular application. The transfer function for each input/output pair can be defined by its poles, zeros, and constant gain. In an ideal linear system, all input/output pairs should share the same poles and only differ by the zeros and constant gain. The identification of the system poles is often important for applications in controls and structural parameter identification.

In our case, it is not necessary to identify a consistent set of system poles. All we want is to simplify our nonparametric frequency response estimates into a convenient transfer function relationship that will allow us to simulate the linear system more efficiently. This allows us to use a simpler method of fitting a single-input single-output (SISO) transfer function to the frequency response estimate for each input/output pair. This method will give a different set of poles, zeros, and a constant for each input/output pair. One advantage of this method is that it is very easy to fit a SISO transfer function, and many methods are readily available.

Another advantage is that it localizes error in the frequency response estimate to the corresponding transfer function estimate. If there is significant error in the frequency response estimate for one of the input/output pairs due to factors such as noise, poor excitation, nonlinearity, or sensor malfunction, the error will only affect the corresponding transfer function estimate, not any other transfer function estimate.

The SISO method chosen for this work uses the *invfreqs* function in MATLAB [40]. This method fits a SISO transfer function of the form

$$H(s) = \frac{B(s)}{A(s)} = \frac{b_1s^n + b_2s^{n-1} + \dots + b_0}{a_1s^m + a_2s^{m-1} + \dots + a_0} \quad (4.27)$$

by minimizing the equation-error cost function

$$J = \sum_k w_t(k) \left| h(k)A(j2\pi f(k)) - B(j2\pi f(k)) \right|^2 \quad (4.28)$$

where $f(k)$ is the frequency at index k , $h(k)$ is the nonparametric frequency response estimate, and $w_i(k)$ is a weighting factor. The problem can be solved easily using least squares regression.

An alternate optimization problem that is more intuitive and often provides better results is to minimize the output-error cost function

$$J = \sum_k w_i(k) \left| h(k) - \frac{B(j2\pi f(k))}{A(j2\pi f(k))} \right|^2 \quad (4.29)$$

Unfortunately, this optimization problem is nonlinear in nature and must be solved using more generic iterative nonlinear programming techniques. The *invfreqs* function uses a Gauss-Newton iterative search. The optimal solution to Equation (4.28) is used as an initial guess for an iterative solution.

4.4 System Identification on Simulated Data

Now that we have developed the required system identification methods, we can then apply them to identify the linear portion of the quarter-car model from our simulated data.

4.4.1 System Identification Results

Two simulated ID experiments were performed for use in the system identification process. Both used the 100-second random signal with a flat spectrum up to 50 Hz described in Section 4.2. The first ID experiment used a linear shock with a damping coefficient of 16.6 lb/(in/s). The second ID experiment used the build 1 shock model with 50 psi gas pressure and 1 click bleed.

The first ID experiment was used to estimate the frequency response \tilde{H}_1 and \tilde{H}_2 . These frequency responses correspond to the effect of the road and aerodynamic inputs on the quarter-car with a linear shock. Twenty data blocks were created from the 100 second data file using a Hamming window and 50% window overlap. This created data blocks

with 9523 samples, corresponding to a 0.105 Hz spectral resolution in the frequency response estimate.

The second ID experiment was then used to estimate the frequency response H_3 using the same settings as the first experiment. This frequency response determines the influence of the portion of the shock force in addition to the force that would be created by the linear shock used in the first ID experiment. The change in shock force input and change in output measurements for the H_1 estimation are calculated using Equations (4.24) and (4.26).

Once the frequency responses \tilde{H}_1 , \tilde{H}_2 , and H_3 have been estimated, a parametric model can be fitted to each input/output pair using the *invfreqs* function. Since we know our system is 4th order and some sensor equations have direct passthrough, we set both the numerator and denominator polynomial orders to 4. We also use a weighting function of 1 below 20 Hz and 0 above 20 Hz. This weighting function could be further adjusted if a frequency response estimate was poor in a particular area or more accuracy was desired in a particular frequency range.

The results for the nonparametric and parametric ID are shown in Figure 4.10 and 4.11 for the tire force and ride height. Each plot shows the ideal response of the linear system with the linear shock from ID experiment 1, the H_1 frequency response estimate, and the parametric model fit. As these figures illustrate, both the H_1 estimate and the parametric estimate closely match the ideal.




Figure 4.10. ID results for tire force




Figure 4.11. ID results ride height

A comparison of the natural frequencies and damping ratios for the transfer function between road input and chassis position is shown in Table 4.2. This shows good agreement between the ideal and ID model.

Table 4.2. Comparison of natural frequency and damping ratios

Mode	Natural Frequency, Hz		Damping Ratio, %	
	Ideal	ID Model	Ideal	ID Model
1	3.04	3.04	6.10	6.38
2	18.44	18.55	42.19	41.54

4.4.2 Simulation using the ID model

Now that we have identified the linear dynamics of our simulated quarter-car data, we can now use the ID model for simulation of the effect of changing shock absorbers. This simulation is almost identical to the simulation that was used to generate the simulated data, with two exceptions. First, the known linear dynamic model is replaced by the ID model. Second, the linear damping force from the first ID experiment must be subtracted from the shock model force. The ID model is discretized using MATLAB's *c2d* function with a sample rate of 500 Hz. The Simulink layout for simulation using the ID model is shown in Figure 4.12.

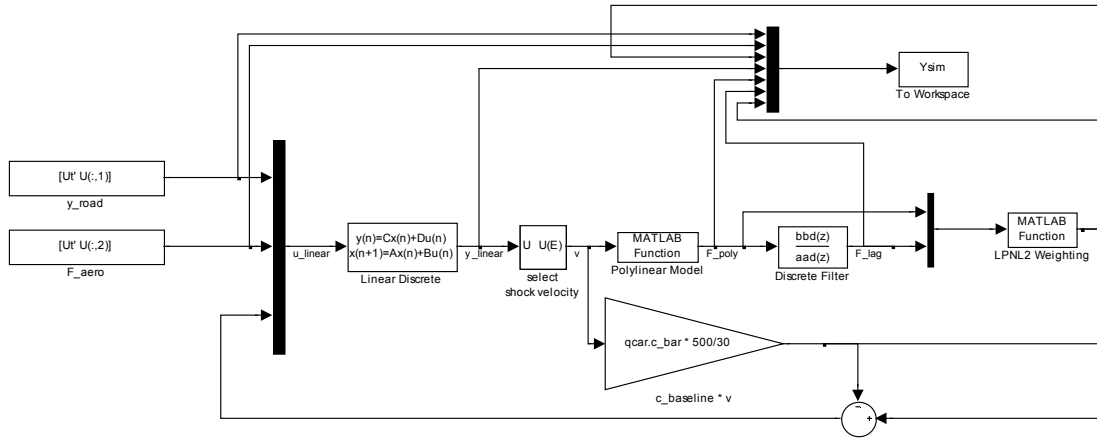


Figure 4.12. Simulink layout for simulation using the ID model

This model can then be used to simulate the response of the quarter-car for any desired drivefile or shock model. A random drivefile was created using the relative amplitude spectrum shown in Figure 4.13, with an amplitude of 1 inch on the wheelloader and an amplitude of 100 lb and a mean of 500 lb on the aerolader. This drivefile was run for all 12 different shock configurations described in Chapter 3 using both the original model and the ID model. To compare the responses from the two models, the standard deviation was calculated for each signal. A comparison of the tire force standard deviation for the simulated data and the ID model for all 12 shock configurations is shown in Figure 4.14. As this figure shows, the tire force standard deviation trend is accurately predicted by the ID model. The error in tire force standard deviation is approximately 4 lb, which is less than 1% error. Similar trends exist for the other signals. This shows that the ID model can be used to accurately predict the response measures and aid in shock selection.




Figure 4.13. Relative amplitude spectra for drivefile




Figure 4.14. Comparison of tire force standard deviation

4.5 System Identification for Quarter-Car Rig Data

The ID methods were developed in Section 4.4 using ideal simulated data sets from a known linear system with no sensor noise. To further prepare the method for real 8-post rig data, the method was applied to quarter-car rig data. A series of tests were performed

using an existing quarter-car rig, 6 different builds of the Penske 7300 shock absorber, and instrumentation to measure the quarter-car rig response. The data collected from the tests was then used to identify a quarter-car model using the methods developed in Section 4.4. To confirm that the model can be used to accurately predict the response of the quarter-car rig, the model was used to predict the response for different shock builds, which was then compared to test data.

4.5.1 Experimental Setup

The quarter-car rig tests were performed on an existing quarter-car rig, shown in Figure 4.14, available at the Center for Vehicle Systems and Safety (CVeSS) of Virginia Tech. The quarter-car rig consists of a fixed external frame and two floating frames, which slide up and down inside the fixed frame using roller bearings and Teflon bearings. Both the fixed frame and the floating frames are constructed from 80/20[®] extruded aluminum. The upper and lower floating frames are meant to represent the sprung and unsprung masses. At the bottom of the rig, there is a hydraulic actuator, which can be used to simulate a road excitation. The hydraulic actuator attaches to the unsprung mass through a series of rubber bushings, selected to represent the stiffness of a passenger car tire.

The sprung mass is supported by the unsprung mass using two air springs. To provide a larger compressible air volume, which results in a more linear air spring behavior, the air springs were attached to a 30 gallon air tank. For our study, we also mounted a Penske 7300 shock absorber between the air springs, as shown in Figure 4.15. Before each test, the air springs were used to set the suspension sag so that the shock had 4.5 inches of shaft exposed. There was space to add eight 50-lb weights to the sprung mass, bringing the total sprung mass up to about 575 lb. While this is lower than the estimated quarter-car sprung mass of 675 lb for a NASCAR cup car, the rig can still be used to prove that the system identification methods will work on experimental data.





Figure 4.15. Quarter-car rig at CVeSS (photo by author, 2009)

Sensors were mounted on the quarter-car rig to measure the rig's response during each test. A shock potentiometer was mounted across the shock directly to the shock bolts, as shown in Figure 4.16b. An accelerometer was mounted on the sprung and unsprung masses, as shown in Figure 4.16a. A string potentiometer with an integrated velocity sensor was mounted between the sprung mass and the fixed frame, as shown in Figure 4.16c. Another string potentiometer was mounted between the unsprung mass and the floor. The hydraulic actuator also had an integrated LVDT and a delta-P cell.



(a)

(b)

(c)

Figure 4.16. Sensor mounting: (a) Accelerometer, (b) Shock potentiometer, (c) String potentiometer with integrated tachometer (photos by author, 2009)

All data acquisition, signal monitoring, and signal generation was done using a dSPACE AutoBox connected to a PC running Controldesk software. The AutoBox is a real-time system, which has 40 A/D channels and 16 D/A channels. It can also be connected to a PC using a fiberoptic cable to transfer information between the Controldesk software user interface on the PC and the real-time AutoBox.

All sensor signals were routed to the A/D channels on the AutoBox for data acquisition shown in Figure 4.17. Controldesk was used to define different drivefiles to be used as a position command for the hydraulic actuator. The signal defined in Controldesk was then converted to an analog signal through the dSPACE D/A, and the analog position command was sent to the MTS 407 controller, which controls the hydraulic actuator.

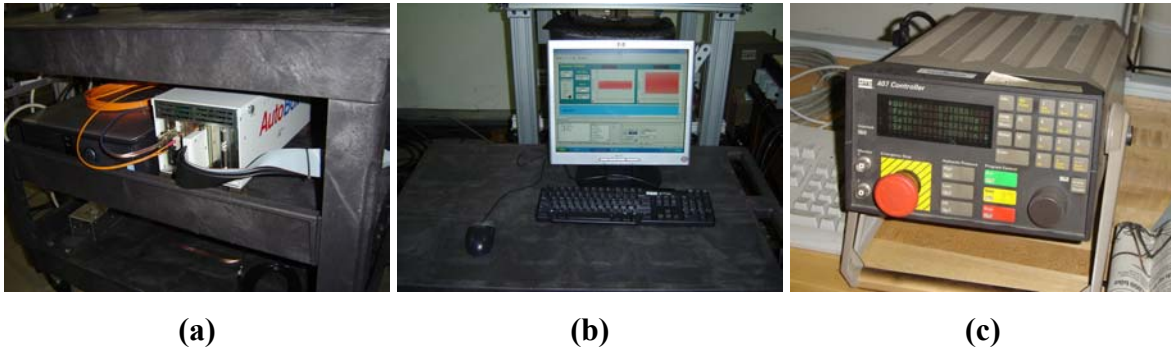


Figure 4.17. Other components: (a) dSPACE AutoBox, (b) dSPACE ControlDesk software, (c) MTS 407 controller (photos by author, 2009)

A series of drivefiles were run for each shock tested. Before each series of tests, a warmup sine test was run at 0.6 inch amplitude and 2.8 Hz to heat up the shock. The first test that was run was a 120 second random signal created by the multisine method described in Chapter 3, with the relative amplitude spectrum shown in Table 4.3. The spectrum was adjusted to be large enough to excite all the desired frequencies, while not causing excessive vibrations in the rig.

Table 4.3. Relative amplitude spectrum for random drivefile

Frequency Hz	Relative Amplitude
0.1	1
1	1
2	0.8
3	0.7
5	0.5
8	0.3
10	0.3
15	0.3
20	0.1

To explore the steady-state response at specific frequencies, six sine tests were run at the amplitudes and frequencies shown in Table 4.4. These frequencies were chosen to include frequencies near, above, and below the sprung and unsprung resonance frequencies. During sine tests, the rig was allowed to reach steady-state, then a 20 second data block was recorded.

Table 4.4. Sine wave tests

Frequency Hz	Amplitude in
1	1
2.8	0.6
3.5	0.5
5	0.4
9	0.4
12	0.3

A bump test was also conducted by collecting a 60 second data block a using a 1 inch square wave with a period of 20 seconds.

Six different shock configurations were tested, as shown in Table 4.5. These include all three shock builds tested in Chapter 3 at both 1 click and 14 clicks on the bleed adjuster and 50 psi gas pressure.

Table 4.5. Shock configurations tested

	Bleed	Comp	HV	Gas
1 Build 1	soft	soft	hard	50 psi
2 Build 1	hard	soft	hard	50 psi
3 Build 2	soft	soft	soft	50 psi
4 Build 2	hard	soft	soft	50 psi
5 Build 4	soft	hard	0	50 psi
6 Build 4	hard	hard	0	50 psi

After the data was collected, it was then processed to prepare it for the ID process. The main processing steps were to estimate the shock velocity and shock force for each test. To make sure the shock velocity was calculated correctly, it was calculated using three different methods

1. Filter the shock potentiometer and differentiate,
2. Calculate the shock displacement using the string potentiometers on the sprung and unsprung masses, differentiate, and filter, and
3. Calculate the shock velocity using the string potentiometer velocity and filter.

Filtering was done by applying an 8th order, zero-phase, Butterworth lowpass filter using MATLAB's *filtfilt* command, which filters the signal forwards and backwards to cancel

phase lag. Numerical differentiation was performed using MATLAB's *gradient* command, which uses central differences to reduce phase shift. These three methods provided very consistent estimates of the shock velocity. Since the velocity calculated from the shock potentiometer measurement will be available on the 8-post rig, we will use it for further analysis. The shock velocity was then used as an input to the corresponding shock model for the test, providing an estimate of the shock force.

4.5.2 Experimental ID

The ID method developed in Section 4.3 requires data from a test with a linear shock and a test with a different shock. Since shock configuration 1 exhibits a near-linear behavior, the data from the random test with shock configuration test was used as the linear data set. The linear damping for shock configuration 1 was estimated to be 31.03 lb/(in/s) using linear regression on the shock dynamometer force-velocity plot. The data from the random test with shock configuration 1 was then used to identify the frequency response between the actuator position and the sensor output with shock configuration 1 installed.

Once this transfer function was determined, the data from the random test for shock configuration 5 was used to determine the transfer function between the change in shock force and the sensor output.

The FRF estimates were calculated using 50 data blocks with a Hamming window using 50% overlap from the 120 second random signal. This corresponds to a spectral resolution of 0.22 Hz. The parametric transfer function estimate was calculated using *invfreqs* on each input-output pair. The numerator and denominator polynomials were chosen to be 4th order, and the coefficients were fitted using the FRF estimate between 0.5 and 15 Hz, corresponding to the frequency range where the FRF estimates had good coherence. The resulting estimates for the sprung and unsprung mass positions are shown in Figure 4.18. The coherence for the FRF estimates is shown in Figure 4.19. The characteristic equation for the transfer function estimate between wheelloader position and chassis position indicated that the first mode was at 3.9 Hz with 54% damping, and the second mode was at 8.6 Hz with 32% damping.

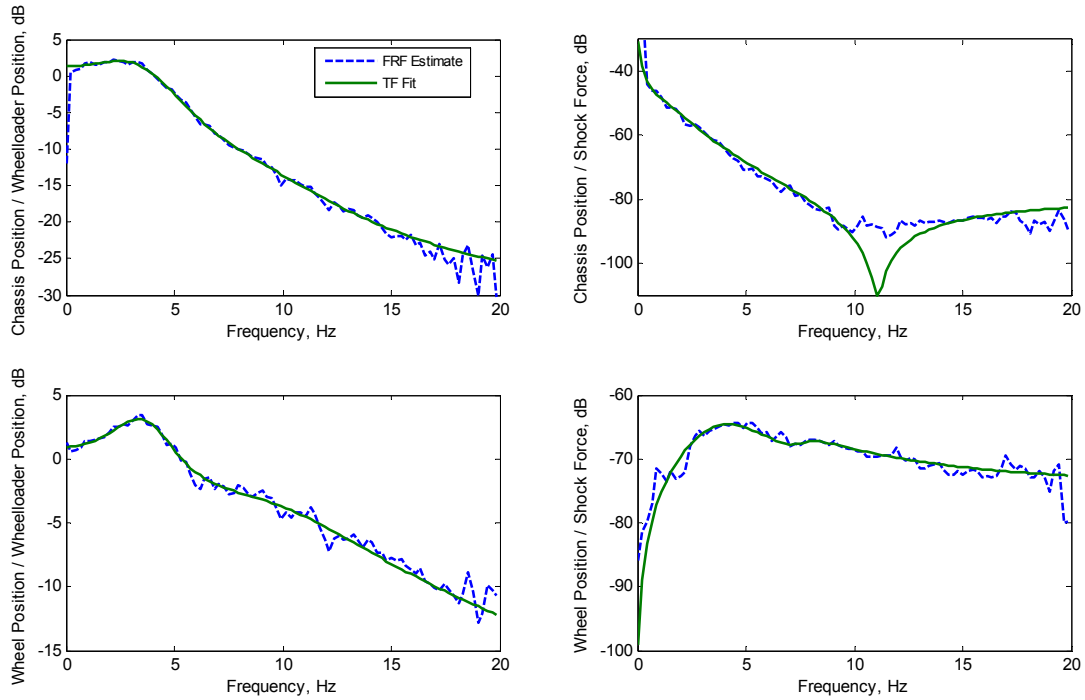


Figure 4.18. Quarter-car rig FRF estimates




Figure 4.19. Coherence for FRF estimates

4.5.3 Validation

Now that we have used quarter-car rig data to create a model of the rig, we can now use the model to predict the response of the quarter-car rig. To validate the model's ability to predict rig response, the predictions are compared against test data. The model was used to run simulations using the same Simulink layout shown in Figure 4.12 using the same drivefiles used on the quarter-car rig.

First, we compare the time signals directly. Since shock configurations 1 and 5 were used to generate the model, we consider shock configuration 3 for validation purposes. Figure 4.20 compares simulation versus test data for a 5 second time signal of the chassis position for the random test with shock configuration 3. As this figure shows, the simulation shows good agreement with the experimental data. Similar agreement exists with the other sensors.




Figure 4.20. Comparison of simulated and measured time signals

Next, we compare how well our model is able to predict the change in different response measures as the shocks are changed. Figure 4.21 compares measured and predicted response measures for four different responses. While there is some error in predictions, the simulation is generally effective at predicting the trend. Since the trend is what is most commonly used in rig testing, it seems reasonable to use this model to do preliminary shock studies to make rig testing more efficient.

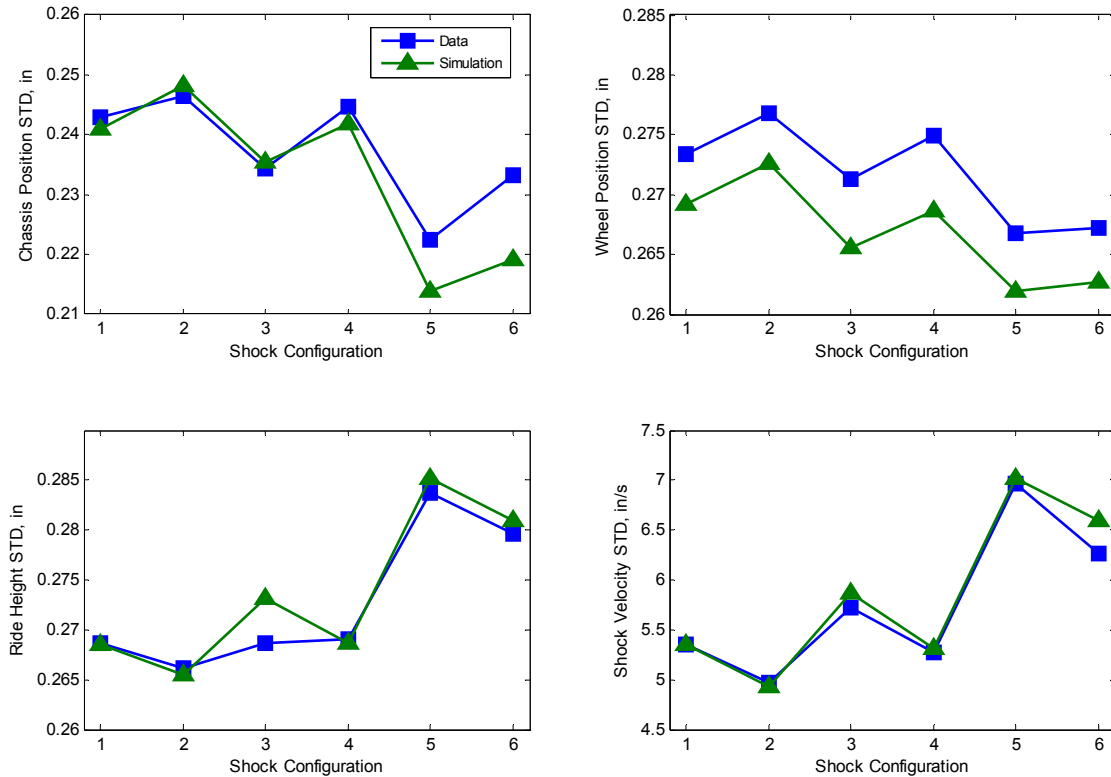


Figure 4.21. Comparison of simulated versus measured response measures

4.6 Issues with Extending to Full-Vehicle Modeling

While the quarter-car rig results presented here are promising, extending this method from this ideal environment to work for 8-post rig test data presents further challenges that should be considered. First, there are many more potential sources of nonlinearity on the 8-post rig. These include tire separation, coil bind, coil separation, progressive spring rates, bump stops, nonlinear motion ratios, tire and joint friction, and play in the suspension. Also a concern is that there are more inputs to consider, including inputs that will be filtered across the chassis and four shocks whose force is indirectly excited from the rig actuators.

Despite these potential concerns, there are reasons to suggest that we may overcome these issues and be successful modeling the 8-post rig. First, the RMS-based response measures are gross dynamic measures of the response, so even our model does not match

the ideal response exactly, it is still possible that our RMS trends will match. Linear modeling has already been used to model the response of an 8-post rig test with positive results during drivefile iteration. The weight of the vehicle and aerolader downforce will act to preload the tires, reducing the amount of time that the tires will separate from the wheelloaders. Suspension joint friction and slop will often be minimal to provide good performance.

One area where we will improve as we move to 8-post rig scale is the shock force estimates. In this chapter, significant error may be caused by the estimation of the shock force used in the ID process. Any error in the shock model or the velocity estimate would cause an error in the shock force estimate, which would create an error in the model. For the 8-post study we will directly measure shock force using a shock load cell.

Another area of improvement is the requirement that the baseline shocks are linear. We will develop the identification method such that it will work well for both linear and nonlinear shocks.

Chapter 5

Full-Vehicle Modeling and Identification

Chapters 5 and 6 extend the quarter-car modeling and identification work presented in Chapter 4 to full-vehicle scale. This chapter focuses on performing system identification using data generated from a known simulation model. Since we know the ideal model, we will be able to verify that our identification methods accurately match the ideal model behavior. Once we have shown that our identification methods can accurately match the ideal model, we will then apply our methods to data collected from an 8-post rig experiment in Chapter 6.

5.1 Full-Vehicle Model

This section presents the full-vehicle model which will be used as the ideal model in this chapter. First, the linear full-vehicle model is derived. Important model dynamics are then characterized by showing select frequency response plots. Next, we incorporate the nonlinear shock models developed in Chapter 3 into our model. Once our model has been fully assembled, we will then perform simulations using the model. The resulting data sets will be used to perform system identification to generate a model that will be compared with the original model.

5.1.1 Model Derivation

The vehicle may be modeled as 5 rigid bodies, representing the chassis and the 4 wheels, which are connected to the chassis through the suspension members. Suspensions are typically designed to allow for only vertical wheel motion relative to the chassis, so the wheel rigid bodies can be simplified to a vertical mass. As discussed before, the 8-post rig applies forces primarily in a vertical direction and constrains vehicle motion in the yaw plane by constraining the tire contact patch to remain within the wheelloader

constraints. As a result of this, the yaw plane dynamics (lateral, longitudinal, yaw) are not excited significantly in comparison with the vertical dynamics (heave, pitch, roll). This suggests treating the chassis as having 3 inertial terms corresponding to the heave, roll, and pitch degrees of freedom. The above reasoning motivates the selection of a 7 degree-of-freedom (7 DOF) model, with degrees of freedom for heave, pitch, and roll of the chassis and vertical motion of the four tires.

The derivation of the 7 DOF model begins by considering the effect of a suspension force on the sprung mass and unsprung mass. Once the general formulation is complete, this will then be applied to model the effect of different suspension components, including primary suspension springs, tires, shocks, and rollbars. Upon completing the formulation for all the desired components, the forces due to each component will be combined using superposition to form the final equations of motion.

We begin with the diagram of the chassis shown in Figure 5.1. The coordinate system is chosen about the center of gravity, which lies a distance a behind the front axle. We assume the center of gravity to lie along the vehicle centerline and the front and rear track widths to be equal. The influence of external forces applied to the 3 DOF chassis model by the aeroloaders may be represented as a heave force, roll moment, and pitch moment. The compressive force due to the suspension at each corner is also shown.




Figure 5.1. Chassis geometry and forces

The geometric transformation between the chassis heave, pitch, and roll and the displacement of the four corners of the chassis is

$$\begin{bmatrix} z_1^{LF} \\ z_1^{RF} \\ z_1^{LR} \\ z_1^{RR} \end{bmatrix} = \begin{bmatrix} 1 & \frac{t}{2} & -a \\ 1 & -\frac{t}{2} & -a \\ 1 & \frac{t}{2} & b \\ 1 & -\frac{t}{2} & b \end{bmatrix} \begin{bmatrix} z \\ \theta \\ \phi \end{bmatrix}$$

(5.1)

$$z_1 = T\Phi$$

Summation of forces and moments on the chassis yields

$$\begin{bmatrix} M & 0 & 0 \\ 0 & J_\theta & 0 \\ 0 & 0 & J_\phi \end{bmatrix} \begin{bmatrix} \ddot{z} \\ \ddot{\theta} \\ \ddot{\phi} \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ \frac{t}{2} & -\frac{t}{2} & \frac{t}{2} & -\frac{t}{2} \\ -a & -a & b & b \end{bmatrix} \begin{bmatrix} F_{LF} \\ F_{RF} \\ F_{LR} \\ F_{RR} \end{bmatrix} + \begin{bmatrix} F_z \\ M_\theta \\ M_\phi \end{bmatrix}$$

(5.2)

$$M_c \ddot{\Phi} = T^T F_c + F_{ext}$$

For the unsprung masses, we define positive displacement to be upwards. This yields the equation

$$\begin{bmatrix} m_t^{LF} & 0 & 0 & 0 \\ 0 & m_t^{RF} & 0 & 0 \\ 0 & 0 & m_t^{LR} & 0 \\ 0 & 0 & 0 & m_t^{RR} \end{bmatrix} \begin{bmatrix} z_2^{LF} \\ z_2^{RF} \\ z_2^{LR} \\ z_2^{RR} \end{bmatrix} = - \begin{bmatrix} F_{LF} \\ F_{RF} \\ F_{LR} \\ F_{RR} \end{bmatrix} \quad (5.3)$$

$$M_t \ddot{z}_2 = -F_c$$

Equations (5.2) and (5.3) define the effect of a generic compressive suspension force on the sprung and unsprung masses. Next, we will consider the equations for linear springs, tires, shocks, and anti-rollbars.

For a linear primary suspension spring, the suspension force is proportional to the suspension displacement. This can be represented for all four corners as

$$\begin{bmatrix} F_{LF} \\ F_{RF} \\ F_{LR} \\ F_{RR} \end{bmatrix} = \begin{bmatrix} k^{LF} & 0 & 0 & 0 \\ 0 & k^{RF} & 0 & 0 \\ 0 & 0 & k^{LR} & 0 \\ 0 & 0 & 0 & k^{RR} \end{bmatrix} \left\{ \begin{bmatrix} z_2^{LF} \\ z_2^{RF} \\ z_2^{LR} \\ z_2^{RR} \end{bmatrix} - \begin{bmatrix} z_1^{LF} \\ z_1^{RF} \\ z_1^{LR} \\ z_1^{RR} \end{bmatrix} \right\} \quad (5.4)$$

$$F_c = K_s (z_2 - z_1)$$

A linear tire applies a force to the unsprung mass proportional to the tire deflection, expressed by

$$\begin{bmatrix} F_t^{LF} \\ F_t^{RF} \\ F_t^{LR} \\ F_t^{RR} \end{bmatrix} = \begin{bmatrix} k_t^{LF} & 0 & 0 & 0 \\ 0 & k_t^{RF} & 0 & 0 \\ 0 & 0 & k_t^{LR} & 0 \\ 0 & 0 & 0 & k_t^{RR} \end{bmatrix} \left\{ \begin{bmatrix} z_R^{LF} \\ z_R^{RF} \\ z_R^{LR} \\ z_R^{RR} \end{bmatrix} - \begin{bmatrix} z_2^{LF} \\ z_2^{RF} \\ z_2^{LR} \\ z_2^{RR} \end{bmatrix} \right\} \quad (5.5)$$

$$F_t = K_t (z_R - z_2)$$

Combining Equations (5.1), (5.2), and (5.4) for the sprung mass yields

$$\begin{aligned}
 M_c \ddot{\Phi} &= -T^T K_s z_1 + T^T K_s z_2 + F_{ext} \\
 &= -T^T K_s T \Phi + T^T K_s z_2 + F_{ext} \\
 M_c \ddot{\Phi} &= \begin{bmatrix} -T^T K_s T & T^T K_s \end{bmatrix} \begin{bmatrix} \Phi \\ z_2 \end{bmatrix} + F_{ext}
 \end{aligned} \tag{5.6}$$

Combining Equations (5.1), (5.3), (5.4), and (5.5) for the unsprung mass gives

$$\begin{aligned}
 M_t \ddot{z}_2 &= -F_c + F_t \\
 &= -K_s (z_2 - z_1) + K_t (z_R - z_2) \\
 &= K_s z_1 - (K_s + K_t) z_2 + K_t z_R \\
 &= K_s T \Phi - (K_s + K_t) z_2 + K_t z_R \\
 M_t \ddot{z}_2 &= \begin{bmatrix} K_s T & -(K_s + K_t) \end{bmatrix} \begin{bmatrix} \Phi \\ z_2 \end{bmatrix} + K_t z_R
 \end{aligned} \tag{5.7}$$

Equations (5.6) and (5.7) can be combined to provide

$$\begin{aligned}
 \begin{bmatrix} M_c & 0 \\ 0 & M_t \end{bmatrix} \begin{bmatrix} \ddot{\Phi} \\ \ddot{z}_2 \end{bmatrix} + \begin{bmatrix} T^T K_s T & -T^T K_s \\ -K_s T & K_s + K_t \end{bmatrix} \begin{bmatrix} \Phi \\ z_2 \end{bmatrix} &= \begin{bmatrix} I & 0 \\ 0 & K_t \end{bmatrix} \begin{bmatrix} F_{ext} \\ z_R \end{bmatrix} \\
 \bar{M} \ddot{Z} + K_p Z &= \begin{bmatrix} I & 0 \\ 0 & K_t \end{bmatrix} \begin{bmatrix} F_{ext} \\ z_R \end{bmatrix}
 \end{aligned} \tag{5.8}$$

The equations for linear shocks in parallel with the suspension and tire springs can easily be derived using similarity to the above equations. Assuming tire damping to be negligible, the damping matrix is

$$C_p = \begin{bmatrix} T^T C_s T & -T^T C_s \\ -C_s T & C_s \end{bmatrix} \tag{5.9}$$

where C_s is a diagonal matrix containing the damping coefficients for the four corners. The resulting equations of motion including linear springs, shocks, and tire stiffnesses are

$$\bar{M}\ddot{Z} + C_p\dot{Z} + K_p Z = \begin{bmatrix} I & 0 \\ 0 & K_t \end{bmatrix} \begin{bmatrix} F_{ext} \\ z_R \end{bmatrix} \quad (5.10)$$

The next component that we wish to include is the rollbar. The purpose of an anti-rollbar is to increase the roll rate of an axle without increasing the vertical ride rate. The anti-rollbar on an axle does this by applying a restoring roll moment proportional to the chassis roll relative to the axle roll. The forces applied by the anti-rollbar for a positive roll angle are shown in Figure 5.2.




Figure 5.2. Forces applied due to anti-rollbar

The relative roll angle is defined as

$$\theta_{rel} = \theta - \theta_{axle} = \theta - \left(\frac{z_2^L - z_2^R}{t} \right) \quad (5.11)$$

$$\theta_{rel} = \begin{bmatrix} 1 & -\frac{1}{t} & \frac{1}{t} \end{bmatrix} \begin{bmatrix} \theta \\ z_2^L \\ z_2^R \end{bmatrix}$$

The force applied to the suspension by the anti-rollbar can be found by

$$F = \frac{1}{t} (k_R \theta_{rel})$$

$$= \begin{bmatrix} \frac{k_R}{t} & -\frac{k_R}{t^2} & \frac{k_R}{t^2} \end{bmatrix} \begin{bmatrix} \theta \\ z_2^L \\ z_2^R \end{bmatrix} \quad (5.12)$$

It can be seen that the anti-rollbar applies a compressive force to the right side of the chassis and a tension force to the left side of the chassis for a positive relative roll angle.

We can then write the suspension forces due to the anti-rollbar in terms of the vehicle position vector as

$$\begin{bmatrix} F_{LF} \\ F_{RF} \\ F_{LR} \\ F_{RR} \end{bmatrix} = \begin{bmatrix} 0 & -\frac{k_R^F}{t} & 0 & \frac{k_R^F}{t^2} & -\frac{k_R^F}{t^2} & 0 & 0 \\ 0 & \frac{k_R^F}{t} & 0 & -\frac{k_R^F}{t^2} & \frac{k_R^F}{t^2} & 0 & 0 \\ 0 & -\frac{k_R^R}{t} & 0 & 0 & 0 & \frac{k_R^R}{t^2} & -\frac{k_R^R}{t^2} \\ 0 & \frac{k_R^R}{t} & 0 & 0 & 0 & -\frac{k_R^R}{t^2} & \frac{k_R^R}{t^2} \end{bmatrix} \begin{bmatrix} \Phi \\ z_2 \end{bmatrix} \quad (5.13)$$

$$F_c = K_1 Z$$

where k_R^F and k_R^R are the front and rear anti-rollbar stiffnesses. Using Equations (5.2) and (5.3), the contribution to the stiffness matrix due to the anti-rollbar is

$$K_{roll} = \begin{bmatrix} -T^T K_1 \\ K_1 \end{bmatrix} \quad (5.14)$$

This stiffness matrix can be added to the existing stiffness matrix in Equation (5.10) to include the effects of rollbars. This gives us the new equations of motion

$$\bar{M}\ddot{Z} + C\dot{Z} + (K_p + K_{roll})Z = \begin{bmatrix} I & 0 \\ 0 & K_t \end{bmatrix} \begin{bmatrix} F_{ext} \\ z_R \end{bmatrix} \quad (5.15)$$

The last component that we wish to include is an arbitrary suspension force. This will be used to include nonlinear shock models in our simulations, but it could also be used for any other arbitrary suspension force. Defining our external suspension force to be F_D , we can find the effect of an arbitrary suspension force on the vehicle response using Equations (5.2) and (5.3) as

$$\begin{bmatrix} M_c & 0 \\ 0 & M_t \end{bmatrix} \begin{bmatrix} \ddot{\Phi} \\ \ddot{z}_2 \end{bmatrix} = \begin{bmatrix} T^T \\ -I \end{bmatrix} F_D \quad (5.16)$$

Combining Equations (5.15) and (5.16) yields the final equation of motion

$$\bar{M}\ddot{Z} + C_p\dot{Z} + (K_p + K_{roll})Z = \begin{bmatrix} I & 0 & T^T \\ 0 & K_t & -I \end{bmatrix} \begin{bmatrix} F_{ext} \\ z_R \\ F_D \end{bmatrix} \quad (5.17)$$

$$\bar{M}\ddot{Z} + C_p\dot{Z} + KZ = Lu$$

Where the stiffness and damping matrices can be found using

$$K = \begin{bmatrix} T^T K_s T & -T^T K_s \\ -K_s T & K_s + K_t \end{bmatrix} + \begin{bmatrix} -T^T K_1 \\ K_1 \end{bmatrix} \quad (5.18)$$

$$C_p = \begin{bmatrix} T^T C_s T & -T^T C_s \\ -C_s T & C_s \end{bmatrix}$$

The matrices K_s , K_t , and C_s are 4x4 diagonal matrices containing the spring stiffness, tire stiffness, and linear damping for the four corners. The matrix K_l is defined in Equation (5.13), and the transformation matrix T is defined in Equation (5.1). The mass matrix is a 7x7 diagonal matrix containing the inertia of the 7 degrees of freedom.

The equations of motion can be rewritten in state space form as

$$x_1 = Z$$

$$x_2 = \dot{Z}$$

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} 0 & I \\ -\bar{M}^{-1}K & -\bar{M}^{-1}C_p \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} 0 \\ \bar{M}^{-1}L \end{bmatrix} u \quad (5.19)$$

$$\dot{x} = Ax + Bu$$

The matrices A and B can be calculated using Equations (5.1), (5.13), and (5.17)-(5.19).

In addition to defining the dynamic equations, we must also define the output equation

$$y = Cx + Du \quad (5.20)$$

Where l is the wheelbase, l_1 is the distance from the front axle to the nose, and l_2 is the distance from the rear axle to the spoiler.

In addition to the output equations shown above, additional output equations were constructed for the state vector, suspension deflections, chassis accelerations, wheel accelerations, and spring forces.

The model parameters used are shown in Table 5.1. The sprung and unsprung masses correspond to a 3000 lb car with 85% sprung mass, and 45% of the unsprung mass in the front. The roll and pitch radii of gyration are 40% and 60% of the track and wheelbase, respectively. The suspension and tire stiffnesses are within the range of typical values. While these parameters do not precisely reflect those of a NASCAR Cup car, they do describe a vehicle whose behavior is qualitatively similar, making it sufficient for ID method development.

Table 5.1. Model parameters

Track Width	70 in
CG-Front Axle	60 in
CG-Rear Axle	50 in
Sprung Mass	2550 lbm
Roll Inertia	1.5E+06 lbm-in ²
Pitch Inertia	5.5E+06 lbm-in ²
Front Unsprung Mass	101 lbm
Rear Unsprung Mass	124 lbm
Front Spring Rate	1280 lbf/in
Rear Spring Rate	588 lbf/in
Tire Stiffness	2000 lbf/in
Linear Damper	30 lbf/(in/s)

5.1.2 Frequency Response

This section discusses the frequency response for the linear vehicle model presented in the previous section. These plots provide insight on system behavior, which can be applied during the identification process. The outputs considered include shock velocity, tire force, and ride height. We only consider 3 of the 11 potential inputs, but the observations below may be generalized to the remaining 8 inputs. The inputs considered

include chassis heave force, left front wheelloader position, and left front external shock force.

The first set of outputs that we consider are the shock velocities, as shown in Figure 5.3. For the chassis force input, 1000 lb chassis heave force gives 10 in/s near resonance, and 1 in/s below 1 Hz and above 9 Hz. Further chassis velocity can be attained by aeroloader force by applying roll and pitch moments.

The left front wheelloader input creates about 1 in/s shock velocity at low frequency and 30 in/s shock velocity near resonance for a 1 inch left front wheelloader input. The left front shock velocity frequency response maintains a high level at higher frequencies due to the proximity of the wheelloader to the shock. All other shock velocities drop about 20 dB at higher frequencies due to the filtering effect of the chassis.

The left front external shock force produces about 1 in/s left front shock velocity for every 100 lb. The other shock velocities have similar levels at lower frequencies, but the level drops 40 dB at higher frequencies due to the filtering effects of the chassis.




Figure 5.3. Shock velocity frequency response

The next set of outputs that we consider is the tire forces. For the chassis heave force input, 1000 lb of chassis heave force causes about 500 lb tire force near resonance, 300 lb of force at low frequency, and 100 lb at 10 Hz.

For the left front wheelloader position input, a 1 inch input produces 100 lb of tire force at low frequency, and as much as 3000 lb near resonance. As frequency increases from resonance, the left front tire force maintains a high level, while the other tire forces drop due to the filtering effect of the chassis.

For the left front external shock force input, a 100-lb shock force input creates 10 lb tire force at low frequency, and as much as 100 lb at resonance. The left front tire force stays at a high level at higher frequencies, while the other tire forces drop 40 dB.




Figure 5.4. Tire force frequency response

The last set of outputs we consider is the front and rear ride heights. For the chassis heave force, 1000 lb of heave force provides 0.3-1 inches of ride height variation from 0-5 Hz. Above 5 Hz, the response rolls off at about 40 dB/decade.

For the left front wheelloader input, the response is small at lower frequencies as the chassis moves in phase with the wheelloaders. At 5 Hz, the response peaks, giving 1 inch front ride height variation and 0.3 inches of rear ride height variation for a 1 inch input. At higher frequencies, the front ride height maintains a high response level, while the rear ride height drops to -40 dB due to the filtering effect of the chassis.




Figure 5.5. Ride height frequency response

These observations can be used to assist our identification efforts for these key outputs. To successfully perform nonparametric frequency response estimation, we need to have significant output energy at all frequencies of interest relative to the noise level. This will

provide a good signal-to-noise ratio, or coherence near 1 for a linear system. To design excitation to provide good output signal levels, the system behavior should be considered.

Notice that the chassis heave force is effective for exciting response below 5 Hz, but the response drops off at higher frequencies. This trend is also true for chassis roll and pitch moments. Conversely, the wheelloaders are often ineffective at exciting outputs below 3 Hz but are useful for higher frequency excitation. Also, the most effective excitation at higher frequency will be the wheelloader closest to the sensor of interest. This suggests the need for excitation from both the closest wheelloader and the aeroloaders if we want a good signal-to-noise ratio from 1 to 20 Hz.

While it is ideal to get high signal-to-noise ratios in all output signals at all frequencies, this is often impractical due to the system's behavior, physical input limitations, and vehicle response limitations. If the output level of a particular sensor due to a particular input will always be low over a range of frequencies for the desired excitation signals to be used with the model, the poor frequency response estimation for that input-output pair in that frequency range may be replaced by a small or zero frequency response with minimal reduction in model accuracy from the ideal model.

Notice that accurate frequency response estimation for aeroloader inputs is only critical at lower frequencies. At higher frequencies, the response is attenuated, so a poor frequency response estimate could be replaced by a small or zero frequency response to reduce modeling error.

For the wheelloader inputs and shock force inputs, the frequency response is highest near resonance and remains at a high level at higher frequencies for input-output that are directly connected within one corner and not connected through the chassis. This indicates that frequency response estimation is most important for wheelloader and shock force inputs near resonance. It is also important for directly-coupled input-output pairs at higher frequencies. Modeling error due to low response levels outside these regions can be replaced with a small or zero frequency response.

This section has summarized the full-vehicle model’s behavior for the key outputs of shock velocity, tire force, and ride height. These results can then be applied to design excitation signals for identification tests and to modify the identification process where frequency response estimation is poor.

5.1.3 Full-Vehicle Simulation

Now that we have developed a full-vehicle model and explored the behavior of the model, we will now apply the model to generate data sets that can be used to develop the identification method. Since we already know the ideal model, performing identification on data generated by the model will allow us to validate our identified model with the ideal model, which will allow us to correct any issues in the identification process.

First, we need to incorporate the shock absorber models developed in Chapter 3 into the vehicle model. This is accomplished by applying the shock velocity calculated from the linear vehicle model to the nonlinear dynamic shock model, as shown in Figure 5.6. The shock force calculated by the shock model is then fed back as an input to the linear vehicle model using the external shock force input. This method is identical to the method in Chapter 4, except there are 4 shock models instead of just one.

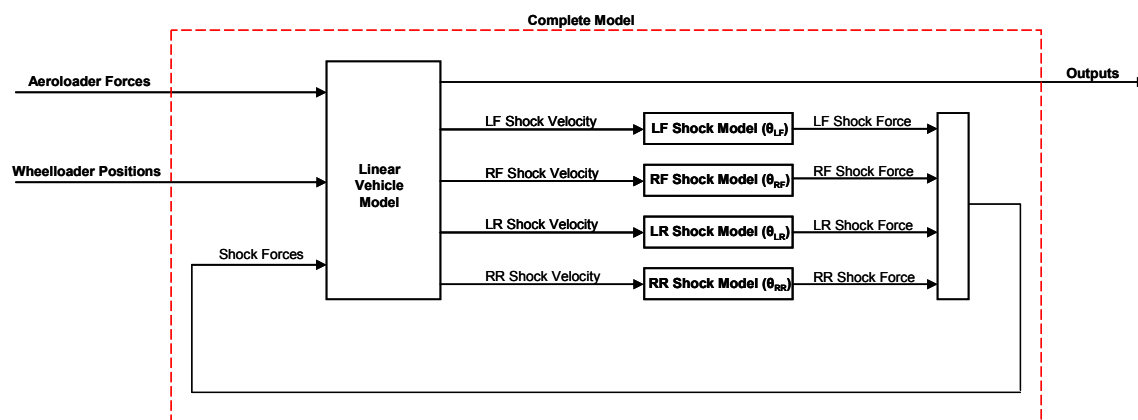


Figure 5.6. Coupling of vehicle and shock absorber models

This model structure can be easily implemented using Simulink. The equations for the LPNL2 nonlinear dynamic shock absorber are described in Chapter 3. To implement the Simulink model in discrete-time, the linear vehicle model and the shock model transfer

functions are discretized at 100 Hz using Tustin's bilinear transformation, available in MATLAB's *c2d* function. To make Simulink model more organized, the shock models are arranged in subsystems, as shown in Figure 5.7.

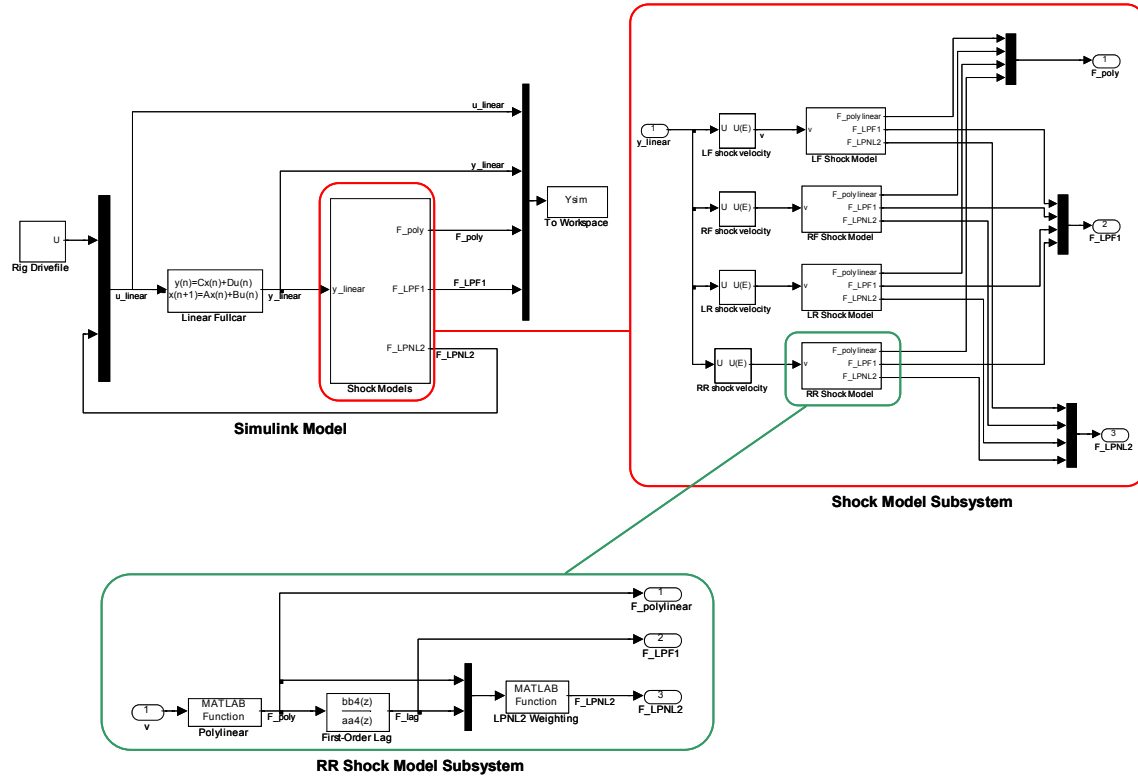


Figure 5.7. Full-vehicle model implemented in Simulink

Five simulations were performed to generate data for use in system identification. An 800 second Gaussian bandlimited white noise drivefile was constructed to excite the Simulink model. Each signal was bandlimited from 0 to 30 Hz using an 8th order discrete-time Butterworth filter. The peak signal amplitudes were set to 1 inch for the wheelloaders and 500 lb for the chassis heave force. The chassis roll moment and chassis pitch moment amplitudes were set to be 3 times the chassis roll and pitch stiffness respectively in in-lb/degree. This drivefile was applied to the Simulink model to generate data that will be used to perform system identification on the linear vehicle model. One baseline simulation was performed with a linear shock with a damping coefficient of 20 lb/(in/s) and four additional simulations were performed where the baseline shock at each corner was replaced by a nonlinear dynamic shock model with a damping coefficient of 8 lb/(in/s).

5.2 System Identification Methods

This section presents system identification methods that will be applied to the data generated using the full-vehicle model. The presentation here builds upon the work presented for the quarter-car study in Section 4.3 and extends it for multiple shocks. To analyze the identification methods, it will be convenient to partition the vehicle model derived in Section 5.1 as shown in Figure 5.8. This diagram partitions the linear system that we wish to identify into the response caused by the actuators and the additional response caused by the external shock force, which is defined by a feedback loop, as shown in Figure 5.8.

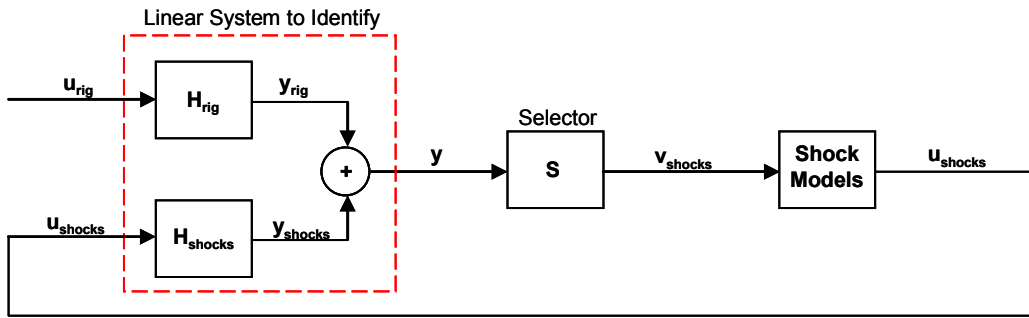


Figure 5.8. Diagram of system to identify

The B matrix in the state-space vehicle model in Equation (5.19) may be partitioned by the columns corresponding to inputs due to the rig and inputs due to the external shock force as

$$\begin{aligned}\dot{x} &= Ax + B \begin{bmatrix} F_{ext} \\ z_R \\ F_D \end{bmatrix} \\ &= Ax + B_{rig} \begin{bmatrix} F_{ext} \\ z_R \end{bmatrix} + B_{shocks} F_D\end{aligned}\tag{5.24}$$

$$\dot{x} = Ax + B_{rig} u_{rig} + B_{shocks} u_{shocks}$$

The D matrix in the state-space output equation in Equation (5.20) may be similarly partitioned as

$$\begin{aligned}
y &= Cx + Du \\
&= Cx + D_{rig}u_{rig} + D_{shocks}u_{shocks}
\end{aligned} \tag{5.25}$$

The transfer functions for this partitioned system may be found by taking the Laplace transform of Equation (5.24) and combining it with Equation (5.25), which gives us

$$\begin{aligned}
y(s) &= \left[C(sI - A)^{-1} B_{rig} + D_{rig} \right] u_{rig} + \left[C(sI - A)^{-1} B_{shocks} + D_{shocks} \right] u_{shocks} \\
&= H_{rig}(s)u_{rig}(s) + H_{shocks}(s)u_{shocks}(s)
\end{aligned} \tag{5.26}$$

The shock velocities may be selected from the output vector using a selection matrix

$$v_{shocks} = Sy \tag{5.27}$$

The above notation will be used to discuss the identification methods.

5.2.1 Nonparametric Frequency Response Estimation

If all the inputs were independent, a common practice is to create a nonparametric estimate of the frequency response matrix, which consists of a complex number for each input, output, and frequency. For each frequency and output, we wish to fit the model

$$y = H_{rig}u_{rig} + H_{shocks}u_{shocks} \tag{5.28}$$

where the inputs and outputs consist of the frequency domain component for a given frequency. To fit the model, we collect several data blocks and select the model coefficients that minimize some measure of model error. Data blocks may be taken from a series of tests or from selecting multiple blocks of data from a longer data set. The measured output in the k^{th} data block will contain some error from the ideal model's estimate, which may be written as

$$y[k] = H_{rig}u_{rig}[k] + H_{shocks}u_{shocks}[k] + \varepsilon[k] \tag{5.29}$$

If we collect K data blocks, we can arrange the results in matrix form as

$$\begin{bmatrix} y[1] \\ \vdots \\ y[K] \end{bmatrix} = \begin{bmatrix} u^T_{rig}[1] & u^T_{shocks}[1] \\ \vdots & \vdots \\ u^T_{rig}[K] & u^T_{shocks}[K] \end{bmatrix} \begin{bmatrix} H^T_{rig} \\ H^T_{shocks} \end{bmatrix} + \begin{bmatrix} \varepsilon[1] \\ \vdots \\ \varepsilon[K] \end{bmatrix} \quad (5.30)$$

$$Y = UH + E$$

We can choose to select our frequency response estimate to minimize the error measure

$$J = E^*E \quad (5.31)$$

This is a least-squares optimization problem, whose solution satisfies the normal equations

$$\begin{aligned} (U^*U)H &= U^*Y \\ S_{uu}H &= S_{uy} \end{aligned} \quad (5.32)$$

The frequency response estimation produced by this method is often called a H_I frequency response estimation. To find a unique solution for the frequency response matrix H in Equation (5.32), the autospectrum matrix must be nonsingular. A necessary condition for the autospectrum matrix to be nonsingular is that the input matrix U defined in Equation (5.30) must be full rank, which is rank 11 in our case. This requires that there must be at least 11 linearly independent input block rows and 11 linearly independent input columns to solve for H .

For the simple case of linear shocks, we can show that the shock force inputs will be linearly dependent on the rig inputs. For linear shocks, the shock force can be written

$$\begin{aligned} u_{shocks} &= \bar{C}v_{shocks} \\ &= \bar{C}Sy \end{aligned} \quad (5.33)$$

where \bar{C} is a 4x4 diagonal matrix whose diagonal entries contain the 4 shock damping coefficients. Substituting Equation (5.28) and solving for the shock force yields

$$\begin{aligned}
u_{shocks} &= \bar{C}Sy \\
&= \bar{C}S(H_{rig}u_{rig} + H_{shocks}u_{shocks}) \\
(I - \bar{C}SH_{shocks})u_{shocks} &= \bar{C}SH_{rig}u_{rig} \\
u_{shocks} &= (I - \bar{C}SH_{shocks})^{-1} \bar{C}SH_{rig}u_{rig} \\
u_{shocks} &= H_A u_{rig}
\end{aligned} \tag{5.34}$$

This shows that the shock force is linearly dependent on the rig inputs when the shocks are linear, so no unique solution to Equation (5.32) exists when this frequency response estimation is applied in a straightforward manner. Physically, this means that since the shock inputs are linearly related to the rig inputs, it is impossible to determine if the outputs are linearly related to the shock inputs or the rig inputs. This argument can be extended to shocks that exhibit linear dynamic behavior by allowing the \bar{C} matrix to be complex.

If the shock is nonlinear, the nonlinear portion of the shock force provides some shock force excitation that is linearly independent from the rig inputs. If the nonlinear contribution is small, the output caused by the independent portion of the shock force will be small, causing poor coherence. This suggests one option to improve coherence is to use highly nonlinear shock absorbers, but it may not be practical to test with a shock with enough nonlinearity to achieve the desired model quality.

The observation that it is impossible to identify the shock input frequency response using conventional procedures and that poor coherence is expected when the shock is close to linear motivate developing alternate identification methods.

5.2.2 Frequency Response Estimation for Shock Force Inputs

To get good estimates of the frequency response between shock force inputs and sensor outputs, more shock force signal energy must be linearly independent from the rig inputs. This section provides a summary of potential methods to get quality shock force input frequency response estimations.

The simplest method of applying uncorrelated shock force is to test with a highly nonlinear shock. This will only be successful if the nonlinear portion of the shock force can excite outputs with good signal-to-noise ratio and the nonlinear shock does not excite vehicle nonlinearities, invalidating the linear vehicle model assumption.

The most direct method of applying linearly independent shock force is to replace the shock with an active or semi-active suspension element and applying independent control inputs. Alternately, different control policies could be applied for different data blocks [38, 39]. This method has the potential for good quality frequency response estimates, but it requires preparation of active or semi-active suspension members to be mounted in the place of the shocks, limiting the convenience of applying this method to an arbitrary vehicle during normal 8-post testing operations.

A convenient method motivated by the option of applying multiple control policies is to test with multiple shocks. This method is equivalent to applying multiple control policies using an active element, where the control policy is limited by the performance that can be achieved by a passive shock. This will produce linearly independent input columns in the U matrix even if the shocks are linear.

To prove that this will allow for unique identification even if the shocks are linear, we will describe an experimental procedure that will create an input matrix that is full rank. First let us partition the shock force into a linear shock force and an additional arbitrary shock force

$$\begin{aligned} u_{shocks} &= \bar{C}Sy + (u_{shocks} - \bar{C}Sy) \\ &= \bar{C}Sy + \Delta F \end{aligned} \tag{5.35}$$

where \bar{C} is a 4x4 diagonal matrix containing baseline linear damping coefficients that are free to be set as desired. The arbitrary portion of the shock force may be written as

$$\Delta F = \begin{bmatrix} \Delta f_1 \\ \Delta f_2 \\ \Delta f_3 \\ \Delta f_4 \end{bmatrix} = \begin{bmatrix} u_{shocks}^1 - \bar{c}_1 v_s^1 \\ u_{shocks}^2 - \bar{c}_2 v_s^2 \\ u_{shocks}^3 - \bar{c}_3 v_s^3 \\ u_{shocks}^4 - \bar{c}_4 v_s^4 \end{bmatrix} \quad (5.36)$$

where u_{shocks}^i , \bar{c}_i , and v_s^i are the shock force, baseline damping coefficient, and the shock velocity of the i^{th} shock. Equation (5.28) can then be rewritten

$$\begin{aligned} y &= H_{rig} u_{rig} + H_{shocks} (\bar{C} S y + \Delta F) \\ (I - H_{shocks} \bar{C} S) y &= H_{rig} u_{rig} + H_{shocks} \Delta F \\ y &= (I - H_{shocks} \bar{C} S)^{-1} (H_{rig} u_{rig} + H_{shocks} \Delta F) \\ y &= \tilde{H}_{rig} u_{rig} + \tilde{H}_{shocks} \Delta F \end{aligned} \quad (5.37)$$

To identify both \tilde{H}_{rig} and \tilde{H}_{shocks} , we need to excite both the rig inputs and the change in shock force input. Let us collect 7 baseline data sets with linear shocks whose damping coefficients are defined by \bar{C} . Further, let the drivefiles for all 7 data sets be linearly independent at all frequencies, such that

$$U_{rig} = \begin{bmatrix} u_{rig}^T [1] \\ \vdots \\ u_{rig}^T [7] \end{bmatrix} \quad (5.38)$$

is full rank for all frequencies. Since the damping force for these 7 baseline data blocks is $\bar{C} S y$, the corresponding ΔF is zero.

Next, we collect 4 additional data sets, using alternate shock selections for each data block. For the first data block, we remove the linear left front shock and replace it with a shock with a different damping coefficient c'_1 , such that

$$\Delta F = \begin{bmatrix} \Delta f_1 \\ 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} \Delta c_1 v_s^1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \quad \Delta c_1 = c_1' - \bar{c}_1 \neq 0 \quad (5.39)$$

The arbitrary contribution of the shock force Δf_1 will be nonzero at all frequencies that v_s^1 is nonzero. It is possible to excite the vehicle with a drivefile u^T_1 such that the shock velocity v_s^1 is nonzero at all frequencies except at frequencies where all inputs have a zero FRF. Since the shock velocity frequency response consists of a ratio of polynomials with a numerator polynomial of order 13 or less, there will be at most 13 frequencies where the shock velocity must be zero. In practice, there will be much less frequencies where the shock velocity must be zero, and the presence of a poor FRF estimate at a few zeroes does not significantly degrade the model quality.

This is repeated for all four shock locations, where the shock at the present location is replaced by an arbitrary shock and all other shocks are the baseline shocks. Collecting the data from all 11 experiments yields the system of equations

$$\begin{bmatrix} y_0[1] \\ \vdots \\ y_0[7] \\ y_1 \\ y_2 \\ y_3 \\ y_4 \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ U_{rig} & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 \\ u^T_1 & \Delta f_1 & 0 & 0 \\ u^T_2 & 0 & \Delta f_2 & 0 \\ u^T_3 & 0 & 0 & \Delta f_3 \\ u^T_4 & 0 & 0 & \Delta f_4 \end{bmatrix} \begin{bmatrix} \tilde{H}_{rig}^T \\ \tilde{H}_{shocks}^T \end{bmatrix} + \begin{bmatrix} \varepsilon_0[1] \\ \vdots \\ \varepsilon_0[7] \\ \varepsilon_1 \\ \varepsilon_2 \\ \varepsilon_3 \\ \varepsilon_4 \end{bmatrix} \quad (5.40)$$

$$Y = \quad \quad \quad UH \quad \quad \quad + \quad E$$

where $y_0[k]$ is the output measured for baseline shock experiment k , and y_i is the output measured from alternate shock experiment i .

Let us restrict the following analysis to frequencies where all actuator inputs do not have a zero for the current shock velocity. Since U_{rig} is full rank, the first 7 columns of the

input matrix are rank 7. Since $\Delta f_i \neq 0$ the last four columns of the input matrix are rank 4. For the input matrix to be rank deficient, at least one of the last four columns must be linearly dependent on the first 7 columns. This implies that the equation

$$\begin{bmatrix} U_{rig} \\ u_{rig}^T \end{bmatrix} d = \begin{bmatrix} 0 \\ \Delta f_i \end{bmatrix} \quad (5.41)$$

must be satisfied for some i . Since U_{rig} is full rank, d must be zero. This requires $\Delta f_i = 0$, which contradicts Equation (5.38). Therefore, an experiment conducted using the above procedure will produce a full rank input matrix at all frequencies where the shock velocity is nonzero, allowing estimation of a unique frequency response estimation.

The least-squares problem defined by the data matrices in Equation (5.40) and the objective function in Equation (5.31) may then be used to solve the normal equations in Equation (5.32) to give the frequency response estimates for the rig inputs and the shock inputs. In practice, more than 11 data sets would be collected to allow for averaging and to provide an estimate of model quality.

In contrast to the above global estimation approach, it is often desirable to first fit a baseline model for the rig inputs, then augmenting the baseline model with component sensitivity models as needed. If the error in the baseline model is small, the results will be very similar. An advantage of this sequential approach is that it allows us to deal with smaller input matrices, making the solution more computationally efficient. Another advantage is that it allows us to calculate coherence for a single-input single-output FRF estimation for the component inputs, giving us a good indication of where the model performs well.

The sequential approach may be written as a sequential least-squares problem by partitioning Equation (5.40) into the rig input tests and the component input tests. The least-squares problem for the rig inputs with K data sets may be written

$$\min J = E_{rig}^* E_{rig}$$

$$\begin{bmatrix} y_0[1] \\ \vdots \\ y_0[K] \end{bmatrix} = \begin{bmatrix} u_{rig}^T[1] \\ \vdots \\ u_{rig}^T[K] \end{bmatrix} \tilde{H}_{rig}^T + \begin{bmatrix} \varepsilon_0[1] \\ \vdots \\ \varepsilon_0[7] \end{bmatrix} \quad (5.42)$$

$$y_{rig} = U_{rig} \tilde{H}_{rig}^T + E_{rig}$$

The solution to this least-squares problem satisfies

$$(U_{rig}^* U_{rig}) \tilde{H}_{rig}^T = U_{rig}^* Y_{rig} \quad (5.43)$$

Once the rig input frequency response has been determined, the shock input frequency responses may be estimated independently. The least-squares problem for the i^{th} shock input with K_i data sets, given the rig input frequency response, may be written

$$\min J_i |_{\tilde{H}_{rig}} = E_i^* E_i$$

$$\begin{aligned} \begin{bmatrix} \Delta y_i[1] \\ \vdots \\ \Delta y_i[K_i] \end{bmatrix} &= \begin{bmatrix} y_i[1] - u_i^T[1] \tilde{H}_{rig}^T \\ \vdots \\ y_i[K_i] - u_i^T[K_i] \tilde{H}_{rig}^T \end{bmatrix} \\ &= \begin{bmatrix} \Delta f_i[1] \\ \vdots \\ \Delta f_i[K_i] \end{bmatrix} \tilde{H}_{shock_i}^T + \begin{bmatrix} \varepsilon_i[1] \\ \vdots \\ \varepsilon_i[K_i] \end{bmatrix} \end{aligned} \quad (5.44)$$

$$\Delta Y_i = \Delta F_i \tilde{H}_{shock_i}^T + E_i$$

where \tilde{H}_{shock_i} is the frequency response for the i^{th} shock input. This least-squares problem satisfies

$$(\Delta F_i^* \Delta F_i) \tilde{H}_{shock_i}^T = \Delta F_i^* \Delta Y_i \quad (5.45)$$

This sequential procedure was applied to the data generated by simulation, providing frequency response estimates for both rig inputs and shock inputs.

If the baseline shocks are nonlinear, a system of equations similar to Equation (5.40) may be derived with more nonzero entries in the last 4 columns of the input matrix. Due to the additional nonzero entries in the input matrix, the equations cannot be decoupled as in Equations (5.22)-(5.25), so the entire system must be solved simultaneously.

5.2.3 Parametric System Identification

Now that the nonparametric frequency responses have been estimated, these estimates may be used to fit a parametric model. The parametric model will have the advantages of smoothing the nonparametric frequency response, improving the model in small frequency ranges where the nonparametric model was poor, and being more efficient for simulation.

There are many different methods of performing parametric system identification, which can be chosen depending on the particular application. The transfer function for each input/output pair can be defined by its poles, zeros, and constant gain. In an ideal linear system, all input/output pairs should share the same poles and only differ by the zeros and constant gain. The identification of the system poles is often important for applications in controls and structural parameter identification.

In our case, it is not necessary to identify a consistent set of system poles. All we want is to simplify our nonparametric frequency response estimates into a convenient transfer function relationship that will allow us to simulate the linear system more efficiently. This allows us to use a simpler method of fitting a single-input single-output (SISO) transfer function to the frequency response estimate for each input/output pair. This method will give a different set of poles, zeros, and a constant for each input/output pair. One advantage of this method is that it is very easy to fit a SISO transfer function, and many methods are readily available.

Another advantage is that it localizes error in the frequency response estimate to the corresponding transfer function estimate. If there is significant error in the frequency

response estimate for one of the input/output pairs due to factors such as noise, poor excitation, nonlinearity, or sensor malfunction, the error will only affect the corresponding transfer function estimate, not any other transfer function estimate.

The SISO method chosen for this work uses the *invfreqs* function in MATLAB [40]. This method fits a SISO transfer function of the form

$$H(s) = \frac{B(s)}{A(s)} = \frac{b_n s^n + b_{n-1} s^{n-1} + \dots + b_0}{s^m + a_{m-1} s^{m-1} + \dots + a_0} \quad (5.46)$$

by minimizing the equation-error cost function

$$J = \sum_k w_i(k) \left| h(k) - \frac{B(j2\pi f(k))}{A(j2\pi f(k))} \right|^2 \quad (5.47)$$

where $f(k)$ is the frequency at index k , $h(k)$ is the nonparametric frequency response estimate, and $w_i(k)$ is a weighting factor. This optimization problem is nonlinear in nature and must be solved using iterative nonlinear programming techniques. The *invfreqs* function uses a Gauss-Newton iterative search to solve the nonlinear programming problem.

5.3 System Identification Results

This section applies the identification methods described in Section 5.2 with the simulated data sets described in Section 5.1 to generate a model that approximates the vehicle behavior. The resulting model is then compared with the original model to determine where the ID process was successful and to help improve the ID method.

5.3.1 Nonparametric Identification Results

Five simulations were performed to generate data for use in system identification. An 800 second Gaussian white noise drivefile, bandlimited from 0 to 30 Hz drivefile was

constructed to excite the Simulink model. One baseline simulation was performed with an empirical shock model with a damping coefficient of 20 lb/(in/s) and four additional simulations were performed where the baseline shocks were replaced by a different empirical shock model with a damping coefficient of 8 lb/(in/s).

The first data set was used to estimate the rig input frequency response. Forty data blocks were created from the 800 second data set using a Hamming window with 50% overlap, giving frequency resolution of 0.03 Hz. The frequency response estimates for the rig inputs was very close to the ideal frequency response, with coherence very close to 1. The only exception was the ride height / roll moment frequency response, which had almost no response due to vehicle symmetry.

The next four data sets were used to estimate the shock input frequency response, using the same settings as the first data set. The coherence and frequency response estimate for the left front shock velocity output for all four shock force inputs is shown in Figure 5.9. The results for the left front shock velocity are consistent with the results for the other three shock velocities. For the left front shock force input, the frequency response estimate is accurate for all frequencies except near the anti-resonance at 6 Hz. For the other shock force inputs, the frequency response estimate is only good near the resonant peaks.

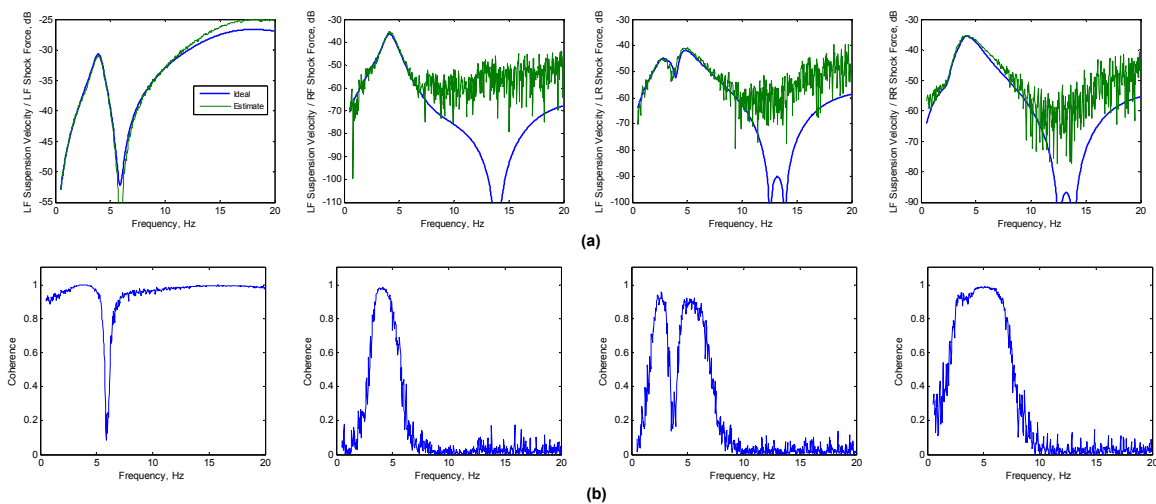


Figure 5.9. Shock velocity FRF estimate: (a) FRF estimate, (b) Coherence

In all cases, the estimate is poor only where there is very low output energy, resulting in poor coherence. This is a significant observation, because if the model is poor when the response level is low, the poor frequency response estimate may be replaced by a small frequency response, which will be close to the ideal model.

Next we show the coherence and frequency response estimate for the left front tire force, which is consistent with the other three tire forces. For the left front shock force, we see that the estimate is accurate except at lower frequencies, corresponding to the low response level. For the other shock force inputs, the estimate is accurate only near the resonance peaks.

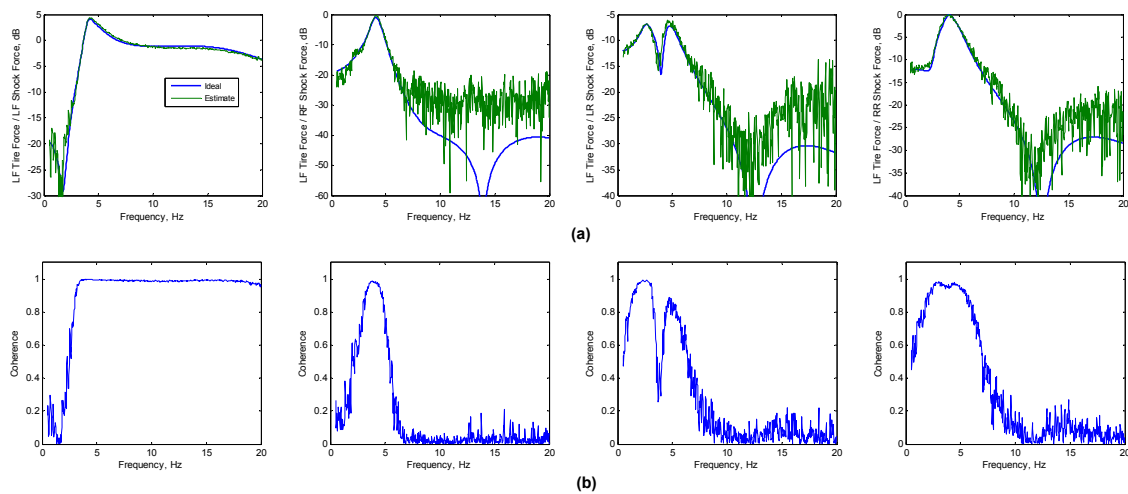


Figure 5.10. Tire force FRF estimate: (a) FRF estimate, (b) Coherence

The last output that we show is the front ride height, whose trends may be generalized to the rear ride height. For the front shock forces, the results are generally good except at higher frequencies, corresponding to a zero in the FRF and high frequency attenuation. For rear shock forces, the estimate is only good near the resonance.

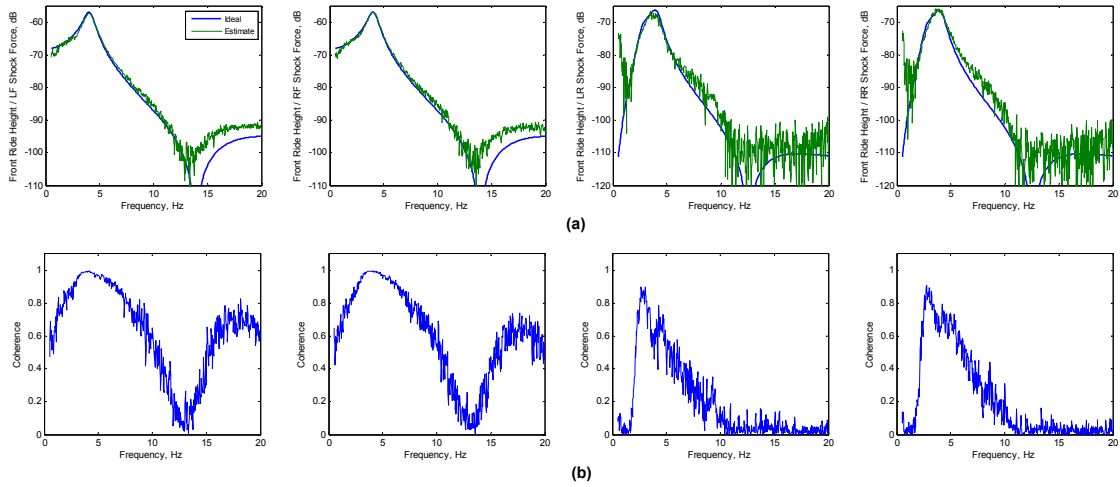


Figure 5.11. Ride height FRF estimate: (a) FRF estimate, (b) Coherence

These results have shown us that our frequency response estimates for shock force inputs are good when the response levels are high, which occurs more often for direct coupling input/output pairs that occur along a block diagonal of the frequency response matrix and for off-diagonal entries near resonance. Estimates are reasonable when the response levels are low. Also the coherence function provides a useful figure of merit for the frequency response estimates, indicating where the model will be poor.

5.3.2 FRF Preconditioning

The FRF estimates described in the previous section may be processed prior to parametric identification to provide better results. First, poor estimates observed for the shock force input at higher frequencies may be replaced by a residual fit of the form

$$H_{high}(f) = \frac{Ke^{j\phi}}{f^N} \quad (5.48)$$

This fit provides a 20N dB/decade rolloff and a constant phase. The form of this residual fit is based on the assumption that the transfer function will be dominated by the highest-order nonzero terms in the numerator and denominator for these higher frequencies.

A similar fit was performed at lower frequencies, which adds a linear phase, of the form

$$H_{low}(f) = \frac{Ke^{j(\phi+bf)}}{f^N} \quad (5.49)$$

Each residual may be fitted by selecting two control points off each FRF where the coherence is good. The frequency, gain, and phase at these two control points then define the residual model constants. If desired, the control points may be manually adjusted to provide the desired residual fit. The high frequency control points were selected to be at 20 and 30 Hz for on-diagonal input/output pairs and were at 6 and 7 Hz for off-diagonal pairs. The low frequency residual control points were placed approximately at 1 and 4 Hz.

To further smooth the FRF estimate prior to parametric identification, the FRFs are lowpass filtered. Filtering a complex vector may be performed analogously to filtering a real vector, with real filters being applied to the magnitude in dB and the unwrapped phase. A 4th order Butterworth filter was constructed to attenuate FRF variations with a width of 0.6 Hz or less, corresponding to 18 frequency samples. This filter was then applied to the magnitude in dB and the unwrapped phase using the MATLAB function *filtfilt* to avoid frequency shift.

A comparison of the ideal FRF, the original FRF estimate, and the modified estimate is shown in Figure 5.12. As this figure shows, the modified FRF estimate is much smoother, making it easier to fit a low-order transfer function.

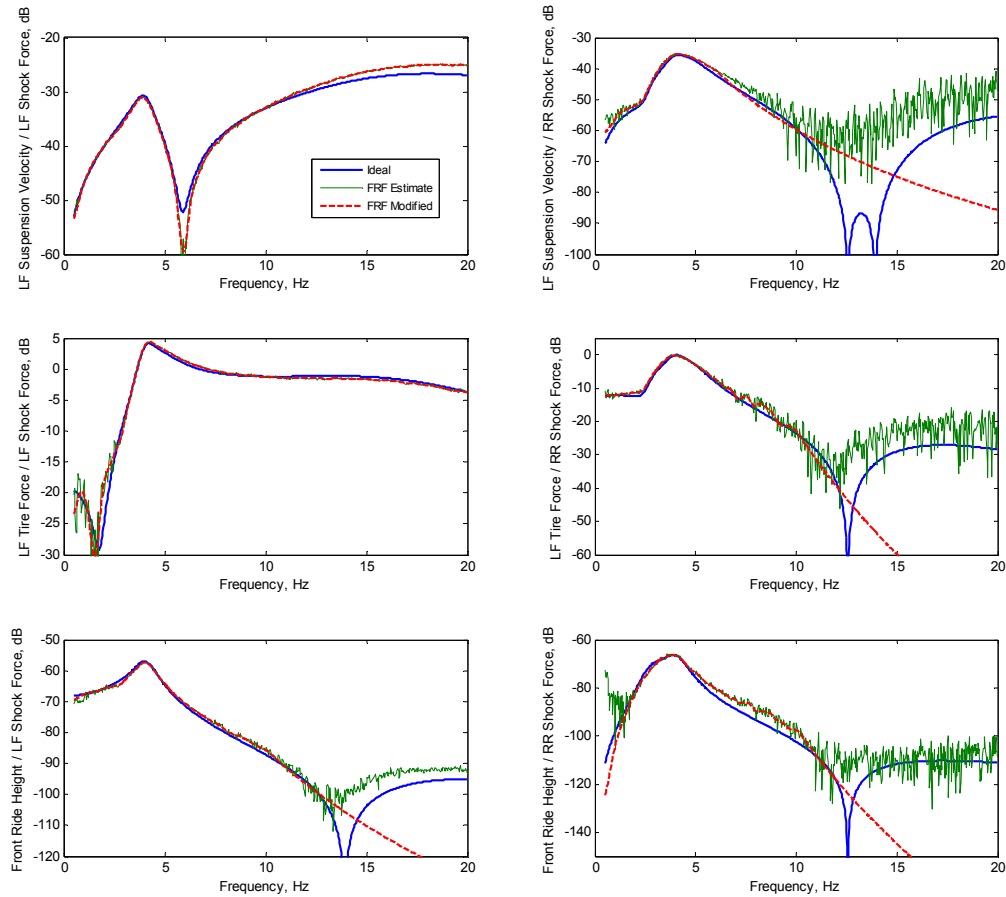


Figure 5.12. Comparison of ideal FRF, FRF estimate, and modified FRF estimate

5.3.3 Parametric Identification Results

Now that the nonparametric frequency response estimates have been calculated, the next step is to use these results to fit a parametric model to each input/output pair using the *invfreqs* function. The denominator polynomial orders were increased until good matching was achieved, resulting in denominator orders ranging from 6 to 8. The tire force and ride height outputs have direct feedthrough from the nearby wheelloaders, so their numerator order is set equal to the denominator order, while all other numerator orders are set to one order lower than the denominator. A weighting function of 1 was used between 0.5 and 10 Hz, while a weighting function of 0.5 was used from 10 to 20 Hz.

The results of the parametric fit created by *invfreqs* are compared with the ideal response and the modified FRF estimate in Figure 5.13. As desired, all transfer function fits closely approximate the ideal FRF where response levels are high, while the transfer function levels are low when the response level is low.

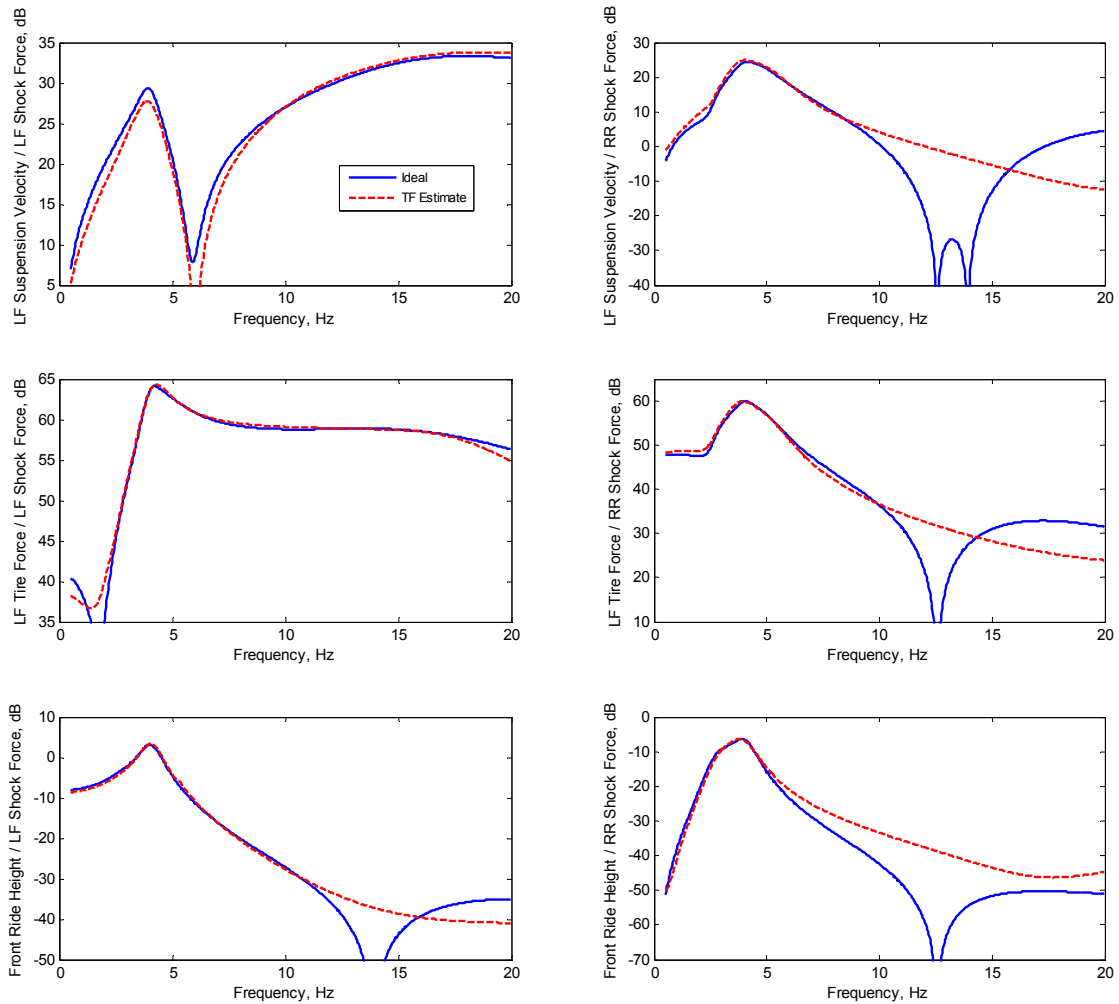


Figure 5.13. Result of parametric identification

5.3.4 Simulation with Identified Model

The transfer function model identified in Section 5.3.3 can now be applied to perform full-vehicle simulations using the Simulink model shown in Figure 5.14.

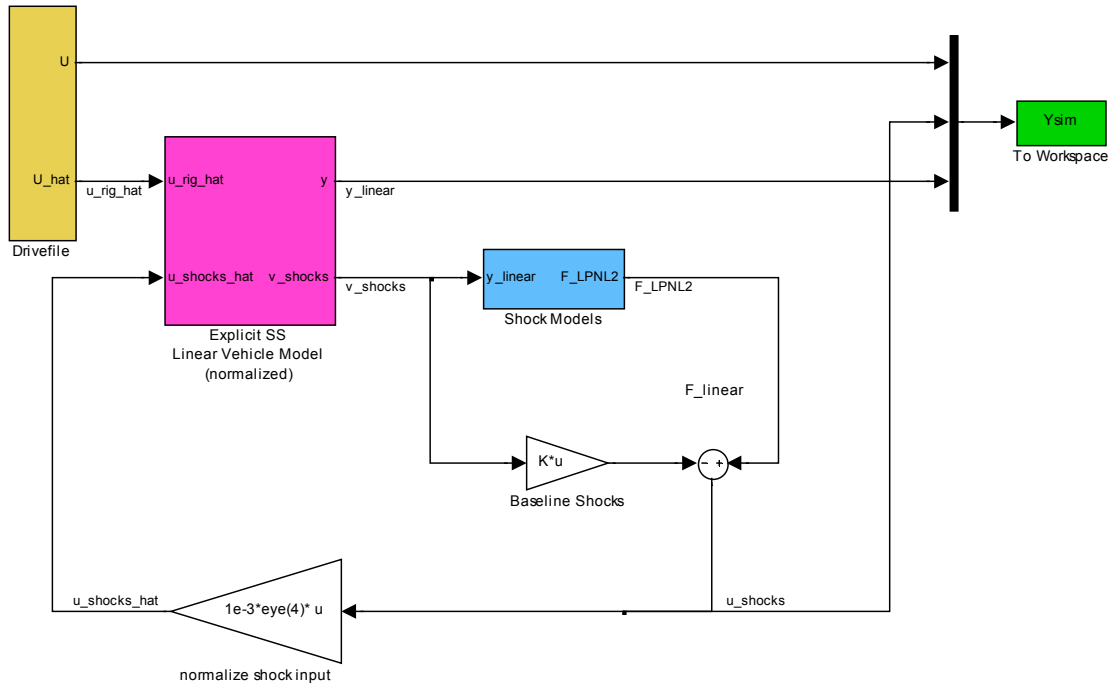


Figure 5.14. Simulink model

The continuous-time transfer function model is converted to a discrete-time state space model sampled at 100 Hz using Tustin's bilinear transformation in the MATLAB function *c2d*. The shock velocity calculated as an output of the linear model is then used to calculate the shock force predicted by the LPNL2 shock models and the shock force that would be generated by the baseline shock at that velocity. The difference in the shock force from the baseline shock force is then normalized and fed back as an input to the linear vehicle model.

Five simulations were performed with the identified model to compare with simulations using the ideal model. The 800 second Gaussian white noise bandlimited to 30 Hz drivefile described in section 5.1.3 was used to excite the Simulink model. One baseline simulation was performed with 4 nonlinear dynamic baseline shock models with a damping coefficient of 20 lb/(in/s) and four additional simulations were performed where the baseline shock at each corner was replaced by a nonlinear dynamic shock model with a damping coefficient of 8 lb/(in/s).

The results from each simulation were bandpass filtered from 1 to 20 Hz using a 4th order Butterworth filter. A comparison of the tire force and front ride height using the ideal model and the identified model with the baseline shocks, the left front shock changed, and the right front shock changed is shown in Figure 5.15. These plots show that the simulation using the ID model match the original data well. Errors in ride height are small, while there is good matching of the tire force at the dominant frequency band of 3-5 Hz. Peak errors in ride height occur when the signal is changing rapidly, suggesting phase error. The largest error occurs with the left front tire force when the left front shock is changed, where the error magnitude is roughly 25% of the signal level.




Figure 5.15. Simulation time trace comparison

The error can be analyzed in further detail in the frequency domain. Figure 5.16 shows the power spectral density (PSD) for the time signals in Figure 5.15. For the ride height, the error is 50 dB lower than the signal level. For the tire force, the error is more than 20 dB lower than the signal level, except for the change in left front shock from 7 to 17 Hz. In this frequency band, the peak error is 40 dB lower than the peak signal level.

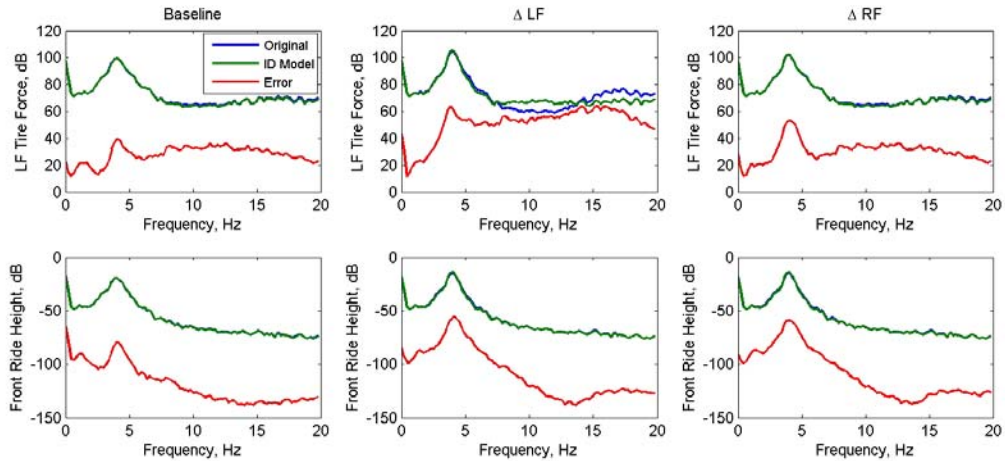


Figure 5.16. Simulation PSD comparison

The main purpose of this model is to accurately predict trends in response measures as shock setups change. Since the PSDs for the ideal simulation and the ID simulation match well, we would expect RMS response metrics to match well. This is the case, as Figure 5.17 illustrates.




Figure 5.17. Metric comparison

5.4 Summary

This chapter has developed a method to identify a linear full-vehicle model from sensor measurements collected during vehicle testing that allows investigation of the effect of shock selection in simulation. To develop the method, a seven degree of freedom vehicle state space model was constructed to generate data sets for the identification process. The identification process can then be judged acceptable if the identified model is similar to the known ideal model.

The frequency response function (FRF) of the state space model was then explored to provide insight prior to the identification process. This analysis showed the importance of both the aeroloaders and the wheelloaders to excite shock velocity at low and high frequencies, respectively. It also showed frequency ranges where we might expect to have poor coherence for certain input/output pairs due to low signal levels.

If the shock force is almost linearly dependent on the wheelloader and aerolader inputs, the typical FRF estimation process will fail. A modified FRF estimation process was described which uses test data from several different shock configurations to address the shock force dependence issue. This method provided FRF estimates, which matched the ideal FRF where the coherence was high. When the coherence was poor due to low shock force or output levels at low or high frequency, the FRF estimate was replaced by a residual fit with a linear phase and linear dB/decade rolloff. The FRF estimate was further smoothed by applying a filter, reducing FRF variations with a width of 0.6 Hz or less.

Once an acceptable FRF estimate was calculated, a parametric model was identified by fitting transfer functions to each input-output pair. This parametric model was then implemented in Simulink using the shock models described in Chapter 3. Simulations were run for 5 different shock configurations and compared to simulation results from the ideal model, showing that the identified model could accurately reproduce the original response and predict trends in response metrics.

Now that this identification process has been shown to accurately reproduce the original system behavior, it can now be applied to data collected from actual 8-post rig testing.

Chapter 6

8-Post Rig Testing and Identification

This chapter discusses the process of conducting an 8-post rig test, using the results to identify a model to predict the results of future tests, and to validate the model accuracy. First, the experimental setup for the 8-post rig tests is described. Next, the results for the preliminary testing will be used to evaluate the validity of linear modeling for each input-output pair, which will provide guidance during the ID process. The system identification methods and results are then presented, followed by model validation.

6.1 Experimental Setup

The 8-post rig test in this study is comprised of four major building blocks: the 8-post rig itself, the instrumented test vehicle installed on the rig, the drivefiles used to command the rig, and the shock configurations installed on the vehicle. Once these building blocks are defined, test procedures can be easily defined as needed. The following sections detail each of these four building blocks independently then describe the actual tests that were conducted.

6.1.1 8-Post Rig

All rig testing was conducted on the 8-post rig at the Virginia Institute for Performance Engineering and Research (VIPER). During an 8-post rig test, the tires are supported by 4 hydraulic actuators, or wheelloaders, that simulate inputs from the track surface. Four pneumatic actuators, or aeroloaders, attach to the vehicle's chassis to provide heave force, roll moment, pitch moment, and warp moment, which simulate the effect of inertial and aerodynamic forces present during a track test. The vehicle is shaken on the rig to

simulate how the car would respond at a particular track or to characterize the vehicle response to more general waveforms such as sine waves or broadband signals.

The test vehicle attached to the rig is shown in Figure 6.1a. From the 8-post lab, the view of the components of the rig is obscured by a reconfigurable floor structure, known as a pit cover, which allows vehicles to be driven directly onto the rig. The rig can be clearly seen with the pit cover removed, as shown in Figure 6.1b. The rig rests on a seismic mass, which consists of 400 tons of concrete and steel isolated from the building foundation, ensuring that most of the actuator energy is used to accelerate the test vehicle and not the ground underneath. A steel base plate is mounted to the seismic mass, which is used to mount the wheelloaders and aeroloaders, as shown in Figure 6.1c. The wheelloader is clamped to the baseplate, while the aeroloader is attached to a swivel joint that mounts to a pedestal, which is then clamped to the baseplate.




Figure 6.1. 8-post rig: (a) NASCAR Cup car on rig, (b) Rig with pit cover removed, (c) Wheelloader and aeroloaders mounted under pit cover (photos by author, 2009)

The rig is installed below the ground in a recess in the seismic mass, known as the pit. This allows the wheel platens to be at ground level, so a vehicle can drive across the pit cover and onto the wheel platens.

When a vehicle is installed on the rig, each tire is supported by a wheel platen, as shown in Figure 6.2a. The wheel platen is mounted to the wheelloader, which provides the interface between the wheelloader and the tire. On the top of the wheel platen is a white Teflon plate, which reduces friction between the tire and the wheel platen. A low friction tire-wheelloader interface reduces the potential for variations in the vehicle's equilibrium

point due to friction and also reduces jacking forces. There are also lateral and longitudinal restraint bars around the perimeter of the wheel platen, which prevent the vehicle from falling off the wheel platen. These restraint bars do not touch the tire during normal testing. The force applied by the wheelloader is measured by four triaxial load cells, as shown in Figure 6.2b. The vertical force measured by the four load cells is connected to a summing box to allow direct measurement of the vertical wheel platen force.




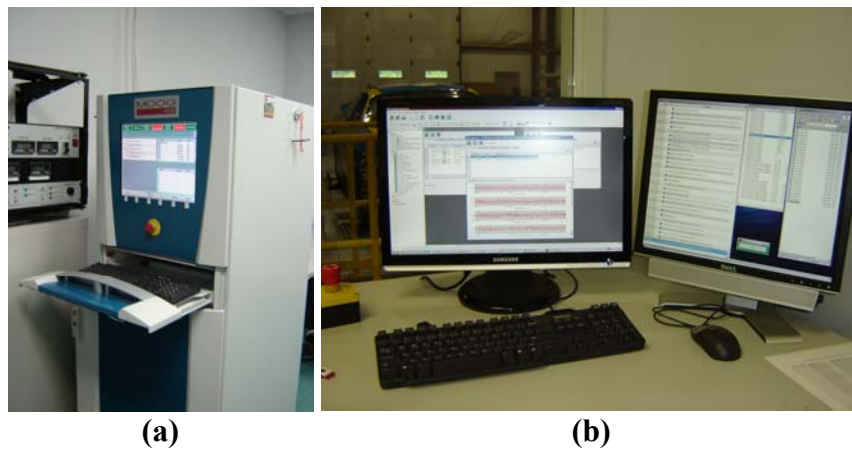
Figure 6.2. Rig attachments to vehicle: (a) Wheel platen and wheel restraint, (b) Wheel platen load cells, (c) Aeroloader mounting and load cell (photos by author, 2009)

Four pneumatic actuators, or aeroloaders, are also attached to the vehicle's frame, as shown in Figure 6.2c. The aeroloader is attached at the bottom using a swivel joint and at the top using a rod end. These joints ideally produce free-free boundary conditions across the aeroloader, which would make the aeroloader an ideal two-force member that provides only axial force along the line between the two pivots. A blue load cell is mounted on the aeroloader shaft to measure the applied force, as shown in Figure 6.2c.

Each wheelloader consists of a hydraulic actuator which has a stroke of 12.4 inches, a control bandwidth of 125 Hz, can produce 7,400 lb of dynamic force, and travels at speeds up to 180 in/s. Each actuator includes a Macro Sensors linear variable differential transformer (LVDT) for position measurement, a PCB accelerometer, and four Michigan Scientific triaxial loadcells to measure the three forces and three moments transferred to the wheel platen. The four wheelloaders are powered by a 200 HP, 100 GPM, 3000 psi hydraulic pump with over 30 gallons of accumulation to prevent transient pressure drops.

Each aeroloader consists of a pneumatic actuator which has a stroke of 24 inches, a control bandwidth of 9 Hz, can produce 4300 lb of force, and control at velocities up to 12 in/s. Each actuator includes a Futek load cell and an Interface LVDT. The four aeroloaders are powered by a 50 HP air compressor, which provides refrigerator-dried air at a working pressure of 120 psi.

The 8-post rig is controlled by the Moog-FCS SmarTest controller shown in Figure 6.3a. The controller receives the analog sensor signals from the rig and converts them to a digital signal, which is then available to the controller's real-time processor. The real-time processor then checks to see if any safety limits have been exceeded requiring rig shutdown. If no safety limits have been tripped, the processor applies a control policy to calculate the actuator command signals, which are then converted to an analog signal through an A/D and routed to the actuator control valves. The controller also includes 24 additional channels of analog input and data logging capability.



**Figure 6.3. Test control: (a) SmarTest controller, (b) FasTest PC interface
(photos by author, 2009)**

Test management is performed using Moog-FCS's FasTest software on a PC networked to the controller as shown in Figure 6.3b. The software allows the user to define a sequence of instructions that define the actuator commands or drivefile, data acquisition, and safety limits. After a sequence of instructions has been constructed, the sequence can be downloaded to the real-time processor for execution.

For each test, we used a sequence of instructions that played a drivefile and recorded the vehicle and rig sensor measurements. The vehicle instrumentation and drivefiles that were used are described in the next two sections.

6.1.2 Test Vehicle Setup

The test vehicle used for the 8-post rig tests, as shown in Figure 6.1a, is a 2007 NASCAR Cup car donated to VIPER by Petty Enterprises. This car's suspension consists of an independent short-long arm (SLA) front suspension and a solid-axle trailing-arm rear suspension, as shown in Figure 6.4. At all four corners, there is a coil spring and a Penske 7300 shock absorber installed. In the front, there is also an anti-roll bar, which provides some coupling across the front axle.




Figure 6.4. Test vehicle suspension: (a) Independent SLA front suspension, (b) Solid-axle trailing-arm rear suspension (photos by author, 2009)

In addition to the sensors already present on the 8-post rig actuators, instrumentation was installed throughout the vehicle to characterize the suspension performance. Each shock has a Penny and Giles shock potentiometer mounted across it with a motion ratio close to 1 with respect to the shock travel. The right front shock potentiometer is shown in Figures 6.5a. During the identification tests, shock force will be measured using the

Penske strain gauge eyelet and Beru F1 inline amplifier shown in Figure 6.5b. The strain gauge eyelet threads onto the shock shaft, replacing the stock eyelet.

PCB accelerometers were also used near the suspension and in the cockpit, as shown in Figure 6.6. Accelerometers were used on each corner to measure the wheel acceleration and the frame acceleration above the suspension. A triaxial accelerometer was used to measure acceleration near the vehicle center of gravity and a single axis accelerometer was used to measure acceleration near the driver's torso, which may be unfavorable to driver comfort. All accelerometers were conditioned using a 16 channel PCB signal conditioner with programmable gain and 8th order elliptical filter, which was programmed to give the accelerometers a sensitivity of approximately 1 V/g and a cutoff frequency of 250 Hz.




Figure 6.5. Shock instrumentation: (a) Shock potentiometer, (b) Shock load cell (photos by author, 2009)




Figure 6.6. Accelerometers: (a) Wheel and chassis accelerometers, (b) Cockpit accelerometers (photos by author, 2009)

In addition to the measured channels described above, additional math channels are calculated during post-processing that will be used for analysis. A summary of all rig channels, vehicle sensor channels, and math channels is shown in Table 6.1. Calculation of math channels is described below.

Table 6.1. Summary of measured and calculated channels

24 Rig Signals		
1-4	Wheelloader	Position
5-8	Aeroloader	Position
9-12	Wheelloader	Force
13-16	Aeroloader	Force
17-20	Wheelloader	Command
21-24	Aeroloader	Command
24 Vehicle Sensors		
25-28	Shock Potentiometer	
29-32	Shock Force	
33-36	Wheel Accelerometer	
37-40	Chassis Accelerometer	
41-43	CG Triaxial Accelerometer	
44	Driver Accelerometer	
18 Math Channels		
45-48	Aeroloader Heave/Pitch/Roll/Warp	Position
49-52	Aeroloader Heave/Pitch/Roll/Warp	Force
53-56	Aeroloader Heave/Pitch/Roll/Warp	Command
57-58	Front and Rear Ride	Heights
59-62	Shock	Velocity

It is often desirable to describe the aeroloader position and force in terms of heave, pitch, roll, and warp modes instead of on a single actuator basis. Assuming left-right aeroloader symmetry, the transformation between aeroloader displacement and mode displacement is

$$X_m = T_x X \quad (6.1)$$

$$\begin{bmatrix} \text{Heave Displacement} \\ \text{Pitch Angle} \\ \text{Roll Angle} \\ \text{Warp Angle} \end{bmatrix} = \begin{bmatrix} \frac{b'}{2L'} & \frac{b'}{2L'} & \frac{a'}{2L'} & \frac{a'}{2L'} \\ \frac{-1}{2L'} & \frac{-1}{2L'} & \frac{1}{2L'} & \frac{1}{2L'} \\ \frac{b'}{L't_f'} & \frac{-b'}{L't_f'} & \frac{a'}{L't_r'} & \frac{-a'}{L't_r'} \\ \frac{1}{t_f} & \frac{-1}{t_f} & \frac{-1}{t_r} & \frac{1}{t_r} \end{bmatrix} \begin{bmatrix} z_{LF}' \\ z_{RF}' \\ z_{LR}' \\ z_{RR}' \end{bmatrix}$$

where X_m is the vector of mode displacements, X is the vector of actuator displacements, and T_x is the transformation matrix. The constants a' , b' , L' , t_f' , t_r' define the aeroloader geometry about a reference point along the vehicle centerline. Roll angle is defined to be the roll angle at the reference point, assuming roll angle varies linearly along the length of the chassis. Warp angle is defined to be the difference between front and rear roll angles. Since the aeroloaders attach to the chassis, these math channels also describe the chassis heave, pitch, roll, and warp. Similarly, the transformation between aeroloader forces and aeroloader mode forces can be written

$$F_m = T_f F \quad (6.2)$$

$$\begin{bmatrix} \text{Heave Force} \\ \text{Pitch Moment} \\ \text{Roll Moment} \\ \text{Warp Moment} \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ -a' & -a' & b' & b' \\ \frac{t_f'}{2} & \frac{-t_f'}{2} & \frac{t_r'}{2} & \frac{-t_r'}{2} \\ \frac{t_f'}{2} & \frac{-t_f'}{2} & \frac{-t_r'}{2} & \frac{t_r'}{2} \end{bmatrix} \begin{bmatrix} F_{LF}' \\ F_{RF}' \\ F_{LR}' \\ F_{RR}' \end{bmatrix}$$

where warp moment is defined to be the difference between the front and rear roll moments. Since the aeroloaders will be run in force control for this study, the transformation between aeroloader commands and mode commands can be written using the transformation matrix T_f .

Another important measurement for aerodynamic performance is the chassis ride height at the front valence and rear spoiler height. Deviation from an optimal chassis position may reduce total aerodynamic downforce, increase drag, or disturb the aerodynamic downforce distribution. The front and rear ride heights can be easily calculated using the chassis heave and pitch.

Since our shock model requires shock velocity as an input, one key element of our vehicle model is to predict the shock velocity. We estimate the shock velocity by numerically differentiating the lowpass filtered shock potentiometer measurement using MATLAB's *gradient* function. This function uses a central difference approximation to calculate the derivative.

6.1.3 Drivefiles

Several different drivefiles were used during 8-post rig testing, for various purposes including better understanding the vehicle and acceptable excitation levels, identifying the vehicle model, model validation, and suspension evaluation.

The first drivefiles tested were a series of sine tests. Sine tests were performed to determine acceptable excitation levels at different frequencies and to provide model validation data. The wheelloaders were commanded sine wave positions in heave, pitch, and roll with frequencies chosen to capture the dynamics before, during, and after resonance of that particular mode.

Table 6.2. Sine test frequencies and amplitudes

Heave		Pitch and Roll	
Frequency Hz	Amplitude in	Frequency Hz	Amplitude in
1	1	1	1
2	1	2	0.7
4	0.7	3	0.5
6	0.5	5	0.5
8	0.3	6	0.3
10	0.3		
15	0.1		
20	0.05		

After sine testing was completed, a random drivefile was designed for use in system identification. Each actuator command within the drivefile was constructed by first generating a Gaussian time sequence, filtering it to achieve a desired spectrum shape, and then adjusting the scale and mean value to desired values. To perform both multiple-input and single-input frequency response estimation, an 800 second drivefile with all 8 actuators active and 8 100 second drivefiles with only one actuator active in each drivefile were created. Inactive aeroloaders were commanded with a DC downforce that provided approximately 3 inches of travel at the front aeroloaders and 2.5 inches of travel at the rear aeroloaders, while inactive wheelloaders were commanded zero displacement.

The relative amplitude spectra for the aeroloaders and the wheelloaders are shown in Table 6.3. The aeroloader spectrum is 100% up to 3 Hz, 10% at 4 Hz, 1% at 5 Hz, and zero above 5 Hz due to the aeroloader bandwidth of 5 Hz. The wheelloader spectrum is similar to the heave sine limits found during heave sine testing, and is similar to the random spectrum used for shock testing and quarter-car testing. These spectra were used to fit a digital filter with similar frequency response using the Yule-Walker method. The filter was then applied to the Gaussian time sequence to provide approximately the desired spectra.

Table 6.3. Random signal relative amplitude spectrum

Aeroloader		Wheelloader	
Frequency Hz	Amplitude in	Frequency Hz	Amplitude in
0	1	0	1
0.5	1	0.5	1
3	1	1	1
4	0.1	2	0.8
5	0.01	3	0.7
		5	0.5
		8	0.3
		10	0.3
		15	0.3
		20	0.1
		30	0.01

The signals were then adjusted to have the desired mean value and scale. The wheelloaders were adjusted to zero mean and a peak value of 0.4 inches, providing an

RMS of approximately 0.11 inches. The mean aeroloader force commands were adjusted to have approximately 3 inches of deflection in the front and 2.5 inches of deflection in the rear. The aeroloader scale was set to a peak variation from the mean force of 560 lb, which provided aeroloader force commands of approximately 160 lb RMS, not including the mean downforce.

The next drivefile used was a 50 second, 0.4 inch peak-to-peak square wave with a period of 10 seconds to all four wheelloaders, while the aeroloaders were commanded a DC downforce. This drivefile will be used to evaluate the transient behavior of the vehicle, determine the rate of shock absorber recovery, and to validate the model for transient events.

The next drivefile used consists of a series of static aeroloader force commands, which will be used for spring rating the vehicle and providing a model at DC. The aeroloader forces are varied from zero to a maximum force determined to be near maximum suspension travel. The four aeroloader force commands are varied from zero force to max force and held constant at 5 different forces, as shown in Figure 6.7. This gives force-deflection data during motion dominated by chassis heave. The process is repeated moving from max force to zero force to determine hysteresis. Additional tests were run in a similar manner, as shown in Figure 6.6 to excite chassis pitch, roll, and warp. Single actuator tests were also run to test single-corner loading.

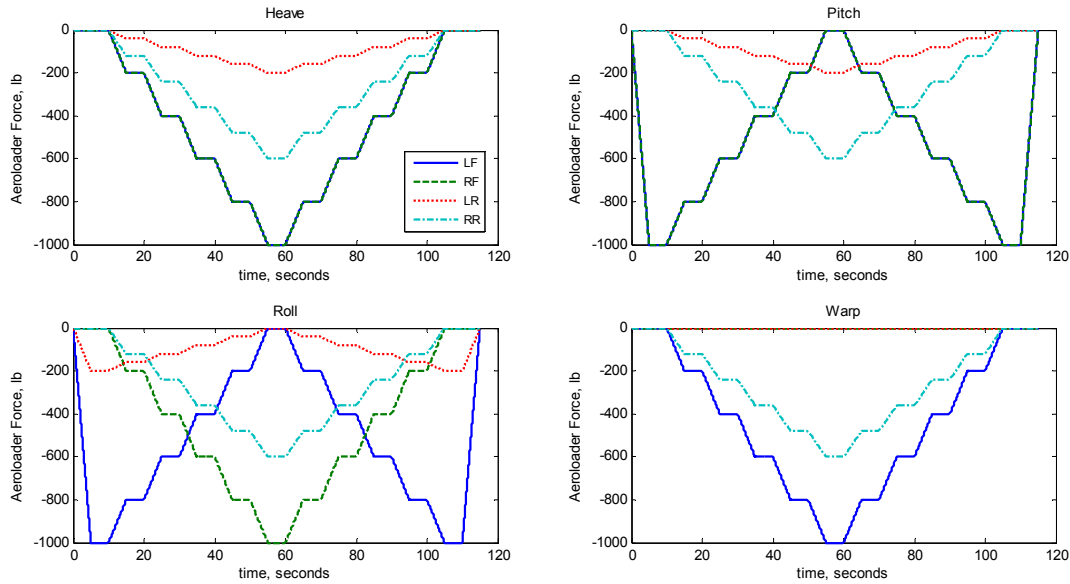


Figure 6.7. Chassis spring rating tests

The last drivefile used is a drivefile that was created using drivefile iteration to represent a lap in a NASCAR Cup car at Richmond International Raceway. This type of track-based drivefile is what a race team would typically use to evaluate their vehicle's performance for a particular track. In this study, we will use it both to evaluate the quality of different shock selections and to validate our model's ability to predict quality of different shock selections.

6.1.4 Shock Builds

This section describes the shock builds that will be used for system identification and model validation. First, a baseline shock build with a damping coefficient of approximately 20 lb/(in/s) was chosen. Four of these baseline shocks were built and installed on the car as a baseline setup. During the identification process, shocks must be changed to separate the response due to shock force from the response due to actuator inputs. To ensure that the change in shock force and change in sensor response was significant, a shock build with a damping coefficient of approximately 9 lb/(in/s) was chosen.

The four baseline builds and the alternate build were tested and modeled using the procedures detailed in Chapter 3. The force-velocity plots for these shocks are shown in Figure 6.8.




Figure 6.8. Shock builds: (a) Baseline, (b) Alternate

6.1.5 Experimental Procedure

The testing that was performed for this chapter will be used to identify a model of a vehicle on the rig and to validate that model. Now that the building blocks for conducting an 8-post rig test have been described, the experimental procedure may be concisely defined. During each test, data was collected at 1250 Hz using the sensors described in Section 6.1.2. Unless otherwise stated, the baseline shock builds A-D were used for testing.

Identification tests were performed using the random excitation described in Section 6.1.3. First, a baseline test was run with the 800 second 8-input random drivefile and the baseline shocks installed. The data from the baseline test will allow us to model the relationship between the actuators and the sensors. Next, an alternate test is run with the same drivefile and the left front shock replaced with the alternate shock build. The alternate shock build provides about half the damping that the baseline shock provides, which provides a significant change in the shock force. The data collected from the alternate test along with data from the baseline test will allow us to model the relationship

between shock force and sensor response. Alternate tests are repeated for each corner, where the baseline shock for that corner is replaced by the alternate shock, while all other shocks are the baseline shocks. Once all four alternate tests are completed, we can determine the influence of all four shock forces on sensor response.

The baseline identification tests are repeated using a series of single-input tests to allow us to analyze the behavior of each input/output pair in further detail. Eight tests are performed using the baseline setup tests where one actuator is excited at a time using a 100 second random drivefile. This will provide a clear picture of how actuator input energy affects the sensor response, which will aid in model development.

To validate the model's ability to predict the response to a drivefile representing a particular test track, the vehicle was excited in the baseline shock configuration using the Richmond International Raceway track drivefile. To validate the model's ability to predict the change in response due to shock selection, the drivefile was repeated while each shock was replaced by the alternate shock.

Further validation tests were performed on the baseline setup including sine tests, bump tests, and static tests. These will allow us to evaluate the model's ability to predict the vehicle response for different drivefiles to steady-state inputs at different frequencies, to transient inputs, and to static inputs.

6.2 Linear Model Validity

This section evaluates the validity of a linear model to approximate each input-output pair. For the wheelloader and aerolader inputs, we consider the results for the one input at a time tests for the linear baseline setup. For the shock inputs, we consider the change in response from the baseline linear response for the multiple-input random test when one shock at a time is replaced by an alternate shock. These experiments give us valuable information about how each input contributes to each output and if the input-output pair can be accurately modeled using a linear model.

For this analysis, we will consider the combined importance of response energy and coherence in linear system identification, as summarized in Table 6.4. Ideally, the coherence will be close to 1 for each input-output pair across the frequency range of interest. This means the input-output relationship observed in the data is accurately described by a linear relationship and it is reasonable to assume that a linear model is appropriate.

Table 6.4. Importance of coherence and output energy for linear modeling

		Coherence	
		High	Low
Response Energy	High	Condition I Good linear estimate Model predicts response	Condition III Poor linear estimate Cannot predict response
	Low	Condition II Good linear estimate Model predicts low response	Condition IV Poor linear estimate We can choose low response

Unfortunately, poor coherence may be caused by various factors including poor signal-to-noise ratio or system nonlinearity. If the coherence is low, for example below 0.5, a linear model only describes 50% or less of the response energy observed in the data. While this is not ideal, the coherence, or relative amount of the response energy described by a linear model, does not fully describe the potential accuracy of a linear system model, and can be a bit misleading. If coherence is poor and the response energy is high, then there will be significant amount of absolute error in the output energy predicted by a linear model. If the response energy is small or negligible, however, a poor coherence does not indicate a large absolute error in the output energy predicted by a linear model. When the output energy is small, we only need to ensure that the predicted response levels are small or zero, and there will be small absolute error.

For brevity, this section only displays the results for the shock velocity, tire force at the right front and left rear locations. Also shown are the results for front and rear ride height. Similar trends have been found for the other sensor outputs and locations.

Figure 6.9 shows the results for the right front and left rear shock velocity resulting from wheelloader inputs. Figure 6.9a shows two seconds of the time response as one wheelloader is applied at a time. For the right front shock velocity, the only time the response was significant was when the nearby right front wheelloader was excited. For the left rear shock velocity, the most significant response was due to the nearby left rear wheelloader, with a smaller but significant response due to the right rear wheelloader. This indicates that the shock velocity in the independent front suspension is only significantly excited by the closest wheelloader, while for the solid axle rear suspension the shock velocity is heavily influenced by the nearby wheelloader with additional significant contribution from the opposite side of the solid axle.

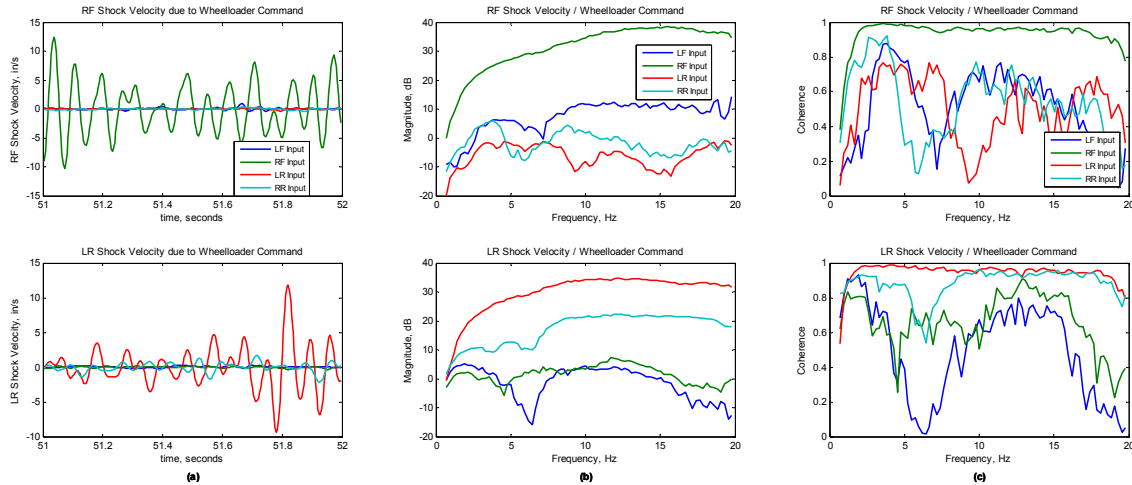


Figure 6.9. Shock velocity due to wheelloader input: (a) Time response, (b) FRF magnitude, (c) Coherence

This result is also evident when looking at the FRF magnitude in Figure 6.9b. For the right front shock velocity, the FRF for the closest wheelloader is 20-30 dB higher than the other wheelloader inputs over most of the frequency band. For the left rear shock velocity, the FRF is 20-30 dB lower for the front inputs and about 10 dB lower for the right rear input compared to the nearby left rear wheelloader input. These observations clearly indicate that the shock velocity response level is only significant in the front due to the closest wheelloader, while in the back the response is only significant for the rear wheelloaders. Also note that the magnitude response drops toward zero at low

frequency, since no shock velocity is excited at low frequency as the wheels move with the chassis in a rigid body mode.

The coherence for these outputs is shown in Figure 6.9c. As expected, the coherence for the nearby wheelloader is high, except at low frequency where there is little shock velocity. The coherence is also high for the left rear shock velocity for the right rear wheelloader input except at a dip in the magnitude near 7 Hz. The coherence for the remaining input-output pairs ranges from good to poor.

A key observation from this plot is that at frequencies where an input-output pair produces a significant response level, the coherence is good. For input-output pairs where the response levels are low, the transfer function may be approximated as zero with minimal error.

Figure 6.10 shows the shock velocity due to aerolader inputs. The coherence is generally high for all inputs, with the possible exception of the right front wheelloader input for the left rear shock velocity. For this input-output pair, the coherence is low when the response level is low.

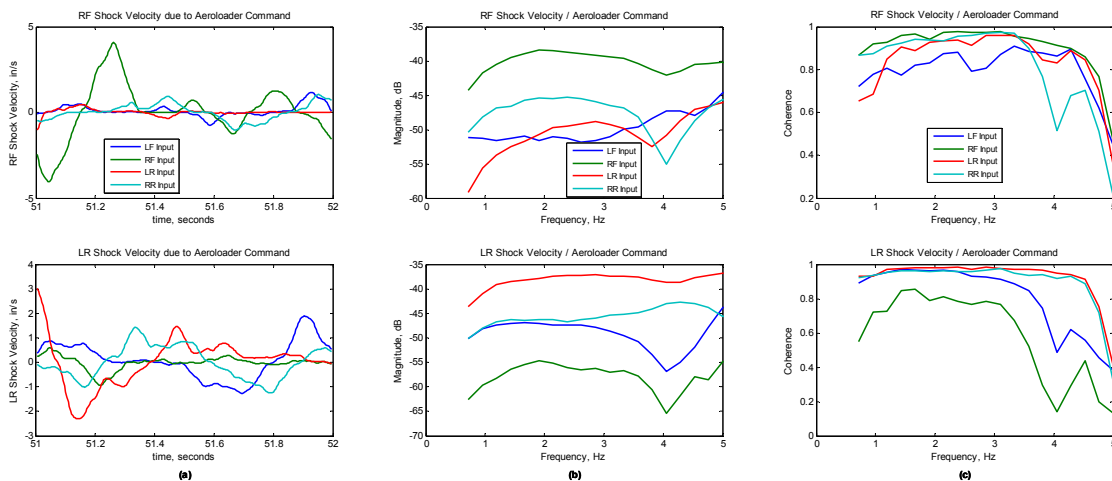


Figure 6.10. Shock velocity due to aerolader input: (a) Time response, (b) FRF magnitude, (c) Coherence

Next, we consider the effect of wheelloader inputs on tire force in Figure 6.11. For each wheelloader force, the most significant response is due to the nearby wheelloader, while

the rear wheelloader force has an additional significant contribution due to the wheelloader across the rear axle. The coherence is good for the nearby wheelloader input for the front shock velocity, and for both rear inputs for the rear shock velocity.

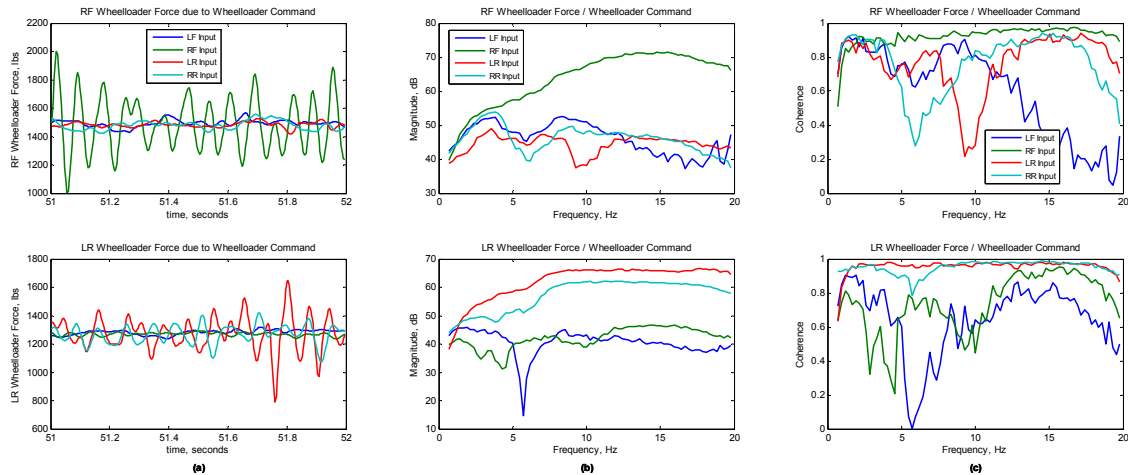


Figure 6.11. Tire force due to wheelloader input: (a) Time response, (b) FRF magnitude, (c) Coherence

The results for tire force due to aeroloader force are shown in Figure 6.12. The coherence is generally good, except for the front aeroloaders for the tire forces on the opposite side of the car. The influence of aeroloader force on tire force can be better understood by considering the aeroloader configuration and the suspension load paths, illustrated in Figure 6.13. The aeroloaders are located laterally within the tire contact patch, while they are located close to the tires but between the axles longitudinally. Since the aeroloaders are close to the tires, the closest tire reacts most of the aeroloader force. Since the aeroloaders are offset from the tires in the longitudinal direction, a small amount of the aeroloader force will be supported by the distant axle. Since the aeroloaders have limited lateral offset from the tires, however, it might be expected that very little aeroloader force would be transmitted from one side of the car to the other. This would be true if all four corners had a fully independent suspension where the primary suspension stiffness could be represented directly above the tire. In the case of the solid rear axle, the springs that react to the aeroloader force are connected to the solid axle inboard of the tires. Since the spring is inboard, when a rear aeroloader pulls down on the car, it also causes a moment about the nearest spring, causing the opposite side of

the car to rise. This causes transmission of the rear aeroloader force to the opposite side of the car. In the case of the independent front suspension, the moment described above is reacted by the control arms where they attach to the chassis.




Figure 6.12. Tire force due to aeroloader input: (a) Time response, (b) FRF magnitude, (c) Coherence




Figure 6.13. (a) Aeroloader and rear spring locations, (b) Simplified rear axle model

Next, we consider the effect of wheelloaders on ride height, as shown in Figure 6.14. Ride height is defined as the vertical distance from a chosen point on the chassis to a point directly below the chassis point on a virtual ground plane defined by the wheelloader positions. Front and rear ride heights are calculated using chassis points along the vehicle centerline at the nose and spoiler, respectively. Like shock velocity, the wheelloaders do not change the ride height at low frequency, as the chassis moves like a rigid body. The coherence is good for the wheelloaders for the axle closest to the ride

height measurement, which also causes the highest response levels. The wheelloader inputs for the opposite axle also have good coherence, except at low frequency, which corresponds to low response levels.

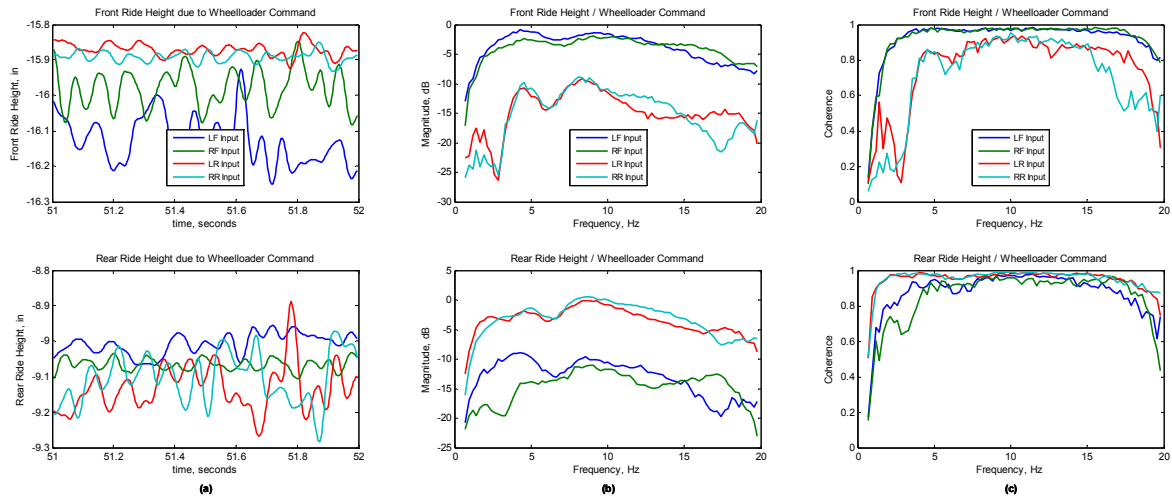


Figure 6.14. Ride height due to wheelloader input: (a) Time response, (b) FRF magnitude, (c) Coherence

Next, we consider the ride height due to aeroloader inputs in Figure 6.15. The response levels are highest for the aeroloaders closest to the ride height measurement. The coherence is generally good, with the exception of the influence of front aeroloaders on rear ride height, which corresponds to low response levels.




Figure 6.15. Ride height due to aeroloader input: (a) Time response, (b) FRF magnitude, (c) Coherence

Next, we consider the shock force inputs. Since we do not directly excite shock force, it is harder to estimate the single-input response. To estimate a single-input response, we consider the change in response from the baseline linear response for the multiple input random tests when one shock at a time is replaced by an alternate shock. The change in output, change in shock force, and FRF are calculated using the procedure described in Chapter 5. If all the baseline shocks are linear, this should give us a single-input FRF estimate and coherence. In reality, all the shocks are slightly nonlinear, and this small nonlinearity will act as additional inputs that are not accounted for in the single input model, slightly reducing coherence and FRF accuracy. Due to this, we will only use the procedure described in Chapter 5 to give us a qualitative estimate of the validity of the linear model for various input-output pairs, and we will adopt a different FRF estimation strategy which will be described in Section 6.3.1 for system identification.

Figure 6.16 shows the results for change in shock velocity due to change in shock force. The influence of changing the shock force on shock velocity is most evident at the corner the shock force is changed. The coherence for this collocated input-output pair is high, except near an anti-resonance at 4 Hz. The rear shock velocity coherence is also high at all frequencies except for the anti-resonance for both rear shock force inputs. The coherence is good at low frequencies for all input-output pairs, except for the shock

velocity due to shock force across the diagonals of the car, which also corresponds to low output levels.

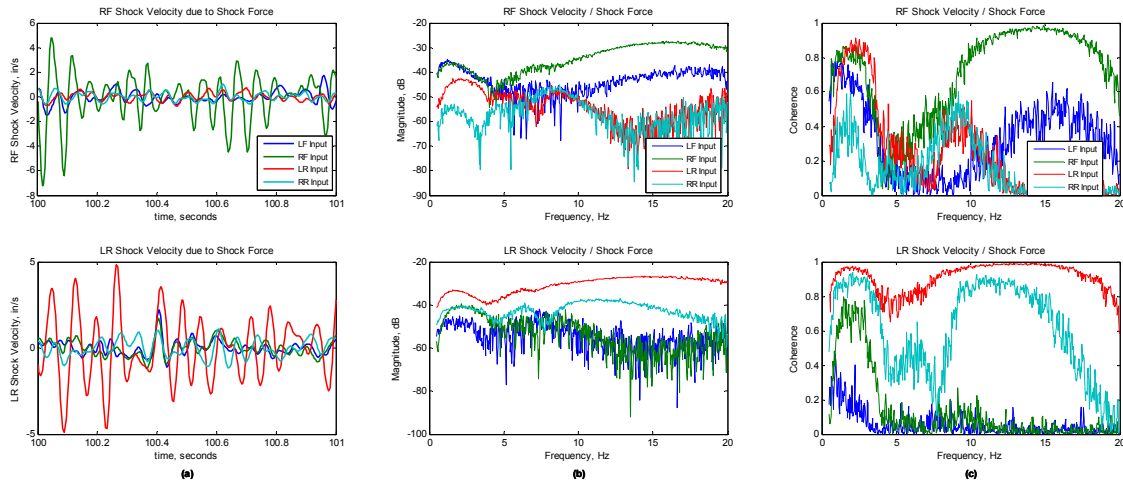


Figure 6.16. Shock velocity due to shock force input: (a) Change in time response, (b) FRF magnitude, (c) Coherence

Next, we consider the influence of shock force on wheelloader force in Figure 6.17. The coherence is good for tire force closest to the shock force and for rear tire force due to rear shock force. For the front shock force influence on the opposite side of the car, the coherence is low despite significant output energy, especially for the rear wheelloader force. This is important to note because this is the first time we have observed an input-output pair where the coherence was poor at all frequencies but the response level was still significant. For the remaining inputs, the coherence is low, except from 2-4 Hz and 8-14 Hz, corresponding to peaks in the FRFs.

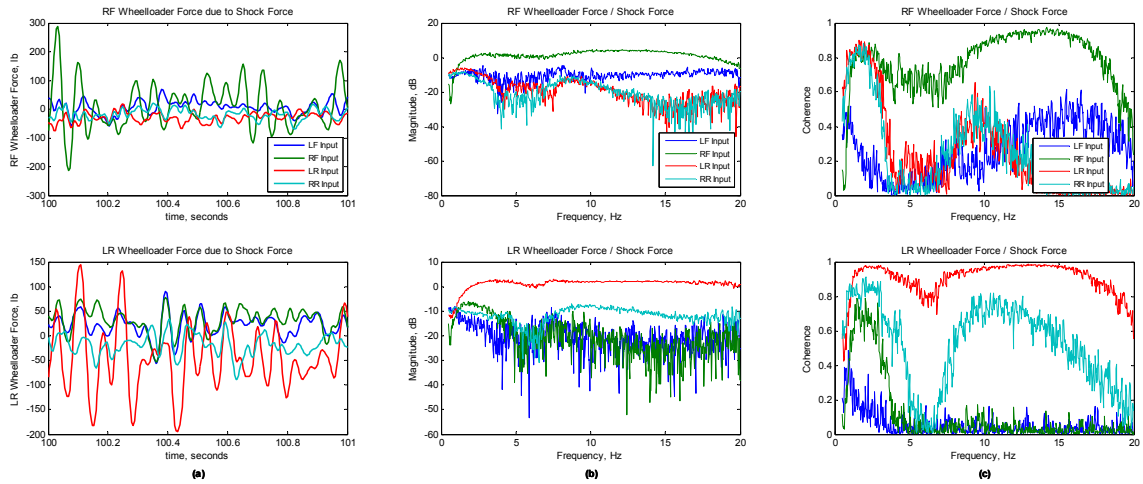


Figure 6.17. Tire force due to shock force input: (a) Change in time response, (b) FRF magnitude, (c) Coherence

The last set of results we consider is the ride height due to shock force, shown in Figure 6.18. The coherence for front ride height is good for all inputs from 1-4 Hz. Above 4 Hz, the FRF magnitude is below -70 dB, meaning 100 lb change in shock force gives less than 0.03 inches change in ride height. For the rear ride height, the results are similar.

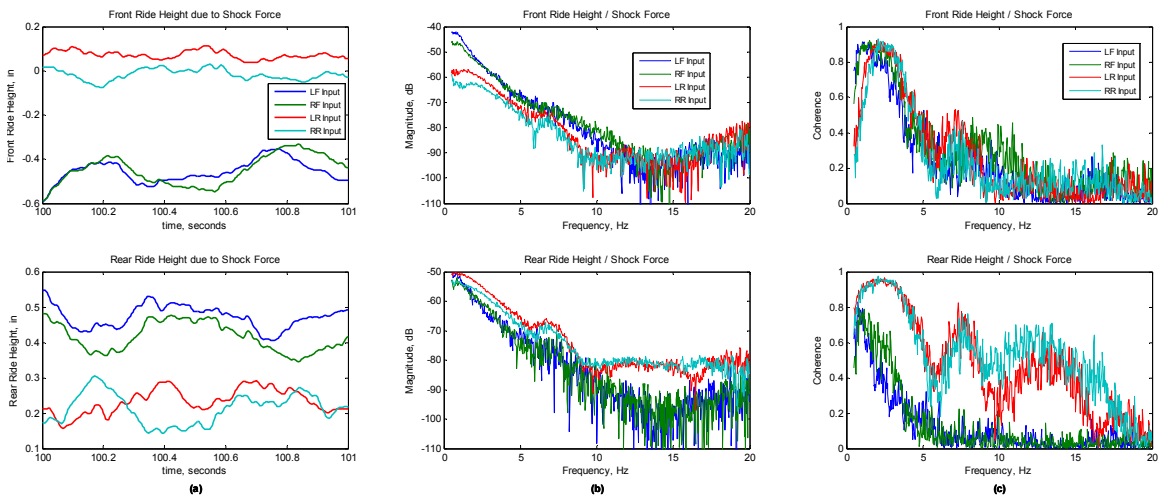


Figure 6.18. Ride height due to shock force input: (a) Change in time response, (b) FRF magnitude, (c) Coherence

This section has reviewed the accuracy of linear modeling for each input-output pair in terms of response level and coherence. This analysis has shown that aside from a few

exceptions, coherence is high when the response levels are high. This indicates that a linear model that has a frequency response similar to the FRF estimates where the coherence is good and has a low frequency response where the coherence is low should capture a large portion of the output energy. Other trends worth noting include that more energy was transmitted for input-output pairs located within the same corner or across the solid rear axle. This increased coupling also led to the best coherence levels. This information can now be applied in the system identification process to address frequency ranges where the FRF estimate is poor.

6.3 8-Post Rig System Identification

This section describes the methods and results for 8-post system identification. First, we describe the FRF estimation method used. Next, the results for FRF estimation are shown. Finally, the FRF estimate is used to fit transfer functions, which will be used to simulate the results of future rig tests.

6.3.1 FRF Estimation Method

In Chapter 5, we developed a FRF estimation method for wheelloader, aeroloader, and shock force inputs. It was shown that when the shocks are linear, the shock force is linearly dependent on the wheelloader and aeroloader inputs, making it impossible to identify the FRF. This issue was addressed by first performing a baseline experiment with 4 linear shocks, then performing 4 additional experiments, where one shock at a time was replaced by a second shock. This method allowed the FRF to be accurately identified, with the FRF estimate matching the known theoretical result.

During simulation, we can choose our shocks to be exactly linear. In practice, all shocks will exhibit some degree of nonlinearity. If we assume our shocks are linear, any variation from this linear model will cause errors in the FRF estimation and the resulting model. To accommodate realistic shock performance, the FRF estimation method is extended to accommodate arbitrary shock forces.

Start with the frequency response model

$$y = H_{rig} u_{rig} + H_{shocks} u_{shocks} \quad (6.3)$$

Assuming that most of the system energy dissipation is caused by the shocks, the FRF matrices H_{rig} and H_{shocks} represent a lightly damped system. This undamped system representation has a slowly-decaying impulse response, causing leakage errors when the data blocks are windowed out of a larger data file. The damping of the identified system may be arbitrarily increased by redistributing the energy dissipated by the shocks from the arbitrary shock force input to a linear shock force which is included in the FRF. The shock force input may be rewritten as a combination of an arbitrary linear shock forces and an additional arbitrary force

$$u_{shocks} = \bar{C} v_{shocks} + \Delta F \quad (6.4)$$

where \bar{C} is an arbitrary 4x4 diagonal matrix defining the linear shocks. The frequency response model may now be written as

$$\begin{aligned} y &= (I - H_{shocks} \bar{C} S)^{-1} (H_{rig} u_{rig} + H_{shocks} \Delta F) \\ y &= \tilde{H}_{rig} u_{rig} + \tilde{H}_{shocks} \Delta F \end{aligned} \quad (6.5)$$

where S is a matrix that selects the shock velocities from the output vector y . The FRF matrix \tilde{H}_{rig} now represents the behavior of a damped system with linear shocks defined by \bar{C} , while the FRF matrix \tilde{H}_{shocks} represents the influence of any additional damping force.

The vehicle must be tested using a linearly independent drivefile and for at least two different shocks at each corner. After all the data is collected for each experiment, ΔF is calculated at each time step. The data for each setup can be divided into data blocks and the FFT may be calculated. The measured response for the i^{th} setup and k^{th} data block may be written as

$$y^{(i)}[k] = \tilde{H}_{rig} u_{rig}[k] + \tilde{H}_{shocks} \Delta F^{(i)}[k] + \varepsilon^{(i)}[k] \quad (6.6)$$

The results can be rearranged in matrix form for I shock setups and K data block per setup as

$$\begin{bmatrix} y^{(1)}[1] \\ \vdots \\ y^{(1)}[K] \\ \vdots \\ y^{(I)}[1] \\ \vdots \\ y^{(I)}[K] \end{bmatrix} = \begin{bmatrix} u_{rig}^T[1] & (\Delta F^{(1)}[1])^T \\ \vdots & \vdots \\ u_{rig}^T[K] & (\Delta F^{(1)}[K])^T \\ \vdots & \vdots \\ u_{rig}^T[1] & (\Delta F^{(I)}[1])^T \\ \vdots & \vdots \\ u_{rig}^T[K] & (\Delta F^{(I)}[K])^T \end{bmatrix} \begin{bmatrix} \tilde{H}_{rig}^T \\ \tilde{H}_{shocks}^T \end{bmatrix} + \begin{bmatrix} \varepsilon^{(1)}[1] \\ \vdots \\ \varepsilon^{(1)}[K] \\ \vdots \\ \varepsilon^{(I)}[1] \\ \vdots \\ \varepsilon^{(I)}[K] \end{bmatrix} \quad (6.7)$$

$$Y = UH + E$$

This system of equations along with the error measure in Equation (5.31) yields a least-squares optimization problem whose solution satisfies the normal equations in Equation (5.32). This problem will have a unique solution when the input autospectrum matrix is invertible, which requires the input matrix U to be full-rank. This will tend to be true if the shocks tested are nonlinear, but the condition number of the input matrix may be large if the shocks tested are not highly nonlinear. This indicates that the matrix is almost rank-deficient, which will make the FRF estimates highly sensitive to noise.

One experimental procedure that will encourage good conditioning of the input matrix is based on the procedure presented in Chapter 5. First, we test with a baseline setup with near-linear shocks. If we define the diagonal of \bar{C} to be the slope of the best-fit line for the force-velocity plot for these shocks, then the ΔF terms will be almost zero for the baseline setup. Next, we change the left front shock with a significantly different shock and we repeat the test. This will make the ΔF term for the left front shock large, while the remaining ΔF terms will be small. Repeating this process for all four corners gives us an input matrix similar to Equation (5.40), having a significant amount of linearly independent input energy in each input column.

This procedure was used to collect the data for FRF estimation. First, a baseline setup was tested using the 800s random drivefile with almost linear shocks with a damping coefficient of 20.74 lb/(in/s). Next, the left front shock was swapped for an alternate shock with a damping coefficient of 9 lb/(in/s) and the 800s random drivefile was repeated. This process was repeated at the remaining three corners, where the alternate shock was run at the corner of interest, while all other shocks were the baseline shocks. This data gives us the information needed to perform FRF estimation.

6.3.2 FRF Estimation

The data collected from the 800 second drivefile for one baseline setup and four alternate setups was used to perform FRF estimation. Shock velocity was then calculated from shock pot data as described above. This shock velocity was used to calculate a linear shock force with damping coefficient of 20.74 lb/(in/s), which was subtracted from the measured shock force to obtain the change in shock force defined in Equation (6.4). Each 800 second data set was then divided into 50 data blocks using a Hamming window with 50% overlap, providing a frequency resolution of 0.03 Hz. The FFT was calculated for each data block, and this frequency-domain data was used to form Equation (6.7). The FRF was then estimated by solving the normal equations defined in Chapter 5. The total coherence for all outputs was generally high, with a drop in coherence from 4-6 Hz for the front tire force, which corresponded to an anti-resonance.

To prepare the FRF estimates for transfer function fitting, the FRFs were modified in regions where the FRF estimate is poor, as discussed in Section 6.2. First, the FRF for several input-output pair paths were zeroed due to low response levels and poor coherence. The zeroed paths are denoted by a “0” in Table 6.5. In cases when it was unclear if the path should be zeroed, the identification and validation process was carried out both with and without the path present in the mode, and the results were compared to determine if the path could be zeroed.




Figure 6.19. Total coherence for select outputs

Table 6.5. Zeroed input-output pairs

		Inputs											
		Wheelloaders				Aeroloaders				Shocks			
		LF	RF	LR	RR	LF	RF	LR	RR	LF	RF	LR	RR
Shock Velocity	LF	0	0	0								0	
	RF	0		0	0								0
	LR	0	0									0	
	RR	0	0								0		
Tire Force	LF	0	0	0	0		0				0		
	RF	0		0	0	0				0			
	LR	0	0				0			0			
	RR	0	0			0					0		
Ride Height	Front			0	0								
	Rear	0	0										
Hub Accels	LF			0	0	0	0	0	0			0	0
	RF			0	0	0	0	0	0			0	0
	LR	0	0			0	0	0	0	0	0		
	RR	0	0			0	0	0	0	0	0		

An important consideration in shock selection is if shock selection at a particular corner has any influence on sensor response. By zeroing paths for shock inputs in the matrix \tilde{H}_{shocks} , it is possible that the influence of shock selection at a particular corner on a particular sensor response may be eliminated. This is more complex than it may seem at first glance, since changes in shock force input at one corner can also change the shock

velocity and shock force at a different corner. For example, while the entries in the FRF matrix \tilde{H}_{shocks} corresponding to the shock velocity due to left rear shock force are zeroed, the left rear shock force does affect the shock velocity and shock force at the other three corners, which do influence the left front shock velocity. Define a logical matrix $[h_{ij}]$ which has zeroes if the corresponding entry in \tilde{H}_{shocks} has been zeroed. The number of nonzero shock velocity feedback paths for the i^{th} output and j^{th} shock force is

$$\sum_{k=9}^{12} h_{ik} h_{kj} \quad (6.8)$$

Applying this method for the zeroing described in Table 6.5 yields the matrix shown in Table 6.6. This shows that despite zeroing some FRF pairs, all shock changes will influence each output.

Table 6.6. Number of active feedback paths

		Inputs				
		Shocks				
		LF	RF	LR	RR	
Outputs	Shock Velocity	LF	3	2	2	2
		RF	2	3	2	2
		LR	2	2	3	2
		RR	2	2	2	3
	Tire Force	LF	2	2	2	3
		RF	2	2	3	2
		LR	2	2	3	2
		RR	2	2	2	3
	Ride Height	Front	3	3	3	3
		Rear	3	3	3	3
	Hub Accels	LF	2	2	1	1
		RF	2	2	1	1
		LR	1	1	2	2
		RR	1	1	2	2

In addition to the input-output pairs that were zeroed because they did not significantly contribute to the response, the FRF estimate was modified in frequency ranges where it was poor due to low response level in the given frequency range. Specifically, the FRF for aeroloader inputs above 5 Hz and for other inputs for the off-diagonal paths at higher frequencies was replaced with a two-point high frequency residual fit defined in Equation

(5.48). This residual fit replaces the poor FRF estimate with a FRF estimate with low response level, which will aid in fitting transfer functions.

To further smooth the FRF estimate prior to parametric identification, the FRFs are lowpass filtered. Filtering a complex vector may be performed analogously to filtering a real vector, with real filters being applied to the magnitude in dB and the unwrapped phase. A 4th order Butterworth filter was constructed to attenuate FRF variations with a width of 0.6 Hz or less. This filter was then applied to the magnitude in dB and the unwrapped phase using the MATLAB function *filtfilt* to avoid frequency shift. This eliminates large jumps in the FRF estimates, forcing the independently estimated FRF at each frequency to behave more smoothly.

The FRF estimates and modified estimates are shown in Figures 6.20-6.23 for shock velocity, tire force, ride height, and hub acceleration. For brevity, only the results for the right front and left rear locations are shown for the shock velocity, tire force, and ride height. The modified FRF estimate is not shown for paths that are zeroed.

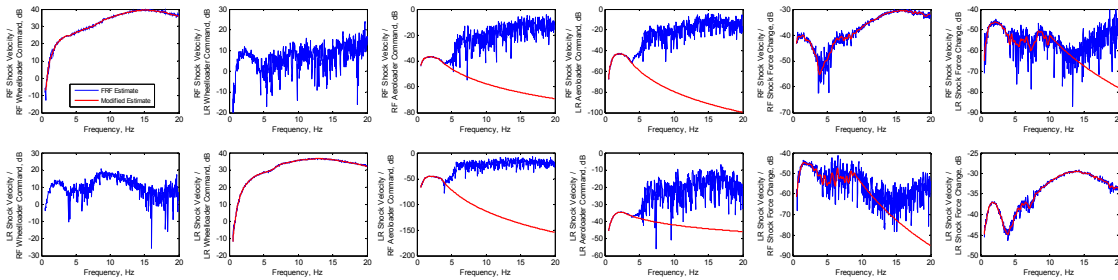


Figure 6.20. Shock velocity FRF estimate

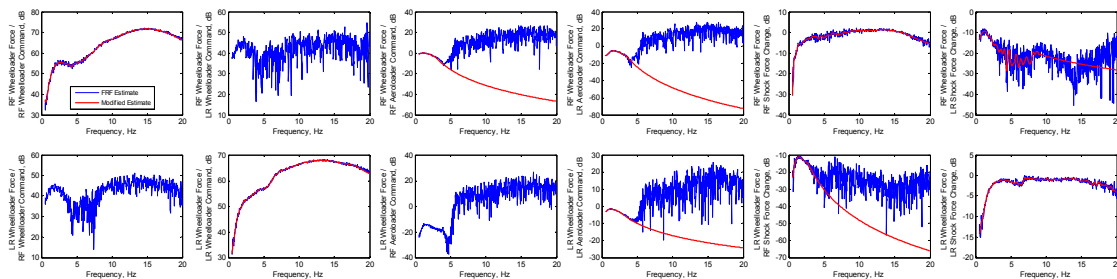


Figure 6.21. Tire force FRF estimate

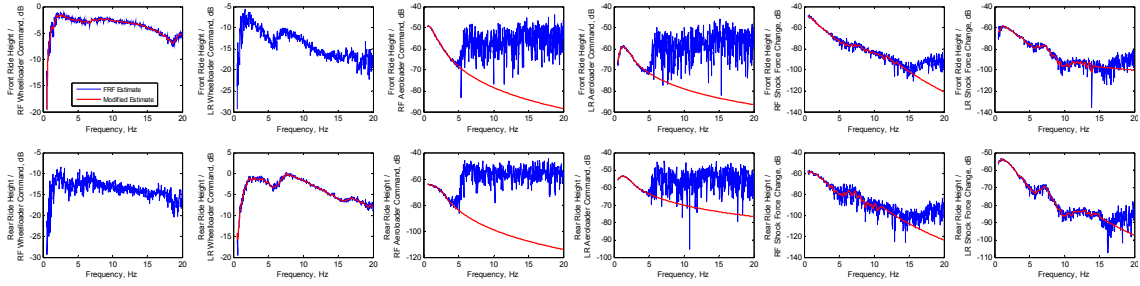


Figure 6.22. Ride height FRF estimate

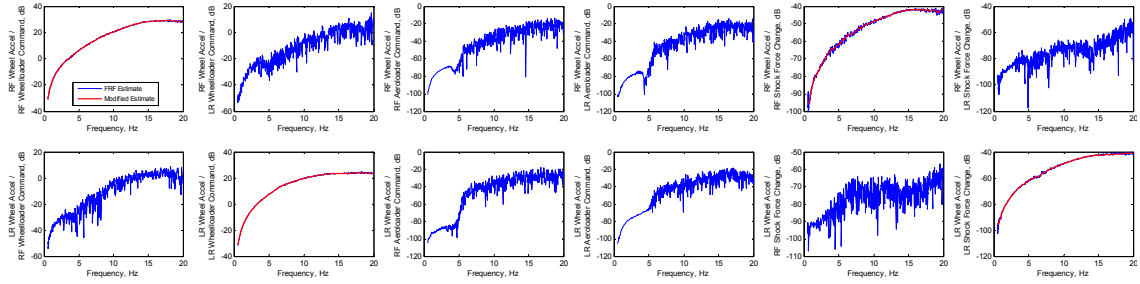


Figure 6.23. Hub acceleration FRF estimate

6.3.3 Parametric Identification

After performing FRF estimation, the next step is to fit a transfer function to each input-output pair using *invfreqz*. The denominator polynomial order was increased from 4 until acceptable matching was achieved without creating artificial modes. This resulted in denominator orders ranging from as low as 4 for aero/loader inputs to as high as 14 for wheel/loader and shock inputs. The weighting functions were

$$w_{aero}(f) = \begin{cases} 1 & f \in [0.1, 5) \\ 0.5 & f \in [5, 7) \\ 0.1 & f \in [7, 12) \\ 0 & \text{otherwise} \end{cases} \quad w_{wheels}(f) = w_{shocks}(f) = \begin{cases} 1 & f \in [0.5, 10) \\ 0.5 & f \in [10, 15) \\ 0.1 & f \in [15, 20) \\ 0 & \text{otherwise} \end{cases} \quad (6.9)$$

The weighting functions were adjusted slightly as needed to provide better matching in different frequency ranges. The *invfreqz* function provided discrete-time transfer functions with a sample rate of 125 Hz, matching the desired simulation rate. The

resulting transfer functions matched the modified FRF estimates well. An example of the resulting transfer function fits is shown in Figure 6.24.

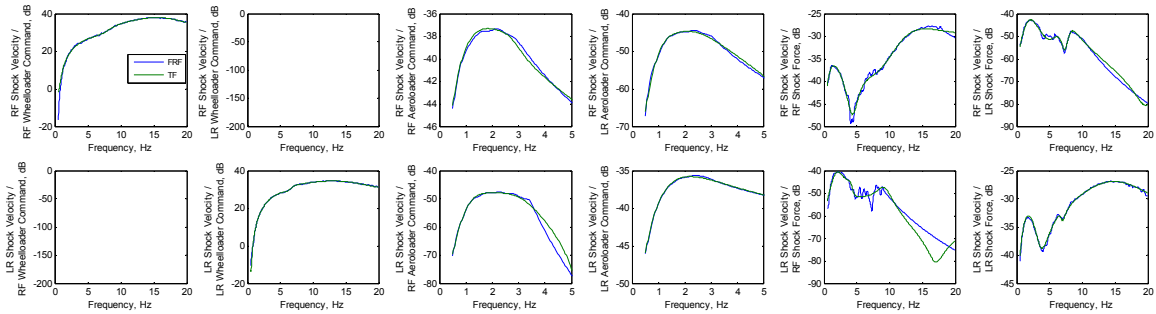


Figure 6.24. Transfer function fits for shock velocity

6.4 Model Validation

Now that we have identified a vehicle model that predicts the influence of drivefile selection and shock forces on sensor response, the vehicle model can be combined with our shock models to provide simulations of 8-post rig tests. First, we describe the modeling structure created in Simulink. Next, we validate the model accuracy by comparing simulation and experimental results.

6.4.1 Simulations with Identified Model

The discrete-time transfer function model identified in the previous section can now be applied to perform simulations to predict the outcome of future 8-post rig tests for varying drivefiles and shock selections. The transfer function model is converted to state-space and is then combined with shock models in Simulink, as shown in Figure 6.25.

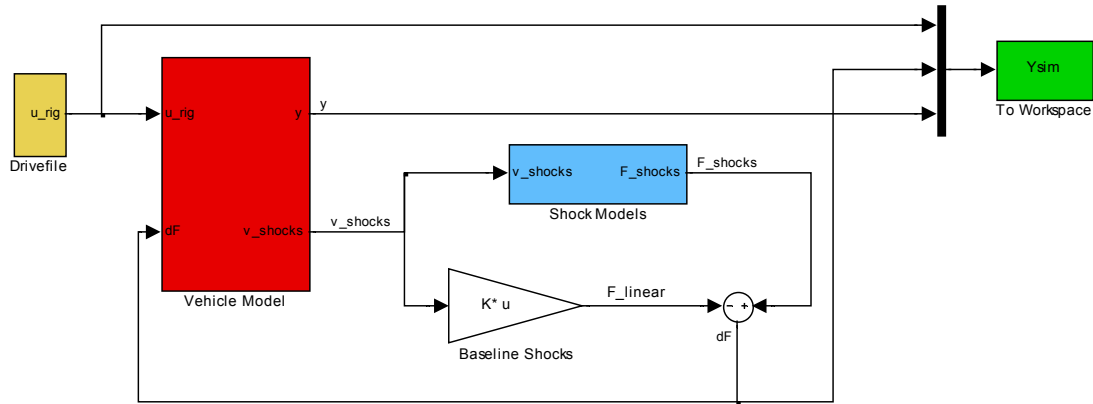


Figure 6.25. Simulink model

Simulations were run with a sample rate of 125 Hz, consistent with the sample rate of the identified discrete-time model. Drivefiles that were run on the 8-post rig were decimated from the controller frequency of 1250 Hz to 125 Hz using the MATLAB function `decimate`, which first lowpass filters the original signal using an 8th order Chebyshev Type I filter with a cutoff frequency of 80% of the reduced sample rate's Nyquist frequency before downsampling the signal. The 8-post rig tests described in Section 6.1 are then reproduced by applying the same drivefiles and using shock models for the same shock setups. The simulation results are stored for comparison with experimental data

6.4.2 Track-Based Drivefile Testing

The first set of validation results that we show are for the Richmond International Raceway track drivefile. This is the type of drivefile that race teams would commonly use during an 8-post rig testing session to characterize their vehicle's performance on a particular racetrack as components such as shocks are changed. First, we show the time response for 4 different signals for the baseline setup in Figure 6.26. As these plots show, the simulation results show generally good agreement with the experimental data. The shock velocity shows excellent agreement, which is critical since we will use this simulated shock velocity to calculate shock force.




Figure 6.26. Baseline response for Richmond drivefile

While accurately predicting the time signals is important, it is critical that the model is able to accurately predict relative trends in the response as the vehicle setup is changed. RMS-based signal measures are commonly used to characterize the relative behavior of different setups. Figure 6.27 shows the RMS trending for the signals shown in Figure 6.27 as the shock setup is changed from the baseline setup to the four alternate setups. The response level for the left rear shock velocity changes the most for a change in left rear shock, does not change from the baseline significantly for a change in the front shocks, and only changes slightly for a change in the right rear shock. This trending is accurately predicted by the model. The simulation trending for the total tire force and front ride height trend similarly with the data, with the exception of the change in right front shock not trending as well. The relative trending for the left rear accelerometer matches extremely well, with the simulation absolute results about 0.4 g higher than the data.

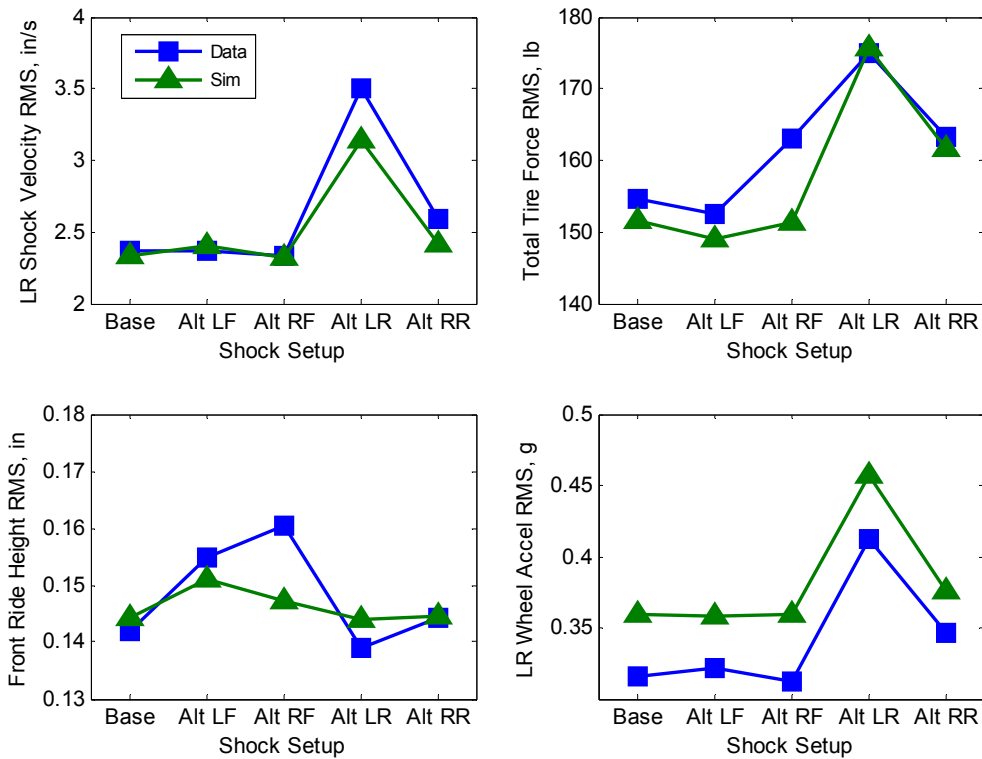


Figure 6.27. Shock setup RMS trending for Richmond drivefile

The intended purpose of this model is to determine which setups will tend to have a high RMS level and which will have a low RMS level for various RMS-based measures. These RMS level predictions can be used to pre-screen potential setups before 8-post rig testing, eliminating poor setups, selecting good setups, and identifying trends to improve efficiency of 8-post testing days. The classification of setups as good or bad is based on the relative RMS response levels, not the absolute RMS levels. If the relationship between the simulated RMS levels and the measured RMS levels are monotonic, response levels will be accurately categorized. To clearly see how the simulation and data trending agree, the RMS response levels for the data is plotted against the simulation results in Figure 6.28. These plots show that the simulation accurately predicted the trend except at one point for the total tire force and one point for the front ride height. These points correspond to the left rear setup.

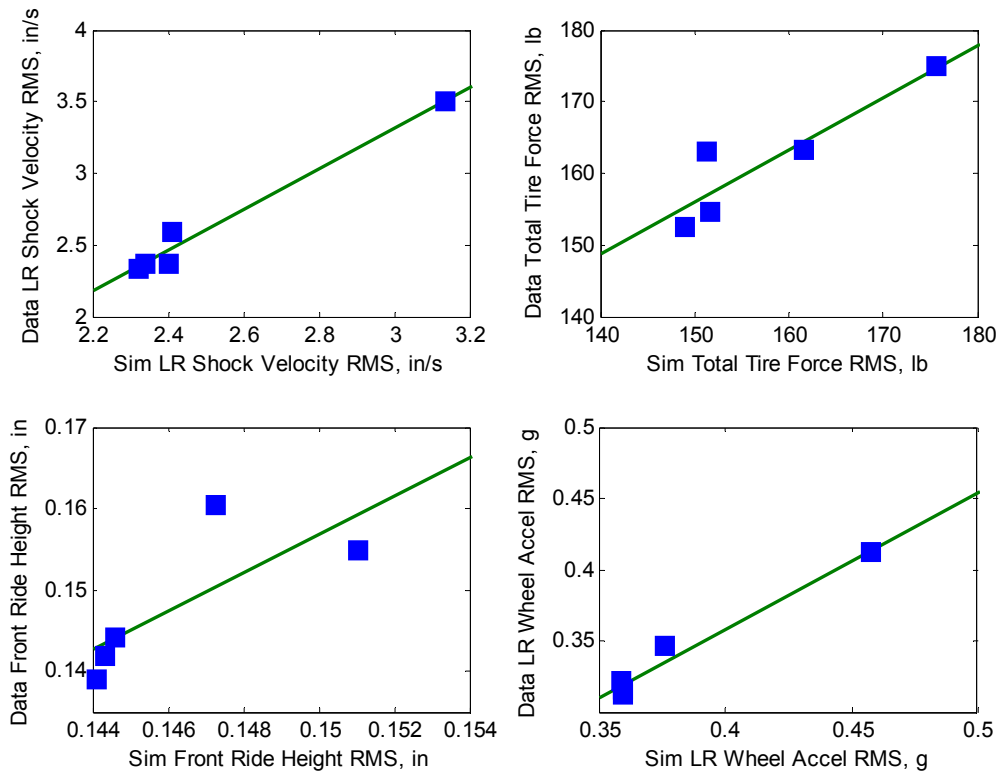


Figure 6.28. Shock setup RMS correlation for Richmond drivefile

6.4.3 Sine and Bump Testing

The next validation is wheelloader sine testing in heave, pitch, and roll on the baseline setup. This allows us to see how well our linear model can match the gain and phase of the experimental results. Figure 6.29 shows the sine testing heave time signals for 7 frequencies and 4 signals, while Figure 6.30 shows the corresponding RMS amplitudes. The simulation results generally look similar to the experimental data, with some phase or amplitude error. The tire force and wheel accelerometer show a factor of 3 harmonic below about 6-8 Hz, while the simulation looks sinusoidal. This indicates that the model does not have sufficient nonlinearity to precisely reproduce the response of these signals in this frequency range.

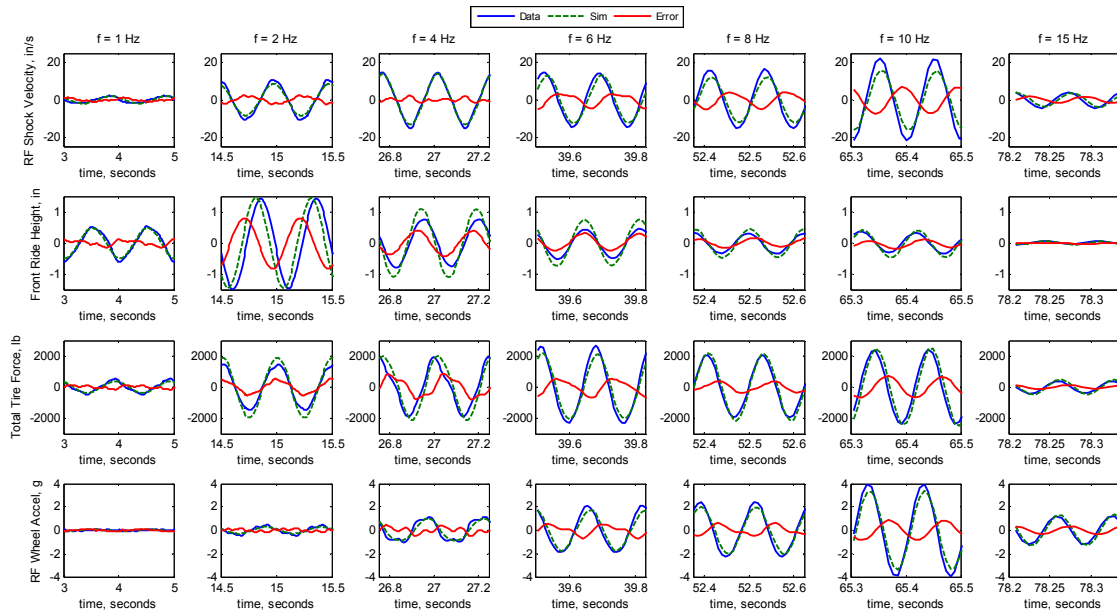


Figure 6.29. Sine heave time results

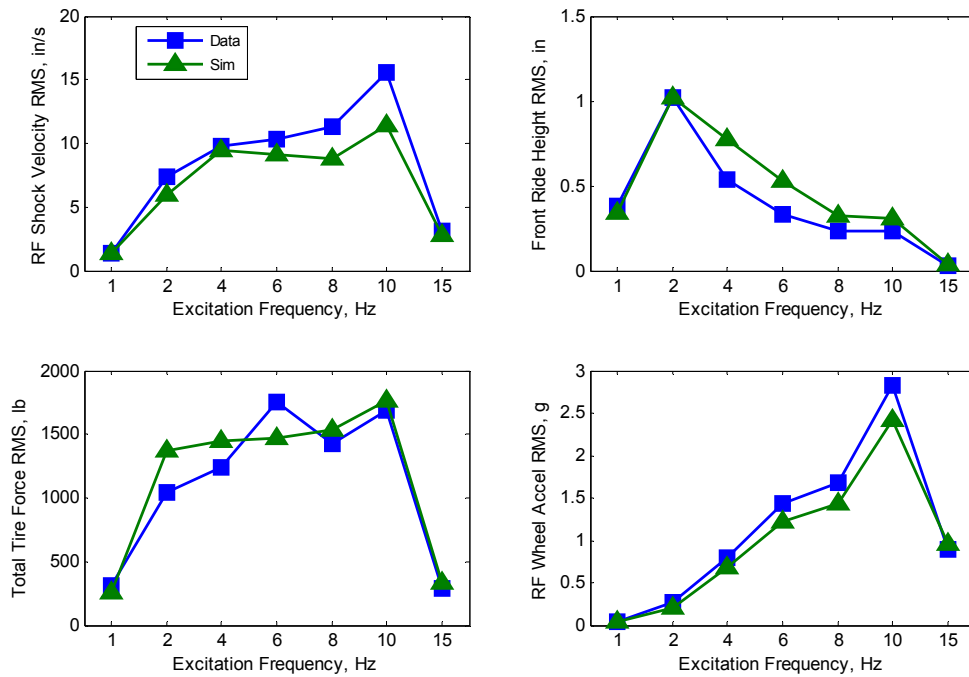


Figure 6.30. Sine heave testing RMS results

Next, we show the results for bump testing in Figure 6.31 for 4 signals during an upward wheelloader bump in heave. The data was lowpass filtered with a cutoff frequency of 60 Hz to remove high frequency content. The simulations show good agreement with the

measured data for this 0.4 inch peak-peak square wave, indicating that the model provides a reasonable estimate of vehicle response to short wavelength track inputs commonly found on the track surface.

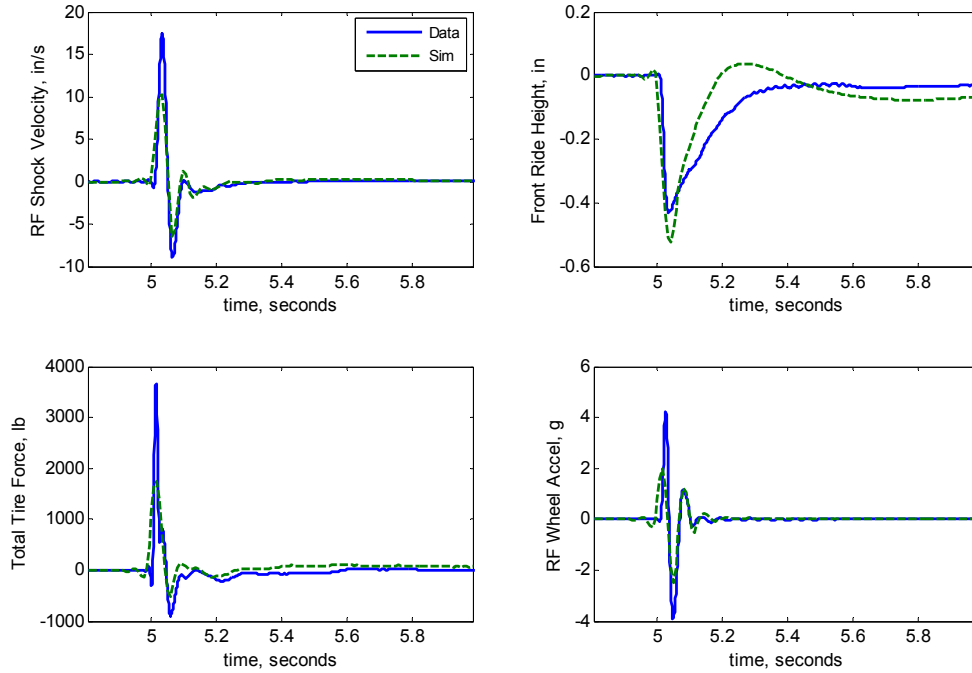


Figure 6.31. Bump testing results

6.4.4 Static Testing

Static testing was performed by loading and unloading the chassis using a series of 5 static aerolader forces levels held constant for 5 seconds in heave, pitch, roll, warp, and single-actuator input configurations. This provides a large dataset to estimate and validate the static behavior of the vehicle. The average force and position for each static position was calculated and the data was used to create a least-squares estimate of a vehicle spring matrix of the form

$$F_{aero} = KX_{aero} + F_{zero} \quad (6.10)$$

where F_{aero} and X_{aero} are the force and position of the four aeroloaders, while K and F_{zero} are the estimated spring matrix and force at zero position. The spring matrix may then be transformed if desired using Equations (6.1) and (6.2) to describe the relationship

between actuator mode displacements and mode forces. This representation of the stiffness matrix was found to be

$$K_m = T_f K T_x^{-1} = \begin{bmatrix} 0.71 & 0.00 & -0.11 & -1.20 \\ -0.14 & 3.21 & -0.07 & 0.57 \\ -0.47 & -0.06 & 0.66 & -3.62 \\ 1.38 & 1.92 & 0.98 & 27.90 \end{bmatrix} \times 10^3 \quad (6.11)$$

where heave force is in lb, heave/pitch/roll moments are in ft-lb, heave displacement is in inches, and heave/pitch/roll angles are in degrees. The estimated force can then be compared with the measured force, as shown in Figure 6.32 for the static heave test. Similar results exist for the other tests.




Figure 6.32. Static testing results

The validation results show that the spring model show generally good agreement with experimental data. The static positions in the left front, right front, and left rear showed some hysteresis between loading and unloading, possibly due to friction in the suspension or at the tire-wheel platen interface. A simple spring model cannot account for this hysteresis, but the weighting on the objective function used in the linear regression could be adjusted to redistribute the hysteresis error as desired.

If DC response was desired for analysis, it could be added to the simulation. Since we are investigating how shock selection influences the trends in the vehicle response and

the DC actuator inputs do not influence shock velocity, we will not include this in our model.

6.5 Summary

This section has developed and applied a method to identify a vehicle model using data collected from an actual 8-post rig test to predict the outcome of future 8-post rig tests for different drivefile and shock selections. The model was then validated by using it to perform simulations for several different drivefiles and shock setups and comparing it to experimental data. This analysis showed us that the model generally does a good job of predicting the amplitude trends of experiments, with some amplitude and phase error. There was also nonlinearity not accurately captured by the model in tire force and hub accelerometer measurements.

Chapter 7

Predicting the Influence of Shock Absorber Setup

This chapter applies the vehicle model identified using 8-post rig data in Chapter 6 to predict the influence of shock absorber setup for a large number of shock setups. After first describing the shock build database included in the shock selection process, the results of simulation are shown. The predicted trends are then verified by testing a select number of setups on the 8-post rig that have been determined to exhibit either a low level or high level of various RMS-based response measures by the simulation.

7.1 Shock Build Database

To illustrate the ability of our identified model to predict the response levels of various RMS-based response measures, we must first define the search space of potential shock setups. The empirical shock model developed in Chapter 3 fits a model to shock dynamometer data collected from a specific shock build, giving us a model that accurately represents the performance of a real shock absorber. This shock modeling structure leads us to use a discrete shock build selection space. If we model N shock builds and we assume that each shock build can be placed at as many locations as desired, we have N^4 possible shock setups to choose from. Since each simulation of the Richmond International Raceway takes approximately 30 seconds, an estimate of the time to run all possible combinations may be calculated, as shown in Table 7.1.


Table 7.1. Number of shock builds and simulation times

Shocks	Setups	Sim Time, h
1	1	0.01
2	16	0.13
3	81	0.68
4	256	2.13
5	625	5.21
6	1,296	10.80
7	2,401	20.01
8	4,096	34.13
9	6,561	54.68
10	10,000	83.33

In practical applications, it will often be desired to optimize a chosen objective function over a large number of shock builds to find the best setups. While some of the potential combinations may be eliminated if their objective function shows limited sensitivity to certain shock selections, it will often be impractical to run all possible combinations. When the simulation time is too large to consider exhaustive search, more efficient optimization methods that can deal with discrete search spaces can be considered, such as genetic algorithms. Since we do not have a desired objective function and our goal is only to show that our model can distinguish between shock setups that provide low and high response levels of various response measures, we will choose a smaller shock build space that will allow us to perform an exhaustive search.

Since we are not trying to cover the entire space of all potential shock builds with our shock database, our only goal in defining our shock database is to assure that our different shock builds provide a noticeable change in response to make each simulation meaningful. To provide good variation between our different shock builds, we first define low, medium, and high damping levels. Standard shim sets are denoted by letters, which correspond to a specific shim thickness, as shown in Table 7.2. Each shim set consists of four shims with the specified stiffness inches and 4 different outer diameters arranged in a pyramid. In this chapter, we will denote shim settings by the compression shim stack letter followed by the rebound shim stack letter. For example, a C/B build has a C shim stack in compression and a B shim stack in rebound.

Table 7.2. Shim sizes (adapted from [9], used with permission of Randy Lawrence, President, Penske Racing Shocks, 2009)



Part #						
VS-AA	AA	.004	.004	.004	.004	Constant
VS-AAP	AA+	.004	.004	.006	.006	Constant
VS-AM	A-	.006	.006	.004	.004	Constant
VS-A	A	.006	.006	.006	.006	Constant
VS-AP	A+	.006	.006	.008	.008	Constant
VS-BM	B-	.008	.008	.006	.006	Constant
VS-B	B	.008	.008	.008	.008	Constant
VS-BP	B+	.008	.008	.010	.010	Constant
VS-CM	C-	.010	.010	.008	.008	Constant
VS-C	C	.010	.010	.010	.010	Constant
VS-CP	C+	.010	.010	.012	.012	Constant
VS-DM	D-	.012	.012	.010	.010	Constant
VS-D	D	.012	.012	.012	.012	Constant
VS-DP	D+	.012	.012	.015	.015	Constant
VS-EM	E-	.015	.015	.012	.012	Constant
VS-E	E	.015	.015	.015	.015	Constant
VS-EP	E+	.015	.015	.020	.020	Constant
VS-FM	F-	.020	.020	.015	.015	Constant
VS-F	F	.020	.020	.020	.020	Constant

The smallest shim set commonly available is an AA shim set, which consists of four shims with a thickness of 0.004 inches each. Using this shim set on a linear piston provides a nominal damping coefficient of approximately 9 lb/(in/s). The alternate shock used in Chapter 6 used the AA shim stack in both compression and rebound, which we denote as an AA/AA build. The AA shim stack will be used as our low shim setting.

Next, we define middle and high levels of damping. The B shim stack provides a nominal damping coefficient of approximately 20 lb/(in/s), while the C shim stack provides a nominal damping coefficient of approximately 30 lb/(in/s). These shim stacks provide more than twice and three times the low shim stack damping. These shim stacks will be used for the middle and high settings. The B shim is used in the baseline shocks use in Chapter 6 in both compression and rebound, which is denoted as a B/B build.

Now that the low, medium, and high levels of damping have been defined, we will now select our shock absorbers to be included in our shock database. If we ran all possible

combinations of low, medium, and high shim stacks in both compression and rebound, we would have 9 shocks in our database. Eliminating the combinations low/high and high/low gives us 7 shocks within our shock database. An exhaustive search with 7 shocks requires 2401 simulations and an estimated simulation time of 20 hours.

Table 7.3. Shock build database

		Rebound Shims		
		AA	B	C
Compression Shims	AA	X	X	
	B	X	X	X
	C		X	X

These 7 shocks were built and tested on the Roehrig shock dynamometer using a drive profile generated from the left front shock displacement measured during 8-post rig testing for the Richmond drivefile and the baseline shock setup. Shock models were then created for each shock build to the shock dynamometer data using the methods described in Chapter 3. These shock models along with the vehicle model were used to perform all 2401 simulations of every shock setup possible with the 7 shocks using the Richmond drivefile.

7.2 Simulated Shock Trends

After all 2401 simulations were performed, RMS levels of all signals in each simulation were calculated for comparison. In this section, we consider the trending for RMS response levels for the front ride height, total tire force, and right front hub acceleration. All RMS response levels have been normalized to range from 0 to 1, 0 being the lowest RMS level seen in the 2401 simulations for that signal, 1 being the highest. The extreme RMS values for each signal used for normalization are shown in Table 7.4

Table 7.4. RMS level extremes

	min	max
Front Ride Height, in	0.158	0.180
Rear Ride Height, in	0.085	0.099
RF Wheel Accel, g	0.20	0.28
LR Wheel Accel, g	0.32	0.47
Total Tire Force, lb	149	204

Figure 7.1 shows the RMS response levels for the front ride height and the associated shock setups. This figure clearly shows that the best shock setups for front ride height variation use either a C/C or a C/B shock build. No other front shocks give a front ride height variation within the lowest 7%. No trends can be noticed for the rear shocks, indicating that the front ride height has very low sensitivity to rear shock selection.

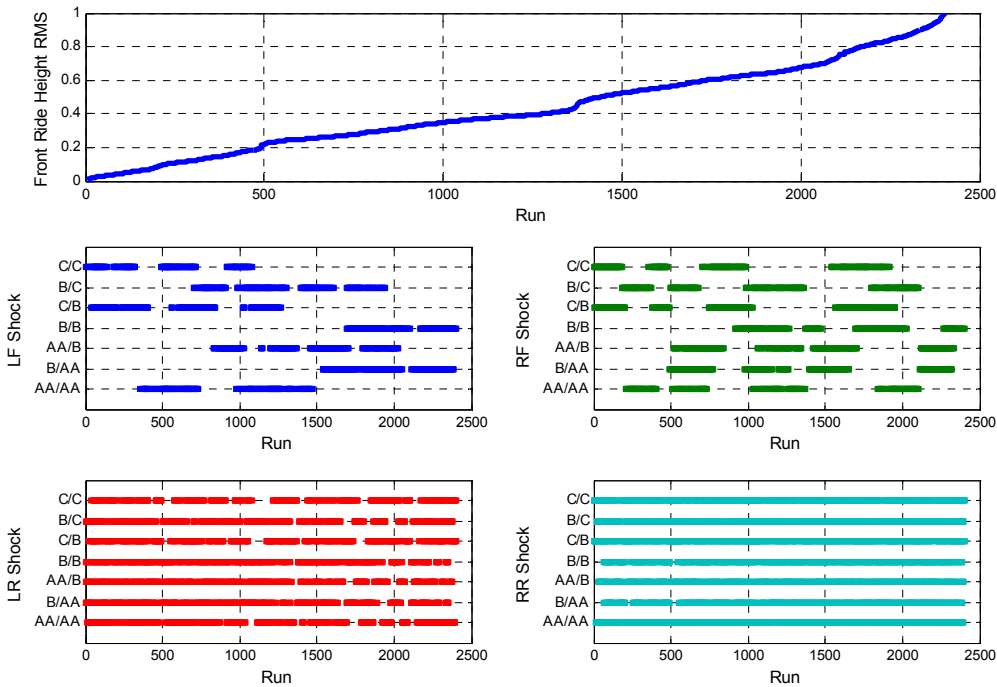


Figure 7.1. Front ride height simulation trending

To see how other response measures can vary given the front ride height RMS level, four response levels are plotted against front ride height RMS in Figure 7.2. While Figure 7.1 indicated that the rear shock selection had limited influence on front ride height, Figure 7.2 shows that tuning front ride height below 5% forces rear ride height variation to exceed 30%. If both front and rear ride height were key response measures, a compromise must be made. There is a similar tradeoff with total tire force – if front ride

height RMS is tuned below 2%, the tire force RMS cannot be less than 10%. There is a strong positive correlation for front ride height and right front hub acceleration, indicating that both measures will be reduced simultaneously. The left rear hub acceleration seems to be independent of front ride height, indicating that the two measures can be tuned independently.

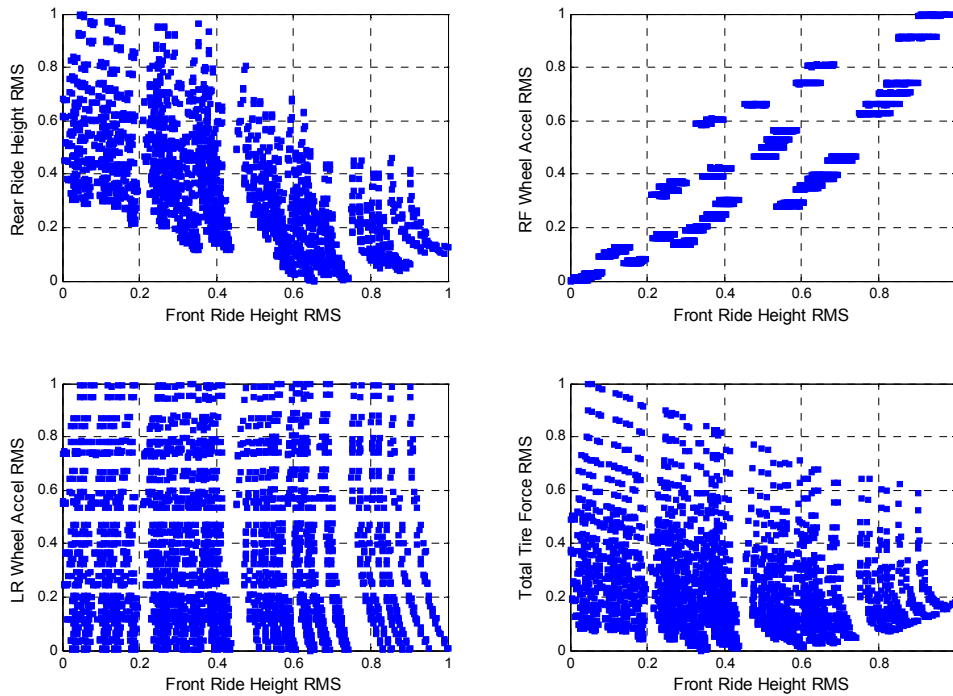


Figure 7.2. Front ride height simulation tradeoffs

The next response measure that we consider is the RMS total tire force. This shows that the worst shock builds for the left front shock are the C/C and C/B. These shock builds were shown to be the best for front ride height. In contrast, these shocks are amongst the best rear shocks for total tire force.

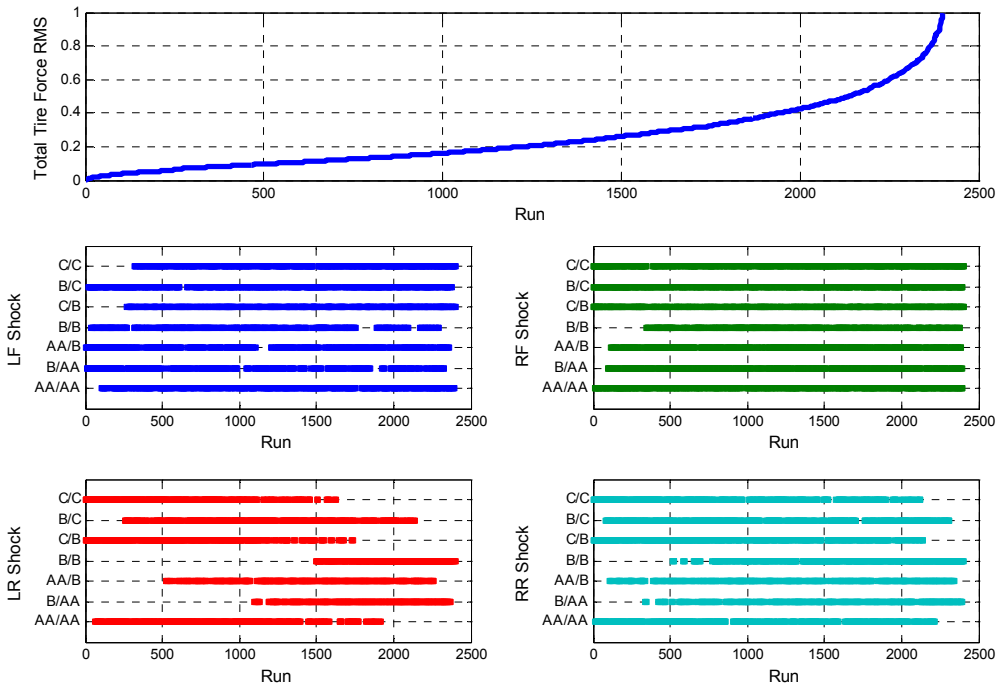


Figure 7.3. Total tire force simulation trending

Figure 7.4 illustrates the tradeoffs between total tire force and other response measures. It has already been shown that there is a tradeoff between total tire force and front ride height. This plot shows us that total tire force, rear ride height, and left rear hub acceleration can be tuned simultaneously with minimal tradeoff. The front ride height and right front hub acceleration exhibit a tradeoff with the total tire force.

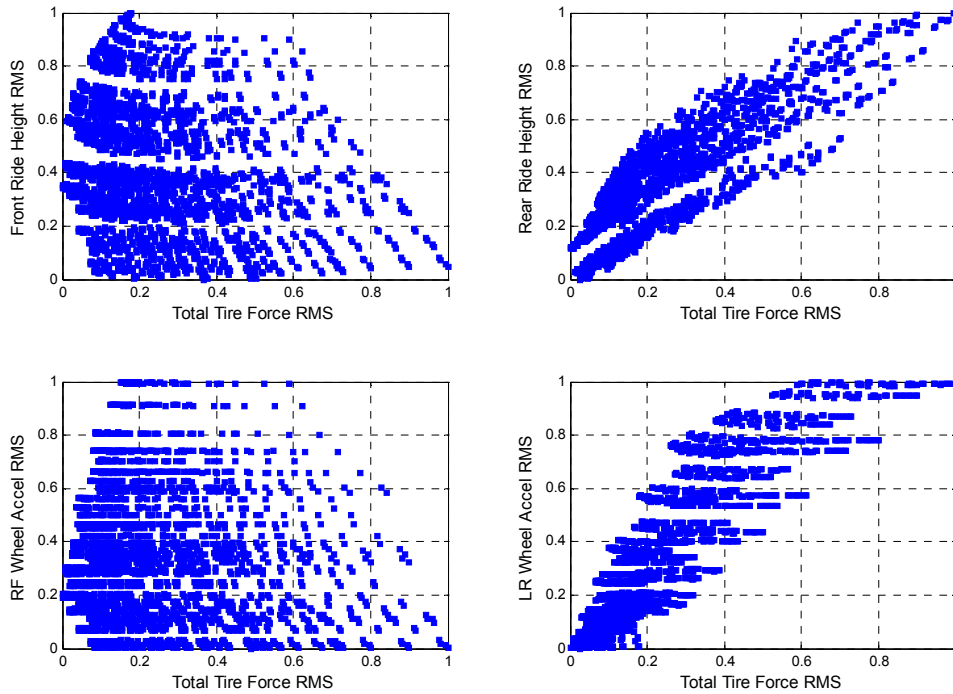


Figure 7.4. Total tire force simulation tradeoffs

The last signal that we consider is the right front hub acceleration in Figures 7.5 and 7.6. The best front shocks for right front hub acceleration are *C/C*, *C/B*, and *AA/AA*. The rear shocks have a limited influence. It has already been shown that the right front hub accelerometer is correlated with front ride height and exhibits a tradeoff with the total tire force. Also shown in Figure 7.6 is that there is a significant tradeoff with rear ride height and that the two hub accelerations can be tuned independently.

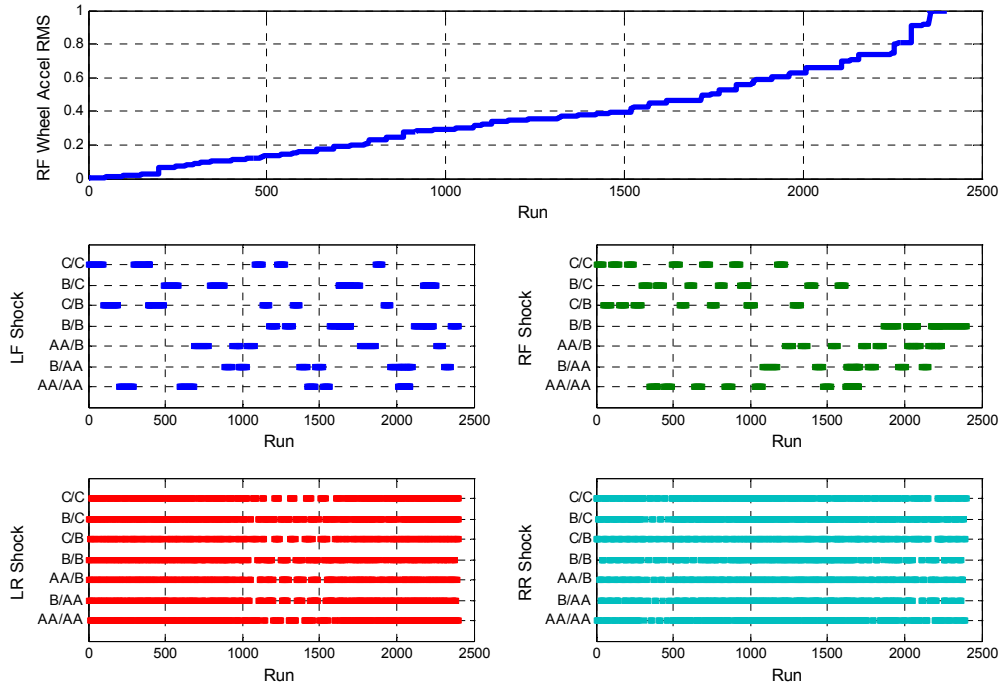


Figure 7.5. Right front hub acceleration simulation trending

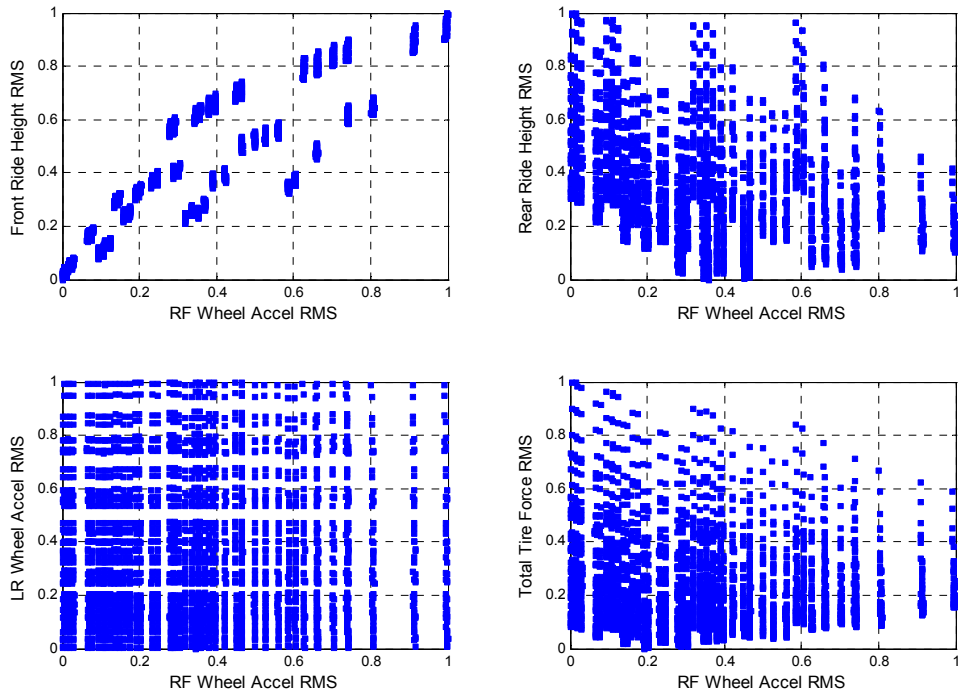


Figure 7.6. Right front hub acceleration simulation tradeoffs

7.3 Experimental Results

Twenty shock setups were selected for 8-post rig testing that were predicted to have either high or low levels of various response measures, as shown in Table 7.5. Four low setups and four high setups were selected for front ride height RMS. Four low setups and four high setups were also selected for total tire force RMS. Two low setups and two high setups were selected for right front hub RMS. These 20 setups were tested on the 8-post rig using the Richmond drivefile and the experimental results were compared with the simulation. To compare the ability to predict trends to the run-to-run variability, the standard deviation of the RMS measures calculated for 5 consecutive runs of the same Richmond drivefile on the baseline setup. This run-run variation is used to apply error bars to the comparison plots.

Table 7.5. Shock setups for 8-post rig testing

Low Runs					High Runs						
		LF	RF	LR	RR			LF	RF	LR	RR
Front Ride Height	1	C/B	C/B	AA/B	C/C	Front Ride Height	11	AA/AA	AA/AA	C/B	C/B
	2	C/B	C/B	AA/AA	C/C		12	AA/AA	AA/AA	C/B	C/C
	3	C/B	C/B	AA/AA	C/B		13	AA/AA	AA/AA	C/C	C/C
	4	C/B	C/B	AA/B	C/B		14	AA/AA	AA/AA	C/C	C/B
Total Tire Force	5	B/AA	C/B	C/B	C/B	Total Tire Force	15	C/B	C/B	AA/AA	AA/AA
	6	B/AA	C/C	C/B	C/B		16	C/B	C/C	AA/AA	AA/AA
	7	B/AA	C/C	C/C	C/B		17	C/C	C/C	AA/AA	AA/AA
	8	B/AA	C/B	C/C	C/B		18	C/C	C/B	AA/AA	AA/AA
RF Accel	9	C/B	C/B	AA/AA	B/AA	RF Accel	19	B/AA	B/AA	C/B	C/B
	10	C/B	C/C	AA/AA	B/AA		20	AA/AA	AA/AA	C/B	C/B

A comparison of the simulated and measured results for front ride height, total tire force, and right front hub acceleration is shown in Figure 7.7. For each response measure, only the setups identified for each response measure as a potential high or low setup in Table 7.5 are shown. For the front ride height, the absolute response levels do not match accurately, but the relative trend does match. This is clearly illustrated when the simulation and measured front ride height RMS levels are plotted against each other. Also plotted is a best-fit line for the simulated versus data plot and ± 2 standard deviations from the best fit line. As this plot shows, the data is close to being within the 2 sigma error bars.

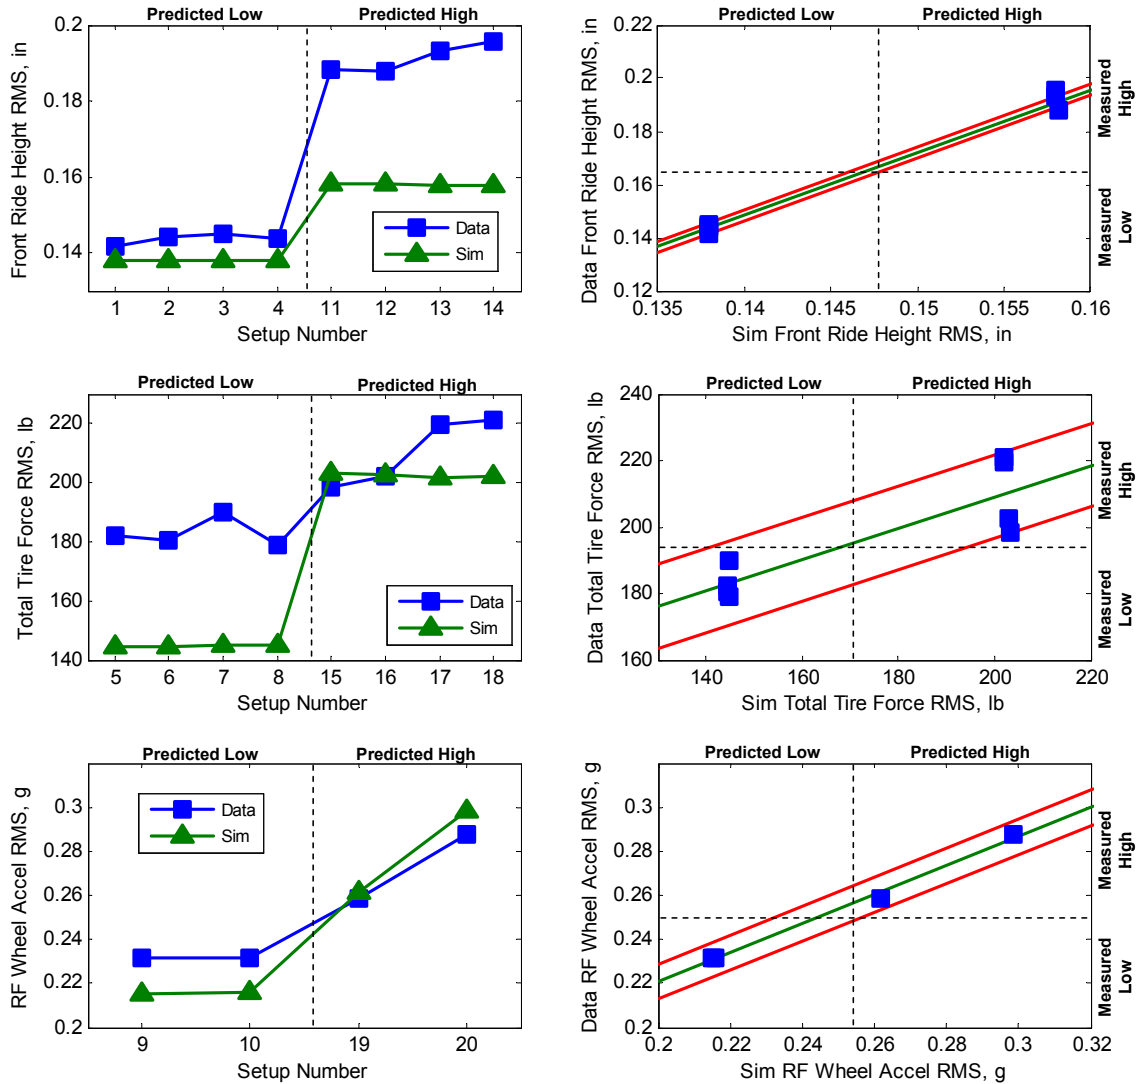


Figure 7.7. Comparison of simulated and measured shock trending

For the total tire force, the absolute trends also do not match well, but the relative trends do match. The data also lies within the ± 2 standard deviation error bars.

The right front rear acceleration data appears to match the simulation well, both in absolute and in relative terms. The data also falls within the 2 sigma error bars.

7.4 Summary

This chapter has applied the vehicle and shock models developed in this research to find the best and worst setups in terms of three different RMS measures. A database of 7 different shock builds was constructed to provide a range of dynamic responses. All 2401 possible shock setups were simulated using the models developed in this research, and the trends in the RMS response levels were observed for a select group of signals. Twenty setups were then selected for 8-post rig testing, with each having either a high or low response level of one of the signals predicted by the simulation. Comparison of the trends predicted by the simulation and measured on the 8-post rig showed that the response levels often did not match on an absolute basis, but showed good agreement on a relative scale. The simulation error was found to be bounded near or within ± 2 standard deviations of the baseline response levels.

Chapter 8

Conclusions

This chapter summarizes this research and presents conclusions based on that research. It also describes recommendations for future research that builds on this research or complementary research that would improve the application of this research.

8.1 Summary and Conclusions

This research has developed a method that uses experimental data to identify a linear vehicle model that predicts the effect of actuator inputs and shock force on sensor measurements during an 8-post rig test. This model can be coupled with suspension component models, such as linear or nonlinear shock absorber models, to simulate how the vehicle response changes with shock selection. The trends observed in simulation as components are changed can be used to increase productivity during 8-post rig testing. Some ways the simulation results may be applied to improve 8-post rig testing efficiency are: suggesting or eliminating setups to test, identifying tradeoffs in vehicle performance, identifying sensitivity of setup parameters, and identifying setup parameters which have limited influence on the response.

This approach differs from previous vehicle modeling studies with component tuning in that it does not make any assumptions about the physics of the vehicle, other than assuming the relationship is linear between the shaker rig actuators, shock forces, and the vehicle sensors.

Chapter 3 developed a physically-motivated nonlinear dynamic shock absorber model that can be quickly fit to experimental data and implemented in simulation studies. This model is based on the understanding that the shock is dominantly a velocity-dependent device, with lag due to compressibility effects. Mechanically, this behavior was shown to

be similar to a damper in series and in parallel with a stiffening spring in series. The effective time constant of such a system can become very small at higher velocities, making numerical integration very slow. To avoid these numerical efficiency issues, a slightly different model structure was suggested. The model consists of an algebraic backbone, which is an algebraic function of velocity, and a nonlinear low-pass filter, which has been designed based on the observation that shocks often exhibit less hysteresis at higher velocities. Due to the simplicity of the model, it can be fitted with data and evaluated quickly, making it ideal for simulation studies.

After fitting several shock models to experimental data, the models were validated using additional data sets. This analysis showed that the RMS force error in the dynamic models was 45% to 75% less than the basic polynomial curve fit model, as shown in Figure 8.1.

Chapter 4 develops the methods needed to develop models of vehicles on the 8-post rig on a quarter-car scale. By developing our method on this smaller scale, it allowed us to analyze the method more fully. Chapter 4 develops the vehicle model identification method at a quarter-car scale. The method was first developed using simulated quarter-car data with a known model. This allowed comparisons between the ideal model and the identified model to verify the ID method was accurate. Comparison of the ideal FRF and the transfer function fit showed perfect FRF matching except at low frequency where the shock force excitation was low.

After developing the method on simulated quarter-car data, the method was applied to experimental data collected from a laboratory quarter-car rig. The FRF estimates had good coherence except where excitation levels were low. The identified model was then used to perform simulations for different shocks, which was then compared to experimental results. This analysis showed that the method could predict the quarter-car rig response trends.

Chapter 5 extends the method to full-vehicle scale in simulation. A seven degree of freedom vehicle state space model was constructed to generate simulated data sets for the

identification process. The identification process can then be judged acceptable if the identified model is similar to the known ideal model.

The frequency response function (FRF) of the state space model was then explored to provide insight prior to the identification process. This analysis showed the importance of both the aeroloaders and the wheelloaders to excite sensor response at low and high frequencies, respectively.

If the shock force is almost linearly dependent on the wheelloader and aerolader inputs, the typical FRF estimation process will fail. A modified FRF estimation process was described which uses test data from several different shock configurations to address the shock force dependence issue. This method provided FRF estimates, which matched the ideal FRF where the coherence was high. When the coherence was poor due to low shock force or output levels at low or high frequency, the FRF estimate was replaced by a residual fit with a linear phase and linear dB/decade roll-off. The FRF estimate was further smoothed by applying a filter, reducing FRF variations with a width of 0.6 Hz or less.

Once an acceptable FRF estimate was calculated, a parametric model was identified by fitting transfer functions to each input-output pair. Simulations were run for 5 different shock configurations and compared to simulation results from the ideal model, showing that the identified model could accurately reproduce the original response and predict trends in response metrics.

Chapter 6 develops and applies the vehicle identification method on to data collected from actual 8-post rig test to predict the outcome of future 8-post rig tests for different drivefile and shock selections. The coherence for the FRF estimates was generally high when the response levels were high, with. If the response level was determined to be low across all frequencies of interest, the corresponding transfer function was zeroed. If the response level was low over a range of frequencies, the poor FRF estimate was replaced with a low FRF estimate with linear phase and a linear dB/decade magnitude roll-off before fitting the transfer function.

The model was then validated by performing simulations for several different drivefiles and shock setups and comparing it to experimental data. This analysis showed us that the model generally does a good job of predicting the amplitude trends of experiments, with some amplitude and phase error. There was also nonlinearity not accurately captured by the model in tire force and hub accelerometer measurements.

Chapter 7 applies the vehicle and shock models developed in this research to find the best and worst setups in terms of three different RMS measures. A database of 7 different shock builds was constructed to provide a range of dynamic responses. All 2401 possible shock setups were simulated using the models developed in this research, and the trends in the RMS response levels were observed for a select group of signals. Twenty setups were then selected for 8-post rig testing, with each having either a high or low response level of one of the signals predicted by the simulation. Comparison of the trends predicted by the simulation and measured on the 8-post rig showed that the response levels often did not match on an absolute basis, but showed good agreement on a relative scale. The simulation error was found to be bounded near or within +/- 2 standard deviations of the baseline response levels, indicating that the simulated trends match the measured trends within expected run-to-run variation limits.

8.2 Recommendations for Future Research

Since there has been little research published in the field of dynamic shaker rig testing for vehicle dynamic applications, this research serves as a starting point. This research has been a successful first step in identifying dynamic system models from laboratory shaker rig test data, and it can be used as a starting point for future research. After first discussing direct improvements and extensions of this research, we will then discuss complementary research that will improve the value of this research path.

The first suggested improvement of this research is to improve the linear system identification method used to identify the linear model. This study uses 8-post rig test data from multiple setups, estimates the FRF, and fits single-input single-output transfer

functions to each input-output pair. While this method is easy to implement, it has several weaknesses.

First, since the transfer function for each input-output pair was fitted independently, this created problems fitting transfer functions where the coherence was poor. This model structure has 7-28 transfer function polynomial coefficients for each input-output pair, allowing it to approximate any FRF desired. This would be good if we always had good FRF estimates, but this creates issues when coherence is poor. In these cases, there is not enough information on the true FRF to accurately fit the model. This was addressed in this study by zeroing transfer functions or replacing the FRF with a low FRF magnitude. A better approach would be to use a model structure that has less freedom across all input-output pairs to approximate an arbitrary linear system, and use the reduced degrees of freedom to fit the model. Multiple input linear system identification methods such as Eigensystem Realization Algorithm (ERA) and subspace methods identify a linear state space model with consistent poles, which reduces the sensitivity of modeling error to poor coherence.

Another interesting observation from this research is that the model was often able to predict relative trends, but not absolute trends. The measured versus simulated response level plots often trended along a straight line, indicating the model may be biased by a scale factor and an offset. In this study, we performed FRF estimation, fitted transfer functions, and then performed simulation incorporating a nonlinear shock absorber model in a feedback loop. While this method should provide good results if there were no modeling error present in these building blocks of the simulation, this process may cause bias in the model due to the feedback dynamics when modeling error is present. One way method to improve modeling error is to use an identification method that selects model parameters to minimize the error between simulation results and data. This will distribute the error between the FRF building blocks to minimize output error, not try to minimize error for just one building block at a time. Due to the nonlinear nature of the shock model, this type of method would require nonlinear identification unless the shocks were linear.

Further consideration should be given to the effect of uncertainties in the input signals. The H_1 FRF estimation method we used minimizes expected modeling error assuming that the input signals are known and the outputs signals are noisy. In the case of the wheelloader and aeroloader inputs, this is the case since we define those signals. For the shock force inputs, however, the measured shock force may have significant noise, which will bias the FRF estimates. Since the assumption that the input measurements are noise-free may be poor for shock absorber force inputs, use of methods that account for input measurement noise, such as H_2 or H_α FRF estimation, should be considered.

Another class of model structures that should be considered for several reasons is physics-based model structures. While these model structures can be made very complex, a good starting point is the 7 degree of freedom model presented in Chapter 5, modified to include the increased coupling present in the rear solid-axle suspension. This model structure will provide a smaller set of parameters compared to generic linear model structures, making it easier to observe the influence of each parameter on modeling error. This model structure can also be used for vehicle parameter estimation, and the estimated parameters can be used in more complex simulations than experienced in the experimental data, such as full-vehicle simulations.

This research only verified that the model was successful for one specific vehicle within one vehicle class over a limited sampling of commonly-used drivefiles and shock setups. Further testing should be conducted to determine under what testing conditions we should expect this model to be effective and when it may break down. This work should also provide guidance for developing new methods that are successful when the present method is not.

Another extension of this research is to accommodate nonlinearities that may be present in 8-post rig testing, such as progressive spring rates, coil bind and separation, bump stops, tire separation, and suspension and tire friction. It is possible in theory to address progressive spring rates, coil bind and separation, and bump stops using the method developed in this research given the kinematic relationship between shock travel and spring travel and a component model. In practice this may be difficult due to the high

sensitivity to suspension position. It is also theoretically possible to apply our substructuring approach to address tire separation, but this will be difficult on the 8-post rig since there are 6 forces and moments acting as inputs to the tire contact patch that may influence sensor response without the sufficient richness of inputs required to perform identification. Suspension friction should be minimized for good performance, and tire friction should be minimized during the 8-post test to improve testing consistency.

One limitation of our shock setup selection method was that we were optimizing the shock setup over a discrete set of shock models. The advantage of this approach is that all shock models represent real shocks that can actually be put on a car. The first disadvantage is that all potential shocks must be built, dynamometer tested, and modeled before selection can begin. This may not be a huge disadvantage for race teams that have a large database of shocks data, but it does limit exploring new shock builds. Second, optimizing over a discrete search space is much more computationally expensive than searching over a continuous search space. Research to develop a continuous space of shock models should be considered for various applications. This would allow searching over a larger search space more efficiently. While the results of the optimization over a continuous shock database may not provide a setup consisting of four known shock builds, it will provide a target shock force-velocity curve, which could be used to aid shock building and design.

Some options for continuous shock parameterizations useful for optimization include

1. Physics-based shock models,
2. Interpolation of shock models for existing builds, and
3. Shape functions.

Physics-based models are useful since their parameters reflect physical properties that can be used to design shock internals, and the limitations in parameter values lead to practical limitations in shock behavior, but they are often computationally inefficient. If there is a mesh of shock models for existing builds, it may be possible to interpolate to provide an estimate of shock performance within the mesh. For example, we modeled an

AA/AA and a B/B build, which have a shim thickness of 0.004 and 0.008 inches respectively. Without further data, it would seem reasonable to approximate the shock force from a 0.006 inch A/A shock build as an average of the AA/AA and B/B shock model force. This type of model would provide the advantages of continuous parameterization, physical significance of each model, and numerical efficiency. The main limitation of this approach is the requirements for a mesh of shock model wide enough to find the best shocks but fine enough to accurately model performance. To explore trends with no data requirements, a shape function model is ideal. A shape function model is an algebraic function whose parameters control specific portions of the curve. Using a shape function model can be useful to explore trends, but the curves often do not reflect physical constraints of real shock absorbers, so they may lead to results that are not achievable with real shocks if the model structure and parameters are not properly constrained.

One critical aspect of 8-post rig testing that was not addressed in this research is how to apply 8-post rig testing to improve vehicle dynamics on the track. Several studies to determine how to apply rig data to tune track performance are needed. As discussed in Chapter 2, since the dynamics on the 8-post rig and on the track are not consistent, assuming that the sensor measurements on the 8-post rig and on the track will be the same may lead to erroneous conclusions. A study should be conducted to determine to what extent and under what conditions trends in rig sensor measurements are correlated with track measurements. Different 8-post rig drivefile generation methods should be used, which may include drivefile iteration for a baseline setup, drivefile iteration for multiple setups, terrain profile-based wheelloader excitation, and aerodynamic model-based aerolader force modification.

Also since yaw dynamics are not significant on the rig, it is challenging to optimize the true metric of race car performance – lap time. Work should be done to integrate the vehicle behavior observed on the 8-post rig to the additional dynamics observed on the 8-post rig to determine how changes on the 8-post rig influence track performance. This can be done using the 8-post rig as a hardware-in-the-loop test with the remaining

dynamics running on a real-time processor, or data collected from the 8-post can be used to perform system or parameter identification for use in simulation.

References

1. Kelly, J. Kowalczyk, H. Oral, H. Track Simulation and Vehicle Characterization with 7 Post Testing. SAE Technical Papers 2002-01-3307.
2. Menzer, J. Seven-post shakers stir up Cup Series' team hierarchy. NASCAR.com, June 2007.
3. Craig, R. Blades, E. Substructure System Identification: Reduced-Order Models. Dynamics and Control of Space Structures III, pp 365-379, 1996.
4. Su, T. Juang, J. Substructure System Identification and Synthesis. Journal of Guidance, Control, and Dynamics, vol 17, no 5, 1994.
5. Lee, D. Hwang, W. Parametric Optimization of Complex Systems Using a Multi-Domain FRF-Based Substructuring Method. Computers and Structures, vol 81, pp 2249-2257, 2003.
6. Sjovall, P. Abrahamsson, T. Substructure System Identification from Coupled System Test Data. Mechanical Systems and Signal Processing, vol 22, pp 15-33, 2008.
7. Kasprzak, J. Floyd, R. Use of Simulation to Tune Race Car Dampers. SAE Technical Papers 942504.
8. Warner, B. Rakheja, S. An Investigation of the Influence of High Performance Dampers on the Suspension Performance of a Quarter Vehicle. SAE Technical Papers 962552.
9. Penske Racing Shocks 7000 Series Technical Manual, 2007.
10. Lang, H. Automotive Dampers at High Stroking Frequency. Ph.D. Thesis, University of Michigan, 1977.

11. Reybrouck, K. A Non Linear Parametric Model of an Automotive Shock Absorber. Vehicle Suspension System Advancements, SAE SP-1031, February 1994.
12. Talbott, M. An Experimentally Validated Physical Model of a High-Performance Automotive Monotube Damper. Master's Thesis, Purdue University, 2002.
13. Emmons, S. Boggs, C. Ahmadian, M. Parametric Modeling of a Highly-Adjustable Race Damper. ASME International Mechanical Engineering Congress and Exposition, 2006.
14. Rao, M. Gruenberg, S. Torab, H. Measurement of Dynamic Properties of Shock Absorbers for NVH. SAE Technical Papers 1999-01-1840.
15. Kowalski, D. Rao, M. Blouh, J. Gruenberg, S. Dynamic Testing of Shock Absorbers under Non-Sinusoidal Conditions. Journal of Automobile Engineering, vol 216, pp 373-384, 2002.
16. Cafferty, S. Characterization of Automotive Dampers using Higher Order Frequency Response Functions. Journal of Automobile Engineering, vol 211, pp 181-203, 1997.
17. Yung, V. Wavelet Analysis of High-Frequency Damper Behavior. Journal of Automobile Engineering, vol 219, pp 977-988, 2005.
18. Kasteel, R. Cheng-guo, W. Lixin, Q. Jin-zhao, L. Wen-zhang, Z. A New Shock Absorber Model with an Application in Vehicle Dynamic Studies. SAE Technical Papers 2003-01-3411.
19. Zhou, W. Chew, C. Hong, G. Inverse Dynamics Control for Series Damper Actuator Based on MR Fluid Damper. IEEE/ASME International Conference on Advanced Intelligent Mechatronics, 2005.
20. Barber, A. Accurate Models for Complex Vehicle Components using Empirical Methods. SAE Technical Papers 2000-01-1625.

21. Leser, C. Renner, T. Salmon, D. Accurate Shock Absorber load Modeling in an All Terrain Vehicle Using Black Box Neural Network Techniques. SAE Technical Papers 2002-01-0581.
22. Savaresi, S. Silani, E. Montiglio, M. Stefanini, A. Previdi, F. A Comparative Analysis of High-Accuracy Black-Box and Grey-Box Models of MR-Dampers for Vehicle Control. SAE Technical Papers 2004-01-2066.
23. Miller, A. Testing a Formula SAE Racecar on a Seven-Poster Vehicle Dynamics Simulator. SAE Technical Papers 2002-01-3309.
24. Kowalczyk, H. Damper Tuning with the use of a Seven Post Shaker Rig. SAE Technical Papers 2002-01-0804.
25. Chan, B. Sandu, C. Development of a Virtual 7-Post Test Rig for Experimental Validation. ASME International Mechanical Engineering Congress and Exposition, 2006.
26. Vilela, D. Vehicle Dynamics Simulation Correlation with Field Maneuvers. SAE Technical Papers 2001-01-3799.
27. Dittman, K. Albright, F. Leser, C. Validation of Virtual Prototypes via a Virtual Test Laboratory. MTS White Paper.
28. Hu, H. Experimental Validation of a Half-Vehicle Suspension Model. SAE Technical Papers 931966.
29. Ballard, R. Overview of Emerging Technology for Laboratory Based Measurement of Vehicle Performance. SAE Technical Papers 921068.
30. Cheli, F. Sabbioni, E. A Dynamic Light-Duty Vehicle Model: Validation with Indoor and Outdoor Experimental Tests. Proceedings of the ASME 2007 IDETC/CIE 2007 DETC2007-34753.

31. Andersen, Erik. Multibody Dynamics Modeling and System Identification for a Quarter-Car Test Rig with McPherson Strut Suspension. Master's Thesis, Virginia Tech, May 2007.
32. Ziegenmeyer, J. Estimation of Disturbance Inputs to a Tire Coupled Quarter-Car Suspension Test Rig. Master's Thesis, Virginia Tech, May 2007.
33. Alasty, A. Ramezani, A. Genetic Algorithm-Based Parameter Identification of a Nonlinear Full-Vehicle Ride Model. SAE Technical Papers 2002-01-1583.
34. Juang, J. Applied System Identification. Prentice Hall, 1993.
35. Roehrig Engineering. 2K Electro-Magnetic Actuator Product Description. www.roehrig.com
36. Boggs, C. Borg, L. Ostanek, J. Ahmadian, M. Efficient Test Procedures for Characterizing MR Dampers. ASME International Mechanical Engineering Congress and Exposition, 2006.
37. Oppenheim, A. Discrete-Time Signal Processing. Prentice Hall, 1999.
38. Soderstrom, T. Ljung, L. Gustavsson, I. Identifiability Conditions for Linear Multivariable Systems Operating Under Feedback. IEEE Transactions on Automatic Control, vol 21, no 6, pp 837-840.
39. Ljung, L. System Identification: Theory for the User. Prentice Hall, 1999.
40. MATLAB R2006b Documentation.