

AFIT/GAE/ENY/95D-02

APPLICATION OF MIXED-NORM OPTIMAL
CONTROL TO A MULTI-OBJECTIVE
ACTIVE SUSPENSION PROBLEM

THESIS

John B. Allison
Captain, USAF

AFIT/GAE/ENY/95D-02

19960327 038

Approved for public release; distribution unlimited

The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U. S. Government.

AFIT/GAE/ENY/95D-02

APPLICATION OF MIXED-NORM OPTIMAL CONTROL TO A
MULTI-OBJECTIVE ACTIVE SUSPENSION PROBLEM

THESIS

Presented to the Faculty of the School of Engineering
of the Air Force Institute of Technology

Air University

In Partial Fulfillment of the
Requirements for the Degree of
Master of Science in Aeronautical Engineering

John B. Allison, BS

Captain, USAF

December, 1995

Approved for public release; distribution unlimited

Acknowledgements

Without my family and friends, I would have never completed this work. I owe them all a debt of gratitude for their never ending support. Mostly, I want to thank my parents for guiding me in the right direction and my sisters, Ann, Amy, and Jill for their encouragement and humor. It has been a long time since I've seen any of my family, and I look forward to once again being a part of their lives.

There is no doubt that I was lucky to have Dr. Brett Ridgely as my advisor. He constantly was willing to answer my stupid questions, and to point me in the right direction. I doubt this thesis turned out the way either of us imagined it would in the beginning. Even though the results were not positive, I believe I've learned more from my failures than I ever could if this problem was solved easily.

Here at Wright-Patterson, I have met some of the most wonderful friends anyone could ask for. I'd like to thank W.T. and Karen for being the closest things to parents I've had here in Ohio. I'd also like to thank Don, Steve, Connie, Jo Lynn, Sid, Mike, Joe, Charley, and Stephanie. It was a pleasure to work with all of you and I'm glad I can call you friends.

While at AFIT, I have made some great friends. We suffered through the classes together, and were always there to find some humor in the situation. Thanks to Lori for rewriting thesis.sty and being the comic relief for the rest of us in the lab. Thanks to Dr. Dave for showing me that the world really works in discrete time. Thanks to Mike and Jenn for being friends and letting me be a part of your lives (even though you use a Mac). Thanks to Dave, Brad, and Rich listening to me gripe, and offering me some hope that the end will come. Without Kris our computers and sense of humor would be down constantly. Thanks for keeping both of them up.

To relieve stress and to get back into shape, a few years ago I took up aerobics. I have met some of the most interesting and wonderfully funny people there. I owe a huge debt to Debby and Brad for showing me that exercise can be a fun part of my life. I'd also like to thank Krista, Chris, Steve, Tim, Tony, Liz, Dave, Jane, Kim, Paula, Lisa, Joe, Larry, Steph, Marge, Rob, Mike, Star, Susan, Sandy, Deb, James, Lori, Becky, and the rest of the aerobics crowd for giving me a forum to blow off some steam, and to let my funny side out.

A lot has happened in my life since I came to Ohio. The most important person in my life here was Christa. I want to thank her for encouraging me to go to AFIT and the support she gave me while here. Although we are not together now, I will always remember the time we spent together with a smile.

John B. Allison

Table of Contents

	Page
Acknowledgements	ii
List of Figures	vii
List of Tables	xiv
Abstract	xv
 I. Introduction	 1-1
1.1 Overview	1-1
1.2 Previous Work	1-1
1.3 Research Objectives	1-3
1.4 Approach	1-3
1.5 Outline	1-4
 II. Norm Based Optimization	 2-1
2.1 Introduction	2-1
2.2 Mathematical Preliminaries	2-1
2.2.1 State-Space Representation	2-1
2.2.2 Controllability	2-2
2.2.3 Observability	2-2
2.3 Transfer Function Norms	2-3
2.3.1 Two-Norm	2-3
2.3.2 Infinity-Norm	2-4
2.3.3 One-Norm	2-4
2.4 H_2 Optimization	2-5
2.5 H_∞ Optimization	2-8
 III. Mixed-Norm Optimization	 3-1
3.1 Introduction	3-1
3.2 Nomenclature	3-3
3.3 Problem Setup	3-4
3.4 Numerical Solution	3-6
3.5 Optimization Results	3-8

IV.	Suspension Problem	4-1
4.1	Introduction	4-1
4.2	Outline	4-1
4.3	System Models	4-2
4.3.1	Six DOF Model	4-3
4.3.2	Four DOF Model	4-3
4.3.3	Two DOF Model	4-3
4.4	System Inputs	4-4
4.4.1	Stochastic Road Noise	4-4
4.4.2	Deterministic Road Irregularities	4-5
4.5	Performance Specifications	4-6
4.6	Passive Dynamics	4-7
4.7	Active Open-Loop Dynamics	4-10
4.8	Controllability and Observability	4-11
4.8.1	Ten DOF Model	4-12
4.8.2	Six DOF Model	4-13
4.8.3	Four DOF Model	4-14
4.8.4	Two DOF Model	4-15
4.9	Original Approach	4-16
4.9.1	Problem One - Controllability	4-16
4.9.2	Problem Two - Digitizing the system	4-19
4.10	Modified Approach	4-24
V.	H_2 Optimization	5-1
5.1	Introduction	5-1
5.2	Problem Set-Up	5-1
5.3	H_2 Design 1	5-2
5.4	H_2 Design 2	5-7
5.5	Conclusions	5-14
VI.	H_∞ Optimization	6-1
6.1	Introduction	6-1
6.2	Problem Set-Up	6-1
6.3	H_∞ Design 1	6-1
6.4	H_∞ Design 2	6-6
6.5	Conclusions	6-13

VII.	Performance Based Mixed-Design	7-1
7.1	Introduction	7-1
7.2	$\ T_{ed1}\ _{\infty}$ - Tire Deflection Constraint	7-2
7.3	$\ T_{ed2}\ _{\infty}$ - Suspension Deflection Constraint	7-6
7.4	$\ T_{ed3}\ _{\infty}$ - Weighted Control Usage Constraint	7-6
7.5	Mixed-Norm Design 1	7-11
7.6	Conclusions	7-26
VIII.	Robustness Based Mixed-Design	8-1
8.1	Introduction	8-1
8.2	Robustness Constraint	8-3
8.3	$\ T_{ed1}\ _{\infty}$ - Tire Deflection Constraint	8-4
8.4	$\ T_{ed2}\ _{\infty}$ - Suspension Deflection Constraint	8-8
8.5	$\ T_{ed3}\ _{\infty}$ - Weighted Control Usage Constraint	8-11
8.6	$\ T_{ed4}\ _{\infty}$ - Robustness Weighting Constraint	8-15
8.7	Robust Design	8-17
8.8	Results	8-24
8.9	Conclusions	8-28
IX.	Conclusions and Recommendations	9-1
9.1	Introduction	9-1
9.2	Conclusions	9-1
9.2.1	Performance	9-1
9.2.2	Modeling	9-3
9.2.3	General	9-4
9.3	Recommendations	9-5
Appendix A.	Removing the $D_{yu} = 0$ Assumption	A-1
Appendix B.	Two DOF Model	B-1
Appendix C.	Ten DOF Passive Suspension Responses	C-1
Appendix D.	Two DOF Passive Suspension Responses	D-1
Appendix E.	Ten DOF Active Open-Loop Responses	E-1
Appendix F.	Two DOF Active Open-Loop Responses	F-1
Appendix G.	Mixed-Norm Design 2 - Empty Configuration Simulations	G-1

Appendix H.	Mixed-Norm Design 2 - Loaded Configuration Simulations	H-1
Appendix I.	Mixed-Norm Design 2 - Ten DOF Medium Configuration Simulations	I-1
Bibliography		BIB-1
Vita		VITA-1

List of Figures

Figure	Page
2.1 H_2 System	2-6
2.2 H_∞ System	2-8
2.3 H_∞ Parameterized System	2-9
2.4 Closed-Loop T_{ed} System with Uncertainty Δ	2-11
2.5 H_∞ System with Additive Uncertainty	2-12
3.1 H_2 MIMO System	3-2
3.2 Mixed-norm feedback system	3-4
3.3 Sample Edgeworth-Pareto Curve	3-9
3.4 Edgeworth-Pareto Surface	3-9
4.1 Ten DOF Truck Model	4-2
4.2 Two DOF Model	4-3
4.3 Road Power Spectral Density Comparison	4-5
4.4 Rounded Pulse Inputs	4-6
4.5 Passive System Eigenvalues (10 DOF)	4-8
4.6 Passive System Eigenvalues (2 DOF)	4-9
4.7 Open-Loop System Eigenvalues (10 DOF)	4-10
4.8 Open-Loop System Eigenvalues (2DOF)	4-11
4.9 Continuous (solid)/Discrete (dotted) Bode Magnitude Plot - Empty Configuration (46.5 Hz)	4-20
4.10 Continuous (solid)/Discrete (dotted) Bode Magnitude Plot - Loaded Configuration (46.5 Hz)	4-21
4.11 Continuous (solid)/Discrete (dotted) Bode Magnitude Plot - Empty Configuration (60 Hz)	4-21
4.12 Continuous (solid)/Discrete (dotted) Bode Magnitude Plot - Loaded Configuration (60 Hz)	4-22
4.13 Continuous (solid)/Discrete (dotted) Bode Magnitude Plot - Empty Configuration (93.1 Hz)	4-22
4.14 Continuous (solid)/Discrete (dotted) Bode Magnitude Plot - Loaded Configuration (93.1 Hz)	4-23
5.1 H_2 Problem Set-Up	5-2
5.2 H_2 Design 1 - Loaded Configuration - Noise Input	5-4
5.3 H_2 Design 1 - Loaded Configuration - Tiny Rounded Pulse Input	5-4

5.4	H_2 Design 1 - Loaded Configuration - Small Rounded Pulse Input	5-5
5.5	H_2 Design 1 - Loaded Configuration - Medium Rounded Pulse Input	5-5
5.6	H_2 Design 1 - Loaded Configuration - Large Rounded Pulse Input	5-6
5.7	H_2 Design 1 - Loaded Configuration - Huge Rounded Pulse Input	5-6
5.8	H_2 Design 2 - Empty Configuration - Noise Input	5-8
5.9	H_2 Design 2 - Empty Configuration - Tiny Rounded Pulse Input	5-8
5.10	H_2 Design 2 - Empty Configuration - Small Rounded Pulse Input	5-9
5.11	H_2 Design 2 - Empty Configuration - Medium Rounded Pulse Input	5-9
5.12	H_2 Design 2 - Empty Configuration - Large Rounded Pulse Input	5-10
5.13	H_2 Design 2 - Empty Configuration - Huge Rounded Pulse Input	5-10
5.14	H_2 Design 2 - Loaded Configuration - Noise Input	5-11
5.15	H_2 Design 2 - Loaded Configuration - Tiny Rounded Pulse Input	5-11
5.16	H_2 Design 2 - Loaded Configuration - Small Rounded Pulse Input	5-12
5.17	H_2 Design 2 - Loaded Configuration - Medium Rounded Pulse Input	5-12
5.18	H_2 Design 2 - Loaded Configuration - Large Rounded Pulse Input	5-13
5.19	H_2 Design 2 - Loaded Configuration - Huge Rounded Pulse Input	5-13
6.1	H_∞ Problem Set-Up	6-2
6.2	H_∞ Design 1 - Loaded Configuration - Noise Input	6-3
6.3	H_∞ Design 1 - Loaded Configuration - Tiny Rounded Pulse Input	6-4
6.4	H_∞ Design 1 - Loaded Configuration - Small Rounded Pulse Input	6-4
6.5	H_∞ Design 1 - Loaded Configuration - Medium Rounded Pulse Input	6-5
6.6	H_∞ Design 1 - Loaded Configuration - Large Rounded Pulse Input	6-5
6.7	H_∞ Design 1 - Loaded Configuration - Huge Rounded Pulse Input	6-6
6.8	H_∞ Design 2 - Empty Configuration - Noise Input	6-7
6.9	H_∞ Design 2 - Empty Configuration - Tiny Rounded Pulse Input	6-8
6.10	H_∞ Design 2 - Empty Configuration - Small Rounded Pulse Input	6-8
6.11	H_∞ Design 2 - Empty Configuration - Medium Rounded Pulse Input	6-9
6.12	H_∞ Design 2 - Empty Configuration - Large Rounded Pulse Input	6-9
6.13	H_∞ Design 2 - Empty Configuration - Huge Rounded Pulse Input	6-10
6.14	H_∞ Design 2 - Loaded Configuration - Noise Input	6-10
6.15	H_∞ Design 2 - Loaded Configuration - Tiny Rounded Pulse Input	6-11
6.16	H_∞ Design 2 - Loaded Configuration - Small Rounded Pulse Input	6-11
6.17	H_∞ Design 2 - Loaded Configuration - Medium Rounded Pulse Input	6-12
6.18	H_∞ Design 2 - Loaded Configuration - Large Rounded Pulse Input	6-12
6.19	H_∞ Design 2 - Loaded Configuration - Huge Rounded Pulse Input	6-13
7.1	Edgeworth-Pareto Optimal Curve for Tire Deflection	7-3

7.2	Edgeworth-Pareto Optimal Curve for Tire Deflection, Expanded	7-4
7.3	Singular Value Plot of $\ T_{ed_1}\ _\infty$ - Tire Deflection	7-4
7.4	$\ T_{ed_2}\ _\infty$ versus $\ T_{ed_1}\ _\infty$	7-5
7.5	$\ T_{ed_3}\ _\infty$ versus $\ T_{ed_1}\ _\infty$	7-5
7.6	Edgeworth-Pareto Optimal Curve for Suspension Deflection	7-6
7.7	Edgeworth-Pareto Optimal Curve for Suspension Deflection, Expanded . .	7-7
7.8	Singular Value Plot of $\ T_{ed_2}\ _\infty$ - Suspension Deflection	7-7
7.9	$\ T_{ed_1}\ _\infty$ versus $\ T_{ed_2}\ _\infty$	7-8
7.10	$\ T_{ed_3}\ _\infty$ versus $\ T_{ed_2}\ _\infty$, Expanded	7-8
7.11	Edgeworth-Pareto Optimal Curve for Weighted Control Usage	7-9
7.12	Edgeworth-Pareto Optimal Curve for Weighted Control Usage, Expanded .	7-10
7.13	Singular Value Plot of $\ T_{ed_3}\ _\infty$ - Weighted Control Usage	7-10
7.14	$\ T_{ed_1}\ _\infty$ versus $\ T_{ed_3}\ _\infty$	7-11
7.15	$\ T_{ed_2}\ _\infty$ versus $\ T_{ed_3}\ _\infty$	7-11
7.16	$\ T_{ed_2}\ _\infty$ versus $\ T_{ed_3}\ _\infty$, Expanded	7-12
7.17	Mixed-Norm Design 1 - Kinit Response - Noise Input	7-13
7.18	Mixed-Norm Design 1 - Kinit Response - Tiny Rounded Pulse Input	7-13
7.19	Mixed-Norm Design 1 - Kinit Response - Small Rounded Pulse Input . . .	7-14
7.20	Mixed-Norm Design 1 - Kinit Response - Medium Rounded Pulse Input . .	7-14
7.21	Mixed-Norm Design 1 - Kinit Response - Large Rounded Pulse Input . . .	7-15
7.22	Mixed-Norm Design 1 - Kinit Response - Huge Rounded Pulse Input	7-15
7.23	Mixed-Norm Design 1 - EP Curve	7-17
7.24	Mixed-Norm Design 1 - Singular Values Plot	7-17
7.25	Mixed-Norm Design 1 - $\ T_{ed_2}\ _\infty$ versus $\ T_{ed_1}\ _\infty$	7-18
7.26	Mixed-Norm Design 1 - $\ T_{ed_3}\ _\infty$ versus $\ T_{ed_1}\ _\infty$	7-18
7.27	Mixed-Norm Design Candidates - Huge Rounded Pulse Response - (Dotted - Initial, Dashed - A, Solid - Final)	7-19
7.28	Mixed-Norm Design Candidates - Large Rounded Pulse - Suspension Responses	7-20
7.29	Mixed-Norm Design Candidates - Large Rounded Pulse - Tire Responses .	7-21
7.30	Mixed-Norm Design Candidates - Medium Rounded Pulse Control Usage .	7-22
7.31	Mixed-Norm Design 1 Controller - Noise Input	7-23
7.32	Mixed-Norm Design 1 Controller - Tiny Rounded Pulse Input	7-24
7.33	Mixed-Norm Design 1 Controller - Small Rounded Pulse Input	7-24
7.34	Mixed-Norm Design 1 Controller - Medium Rounded Pulse Input	7-25
7.35	Mixed-Norm Design 1 Controller - Large Rounded Pulse Input	7-25
7.36	Mixed-Norm Design 1 Controller - Huge Rounded Pulse Input	7-26
8.1	H_∞ System with Additive Uncertainty	8-3

8.2	Robustness Weighting Functions	8-3
8.3	Robustness Subproblem	8-4
8.4	Edgeworth-Pareto Optimal Curve for Tire Deflection	8-5
8.5	Edgeworth-Pareto Optimal Curve for Tire Deflection, Expanded	8-5
8.6	Singular Value Plot of $\ T_{ed1}\ _{\infty}$ - Tire Deflection	8-6
8.7	$\ T_{ed2}\ _{\infty}$ versus $\ T_{ed1}\ _{\infty}$	8-6
8.8	$\ T_{ed3}\ _{\infty}$ versus $\ T_{ed1}\ _{\infty}$	8-7
8.9	$\ T_{ed4}\ _{\infty}$ versus $\ T_{ed1}\ _{\infty}$	8-7
8.10	Edgeworth-Pareto Optimal Curve for Suspension Deflection	8-8
8.11	Edgeworth-Pareto Optimal Curve for Suspension Deflection, Expanded	8-9
8.12	Singular Value Plot of $\ T_{ed2}\ _{\infty}$ - Suspension Deflection	8-9
8.13	$\ T_{ed1}\ _{\infty}$ versus $\ T_{ed2}\ _{\infty}$	8-10
8.14	$\ T_{ed3}\ _{\infty}$ versus $\ T_{ed2}\ _{\infty}$	8-10
8.15	$\ T_{ed4}\ _{\infty}$ versus $\ T_{ed2}\ _{\infty}$	8-11
8.16	Edgeworth-Pareto Optimal Curve for Weighted Control Usage	8-12
8.17	Edgeworth-Pareto Optimal Curve for Weighted Control Usage, Expanded	8-12
8.18	Singular Value Plot of $\ T_{ed3}\ _{\infty}$ - Weighted Control Usage	8-13
8.19	$\ T_{ed1}\ _{\infty}$ versus $\ T_{ed3}\ _{\infty}$	8-13
8.20	$\ T_{ed2}\ _{\infty}$ versus $\ T_{ed3}\ _{\infty}$	8-14
8.21	$\ T_{ed4}\ _{\infty}$ versus $\ T_{ed3}\ _{\infty}$	8-14
8.22	Singular Value Plot of T_{ed4} for K_{2opt} with MATLAB Generated Points	8-15
8.23	Singular Value Plot of T_{ed4} for K_{2opt} with 8000 Additional Points	8-15
8.24	Singular Value Plot of T_{ed4} for new K_{init}	8-16
8.25	Mixed-Norm Design 2 - Kinit Response - Noise Input	8-18
8.26	Mixed-Norm Design 2 - Kinit Response - Tiny Rounded Pulse Input	8-18
8.27	Mixed-Norm Design 2 - Kinit Response - Small Rounded Pulse Input	8-19
8.28	Mixed-Norm Design 2 - Kinit Response - Medium Rounded Pulse Input	8-19
8.29	Mixed-Norm Design 2 - Kinit Response - Large Rounded Pulse Input	8-20
8.30	Mixed-Norm Design 2 - Kinit Response - Huge Rounded Pulse Input	8-20
8.31	Edgeworth-Pareto Optimal Curve for Tire Deflection	8-22
8.32	Edgeworth-Pareto Optimal Curve for Suspension Deflection	8-23
8.33	Response of $\ T_{ed3}\ _{\infty}$ to Multiple Constraints Enforced	8-23
8.34	Edgeworth-Pareto Optimal Curve for Robustness Weighting	8-24
8.35	Mixed-Design 2 - Medium Configuration - Noise Input	8-25
8.36	Mixed-Design 2 - Medium Configuration - Tiny Rounded Pulse Input	8-26
8.37	Mixed-Design 2 - Medium Configuration - Small Rounded Pulse Input	8-26
8.38	Mixed-Design 2 - Medium Configuration - Medium Rounded Pulse Input	8-27

8.39	Mixed-Design 2 - Medium Configuration - Large Rounded Pulse Input . . .	8-27
8.40	Mixed-Design 2 - Medium Configuration - Huge Rounded Pulse Input . . .	8-28
B.1	Two DOF Model	B-1
C.1	Ten DOF Passive - Empty Configuration - Noise Input	C-1
C.2	Ten DOF Passive - Empty Configuration - Tiny Rounded Pulse Input . . .	C-2
C.3	Ten DOF Passive - Empty Configuration - Small Rounded Pulse Input . . .	C-2
C.4	Ten DOF Passive - Empty Configuration - Medium Rounded Pulse Input .	C-3
C.5	Ten DOF Passive - Empty Configuration - Large Rounded Pulse Input . . .	C-3
C.6	Ten DOF Passive - Empty Configuration - Huge Rounded Pulse Input . . .	C-4
C.7	Ten DOF Passive - Loaded Configuration - Noise Input	C-4
C.8	Ten DOF Passive - Loaded Configuration - Tiny Rounded Pulse Input . . .	C-5
C.9	Ten DOF Passive - Loaded Configuration - Small Rounded Pulse Input . .	C-5
C.10	Ten DOF Passive - Loaded Configuration - Medium Rounded Pulse Input .	C-6
C.11	Ten DOF Passive - Loaded Configuration - Large Rounded Pulse Input . .	C-6
C.12	Ten DOF Passive - Loaded Configuration - Huge Rounded Pulse Input . .	C-7
D.1	Two DOF Passive - Empty Configuration - Noise Input	D-1
D.2	Two DOF Passive - Empty Configuration - Tiny Rounded Pulse Input . . .	D-2
D.3	Two DOF Passive - Empty Configuration - Small Rounded Pulse Input . .	D-2
D.4	Two DOF Passive - Empty Configuration - Medium Rounded Pulse Input .	D-3
D.5	Two DOF Passive - Empty Configuration - Large Rounded Pulse Input . .	D-3
D.6	Two DOF Passive - Empty Configuration - Huge Rounded Pulse Input . .	D-4
D.7	Two DOF Passive - Loaded Configuration - Noise Input	D-4
D.8	Two DOF Passive - Loaded Configuration - Tiny Rounded Pulse Input . .	D-5
D.9	Two DOF Passive - Loaded Configuration - Small Rounded Pulse Input . .	D-5
D.10	Two DOF Passive - Loaded Configuration - Medium Rounded Pulse Input	D-6
D.11	Two DOF Passive - Loaded Configuration - Large Rounded Pulse Input . .	D-6
D.12	Two DOF Passive - Loaded Configuration - Huge Rounded Pulse Input . .	D-7
E.1	Ten DOF Open-Loop - Empty Configuration - Noise Input	E-1
E.2	Ten DOF Open-Loop - Empty Configuration - Tiny Rounded Pulse Input .	E-2
E.3	Ten DOF Open-Loop - Empty Configuration - Small Rounded Pulse Input	E-2
E.4	Ten DOF Open-Loop - Empty Configuration - Medium Rounded Pulse Input	E-3
E.5	Ten DOF Open-Loop - Empty Configuration - Large Rounded Pulse Input	E-3
E.6	Ten DOF Open-Loop - Empty Configuration - Huge Rounded Pulse Input .	E-4
E.7	Ten DOF Open-Loop - Loaded Configuration - Noise Input	E-4
E.8	Ten DOF Open-Loop - Loaded Configuration - Tiny Rounded Pulse Input .	E-5

E.9	Ten DOF Open-Loop - Loaded Configuration - Small Rounded Pulse Input	E-5
E.10	Ten DOF Open-Loop - Loaded Configuration - Medium Rounded Pulse Input	E-6
E.11	Ten DOF Open-Loop - Loaded Configuration - Large Rounded Pulse Input	E-6
E.12	Ten DOF Open-Loop - Loaded Configuration - Huge Rounded Pulse Input	E-7
F.1	Two DOF Open-Loop - Empty Configuration - Noise Input	F-1
F.2	Two DOF Open-Loop - Empty Configuration - Tiny Rounded Pulse Input	F-2
F.3	Two DOF Open-Loop - Empty Configuration - Small Rounded Pulse Input	F-2
F.4	Two DOF Open-Loop - Empty Configuration - Medium Rounded Pulse Input	F-3
F.5	Two DOF Open-Loop - Empty Configuration - Large Rounded Pulse Input	F-3
F.6	Two DOF Open-Loop - Empty Configuration - Huge Rounded Pulse Input	F-4
F.7	Two DOF Open-Loop - Loaded Configuration - Noise Input	F-4
F.8	Two DOF Open-Loop - Loaded Configuration - Tiny Rounded Pulse Input	F-5
F.9	Two DOF Open-Loop - Loaded Configuration - Small Rounded Pulse Input	F-5
F.10	Two DOF Open-Loop - Loaded Configuration - Medium Rounded Pulse Input	F-6
F.11	Two DOF Open-Loop - Loaded Configuration - Large Rounded Pulse Input	F-6
F.12	Two DOF Open-Loop - Loaded Configuration - Huge Rounded Pulse Input	F-7
G.1	Mixed-Design 2 - Empty Configuration - Noise Input	G-1
G.2	Mixed-Design 2 - Empty Configuration - Tiny Rounded Pulse Input	G-2
G.3	Mixed-Design 2 - Empty Configuration - Small Rounded Pulse Input	G-2
G.4	Mixed-Design 2 - Empty Configuration - Medium Rounded Pulse Input . .	G-3
G.5	Mixed-Design 2 - Empty Configuration - Large Rounded Pulse Input	G-3
G.6	Mixed-Design 2 - Empty Configuration - Huge Rounded Pulse Input	G-4
H.1	Mixed-Design 2 - Loaded Configuration - Noise Input	H-1
H.2	Mixed-Design 2 - Loaded Configuration - Tiny Rounded Pulse Input	H-2
H.3	Mixed-Design 2 - Loaded Configuration - Small Rounded Pulse Input . . .	H-2
H.4	Mixed-Design 2 - Loaded Configuration - Medium Rounded Pulse Input . .	H-3
H.5	Mixed-Design 2 - Loaded Configuration - Large Rounded Pulse Input . . .	H-3
H.6	Mixed-Design 2 - Loaded Configuration - Huge Rounded Pulse Input	H-4
I.1	Mixed-Design 2 - Ten DOF Medium Configuration - Noise Input	I-1
I.2	Mixed-Design 2 - Ten DOF Medium Configuration - Tiny Rounded Pulse Input	I-2
I.3	Mixed-Design 2 - Ten DOF Medium Configuration - Small Rounded Pulse Input	I-2
I.4	Mixed-Design 2 - Ten DOF Medium Configuration - Medium Rounded Pulse Input	I-3

I.5	Mixed-Design 2 - Ten DOF Medium Configuration - Large Rounded Pulse	
	Input	I-3
I.6	Mixed-Design 2 - Ten DOF Medium Configuration - Huge Rounded Pulse	
	Input	I-4

List of Tables

Table	Page
3.1 Mixed-Norm Nomenclature	3-3
4.1 Road Parameters for Power Spectral Density	4-4
4.2 Rounded Pulse Parameters	4-5
4.3 Suspension Travel Limits	4-7
4.4 Deflection Limits for Tires	4-7
4.5 Ten DOF Controllability Tests	4-13
4.6 Ten DOF Observability Tests	4-14
4.7 Six DOF Controllability Tests	4-15
4.8 Six DOF Observability Tests	4-16
4.9 Four DOF Controllability Tests	4-17
4.10 Four DOF Observability Tests	4-17
4.11 Two DOF Controllability Tests	4-18
4.12 Two DOF Observability Tests	4-18
4.13 Magnitude of Discrete Two DOF Poles	4-23
5.1 H_2 Design 1 Results	5-3
5.2 H_2 Design 2 Results	5-7
6.1 H_∞ Design 1 Results	6-2
6.2 H_∞ Design 2 Results	6-7
7.1 Mixed-Norm Design 1 - Results	7-16
B.1 Two DOF Model Parameters	B-1

Abstract

Mixed-norm optimal control synthesis is used to solve a multi-objective suspension problem. The objective is to develop a controller for an active suspension system on-board a tractor-semitrailer vehicle. The problem is first approached by using H_2 and H_∞ optimization. It is shown that by combining both techniques into one mixed-norm optimization method, it is possible to exploit the strengths of each norm to provide superior performance. Two H_2/H_∞ designs were completed. One design concentrated on optimal performance at one design condition. The second design concentrated on providing the best performance possible at a medium load configuration, while maintaining robust stability at the extreme load configurations.

APPLICATION OF MIXED-NORM OPTIMAL CONTROL TO A MULTI-OBJECTIVE ACTIVE SUSPENSION PROBLEM

I. Introduction

1.1 Overview

In the past few years significant process has been made in control theory concerning the application of H_2 , H_∞ and ℓ_1 norm based control design. Each of these single-norm control designs methodologies optimizes the controller for a certain set of inputs and outputs appropriate to the norm. If the problem at hand falls into one these categories, then the application of a single-norm control law synthesis is appropriate. However, this is often not the case in real world applications.

Unfortunately, there is little the control engineer can do if the problem he faces doesn't neatly fall into one of the single-norm methods. He can either change the problem to fit one of the methods, use the method that comes closest to meeting his needs, or find a new method. Both of these methods compromise the performance of the design because the engineer must solve the problem within limitations imposed by using a single norm.

Current research into the application of using a mixed-norm approach has addressed these issues. If the problem can be broken down into a set of subproblems with each subproblem designed for a specific norm, it can be solved numerically. This is the approach taken in this thesis. The problem addressed is broken into a set of subproblems each designed to take advantage of a norm appropriate to the subproblem. Then the whole problem is solved numerically through the use of the Mixed Norm Toolbox (MXTOOLS) developed by the Air Force Institute of Technology (AFIT) [JRCS95].

1.2 Previous Work

Although the theory behind single-norm problems has been in place for a number of years, and analytical solutions have been identified, the same can not be said about mixed-norm

optimal control. A major step forward was accomplished by Walker [Wal94] when he was able to refine a method originally proposed by Ridgely [Rid92] to numerically solve a general H_2/H_∞ problem for a fixed order compensator. His method was superior to the earlier attempts because it required fewer restrictions on the open loop plants of the H_2 and H_∞ subproblems, P_{zw} and P_{ed} . He showed that the free order optimal controller is unique and that it was in general a non-rational H_2 function. This drove the necessity for a fixed order approach, since implementation of a non-rational H_2 function is currently impossible. Walker concentrated his efforts on the H_2/H_∞ problem; however, he did show that the results were valid for any set of convex constraints.

Spillman followed Walker by approaching the H_2/L_1 and H_2/ℓ_1 problems [Spi94]. Since the L_1 norm is a convex functional, Walker's approach could be used. Spillman derived the analytical gradients of the L_1 norm and the ℓ_1 norm with respect to the design variables (the compensator matrices). This dramatically improved the speed and accuracy over the previous central differencing algorithm that was being used.

Additional improvements to the solution of the H_2/H_∞ problem were made by Smith [Smi94]. She was able to improve the existing H_∞ gradient routines and allowed for the capture of multiple peaks on the singular value plot of the H_∞ constraint function. The addition of the multiple peaks routine allowed the optimization routine to treat each peak as a constraint. This reduces the chances that additional peaks will occur during the optimization process.

The transition into discrete time was accomplished by Jacques [Jac95]. He was able to show that Walker's work held for discrete time. This was a significant step forward since it allowed for the construction of mixed-norm discrete controllers. Since most controllers are implemented digitally his method allowed for the implementation of discrete controllers without any loss of performance from digitization. This did, however, require that a digital model of the plant be used for the optimization process, so some performance may be lost if the plant must be digitized.

The theory behind mixed-norm optimal control is well summarized by Ridgely and Walker [RW95]. They discuss only the continuous H_2/H_∞ problem and its solution. The basic theory holds for H_2/L_1 problems as well, since L_1 is a convex constraint function

[Wal94, Spi94]. The discrete equivalent and an example of mixed-norm optimal control is well summarized by Jacques, Ridgely, and Canfield [JRC95]. The numerical approach behind the mixed-norm solution is well documented by Canfield, Ridgely, and Smith [CRS95].

The result of all the aforementioned work on mixed-norm optimal control has been incorporated into MXTOOLS which runs on MATLABTM[MAT]. The MXTOOLS toolbox solves for the fixed-order optimal compensator for the mixed-norm problem where the objective function is an H_2 transfer function, and the constraint function is either an H_2 , H_∞ , or L_1 (ℓ_1 for discrete) transfer function. It also allows for multiple constraints, thereby allowing the $H_2/L_1/H_\infty$ or $H_2/\ell_1/H_\infty$ problem to be solved.

1.3 Research Objectives

The effort in mixed-norm control at AFIT resulted in MXTOOLS [Wal94, Jac95, Spi94, Smi94]. However, there are only a few examples that have been exercised through MXTOOLS. Most of the effort has been in the creation of the toolbox itself and most of the examples that have been run are academic in nature and simplicity. Now that MXTOOLS has evolved, more difficult examples are being run to fully explore the power of the mixed-norm approach.

The objective of this thesis is to take a complex problem and to attempt to solve it with MXTOOLS. The problem was presented by de Jager as a potential baseline problem for the control community [dJ95]. The problem is based on designing an active suspension system for a tractor semitrailer vehicle. Both single-norm and mixed-norm methods will be used to solve the problem.

1.4 Approach

The goal of this thesis is to develop an active control system that performs better than the current passive suspension system. The model allows for four configurations: passive-empty, passive-loaded, active-empty, and active-loaded. Since the dampers in the passive suspension system are replaced by the actuators, the dynamics of the system change dramatically between the passive and the active systems. There are also four models that can be used

in analysis. The most complex model has ten degrees of freedom (DOF), which results in a twenty state system. This model was considered to be the truth model and used to test the effectiveness of the controllers designed on the lower order models of the system. There are also six, four, and two DOF models, which were evaluated for possible use in the development of the compensator. These models are simplifications of the ten DOF model. The 2 DOF model was used in development of the controller to minimize the order of the compensator and to provide assurances in the area of controllability and observability.

The analysis examined the following aspects of the problem:

- Current passive suspension system
- Impact of introducing an actuator into the system
- Original design approach
- Modified design approach
- Single norm design approaches
 - H_2 design of active suspension system
 - H_∞ design of active suspension system
- Mixed-Norm design approaches
 - H_2/H_∞ performance based design
 - H_2/H_∞ robustness based design

1.5 *Outline*

This thesis is broken into nine chapters. Chapter 2 introduces H_2 and H_∞ optimization. The solution to the mixed-norm problem is discussed in Chapter 3. Chapter 4 introduces the suspension problem and highlights the differences between the passive and active suspensions. The H_2 solution is discussed in Chapter 5. Chapter 6 includes the H_∞ solution. The mixed-norm problems are formulated and solved in Chapters 7 and 8. Finally, conclusions and recommendations are covered in Chapter 9. There are eight appendices. Appendix A

describes how to remove the assumption on no direct feedforward terms in the H_2 subproblem. Appendix B outlines the two degree of freedom model. Appendices C and D are the ten and two DOF passive suspension responses, respectively. Appendices E and F are the open-loop suspension responses for the ten and two DOF models, respectively. Appendices G and H are the empty and loaded configuration simulations for Mixed-Norm Design 2. Finally, Appendix I is the ten DOF simulation for Mixed-Norm Design 2.

II. Norm Based Optimization

2.1 Introduction

The first part of this chapter introduces the mathematical background required to understand the application of norm based optimization. The second part is an introduction into single norm optimization including H_2 and H_∞ norm optimization.

2.2 Mathematical Preliminaries

2.2.1 State-Space Representation

Consider a continuous linear time-invariant (LTI) system governed by the following state-space representation

$$\begin{aligned}\dot{x}(t) &= Ax(t) + Bu(t) \\ y(t) &= Cx(t) + Du(t)\end{aligned}\tag{2.1}$$

This state space represents how the states, x , of the system change with time, t . The control input is u and the measured output is y . The state space can also be realized as a transfer function in the s -domain through the Laplace transform

$$G(s) \equiv C(sI - A)^{-1}B + D\tag{2.2}$$

The following shorthand notation for the above will also be used throughout this thesis to represent the system

$$G(s) \equiv C(sI - A)^{-1}B + D \equiv \left[\begin{array}{c|c} A & B \\ \hline C & D \end{array} \right]\tag{2.3}$$

This thesis is based primarily on continuous theory; however, it should be understood that there is a full set of equivalent expressions for discrete time, which can be used to develop the mathematical background for the discrete problem.

2.2.2 Controllability

Controllability of the system implies that each state can be moved to an arbitrary location with a finite amount of control power. There are many tests of controllability; one of the most common is to check the rank of the controllability matrix

$$Co = [B \ AB \ A^2B \ \dots \ A^{n-1}B] \quad (2.4)$$

If the rank(Co) $\neq$ the number of states, n , then the system is not completely controllable. This implies that no matter what the compensator does, it will have not be able to control the uncontrollable modes of the system. If the mode is not important and is well behaved, this may not be an issue, but if it is important to control or is unstable, the problem is ill-posed. Rectifying the problem may include adding more actuators to provide additional control power or changing the location of the current actuators. For a system that is almost uncontrollable, balancing the system may improve the numerical conditioning.

2.2.3 Observability

The dual of controllability is observability. Observability implies that for any $x(0)$, there is a finite $T > t$ such that $x(0)$ can be computed from the measurements $y(0), y(t), \dots, y(T)$. The most common test for observability is similar to the test of controllability, except the rank of the observability matrix

$$Ob = \begin{bmatrix} C \\ CA \\ CA^2 \\ \vdots \\ CA^{n-1} \end{bmatrix} \quad (2.5)$$

is checked. If the rank(Ob) $\neq$ the number of states, n , then the system is not completely observable. This implies that there are states that exist that one can not identify from the given measurements. To correct this problem it may be necessary to move the sensors or add additional sensors to the system.

2.3 Transfer Function Norms

In norm based optimization, there are generally three norms used. They are the two-norm, infinity-norm, and the one-norm. This section covers the meaning behind the norms, both from a mathematical and application point of view.

2.3.1 Two-Norm

The motivation behind H_2 optimization is that the two-norm provides a measure of how large the energy of an output can be when the system is driven by an unit intensity white Gaussian noise (u.i.w.G.n). Given a system $G(j\omega)$, the two norm is defined as

$$\langle G(j\omega), G(j\omega) \rangle = \|G(j\omega)\|_2^2 \equiv \frac{1}{2\pi} \int_{-\infty}^{\infty} \text{tr} [G^*(j\omega)G(j\omega)] d\omega \quad (2.6)$$

In order for the integral to converge, the transfer function $G(j\omega)$ is restricted to be stable and strictly proper. Equation 2.6 is not easy to solve directly, but there is a method using Lyapunov equations to solve for the two-norm [DFT92].

The two-norm can be computed by

$$\|G(j\omega)\|_2^2 = \text{tr}[B^T L_o B] = \text{tr}[C L_c C^T] \quad (2.7)$$

with L_o and L_c being the positive semidefinite solutions to their respective Lyapunov equations

$$0 = A L_c + L_c A^T + B B^T \quad (2.8)$$

$$0 = A^T L_o + L_o A + C^T C \quad (2.9)$$

This method provides an easy way to compute the two-norm of a transfer function. If the transfer function is unstable, the two-norm is the square root of the sum of the squares of the two-norms of the stable/anti-stable projections.

2.3.2 Infinity-Norm

The motivation behind the infinity-norm is based on its physical meaning. Given a system with a bounded energy input, the infinity-norm provides a measurement of how large the energy of the output can be. Mathematically, the infinity-norm is defined as

$$\|G(j\omega)\|_{\infty} \equiv \sup_{\omega} \bar{\sigma}[G(j\omega)] \quad (2.10)$$

Equation 2.10 is intrinsically based on a search over frequency ω . The result is based on how fine of a frequency search is used to find the maximum singular value. Searching over frequency is one method to find the infinity-norm. The second method is based on an iterative approach using the Hamiltonian matrix

$$H = \begin{bmatrix} A + BR^{-1}D^TC & BR^{-1}B^T \\ -C^T(I + DD^T)^{-1}C & -(A + BR^{-1}D^TC)^T \end{bmatrix} \quad (2.11)$$

where $R \equiv \gamma^2 I - D^T D$. To find the infinity-norm, γ is reduced until the eigenvalues of H are some ϵ distance from the imaginary axis. The accuracy of this calculation will depend on how small ϵ is chosen; the closer ϵ is to zero, the closer the value of γ is to the infinity-norm.

2.3.3 One-Norm

Unlike H_2 and H_{∞} , ℓ_1 is not based in the frequency domain. Instead, it is based on the time response of the system. The ℓ_1 norm is a measure of the maximum magnitude of the output of a system given a bounded magnitude input. This norm can be applied to systems where the peak magnitude of signals is critical, rather than energy.

The one-norm is not well handled in continuous time because it is based on an infinite integral over time. However, an approximation to the one-norm can be made if the system $G(s)$ is transformed by the z-transform into $G_d(z)$. This produces the discrete equivalent system to $G(s)$

$$G_d(z) = \left[\begin{array}{c|c} A_d & B_d \\ \hline C_d & D_d \end{array} \right] \quad (2.12)$$

The one-norm for a SISO system is computed by expanding the pulse response of the system, which is

$$G_d(k) = \begin{cases} D_d & \text{for } k = 0 \\ C_d A_d^{k-1} B_d & \text{for } k \geq 1 \end{cases} \quad (2.13)$$

For a finite one-norm the system must be stable so that the pulse response will decay as $k \rightarrow \infty$. At some truncation level, N , the remainder of the pulse response is negligible. This results in a finite sum in place of the infinite summation and thus the one-norm can be approximated directly. For a SISO transfer function the one-norm is

$$\|G_d\|_1 = \lim_{N \rightarrow \infty} \sum_{k=0}^N |C_d A_d^k B_d| + |D_d| \quad (2.14)$$

This can be expanded to the MIMO case, where the MIMO one-norm of a transfer function matrix is the maximum row sum of the individual SISO transfer function one-norms. For further explanation of the one-norm, see Spillman [Spi94] or Dahleh and Diaz-Bobillo [DDB95].

2.4 H_2 Optimization

The H_2 problem is shown in Figure 2.1, where the input w is u.i.w.G.n. and z is the output we wish to control. The goal of H_2 optimization is to minimize $\|z\|_2$ given a u.i.w.G.n. w , which is equivalent to

$$\inf_{K_{\text{stabilizing}}} \|T_{zw}\|_2 \equiv \alpha_o \quad (2.15)$$

The state space realization of P_2 is

$$P_2 = \left[\begin{array}{c|cc} A_2 & B_w & B_2 \\ \hline C_z & D_{zw} & D_{zu} \\ C_2 & D_{yw} & D_{yu} \end{array} \right] \quad (2.16)$$

The notation used is slightly modified from the normal H_2 notation, but the reason for the modification will become clear in the mixed-norm discussion.

In order to solve the H_2 problem, the following assumptions are made:

1. $D_{zw} = 0$
2. $D_{yu} = 0$




Figure 2.1: H_2 System

3. (A_2, B_2) is stabilizable and (C_2, A_2) is detectable
4. $D_{zu}^T D_{zu}$ and $D_{yw} D_{yw}^T$ have full rank
5. $\begin{bmatrix} A - j\omega I & B_2 \\ C_z & D_{zu} \end{bmatrix}$ has full column rank for all ω
6. $\begin{bmatrix} A - j\omega I & B_w \\ C_2 & D_{yw} \end{bmatrix}$ has full row rank for all ω

Assumption 1 is required to ensure a finite two-norm of the closed-loop system. This assumption can be lifted for special cases and the motivation for this is discussed at the end of this section. Assumption 2 is not required but simplifies the computations. It can be lifted through a shifting technique shown in Appendix A. Assumption 3 is necessary to ensure that a stabilizing compensator does exist. Assumption 4 requires the problem to be regular, thereby penalizing control and assuring that there are no perfect measurements. Lastly, Assumptions 5 and 6 are required to ensure the existence of stabilizing solutions to the algebraic Riccati equations (AREs) used to solve the H_2 problem.

The resulting H_2 optimal controller is then given as

$$K_{2_{opt}}(s) = \left[\frac{A_2 - K_f C_2 - B_2 K_c}{-K_c} \mid \frac{K_f}{0} \right] \quad (2.17)$$

where

$$K_c = B_2^T X_2 + D_{zu}^T C_z \quad (2.18)$$

$$K_f = Y_2 C_2^T + B_w D_{yw}^T \quad (2.19)$$

X_2 and Y_2 are the real, unique, symmetric, positive semidefinite solutions to the AREs

$$(A_2 - B_2 D_{zu}^T C_z)^T X_2 + X_2 (A_2 - B_2 D_{zu}^T C_z) - X_2 B_2 B_2^T X_2 + \hat{C}_z^T \hat{C}_z = 0 \quad (2.20)$$

where

$$\hat{C}_z = (I - D_{zu} D_{zu}^T) C_z \quad (2.21)$$

and

$$(A_2 - B_w D_{yw}^T C_2) Y_2 + Y_2 (A_2 - B_w D_{yw}^T C_2)^T - Y_2 C_2^T C_2 Y_2 + \hat{B}_w \hat{B}_w^T = 0 \quad (2.22)$$

where

$$\hat{B}_w = B_w (I - D_{yw}^T D_{yw}) \quad (2.23)$$

It should be noted that this solution assumes that $D_{zw} = 0$; for the most general solution, this is not required. If $D_{zw} \neq 0$, then the compensator will have a non-zero D_c term as well as some additional terms. The D_c term is fixed and given by

$$D_c = -D_{zu}^\dagger D_{zw} D_{yw}^\dagger \quad (2.24)$$

such that

$$D_{zw} - D_{zu} D_{zu}^\dagger D_{zw} D_{yw}^\dagger D_{yw} = 0 \quad (2.25)$$

where $(\cdot)^\dagger$ denotes pseudoinverse.

However, this restriction on D_{zw} is very strict, and in most cases requires $D_{zw} = 0$ which implies $D_c = 0$ [RW95]. This requirement does not exist for the discrete case, where the closed-loop transfer function T_{zw} may have a non-zero D term. Therefore, the continuous H_2 optimal controller is strictly causal, whereas the discrete H_2 optimal controller must be causal but not necessarily strictly causal.




Figure 2.2: H_∞ System

2.5 H_∞ Optimization

The goal of H_∞ optimization is to minimize the maximum energy of the output e to a bounded energy input d . This is equivalent to

$$\inf_{K_{\text{stabilizing}}} \|T_{ed}\|_\infty \equiv \gamma_o \quad (2.26)$$

The H_∞ problem is shown in Figure 2.2. The state space representation of the H_∞ problem is

$$P_\infty = \left[\begin{array}{c|cc} A_\infty & B_d & B_\infty \\ \hline C_e & D_{ed} & D_{eu} \\ C_\infty & D_{yd} & D_{yu} \end{array} \right] \quad (2.27)$$

and is very similar to the set-up of the H_2 problem. The following assumptions are made for H_∞ optimization:

1. $D_{ed} = 0$
2. $D_{yu} = 0$
3. (A_∞, B_∞) is stabilizable and (C_∞, A_∞) is detectable
4. $D_{eu}^T D_{eu}$ and $D_{yd} D_{yd}^T$ full rank
5. $\begin{bmatrix} A_\infty - j\omega I & B_\infty \\ C_e & D_{eu} \end{bmatrix}$ has full column rank for all ω
6. $\begin{bmatrix} A_\infty - j\omega I & B_d \\ C_\infty & D_{yd} \end{bmatrix}$ has full row rank for all ω

Assumptions 1 and 2 are not necessary but make the derivation simpler. A description of the shifting technique required to remove these assumptions can be found in [Dai90]. Assumption 3 is required to ensure that a stabilizing compensator does exist. Assumption 4 ensures that all of the controls are penalized and that there are no perfect measurements. This assumption is strengthened through scaling such that $D_{eu}^T D_{eu} = I$ and $D_{yd} D_{yd}^T = I$. Lastly, Assumptions 5 and 6 are required to ensure that there is a solution to the Ricatti equations for the H_∞ norm calculations.

To find the H_∞ controller, an iterative process is used. The method is based on solution of two AREs and a coupling condition. The method parameterizes all sub-optimal controllers such that

$$\|T_{ed}\|_\infty < \gamma, \quad \gamma > \gamma_o \quad (2.28)$$

where γ_o is the actual infinity-norm of T_{ed} .




Figure 2.3: H_∞ Parameterized System

The H_∞ system with the parameterized controller is shown in Figure 2.3. The parameterization is accomplished through J where

$$J(s) = \begin{bmatrix} J_{yu} & J_{ur} \\ J_{vy} & J_{vr} \end{bmatrix} = \left[\begin{array}{c|cc} A_j & K_f & K_{fl} \\ \hline -K_c & 0 & I \\ K_{cl} & I & 0 \end{array} \right] \quad (2.29)$$

and

$$A_j = A_\infty - K_f C_\infty - B_\infty K_c + \gamma^{-2} Y_\infty C_e^T (C_e - D_{eu} K_c) \quad (2.30)$$

$$K_c = (B_\infty X_\infty + D_{eu}^T C_e)(I - \gamma^{-2} Y_\infty X_\infty)^{-1} \quad (2.31)$$

$$K_f = Y_\infty C_\infty^T + B_d D_{yd}^T \quad (2.32)$$

$$K_{cl} = -(\gamma^{-2} D_{yd} B_d^T X_\infty + C_\infty)(I - \gamma^{-2} Y_\infty X_\infty)^{-1} \quad (2.33)$$

$$K_{fl} = \gamma^{-2} Y_\infty C_e^T D_{eu} + B_\infty \quad (2.34)$$

X_∞ and Y_∞ are the solutions to the AREs

$$\begin{aligned} (A_\infty - B_\infty D_{eu}^T C_e)^T X_\infty + X_\infty (A_\infty - B_\infty D_{eu}^T C_e) \\ + X_\infty (\gamma^{-2} B_d B_d^T - B_\infty B_\infty^T) X_\infty + \hat{C}_e^T \hat{C}_e = 0 \end{aligned} \quad (2.35)$$

where

$$\hat{C}_e = (I - D_{eu} D_{eu}^T) C_e \quad (2.36)$$

and

$$\begin{aligned} (A_\infty - B_d D_{yd}^T C_\infty) Y_\infty + Y_\infty (A_\infty - B_d D_{yd}^T C_\infty)^T \\ + Y_\infty (\gamma^{-2} C_e^T C_e - C_\infty^T C_\infty) Y_\infty + \hat{B}_d \hat{B}_d^T = 0 \end{aligned} \quad (2.37)$$

where

$$\hat{B}_d = B_d (I - D_{yd}^T D_{yd}) \quad (2.38)$$

The controller can be formed by choosing any stable, proper Q such that

$$\|Q\|_\infty < \gamma \quad (2.39)$$

The parameterization of a controller K is valid if and only if the following three conditions hold:

1. $H_X \in \text{dom}(\text{Ric})$ with $X_\infty = \text{Ric}(H_X) \geq 0$
2. $H_Y \in \text{dom}(\text{Ric})$ with $Y_\infty = \text{Ric}(H_Y) \geq 0$

$$3. \rho(Y_\infty X_\infty) < \gamma^2$$

where

$$H_X = \begin{bmatrix} A_\infty - B_\infty D_{eu}^T C_e & \gamma^{-2} B_d B_d^T - B_\infty B_\infty^T \\ -\hat{C}_e^T \hat{C}_e & -(A_\infty - B_\infty D_{eu}^T C_e)^T \end{bmatrix} \quad (2.40)$$

$$H_Y = \begin{bmatrix} A_\infty - B_d D_{yd}^T C_\infty & \gamma^{-2} C_e^T C_e - C_\infty^T C_\infty \\ -\hat{B}_d^T & -(A_\infty - B_d D_{yd}^T C_\infty)^T \end{bmatrix} \quad (2.41)$$

are the Hamiltonians associated with Equations 2.35 and 2.37, respectively.

To solve for the controller, choose a γ and check the above conditions. Repeat the process while increasing γ until none of the conditions fail. Iterate between the last two points until the desired accuracy (within numerical accuracy) is satisfied.

One of the advantages H_∞ optimization has over H_2 optimization is that, due to the submultiplicative property of the H_∞ solution, it can be used to guarantee robust stability. If the closed-loop system T_{ed} is modeled with uncertainty as in Figure 2.4, then the Small Gain Theorem can be used.

Theorem 2.1 Small Gain Theorem

Let $T_{ed} \in H_\infty$. Assume Δ is connected to T_{ed} as in Figure 2.4. Then the closed-loop system is internally stable if

$$\|T_{ed}\Delta\|_\infty \leq \|T_{ed}\|_\infty \cdot \|\Delta\|_\infty < 1 \quad (2.42)$$

Proof: [Zam66]




Figure 2.4: Closed-Loop T_{ed} System with Uncertainty Δ

This implies that the lower $\|T_{ed}\|_\infty$ is made, the more uncertainty the system can handle before becoming unstable. An extension to Equation 2.42 is

$$\|T_{ed}\|_\infty < \frac{1}{\|\Delta\|_\infty} \quad (2.43)$$

If the system uncertainty can be modeled, then a weighting function W can be derived such that it is possible to provide robust stability against the uncertainty that is modeled. This does assume that it is possible to find a controller such that $\|T_{ed}\|_\infty \leq \frac{1}{\|\Delta\|_\infty}$, and if such a controller can not be found, it is not possible to guarantee robust stability.

In this thesis, an additive uncertainty model is used to guarantee robust stability between the empty and loaded configurations. Figure 2.5 shows how the additive uncertainty is incorporated into the system. P represents the nominal plant and $\tilde{P}$ is the off-nominal plant at which robust stability can be guaranteed for all $\|\Delta\|_\infty < 1$ if

$$\|WKS\|_\infty \leq 1 \quad (2.44)$$

where

$$W = \tilde{P} - P \quad (2.45)$$

and S is the sensitivity function, $(I - PK)^{-1}$. Further development of this technique is left for Chapter 8. There are other ways to model uncertainty, and for more information on them see [DFT92].




Figure 2.5: H_∞ System with Additive Uncertainty

This property of H_∞ optimization implies that H_∞ optimization can be used for both performance and to provide robust stability at off-nominal conditions. It should be noted that ℓ_1 optimization also has the submultiplicative property, and therefore can be used in a similar fashion to provide performance and robustness. The choice of which norm to use will depend on the problem and how it is best modeled.

III. Mixed-Norm Optimization

3.1 Introduction

In the previous chapter, H_2 and H_∞ norm based control synthesis methods were discussed. Both of these are designed to provide an optimal controller for a specific type of input and output. Unfortunately, most problems do not fit the mold of a single norm approach. This results in a compromise that must be made by the designer. Either the problem must be modified in some way to make it fit into one of the norms, choose the single-norm synthesis method that is closest to being appropriate for the problem, or find another method that works better. Either of the norm based methods are a compromise in design forced on the designer because of the limitations of the single-norm methods.

Recently, a synthesis method has been developed to design an output feedback controller which minimizes the two-norm of one transfer function, while constraining the one-norm, two-norm, and/or infinity-norm of possibly dissimilar transfer functions. This method does not guarantee global optimality, but has been shown to produce good results [Jac95, Spi94, Smi94]. This is the method used in this thesis, and which will be described in this chapter.

The method allows for a single H_2 objective function, and multiple constraints that could be H_2 , H_∞ , and/or ℓ_1 in nature. In essence, it provides a mixed-norm approach to the problem which allows for the exploitation of the best norm applicable to the constraint functions, with an H_2 function being the objective.

This approach allows the control engineer great insight into the problem. In the traditional single-norm approach, the problem was MIMO in nature because regularity conditions often added sensor noise and control usage as an additional input and output. Therefore, weightings on the inputs and the controlled outputs must be chosen. The norm was driven in large part by the ratio of the weightings, and there was no clear way to see what happened to the norm from just changing a single weighting, since multiple outputs affected the norm. This problem is removed in the mixed-norm approach because each output can be put into its own constraint function. This eliminates the iteration process in the choice of the ratio

of weightings and provides a clear picture of what happens to the two-norm of the objective function as one norm of a particular subproblem is constrained.

Another advantage is that with the mixed-norm approach the cross-weighting between objectives are eliminated. For example, take a simple MIMO H_2 problem with two u.i.w.G.n inputs and two controlled outputs as in Figure 3.1. For this example, the design goal is to minimize the two-norm of z_1 with respect to w_1 , and z_2 with respect to w_2 .




Figure 3.1: H_2 MIMO System

The closed-loop transfer function is T_{zw} , which can be partitioned into

$$T_{zw} = \begin{bmatrix} T_{z_1 w_1} & T_{z_1 w_2} \\ T_{z_2 w_1} & T_{z_2 w_2} \end{bmatrix} \quad (3.1)$$

Therefore, the two-norm of T_{zw} is influenced by the cross-term transfer functions $T_{z_1 w_2}$ and $T_{z_2 w_1}$. $T_{z_1 w_2}$ is the closed-loop transfer function from the second input to the first output, and $T_{z_2 w_1}$ is the closed-loop transfer function from the first input to the second output. However, these transfer functions are not a part of the design goal, yet they influence the optimization results. This implies that depending on the behavior of the cross-term transfer functions, the resulting controller may not be optimal for the design goals.

The influence of the cross-term transfer functions is eliminated in the mixed design because each transfer function can be separated into a subproblem. For this example, the mixed problem could be set up to reduce $\|T_{z_1 w_1}\|_2$ while constraining $\|T_{z_2 w_2}\|_2$ equal to or below some set value. In addition, the norm on the constraint function is not limited to the two-norm, so if an infinity-norm is a better norm for the constraint, an infinity-norm can be used.

This thesis deals primarily with the application of H_∞ constraints, but for completeness, this chapter will also include the notation required for the application of ℓ_1 constraints in conjunction with H_∞ constraints. To date, MXTOOLS has not been used with H_2 constraints.

3.2 Nomenclature

The notation used in the previous chapter was not standard for single norm optimization, but was the notation adapted for the mixed-norm method. This will become clearer as the mixed-norm synthesis method is discussed. Table 3.1 lists the nomenclature used in mixed-norm synthesis:

Symbol	Definition
T_{zw}	Objective H_2 norm transfer function
T_{mr}	Constraint ℓ_1 norm transfer function
T_{ed}	Constraint H_∞ norm transfer function
α_o	Minimum free order H_2 norm of T_{zw}
$\underline{\alpha}$	Minimum fixed order H_2 norm of T_{zw}
ν_o	Minimum free order ℓ_1 norm of T_{mr}
$\underline{\nu}$	Minimum fixed order ℓ_1 norm of T_{mr}
γ_o	Minimum free order H_∞ norm of T_{ed}
$\underline{\gamma}$	Minimum fixed order H_∞ norm of T_{ed}
K_{2opt}	The unique $K(s)$ that makes $\ T_{zw}\ _2 = \alpha_o$
$\bar{\nu}$	Value of $\ T_{mr}\ _1$ when $K(s) = K_{2opt}$
$\bar{\gamma}$	Value of $\ T_{ed}\ _\infty$ when $K(s) = K_{2opt}$
K_{mix}	A $K(s)$ which minimizes $\ T_{zw}\ _2$ while constraining $\ T_{mr}\ _1 \leq \nu$ and $\ T_{ed}\ _\infty \leq \gamma$
α^*	Value of $\ T_{zw}\ _2$ when $K(s) = K_{mix}$
ν^*	Value of $\ T_{mr}\ _1$ when $K(s) = K_{mix}$
γ^*	Value of $\ T_{ed}\ _\infty$ when $K(s) = K_{mix}$
κ	Vectorized compensator

Table 3.1: Mixed-Norm Nomenclature

3.3 Problem Setup

The mixed-norm problem is defined as

$$\inf_{K_{admissible}} \|T_{zw}\|_2 \quad (3.2)$$

subject to

$$\|T_{mr}\|_1 \leq \nu$$

$$\|T_{ed}\|_\infty \leq \gamma$$

where $K_{admissible}$ is the set of stabilizing controllers of some fixed order. The mixed-norm problem can have a single constraint transfer function, or multiple constraint transfer functions. Figure 3.2 shows the set-up of the mixed-norm problem.




Figure 3.2: Mixed-norm feedback system

P can be partitioned as

$$P = \begin{bmatrix} P_{zw} & P_{zr} & P_{zd} & P_{zu} \\ P_{mw} & P_{mr} & P_{md} & P_{mu} \\ P_{ew} & P_{er} & P_{ed} & P_{eu} \\ P_{yw} & P_{yr} & P_{yd} & P_{yu} \end{bmatrix} \quad (3.3)$$

For this general plant P , the inputs w , r , and d are as follows:

- w : unit intensity white Gaussian noise
- r : unknown but bounded magnitude input with $\|r\|_\infty \leq 1$

- d : unknown but bounded energy input with $\|d\|_2 \leq 1$

The outputs z , m , and r are all outputs whose norms are either minimized or constrained. Ignoring the cross transfer functions (they can be specified in the objective function or one of the constraint functions directly if necessary), the state space for the mixed problem is

$$\begin{aligned} x_2(t) &= A_2 x_2(t) + B_w w(t) + B_2 u(t) \\ z(t) &= C_z x_2(t) + D_{zw} w(t) + D_{zu} u(t) \\ y_2(t) &= C_2 x_2(t) + D_{yw} w(t) + D_{yu} u(t) \end{aligned} \quad (3.4)$$

$$\begin{aligned} x_1(t) &= A_1 x_1(t) + B_r r(t) + B_1 u(t) \\ m(t) &= C_m x_1(t) + D_{mr} r(t) + D_{mu} u(t) \\ y_1(t) &= C_1 x_1(t) + D_{yr} r(t) + D_{yu} u(t) \end{aligned} \quad (3.5)$$

$$\begin{aligned} x_\infty(t) &= A_\infty x_\infty(t) + B_d d(t) + B_\infty u(t) \\ e(t) &= C_e x_\infty(t) + D_{ed} d(t) + D_{eu} u(t) \\ y_\infty(t) &= C_\infty x_\infty(t) + D_{yd} d(t) + D_{yu} u(t) \end{aligned} \quad (3.6)$$

Equations 3.4 correspond to the H_2 objective transfer function. Equations 3.5 and 3.6 correspond to the ℓ_1 and H_∞ constraint transfer functions, respectively. The notation used is consistent with that from the previous chapter.

The objective transfer function and the constraint transfer functions may or may not have states in common. Since each subproblem is based on the same set of measurements, when $w = r = d = 0$, we have $y_2 = y_1 = y_\infty$. This does require D_{yu} to be the same for every subproblem. The following assumptions are made for the mixed-norm problem:

1. (A_2, B_2) stabilizable, (C_2, A_2) detectable
2. (A_1, B_1) stabilizable, (C_1, A_1) detectable
3. (A_∞, B_∞) stabilizable, (C_∞, A_∞) detectable
4. $D_{zu}^T D_{zu}$ full rank, $D_{yw} D_{yw}^T$ full rank
5. $\begin{bmatrix} A_2 - j\omega I & B_2 \\ C_z & D_{zu} \end{bmatrix}$ full column rank $\forall \omega$
6. $\begin{bmatrix} A_2 - j\omega I & B_w \\ C_2 & D_{yw} \end{bmatrix}$ full row rank $\forall \omega$

$$7. D_{yu} = 0$$

$$8. D_{zw} = 0$$

Assumptions 1-3 are required to ensure the existence of at least one stabilizing controller. As long as the weightings in the ℓ_1 and H_∞ problem are stable, it can be shown that the H_2 stabilizing controller also stabilizes the ℓ_1 and H_∞ subproblems [Wal94]. Assumptions 4 - 6 are required to ensure that the H_2 problem isn't singular. A singular H_2 problem is acceptable; however, since an initial stabilizing controller is required to start the optimization process, it is easier to have a regular H_2 problem that can be solved separately to provide an initial stabilizing controller, $K_{2_{opt}}$. When the H_2 problem isn't regular, the overall optimization problem must be regular, so that all controls are penalized and all measurements are corrupted by noise. Assumption 7 is not required but simplifies the derivation and can be lifted through a shifting technique shown in Appendix A. Lastly, Assumption 8 is required to ensure a finite $\|T_{zw}\|_2$. This can be lifted for special cases but version 1.0 of MXTOOLS does not account for these special cases.

3.4 Numerical Solution

The goal of the mixed-norm optimization process is to find a controller K_{mix} that will minimize $\|T_{zw}\|_2$ while meeting the constraints on $\|T_{mr}\|_1$ and $\|T_{ed}\|_\infty$. The state space representation of the desired controller is

$$K_{mix}(s) = \left[\begin{array}{c|c} A_c & B_c \\ \hline C_c & D_c \end{array} \right] \quad (3.7)$$

where the inputs into the controller are the measurements, y , and the output is the control signal, u . It should be noted that in the continuous case, D_c is not a design variable. Currently, MXTOOLS assumes $D_{zw} = 0$, and therefore $D_c = 0$. For the discrete case, D_{zw} is not assumed equal to zero, and thus D_c is an additional design variable used in the optimization.

The solution of the mixed-norm problem starts with an initial stabilizing controller, K_{init} . Most often it is the H_2 optimal controller for T_{zw} , $K_{2_{opt}}$. This sets the number of states of the controller to the same as the number of states in the H_2 objective subproblem.

The optimization routine is not limited to the size of $K_{2_{opt}}$, and can be run for both larger and smaller controllers. A larger K_{init} can be formulated by augmenting $K_{2_{opt}}$ with near pole/zero cancelations or roll-off poles. A smaller K_{init} can be found through various model reduction techniques. The only restriction is that this reduced K_{init} must be stabilizing. There is also no guarantee that there exists a K_{mix} derived from a reduced K_{init} which will be able to satisfy the same constraints that a higher order K_{mix} can.

MXTOOLS requires the user to provide K_{init} , and if it isn't provided, it will automatically use $K_{2_{opt}}$ as K_{init} . The next step in the solution process is to put K_{init} into modal form. This reduces the number of variables that must be solved for while preserving the characteristics of K_{init} . However, the current formulation doesn't allow for repeated eigenvalues in the compensator. This has not proven to be a significant restriction for the work performed up to this point [RW95, JRC95, Spi94, CRS95]. If necessary, MXTOOLS could be modified to use a block Jordan form which would allow for repeated eigenvalues. The next step in solving for K_{mix} is the vectorization of K_{init} . This puts all of the elements of K_{init} into a single vector, κ . The optimization is then done on κ by a Sequential Quadratic Programming (SQP) algorithm which incorporates the analytical gradients for the H_2 , ℓ_1 , and H_∞ norm calculations for stabilizing compensators and a central differencing algorithm for non-stabilizing compensators [CRS95].

The optimization is done on κ until the constraints provided by the user are met, or the optimization routine fails. The user can provide a single set of constraints or multiple cases of constraints for the optimization routine. If multiple cases are provided, the Hessian from each solution is passed to the next optimization process. This is one of the advantages of the SQP routine over the CONSTR routine native to the Optimization Toolbox of MATLAB.

There is no known limit to the size of κ and examples up to 100 variables have been successfully done [Jac95]. However, as the number of variables increase the computational time it takes to find the optimal solution increases as well. Implementation of high order controllers also poses additional problems. Therefore, it is desirable to optimize on the lowest order controller possible that will still provide adequate results. As both the ability to optimize over a larger number of variables and to implement higher order controllers improve, compensators closer to the free order optimal performance will be possible. For

details on how the SQP algorithm works within the framework of MXTOOLS, the reader is encouraged to read [CRS95].

3.5 Optimization Results

To simplify the discussion, assume that the problem has two infinity-norm constraints, γ_1 and γ_2 . This discussion can be expanded with only notational changes for more constraints or the use of a different norm.

If the constraints are set such that $\gamma_1 \geq \bar{\gamma}_1$, and $\gamma_2 \geq \bar{\gamma}_2$ the resulting optimal controller is $K_{mix} = K_{2_{opt}}$. Since $K_{2_{opt}}$ results in the lower bound on α , it is not possible to reduce α any further by relaxing the constraints. This results in the constraints being satisfied with $\gamma_1^* = \bar{\gamma}_1$ and $\gamma_2^* = \bar{\gamma}_2$. This is a rather uninteresting problem because it provides the same result as doing H_2 optimization on P_2 without any constraints.

If the constraints are set such that $\gamma_1 < \bar{\gamma}_1$ and $\gamma_2 < \bar{\gamma}_2$ and MXTOOLS can find a mixed-norm optimal compensator, at least one of the constraints will be satisfied with equality within numerical tolerance. This is due to the Kuhn-Tucker conditions at optimality [Wal94].

The result of MXTOOLS is a set of controllers that satisfy a series of constraints while minimizing the two-norm of the objective function. One of the outputs of MXTOOLS is an Edgeworth-Pareto (EP) curve that can be formed by reducing the norm on one constraint while letting the other constraint be fixed, unconstrained, or nonexistent. This curve shows the direct trade-off between improvements (reduction) in the constrained norm and the degradation (increase) of the objective two-norm. A sample EP curve is shown in Figure 3.3. The EP curve monotonically increases as the constraint is lowered.

With two infinity-norm constraints, an EP surface of controllers can be formed. This is accomplished by tracing the EP curve for one constraint while holding the other constant, then moving the constant constraint and tracing another EP curve. With enough controllers calculated, a surface is formed which shows how the two-norm of the objective function changes with both constraints. A portion of an EP surface developed by Jacques [Jac95] is shown in Figure 3.4. This surface was developed by finding approximately 400 mixed-norm controllers for a problem with two constraints; one H_∞ and one ℓ_1 .




Figure 3.3: Sample Edgeworth-Pareto Curve




Figure 3.4: Edgeworth-Pareto Surface

Creating a surface for the two constraints will provide insight into how the objective function changes; however, often that much information isn't necessary to solve the problem. One other approach is to reduce one constraint until it meets a certain goal. Then hold that constraint constant while reducing the other constraint to meet a different goal. This method provides a systematic way to meet both goals (if possible) while keeping the objective two-norm minimized. In many ways this is analogous to finding the correct weightings within a single-norm approach, but through MXTOOLS the weightings are found through a series of trade-offs and may include the norm best suited for each constraint.

IV. Suspension Problem

4.1 Introduction

The problem approached in this thesis was proposed by de Jager [dJ95] as a possible benchmark mixed-norm problem for the controls community. The problem addresses the challenge of placing an active suspension system on board a tractor semi-trailer vehicle. The goal is to minimize the accelerations on the driver and cargo while keeping suspension deflection and control usage within specifications. This chapter is designed to introduce the problem, and to identify the challenges faced in finding a solution.

4.2 Outline

This chapter is broken down into ten sections, with each section discussing one aspect of the problem. Following the outline in this second section, the third section discusses the models provided by de Jager. The development of the equations for the model is detailed in [dJ95]. The fourth section describes the inputs to the model. The fifth section covers the specifications placed on the performance of the system. The sixth section discusses the passive dynamics of the system. It is these passive dynamics that we wish to improve upon. The seventh section looks at the open-loop dynamics of the system after the actuator is included in the model.

Controllability and observability play a heavy role in the solution to this problem. In section eight, the controllability and observability of the open-loop models is discussed. Also, the improvements from balancing the system are shown.

The last two sections discuss the original approach taken for this problem and the rationale for changing that approach, resulting in the modified approach. As will be seen, the challenges posed by this problem unknown at the beginning of this research forced a significant departure from the original approach.




Figure 4.1: Ten DOF Truck Model

4.3 System Models

All of the system models are based on the ten DOF model as seen in Figure 4.1. The model was provided to de Jager by DAF Trucks Incorporated, located in the Netherlands. The default location of the actuator is in place of the damper located on the rear suspension. There are other possible actuator locations and multiple actuators are permissible. There are several sensor positions available on the ten DOF model. The four locations suggested by de Jager are sensors on the front and rear suspension for deflection measurements and accelerometers on the front and rear axle. For each model, the number of states is equal to twice the degrees of freedom. The states include the displacement and velocity for each degree of freedom. This results in a twentieth order plant matrix for the ten DOF model. Since most synthesis methods result in a compensator with an equal or higher number of states as the plant, using the ten DOF model for development of the compensator would result in a very high order controller. For implementation of the controller, the higher order controller would have to be reduced. As an alternative, the compensator could be developed using a lower order plant model or a reduced order controller would need to be designed for the higher order plant. There are three lower order models included by de Jager. They consist of a six, a four, and a two DOF model, discussed next.

4.3.1 Six DOF Model

The first reduction in complexity from the ten DOF model is the six DOF model. The six DOF model assumes that the motor, cabin, and chassis are all rigidly connected. This results in a twelfth order system. The same measurements from the ten DOF model are used.

4.3.2 Four DOF Model

The four DOF model is a further reduction from the six DOF model. This model assumes that the trailer is rigidly connected to the chassis and that the trailer suspension is rigid. The result is a single mass with two suspensions, the front and rear. The measurements are still the same as the ten DOF model. This results in an eighth order plant.

4.3.3 Two DOF Model

The two DOF model assumes that the front suspension is rigid, thereby reducing the system to a single truck mass over the rear suspension. This is the most simplistic representation of the system and reduces the plant to fourth order. The two DOF model is shown in Figure 4.2. Due to the loss of the front suspension, the measurements change for this model. The measurements are the suspension deflection, the acceleration of the axle, and the acceleration of the truck mass. For reasons discussed later in this chapter, the two DOF model was used as the basis for controller design. Appendix B outlines the basic parameters used in developing the two DOF model.




Figure 4.2: Two DOF Model

4.4 System Inputs

There are two types of inputs for vehicles traveling on a road. The first input is road noise. This is a stochastic process and is modeled as such. The second type of input represents bumps in the road and is deterministic in nature. This section discusses both types of inputs and how they are modeled in the simulation process of the suspension system.

4.4.1 Stochastic Road Noise

Road noise is modeled as a random process to represent the irregularities in the road from potholes and damaged roadway. This road model is used to describe two types of road, a fair roadway and a minor roadway. The fair roadway is considered typical of highways whereas the minor roadway is considered typical of back roads.

Both types of roads are describes by a power spectrum density function based on spatial frequency Ω . Noise seen by the vehicle due to irregular roads is made a function of time by assuming a constant forward velocity, v . This allows Ω to be written as a function of temporal frequency, ω , where $\omega = v\Omega$. The resulting equation for the power spectral density is

$$\Phi_{\omega}(\omega) = \frac{\Phi_{\Omega}^o v^{k-1}}{\omega^k} \quad (4.1)$$

Therefore, the road can be modeled by two parameters, Φ_{Ω}^o and k . Table 4.1 contains the parameters for both the fair roadway and the minor roadway.

Type of road	Φ_{Ω}^o [$m^2(rad/m)^{k-1}$]	k [-]
Fair roadway	1×10^{-6}	2.0
Minor roadway	18.8×10^{-6}	2.5

Table 4.1: Road Parameters for Power Spectral Density

In the analysis, a velocity of 25 meters per second was used. For simulation purposes, a white noise was filtered with a second order Butterworth filter such that it provided an overbound to both types of roads and could be used for comparison purposes. Figure 4.3 shows how the filtered noise is an overbound.




Figure 4.3: Road Power Spectral Density Comparison

4.4.2 Deterministic Road Irregularities

The second model of a road profile is deterministic. It consist of a rounded pulse input, with pulse intensities classified as: tiny, small, medium, large, and huge. This model is representative of bumps found in the road and is dependent on two parameters q_{max} and l_d/v . Table 4.2 contains the two parameters defining the different rounded pulses assuming a constant speed of 25 m/s.

Rounded Pulse	q_{max} [m]	$\frac{l_d}{v}$ [s]
Tiny	0.0695	0.0116
Small	0.0315	0.049
Medium	0.0909	0.237
Large	0.1216	0.500
Huge	0.1886	1.000

Table 4.2: Rounded Pulse Parameters

The height of the pulse is defined as q and can be written as

$$q(t) = q_{max} \left(\frac{e^2}{4} \right) \left(2\pi \frac{t}{\frac{l_d}{v}} \right)^2 e^{\left(-2\pi \frac{t}{\frac{l_d}{v}} \right)} \quad (4.2)$$




Figure 4.4: Rounded Pulse Inputs

Equation 4.2 is plotted in Figure 4.4 for the five sets of values in Table 4.2 for a constant speed of 25 m/s. This shows that the magnitude of the pulses and their duration vary significantly and provide a wide range of inputs to the system.

4.5 Performance Specifications

There are two types of specifications imposed on this problem. The first is limits on the suspension travel and tire deflection, and the second is limits based on International Standards Organization (ISO) 2631, which is based on frequency dependent acceleration limits for human exposure. The second set of limits was not taken into account in this thesis because it will be shown that the suspension travel limits are overly restrictive for this design and can not all be met through the methods used in this thesis. Since this is the case, the frequency based limits were not evaluated. For all of the limits the values are based on positive being upward (or expansion) and negative being downward (or compression), with zero being the static case.

There are five suspension systems in the ten DOF model. These include two for the cabin, two connected to the chassis, and one on the trailer. Table 4.3 outlines the limits on the travel of each suspension. The actuator is replacing the rear suspension; therefore,

Suspension	Working Space [m]	Extension [m]	Compression [m]
Front	0.23	0.14	-0.09
Rear	0.23	0.14	-0.09
Trailer	0.23	0.14	-0.09
Cabin Front	0.10	0.05	-0.05
Cabin Rear	0.10	0.05	-0.05

Table 4.3: Suspension Travel Limits

Tire	Working Space [m]	Extension [m]	Compression [m]
Front	0.103	0.034	-0.069
Rear	0.103	0.024	-0.079
Trailer	0.103	0.020	-0.083

Table 4.4: Deflection Limits for Tires

its allowable travel distance is the same as the rear suspension. In addition, the actuator is limited to a maximum force of 100kN.

The second set of deflection limits is based on the compression and extension that the tires can have while maintaining handling qualities. In the actual system, it is desirable that at all times at least one tire is on the ground. For modeling purposes, it is necessary to assume that all the tires are on the ground at all times because once any of them leave the ground, the linearity assumptions in the model are violated. Table 4.4 contains the limits imposed on the deflection of the tires.

With the specifications defined, it will be the goal of this thesis to meet these specifications while minimizing the acceleration on the driver and cargo.

4.6 *Passive Dynamics*

The current (pre-control) configuration is a passive suspension system without any actuation. The dynamics of the truck are dependent on the suspension configuration and the load configuration. Two sets of dynamics were provided by de Jager that included an empty and a loaded configuration. On the axles there are springs and shock absorbers that provide damping to the system. The poles of the system are shown in Figure 4.5 for both the loaded




Figure 4.5: Passive System Eigenvalues (10 DOF)

and empty configuration. From the location of the poles it can be seen that the system is lightly damped. It is also interesting to note how four poles of the empty configuration leave the real axis when the vehicle is loaded. This is understandable as the loaded configuration is going to respond more slowly and with previously overdamped modes oscillating under the increased mass.

In both configurations, there are underdamped poles. For the empty configuration there are two pairs of poles severely underdamped at 5.6% damping and 7.7% damping. The loaded configuration is worse, with the two pairs having 4.2% and 5.8% damping. The passive system has a problem with its current damping levels, as can be seen from the response of the system to the road noise and the rounded pulse inputs. These responses can be found in Appendix C.

In addition to having problems with damping, the system also has modes at significantly different frequencies. For the empty system the frequencies vary from 7.7 rad/s to 113 rad/s. The average frequency for the empty modes is 37 rad/s. For the loaded configuration, the frequencies vary from 6.75 rad/s to 74.5 rad/s, with the average being 32 rad/s. With such a wide spread of frequencies there will be a problem if the system is digitized and this will be addressed later in this chapter.




Figure 4.6: Passive System Eigenvalues (2 DOF)

Since the ten DOF model has both damping problems and widely spread frequencies, it is important to see how this compares to the two DOF model. The two DOF model is the most simplified model, and therefore only represents the major modes of motion. With the reduced order model, the numerics will improve because of the simplicity.

Figure 4.6 shows the pole locations for the empty and loaded configurations using the two DOF model. The damping of the system has improved, with the empty configuration's damping being 26% and 96%. The loaded configuration damping is also improved to a damping of 29% and 32%. Both configurations are significantly more damped than the ten DOF model. For the empty configuration the poles are located at 50 rad/s and 22 rad/s. This a significant reduction in the spread of the empty poles; however, for the loaded configuration the improvement is much less with the poles being at 7 rad/s and 59 rad/s. Open-loop simulation results with the two DOF passive model can be found in Appendix D.

It is important to realize that the passive system is the current configuration of the suspension system, and not the starting point for the controller design. The passive system can only be used as a benchmark for comparison on how well the active system performs. Unless it is possible to improve on the passive system's performance, there is no reason to implement the active control system.




Figure 4.7: Open-Loop System Eigenvalues (10 DOF)

4.7 Active Open-Loop Dynamics

To implement an active suspension system, the shock absorber on the rear suspension is replaced by the computer controlled actuator. In addition, the spring is also replaced by a lighter spring. The result is that the damping of the rear suspension reduced by 95% and the spring constant is reduced by 75%.

The impact is that the already underdamped system gets worse. This can be seen in the pole positions of the open-loop active suspension system. The ten DOF models poles are seen in Figure 4.7. The system has a similar pattern of underdamped poles and wide spread frequencies as the passive system. For the empty configuration there are four pairs of poles with damping less than 10%. The frequencies also vary dramatically from 9.1 rad/s to 113 rad/s. The loaded system also has four pairs of poles with damping less than 10% and frequencies from 3.6 rad/s to 75 rad/s. The performance of the system is very sensitive to the damping within the suspension system [KG93] and this can be seen from the open-loop simulations found in Appendix E.

To see if the underdamped modes and the wide spread frequencies resemble the major modes of motion, the two DOF model is examined. The poles can be seen in Figure 4.8. It is important to note the axis on this figure; the magnitude of the real part of the poles




Figure 4.8: Open-Loop System Eigenvalues (2DOF)

are less than unity with the imaginary part of one set of poles around sixty. For the empty configuration, the poles are damped at 1.5% and 8.9%. The frequencies of the poles are 9.1 and 58.5 rad/s. The loaded configurations damping is 1.4% and 3.5%, with the frequencies ranging from 3.5 to 58.5 rad/s. The two DOF open-loop response can be seen in Appendix F.

The result is that the open-loop system is severely underdamped and has wide spread frequencies. It is important to remember that this is the open-loop system, to be used as the starting point for control. It will be the goal of the active suspension system to overcome the deficiencies in the damping until the system performs better than the passive system. It will be seen later that the characteristics of the open-loop system will be a major factor in the approach taken to solve the problem.

4.8 Controllability and Observability

Before a control problem can be solved, it is important to determine that a solution is possible. For this problem, this involved determining the controllability and observability of the system.

If a system is barely controllable, the system can be balanced as one way to try to improve the controllability. This is a similarity transformation on the system that tries to

equalize the controllability and observability grammians. In essence, it tries to balance the observability and controllability of the system.

Controllability and observability analyses were performed on the models to determine the ability of the actuator to control the system, and the ability of the sensors to determine the states. The tests started with the controllability and observability test discussed earlier in Sections 2.2.2 and 2.2.3, and were expanded to include additional tests concerning the observability and controllability grammians. The tests included:

1. Rank of Controllability (Observability) matrix
2. Zero singular values of the Controllability (Observability) matrix
3. Rank of Controllability (Observability) grammian matrix
4. Orders of magnitude between minimum and maximum singular values of the Controllability (Observability) grammian

There are additional tests possible; however, these four tests should provide a good understanding as to the controllability and observability of the model.

4.8.1 Ten DOF Model

The results of the controllability and observability analysis are summarized in Table 4.5.

The conclusions are mixed. Some of the tests imply controllability, and others imply uncontrollability. This puts into serious doubt the controllability of the system with the single actuator. This implies that unless more actuators are added to the system such that the ten DOF system is clearly controllable, the ten DOF model should not be used in the development of the controllers. This is one motivation to base the controllers off a lower order model.

The same type of analysis is done for the observability of the system. Table 4.6 has the results from the observability tests. Like the controllability tests, the results were mixed as well. It is more likely that the system is observable than it is controllable, but in the design of an output feedback controller, both are required in order to do a good design when poles are in undesirable locations.

Configuration	Test	Results	Controllable?
Empty - Unbalanced	1	5, not full rank	No
	2	Yes	No
	3	20, full rank	Yes
	4	14 orders ($10^{-7} - 10^{-21}$)	Borderline
Empty - Balanced	1	5, not full rank	No
	2	Yes	No
	3	20, full rank	Yes
	4	9 orders ($10^{-2} - 10^{-11}$)	Yes
Loaded - Unbalanced	1	7, not full rank	No
	2	Yes	No
	3	20, full rank	Yes
	4	12 orders ($10^{-7} - 10^{-19}$)	Borderline
Loaded - Balanced	1	7, not full rank	No
	2	Yes	No
	3	20, full rank	Yes
	4	8 orders ($10^{-2} - 10^{-10}$)	Yes

Table 4.5: Ten DOF Controllability Tests

There is a conflict in the tests results for both controllability and observability. If there are any uncontrollable modes in the system, they should be able to be identified by either diagonalizing the system or checking $\text{rank}([\lambda_i I - A, B])$. Both of these were done and resulted in identifying no uncontrollable modes. This would imply that the system is controllable, but this is in direct conflict with the results from other tests. This may imply that for this system, the numerical algorithms within MATLAB that were used to test controllability and observability failed. Doing the same analysis for the lower order models may provide better insight into the controllability and observability of the system, and this is what is done next.

4.8.2 Six DOF Model

The follow-on to the tests of the ten DOF model is the same tests on the six DOF model. The results of the controllability tests are shown in Table 4.7. These are very similar to the results for the ten DOF model. The six DOF model is questionably controllable. The results of the observability analysis of the six DOF model are in Table 4.8.

Configuration	Test	Results	Observable?
Empty - Unbalanced	1	9, not full rank	No
	2	No	Yes
	3	20, full rank	Yes
	4	9 orders ($10^6 - 10^{-5}$)	Yes
Empty - Balanced	1	10, not full rank	No
	2	No	Yes
	3	20, full rank	Yes
	4	10 orders ($10^{-2} - 10^{-12}$)	Yes
Loaded - Unbalanced	1	10, not full rank	No
	2	No	Yes
	3	20, full rank	Yes
	4	9 orders ($10^6 - 10^{-5}$)	Yes
Loaded - Balanced	1	10, not full rank	No
	2	No	Yes
	3	20, full rank	Yes
	4	8 orders ($10^{-2} - 10^{-10}$)	Yes

Table 4.6: Ten DOF Observability Tests

The conclusion is that there are similar questions on the controllability and observability of the six DOF system as there are with the ten DOF model. Further tests again led to the belief that MATLAB's numerics were breaking down. This leads into looking at the controllability and observability of the four DOF model.

4.8.3 Four DOF Model

The same analysis that was performed on the ten and six DOF models was done on the four DOF model. The controllability results are shown in Table 4.9. The controllability results are encouraging, as the loaded configuration is clearly controllable. However, there is still some doubt regarding the empty configuration. The observability tests are shown in Table 4.10. Both the empty and loaded configurations are observable.

The conclusion for the four DOF controllability and observability test is that the loaded system is both controllable and observable. This may or may not be the case with the empty system since the results of the controllability tests are mixed. This leads us into the two DOF model.

Configuration	Test	Results	Controllable?
Empty - Unbalanced	1	6, not full rank	No
	2	Yes	No
	3	12, full rank	Yes
	4	12 orders ($10^{-7} - 10^{-19}$)	Borderline
Empty - Balanced	1	6, not full rank	No
	2	Yes	No
	3	12, full rank	Yes
	4	9 orders ($10^{-2} - 10^{-11}$)	Yes
Loaded - Unbalanced	1	6, not full rank	No
	2	Yes	No
	3	12, full rank	Yes
	4	12 orders ($10^{-7} - 10^{-19}$)	Borderline
Loaded - Balanced	1	6, not full rank	No
	2	Yes	No
	3	12, full rank	Yes
	4	8 orders ($10^{-2} - 10^{-10}$)	Yes

Table 4.7: Six DOF Controllability Tests

4.8.4 Two DOF Model

The two DOF model is the simplest of all the models, and therefore most likely to be controllable and observable due to its physical dynamics. The system is basically two masses with a spring and damper between them. When an actuator is placed between the two masses, it would be expected to have the system be controllable. The results of the controllability analysis can be found in Table 4.11. It is clear that the two DOF model is controllable. Balancing improves the conditioning of the system. The observability results are found in Table 4.12, implying observability.

The conclusion is that out of all of the models, the only model with certain controllability and observability is the two DOF model. The problems with the higher order models are believed to be due to MATLAB's inability to deal with large, widely separated in frequency systems. Because of this, the two DOF model was chosen as the baseline for controller development. This accomplished two tasks: first, it guaranteed that the system could be controlled, and second, it served to reduce the order of the compensator. Since the plant matrix is fourth order, as long as no dynamic weights are added, the resulting K_{2opt} compensator will be fourth order.

Configuration	Test	Results	Observable?
Empty - Unbalanced	1	7, not full rank	No
	2	No	Yes
	3	12, full rank	Yes
	4	11 orders ($10^6 - 10^{-5}$)	Yes
Empty - Balanced	1	10, not full rank	No
	2	No	Yes
	3	12, full rank	Yes
	4	8 orders ($10^{-2} - 10^{-10}$)	Yes
Loaded - Unbalanced	1	8, not full rank	No
	2	No	Yes
	3	12, full rank	Yes
	4	10 orders ($10^6 - 10^{-4}$)	Yes
Loaded - Balanced	1	9, not full rank	No
	2	No	Yes
	3	12, full rank	Yes
	4	8 orders ($10^{-2} - 10^{-10}$)	Yes

Table 4.8: Six DOF Observability Tests

4.9 Original Approach

The original approach to the problem was to digitize the plant and use $H_2/\ell_1/H_\infty$ mixed-norm optimization. The original approach broke the mixed-norm problem into one H_2 objective function, two ℓ_1 constraint functions, and one H_∞ constraint function. The H_2 objective function was the rms of the cabin acceleration in response to the road noise. The first ℓ_1 constraint was based on the suspension travel. The second ℓ_1 constraint was the tire deflection. The final constraint was an H_∞ constraint used to guarantee robustness between the empty and loaded configurations.

4.9.1 Problem One - Controllability

The first problem faced was the fact that none of the models besides the two DOF model are clearly controllable. The four DOF model is controllable in the loaded configuration, but unclear in the empty configuration. The mere fact that the two DOF model is controllable does not guarantee that a controller developed based on the two DOF model will work on the ten, six, or four DOF models. In fact, it raises serious questions about the ability to

Configuration	Test	Results	Controllable?
Empty - Unbalanced	1	6, not full rank	No
	2	No	Yes
	3	8, full rank	Yes
	4	13 orders ($10^{-7} - 10^{-20}$)	Borderline
Empty - Balanced	1	7, not full rank	No
	2	No	Yes
	3	8, full rank	Yes
	4	6 orders ($10^{-2} - 10^{-8}$)	Yes
Loaded - Unbalanced	1	8, full rank	Yes
	2	No	Yes
	3	8, full rank	Yes
	4	8 orders ($10^{-7} - 10^{-15}$)	Yes
Loaded - Balanced	1	8, full rank	Yes
	2	No	Yes
	3	8, full rank	Yes
	4	4 orders ($10^{-2} - 10^{-6}$)	Yes

Table 4.9: Four DOF Controllability Tests

Configuration	Test	Results	Observable?
Empty - Unbalanced	1	8, full rank	Yes
	2	No	Yes
	3	8, full rank	Yes
	4	7 orders ($10^6 - 10^{-1}$)	Yes
Empty - Balanced	1	8, full rank	Yes
	2	No	Yes
	3	8, full rank	Yes
	4	6 orders ($10^{-2} - 10^{-8}$)	Yes
Loaded - Unbalanced	1	8, full rank	Yes
	2	No	Yes
	3	8, full rank	Yes
	4	7 orders ($10^6 - 10^{-1}$)	Yes
Loaded - Balanced	1	8, full rank	Yes
	2	No	Yes
	3	8, full rank	Yes
	4	4 orders ($10^{-2} - 10^{-6}$)	Yes

Table 4.10: Four DOF Observability Tests

Configuration	Test	Results	Controllable?
Empty - Unbalanced	1	4, full rank	Yes
	2	No	Yes
	3	4, full rank	Yes
	4	4 orders ($10^{-7} - 10^{-11}$)	Yes
Empty - Balanced	1	4, full rank	Yes
	2	No	Yes
	3	4, full rank	Yes
	4	1 order ($10^{-2} - 10^{-3}$)	Yes
Loaded - Unbalanced	1	4, full rank	Yes
	2	No	Yes
	3	4, full rank	Yes
	4	4 orders ($10^{-7} - 10^{-11}$)	Yes
Loaded - Balanced	1	4, full rank	Yes
	2	No	Yes
	3	4, full rank	Yes
	4	2 orders ($10^{-2} - 10^{-4}$)	Yes

Table 4.11: Two DOF Controllability Tests

Configuration	Test	Results	Observable?
Empty - Unbalanced	1	4, full rank	Yes
	2	No	Yes
	3	4, full rank	Yes
	4	5 orders ($10^6 - 10^1$)	Yes
Empty - Balanced	1	4, full rank	Yes
	2	No	Yes
	3	4, full rank	Yes
	4	1 order ($10^{-2} - 10^{-3}$)	Yes
Loaded - Unbalanced	1	4, full rank	Yes
	2	No	Yes
	3	4, full rank	Yes
	4	5 orders ($10^6 - 10^1$)	Yes
Loaded - Balanced	1	4, full rank	Yes
	2	No	Yes
	3	4, full rank	Yes
	4	2 orders ($10^{-2} - 10^{-4}$)	Yes

Table 4.12: Two DOF Observability Tests

find a solution to this problem at all since the model that is closest to reality (the ten DOF) is marginally controllable.

This left two options. The first was to find a combination of actuators that would result in a controllable system. This is possible with the way the model is set up; however, the impacts of adding additional actuators are not explained by de Jager nor are the location of the actuators documented. The impact of adding just the one actuator included a reduction in the damping and spring constant of the rear suspension. No information was provided that accounted for how the other suspensions would be impacted by additional actuators.

The second option was to base the designs on the two DOF model, and to optimize about the specifications based on the two DOF model. This is the option that was selected. It accomplished the goal of reducing the order of the controller as well as providing a method to judge the different controllers. After the optimization was finished, the controllers would be tested with the ten DOF model to see if they were able to control the system.

4.9.2 Problem Two - Digitizing the system

For implementation, any controller found would need to be digitized and coded into a computer onboard the vehicle. There are two approaches to this: either design a discrete controller from the beginning or digitize a continuous controller. Since the discrete optimization allowed for more freedom for the optimization routine and the fact that ℓ_1 optimization is best accomplished through discrete systems, digitizing the plant was the preferred method.

There are various ways to digitize a continuous system, with each attempting to maintain the essential frequency information about the system. In most cases, this is accomplished by choosing a sample rate between five and ten times faster than the fastest mode of the system. For this problem, the sample rate is also limited in the specifications as

$$T_s \geq .002 + .00004n_c^2 \quad (4.3)$$

which for a fourth order controller results in $T_s \geq 0.0026$ seconds, or less than 378 Hz. For both of the configurations, the fastest modes of the two DOF model are around 58.5 rad/s. This implies that to capture both modes the sample frequency should range from 46.5 to




Figure 4.9: Continuous (solid)/Discrete (dotted) Bode Magnitude Plot - Empty Configuration (46.5 Hz)

93.1 Hz. This is only a rule of thumb, and the specific responses of the chosen sample frequency must be looked at.

Three sample frequencies were looked at; 46.5, 60, and 93.1 Hz to see how well the system can be digitized. The Bode magnitude plots comparing the continuous to discrete systems can be seen in Figures 4.9 - 4.14. It can be seen that the lower frequency response of the suspension deflection transfer function is dramatically degraded through digitization. Increasing the sampling frequency does improve the capture of the lower frequency pole of the empty configuration, but does not improve the lower frequency response of the suspension deflection. These were all digitized by a zero-order hold method, which was chosen because it is the same method used to solve the L_1 problem [Spi94].

The continuous system has poles near the imaginary axis. For discrete systems, the stability boundary is the unit circle, with poles outside of the unit circle being unstable, and poles within the unit circle being stable. Unfortunately, the digitization of the system causes the poles to become very close to the unit circle, and the higher the sampling frequency, the closer they come. A summary of how the poles moved due to digitization is found in Table 4.13.




Figure 4.10: Continuous (solid)/Discrete (dotted) Bode Magnitude Plot - Loaded Configuration (46.5 Hz)




Figure 4.11: Continuous (solid)/Discrete (dotted) Bode Magnitude Plot - Empty Configuration (60 Hz)




Figure 4.12: Continuous (solid)/Discrete (dotted) Bode Magnitude Plot - Loaded Configuration (60 Hz)




Figure 4.13: Continuous (solid)/Discrete (dotted) Bode Magnitude Plot - Empty Configuration (93.1 Hz)




Figure 4.14: Continuous (solid)/Discrete (dotted) Bode Magnitude Plot - Loaded Configuration (93.1 Hz)

Configuration	46.5 Hz	60 Hz	93.1 Hz
Empty	0.9813	0.9855	0.9906
	0.9826	0.9865	0.9913
Loaded	0.9821	0.9861	0.9910
	0.9973	0.9979	0.9987

Table 4.13: Magnitude of Discrete Two DOF Poles

With the poles being so close to the unit circle, the discrete H_2 solver failed to provide stabilizing solutions. It is unknown if the cause is from poor numerical conditioning within the H_2 solver, or if it is due to numerical problems within the Discrete Algebraic Ricatti Equation (DARE) solver. The conclusion was that the method chosen to proceed would have to be based on a continuous plant. This effectively eliminated the possibility of using the ℓ_1 constraints. One possibility would be to use L_1 constraints; however, since the solution to L_1 is based on digitization of the continuous plant and solving the discrete problem, it most likely would suffer from ill-conditioned numerics as well.

4.10 Modified Approach

The problems in the original approach caused a significant change in the direction of the research. No longer could the higher order models be used to synthesize controllers. No longer could the method be based on a discrete plant with ℓ_1 constraints.

The solution to this dilemma was to use the two DOF model and design continuous controllers. Therefore, the ℓ_1 constraints were replaced by H_∞ constraints. ℓ_1 constraints are more appropriate for this problem, since the ℓ_1 norm is based on maximum magnitude, and is most appropriate for constraints based on deflection limits. However, the problem may be solvable using H_∞ constraints. The H_∞ norm is based on the energy of the output and tries to minimize the average output over frequency. The H_∞ norm does not limit magnitude, but if the constraint norm is sufficiently restricted, it is hoped that the maximum magnitudes will be within limits.

A modified approach to the problem was developed in response to the problems with controllability and digitization of the system. The problem will first be solved through traditional H_2 and H_∞ methods. Then the problem will be put into the mixed-norm form, and solved. Two mixed-norm goals were developed. The first was to optimize performance at one design point. This would attempt to design a controller to get the best response possible with mixed-norm methods. The second mixed-norm design was based on providing the best response at a middle load configuration, while guaranteeing robust stability at the extremes. If the responses at the extreme load configurations are acceptable, it would indicate that no gain scheduling would be necessary. If not acceptable, it would indicate the contrary. The two DOF unbalanced model was used with the inputs scaled from Newtons to kiloNewtons. The scaling was done to improve the numerics of the problem. These tasks and their results are discussed in the following chapters.

V. H_2 Optimization

5.1 Introduction

This chapter covers the design, results, and conclusions of the H_2 single-norm synthesis method. Single-norm methods were used in order to provide a comparison to the mixed-norm results.

Two designs were accomplished using MATLAB's H2SYN command. Both designs were based on the two DOF model without balancing. The first design was based on achieving the best performance at the loaded configuration while maintaining actuator force limits. It will be shown that this design is destabilizing at the empty configuration. The second design was based on achieving the best response possible at the empty configuration, which also stabilizes the loaded configuration. This design was done without considering actuator or deflection limits in order to determine what the best possible response was, given no limitations on the system.

5.2 Problem Set-Up

Both designs were based on Figure 5.1, but differed in the weightings used. It is unusual for an H_2 design to have some of the indicated feedforward terms. D_{yw} is non-zero because the measurement of the axle acceleration is directly impacted by the road disturbance. D_{yu} and D_{zu} are non-zero because the actuator force directly impacts the acceleration of the vehicle and axle. For the first design, the output vector z , which we wish to minimize the two-norm of, was

$$z = \begin{bmatrix} \text{Suspension Deflection} \times 900 \\ \text{Truck Mass Acceleration} \times 1 \\ \text{Control Usage} \times 0.3 \end{bmatrix} \quad (5.1)$$

The output weightings were determined through repeated designs. The input weighting on the road noise (W_w) was 1, and the weightings on the measurement noise (W_n) were 0.01 on the deflection measurement, and 0.001 on the accelerations measurements. The actual measurement noises for the problem were not presented by de Jager and are unknown. The numbers used were chosen based on engineering judgment.

The second design used the same input weightings, but used the output vector

$$z = \begin{bmatrix} \text{Truck Mass Acceleration} \times 1 \\ \text{Control Usage} \times 1.0 \times 10^{-7} \end{bmatrix} \quad (5.2)$$

This design effectively penalizes only truck mass acceleration and lets the control usage be very large (but finite). It should be noted that neither of the designs minimized the tire deflection. This is one of the limitations of H_2 optimization. The tire deflection has a feedforward term resulting in a non-zero D_{zw} term. Therefore, in order to meet all of the assumptions necessary to do an H_2 design, the tire deflection was left out of the output vector.




Figure 5.1: H_2 Problem Set-Up

5.3 H_2 Design 1

The first design was based on the two DOF model in the loaded configuration. Since no dynamic weightings were used, the open-loop plant P_2 remained fourth order. This resulted in a fourth order K_{2opt} and an eighth order T_{zw} . The closed-loop poles and their damping are shown in Table 5.1.

Configuration	Pole	Damping
Loaded	$-0.018 \pm 0.018j$	71 %
	$-1.66 \pm 3.98j$	39 %
	$-7.67 \pm 14.2j$	48 %
	$-16.5 \pm 59.6j$	27 %
$\ T_{zw}(\text{Loaded})\ _2 = 8,712$		
Empty	185	(-)
	$1.64 \pm 62.1j$	(-)
	9.86	(-)
	$-3.03 \pm 2.82j$	77 %
	$0.018 \pm 0.018j$	71 %
$\ T_{zw}(\text{Empty})\ _2 = \infty$		

Table 5.1: H_2 Design 1 Results

Figures 5.2 - 5.7 are the loaded configuration responses. Please note that the acceleration plot in the figures are for the truck mass. The vector gain margin of the loaded configuration is $[-0.018, 0.0176]$ dB and the vector phase margin is $\pm 0.11^\circ$. The empty configuration system is not closed-loop stable using this controller.

The performance of this design is good, with only two specifications not met. The first is the tire deflection with the medium rounded pulse, and the second is the suspension deflection of the huge rounded pulse. Both are minor violations. These may have been avoided if the actuator force was allowed to exceed its limit of 100 kN. The actuator was at its limit in the medium rounded pulse case. This was the motivation for the next design.




Figure 5.2: H_2 Design 1 - Loaded Configuration - Noise Input




Figure 5.3: H_2 Design 1 - Loaded Configuration - Tiny Rounded Pulse Input




Figure 5.4: H_2 Design 1 - Loaded Configuration - Small Rounded Pulse Input




Figure 5.5: H_2 Design 1 - Loaded Configuration - Medium Rounded Pulse Input




Figure 5.6: H_2 Design 1 - Loaded Configuration - Large Rounded Pulse Input




Figure 5.7: H_2 Design 1 - Loaded Configuration - Huge Rounded Pulse Input

5.4 H_2 Design 2

The second H_2 design was based on the empty configuration with no (effective) limitations placed on control usage or suspension deflections. The resulting pole locations and damping information is found in Table 5.2. Figures 5.8 - 5.19 are simulations of both the empty and loaded configurations for Design 2. The vector gain margins for this design is $[-0.37, 0.37] \times 10^{-6}$ dB and the vector phase margin is 2.5×10^{-6} degrees.

Configuration	Pole	Damping
Empty	$-1.6 \times 10^{-7} \pm 57.7j$	$2.76 \times 10^{-7}\%$
	$-0.0025 \pm 0.0025j$	71 %
	$-0.018 \pm 0.018j$	71 %
	$-4.38 \pm 10.1j$	40 %
$\ T_{zw}(\text{Empty})\ _2 = 2.798$		
Loaded	$-1.54 \times 10^{-7} \pm 57.7j$	$2.67 \times 10^{-7}\%$
	$-0.0025 \pm 0.0025j$	71 %
	$-0.018 \pm 0.018j$	71 %
	$-0.673 \pm 4.26j$	16 %
$\ T_{zw}(\text{Loaded})\ _2 = 1.433$		

Table 5.2: H_2 Design 2 Results

This indicates a system that is on the verge of instability, which is evidenced by the pair of complex poles that are virtually on the imaginary axis. Interestingly enough, the loaded configuration is stable using this controller. This clearly indicates a problem in interpretation of stability margins - while sufficient to guarantee stability, they are not necessarily sufficient for any given perturbation.




Figure 5.8: H_2 Design 2 - Empty Configuration - Noise Input




Figure 5.9: H_2 Design 2 - Empty Configuration - Tiny Rounded Pulse Input




Figure 5.10: H_2 Design 2 - Empty Configuration - Small Rounded Pulse Input




Figure 5.11: H_2 Design 2 - Empty Configuration - Medium Rounded Pulse Input




Figure 5.12: H_2 Design 2 - Empty Configuration - Large Rounded Pulse Input




Figure 5.13: H_2 Design 2 - Empty Configuration - Huge Rounded Pulse Input




Figure 5.14: H_2 Design 2 - Loaded Configuration - Noise Input




Figure 5.15: H_2 Design 2 - Loaded Configuration - Tiny Rounded Pulse Input




Figure 5.16: H_2 Design 2 - Loaded Configuration - Small Rounded Pulse Input




Figure 5.17: H_2 Design 2 - Loaded Configuration - Medium Rounded Pulse Input




Figure 5.18: H_2 Design 2 - Loaded Configuration - Large Rounded Pulse Input




Figure 5.19: H_2 Design 2 - Loaded Configuration - Huge Rounded Pulse Input

5.5 Conclusions

From H_2 Design 2 we can determine that the minimum truck mass accelerations during the simulation vary from -0.05 to 0.05 m/s^2 for the empty configuration and from -0.01 to 0.01 m/s^2 for the loaded configuration. From H_2 Design 1, H_2 optimization is a good method of controlling the system. There is significant penalty in the accelerations by keeping the control usage and deflections within limits. One drawback is that the tire deflection can not be directly controlled. It turns out that the tire deflection is minimized by minimizing the suspension deflection. This is understandable since the suspension is connected directly to the tire.

The second drawback is the lack of robustness in the design. Whereas the empty configuration design is robust at the loaded configuration, the reverse is not true. This is not indicated by the margins of the system. The margins are larger for Design 1 than the margins for Design 2. Since the margins are based on an output disturbance, this would indicate that a better model of the influence of changing the load may provide better robustness information. It is intuitive that adding mass to the system is stabilizing, since adding mass will always slow down the system response. Removing mass will speed up the response and the actuator can then destabilize the system since it is expecting the system to respond slower. Unfortunately, since the H_2 single-norm method does not have the submultiplicative property, there is no way to guarantee stability at either extreme by adding another output to the output vector z . It will be seen later that the margins can be greatly improved by using the mixed-norm approach with a robustness constraint.

VI. H_∞ Optimization

6.1 Introduction

This chapter covers the design, results, and conclusions of two H_∞ optimization designs. The first was based on the loaded configuration with a goal of performance while meeting control usage specifications. The second design was based on the empty configuration with the same goal. This differs from the H_2 optimization objectives for the second design. In the H_2 optimization designs, the tire deflection could not be forced to meet the specification, and this was especially true in the empty configuration. That fact drove the objective for that design. In the H_∞ design, the restriction of no feedforward terms is lifted, thereby allowing direct control of the tire deflection through weighting.

6.2 Problem Set-Up

Both designs were based on Figure 6.1. The same input weightings were used for the road disturbance and the sensor noise as for the H_2 design. For the first design the output vector was

$$e = \begin{bmatrix} \text{Tire Deflection} \times 1.5 \\ \text{Suspension Deflection} \times 1,000 \\ \text{Truck Mass Acceleration} \times 1 \\ \text{Control Usage} \times 1.6 \end{bmatrix} \quad (6.1)$$

and for the second design the output vector was

$$e = \begin{bmatrix} \text{Tire Deflection} \times 0.4 \\ \text{Suspension Deflection} \times 10 \\ \text{Truck Mass Acceleration} \times (1 \times 10^{-4}) \\ \text{Control Usage} \times (1 \times 10^{-4}) \end{bmatrix} \quad (6.2)$$

6.3 H_∞ Design 1

The first design was based on the loaded configuration. The resulting poles and damping are seen in Table 6.1. The result is that the controller is able to keep the deflections within specifications with the exception of the tire deflection with the medium rounded pulse input. The tire deflection could have been brought within specifications; however, that would have




Figure 6.1: H_∞ Problem Set-Up

Configuration	Pole	Damping
Loaded	$-0.018 \pm 0.018j$	71 %
	$-1.66 \pm 3.98j$	39 %
	$-4.71 \pm 6.50j$	58 %
	$-16.2 \pm 56.6j$	28 %
$\ T_{ed}(\text{Loaded})\ _\infty = 7,780$		
Empty	31.3	(-)
	13.8	(-)
	$-0.018 \pm 0.018j$	71 %
	$-2.95 \pm 2.21j$	80 %
	$-8.30 \pm 67.3j$	12 %
$\ T_{ed}(\text{Empty})\ _\infty = 14,418$		

Table 6.1: H_∞ Design 1 Results

violated the control usage specifications. The controller is not robust and destabilizes the closed-loop system in the empty configuration. For the loaded configuration, the controller provides excellent damping, but at the cost of a lack of noise rejection compared to the passive suspension system and actuator deflections that may be quicker than possible to accomplish.

H_∞ Design 1 could have included an additional term in the output vector that could have provided improved margins. This would have been in the form of one of the robustness tests. However, since H_∞ optimization minimizes the entire output vector, the maximum singular value due to the robustness output may not be minimized if the other singular values have higher magnitude. The vector gain margins of H_∞ Design 1 are $[-0.54, 0.58]dB$ and the vector phase margin is $\pm 3.69^\circ$. Figures 6.2 - 6.7 show the response of the system at the loaded configuration.




Figure 6.2: H_∞ Design 1 - Loaded Configuration - Noise Input




Figure 6.3: H_∞ Design 1 - Loaded Configuration - Tiny Rounded Pulse Input




Figure 6.4: H_∞ Design 1 - Loaded Configuration - Small Rounded Pulse Input




Figure 6.5: H_{∞} Design 1 - Loaded Configuration - Medium Rounded Pulse Input




Figure 6.6: H_{∞} Design 1 - Loaded Configuration - Large Rounded Pulse Input




Figure 6.7: H_∞ Design 1 - Loaded Configuration - Huge Rounded Pulse Input

6.4 H_∞ Design 2

The second design was based on the empty configuration. The resulting poles and their damping are shown in Table 6.2. Figures 6.8 - 6.19 show the simulations of both the empty and loaded configurations. The limiting factor in the design is the actuator power on the small and medium pulse inputs. For the empty configuration, there is only one violation of the specifications and that is the tire deflection for the medium pulse. It should be noted that the H_∞ controller is very aggressive in order to meet the suspension and tire specifications. This results in extreme accelerations for all of the rounded pulse inputs that would most likely kill the driver and destroy the cargo. The vector gain margins for this design are $[-0.033, 0.033]$ dB and the vector phase margin is $\pm 0.217^\circ$. Again, these margins may have been improved by the inclusion of a robustness output. A robustness output will be used in the second mixed-norm design, and it will be shown to provide much improved margins with guaranteed robust stability, which can not be done as directly in the single-norm approach.

Configuration	Pole	Damping
Empty	$-0.018 \pm 0.018j$	71 %
	$-4.39 \pm 10.1j$	40 %
	$-18.4 \pm 37.0j$	45 %
	$-0.018 \pm 0.018j$	71 %
	$-234 \pm 238j$	70 %
$\ T_{ed}(\text{Empty})\ _{\infty} = 20.3$		
Loaded	$-0.018 \pm 0.018j$	71%
	$-0.27 \pm 3.47j$	7.7 %
	$-10.2 \pm 62.1j$	16 %
	-20.8	100 %
	-3553	100 %
$\ T_{ed}(\text{Loaded})\ _{\infty} = 74.6$		

Table 6.2: H_{∞} Design 2 Results

The controller, when used with the loaded configuration, resulted in less severe accelerations for the pulses, but failed to meet the specifications. This is especially true for the suspension deflections driven by large and huge rounded pulse inputs.




Figure 6.8: H_{∞} Design 2 - Empty Configuration - Noise Input




Figure 6.9: H_{∞} Design 2 - Empty Configuration - Tiny Rounded Pulse Input




Figure 6.10: H_{∞} Design 2 - Empty Configuration - Small Rounded Pulse Input




Figure 6.11: H_{∞} Design 2 - Empty Configuration - Medium Rounded Pulse Input




Figure 6.12: H_{∞} Design 2 - Empty Configuration - Large Rounded Pulse Input




Figure 6.13: H_∞ Design 2 - Empty Configuration - Huge Rounded Pulse Input




Figure 6.14: H_∞ Design 2 - Loaded Configuration - Noise Input




Figure 6.15: H_∞ Design 2 - Loaded Configuration - Tiny Rounded Pulse Input




Figure 6.16: H_∞ Design 2 - Loaded Configuration - Small Rounded Pulse Input




Figure 6.17: H_∞ Design 2 - Loaded Configuration - Medium Rounded Pulse Input




Figure 6.18: H_∞ Design 2 - Loaded Configuration - Large Rounded Pulse Input




Figure 6.19: H_∞ Design 2 - Loaded Configuration - Huge Rounded Pulse Input

6.5 Conclusions

It is shown that the H_∞ optimal controller can control the suspension and tire deflections. The controllers are robust from the empty to loaded configuration, but not the other way. Similar to the H_2 results, this is not indicated by the margins, which are larger for the loaded configuration than the empty configuration.

The H_∞ controllers have an advantage over the H_2 design in that they can directly control the tire deflection. Their disadvantage is that when designed around the empty configuration, the H_∞ optimization is too aggressive resulting in excessive accelerations that would not be tolerable in an implemented system. H_∞ Design 1 controller has better noise rejection performance than H_2 Design 1 as seen in the simulations. This was unexpected but is a result of the damping of one set of poles improving over the H_2 design by 11%, increasing from 48% in the H_2 Design 1 (Table 5.1) to 59% in the H_∞ Design 1 (Table 6.1).

VII. Performance Based Mixed-Design

7.1 Introduction

The first mixed-norm optimal design was based on the loaded configuration, with the goal of achieving the best performance possible. The problem was broken down into four sub-problems

1. H_2 objective function on truck mass acceleration
2. H_∞ constraint function on tire deflection
3. H_∞ constraint function on suspension deflection
4. H_∞ constraint function on weighted control usage

During the single-norm design, it was noticed that the controllers demanded very fast actuator movements. The actuator is modeled as a perfect actuator and therefore no actuator dynamics are taken into account. The weighting function

$$W_u(s) = \frac{10(s+1)}{(s+10)} \quad (7.1)$$

was used to penalize high frequency actuation. As actuator dynamics become available for this problem, they should be included in the problem set-up and simulation.

The H_2 objective function was set up identically to the single-norm H_2 design with the output vector

$$z = \begin{bmatrix} \text{Truck Mass Acceleration} \times 1 \\ \text{Control Usage} \times (1 \times 10^{-4}) \end{bmatrix} \quad (7.2)$$

The control usage was included in the output vector to ensure that the H_2 subproblem was regular. This allowed $K_{2_{opt}}$ to be used as the initial stabilizing controller, K_{init} .

The three H_∞ constraint subproblems were set up identically to the H_∞ designs of the previous chapter. However, each subproblem consisted of a scalar output and there were no sensor noises. This caused each constraint subproblem to be singular, but this is allowable in the mixed-norm approach as long as the entire problem is regular. All of the output

weightings were set to one, with the exception of the control usage which is now weighted according to Equation 7.1. This resulted in a mixed-norm problem with three constraint functions that can be written as

$$\inf_{K_{admissible}} \|T_{zw}\|_2 \quad (7.3)$$

subject to

$$\|T_{ed1}\|_\infty \leq \gamma_1$$

$$\|T_{ed2}\|_\infty \leq \gamma_2$$

$$\|T_{ed3}\|_\infty \leq \gamma_3$$

where $K_{admissible}$ is the set of all stabilizing fourth order controllers. The procedure followed was to first identify the Edgeworth-Pareto (EP) curves for each constraint separately, and examine the trade-offs with the other constraints. The final controller for this design will be based on the best performing controller from one of these single constraint designs. The other constraints will then be reduced further to see if any performance improvements are possible.

The norms based on K_{2opt} were

$$\underline{\alpha} = 24.1$$

$$\bar{\gamma}_1 = 27,331$$

$$\bar{\gamma}_2 = 27,600$$

$$\bar{\gamma}_3 = 5.15 \times 10^7$$

and provide the starting point for the designs.

7.2 $\|T_{ed1}\|_\infty$ - Tire Deflection Constraint

By restricting $\|T_{ed1}\|_\infty$ (denoted by γ_1^*) it is possible to control the tire deflection. MX-TOOLS was able to reduce γ_1^* from the initial 27,331 to 1.2. When the constraint was set




Figure 7.1: Edgeworth-Pareto Optimal Curve for Tire Deflection

lower than 1.2, the optimization routine would fail to find a controller. By plotting α^* versus γ_1^* it is possible to develop an EP curve. Figure 7.1 is the full EP curve and Figure 7.2 is a close-up of the knee of the curve.

Since the constraint function is an H_∞ constraint, there is a singular value plot associated with the EP curve. As the constraint is lowered, the peaks on the singular value plot are also lowered. This can be seen in Figure 7.3.

The last aspect of the optimization process is to determine what happens to the other closed-loop transfer function norms as the first constraint is reduced. The other norms were left unconstrained and therefore are not optimal. Both of the other two constraints' infinity norms initially appear to reduce linearly with the γ_1 constraint. This is true until γ_1^* is less than 100. Figures 7.4 and 7.5 show how their respective norms change as γ_1^* is reduced.

It is seen that γ_2 appears to be very sensitive to γ_1^* as γ_1^* is reduced below 5. This is shown by the discontinuous jump in γ_2 , and it may suggest that at lower values, the tire deflection and suspension deflection infinity norms may compete against each other. The behavior of γ_3 shows a definite trend that as γ_1^* is reduced below a certain value, the control usage goes up.




Figure 7.2: Edgeworth-Pareto Optimal Curve for Tire Deflection, Expanded




Figure 7.3: Singular Value Plot of $\|T_{ed1}\|_{\infty}$ - Tire Deflection




Figure 7.4: $\|T_{ed2}\|_{\infty}$ versus $\|T_{ed1}\|_{\infty}$




Figure 7.5: $\|T_{ed3}\|_{\infty}$ versus $\|T_{ed1}\|_{\infty}$




Figure 7.6: Edgeworth-Pareto Optimal Curve for Suspension Deflection

7.3 $\|T_{ed_2}\|_{\infty}$ - Suspension Deflection Constraint

The suspension and tire deflections have the same behavior. This is understandable since the suspension is directly connected to the tire, and it is the tire's deflection that drives the requirement for the suspension to move. Therefore, if one is reduced then so is the other. This trend continues with the reduction of γ_2^* . The lowest γ_2^* achievable by constraining γ_2 alone with MXTOOLS was 3.8. It was found that by constraining γ_1 as well, it was possible to achieve a lower value of γ_2^* , with the lowest value being 3.6277. This is simply due to numerics. The EP curve is shown in Figures 7.6 and 7.7. The singular value plot is shown in Figure 7.8.

The tire deflection norm reduces linearly with the reduction in the suspension deflection. Unlike reducing the tire deflection, there is no erratic behavior as γ_2^* is reduced below 5. There is a similar upswing in γ_3 as γ_2^* is restricted below 20. Figures 7.9 and 7.10 show how the unconstrained γ_1 and γ_3 behave as γ_2^* is reduced.

7.4 $\|T_{ed_3}\|_{\infty}$ - Weighted Control Usage Constraint

The third constraint was weighted control usage. Since the objective function is almost singular, the initial $\bar{\gamma}_3$ is very high. It was possible to reduce γ_3^* from 51 million to 16,000




Figure 7.7: Edgeworth-Pareto Optimal Curve for Suspension Deflection, Expanded




Figure 7.8: Singular Value Plot of $\|T_{ed2}\|_{\infty}$ - Suspension Deflection




Figure 7.9: $\|T_{ed1}\|_{\infty}$ versus $\|T_{ed2}\|_{\infty}$




Figure 7.10: $\|T_{ed3}\|_{\infty}$ versus $\|T_{ed2}\|_{\infty}$, Expanded




Figure 7.11: Edgeworth-Pareto Optimal Curve for Weighted Control Usage

with MXTOOLS. Figure 7.11 is the full EP curve and Figure 7.12 is the knee of the EP curve. It can be seen that this constraint can be pushed a great distance with only a small amount of increase in α^* . For example, near the end of the curve γ_3^* is reduced about 20,000 whereas the increase in α^* is only around 175.

Numerically, this is a difficult function to optimize. This is seen best in the singular values of $\|T_{ed3}\|_\infty$, shown in Figure 7.13. The numerical problems are related to the differences in the frequency search and the Ricatti solutions for γ_3^* . The spike around 60 rad/s is very narrow, which causes problems if the frequency search isn't fine enough. For example, from the initial controller ($K_{2_{opt}}$) we know that the lower bound of $\bar{\gamma}_3 = 5.1 \times 10^7$; however, the frequency search used to develop Figure 7.13 identified $\bar{\gamma}_3 = 4.8 \times 10^6$, which is wrong. The plot was done with 500 points between 10^{-3} and 10^3 rad/s. Fortunately, this problem is reduced when the constraint is lowered since the lower γ_3^* is, the flatter the curve and the less likely that the spike will be missed in a frequency search.

As the weighted control usage is reduced, the tire deflection reduces as well. This can be seen in Figure 7.14. This is not the case for the suspension deflection though. Initially the suspension deflection is reduced with the reduction of weighted control usage, but near the lower portion of the curve, the suspension deflection becomes erratic. The full curve is seen in Figure 7.15. A close-up of the erratic behavior is seen in Figure 7.16. There were




Figure 7.12: Edgeworth-Pareto Optimal Curve for Weighted Control Usage, Expanded




Figure 7.13: Singular Value Plot of $\|T_{ed3}\|_{\infty}$ - Weighted Control Usage

no restrictions placed on $\|T_{ed2}\|_\infty$ during the optimization, but this may imply competition between γ_2 and γ_3^* at the lower end of the possible γ_3^* curve.




Figure 7.14: $\|T_{ed1}\|_\infty$ versus $\|T_{ed3}\|_\infty$




Figure 7.15: $\|T_{ed2}\|_\infty$ versus $\|T_{ed3}\|_\infty$

7.5 Mixed-Norm Design 1

Mixed-Norm Design 1 was the result of further constraining the last controller found by reducing the single constraint on suspension deflection (T_{ed2}) as shown in Figure 7.7. This




Figure 7.16: $\|T_{ed2}\|_{\infty}$ versus $\|T_{ed3}\|_{\infty}$, Expanded

controller was chosen because its performance was closest to meeting the specifications. This controller is defined as K_{init} and resulted in the initial norms of

$$\alpha^* = 1,459$$

$$\gamma_1 = 3.92$$

$$\gamma_2^* = 3.8$$

$$\gamma_3 = 43,744$$

It was found that by holding $\gamma_2 \leq 3.8$ and reducing the constraint on γ_1 , the performance of the system could be improved slightly (the system is closer to meeting all of the specifications). In addition, since the actuator performance was well within specifications, the control usage was left unconstrained. It will be seen that this is not the case for Mixed-Norm Design 2, where it will need to be constrained. The results of the control usage constraint could be left out of the analysis for Mixed-Norm Design 1, but are included for completeness. Figures 7.17 - 7.22 show the simulation results based on K_{init} .

From the simulation results it is clear that even at the lowest γ_2^* reachable with MX-TOOLS by constraining γ_2 alone, the specification on suspension deflection can not be met for all pulse inputs. In particular, there are violations for the medium, large, and huge




Figure 7.17: Mixed-Norm Design 1 - Kinit Response - Noise Input




Figure 7.18: Mixed-Norm Design 1 - Kinit Response - Tiny Rounded Pulse Input




Figure 7.19: Mixed-Norm Design 1 - Kinit Response - Small Rounded Pulse Input




Figure 7.20: Mixed-Norm Design 1 - Kinit Response - Medium Rounded Pulse Input

Iteration	α^*	γ_1^*	γ_2^*	γ_3
Initial	1,459	3.9201	3.7999	43,744
A	1,462	3.9004	3.7948	43,754
B	1,464	3.8900	3.8000	43,766
C	1,469	3.8700	3.8000	43,790
D	1,473	3.8500	3.8000	43,815
E	1,476	3.8300	3.8000	43,839
F	1,483	3.8000	3.7532	43,877
G	1,494	3.7500	3.6277	43,942
Final	1,505	3.7000	3.7265	44,009

Table 7.1: Mixed-Norm Design 1 - Results

rounded pulse inputs. It may be possible to reduce γ_2^* further, and in fact it was shown that by reducing γ_1^* there was a reduction in γ_2^* as well.

The multiple constraint mixed-design consisted of reducing γ_1 , while holding the constraint for γ_2 constant ($\gamma_2 \leq 3.8$). This allowed for improvement in the tire deflection while maintaining (or improving) the suspension deflection. Table 7.1 shows how γ_1^* was reduced and its impacts on the norms of the other transfer functions.

It can be seen that the γ_1 constraint is always met with equality (within numerical accuracy), whereas the γ_2 constraint is only met with equality for some of the design points. When the γ_1 constraint was set lower than 3.7, MXTOOLS failed to find a solution. Figure 7.23 is the EP Optimal curve. Figure 7.24 is the singular value plot. Figures 7.25 - 7.26 show the resulting values of γ_2^* and γ_3 as γ_1^* is reduced.

The improvement in both γ_1^* and γ_2^* from the initial iteration to the final iteration did not improve the performance of the system significantly. The largest difference was noticed in the response to the huge rounded pulse. This can be seen in Figure 7.27. The initial reduction of 0.02 in γ_1^* ($K_{init} \rightarrow K_A$) does improve the damping of the tire response, but the improvements between iteration A from Table 7.1 and the Final iteration did not further improve the response. In fact, the further reduction in γ_1^* causes the suspension response to worsen slightly. This is best seen in the large rounded pulse response. The initial oscillation of the suspension response improved slightly, but caused overshoot in the second oscillation. A similar response is also seen in the tire deflection. Figure 7.28 has the response of the




Figure 7.23: Mixed-Norm Design 1 - EP Curve




Figure 7.24: Mixed-Norm Design 1 - Singular Values Plot




Figure 7.25: Mixed-Norm Design 1 - $\|T_{ed2}\|_{\infty}$ versus $\|T_{ed1}\|_{\infty}$




Figure 7.26: Mixed-Norm Design 1 - $\|T_{ed3}\|_{\infty}$ versus $\|T_{ed1}\|_{\infty}$




Figure 7.27: Mixed-Norm Design Candidates - Huge Rounded Pulse Response - (Dotted - Initial, Dashed - A, Solid - Final)

suspension for selected controllers from Table 7.1, and Figure 7.29 is the tire response for the same set of controllers.

The benefits of reducing γ_1^* and γ_2^* are not well represented in the simulations. This may be due to the nature of the H_∞ constraints. The H_∞ constraints are geared towards reducing the energy of the output, and this does appear to be happening as the settling time is reducing, but at the cost of higher maximum magnitude. This is why the ℓ_1 norm was initially desired for these constraints. The controller also appears to be increasingly aggressive as γ_1^* is reduced. This can be seen in Figure 7.30, which is the control usage for the medium rounded pulse. It can be seen that the control usage is increased as γ_1^* is




Figure 7.28: Mixed-Norm Design Candidates - Large Rounded Pulse - Suspension Responses




Figure 7.29: Mixed-Norm Design Candidates - Large Rounded Pulse - Tire Responses




Figure 7.30: Mixed-Norm Design Candidates - Medium Rounded Pulse Control Usage

reduced; however, the control usage decreases after the first reaction which allows the tire and suspension deflection to increase as well.

Based on the simulations, there is little improvement by reducing γ_1^* or γ_2^* beyond iteration A; therefore, the A controller was chosen as K_{mix} for Mixed-Norm Design 1. Simulations of the closed-loop system with K_{mix} can be found in Figures 7.31 - 7.36. Much like the single-norm designs based on the loaded configuration, this controller failed to stabilize the closed-loop system in the empty configuration. The vector gain margin for this design was $[-0.025, 0.025]dB$ and the vector phase margin was $\pm 0.165^\circ$. The small margins are similar to those of the H_2 designs and the H_∞ designs. The next chapter will include an additional constraint that will dramatically improve the margins.




Figure 7.31: Mixed-Norm Design 1 Controller - Noise Input




Figure 7.32: Mixed-Norm Design 1 Controller - Tiny Rounded Pulse Input




Figure 7.33: Mixed-Norm Design 1 Controller - Small Rounded Pulse Input




Figure 7.34: Mixed-Norm Design 1 Controller - Medium Rounded Pulse Input




Figure 7.35: Mixed-Norm Design 1 Controller - Large Rounded Pulse Input




Figure 7.36: Mixed-Norm Design 1 Controller - Huge Rounded Pulse Input

7.6 Conclusions

This chapter involved creating the best mixed-norm controller for a single design point. The design highlighted the process that is involved in finding a mixed-norm optimal controller. It was seen how each constraint could be reduced separately to show how the other norms of the system react. This provided the foundation for concurrent multiple active constraints. One interesting note was that by reducing γ_1 it was possible to reduce γ_2 further than reducing γ_2 by itself. Any method based on numerical optimization will have interesting facets such as this because of the nature of numerical optimization.

One of the key conclusions from this design is the fact that the best performing controller is not necessarily the last controller found. It was shown how reducing the infinity-norms of the suspension and tire deflections beyond a certain point worsened the time responses. It is important to know how far you can reduce the norms, but that knowledge does not replace the necessity of simulating the system with the various controllers to see how well they perform. The ability to trace out the EP curve does give the designer the ability to easily

perform trade-off analyses with the norms, which in turn may indicate which controllers perform better than the others.

VIII. Robustness Based Mixed-Design

8.1 Introduction

The second mixed-norm design was based on a medium (half-full) load configuration, with the goal of providing the best performance at that configuration while guaranteeing robust stability at both the loaded and empty configurations. The design is fundamentally the same as the previous mixed design, but with one additional constraint that, if satisfied, guarantees robustness.

In a single-norm H_∞ problem, robustness can be guaranteed because of the Small Gain Theorem. However, implementation is very difficult because the robustness weighting must be added to the performance objectives as well. Since the infinity-norm is the maximum singular value, it will be influenced by both the performance transfer functions and the robustness transfer functions. By forcing $\|T_{ed}\|_\infty \leq 1$ robustness will be guaranteed, but this may not provide the best performance while still being robust.

In the mixed-norm approach, implementation involves adding an additional constraint for robustness, then setting it to be less than or equal to one. This resulted in a mixed-norm problem with five subproblems:

1. H_2 objective function on truck mass acceleration
2. H_∞ constraint function on tire deflection
3. H_∞ constraint function on suspension deflection
4. H_∞ constraint function on weighted control usage
5. H_∞ constraint function for robustness

Each of the first four subproblems are identical to the first mixed-norm design, but using the plant matrix from the medium load configuration. This resulted in the mixed-norm problem

$$\inf_{K_{admissible}} \|T_{zw}\|_2 \quad (8.1)$$

subject to

$$\|T_{ed1}\|_\infty \leq \gamma_1$$

$$\|T_{ed2}\|_\infty \leq \gamma_2$$

$$\|T_{ed3}\|_\infty \leq \gamma_3$$

$$\|T_{ed4}\|_\infty \leq \gamma_4$$

By using $K_{2_{opt}}$ for K_{init} , $\underline{\alpha}$ and the various $\bar{\gamma}'s$ are defined. For the medium load configuration, the initial values were:

$$\underline{\alpha} = 42.75$$

$$\bar{\gamma}_1 = 86,240$$

$$\bar{\gamma}_2 = 83,760$$

$$\bar{\gamma}_3 = 1.58 \times 10^8$$

$$\bar{\gamma}_4 = 1,541$$

In similar fashion to the first mixed-norm design, the first step in the solution was to identify the behavior of the objective function and the constraint functions as each constraint is separately reduced. Then the best controller was selected that met the robustness test. That controller was then improved upon by further reducing the constraints. This resulted in the best controller for the medium configuration, which was then shown to be stable for both the empty and loaded configurations. In addition, the controller was then used on a 10 DOF simulation to see what, if any, control it had over the system.

8.2 Robustness Constraint

The robustness constraint was formed by using an additive perturbation model for the changes in the plant matrix. The additive uncertainty model is shown in Figure 8.1.




Figure 8.1: H_∞ System with Additive Uncertainty

Here, $\|\Delta\|_\infty < 1$ will be assumed. Taking the worst case with $\Delta = I$, we can solve for W knowing P and $\tilde{P}$, as

$$W = \tilde{P} - P \quad (8.2)$$

where $\tilde{P}$ represents either the loaded or empty plant, and P is the medium load plant. Since there are two off-nominal conditions that must be met, there are two weighting functions. This results in the need for an overbound function to be created. Figure 8.2 shows the two weighting functions and the overbound function used as W .




Figure 8.2: Robustness Weighting Functions

The last subproblem was then set up as in Figure 8.3, where W is the transfer function that corresponds to the overbound weighting function. The overbound weighting function was

$$W(s) = \frac{1.62(s^2 + 0.12s + 0.01)(s^2 + 0.5s + 0.25)}{(s^2 + 0.9308s + 12.8164)(s^2 + 4.0s + 64.0)} \quad (8.3)$$

By using W as the weighting function, if $\|T_{ed_4}\|_\infty \leq 1$, the system is guaranteed to be robustly stable for the loaded and empty configurations.




Figure 8.3: Robustness Subproblem

8.3 $\|T_{ed_1}\|_\infty$ - Tire Deflection Constraint

Reducing $\|T_{ed_1}\|_\infty$ (denoted by γ_1^*) provided the ability to control the tire deflection. The full EP curve is shown in Figure 8.4, and a close-up at the knee of the curve is shown in Figure 8.5. Figure 8.6 is the singular value plot for the reduction of γ_1^* . MXTOOLS was able to find solutions down to $\gamma_1^* = 2$. The infinity-norm of T_{ed_2} reduced linearly with the reduction in γ_1^* . This was true until $\gamma_1^* < 10$, where γ_2 became erratic and finally exploded to 9,164 when $\gamma_1^* = 2$. This can be seen in Figure 8.7. The weighted control usage also reduced with a reduction in γ_1^* ; however, it started to increase when $\gamma_1^* < 20$, and increased to 1.1×10^6 when $\gamma_1^* = 2$. This can be seen in Figure 8.8. Finally, the robustness weighting infinity-norm dropped initially to around ten and then remained at ten until γ_1^* was reduced

below twenty, and then shot up to a final value of 715. This does imply that none of these designs could guarantee robust stability. Figure 8.9 shows how γ_4 changed with γ_1^* .




Figure 8.4: Edgeworth-Pareto Optimal Curve for Tire Deflection




Figure 8.5: Edgeworth-Pareto Optimal Curve for Tire Deflection, Expanded




Figure 8.6: Singular Value Plot of $\|T_{ed1}\|_{\infty}$ - Tire Deflection




Figure 8.7: $\|T_{ed2}\|_{\infty}$ versus $\|T_{ed1}\|_{\infty}$




Figure 8.8: $\|T_{ed3}\|_{\infty}$ versus $\|T_{ed1}\|_{\infty}$




Figure 8.9: $\|T_{ed4}\|_{\infty}$ versus $\|T_{ed1}\|_{\infty}$

8.4 $\|T_{ed_2}\|_\infty$ - Suspension Deflection Constraint

Reducing $\|T_{ed_2}\|_\infty$ (denoted by γ_2^*) provided the ability to control the suspension deflection. The full EP curve is shown in Figure 8.10, and a close-up at the knee of the curve is shown in Figure 8.11. Figure 8.12 is the singular value plot for the reduction of γ_2^* . MXTOOLS was able to find solutions down to $\gamma_2^* = 1.8$. The infinity-norm of T_{ed_1} reduced linearly with the reduction in γ_2^* . This can be seen in Figure 8.13. The weighted control usage also reduced with a reduction in γ_2^* , however it started to increase when $\gamma_2^* < 10$, and then increased. This can be seen in Figure 8.14. Finally, the robustness weighting infinity-norm initially dropped to around ten and remained there until γ_2^* was reduced below thirty. Then γ_4 increased to over 100 at $\gamma_2^* = 5$, and fell back down to around 20 when $\gamma_2^* = 1.8$. This again implies that none of these designs could guarantee robust stability. Figure 8.15 shows how γ_4 changed with γ_2^* .




Figure 8.10: Edgeworth-Pareto Optimal Curve for Suspension Deflection




Figure 8.11: Edgeworth-Pareto Optimal Curve for Suspension Deflection, Expanded




Figure 8.12: Singular Value Plot of $\|T_{ed2}\|_{\infty}$ - Suspension Deflection




Figure 8.13: $\|T_{ed1}\|_{\infty}$ versus $\|T_{ed2}\|_{\infty}$




Figure 8.14: $\|T_{ed3}\|_{\infty}$ versus $\|T_{ed2}\|_{\infty}$




Figure 8.15: $\|T_{ed4}\|_{\infty}$ versus $\|T_{ed2}\|_{\infty}$

8.5 $\|T_{ed3}\|_{\infty}$ - Weighted Control Usage Constraint

Reducing $\|T_{ed3}\|_{\infty}$ (denoted by γ_3^*) provided the ability to control the usage of the actuator. The full EP curve is shown in Figure 8.16, and a close-up at the lower end of the curve is shown in Figure 8.17. Figure 8.18 is the singular value plot for the reduction of γ_3^* . MXTOOLS was able to find solutions down to $\gamma_3^* = 29,500$. The infinity-norm of T_{ed1} reduced linearly with the reduction in γ_3^* . This can be seen in Figure 8.19. The suspension deflection also reduced with a reduction in γ_3^* ; however, it behaved erratically when γ_3^* was below 60,000. This included a small dip and a spike. This can be seen in Figure 8.20. Finally, the robustness weighting infinity-norm was very erratic with one very large spike that corresponds to a similar spike in γ_2 . All of the values of γ_4 were larger than one. This again implies that none of these designs could guarantee robust stability. Figure 8.21 shows how γ_4 changed with γ_3^* .




Figure 8.16: Edgeworth-Pareto Optimal Curve for Weighted Control Usage




Figure 8.17: Edgeworth-Pareto Optimal Curve for Weighted Control Usage, Expanded




Figure 8.18: Singular Value Plot of $\|T_{ed3}\|_{\infty}$ - Weighted Control Usage




Figure 8.19: $\|T_{ed1}\|_{\infty}$ versus $\|T_{ed3}\|_{\infty}$




Figure 8.20: $\|T_{ed2}\|_{\infty}$ versus $\|T_{ed3}\|_{\infty}$




Figure 8.21: $\|T_{ed4}\|_{\infty}$ versus $\|T_{ed3}\|_{\infty}$

8.6 $\|T_{ed_4}\|_\infty$ - Robustness Weighting Constraint

Unlike the other constraint functions, the robustness weighting constraint was nearly impossible to optimize alone. This is due to an extreme spike in the singular value plot of T_{ed_4} . This spike is missed by the optimization routine because of its narrow frequency band. Figure 8.22 shows how MATLAB's SIGMA function plots the singular values of T_{ed_4} at the medium configuration with $K_{2_{opt}}$. Figure 8.23 is the same plot, but with an additional 8000 points plotted between 55 and 60 rad/s. MXTOOLS incorrectly identified the peak on the left as the maximum peak resulting in an incorrect $\|T_{ed_4}\|_\infty$ value of 9.6 (rather than the actual 1,541).




Figure 8.22: Singular Value Plot of T_{ed_4} for $K_{2_{opt}}$ with MATLAB Generated Points




Figure 8.23: Singular Value Plot of T_{ed_4} for $K_{2_{opt}}$ with 8000 Additional Points




Figure 8.24: Singular Value Plot of T_{ed_4} for new K_{init}

MXTOOLS was able to find a controller with $\|T_{ed_4}\|_{\infty} \leq 1$ by pushing down on the two peaks on the left and ignoring the peak at 58 rad/s. Reducing the two peaks on the left caused the right hand peak to reduce as well. This resulted in a controller with guaranteed robust stability, which was then used as the starting point for the mixed-norm design. Figure 8.24 is the singular value plot for this controller. As can be seen, the sharp peak on the right hand side is now much lower than the other peaks. This figure was done with the same frequencies as in Figure 8.23. This allows the optimization to continue; however, since the sharp peak is ignored in the optimization routine, this controller is not necessarily EP optimal. This implies that if the third peak was included, a lower two-norm may have been found if the optimal controller caused the peak at 58 rad/s to be equal to one.

With this initial robust stabilizing controller (a new K_{init}), it is possible to push down on the other constraints to improve performance. This is done in the next section.

8.7 Robust Design

The robust design started with the controller found in the last section, now defined as K_{init} .

The initial norms for this design were:

$$\alpha = 1,907$$

$$\gamma_1 = 42.29$$

$$\gamma_2 = 42.75$$

$$\gamma_3 = 13,440$$

$$\gamma_4 = 1.00$$

The margins for this initial controller were $VGM_s = [-4.39, 9.31]dB$ and $VGM_t = [-11.73, 4.81]dB$ with $VPM = \pm 43.50^\circ$. Because of the robustness constraint, the margins are much higher than any other previous design.

The simulations with the initial controller can be seen in Figures 8.25 - 8.30. There are multiple violations of the specifications, with only the tiny rounded pulse and road noise inputs producing acceptable results. The actuator is not used much, resulting in longer settling times than before. It is also noticeable that the medium configuration results in the quick oscillations that are similar to the empty configuration. The goal of this design will be to improve on this initial response by reducing γ_1 and γ_2 . The weighted control usage isn't initially constrained, since the control usage is well within limits.




Figure 8.25: Mixed-Norm Design 2 - Kinit Response - Noise Input




Figure 8.26: Mixed-Norm Design 2 - Kinit Response - Tiny Rounded Pulse Input




Figure 8.27: Mixed-Norm Design 2 - Kinit Response - Small Rounded Pulse Input




Figure 8.28: Mixed-Norm Design 2 - Kinit Response - Medium Rounded Pulse Input




Figure 8.29: Mixed-Norm Design 2 - Kinit Response - Large Rounded Pulse Input




Figure 8.30: Mixed-Norm Design 2 - Kinit Response - Huge Rounded Pulse Input

Using MXTOOLS, $\|T_{ed_1}\|_\infty$ and $\|T_{ed_2}\|_\infty$ were reduced from 42.29 to 5.1 and 42.75 to 5.1, respectively. This resulted in the best performance possible for a robust design without constraining the weighted control usage constraint. This resulted in:

$$\alpha^* = 4,459$$

$$\gamma_1^* = 5.10$$

$$\gamma_2^* = 5.10$$

$$\gamma_3 = 49,910$$

$$\gamma_4^* = 1.00$$

This design, however, resulted in excessive control usage. Therefore, the constraints on tire and suspension deflection were loosened and a constraint on weighted control usage was imposed. This reduced the control usage to back within specifications. The resulting values of the norms were:

$$\alpha^* = 4,352$$

$$\gamma_1^* = 5.25$$

$$\gamma_2^* = 5.25$$

$$\gamma_3^* = 49,300$$

$$\gamma_4^* = 1.00$$

Since the design is based on robustness, the margins are very important. The vector gain margins were $VGM_s = [-5.134, 14.29]dB$ and $VGM_t = [-35.461, 5.947]dB$. The vector phase margins were $VP M_s = \pm 47.60^\circ$ and $VP M_t = \pm 58.89^\circ$. The margins were evaluated at the medium configuration. For most of the development of this controller, there were three constraints imposed on the solution: the tire deflection, the suspension deflection, and the robustness weighting. Each EP optimal controller had at a minimum one active constraint,




Figure 8.31: Edgeworth-Pareto Optimal Curve for Tire Deflection

and for this case, there were always at least two active constraints. Figures 8.31 and 8.32 show how the two-norm of the objective functions changed when the suspension constraints were reduced. In both plots, the x's reflect the controllers found by MXTOOLS, and the 'o' represents the controller chosen after imposing a restriction on the weighted control usage. The final controller is placed on these charts for reference only, in order to show the values of the constraint norms for the final controller in relation to the controllers formed in the creation of these EP curves. The final controller is actually slightly higher than the curve because the additional constraint causes a higher two-norm.

For most of the designs the weighted control usage was not constrained. Figure 8.33 shows how the weighted control usage changed as the other constraints were restricted. The weighted control usage went down initially but then started to increase. This curve is not EP optimal for the weighted control usage. The 'o' on the figure (which is not on the curve) is the final constrained control usage controller.

For the design to be guaranteed robust, the infinity-norm of the fourth constraint must be equal to or less than one. Figure 8.34 shows that the infinity-norm always remained at one with small variations each way. These variations are the result of the difference in the frequency search and a Ricatti solution for the H_∞ norm. The optimization routine uses a frequency search, and if the infinity-norm found during the frequency search is within




Figure 8.32: Edgeworth-Pareto Optimal Curve for Suspension Deflection




Figure 8.33: Response of $\|T_{ed3}\|_{\infty}$ to Multiple Constraints Enforced




Figure 8.34: Edgeworth-Pareto Optimal Curve for Robustness Weighting

tolerance, the controller is accepted as being EP optimal. All of the controllers in Figure 8.34 met the frequency search tolerance, but when the infinity-norm was calculated through Ricatti methods, the norms changed slightly and that is shown in the figure. The resulting controllers were stabilizing for the two extreme configurations, and therefore further restriction of the robustness weighting constraint wasn't necessary. If the minor deviation above one caused the extreme configurations to become unstable, the robustness weighting constraint could have been reduced to a value less than one such that the variations would not cause the infinity-norm to be greater than one. This would further improve the margins and provide a stronger guarantee of robust stability for the design.

8.8 Results

The closed-loop system was simulated with the medium, empty, and loaded configurations. Figures 8.35 - 8.40 show the simulations for the medium configurations. The empty configuration simulations can be found in Appendix G. The loaded simulations can be found in Appendix H. This controller does in fact stabilize both the empty and loaded configurations and therefore the robustness constraint achieved its goal.

The performance of the controller failed to meet the specifications at the medium configuration. Specifications were only met for the tiny rounded pulse input. The accelerations




Figure 8.35: Mixed-Design 2 - Medium Configuration - Noise Input

of this design often resulted in very short spikes in both positive and negative acceleration and would have resulted in damage to the vehicle or to the driver.

When the controller is used at the empty configuration, the system is stable, but the performance again is out of specification. On the positive side, the accelerations felt by the system due to the rounded pulse inputs are smaller for every case than the passive suspension system. The noise rejection is about the same for both the passive and active systems, with the active system actuation being very fast.

The design was also simulated with the ten DOF model in the medium configuration. The measurements were modified to represent the same measurements that the two DOF model has. The measurements fed back were: cabin vertical accelerations, rear suspension deflection, and rear tire deflection. The results of the simulations can be seen in Appendix I.

The controller designed based on the two DOF model failed to improve the performance of the ten DOF system. This could be accounted for by the lack of controllability of the system, or the choice of feedback variables. The result is that this controller showed promise




Figure 8.36: Mixed-Design 2 - Medium Configuration - Tiny Rounded Pulse Input




Figure 8.37: Mixed-Design 2 - Medium Configuration - Small Rounded Pulse Input




Figure 8.38: Mixed-Design 2 - Medium Configuration - Medium Rounded Pulse Input




Figure 8.39: Mixed-Design 2 - Medium Configuration - Large Rounded Pulse Input




Figure 8.40: Mixed-Design 2 - Medium Configuration - Huge Rounded Pulse Input

on the two DOF model, but it was not capable of achieving the same performance using the ten DOF model.

8.9 Conclusions

This chapter involved creating a robustly stabilizing controller through the use of MX-TOOLS. It was shown that by creating a robust stability constraint, it was possible to generate controllers that stabilized both the empty and loaded configurations while providing the best performance possible at a medium load configuration. This method highlighted the ability of the mixed-norm approach to create a separate robust stability constraint that does not get blended with other constraints, which is the case in the single-norm methods. With the robust stability constraint satisfied, the margins of the system were vastly improved over all of the previous designs.

IX. Conclusions and Recommendations

9.1 Introduction

The goal of this research was to apply mixed-norm optimal control to an active suspension problem to see how well the mixed-norm synthesis method compared to the more traditional H_2 and H_∞ synthesis methods. The hope was to develop an active suspension system that could perform better than the passive suspension system. The problem faced was anything but academic in nature, with multiple configurations and widely varied dynamics. This chapter represents the conclusions drawn from the research and a set of recommendations on a possible alternative framework to find a solution.

9.2 Conclusions

There are several areas of this thesis that offer conclusions. The conclusions fall into three categories: performance, modeling and general. The performance conclusions are the result of comparisons between the different designs in this thesis. The modeling conclusions result from the analysis of the models used in the design and evaluation. Finally, the general conclusions are the combination of both the performance and modeling, and how they relate to each other in the solution to this problem.

9.2.1 Performance

It is possible to compare the performance of the passive system in the loaded configuration, the H_2 Design 1 controller, the H_∞ Design 1 controller, and Mixed-Norm Design 1 controller. Each of the active systems are based on the loaded configuration and had similar objectives. These designs can also be compared to Mixed-Norm Design 2 controller at the loaded configuration as well, even though it wasn't nominally designed for this configuration.

None of these designs, including the passive system, met the specifications. The passive system violated the tire suspension specification with the medium rounded pulse. When the damping in the system is replaced by the actuator, it is not possible to prevent multiple violations of the specifications with any of the control synthesis methods used. Each of the active suspension designs had multiple violations; the H_2 Design 1 having three violations,

the H_∞ Design 1 having two, the Mixed-Norm Design 1 having six, and Mixed-Norm Design 2 having five.

In comparing the different designs in response to the road noise, a full Monte Carlo simulation is required. However, based on the simulation that was run with the same filtered white noise, the indication is that the best performer is the passive system. For the length of the simulations, the passive system's response remained under 0.05 m/s^2 , whereas all of the active systems were much larger in comparison. Between the active systems, H_∞ Design 1, the Mixed-Norm Design 1, and Mixed-Norm Design 2 all had similar results with the responses under 0.1 m/s^2 . Unexpectedly, H_2 Design 1 had the worst performance with the responses under 0.4 m/s^2 . This implies that the H_2 design sacrificed good noise response as a result of trying to get the deflections within specifications.

Comparing the results of the rounded pulse inputs is difficult at best. In many ways it is unfair to compare them at all. If a specification was violated, the comparison of the accelerations would be skewed since the actual system could not respond the same as the simulation. Either the tire or suspension deflections would hit their limits and hold constant or the hardware would break. It was unexpected, though, that the performance of the Mixed-Norm Design 2 compared well with the other designs since it was not optimized around the loaded configuration.

It is possible to compare the designs for the tiny rounded pulse, since all of the designs remain in specifications. In no way can the response to one input define the best controller, but it does offer some insight on the performance of each controller. To start with, the passive system resulted in an acceleration spike of around 7 m/s^2 . Both the H_2 Design 1 and H_∞ Design 1 spiked at 6 m/s^2 . Mixed-Norm Design 1 spiked at 1 m/s^2 , and Mixed-Norm Design 2 spiked at 0.8 m/s^2 . The trade-off for the lower acceleration in the Mixed-Norm Design 2 was a series of several additional oscillations prior to settling down. For slightly more acceleration, the Mixed-Norm Design 1 resulted in a much less oscillatory response. The overall trade-off in the mixed-norm designs was lower acceleration responses in general but with a larger number of specifications broken.

There are a few observations that should be made about the performance of the mixed-norm designs in comparison with the single-norm designs. The mixed-norm design has the

potential to always perform equal to or better than either an H_2 or H_∞ design. If the objective function in a mixed-norm design is identical to that of an H_2 design and the constraints are inactive, then the resulting controller is $K_{2_{opt}}$. If a constraint function is set up identically to that of an H_∞ design and the constraint is set to $\underline{\gamma}$, the resulting K_{mix} will result in the same infinity-norm as the H_∞ design. This does assume that K_{init} is of equal or higher order than the H_2 or H_∞ subproblem. This of course isn't in general how the mixed-norm problem is set up, since it doesn't exploit the reasons to do a mixed-norm design.

This results in the conclusion that the mixed-norm designs can perform as well as their single-norm counterparts, and will in general perform better. For this problem, the single-norm methods provided controllers that violated less specifications than the mixed-norm designs, but at the cost of poorer acceleration responses. The mixed-norm designs were limited by the numerical difficulties surrounding this problem and the inability to use the original approach which would have fully exploited the power of the mixed-norm method.

A second conclusion based on the performance of the designs used in this thesis is that this problem may be better solved if the controller was scheduled in accordance to the load of the system. The dynamics of the system change radically between the empty and loaded configurations. It was shown that the controllers designed at the loaded configurations would not stabilize the empty configuration, and that the controllers designed at the empty configurations did not provide good performance at the loaded configuration. A better method possibly would be to generate controllers based on various loads and schedule between them. This may provide better performance throughout the load envelope.

9.2.2 Modeling

In order for this problem to be solved, the actual tractor semitrailer vehicle must be controllable by the active suspension system. If not fully controllable, at a minimum the tire and suspension deflections as well as the cabin and trailer accelerations must be controllable. It was seen in Chapter 4 that the model closest to reality, the ten DOF model, is marginally controllable and causes MATLAB's numerics to break down. This caused severe restrictions on design using the ten DOF model and forced the designs to be based off the two DOF

model. By implementing the controllers designed from the two DOF model, it was hoped that some control effectiveness would be seen during the simulations of these controllers with the ten DOF model. However, this did not happen and the simulations showed that the system was not well controlled by Mixed-Norm Design 2. This results in the conclusion that the tractor semitrailer vehicle is not practically controllable by the controllers designed in this thesis. Therefore, before this problem can be solved, the controllability of the ten DOF model must be improved.

The original approach to this problem was to use ℓ_1 constraints on the suspension and tire deflections. This was not possible because of the inability to digitize the system, and therefore H_∞ constraints were substituted in place of the ℓ_1 constraints. The designs based on the H_∞ constraints failed to meet the specifications because their maximum magnitude exceeded the limits imposed by the problem. This results in the conclusion that H_∞ constraints are most likely not the best choice as the constraints for this problem. Since it wasn't possible to use ℓ_1 constraints, it is unknown if ℓ_1 would have produced better results. However, due to the nature of the ℓ_1 norm, it is likely that the application of ℓ_1 constraints would have resulted in better results.

One of the assumptions used in this thesis is that the specifications for the ten DOF model could be directly applied to the lower order models. This assumption was made because of the lack of evidence to the contrary. In retrospect, this may be an invalid assumption because the weight distribution of the vehicle is different for each model. This implies that for each model, the deflections caused by the same input would, in general, be different. It also implies that the force required by the actuator would also be different for each model. Before further research is done, the specifications should be defined for each model, such that if they are satisfied at the lower order models, the higher order models will also be satisfied.

9.2.3 *General*

The goal of this thesis was to create an active suspension design that performed better than the current passive suspension system. The main conclusion of this thesis is that it is not possible to achieve the goal with the controllers designed in this thesis. This does not

imply that there does not exist a solution to this problem, but it does imply that a different approach is necessary to find a solution, if one exists.

Designs were done on the two DOF model, and the various methods have been compared. None of the methods clearly outperformed the others. It was shown that the mixed-norm designs provided similar responses to the single-norm controllers, but provided several design advantages over the single-norm approach. In addition, it was shown that the mixed-norm approach allowed significant improvements in the way of margins by adding a robust stability constraint which provided guaranteed robust stability. This was shown in Mixed-Norm Design 2, which is the only design with acceptable margins. The conclusion is that the mixed-norm approach is better than the single-norm approach because the set-up of the problem allows much more flexibility for the designer, and the resulting controllers are optimal with respect to the actual design goals, and not a compromise between competing objectives as in the single-norm approach. Unfortunately, this problem with its modeling difficulties and the lack of being able to use ℓ_1 constraints was not able to highlight the strengths of the mixed-norm approach to its full potential. The next section provides some insight on how this problem may be solved in the future.

9.3 Recommendations

To solve this problem within the given specifications, a different approach is going to be needed. This research has provided several ideas that may provide the key to finding a solution to this problem.

The first problem that needs to be overcome is the controllability of the ten DOF model. It is very difficult to meet tight specifications for a system which is barely controllable. The goal of the suspension system was to improve the driver's acceleration response to the road. The current location of the actuator is in the rear suspension between the rear axle and the chassis. However, control of the rear of the chassis has two major problems. First of all, the dynamics that the actuator must control are directly affected by the load. The kingpin is directly above the rear suspension; therefore, as the load increases, the force required by the actuator to move the chassis will increase as well. This drives the requirement to schedule any actuator at that location. Secondly, the cabin is located on the front of the chassis, yet

the actuator is at the rear of the chassis. Therefore, when a road disturbance comes into the system, the active suspension system at the rear of the chassis is trying to counteract the disturbance and then, while it is trying to regulate the system, it is hit by the same disturbance. This makes it very difficult to control the accelerations of the cabin from the rear of the chassis.

An alternative to the current actuator location would be to place an actuator at the rear of the cabin. If needed, an additional actuator could be placed in the front of the cabin as well. Since it is the cabin that needs to be controlled, it would be better to provide actuation directly on the cabin. This may also reduce or eliminate the need to schedule the controller as well, as the weight of the cabin does not change with the weight of the trailer. Therefore, the influence on the change of the dynamics would be less than the current actuator location.

By placing the active suspension system beneath the cabin, the cabin should be able to be controlled. This would accomplish one of the tasks of the active suspension system. The other task was to minimize the accelerations on the cargo. Again, like the cabin, it would be desirable to isolate the trailer, but this is not possible. With a combination of actuators on the cabin and the rear suspension, it may be possible to have enough control power to improve the performance of the trailer. It should also be possible to add an additional actuator at the rear of the trailer. This would not be a preferred solution, because the power for the actuator would have to be transmitted to the trailer, and the fact that the trailer doesn't necessarily stay with the cabin. However, this could be standardized on a fleet of trucks.

If in fact it was decided to put an actuator or a pair of actuators under the cabin, the ten DOF model would have to be used for controller design. This is because the first assumption to reduce the model was the combination of the cabin, motor, and chassis into one mass. The use of an H_2 controller on the ten DOF model would involve at a minimum a twenty state controller. This exceeds the limitations placed on the sampling rate and therefore a reduced order controller would have to be created. This could lead into an interesting study of how low an order is possible with adequate results. This would also require a synthesis method that allows the designer to pick the order of the controller.

As the problem was posed, it fits well into the ability of MXTOOLS to solve. However, it fits well as an $H_2/\ell_1/H_\infty$ problem. This involves overcoming the second problem, digitizing the system. Currently, the system can't be accurately digitized because of the widely separated poles of the system. One solution would be to separate the system into two systems, one slow and one fast. This is possible through the MATLAB function SLOWFAST. This would create two plants for the system, one with the slow dynamics of the system and one with the fast dynamics. Since this would eliminate the extreme separation of the poles, it should be possible then to digitize each system accurately. A second approach would be to keep the passive damping in the system, and design the controller to only augment the passive system, instead of replacing it.

With the system digitized, it should be possible to use ℓ_1 constraints. The mixed-norm problem would then be a multi-plant problem with multiple ℓ_1 constraints. This would require modification to MXTOOLS to allow for multiple ℓ_1 constraints with possibly different sampling frequencies. MXTOOLS would also need to be modified to place a stability constraint on each subproblem, to further restrict the admissible region to include controllers that stabilize P_2 as well as each of the constraint transfer functions. This is required because the assumption that a controller that stabilizes P_2 also stabilizes P_∞ and P_1 only holds if the core dynamics of each plant matrix are the same. There would be two choices for the plant of the objective subproblem, the fast or the slow plant. Choosing the correct plant for the objective would be one of the analyses that would have to be done after the problem is set up. In addition, a pair of infinity-norm constraints would need to be added to provide robust stability and improve the margins of each system. ℓ_1 constraints could be used, but since it is numerically easier to solve the constrained H_∞ problem, H_∞ constraints may be a better choice.

The resulting problem as stated here would be difficult to solve and would no doubt pose yet unknown challenges itself. The problem would be an excellent test for the full exploitation of MXTOOLS. This problem would be unique in that it would involve several aspects of a mixed-norm problem never considered together in one problem. It would involve finding a reduced order controller for a MIMO multi-plant mixed-norm problem with multiple ℓ_1 and H_∞ constraints. To anybody who tries, good luck.

Appendix A. Removing the $D_{yu} = 0$ Assumption

If $D_{yu} \neq 0$, then a loop shifting technique can be used to solve for the controller [SLC89]. The method shifts the output in order to set $D_{yu} = 0$, solves for the corresponding controller, and then shifts the controller to take into account the non-zero D_{yu} term.

The shift involves redefining the output vector

$$\tilde{y} = y - D_{yu}u \quad (\text{A.1})$$

which for the mixed norm case is generalized to

$$\tilde{y}_2 = y_2 - D_{yu}u \quad (\text{A.2})$$

$$\tilde{y}_1 = y_1 - D_{yu}u \quad (\text{A.3})$$

$$\tilde{y}_\infty = y_\infty - D_{yu}u \quad (\text{A.4})$$

Using the redefined output vectors, the resulting problem satisfies the assumption that $D_{yu} = 0$. Therefore, the mixed-norm problem is solved with the redefined output vector.

The resulting controller is then

$$K = \left[\begin{array}{c|c} A_c & B_c \\ \hline C_c & D_c \end{array} \right] \quad (\text{A.5})$$

This controller is then shifted to form

$$\tilde{K} = \left[\begin{array}{c|c} \tilde{A}_c & \tilde{B}_c \\ \hline \tilde{C}_c & \tilde{D}_c \end{array} \right] \quad (\text{A.6})$$

where

$$\tilde{A}_c = A_c - B_c D_{yu} (I + D_c D_{yu})^{-1} C_c \quad (\text{A.7})$$

$$\tilde{B}_c = B_c [I - D_{yu} (I + D_c D_{yu})^{-1} D_c] \quad (\text{A.8})$$

$$\tilde{C}_c = (I + D_c D_{yu})^{-1} C_c \quad (\text{A.9})$$

$$\tilde{D}_c = (I + D_c D_{yu})^{-1} D_c \quad (\text{A.10})$$

The one restriction to this shifting method is that the system is assumed well-posed, implying that $(I + D_c D_{yu})^{-1}$ exists.

Appendix B. Two DOF Model

This appendix covers the background information on the two DOF model. The two DOF model is shown in Figure B.1.




Figure B.1: Two DOF Model

There are four states for this system:

- x_1 Axle Displacement
- x_2 Truck Mass Displacement
- x_3 Axle Velocity
- x_4 Truck Mass Velocity

There are three configurations for the truck mass: empty, medium, or loaded. Each configuration changes the truck mass, but not the values for the spring constants or the damper constant. The values of the constants are shown in Table B.1.

Constant	Configuration	Value
K_{sr}	All	5,000,000 MN/m
K_{tr}	All	125,000 MN/m
B_{sr}	All	2500 kNs/m
Axle Mass	All	1500 MG
Truck Mass	Empty	1457.9 MG
Truck Mass	Medium	3015.3 MG
Truck Mass	Loaded	9484.8 MG

Table B.1: Two DOF Model Parameters

Appendix C. Ten DOF Passive Suspension Responses

Figure C.1 - C.6 show the open-loop passive responses of the ten DOF system in the empty configuration. Figures C.7 - C.12 show the same for the loaded configuration. The legend for all figures is

Style	Acceleration	Deflection
-	Cabin Vertical	Front
--	Cabin Rotational	Rear
---	Trailer Rotational	Trailer




Figure C.1: Ten DOF Passive - Empty Configuration - Noise Input




Figure C.2: Ten DOF Passive - Empty Configuration - Tiny Rounded Pulse Input




Figure C.3: Ten DOF Passive - Empty Configuration - Small Rounded Pulse Input




Figure C.4: Ten DOF Passive - Empty Configuration - Medium Rounded Pulse Input




Figure C.5: Ten DOF Passive - Empty Configuration - Large Rounded Pulse Input




Figure C.6: Ten DOF Passive - Empty Configuration - Huge Rounded Pulse Input




Figure C.7: Ten DOF Passive - Loaded Configuration - Noise Input




Figure C.8: Ten DOF Passive - Loaded Configuration - Tiny Rounded Pulse Input




Figure C.9: Ten DOF Passive - Loaded Configuration - Small Rounded Pulse Input




Figure C.10: Ten DOF Passive - Loaded Configuration - Medium Rounded Pulse Input




Figure C.11: Ten DOF Passive - Loaded Configuration - Large Rounded Pulse Input




Figure C.12: Ten DOF Passive - Loaded Configuration - Huge Rounded Pulse Input

Appendix D. Two DOF Passive Suspension Responses

Figure D.1 - D.6 show the open-loop passive responses of the two DOF system in the empty configuration. Figures D.7 - D.12 show the same for the loaded configuration.




Figure D.1: Two DOF Passive - Empty Configuration - Noise Input




Figure D.2: Two DOF Passive - Empty Configuration - Tiny Rounded Pulse Input




Figure D.3: Two DOF Passive - Empty Configuration - Small Rounded Pulse Input




Figure D.4: Two DOF Passive - Empty Configuration - Medium Rounded Pulse Input




Figure D.5: Two DOF Passive - Empty Configuration - Large Rounded Pulse Input




Figure D.6: Two DOF Passive - Empty Configuration - Huge Rounded Pulse Input




Figure D.7: Two DOF Passive - Loaded Configuration - Noise Input




Figure D.8: Two DOF Passive - Loaded Configuration - Tiny Rounded Pulse Input




Figure D.9: Two DOF Passive - Loaded Configuration - Small Rounded Pulse Input




Figure D.10: Two DOF Passive - Loaded Configuration - Medium Rounded Pulse Input




Figure D.11: Two DOF Passive - Loaded Configuration - Large Rounded Pulse Input




Figure D.12: Two DOF Passive - Loaded Configuration - Huge Rounded Pulse Input

Appendix E. Ten DOF Active Open-Loop Responses

Figure E.1 - E.6 show the open-loop active responses of the ten DOF system in the empty configuration. Figures E.7 - E.12 show the same for the loaded configuration. The legend for all figures is

Style	Acceleration	Deflection
-	Cabin Vertical	Front
- -	Cabin Rotational	Rear
- . - .	Trailer Rotational	Trailer




Figure E.1: Ten DOF Open-Loop - Empty Configuration - Noise Input




Figure E.2: Ten DOF Open-Loop - Empty Configuration - Tiny Rounded Pulse Input




Figure E.3: Ten DOF Open-Loop - Empty Configuration - Small Rounded Pulse Input




Figure E.4: Ten DOF Open-Loop - Empty Configuration - Medium Rounded Pulse Input




Figure E.5: Ten DOF Open-Loop - Empty Configuration - Large Rounded Pulse Input




Figure E.6: Ten DOF Open-Loop - Empty Configuration - Huge Rounded Pulse Input




Figure E.7: Ten DOF Open-Loop - Loaded Configuration - Noise Input




Figure E.8: Ten DOF Open-Loop - Loaded Configuration - Tiny Rounded Pulse Input




Figure E.9: Ten DOF Open-Loop - Loaded Configuration - Small Rounded Pulse Input




Figure E.10: Ten DOF Open-Loop - Loaded Configuration - Medium Rounded Pulse Input




Figure E.11: Ten DOF Open-Loop - Loaded Configuration - Large Rounded Pulse Input




Figure E.12: Ten DOF Open-Loop - Loaded Configuration - Huge Rounded Pulse Input

Appendix F. Two DOF Active Open-Loop Responses

Figure F.1 - F.6 show the open-loop active responses of the two DOF system in the empty configuration. Figures F.7 - F.12 show the same for the loaded configuration.




Figure F.1: Two DOF Open-Loop - Empty Configuration - Noise Input




Figure F.2: Two DOF Open-Loop - Empty Configuration - Tiny Rounded Pulse Input




Figure F.3: Two DOF Open-Loop - Empty Configuration - Small Rounded Pulse Input




Figure F.4: Two DOF Open-Loop - Empty Configuration - Medium Rounded Pulse Input




Figure F.5: Two DOF Open-Loop - Empty Configuration - Large Rounded Pulse Input




Figure F.6: Two DOF Open-Loop - Empty Configuration - Huge Rounded Pulse Input




Figure F.7: Two DOF Open-Loop - Loaded Configuration - Noise Input




Figure F.8: Two DOF Open-Loop - Loaded Configuration - Tiny Rounded Pulse Input




Figure F.9: Two DOF Open-Loop - Loaded Configuration - Small Rounded Pulse Input




Figure F.10: Two DOF Open-Loop - Loaded Configuration - Medium Rounded Pulse Input




Figure F.11: Two DOF Open-Loop - Loaded Configuration - Large Rounded Pulse Input




Figure F.12: Two DOF Open-Loop - Loaded Configuration - Huge Rounded Pulse Input

Appendix G. Mixed-Norm Design 2 - Empty Configuration Simulations

Figure G.1 - G.6 show the closed-loop responses of the two DOF system in the empty configuration. The controller used is the Mixed-Norm Design 2 controller.




Figure G.1: Mixed-Design 2 - Empty Configuration - Noise Input




Figure G.2: Mixed-Design 2 - Empty Configuration - Tiny Rounded Pulse Input




Figure G.3: Mixed-Design 2 - Empty Configuration - Small Rounded Pulse Input




Figure G.4: Mixed-Design 2 - Empty Configuration - Medium Rounded Pulse Input




Figure G.5: Mixed-Design 2 - Empty Configuration - Large Rounded Pulse Input




Figure G.6: Mixed-Design 2 - Empty Configuration - Huge Rounded Pulse Input

Appendix H. Mixed-Norm Design 2 - Loaded Configuration Simulations

Figure H.1 - H.6 show the closed-loop responses of the two DOF system in the loaded configuration. The controller used is the Mixed-Norm Design 2 controller.




Figure H.1: Mixed-Design 2 - Loaded Configuration - Noise Input




Figure H.2: Mixed-Design 2 - Loaded Configuration - Tiny Rounded Pulse Input




Figure H.3: Mixed-Design 2 - Loaded Configuration - Small Rounded Pulse Input




Figure H.4: Mixed-Design 2 - Loaded Configuration - Medium Rounded Pulse Input




Figure H.5: Mixed-Design 2 - Loaded Configuration - Large Rounded Pulse Input




Figure H.6: Mixed-Design 2 - Loaded Configuration - Huge Rounded Pulse Input

Appendix I. Mixed-Norm Design 2 - Ten DOF Medium Configuration Simulations

Figure I.1 - I.6 show the closed-loop responses of the ten DOF system in the medium configuration. The controller used is the Mixed-Norm Design 2 controller. The legend for all figures is

Style	Acceleration	Deflection
-	Cabin Vertical	Front
--	Cabin Rotational	Rear
---	Trailer Rotational	Trailer




Figure I.1: Mixed-Design 2 - Ten DOF Medium Configuration - Noise Input




Figure I.2: Mixed-Design 2 - Ten DOF Medium Configuration - Tiny Rounded Pulse Input




Figure I.3: Mixed-Design 2 - Ten DOF Medium Configuration - Small Rounded Pulse Input




Figure I.4: Mixed-Design 2 - Ten DOF Medium Configuration - Medium Rounded Pulse Input




Figure I.5: Mixed-Design 2 - Ten DOF Medium Configuration - Large Rounded Pulse Input




Figure I.6: Mixed-Design 2 - Ten DOF Medium Configuration - Huge Rounded Pulse Input

Bibliography

- [CRS95] R. Canfield, B. Ridgely, and L. Smith. Numerical solution to the fixed order, general mixed H_2/H_∞ control problem. Submitted to *Automatica*, 1995.
- [Dai90] R. L. Dailey. Lecture notes for the workshop on H_∞ and μ methods of robust control, 1990. In conjunction with the *American Control Conference*, San Diego, CA.
- [DDB95] M. A. Dahleh and I. J. Diaz-Bobillo. *Control of Uncertain Systems*. Prentice-Hall, Englewood Cliffs, NJ, 1995.
- [DFT92] J. C. Doyle, B. A. Francis, and A. R. Tannenbaum. *Feedback Control Theory*. Macmillan, 1992.
- [dJ95] B. de Jager. Multi-objective suspension control problem. Submitted to the *1995 IEEE Conference on Decision and Control*, 1995.
- [Jac95] D. R. Jacques. *Optimal Mixed-Norm Control Synthesis for Discrete Linear Systems*. PhD thesis, Air Force Institute of Technology, June 1995.
- [JRC95] D. Jacques, B. Ridgely, and R. Canfield. Discrete-time, mixed-norm control synthesis applied to aircraft terrain following. Submitted to the *AIAA Journal of Guidance, Control and Dynamics*, 1995.
- [JRCS95] D. R. Jacques, D. B. Ridgely, R. A. Canfield, and M. A. Spillman. A MATLAB toolbox for fixed-order, mixed-norm control synthesis. In *1995 IEEE Conference on Control Applications*, pages 470–475, Albany, NY, 1995.
- [KG93] S. M. Karamihas and T. D. Gillespie. Characterizing trucks for dynamic load prediction. *Heavy Vehicle Systems, Int. J. of Vehicle Design*, 1(1):3–19, 1993.
- [MAT] MATLAB: High performance numeric computation and visualization software. The Math Works, Inc., Natick MA, 1994.
- [Rid92] D. B. Ridgely. *A Nonconservative Solution to the General Mixed H_2/H_∞ Optimization Problem*. PhD thesis, Massachusetts Institute of Technology, 1992.
- [RW95] B. Ridgely and D. Walker. The general mixed H_2/H_∞ control problem. Submitted to *Automatica*, 1995.
- [SLC89] M. G. Safonov, D. J. N. Limebeer, and R. Y. Chiang. Simplifying the H_∞ theory via loop-shifting, matrix-pencils and descriptor concepts. *International Journal of Control*, 50(6):2467–2488, 1989.
- [Smi94] L. D. Smith. Improved numerical solution of mixed H_2/H_∞ with applications. Master's thesis, Air Force Institute of Technology, December 1994.
- [Spi94] M. Spillman. Applications of ℓ_1 and mixed H_2/ℓ_1 optimization. Master's thesis, Air Force Institute of Technology, December 1994.
- [Wal94] D. E. Walker. *H_2 Optimal Control with H_∞ , μ , and L_1 Constraints*. PhD thesis, Air Force Institute of Technology, June 1994.
- [Zam66] G. Zames. On the input-output stability of time-varying nonlinear feedback systems, Part I: Conditions derived using concepts of loop gain, conicity, and positivity. *IEEE Trans. Auto. Control*, AC-11:228–238, April 1966.

Vita

Captain John B. Allison was born [REDACTED] 25 November 1968 in [REDACTED] Morro Bay, California. He graduated from Morro Bay High School, Morro Bay, California in 1987 and subsequently attended the University of Southern California on an AFROTC scholarship. In May 1991, he received a Bachelor of Science in Aerospace Engineering and was commissioned in the USAF. Capt Allison began active duty service in October 1991 at Wright-Patterson AFB, Ohio. He was assigned to the Aeronautical System Center, working at the Propulsion System Program Office as a logistical engineer and cost analyst. He entered the Air Force Institute of Technology's Graduate School of Engineering in May 1994 and, upon graduation, will be assigned to Phillips Laboratory, Kirtland AFB, New Mexico.

Permanent address: 1416 Las Encinas Dr.
Los Osos, CA 93402

VITA-1

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE December 1995		3. REPORT TYPE AND DATES COVERED Master's Thesis
4. TITLE AND SUBTITLE APPLICATION OF MIXED-NORM OPTIMAL CONTROL TO A MULTI-OBJECTIVE ACTIVE SUSPENSION PROBLEM			5. FUNDING NUMBERS	
6. AUTHOR(S) John B. Allison Captain, USAF				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Institute of Technology Wright-Patterson AFB OH 45433-6583			8. PERFORMING ORGANIZATION REPORT NUMBER AFIT/GAE/ENY/95D-02	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Dr. Marc Jacobs AFOSR/NM 110 Duncan Ave, Suite B115 Bolling AFB DC 20332-0001			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution unlimited			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) Abstract Mixed-norm optimal control synthesis is used to solve a multi-objective suspension problem. The objective is to develop a controller for an active suspension system on-board a tractor-semitrailer vehicle. The problem is first approached by using H_2 and H_∞ optimization. It is shown that by combining both techniques into one mixed-norm optimization method, it is possible to exploit the strengths of each norm to provide superior performance. Two H_2/H_∞ designs were completed. One design concentrated on optimal performance at one design condition. The second design concentrated on providing the best performance possible at a medium load configuration, while maintaining robust stability at the extreme load configurations.				
14. SUBJECT TERMS Control Theory, H2 Optimization, Hinf Optimization Mixed H2/Hinf Optimization, Active Suspensions			15. NUMBER OF PAGES 202	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

GENERAL INSTRUCTIONS FOR COMPLETING SF 298

The Report Documentation Page (RDP) is used in announcing and cataloging reports. It is important that this information be consistent with the rest of the report, particularly the cover and title page. Instructions for filling in each block of the form follow. It is important to *stay within the lines* to meet *optical scanning requirements*.

Block 1. Agency Use Only (Leave blank).

Block 2. Report Date. Full publication date including day, month, and year, if available (e.g. 1 Jan 88). Must cite at least the year.

Block 3. Type of Report and Dates Covered. State whether report is interim, final, etc. If applicable, enter inclusive report dates (e.g. 10 Jun 87 - 30 Jun 88).

Block 4. Title and Subtitle. A title is taken from the part of the report that provides the most meaningful and complete information. When a report is prepared in more than one volume, repeat the primary title, add volume number, and include subtitle for the specific volume. On classified documents enter the title classification in parentheses.

Block 5. Funding Numbers. To include contract and grant numbers; may include program element number(s), project number(s), task number(s), and work unit number(s). Use the following labels:

C - Contract	PR - Project
G - Grant	TA - Task
PE - Program Element	WU - Work Unit Accession No.

Block 6. Author(s). Name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. If editor or compiler, this should follow the name(s).

Block 7. Performing Organization Name(s) and Address(es). Self-explanatory.

Block 8. Performing Organization Report Number. Enter the unique alphanumeric report number(s) assigned by the organization performing the report.

Block 9. Sponsoring/Monitoring Agency Name(s) and Address(es). Self-explanatory.

Block 10. Sponsoring/Monitoring Agency Report Number. (If known)

Block 11. Supplementary Notes. Enter information not included elsewhere such as: Prepared in cooperation with...; Trans. of...; To be published in.... When a report is revised, include a statement whether the new report supersedes or supplements the older report.

Block 12a. Distribution/Availability Statement. Denotes public availability or limitations. Cite any availability to the public. Enter additional limitations or special markings in all capitals (e.g. NOFORN, REL, ITAR).

DOD - See DoDD 5230.24, "Distribution Statements on Technical Documents."

DOE - See authorities.

NASA - See Handbook NHB 2200.2.

NTIS - Leave blank.

Block 12b. Distribution Code.

DOD - Leave blank.

DOE - Enter DOE distribution categories from the Standard Distribution for Unclassified Scientific and Technical Reports.

NASA - Leave blank.

NTIS - Leave blank.

Block 13. Abstract. Include a brief (*Maximum 200 words*) factual summary of the most significant information contained in the report.

Block 14. Subject Terms. Keywords or phrases identifying major subjects in the report.

Block 15. Number of Pages. Enter the total number of pages.

Block 16. Price Code. Enter appropriate price code (*NTIS only*).

Blocks 17. - 19. Security Classifications. Self-explanatory. Enter U.S. Security Classification in accordance with U.S. Security Regulations (i.e., UNCLASSIFIED). If form contains classified information, stamp classification on the top and bottom of the page.

Block 20. Limitation of Abstract. This block must be completed to assign a limitation to the abstract. Enter either UL (unlimited) or SAR (same as report). An entry in this block is necessary if the abstract is to be limited. If blank, the abstract is assumed to be unlimited.