

Autonomous Speed Control for Conventional Vehicles

R. Craig Coulter

CMU-RI-TR-96-18

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
in the field of Robotics

The Robotics Institute
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, Pennsylvania 15213

March 7, 1996

© 1996 by R. Craig Coulter. All rights reserved.

This research was sponsored by ARPA under contracts "Perception for Outdoor Navigation" (contract number DACA 76-89-C-0014, monitored by the US Army Topographic Engineering Center) and "Unmanned Ground Vehicle System" (contract number DAAE07-90-C-RO59, monitored by TACOM). Views and conclusions contained in this document are those of the author and should not be interpreted as representing official policies, either expressed or implied, of ARPA or the United States Government.



Robotics

Thesis

Autonomous Speed Control for Conventional Vehicles

R. Craig Coulter

Submitted in Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy
in the field of Robotics

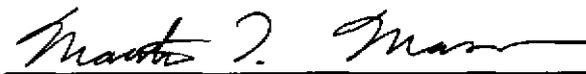
ACCEPTED:



William L. Whittaker Thesis Committee Chair

29 April 1996

Date



Matthew T. Mason Program Chair

May 6, 1996

Date



Raj Reddy Dean

5/6/96

Date

APPROVED:



Paul Christiano Provost

9 May 1996

Date

DEDICATION

In the spring of 1993, near the midpoint of my Ph.D. program, both of my parents lost their jobs to economic downsizing on the very same day. Within two months, my mother fell very ill and would remain so for the better part of a year. My father took part time jobs in construction and excavation to help cover expenses, but it soon became apparent that such stop-gap measures were insufficient. It seemed likely that I would be forced to leave school to support my family.

My father refused to allow me to drop out of the program except as a last resort, and so to provide enough income to make it through the coming winter, he proposed that we cut and sell timber from our farm. That fall, until the first blizzard made the logging roads impassible, we cut poplar and cherry. Each weekend we felled trees, and cut and cleared the logs. During the week while I studied at school, he would load the logs and sell them at a local sawmill. This cycle continued for a four month period and together, we cut enough to meet expenses until spring, when my father found another job. It was a tough time, but we are better for it...

This dissertation is dedicated to my father, Robert C. Coulter, as a concrete realization of his dedication to his children's education, and his belief that through hard work and uncompromising commitment in the face of adversity, a man may better himself and raise his station in life...



ABSTRACT

The complete automation of conventional vehicles requires the automation of each of their degrees of freedom. Most work in mobile robotics research has concentrated on the development of so-called *autonomous navigation systems* - software systems capable of steering mobile robots safely through environments of interest. The problem of controlling the speed of the mobile robot has been of little interest in the field until the recent application of automation technologies to conventional vehicles. Researchers in the field noted the difficulty of performing servo control on such machines, especially at the low speeds and over the various inclinations which typify operation in rugged, unstructured, 3-D terrain. The standard approaches implemented in the automotive industry to perform cruise control (proportional control about velocity) were shown to produce poor performance when adapted to such applications. Acceptable performance is attainable through the implementation of lookup table approaches, though at a significant calibration cost. The need for a robust controls technology is central to the automation of such low speed applications as cross country driving, off-road haulage, and the operation of heavy equipment such as loaders and scrapers. In addition, predictable acceleration performance at highway speeds has future application for lane merging and passing in automated highway systems.

This thesis addresses the problem of engineering an autonomous speed control system for a conventional vehicle operating in rugged, unstructured, 3-D terrain. This work produces two new algorithms which enable robust autonomous speed control. LITHIA is a robust servo-control algorithm formulated and implemented in the power or energy domain. LITHIA enables robust, fast command following of speed profiles even in the presence of the erratic, large-magnitude force disturbances associated with rough terrain navigation. LITHIA may be extended, in a straightforward manner, to perform optimal control / speed planning using the well understood feedforward planning method, where costs are based on both servo-control performance and vehicle stability. Extension of vehicle stability cost calculations to include the dynamic response of the suspension is enabled by a computationally efficient motion evaluation algorithm, called SPECTRE. This new algorithm operates in the frequency space, as opposed to the traditional time-space, enabling significantly faster real time motion evaluation.

This thesis offers two significant contributions to the mobile robotics community. First, the demonstrated solution to the problem of attaining stable, robust speed control of conventional vehicles through the formulation of control laws in the energy or power domain is both novel and unprecedented. The thesis will show significant performance improvements through experimentation, as well as develop the physical and mathematical theory that details the merits of such formulation. Second, the development and implementation of a frequency-space solution to the motion evaluation problem is both novel and unprecedented. The thesis shows that significant computational advantage is had through the implementation of a filter model of vehicle dynamics that eliminates the need for on-line, real-time ODE integration.

<u>PART I - PRELIMINARIES</u>	<u>6</u>
CHAPTER I - INTRODUCTION	7
1. Background	7
2. A Summary of the Autonomous Speed Control Problem	8
3. Results and Contributions	10
4. Acknowledgments	11
CHAPTER II - PRIOR WORK	13
5. Historical Perspective	14
6. Autonomous Speed Control - Prior Work	19
7. Motion Planning - Prior Work	21
8. Motion Evaluation - Prior Work	21
9. Automotive Modelling - Related Work / Perspective	22
10. Simulation Technologies	22
11. Conclusions	23
CHAPTER III - LINEAR CONTROL EXPERIMENTS	25
12. Experimental Procedure	25
13. Conclusions	30
<u>PART II - MODELLING</u>	<u>31</u>
CHAPTER IV - SYSTEM MODELS	32
14. Externally Applied Forces	32
15. Reaction Forces	35
16. The Newton Euler Equations	35
17. Insights from the System Model	40
CHAPTER V - DISTURBANCE FORCE ANALYSIS	41
18. Representations of Motion	41
19. Relative Representations	42
20. Relative Strength Analysis	44
21. Relative Impulse Analysis	45
22. Relative Work Analysis	47
23. Problem-Specific Calculations	48
24. Discussion and Conclusions	55
CHAPTER VI - CONTROL FORCE MODELS	57
25. Propulsion Force	57
26. Torque Converter	60
27. The Combined Powertrain Equations	61
28. Total Dynamical System and Compensation	62
29. State Space Analysis	65
30. Braking Force	67

<u>PART III - AN AUTONOMOUS SPEED CONTROL SYSTEM</u>	<u>70</u>
CHAPTER VII - LITHIA: A MODEL BASED CONTROLLER	71
31. The LITHIA Control Algorithm - First Generation	71
32. Experimentation - First Generation LITHIA	74
33. Deconstructing the Algorithm	83
34. Redesigning LITHIA using Classical Control Techniques	87
35. Distinguishing LITHIA from so-called "Linear Techniques"	94
36. Brake Servo Control	99
37. The Switching Problem	100
38. Optimization	102
39. Feedforward Optimal Control	105
40. Summary	107
CHAPTER VIII - SPECTRE: A FREQUENCY SPACE MOTION EVALUATION SYSTEM	108
41. Modelling	108
42. Development of a Simple Suspension Model	111
43. Solution of the Time Domain Problem	113
44. Time Complexity Analysis	114
45. SPECTRE	120
46. Experimentation	121
47. Summary	123
<u>PART IV - CONCLUSIONS</u>	<u>124</u>
CHAPTER IX - SUMMARY	125
48. Results and Contributions	126
49. Technical Perspectives	128
50. Future Research and Work	132
51. General Perspectives	132
<u>PART V - APPENDICES</u>	<u>134</u>
APPENDIX I - KINEMATIC TRANSFORMS	135
52. Definition of the Kinematic Transformations	135
APPENDIX II - DISTURBANCE REDUCTION	138
53. Relative Analysis - the Case for Weight Reduction	138
54. Cost Functions and Performance Metrics for Reconfiguration	139
55. The Power Weight Spiral	142
56. Application to the NavLab II Mobile Robot	145
57. Concept Configuration Designs	147
58. References	153

Part I

Preliminaries

Part I of this thesis contains background materials necessary and relevant to the main corpus of work. A brief *introduction* to the nature of the problem and its key terminology is provided. A *prior and related work* section ties the development of speed control and related technologies to the historical development of complete autonomous mobile robot systems. Finally, a set of *experiments* previously performed by other researchers is recreated, as the results had neither been recorded nor published.

CHAPTER I - INTRODUCTION

1. BACKGROUND

This thesis considers the problem of controlling the speed of an autonomous conventional vehicle driving in rugged unstructured terrain. To begin, let's define some of the background terms relevant to this problem:

1.1 The Rough Terrain Autonomous Navigation Problem

Autonomous navigation is the problem of producing an intelligent, independently mobile machine. This is the core problem addressed by the mobile robotics field over its 30 year span. It has been solved repeatedly, for a variety of environments, using essentially the same approach - the coupling of intelligent sensing and planning to an electronically controlled machine. The *rough terrain navigation problem* is autonomous navigation applied to unstructured, natural terrain, usually peppered with various sized discrete obstacles, and replete with natural barriers, such as trenches, cliffs, and streams. The general problem can be subdivided into two problems - longitudinal and lateral. The lateral problem is frequently itself called autonomous navigation. The longitudinal problem is called *autonomous speed control*.

1.2 The Autonomous 3-D Speed Control Problem

This thesis considers the problem of *autonomous speed control* over rugged, unstructured terrain. Speed control is the problem of servoing or regulating the speed of the vehicle in the presence of external disturbance forces. The 3-D problem is distinguished by wide and sometimes sudden changes in pitch over the length of the path, and a rugged surface, sufficient to excite the suspension dynamics of the vehicle. The autonomous problem is distinguished by the need for a planning agent to determine safe and / or appropriate speeds, given the context of the operations. The inclusion of a planning agent in the system is driven by the complicated interaction between the speed of the machine and its bodily reaction to the terrain. Because the stability of the machine is a complicated function of speed and terrain inputs, it is insufficient (or at least inefficient) to select a single traversal speed that always meets the minimal stability criteria.

1.3 Conventional Vehicles

This thesis applies 3-D speed control techniques to a *conventional vehicle* - a retro-fitted military vehicle called a HMMWV. A conventional vehicle is an Ackerman steered vehicle, powered by an internal combustion engine, driven by an automatic transmission, and supported by a passive dynamic suspension. It is especially important to note that the application of these technologies to conventional vehicles is not a limitation, but an extension of the problem. Conventional vehicles represent a more complicated, more difficult controls problem due to their inherently non-linear, coupled dynamics.

2. A SUMMARY OF THE AUTONOMOUS SPEED CONTROL PROBLEM

In the case of conventional vehicles, the lateral degree of freedom has been addressed for both on-road and off-road applications. Such research has principally been concerned with problems of perception and motion planning - where motion planning usually refers to the planning of spacial paths. To achieve *full* autonomy of a conventional mobile robot, one must achieve autonomy of each of the degrees of freedom. This thesis concerns itself with the other longitudinal degree of freedom - the autonomous speed control problem. This area has only recently presented interesting research problems for a number of reasons. First, most previous mobile robots simply weren't difficult to control. Large classes of machines, such as electrically powered indoor mobile robots operating on flat floors are dynamically well behaved. Second, until recently, most mobile robots were not required to perform control tasks that are physically difficult to achieve. For example, most road following mobile robots operate at highway speeds, and are required only to maintain that speed - a problem that is easily solved by the cruise control. The recent combination of several requirements has produced a more challenging speed control problem:

- Use of conventional vehicles.
- Operations on rugged terrain.
- Operations at low speed.
- Cooperation with a spacial planning system.
- Requirement to simultaneously ensure stability.

2.1 Importance of the Problem

Conventional vehicles represent the largest class of work vehicles in the world. The automation of such industries as forestry, mining, haulage and agriculture is predicated on the ability to achieve autonomy of the individual work vehicles. The solution of the autonomous speed control problem for the operation of conventional vehicles in unstructured 3-D terrain is key to the automation of these, and other, like industries.

2.2 Principal Issues

Throughout the development of this thesis, the fundamental problems that face autonomous speed controller design will be discussed and developed. However, it may be useful to keep in mind this terse summary of the issues. Three principal issues describe the problem of motion controller design for conventional vehicles performing rough terrain navigation:

Controllability Issue:

The design of motion controllers for rough terrain navigation is difficult because the disturbance forces are of the same order of magnitude as the powertrain propulsion force, while the momentum of the vehicle is low.

Observability Issue:

The use of conventional vehicles aggravates the throttle compensation problem by introducing non-linear, non-observable dynamics. Two coupled states are present, only one of which is normally observable.

Stability Issue:

The response of the vehicle's passive dynamics is a function of velocity. It is this issue that couples motion controller and motion evaluator design.

2.3 Summary of the Solution

This thesis will demonstrate that prior attempts at servo-control produced poor performance due to an unrealistic representation of the physics of the system. In this work, performance is found not to be an issue of control, but one of proper formulation and representation of system dynamics. Accordingly, the dissertation will provide detailed derivations of the physical models used in servo-controller design. The basic result may be simply stated as follows. Traditionally applied proportional control produces an *oscillatory class of behavior*, while properly applied proportional control (as demonstrated in this thesis) produces an *asymptotic class of behavior*.

2.4 Overview of the Dissertation

This dissertation is divided into five parts. **Part I** contains background materials necessary and relevant to the main corpus of work. A brief *introduction* to the nature of the problem and its key terminology is provided. A *prior and related work* section ties the development of speed control and related technologies to the historical development of complete autonomous mobile robot systems. It was necessary to recreate *experiments* previously performed in the field.

Part II develops the models and performs a set of analyses that are used to constrain the design process. The chapter entitled *System Models* develops the physical mechanical model used throughout the thesis. Models and analyses of disturbance forces and control forces are developed in the next two chapters. It was necessary (or perhaps just useful) to develop a set of comparison metrics for various physical mechanics phenomena - these relative analyses provide insight in the control design process. In **Part III**, the autonomous speed control problem is partitioned into two areas - integrated motion control (which may be extended to optimal servo-control), and motion evaluation. These two chapters discuss the design, development and implementation of LITHIA and SPECTRE - the software systems that respectively address each problem. **Part IV** contains a summary, conclusions and some discussion. **Part V** contains a set of appendices including a chapter on *design for disturbance rejection* which is addressed as a mechanical precursor to servo-controls algorithm design.

3. RESULTS AND CONTRIBUTIONS

3.1 Results

This thesis has produced *a novel energy-based servo-controls algorithm* which has demonstrated superior command following and disturbance rejection performance over all previously known algorithms. This energy or power-based algorithm has enabled robust, stable speed control over rugged, unstructured 3-D terrain for extensive distances while maintaining asymptotic convergence. The asymptotic behavior of the energy-based algorithm is not merely a performance improvement over the traditional methods, but is shown to be an entirely new class of response behavior, nearly identical to the response of standard second order systems. Experimentation demonstrates speed control performance that is as good as, or better than, that attainable by a human driving the same terrain.

This thesis has also developed *an optimal speed planner*, based on the feedforward method of motion planning. Speed planning is shown to be analogous to motion planning, with sufficient similarity to allow the combination of speed and motion planning into a single optimal planner. No specific claims are made of this work, other than the demonstration that the autonomous speed control problem may be formulated in a manner entirely analogous to the traditional autonomous navigation problem. Such a formulation leads in a straightforward manner to the eventual integration of autonomous speed control and autonomous navigation systems.

Finally, this thesis has produced *a novel frequency-space motion evaluation methodology* that eliminates the need for expensive real time ODE integration. Extensive cost comparisons are provided which demonstrated a minimum of one order of magnitude difference in run time cost. The most compelling attribute of this new methodology is its constant run-time cost in the face of increasing model complexity and / or improved model resolution. The mathematics of the method are developed and a set of frequency-space images of obstacles are obtained from laser range finder images.

3.2 Principal Contributions

There are two classes of contributions made in this thesis - fundamental and demonstrative. Fundamental contributions are both novel and unique in the field, and have a full grounding in either physical or mathematical theory. Demonstrative contributions lack such grounding, but offer instead a demonstration of either the applicability of a previously untried method, or the importance of an unnoticed phenomena. This thesis offers two fundamental contributions, and two demonstrative contributions.

3.2.1 Linearity of the Energy Domain

This thesis makes a fundamental contribution by proving that linear servo control techniques are appropriate to the throttle control portion of the autonomous speed control problem when the system equations of motion are formulated in the system's *energy space*, rather than the traditional *velocity space*. *This phenomena is shown to arise from the necessity of enforcing a differential equation compatibility between the First Law of Thermodynamics (the governing principal of the*

system forcing function) and Newton's Second Law (the governing principal of the characteristic system equation). This has the effect of:

1. Allowing the full application of linear control techniques.
2. Allowing the formulation of a control law that produces a standard second order transfer function. Such a formulation allows the use of fundamental design techniques such as pole placement.

3.2.2 Frequency Space Obstacle Detection

This thesis makes a second fundamental contribution by developing an alternate, frequency-space, solution to the motion evaluation problem. This new method is shown to be of constant order in model complexity, whereas the time space formulation is guaranteed to be of non-constant order. In even the simplest suspension modelling case, the method is shown to produce a one order of magnitude improvement in run time.

3.2.3 Optimal Control Application

This thesis demonstrates that the classical LQ optimal control method can be applied to the autonomous speed control problem when formulated in the linear energy space rather than velocity space. Questions of the appropriateness of such an application arise, which leads to the ultimate rejection of the method in favor of the so-called action space planning, or feedforward optimal control method. Such a formulation demonstrates the compatibility of integrating autonomous speed control and autonomous navigation problems in a single optimal control planner.

3.2.4 Disturbance Reduction Functions

This thesis demonstrates that the reduction of the magnitude of the disturbance forces through mechanical means is tightly tied to the thermodynamic efficiency of the retro-fit design, as expressed by the Claussius statement of the Second Law of Thermodynamics.

4. ACKNOWLEDGMENTS

This thesis took its form over a three year period in which I worked very closely with Alonzo Kelly, who was, at the time, a fellow graduate student in the RI Ph.D. program. Al had come back to graduate school after some 15 years of industrial experience as an engineer and had a keen insight into the nature of physical problems, which he willingly shared with the rest of the members of our research group. It is unfortunate that he is so quiet and retiring that many of those with whom he worked never came to appreciate the full measure of his abilities. In many ways, Al acted as a mentor during this period. He taught me how to think like an engineer, how to stand outside of accepted or assumed solutions, and most importantly, he forced me into the practice of striving to reduce a problem to a few fundamental statements of mathematics or physics - to seek to understand the nature of the problem, not merely its instance. For these reasons, I have to acknowledge Al Kelly as the single most influential contributor to this work.

I want to thank my advisors, Red Whittaker and Tony Stentz, for supporting me throughout this

work. Red and Tony permitted me the academic freedom that is necessary to the production of such a theoretical thesis. Without their "hands-off" management style, I would likely have succumbed to that empty narcosis of writing software and performing experiments for the sake of showing your advisor something rather than being permitted the time necessary to just sit and think about the problem in depth and detail. The work benefitted immeasurably from their patience and their ability to distinguish activity from productivity.

Part of this work arose from an engineering effort that surrounded the reconfiguration of the NavLab II. I had the pleasure of leading a team of what proved to be a group of highly motivated, highly capable engineers. The team that performed the electro-mechanical reconfiguration included George Mueller, Jeremy Armstrong, Alex MacDonald, Joed Haddad, Jim Frazier, Jim Moody, Bill Ross, Joe Oliver, Travis Schluessler, Bob Shuska, Henein Simon and random others from the Robotics Institute who donated their time and energy - they were much appreciated. Jim Frazier deserves special mention for his unceasing loyalty to the improvement of the NavLab II. Mr. Frazier was one of the last remaining members of the original NavLab team and was key in explaining the reasoning behind the original design as well as suggesting methods for improving it.

Dave Pahnos has been a good and constant friend to me throughout my graduate studies, providing perspective at those times when I lost my own. I appreciate his repeated counsel to finish the dissertation, whatever the cost, and his reminder that I was unlikely to find a suitable wife if I spent all of my time at the Field Robotics Center.

Finally, and most importantly, I would like to thank my most suitable wife Cynthia, who I found far from the FRC, and my family - especially my father, Robert C. Coulter, to whom this dissertation is dedicated. Without Cynthia I would surely have gone mad long ago, and without my father's tenacity, prudence, and sheer bold and angry determination, the economic downturn of 1993 would have forced me to leave graduate school and seek employment to support our family. He will never really understand the depth of my appreciation.

CHAPTER II - PRIOR WORK

The autonomous speed control problem is a *recent addition* to autonomous field robotics research. Most related work in the field has concentrated on the development of autonomous software systems capable of planning and executing paths through increasingly *challenging terrain and / or* less structured environments. The problem of determining where the machine should go took precedence over the problem of determining how fast the machine should travel - and with good reason. The autonomous speed control problem only recently arose as a challenge to the field for three basic reasons:

1. Previous mobile robot systems have not operated on machines that have *complicated powertrain transfer functions*. Control signal to output mappings have generally been controllable through linear techniques, either through *coincidence* (the propulsion system happened to be electrical, with a good linear transfer function mapping) or by construction (engineers replaced entire sections of the powertrain to produce a good linear mapping). The decision to use conventional vehicles forced the acceptance of a less than easily controllable dynamic system.
2. Mobile robots have not operated under conditions where *disturbances prove significant* to the control of the machine. Prior operations have been primarily either *flat-floor (or outdoors, level terrain)* with a few, discrete obstacles. Avoidance, rather than mounting, of obstacles was the preferred navigation mode.
3. Many applications in mobile robotics did not require that *speed be servoed*. Road-following, at least as currently exhibited in research, requires only that the machine maintain a nominal cruising speed. There was also no need for a speed planning agent - the human selected the speed, the robot merely achieved and maintained the speed. For a large segment of problems fitting this description the conventional cruise control was a sufficient solve.

The extension of autonomous systems based on conventional vehicles to unstructured 3-D terrain has challenged current speed control technologies on all of these fronts simultaneously. What has emerged is essentially a *servo control problem of a non-linear system in the presence of large magnitude disturbances*. The problem has additional unique characteristics which largely follow from the requirements of autonomy. Autonomy calls for a *planning agent capable of optimizing* the selection of motions against some predetermined cost function. In this case, the robot must select a speed profile such that it provides for its own physical safety while minimizing command-following error, and perhaps meeting some additional pre-selected constraints (such as a maximum speed limit).

The application of autonomous speed controls technologies to conventional vehicles operating in rugged terrain also gives rise to some tertiary technical problems, such as in motion evaluation. Current motion evaluators are computationally too inefficient to perform the feedforward motion simulations required to predict suspension responses for the vehicle. As part of this thesis, a new

paradigm of motion evaluators was explored, which proved fruitful for operations in rugged terrain.

This chapter provides an overview of prior and related work. A historical perspective is offered for the benefit of readers not familiar with the progression of autonomous navigation technologies, with the intent to show the convergence of the conditions that call for more capable speed control technologies. The section is fairly general in scope and is not especially intended to bolster the case for adopting or rejecting the content of this dissertation. The remainder of the chapter is devoted to the discussion of a variety of prior works. Given the nascence of autonomous speed control technologies, the bulk of the remaining material concerns related rather than prior works. Related works, in this context, are solutions to similar or analogous problems that proved useful to the construction of the solutions presented in this thesis.

5. HISTORICAL PERSPECTIVE

5.1 Seminal Work - Stanford Research Institute

Mobile robotics research began in earnest with the development of Shakey at the Stanford Research Institute (SRI), Nilsson[24]. This mobile robot (or mobile automaton as Nilsson called them) was constructed as a test-bed for the application of artificial intelligence techniques. Shakey provided a proof of concept demonstration of the possibilities of autonomous navigation coupled with higher level reasoning, such as box-herding.¹ SRI approached the problem from the classical AI perspective, extending and adapting then-current vision and planning methods from AI and AI-based computer vision. The planning system was an extension of a theorem prover called QA-3, in which the state of the world and the desired goal state were represented in first-order predicate calculus. A motion plan was developed by “proving” the goal state through the application of rules based on low-level motion commands. The vision problem was formulated as an *image understanding problem*. The scene was reduced to a collection of edge lines, whose possible relationships were considered by an *object identification* program. The possible scene interpretations were then represented as a hypothesis in a decision tree.

SRI’s goal was to produce an automaton capable of a wide variety of tasks, rather than a few specific tasks. To quote Nilsson[24], “The main theme of the research is the integration of the necessary planning systems, models of the world, and sensory processing systems into an efficient whole capable of performing a wide range of tasks in a real environment.” They believed that generality required general methods, and concluded that the way to accomplish this generality was to integrate into one system “many of the abilities that are usually found separately in individual Artificial Intelligence programs.” (Nilsson[24]). Motion control was clearly not of significant interest to this work beyond the need to provide gross motion in flat floor environments. The SRI group shaped the mobile robotics navigation problem by addressing questions basic to the technology:

1. Shakey was capable of finding and herding a number of cardboard boxes into a central location in the room.

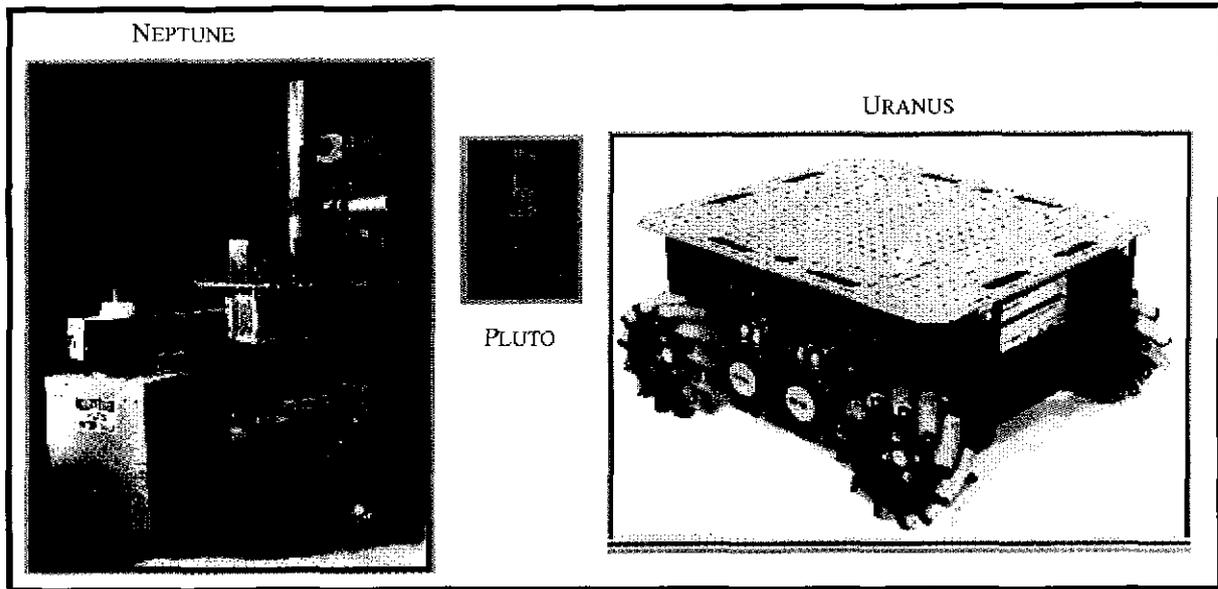
1. Established the required component technologies.
2. *Demonstrated how they should be connected to produce autonomous navigation.*
3. Demonstrated that the Artificial Intelligence approach led to navigation.
4. Established the precedent for testing real robots in the real world, stating that, “the feasibility of a system... for controlling a real robot can only be confronted by actually attempting to control a real robot with such a system.”

5.2 Stanford Artificial Intelligence Laboratory

The Stanford Artificial Intelligence Laboratory produced the Stanford Cart (Moravec[21]), a four wheel mobile robot used for both indoor and outdoor navigation in cluttered environments. The principle objective of the research surrounding this cart was to test new stereo vision techniques and obstacle-avoiding path-planning techniques. The system proved to be reliable, but painfully slow, moving about 20 meters over 5 hours in what Moravec[21] described as “lurches.” In addition to the innovative research work done in vision and planning, Moravec provided a number of key insights into some of the problems of integration in mobile robotics, such as the inability of the position estimation system to maintain accurate enough track of the mobile robot. This caused the vision system to find already located obstacles a second time, effectively filling the map with imaginary obstacles. The position estimation problem was, in turn, coupled to the accuracy of the control system. The cart frequently produced significant motion errors when executing “lurches.” Motion control technologies were not of significant interest in the scope of this work, although the generally poor motion control clearly influenced the design goals of Moravec’s next project.

5.3 CMU Mobile Robot Laboratory (WWW [35])

Moravec’s mixed success with the Stanford Cart (a successful implementation, with innovative research, but disappointing performance, as noted by Moravec) prompted him to build a new, more physically capable machine at Carnegie Mellon University’s Mobile Robot Lab. The basic design had produced reason and motion, but the rate of travel had been too slow for any practical application. The debate over the reasons for the speed deficiencies split into two principle camps: hardware and software. The hardware perspective was that the processors simply weren’t powerful enough to support navigation; predictions of required processor speed were made (Moravec has pointed out their significant underestimation) based on estimates of algorithmic complexity. The software perspective was that the vision processing algorithms could be made more efficient. Both perspectives were found out to be true. The Stanford Cart’s long cycle times thus focussed subsequent research on the problem of developing higher performance subcomponent technologies. The Mobile Robot Lab built two more machines (Neptune[26], Uranus[2]) and members of the group produced a large body of research work including (Elfes[12], Feng and Krogh[13], Muir and Neuman[22],[23]).



5.4 The Carnegie Mellon Wheeled Machines - Early Works, Terragator and NavLab

Moravec's work at SAIL prompted interest from the Defense Advanced Research Project Agency (DARPA), who initiated a program at Carnegie Mellon University called the Strategic Computing Initiative. This program was fairly wide in scope, calling for general advances in computer science & computer engineering. One portion of the program targeted road following navigation research resulting in the development of a hardier class of mobile robots and more robust navigation software.

A rugged mobile robot called Terragator[34] was constructed by the Civil Engineering Robotics Laboratory (CERL), the forerunner to the Field Robotics Center. This same group designed and deployed the teleoperated mobile robots that responded to the Three Mile Island nuclear disaster in the late 1970's. Their construction techniques were far advanced over those used in the construction of previous mobile robot test-beds, consequently Terragator represented a new era in mobile robot design. It was physically rugged and powerful, capable of literally climbing the vertical surfaces of trees. The earliest navigation research work focussed on following the sidewalks on CMU's campus. As the capabilities of these sidewalk-following algorithms steadily grew more robust, the capacity for adaptation to road-following became clear, as did the need for a road-worthy mobile robot. A Chevy van was retrofitted to carry the necessary computing and sensors, provide remote power and carry the researchers into the field. This mobile robot was dubbed NavLab, which stood for Navigational Laboratory, and reflected the philosophy of a mobile laboratory in which scientists could work on developing, testing and debugging navigational programs while out in the real world.

Technical papers from this era still refer to mobile robots as "Artificial Intelligence systems" or "symbol manipulators." Thorpe[33] Their approach was to develop vision processing algorithms

that were connected to the navigation algorithms through a distributed blackboard system. Vision processes performed two functions: color vision processing and 3-D image processing. Color vision processing techniques identified the road in an image, representing it as a set of geometric parameters. Edge detection, clustering and region growing operations were applied to 3-D image data to find obstacles, allowing the scene to be dissected into passible and impassible regions. They made the following key contributions:

1. Their approach recognized the difference between symbolic and real-time interfaces, which they called high-level and low-level reasoning.
2. The central importance of geometry and time were recognized as being "pervasive throughout the system." The constantly changing location of the object was found to be important to the system in a constant motion system.
3. Recognized that as vehicle speeds increase simplifying assumptions of dynamics fell apart.

Thorpe characterized the future needs of road-following research by calling for higher processing speeds (specifically he proposed using parallel processing techniques on an experimental computer called the WARP). Finally, Thorpe called for leaving the road and producing fully integrated on-road / off-road systems.

Speed control on Terragator and the NavLab were performed on powertrains engineered to be easily controllable. The Terragator's powertrain used electrical motors with high low-end torque characteristics. The NavLab's standard transmission was removed and replaced with a hydrostatic transmission to enable speed control (Dowling[11]).



5.5 Dickmanns, Graefe and Zapp ([8],[9],[10])

This work in monocular vision processing techniques stands out (especially considering its time, about one decade ago as of this writing) as an example of integrated systems engineering practice in autonomous mobile robotics. The researchers developed a set of monocular vision processing techniques for road following and engineered the necessary parallel processing hardware to support the real time requirements of a system operating at 100 k.p.h. They approached this work

as a problem of *design optimization*, rather than applying a brute force approach to individual techniques. Their emphasis on the real-time approach led them to emplace significant time constraints, deliberately imposing cycle time limits of 0.1 sec. Their approach demonstrated that through system-wide optimization, high performance need not wait for faster processors.

5.6 The Autonomous Land Vehicle

The Autonomous Land Vehicle (ALV) was the product of a DARPA program investigating autonomous cross country navigation, carried out by Martin Marietta Corporation and the Hughes Artificial Intelligence Center (Daily[7]). The ALV was a custom build eight-wheeled, skid-steered, adaptive suspension vehicle. Both propulsion and suspension were hydraulically controlled. The chassis carried the requisite computers, sensors, electronic support equipment, and power generation equipment, covered by a large fiberglass shell. There was room inside to carry researchers and programmers, making the ALV analogous to a large, custom-built, off-road NavLab.

The ALV program forced researchers to face more difficult terrains, and look more closely at the concept of "obstacle." In open terrain, obstacles were no longer specifically identifiable with objects, such as rocks or chairs. The slope, ditches, and other terrain characteristics were often "obstacles" if approached from a certain angle, or at too great a speed. Obstacleness became dependent upon the vehicle as well as the terrain. Under the concept of admissibility, a patch of terrain was considered admissible for a certain pose of the vehicle. Vehicle models were forward simulated over the terrain in candidate paths, in a search for an acceptable trajectory. The ALV team made more concrete the separation between low-level and high-level processes. In addition to producing innovative approaches to path planning, including weighted arbitration, they reinforced the concept that a complete system could be split into two portions, one that merely kept the vehicle alive by avoiding obstacles, though it might wander aimlessly, and a second that would direct the wandering of the first toward some particular goal. (Payton[25])

5.7 More CMU Robots - FastNav, NavLab II, IIb, III, IV, & V

Carnegie Mellon moved to off-terrain experimentation with the FastNav program (Singh[29]), a software development program on the NavLab mobile robot. The research content of this program was quite broad, including path tracking control (Shin[28]) obstacle avoidance, and the integration of satellite deduced (GPS) position estimates and Inertial Navigation Systems (INS). Road following work (Stentz[30]) continued concurrent with this off-road work, as well as additional work in control (Amidi[1]).

Off-road experience with the NavLab I demonstrated its unworthiness as a testbed. The vehicle lacked the necessary power and suspension to handle anything more aggressive than basically flat terrain. The CMU group moved to a more aggressive platform, a rugged military vehicle called the HMMWV (High Mobility Multipurpose Wheeled Vehicle). The NavLab II was basically a copy of the NavLab I, moved to a different platform. The engineering design philosophy on both of these vehicles was function oriented, generality being a plus. A second HMMWV, the NavLab IV, built commercially by RedZone Robotics for University of Massachusetts was a commercial knock-off

of higher quality. NavLab IIb (Coulter[5], Coulter and Mueller[6]) was a redesign of the NavLab II, retaining the same platform but reducing the payload and powerdraw through the standard application of systems engineering principles. The NavLab III took the road following technology and a minimalist hardware perspective to move to a 50 lbf package on a Honda Accord, and a similar package on a Pontiac Transport - the NavLab V. These last two designs indicate a shift toward performance design.



Road following and controls algorithms from the NavLab I were ported to the NavLab II. Immediate performance improvements followed as the allowable vehicle speed increased to the highway speed limit. (Pomerleau[27]) The use of a rugged off-road platform also allowed for the creation of so-called cross-country navigation algorithms, which allowed the vehicle to operate in essentially barren terrain, performing obstacle avoidance. Early works (Stentz et al.[31]) were plagued by real-time induced problems. Later systems by Langer[19] and Kelly[16] have managed to produce more robust low level terrain navigation at higher speeds. Kelly reports that the currently achievable speeds are in fact a theoretical maximum due to the limitations of the vehicle sensory hardware, and called for the construction of a new generation of laser rangefinders, and more efficient data processing algorithms. Work in global planning resulted in new algorithms from Stentz[32] that offer the ability to perform independent exploration in unknown, unstructured domains.

6. AUTONOMOUS SPEED CONTROL - PRIOR WORK

Most people are familiar with automotive cruise control, and draw this first analogy with autonomous speed control; however, speed control for on-road vehicles (cruise control) is fundamentally a different problem. It is a regulatory problem of a system with a much greater momentum, that sees smaller disturbance forces of slower rate. Cruise controls do not operate under 30 m.p.h. and, until the recent addition of pitch sensors, had difficulty with even moderate road grades. The off-road speed control problem is, by contrast, essentially a servo controls problem. The vehicle is operating at low speed where it has little momentum. Variations in

disturbance forces (such as gravity force) are much greater, changes in pitch tend to be more sudden and of much greater magnitude. The result is a system which spends considerable time occupying its transient region.

6.1 NIST & Martin Marietta Work - the Lookup Table Approach

There are few cases of precedence in rough terrain speed control of conventional vehicles. Karl Murphy, from NIST, is probably the seminal work in the field, unfortunately his work has not been published, nor is it recorded. My understanding of Murphy's work has come from email exchanges with him, and through personal conversations with Patrick Rowe, who worked with Murphy's controllers, and ultimately abandoned them in favor of other techniques. Murphy used a lookup table in combination with a proportional / feedforward control law, in which the gains were scheduled. Gain scheduling is essentially the process of breaking a nonlinear system into linear regions and determining optimal gains for each region. Murphy's control law was a linear combination of three terms:

$$\phi = C + K_1 e_V + K_2 \sin \theta$$

The values of the constant C and the gains K_1 and K_2 were stored in a lookup table indexed on the commanded vehicle speed. Pat Rowe, who is now a graduate student here at RI, took this controller over from Karl Murphy to try to improve its performance. He found it difficult to tune, owing to the large number of experiments that one needed to apply to find the optimal lookup table values. He abandoned the use of this method and was beginning to experiment with a PID controller when he left NIST to come to RI.

Murphy's control law is the same method used by Martin Marietta in their controller. Through successive refinement of the lookup table values, Martin personnel have been able to extract good performance from the method. This is not at all surprising; gain scheduling as a method is virtually guaranteed to work, provided that the system is broken up into sufficiently small, linear portions. The main drawback to the method is the difficulty of tuning it, as noted by Rowe.

6.2 Linear Methods of Integrated Motion Control

The CMU UGV group's first off-road autonomous navigation system (Stentz et al. [31]) interfaced to a controller built by Omead Amidi[1]. The speed control algorithm on this controller was a PID linear controller that closed a single loop around velocity. I used this controller over several months and gradually learned about its limitations and characteristics, which can be summarized as follows:

1. Stability vs. Command Following trade-off: We were unable to find a set of gains that produced both good (fast) transient behavior and stable steady state behavior.
2. Integrator Wind-up - The integral term often stored excessive error, leading to velocity run-away as the vehicle crested a hill.

3. Gain sensitivity - The system bandwidth over which a set of gains provided acceptable performance was small.

4. High speed performance - at highway speed, when operating as a cruise control, the PID controller worked well.

It is interesting to note that although the PID algorithm performed poorly in throttle control, it performed well in brake control. The reasons for this difference will be explored in later chapters.

7. MOTION PLANNING - PRIOR WORK

There are two principal methods of search, which are distinguished by the algorithm used to construct the planned output path. I will call them *generate-from-search* and the *generate-and-test* paradigms. In the *generate-from-search* method, the terrain features are tessellated into a grid, which is searched for an acceptable path. Each cycle of the search process, the search engine is required to evaluate one or more of the terrain grids to determine whether or not the vehicle could occupy or pass through. A cell that is considered acceptable is termed admissible to the search. (Stentz[30]) The space is searched until enough cells can be pieced together to form an acceptable spacial trajectory. The *generate-and-test* paradigm constructs a set of acceptable spacial trajectories first, and then evaluates the motion that the vehicle would execute using techniques of forward simulation. The simulated responses are then graded against a set of performance criteria and the optimal path is chosen for execution.

High speed, rough terrain navigators have moved from the *generate-from-search* method back to the *generate-and-test* method, which originated with the ALV program. Kelly makes the case for the latter by noting that the construction of a path that is executable by a non-holonomic vehicle is the most significant constraint of the planning process. Kelly improved upon the general method by including a more accurate *representation of the steering dynamics* in the candidate trajectory generator. This method of feedforward dynamic simulation, coupled with an optimal control evaluator has also been called *action-space planning* in a recent paper by (Kelly and Singh[18]).

8. MOTION EVALUATION - PRIOR WORK

The process of search requires some basis for discriminating "good" terrain from "bad" terrain. The method of discrimination is dependent, to a degree, upon the method of search. For vehicles executing trajectories across terrain, forward propagation of a vehicle model and evaluation of the dynamic response of the vehicle / terrain / suspension interactions is usual. This method also originated in a pseudo-static manner with the ALV program. Their software system forward propagated a model of the ALV across the candidate trajectory, allowing the suspension to settle to an equilibrium position at each point. The equilibrium pose then offered a stability metric for path evaluation.

At Carnegie Mellon, Coulter[4] offered an alternative, in which the maximum velocity profile of the vehicle is calculated from a kinematically admissible path. The velocity profile is generated by backsolving the differential equations for a given maximum acceleration. Such an evaluation

method is a hybrid between motion evaluation and velocity planning, which suffers from the drawback of generating only the maximum trajectory, not the optimal trajectory, but offers the run time advantage of computing algebraic statements rather than solving differential equations.

9. AUTOMOTIVE MODELLING - RELATED WORK / PERSPECTIVE

The automotive world has produced an immense, voluminous quantity of work in engine and powertrain modelling from which this thesis draws. In general, the automotive literature addresses deep technical issues in each of a number of specialities. General descriptions of physical mechanical issues are no longer of interest. For example, if one were interested in modelling engine dynamics, the combustion process might be of interest - combustion related papers are generally interested in modelling processes at very high levels of fidelity. A combustion model may include a shell ignition process model that describes the development and progression of combustion. Clearly, this level of fidelity is too fine for the needs of the work presented in the thesis - which leads to a central question which pervades - what level of fidelity will suffice the world. Fidelity is the principal issue that distinguishes the models presented in this thesis from those that are found in the automotive literature, and the related issue of computational efficiency makes automotive models inappropriate for use in autonomous speed control. I considered a wide range of automotive models including:

1. Thermodynamic
2. Incompressible Mass Flow
3. Rotational Dynamic
4. Combustion Dynamic
5. Fluid Dynamic

The set of models that I finally produced were the simplest set that I felt captured the fundamentals of the dynamical processes at work. The validity of the choice is supported by the performance of the servo-controls algorithm produced in this thesis.

10. SIMULATION TECHNOLOGIES

Many simulation technologies can appear to be very similar to autonomous speed control technologies when such control is based upon modelling, as is the case in this thesis. Numerous software simulation packages are currently available, often as commercial products, which are capable of representing either engine and powertrain dynamics, or suspension dynamics, or vehicle response. These software tools are used to simulate dynamical processes, usually for the purposes of engineering design and analysis. Even newer, higher fidelity technologies are becoming available which represent "virtual vehicles" capable of being operated in virtual worlds. The emphasis is on developing non-real-time models of the vehicle or vehicle components for evaluation. Analytical simulators emphasize the computation of an accurate solution to a differential equation model.

Mathematically, these simulators tend to integrate high-order differential equation models of physical processes. Preservation of accuracy in integration tends to be highly coupled to the integration time step, forcing the use of small-time step (and computationally intensive) routines. The virtual simulators are extended software systems that integrate many such physical simulators with a virtual reality human interface to produce a realistic driving environment. The primary example of such a virtual simulator is the Iowa Driving Simulator[14]. The motion evaluation process appears to be similar to forward dynamic simulation, but possesses the following distinguishing characteristics:

1. It is often the case that the motion evaluator need only return an ordered set of paths - a description of which path is better is sufficient, specific calculations of vehicle reaction are not necessary.
2. *The motion evaluator often is the central function in the planning cycle.* Its run time can dictate the cycle time of the system, making it important to limit the time spent computing an answer.
3. Motion evaluators are run in real time on a vehicle. Motion simulators are run on high-speed workstations where exact run times are non-critical.

In either case, the production of a more accurate answer typically requires the use of a more refined model, which takes longer to run. The accuracy of an answer is strongly coupled to the run time of the solution.

11. CONCLUSIONS

Very little work has been done in autonomous speed control. Gain scheduling as a control technology has been shown to achieve speed control, at the usual cost of significant tuning efforts. It is reasonable to attempt to move from a discretized state of control commands to a more compact, functionalized representation - this is precisely what Rowe attempted at NIST, and what CMU implemented in earlier speed controllers. Although the linear controls techniques did not prove satisfactory for autonomous speed control in unstructured 3-D terrain, they do provide valuable constraint information for the design of more capable controllers. Thus, the first task of this thesis is to independently implement and study the performance of the proportional controller. Previous experience with proportional controllers has already offered some indication that the dynamics of the powertrain transfer function are nonlinear. Thus modelling of the system is also indicated as a means of identifying the underlying causes of this nonlinear behavior.

The historical section pointed out the recognition of functional separation in autonomous navigation systems between the "reasoning" portion of the system, and the "reacting" portion of the system. I believe that the same functional separation exists in the autonomous speed control problem. I have proposed in this thesis a solution to the reacting portion of the problem. I do not attempt to reason about speed at the mission planning level, but rather I assume that some external agent has already done this and passed on a directive to the system which is implicitly represented in the constraints of the optimal control speed planner. The purpose of the system is then to optimize these directive against the realities of the physical terrain and the capabilities of the

vehicle under its control. Prior work in planning indicates that an optimal control solution might be employable *if the underlying powertrain dynamics were made more controllable*. It seems that we have a similar problem to autonomous navigation (the steering problem) except that the transfer function between the planner output and the system output is significantly more complicated than in the steering case. A reasonable plan of attack is then to first produce an integrated motion control scheme which makes controllable the powertrain, then attempt to apply the optimal control solution to this new closed loop system.

Prior work in planning systems have also indicated that the process of motion evaluation becomes increasingly difficult as the terrain complexity increases. The performance of an optimal controller or planning is greatly dependent upon the run time of the motion evaluator. The “goodness” of a path is highly dependent upon the interaction of the terrain and the vehicle as described by the response of the passive dynamics of the suspension. The motion evaluation software system lies at the heart of the innermost cycle of the planner along with the trajectory generator, making its performance indicative of the cycle performance of the planning system. Faster or higher bandwidth techniques could help to improve the robustness of the system (by improving the reaction time) and the quality of the evaluation (by modelling higher-order dynamics).

CHAPTER III - LINEAR CONTROL EXPERIMENTS

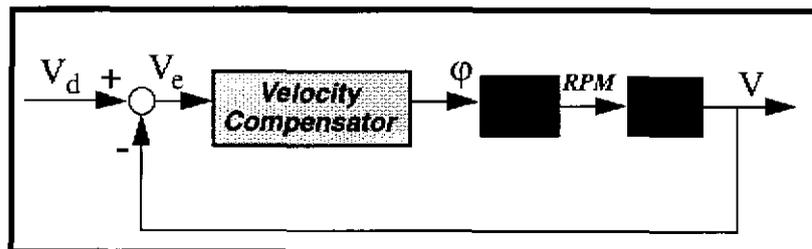
It is natural, when using a discrete, explicit control representation such as a lookup table, to attempt to express the control law in a more compact, functionalized form. Compact representations usually offer greater ease of calibration, and often more insight into the problem. At the commencement of this work, such advantages were sought in the form of linear control techniques, specifically proportional control, without much success. As noted in the prior work section, proportional controllers were found to be capable in a highway-speed, regulatory capacity, but were unable to provide robust control at the low speeds required of cross country navigation. Small changes in slopes caused significant changes in the performance of the controller. The limitations of the proportional controller offer a good starting set of constraints for this research; their limitations are useful in indicating the underlying problems with the algorithm and a direction for their improvement.

12. EXPERIMENTAL PROCEDURE

Each of the controls experiments were run over the same course - a hill in Schenely Park called Flagstaff Hill. Flagstaff is a paved path used for biking and walking through the park. It is fairly inclined, with a slope that varies between level and about 7 degrees. It is positively inclined over its entire length. During each experiment, the NavLab II began from a standstill at the same position at the bottom of the hill. Each algorithm was allowed to run until either the truck came to a stall, the truck completed the course to the top of the hill, or the truck completed the course to a large tree about one third of the way up. The decision as to whether to allow the truck to continue to the top of the hill, or to stop at the tree was usually dependent upon its behavior. If the algorithm was causing large cycles in throttle position (which is often frightening to passers-by) experimentation was stopped at the tree. More well-behaved sets were allowed to continue to the top of the hill.

12.1 Proportional Control

A single loop proportional controller was constructed and tested on Flagstaff Hill. In the single stage feedback controller, the velocity compensator has a single input - the velocity error V_e - and a single output, the throttle plate angle ϕ . In the following graphic, $E(t)$ represents engine

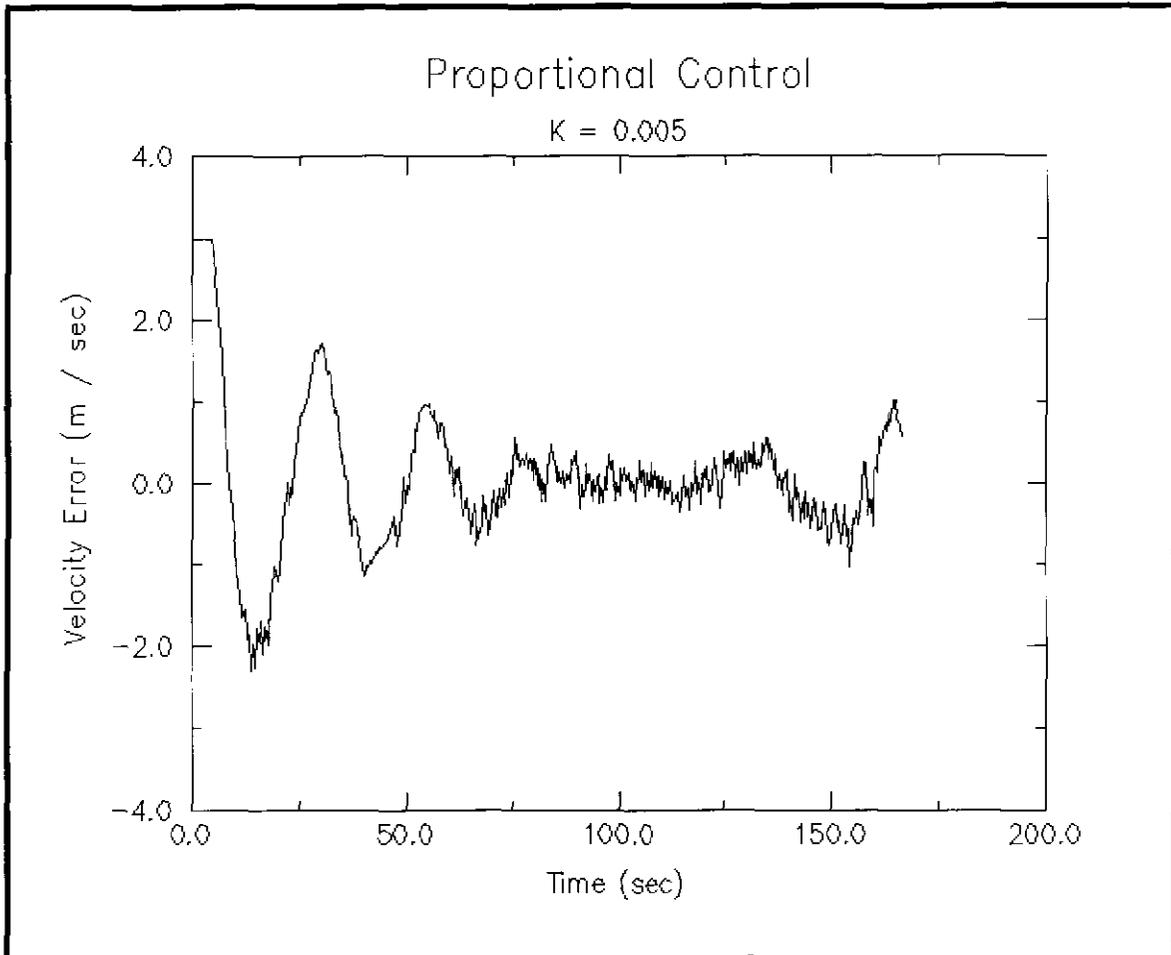


dynamics, and $P(t)$ represents powertrain dynamics.

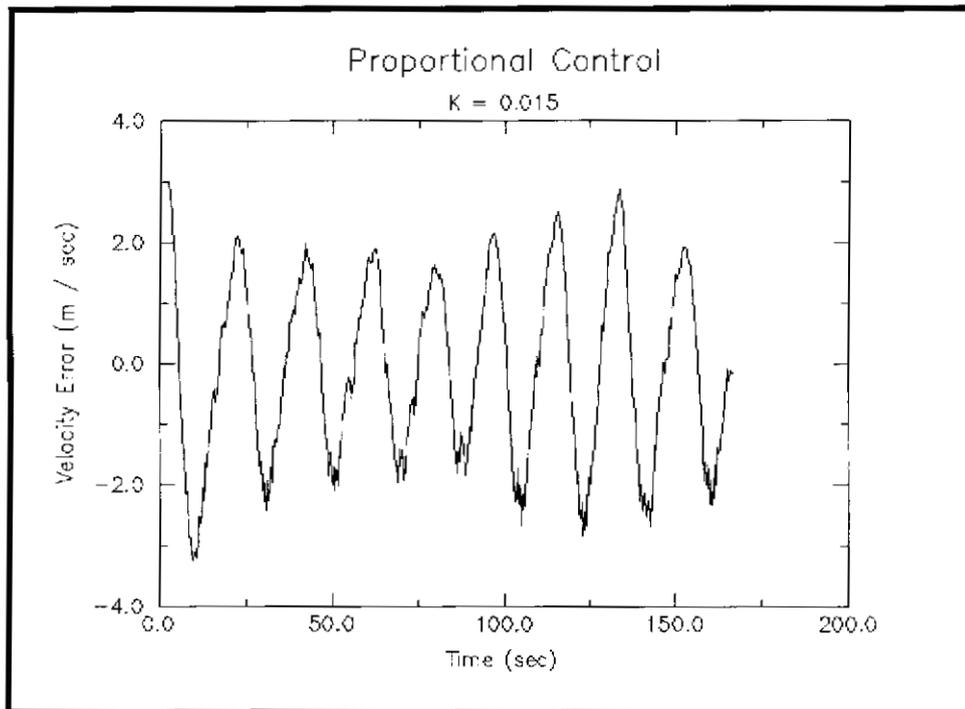
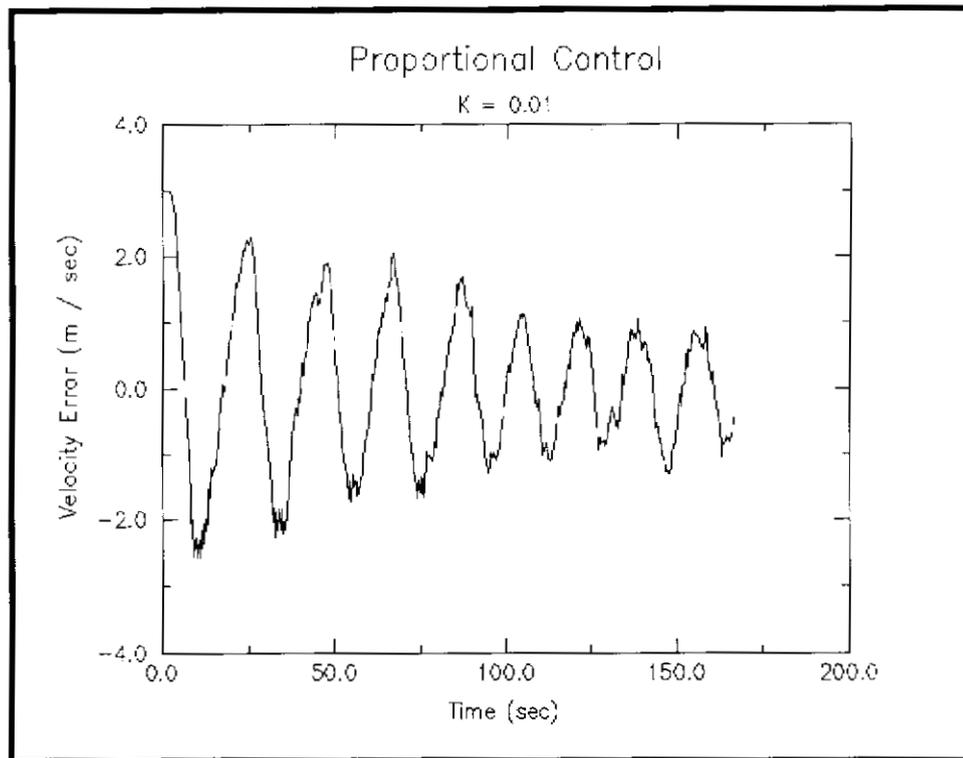
The proportional gain was varied by approximately 1 order of magnitude, and the gross response of the vehicle recorded. The single loop proportional control law is given by:

$$\dot{\phi} = K_p V_e$$

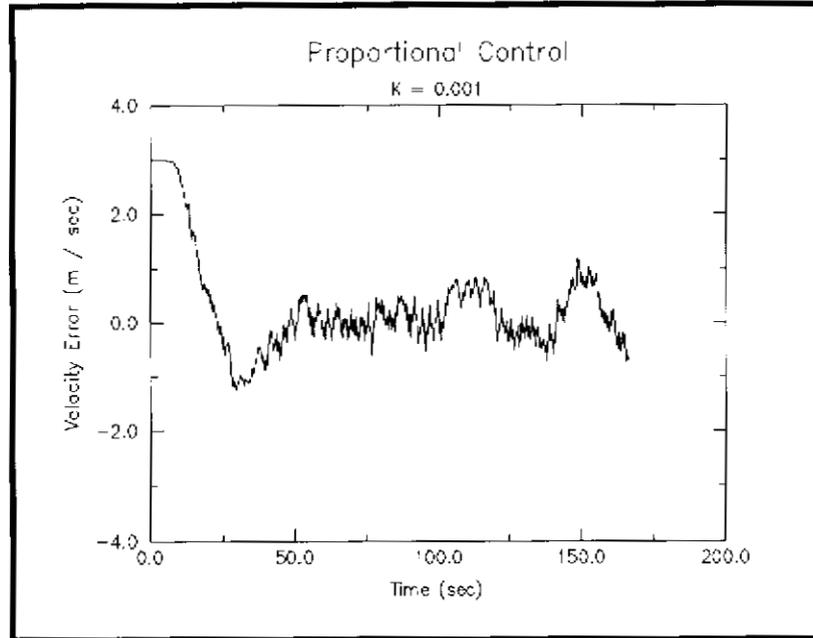
The following graph of velocity error as a function of time is typical of all of the proportional runs. Characteristics of the proportional controller include a tendency to overshoot, and a pseudo-oscillatory behavior.



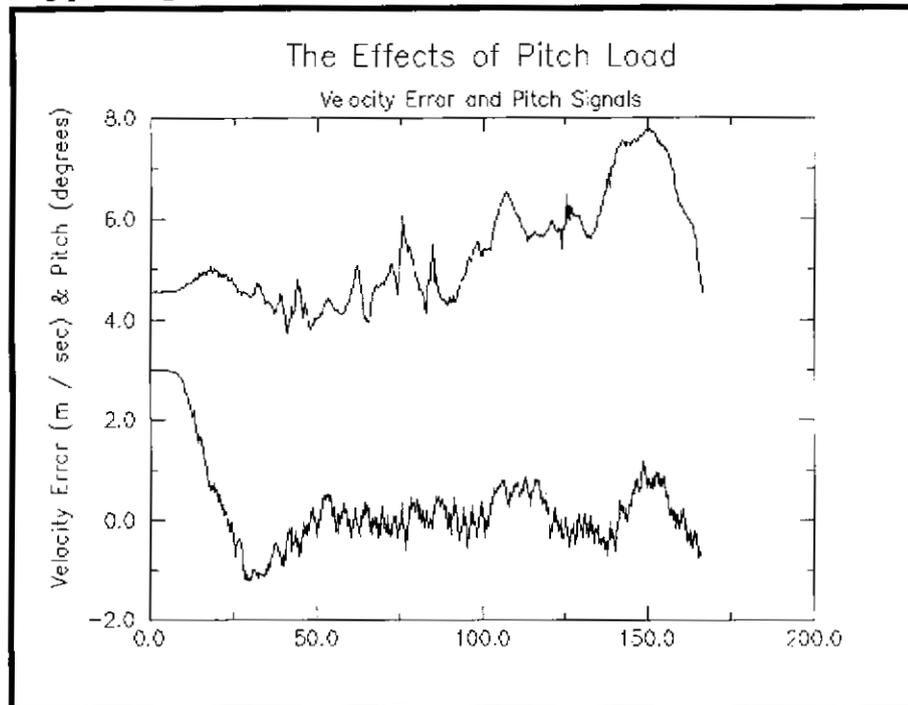
Varying the gain produces little change in the characteristic behavior. Overshoot always occurs - convergence is dependent upon the gain; large gains induce non-converging pseudo-oscillatory responses, while smaller gains produce a more damped pseudo-oscillatory response.



It is possible to “tune” the proportional controller to produce essentially critically damped behavior. Even under the best circumstances, the controller exhibits a high sensitivity to changes in the terrain pitch. In the following graphic, note that the controller overshoots, then settles about zero. The two humps at about 100 and 150 seconds correlate to increases of one or two degrees in the terrain pitch.



The general influence of gravitational loading may be seen when the velocity error plot is displayed with the attending pitch signal:



12.1.1 Pseudo-Oscillatory Response

In the text above, the response of the vehicle to this control law is referred to as *pseudo-oscillatory*. From elementary control theory, the classic response of an underdamped second order system is oscillatory. I wish to distinguish this bilateral response from the sometimes unilateral response of the vehicle. The control input is a throttle plate, which regulates the flow of fuel to the engine - we are able to regulate the amount of energy that is put into the system, but cannot, through this unilateral control input, *remove* energy from the system.

Qualitatively, the response of the vehicle feels like a surging, especially when the overshoot is so large that the control law closes the throttle completely and waits for the frictions of the system to reduce the speed of the vehicle. Mechanically, there is a sudden input of energy into the system, which then is bled out more gradually.

12.1.2 Overshoot

The proportional controller, as constructed, is guaranteed to overshoot. This phenomena can be seen by examination of the control law in the context of our experience driving a car. Recall that the implemented control law is given by:

$$\dot{\phi} = K_p V_e$$

Consider the action of the throttle under the following circumstances:

1. When $V_e = 0$; $d\phi / dt = 0$. Throttle position is held constant.
2. When $V_e > 0$; $d\phi / dt > 0$. Throttle is opening.
3. When $V_e < 0$; $d\phi / dt < 0$. Throttle is closing.

It is clear that the throttle is always opening, though at a progressively slower rate, until the zero crossing is reached, indicating that the algorithm will always cause the vehicle to overshoot its intended velocity. In these experiments, proportional control was implemented as a rate law because of the desired convergence compatibility between the control command and the feedback error. It is typically the case that the control signal should tend to zero as the feedback error tends to zero. In the case of operating an automobile's throttle, it seemed to make more sense to hold a *constant* throttle position as the velocity error tended toward zero.¹

12.1.3 Convergence and a Comparison between Murphy and NavLab II PID Control

Both Murphy's control law and the NavLab II's PID control law make the throttle position proportional to the velocity error. Both laws thus suffer from a sort of mismatch between feedback and control commands - if either law were implemented with only the proportional velocity error term, the throttle would close as the desired speed is approached. Thus each control law must be augmented by another term to ensure that this does not happen.

1. This happened quite by accident. When I originally implemented the proportional controller with $\phi = kV_e$, the throttle closed as the vehicle approached the desired speed. Changing to a rate law allowed the vehicle to attain at least gross bodily motion.

Murphy's control law attempts to functionalize the throttle plate angle explicitly, using a linear combination of three terms:

$$\phi = C + K_1 e_V + K_2 \sin \theta$$

Murphy's controller is essentially a servo-controls design with a proportional feedback loop augmented by a feedforward term. The constant C is not of classical origins, it is functionally required to compensate for frictions. Consider this controller operating in level terrain, where the sine term becomes zero. Now, consider the convergence behavior of this law. As the velocity error draws toward zero, the throttle plate moves to position C . Clearly, for each velocity that the vehicle is to travel, there must be a different value of C chosen. Thus the strength of this law lies in one's ability to encode in C the nominal throttle position - the other two terms offer some adjustment on this nominal value - the proportional term is presumably to counteract random system errors, while the last term counteracts gravitational disturbances. This system is essentially an *augmented lookup table*.

The control law implemented on the NavLab II is a standard PID control law, which is given as follows:

$$\phi = K_P V_e + K_I \int V_e + K_D \dot{V}_e$$

The mismatch between the convergent feedback and command signals is overcome in this case by the use of the integral control. Integral control has the effect of forcing the throttle plate position ever higher as the feedback signal refuses to close to zero. In this sense *integral control is completely equivalent to the control law applied in the previous linear control experiments*.

13. CONCLUSIONS

The exploration of linear control techniques indicates that either the throttle plate angle is not mapping over to force or the system is dominated by its frictions. If the throttle mapped to force, then the proportional control law should provide convergence in non-integral form. The fact that it does provide convergence in integral form seems to indicate that the throttle plate angle maps onto some integral of force. The experiments on Flagstaff Hill also indicated that there is a strong correlation between the output error and the pitch signal, which seems to indicate some need for feedforward control. Subsequent chapters will support both of these hypotheses.

Part II

Modelling

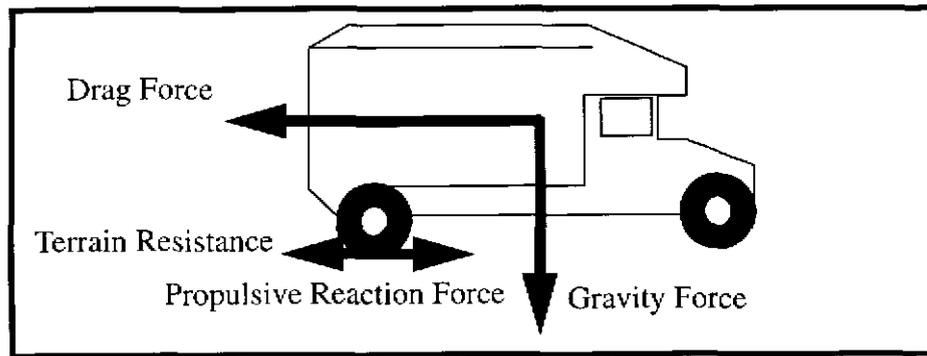
This thesis is heavily dependent upon the production of a physical mechanics model of the system. **Part II** presents the process and results of modelling for this system. In the first chapter, a *system model* is formulated from the externally applied and reactive forces. A subsequent chapter deals explicitly with these externally applied forces as *disturbances*, producing several relative analyses that are useful to the production of the integrated control system in later chapters. Finally, the internally generated *control forces* are modelled.

CHAPTER IV - SYSTEM MODELS

14. EXTERNALLY APPLIED FORCES

Mobile robots are influenced by a set of externally applied forces, produced by one of three means:

- Forces arising from the gravitational field.
- Forces arising from the motion of the machine in a fluid medium.
- Forces arising in reaction to the action of the machine upon the soil.



Note the dissimilarity in nature among the forces: the gravitational force is a *field force*, the drag force arises from *viscous friction*, and the terrain resistance arises from *mechanical friction*.

14.1 Gravity Force

The combined force of gravitational attraction and the Coriolis effect as felt at the Earth's surface is known as the *local gravity force*. Gravity, like gravitation, is treated as a field force. It is represented locally by the vector \mathbf{g} . The magnitude and direction of this force are constant, making its effects upon the vehicle a *function of the orientation*.¹ In conventional vehicles, there are two principal gravitational effects:

- Gravity loading of the powertrain.
- Gravity loading in the case of stability.

The projection of the gravity vector on the vehicle's propulsion axis manifests itself as a *load on the powertrain*. Clearly, this is of concern to vehicle motion control. Gravity loading in the prediction of stability is of concern in motion planning.

1. Note that the value of the Coriolis force varies with latitude, and the magnitude of gravitational attraction varies with altitude, so the magnitude of \mathbf{g} is only locally constant.

14.1.1 Powertrain Loading

The gravitational load that acts on the powertrain is found to be:

$$F_g = mg \sin \Theta$$

where Θ is the pitch of the vehicle.

14.1.2 Rate of Change of Powertrain Load

The time rate of change of this force is given by:

$$\frac{dF_g}{dt} = mg \cos \Theta \frac{d\Theta}{dt} = mg v \cos \Theta \frac{d\Theta}{ds}$$

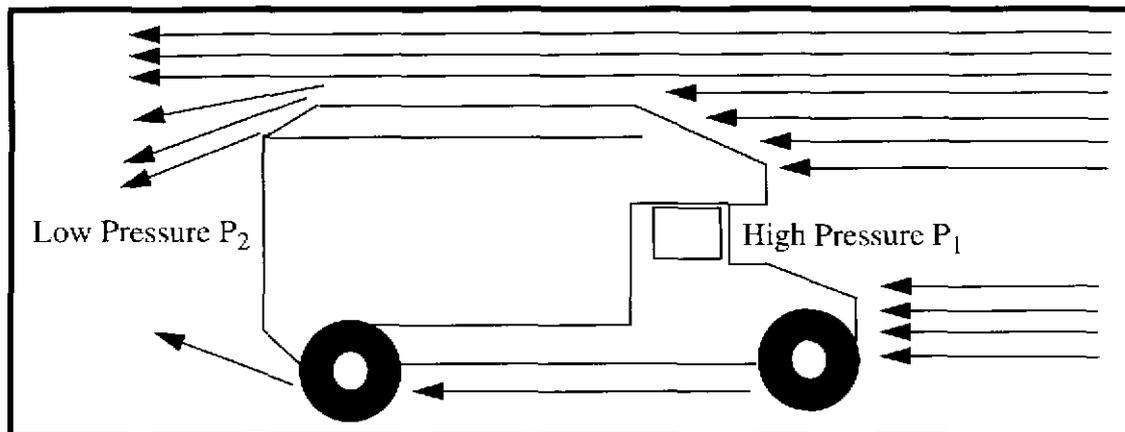
where Θ is the pitch of the vehicle and s measures the distance along the vehicle's path. The time derivative of s is represented by v - the vehicle's speed.

14.2 Aerodynamic Forces

The relative motion of the vehicle and its surrounding fluid gives rise to a set of three aerodynamic forces: *lift*, *drag* and *side force*; and a moment about each axis. For the purposes of this thesis, it will be assumed that the principal relative motion is along the longitudinal axis, and that it results in a drag force only.

14.2.1 Drag Force

The force of viscous friction at the air / vehicle surface acts to retard vehicle motion. This viscous friction is commonly called the *drag force*. The drag force has two sources: *form drag* and *friction drag*. *Form drag* is the result of the pressure differential formed as the vehicle body passes through the fluid. The pressure P_1 in front of the vehicle is higher than the pressure P_2 behind the vehicle, which gives rise to a retarding force $(P_1 - P_2) * A$, where A is the effective cross sectional area. Friction drag results from the contact between the fluid stream and the vehicle surface.



The total drag force is empirically represented by the following formula:

$$F_d = \frac{1}{2} \rho V^2 C_d A$$

where ρ represents the density of air, A is the effective cross sectional area of the vehicle, V is the velocity of the vehicle *relative to the air*, and C_d is the drag coefficient.

14.2.2 Drag Force Sensitivity

The sensitivity of the drag force to changes in vehicle speed is found by taking the first derivative of F_d with respect to V .

$$\frac{dF_d}{dV} = (\rho C_d A) V$$

14.3 Rolling Resistance

The interaction of the vehicle tires with the terrain results in an energy loss, most commonly referred to as rolling resistance. There are many models of rolling resistance. The first order approximation is a coulombic model.

$$F_R = fmg$$

where f is the rolling resistance coefficient. Changes in mechanical configuration, such as tread design, tire inflation, soil compaction, and tire temperature can and do affect the value of the resistance coefficient. However, the rate of change of these variables is often insignificant to the control problem at hand. Rolling resistance is also a function of vehicle speed, though the sensitivity is small below speeds of $\sim 80 - 100$ m.p.h. ([15]Gillespie)

14.4 Orthogonality of the Forms of the Forces

Each of the forces is parameterized in a nearly orthogonal space. The gravity load is a function of terrain geometry and vehicle mass. The aerodynamic load is a function of fluid properties, vehicle geometry and the speed of the vehicle. Rolling resistance is a function of the tire and soil properties, and the mass of the vehicle. The forces are generally orthogonal in their parameterization, the exception being that both mechanical friction and gravity forces are mass dependent.

$$\begin{aligned} F_g &= fn(m, \Theta) \\ F_a &= fn(\rho, C_d, A, V) \\ F_r &= fn(m, f) \end{aligned}$$

15. REACTION FORCES

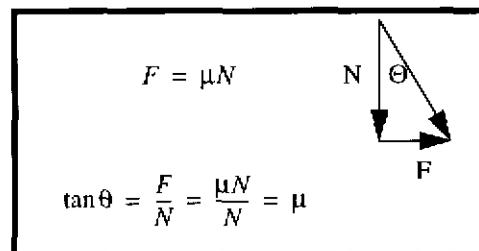
Wheeled vehicles move and maneuver as a result of the frictions generated at the tire / terrain interface. The vertical reaction at this interface acts as the input to the suspension system, generating the vehicle's bodily response to excursion.

15.1 Friction

The tire and possibly the terrain deflect, generating a contact patch of finite area - this patch is the interface through which all of the lateral and longitudinal control forces must be generated. All propulsive, braking, and maneuvering forces are transferred to the terrain at this area of frictional contact. The amount of friction generated is dependent upon mechanical properties of the two surfaces, as well as the magnitude of the compressive normal force.

15.1.1 Friction cone model

The friction cone is a model that describes the set of impressed forces for which a reaction may be generated, under the assumption that the interface friction is Coulombic. The friction cone is generated from the following geometric construction. A vector representation of a normal force N , and the maximum generated frictional force F is shown. The interior angle Θ is found from trigonometric identities to be $\arctan \mu$:



The friction cone describes the set of forces that can be supported through frictional reaction. If the vector sum of the normal and lateral impressed forces describes a line of action that falls within the swept cone, then sufficient reactive forces are generated. The lateral impressed force may be the result of accelerating, braking, steering and any combination. Soil Reaction

The general soil reaction is a problem of deformation mechanics. Soil compression and deformation lead to sinkage and slip. In general these are complicated problems of terramechanics which are considered to be beyond the scope of this thesis. For this reason, the soil is assumed to be rigid and non-compliant, and all experimentation is performed in areas that meet these assumptions.

16. THE NEWTON EULER EQUATIONS

The derivation of the Newton -Euler equations proceeds directly from Newton's Second Law and Euler's equation. First, each of the forces must be represented in the same frame and coordinate system. We will consider inertial forces represented in the moving coordinate system - the first

subsection will detail the kinematic transformations¹ of these externally applied forces. The next section will consider forces arising due to the motion of the rigid body. The third section will consider the effects on translational motion of the rigid body as specified by Newton's second law, while the final section considers the effects of these forces on the rotational motion of the rigid body as specified by Euler's equations.

16.1 Kinematic Transformations of the Externally Applied Forces

16.1.1 Transformation of the Gravity Force

The gravity force is of magnitude mg , in the negative Z direction of the fixed frame.

$$F_g^{fixed} = \begin{bmatrix} 0 \\ 0 \\ -mg \end{bmatrix}$$

The gravity force is represented in the moving coordinates through the following transformation:

$$F_g^{moving} = \begin{bmatrix} c\theta s\phi mg \\ -s\theta mg \\ -c\theta c\phi mg \end{bmatrix} = \begin{bmatrix} (c\psi c\phi - s\psi s\theta s\phi) & (s\psi c\phi + c\psi s\theta s\phi) & -c\theta s\phi \\ -s\psi c\theta & c\psi c\theta & s\theta \\ (c\psi s\phi + s\psi s\theta c\phi) & (s\psi s\phi - c\psi s\theta c\phi) & c\theta c\phi \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ -mg \end{bmatrix}$$

16.1.2 Transformation of the Aerodynamic Resistance Force

The aerodynamic resistance force is reported in the negative y direction of the moving coordinate system.

$$F_{aero}^{moving} = \begin{bmatrix} 0 \\ -\frac{\rho V^2}{2} \\ 0 \end{bmatrix}$$

¹. Please note that the kinematic transforms do not follow the SAE standards for vehicle coordinate systems, but are standard transforms for military grade INS / GPS units. The transform rules are given in the first appendix.

The aerodynamic resistance force is represented in the fixed coordinates through the following transformation:

$$F_{aero}^{fixed} = \begin{bmatrix} -s\psi c\theta \\ c\psi c\theta \\ s\theta \end{bmatrix} \left(-\frac{\rho V^2}{2} \right) = \begin{bmatrix} (c\psi c\phi - s\psi s\theta s\phi) & -s\psi c\theta & (c\psi s\phi + s\psi s\theta c\phi) \\ (s\psi c\phi + c\psi s\theta s\phi) & c\psi c\theta & (s\psi s\phi - c\psi s\theta c\phi) \\ -c\theta s\phi & s\theta & c\theta c\phi \end{bmatrix} \begin{bmatrix} 0 \\ -\frac{\rho V^2}{2} \\ 0 \end{bmatrix}$$

16.1.3 Transformation of the Rolling Resistance Force

The rolling resistance force is reported in the negative y direction of the moving coordinate system.

$$F_{rolling}^{moving} = \begin{bmatrix} 0 \\ -fmg \\ 0 \end{bmatrix}$$

The rolling resistance force is represented in the fixed coordinates through the following transformation:

$$F_{rolling}^{fixed} = \begin{bmatrix} -s\psi c\theta \\ c\psi c\theta \\ s\theta \end{bmatrix} (-fmg) = \begin{bmatrix} (c\psi c\phi - s\psi s\theta s\phi) & -s\psi c\theta & (c\psi s\phi + s\psi s\theta c\phi) \\ (s\psi c\phi + c\psi s\theta s\phi) & c\psi c\theta & (s\psi s\phi - c\psi s\theta c\phi) \\ -c\theta s\phi & s\theta & c\theta c\phi \end{bmatrix} \begin{bmatrix} 0 \\ -fmg \\ 0 \end{bmatrix}$$

16.1.4 Transformation of the Tire Reactive Forces

The tire reactive forces are reported in each direction of the moving coordinate system.

$$F_{tire}^{moving} = \begin{bmatrix} f_{sx} \\ f_{sy} \\ f_{sz} \end{bmatrix}$$

The tire reactive forces can be represented in the fixed coordinates through the following transformation:

$$F_{tire}^{fixed} = \begin{bmatrix} (c\psi c\theta - s\psi s\theta s\phi) & -s\psi c\theta & (c\psi s\phi + s\psi s\theta c\phi) \\ (s\psi c\phi + c\psi s\theta s\phi) & c\psi c\theta & (s\psi s\phi - c\psi s\theta c\phi) \\ -c\theta s\phi & s\theta & c\theta c\phi \end{bmatrix} \begin{bmatrix} f_{sx} \\ f_{sy} \\ f_{sz} \end{bmatrix}$$

16.2 Apparent Forces

Apparent forces arise due strictly to the motion of a non-inertial frame. Recall the frame to frame relationship known as the Coriolis Law:

$$\ddot{a}_f = \ddot{a}_m + \frac{d^2 \hat{p}}{dt^2} + \frac{d\omega}{dt} \times \hat{r}_m + 2\omega \times \dot{r}_m + \omega \times \omega \times \hat{r}_m$$

The vector r_m describes the position of the center of mass of the body in the moving frame. The vectors v_m and a_m are both zero, as the center of mass is immobile within the moving frame. The reduced Coriolis law thus takes the following form:

$$\ddot{a}_f = \frac{d^2 \hat{p}}{dt^2} + \frac{d\omega}{dt} \times \hat{r}_m + \omega \times \omega \times \hat{r}_m$$

The total force is the sum of the externally applied and apparent forces, and is represented by multiplying the reduced Coriolis law by the mass of the body:

$$m\ddot{a}_f = m \frac{d^2 \hat{p}}{dt^2} + m \left(\frac{d\omega}{dt} \times \hat{r}_m \right) + m (\omega \times (\omega \times \hat{r}_m))$$

$$m\ddot{a}_f = \Sigma F_e + m \left(\frac{d\omega}{dt} \times \hat{r}_m \right) + m (\omega \times (\omega \times \hat{r}_m))$$

16.3 Total Force Coordinate System Transformation

The total force equation, as written, reports the total force in the fixed coordinate system. For motion control purposes, these forces are best reported in the moving coordinate system. The term on the left hand side of the following equation is the acceleration with respect to the fixed frame, transformed into the moving coordinate system. There is no similar confusion with respect to the externally applied forces, as the concept of force is independent of the frame of reference, thus only

the coordinate system need be specified. For simplicity, the coordinate transformed acceleration is denoted with the superscript moving, referring to the coordinate system, and the subscript f, for fixed, denoting the frame of reference.

$$m\ddot{a}_f^{moving} = \Sigma \dot{F}_e^{moving} + m \left(\frac{d\dot{\omega}}{dt} \times \hat{r}_m \right) + m (\dot{\omega} \times (\dot{\omega} \times \hat{r}_m))$$

This final equation relates the accelerations that would be read by an instrument on board the vehicle, reported in a coordinate system that is attached to the vehicle.

16.4 Final Vector Matrix Form

The vector form of the externally applied forces (including the reaction forces) can be directly substituted into the vector-matrix form of Newton's second law.

$$ma = \begin{bmatrix} c\theta s\phi mg \\ -s\theta mg \\ -c\theta c\phi mg \end{bmatrix} + \begin{bmatrix} 0 \\ \frac{\rho V^2}{2} \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ -fmg \\ 0 \end{bmatrix} + \begin{bmatrix} f_{sx} \\ f_{sy} \\ f_{sz} \end{bmatrix} + m \left(\frac{d\dot{\omega}}{dt} \times \hat{r}_m \right) + m (\dot{\omega} \times (\dot{\omega} \times \hat{r}_m))$$

From a previous derivation, the relationship between ω and the Euler angles is known. The time derivatives of each may also be taken using the chain rule.

$$\omega = \begin{bmatrix} \omega_x \\ \omega_y \\ \omega_z \end{bmatrix} = \begin{bmatrix} c\phi\dot{\theta} - s\phi c\theta\dot{\psi} \\ \dot{\phi} + s\theta\dot{\psi} \\ s\phi\dot{\theta} + c\phi c\theta\dot{\psi} \end{bmatrix}$$

$$\frac{d\omega}{dt} = \begin{bmatrix} \dot{\omega}_x \\ \dot{\omega}_y \\ \dot{\omega}_z \end{bmatrix} = \begin{bmatrix} c\phi\ddot{\theta} - s\phi\dot{\phi}\dot{\theta} - s\phi c\theta\ddot{\psi} - c\phi c\theta\dot{\phi}\dot{\psi} + s\phi s\theta\dot{\psi}\dot{\theta} \\ \ddot{\phi} + s\theta\ddot{\psi} + c\theta\dot{\psi}\dot{\theta} \\ s\phi\ddot{\theta} + c\phi\dot{\phi}\dot{\theta} + c\phi c\theta\ddot{\psi} - s\phi c\theta\dot{\psi}\dot{\phi} - c\phi s\theta\dot{\psi}\dot{\theta} \end{bmatrix}$$

The following cross products may be calculated; the second is the Euler acceleration:

$$\dot{\omega} \times \hat{r} = \begin{bmatrix} \omega_y r_z - r_y \omega_z \\ -\omega_x r_z + r_x \omega_z \\ \omega_x r_y - r_x \omega_y \end{bmatrix} \quad \frac{d}{dt} \dot{\omega} \times \hat{r} = \begin{bmatrix} \dot{\omega}_y r_z - r_y \dot{\omega}_z \\ -\dot{\omega}_x r_z + r_x \dot{\omega}_z \\ \dot{\omega}_x r_y - r_x \dot{\omega}_y \end{bmatrix}$$

The centripetal acceleration may also be calculated:

$$\vec{\omega} \times (\vec{\omega} \times \vec{r}) = \begin{bmatrix} (-\omega_x r_z + r_x \omega_z) r_z - (\omega_x r_y - r_x \omega_y) r_y \\ -(\omega_y r_z - r_y \omega_z) r_z + (\omega_x r_y - r_x \omega_y) r_x \\ (\omega_y r_z - r_y \omega_z) r_y - (-\omega_x r_z + r_x \omega_z) r_x \end{bmatrix}$$

Substitution of these cross products into the vector-matrix equation yields this final, general form:

$$m\vec{a} = \begin{bmatrix} c\theta s\phi mg \\ -s\theta mg \\ -c\theta c\phi mg \end{bmatrix} + \begin{bmatrix} 0 \\ -\frac{\rho V^2}{2} \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ -fmg \\ 0 \end{bmatrix} + \begin{bmatrix} f_{sx} \\ f_{sy} \\ f_{sz} \end{bmatrix} + m \left(\begin{bmatrix} \omega_y r_z - r_y \omega_z \\ -\omega_x r_z + r_x \omega_z \\ \omega_x r_y - r_x \omega_y \end{bmatrix} + \begin{bmatrix} (-\omega_x r_z + r_x \omega_z) r_z - (\omega_x r_y - r_x \omega_y) r_y \\ -(\omega_y r_z - r_y \omega_z) r_z + (\omega_x r_y - r_x \omega_y) r_x \\ (\omega_y r_z - r_y \omega_z) r_y - (-\omega_x r_z + r_x \omega_z) r_x \end{bmatrix} \right)$$

16.5 Suspension and / or Soil Modelling

Thus far, we have posited the existence of a terrain reactive force acting upon the rigid body at four discrete locations. In reality, these normal forces do not act on a rigid body - they act through yielding terrain and a dynamical system comprised of suspension elements. Both suspensions and terrain response may be modelled using the lumped parameter method; in essence, the dynamical action of these elements is analogous to a filter acting upon the normal force. Regardless of the complexity of these dynamical models, at some point the terrain / tire interface will be specified by a rigid displacement - the magnitude of the terrain / suspension forces then being calculated through the model to find the magnitude of the force acting on the rigid body.

17. INSIGHTS FROM THE SYSTEM MODEL

The system model embodies, in compact form, the physical constraints that the system imposes upon the design of an autonomous speed control system. The model can be used to construct this qualitative description of the development process: The system model contains a number of column vectors, each of which represents an applied or apparent force. The control axes of the machine are represented by the reaction column vector whose components will be altered by the impressed steering and powertrain forces. In general, the forces applied to the system state may be divided into *control* and *disturbance* forces. One part of the autonomous speed control problem then becomes the design of an algorithm that *produces control forces such that they reject disturbance forces while maintaining stability*. More sophisticated layers may be added that optimize these processes against consumables or other costs.

The model does not make clear whether any of the disturbance forces are in fact significant to the speed control problem. The next chapter will produce a set of comparative analyses that indicate if and when each disturbance force becomes significant enough to warrant compensation. The functional dependences of propulsive force generation have also not been developed. Since the significance of disturbance forces are relative to the ability of the powertrain to physically reject them, a model of propulsive force generation will be developed in an upcoming chapter.

CHAPTER V - DISTURBANCE FORCE ANALYSIS

Autonomous speed control in rugged, unstructured terrain is complicated by the presence of disturbance forces whose magnitude is significant when compared to the control force. This chapter develops a means for quantitatively assessing the *significance* of a disturbance to the system. Newton's second law represents the motion of the body; two integrals are considered as possible alternate representations of motion. For each representation, a normalization is applied to produce a comparison of the effects of each disturbance. Dominant regimes may then be identified, and the force and force rates compared to propulsive force magnitudes and generation rates to determine when compensation is necessary.

18. REPRESENTATIONS OF MOTION

The motion of a body is described by Newton's laws. However, Newton's laws may be expressed in a number of different forms. Newton's second law, in its original form, is written:

$$\Sigma \vec{F}_i = \frac{d}{dt} \vec{p}$$

The *impressed forces* act to change what Newton called the *quantity of the motion*, here represented by \vec{p} , and more commonly called the *momentum*. Integrating both sides with respect to time leads to another form, which is called the *impulse-momentum form of Newton's second law*:

$$\vec{p} = \int (\Sigma \vec{F}_i) dt = \Sigma \int \vec{F}_i dt$$

The terms on the right, which are the time integrals of the impressed forces are called *impulses*. If each impulse is denoted \vec{J}_i , then we can more compactly express this form as:

$$\vec{p} = \Sigma \vec{J}_i$$

The impulse momentum equation considers the result of the impression of a force *over time*. A third expression, called the *work-energy form of Newton's second law*, can be obtained by considering the impression of a force *through a distance*. Starting with Newton's second law, rewrite \vec{p} as the product of the mass and the velocity, $m\vec{v}$, and assume that the mass is constant.

$$\Sigma \vec{F}_i = \frac{d}{dt} \vec{p} = \frac{d}{dt} (m\vec{v}) = m \frac{d}{dt} \vec{v}$$

Integrate both sides through the distance dx .

$$\int (\Sigma \vec{F}_i) dx = \Sigma \int \vec{F}_i dx = \int m \frac{d}{dt} v dx$$

Realize that \mathbf{F} and \mathbf{v} have the same vector direction, along the line of integration, which would allow us to reduce the vector integration to a scalar one. This equation is written with a summation; it is only the resultant force of that summation that has the same vector direction as the resultant velocity of the mass. In order to preserve the summation, invoke the principle of superposition, and realize that this integration can be performed for each of the impressed forces and then summed. The summation and the integral can be interchanged.

$$\Sigma \int F_i dx = \int m \frac{dv}{dt} dx$$

The term on the left, in the previous equation, is defined to be the *work done by the force F*. The terms on the right integrate to the *change in kinetic energy of the mass*.

$$\Sigma W_i = \int m \frac{dv}{dt} dx = m \int \frac{dx}{dt} dv = m \int v dv$$

$$\Sigma W_i = \frac{1}{2} m v_2^2 - \frac{1}{2} m v_1^2$$

This work-energy relationship is unique among the three forms because it is the only scalar representation of the result of the action of an impressed force. It describes the result of a force acting through a distance as a change in the mass property called the *kinetic energy*.

19. RELATIVE REPRESENTATIONS

Each of the three forms of Newton's second law may be *normalized* to form a new relative expression.

19.1 Relative Force and Relative Strength

The sum of the forces is equal to the total force, denoted F_t . If the individual forces are normalized by the total force, a set of dimensionless force ratios follows:

$$F_1 + F_2 + F_3 + \dots = F_t$$

$$N_i = \frac{F_i}{F_t}$$

$$\frac{F_1}{F_t} + \frac{F_2}{F_t} + \frac{F_3}{F_t} + \dots = 1$$

$$N_1 + N_2 + N_3 + \dots = 1$$

In a similar vein, let the sum of the absolute value of the forces be called the total strength and denoted S_t . Let the ratio of a force magnitude to the total strength be called the *relative strength*, S_i :

$$\begin{aligned} |F_1| + |F_2| + |F_3| + \dots &= S_t & S_i &= \frac{|F_i|}{S_t} \\ \frac{|F_1|}{S_t} + \frac{|F_2|}{S_t} + \frac{|F_3|}{S_t} + \dots &= 1 & S_1 + S_2 + S_3 + \dots &= 1 \end{aligned}$$

19.2 Relative Impulse Representation

Using the impulse-momentum relationship, the effect of an applied force through time can be examined. The change in momentum of a force F , applied through time is equal to the impulse J :

$$J = \int \vec{F} dt$$

This impulse may be normalized by the momentum of the body. Call this the *normalized impulse*, and define it to be:

$$J_n = \frac{\int \vec{F} dt}{mv}$$

19.3 Relative Energy Representation

The work-kinetic energy relationship quantifies the effect of a force applied through a distance. The change in the kinetic energy of a body produced by a force F , applied through a distance is equal to the work W :

$$\Delta KE = W = \int F dx$$

This work can be normalized by the kinetic energy of the body. Call this the *normalized work*, and define it to be:

$$W_n = \frac{\int F dx}{\frac{1}{2}mv^2} = \frac{2 \int F dx}{mv^2}$$

19.4 Summary

Each of the three relative representations offers a different perspective on the results of an application of a force on a body. The *relative strength of a force* can be used to find regimes in

which particular forces dominate the body. This is especially useful when designing a control system, as it allows the engineer to determine which forces may require special attention, as opposed to those that can be treated as random disturbances. The *normalized momentum* is useful in the analysis of the implication of an “obstacle” during navigation. If the notion of an obstacle is expanded from a solid physical object, like a brick wall, to anything that applies a force that tends to impede progress, then the *significance* of the obstacle relative to the momentum of our vehicle can be calculated. This will be shown in more detail in the next section. The *normalized kinetic energy* similarly allows a useful analysis of an obstacle. The obvious difference between the normalized momentum and the normalized kinetic energy representations is that one considers a force moved through a distance, while the other considers a force applied through time.

20. RELATIVE STRENGTH ANALYSIS

In this section, the relative formulations will be applied to the problem using the three externally applied forces from the second section. Recall that it has been shown that the externally applied forces take the following form:

$$F_r = fmg \quad F_g = mg \sin \Theta \quad F_d = \frac{1}{2} \rho V^2 C_d A$$

20.1 Relative Strength of the Rolling Resistance

The relative strength of the rolling resistance is found to be:

$$S_r = \frac{fmg}{mg(f + |\sin \Theta|) + \frac{1}{2} \rho V^2 C_d A} = \frac{f}{f + |\sin \Theta| + \frac{\rho V^2 C_d A}{2mg}}$$

To simplify the denominator, define a constant γ , and call it the *relative strength coefficient*.

$$\gamma = \frac{\rho C_d A}{2mg}$$

$$S_r = \frac{f}{f + |\sin \Theta| + \gamma V^2}$$

20.2 Relative Strength of the Gravity Load

The relative strength of the gravity load is found to be:

$$S_g = \frac{|\sin\Theta|}{f + |\sin\Theta| + \gamma V^2}$$

20.3 Relative Strength of the Drag Load

The relative strength of the drag load is found to be:

$$S_d = \frac{\gamma V^2}{f + |\sin\Theta| + \gamma V^2}$$

21. RELATIVE IMPULSE ANALYSIS

21.1 Relative Impulse of the Gravity Load

Gravity loads are caused by terrain deviations. Large, long terrain deviations are things like hills, while small terrain deviations are things like rocks and potholes. A constant terrain deviation can be characterized by its angle of inclination Θ and its length L . Assuming that the deviation is to be traversed at the current velocity V , it will do so in time $\Delta t = L/V$. Using this relationship, solve for the relative impulse of a gravity load as follows:

$$I_g = \frac{mg \sin\Theta (\Delta t)}{mV} = \frac{g \sin\Theta (\Delta t)}{V} = \frac{g \sin\Theta L}{V^2}$$

21.2 Relative Impulse of the Drag Force

The drag load is caused by the motion of the body through a viscous fluid, in this case air. After writing the relative impulse, note that a substitution can be made using the non-dimensional term γ , as well as a substitution of the distance travelled L for the product $V\Delta t$.

$$I_d = \frac{\frac{1}{2}\rho V^2 C_d A (\Delta t)}{mV} = \gamma g L$$

This particular ratio is interesting because it is *not dependent upon the speed at which the vehicle travels, or the time of traversal*, but only on the distance travelled. This seems counter-intuitive,

especially considering that the drag force is a function of velocity. You can convince yourself of the truth of the final ratio by considering that the drag force rises quadratically in velocity, while the momentum rises linearly, yielding a linear ratio. Since the definition of impulse requires multiplication by the elapsed time, an expression of distance results. For any distance travelled, if travelled at a higher speed, less time is needed. Thus while the ratio of the drag force to the momentum increases linearly, the elapsed time decreases linearly.

21.3 Relative Impulse of the Rolling Resistance

Rolling resistance is caused by the interaction of the vehicle tires with the soil, neglecting the work done in compacting or bulldozing soil. Writing the relative impulse and substituting, as before, for Δt produces the following ratio:

$$I_r = \frac{fmg(\Delta t)}{mV} = \frac{fgL}{V^2}$$

21.4 Summed Relative Impulse

Summing the individual relative impulses, leads to the following expression.

$$I_T = I_d + I_g + I_r = \gamma gL + \frac{g \sin \Theta L}{V^2} + \frac{fgL}{V^2}$$

Grouping terms and defining the parameter λ , yields:

$$I_T = \gamma gL + \frac{gL}{V^2} (f + \sin \Theta)$$

$$\lambda = \frac{1}{V^2}$$

$$I_T = \gamma gL + \lambda gL (f + \sin \Theta)$$

For the sake of simplicity, let's define the parameters $\hat{\gamma}$ and $\hat{\lambda}$, equal to γg and λg :

$$\hat{\gamma} = \gamma g$$

$$\hat{\lambda} = \lambda g$$

$$I_T = L (\hat{\gamma} + \hat{\lambda} (f + \sin \Theta))$$

22. RELATIVE WORK ANALYSIS

The relative works are derived in a fashion nearly identical to the relative impulses.

22.1 Relative Work of the Gravity Load

The ratio of the work done by the gravity load over a distance L , to the current kinetic energy of the body is found to be:

$$W_g = \frac{mg \sin \Theta L}{\frac{1}{2}mV^2} = \frac{2g \sin \Theta L}{V^2}$$

22.2 Relative Work of the Drag Load

The ratio of the work done by the aerodynamic drag force to the current kinetic energy of the body is found to be:

$$W_d = \frac{\frac{1}{2}\rho V^2 C_d A L}{\frac{1}{2}mV^2} = 2\hat{\gamma}L$$

Again using the previously defined constant $\hat{\gamma}$.

22.3 Relative Work of the Rolling Resistance

The ratio of the work done by the terrain interaction to the current kinetic energy of the body is found to be:

$$W_r = \frac{fmgL}{\frac{1}{2}mV^2} = \frac{2fgL}{V^2}$$

22.4 Summed Relative Works

Sum the relative work done by each external force, and call this the total relative work, W_T .

$$W_T = W_d + W_g + W_r = 2\hat{\gamma}L + \frac{2g \sin \Theta L}{V^2} + \frac{2fgL}{V^2}$$

Now recall the constant $\hat{\lambda}$, defined as follows:

$$\hat{\lambda} = \frac{g}{V^2}$$

$$W_T = 2L(\hat{\gamma} + \hat{\lambda}(f + \sin\Theta))$$

23. PROBLEM-SPECIFIC CALCULATIONS

In this section, calculations specific to the problem at hand, cross country navigation of the NavLab II mobile robot, are presented. This mode of navigation is typified by the following characteristics:

1. Navigation at low speeds, typically 1 m/sec to 10 m/sec. Here I mean to imply that even a human would typically drive the vehicle at these speeds, owing to the discomfort of higher speeds.
2. Navigation across terrain whose vertical geometry rapidly changes. Off-road terrain is filled with both low frequency geometric challenges (*hills and valleys*) and high frequency challenges (*potholes & rocks*).

23.1 Force Analysis

23.1.1 Gravity Loads

The HMMWV weight is approximately 10,200 lbf, thus for every degree of pitch, the load acting upon the powertrain *increases by approximately 790 Newtons or 178 pounds*.

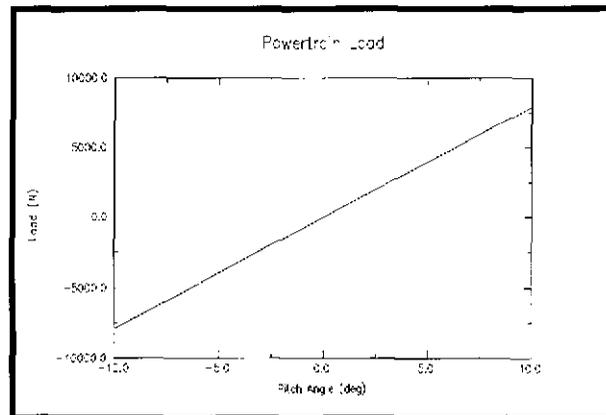


Figure 1 Gravity Load vs. Pitch Angle

The NavLab II weighs approximately five tons and has a wheelbase of approximately 3 meters. If it is assumed that the vehicle is driven along a horizontal plane and transitions to a plane inclined

by 1 degree, then the following curve results. Note that for every meter / sec of vehicle speed, the powertrain load increases at a rate of approximately 265 N/sec or 60 lbf/sec .

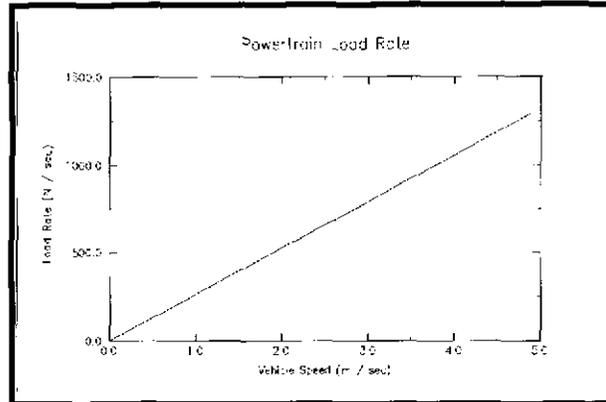


Figure 2 Gravity Load Rate vs. Vehicle Speed

23.1.2 Aerodynamic Loads

In this section, the aerodynamic forces and their time derivatives are calculated using the following numerical estimates for these constants:

$$\rho = 1.22 \frac{\text{kg}}{\text{m}^3} \quad A = 6 \text{m}^2 \quad C_d = 1.0$$

The cross-sectional area is a slight overestimate of that of the true vehicle. The drag coefficient is difficult to accurately obtain. For this analysis, I justify the use of the value 1 as follows: the nominal value of the drag coefficient for modern automobiles ranges from ~ 0.32 for sports cars to ~ 0.45 for pickup trucks. The HMMWV is not noted for its aerodynamic styling, and it has a number of appendages such as the air conditioning unit, the ERIM and the Staget that further detract from its shapely profile, thus its drag coefficient is surely larger than 0.45. An upper bound for the drag coefficient can be had from the drag coefficient for a flat plate whose surface normal is parallel to the flow, 1.95. The value of 1 is a good overestimate of a vehicle drag coefficient. Substituting these values into the given equation yields a relationship between vehicle speed and drag force:

$$F_d = 3.66V^2$$

The following graphs illustrate the growth of drag force as a function of velocity. The range of the first graph is 0 to 30 m / sec (0 to 67 m.p.h), the range of the second graph is 0 to 10 m / sec (0 to 22 m.p.h).¹

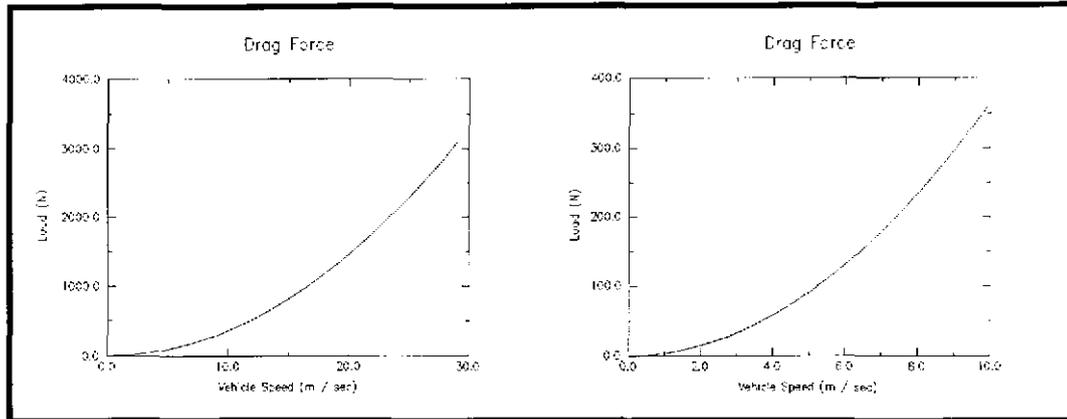


Figure 3 Aerodynamic Load vs. Vehicle Speed

It is unlikely that the HMMWV can accelerate at better than 0.3 g's. The change in drag force associated with this acceleration at 10 m/sec is 215 N / sec, (48 lbf / sec).

23.1.3 Rolling Resistance Loads

For large trucks operating in moderately packed soil, f ranges from 0.15 to 0.20. This implies that the rolling resistance is of the order of 2000 lbf.

23.2 Relative Strength Analysis

23.2.1 Relative Strength Coefficient

In a previous section, the concept of the relative strengths of the three forces were defined. The proportionality constant among these forces was called the relative strength coefficient γ . For the case of the HMMWV, the coefficient is calculated to be:

$$\gamma = \frac{\rho C_d A}{2mg} = \frac{\left(1.22 \frac{\text{kg}}{\text{m}^3}\right) (1) (6\text{m}^2)}{10200\text{lbf} \times 4.45 \left(\frac{\text{N}}{\text{lbf}}\right)} = 1.61 \times 10^{-4} \left(\frac{\text{s}^2}{\text{m}^2}\right)$$

1. The range of the first graph corresponds approximately with that of the on-road navigation problem, while that of the second exceeds the speed range of the current off-road navigation problem.

This coefficient is used to determine the speed at which the aerodynamic forces approach the same order as the other forces. Recall the denominator of the relative strength ratios:

$$f + |\sin \Theta| + \gamma V^2$$

For the domain of operation, f is approximately 0.20, while the sin term is in the range of 0 to 0.349 (0 to 20 degrees). The aerodynamic portion of the denominator approaches the same order when:

$$\gamma V^2 = 0.20$$

$$V = \sqrt{\frac{0.20}{\gamma}} = \sqrt{1242} = 35 \frac{m}{s}$$

We might say that the aerodynamic forces become significant at a much lower threshold, say 10% of the rolling resistance, in which case the corresponding velocity would be 11 m/sec.

23.2.2 Gravity Relative Strength

Recall the equation for the gravity relative strength:

$$S_g = \frac{|\sin \Theta|}{f + |\sin \Theta| + \gamma V^2}$$

Using the previously determined values for f and γ , and assuming a velocity of 5 m/sec, the equation reduces to:

$$S_g = \frac{|\sin \Theta|}{0.2 + |\sin \Theta|}$$

Plotting this equation for angles between 0 and 20 degrees:

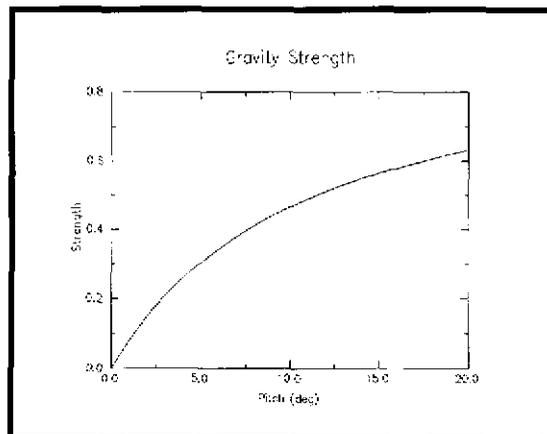


Figure 4 Gravity Strength vs. Pitch Angle

23.2.3 Aerodynamic Relative Strength

Recall the equation for the aerodynamic relative strength:

$$S_d = \frac{\gamma V^2}{f + |\sin \Theta| + \gamma V^2}$$

Using the previously selected values for f and γ , and choosing a velocity of 5 m / sec, this equation reduces to:

$$S_d = \frac{.0161}{0.21 + |\sin \Theta|}$$

Plotting this equation for angles between 0 and 20 degrees:

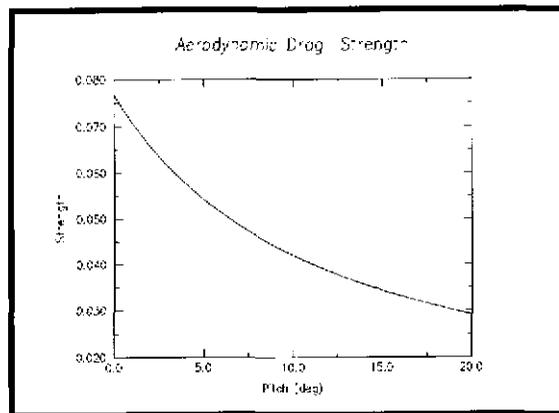


Figure 5 Aerodynamic Strength vs. Pitch

23.2.4 Rolling Friction Strength

Recall the equation for the aerodynamic relative strength:

$$S_r = \frac{f}{f + |\sin \Theta| + \gamma V^2}$$

Using the previously selected values for f and γ , and choosing a velocity of 5 m / sec, this equation reduces to:

$$S_d = \frac{0.2}{0.21 + |\sin \Theta|}$$

Plotting this equation for angles between 0 and 20 degrees:

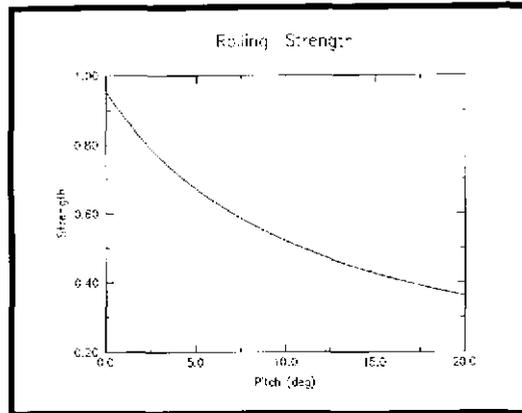


Figure 6 Rolling Friction Strength vs. Pitch Angle

23.2.5 Comparative Relative Strengths

If the three strength functions are plotted on the same axes, a direct comparison can be made:

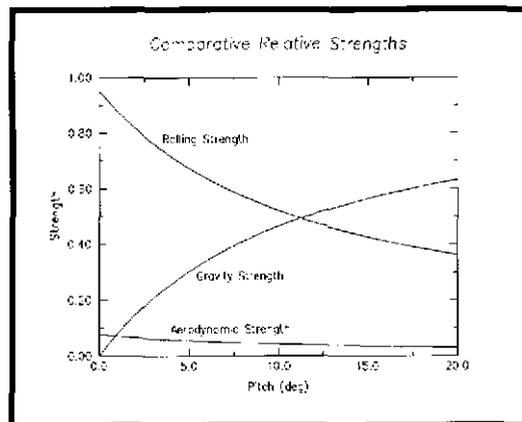


Figure 7 Comparison of Relative Strengths

Note the crossover point for rolling strength and gravity strength at approximately 10° of pitch. Under these conditions, for pitch angles less than 10° the rolling resistance dominates, for pitch angles greater than 10° , the gravity strength dominates. If the environment in which the vehicle will operate tends to be flat, then it is the case that the rolling forces always dominate the other two. Significant slopes indicate that a trade-off between dominance may occur between the two, possibly requiring that the control solution compensate for the change in dynamics. We should note that 10° is a rather extreme pitch angle.

23.3 Relative Impulse Analysis

23.3.1 Gravity Relative Impulse

Recall the gravity relative impulse equation:

$$I_g = \frac{g \sin \Theta L}{V^2}$$

The equation can be written on a distance-specific basis:

$$\frac{I_g}{L} = \frac{g \sin \Theta}{V^2}$$

The following plot contains a set of graphs of relative impulse for pitch angles between 0 and 20 degrees, at five different speeds.

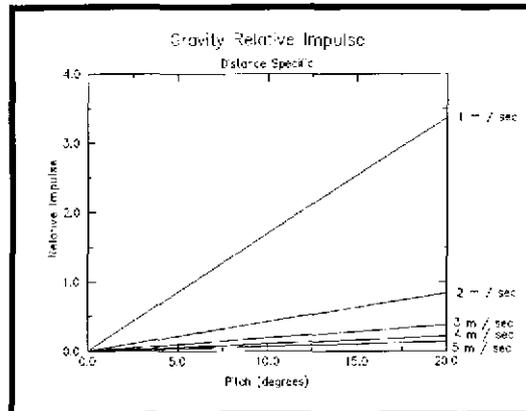


Figure 8 Gravity Relative Impulse vs. Pitch Angle

Note that when the relative impulse exceeds unity, the obstacle presents an impediment equal to all of the momentum of the vehicle. Without the addition of vehicle momentum, the vehicle cannot mount the obstacle. The relative impulse gives a quantitative measurement of the amount of additional momentum necessary to mount the obstacle.

23.3.2 Drag Force Relative Impulse

Recall the drag force relative impulse equation:

$$I_d = \gamma g L$$

The equation can be written on a distance-specific basis, and its constant value calculated:

$$\frac{I_d}{L} = \gamma g = 0.001$$

23.3.3 Rolling Resistance Relative Impulse

Recall the rolling resistance relative impulse equation:

$$I_r = \frac{fgL}{v^2}$$

The equation can be written on a distance-specific basis, and its constant value calculated:

$$\frac{I_r}{L} = \frac{fg}{v^2}$$

The following graphic plots the relative impulse as a function of the velocity of the vehicle:

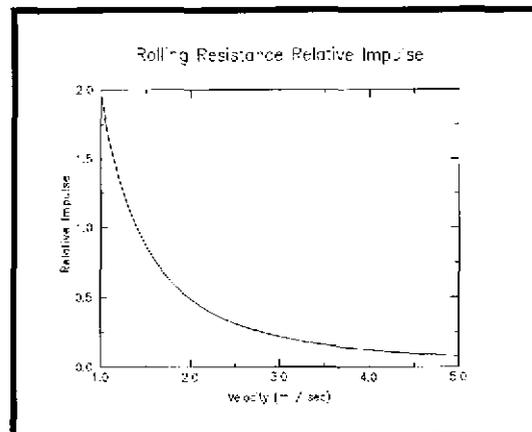


Figure 9 Rolling Resistance Relative Impulse vs. Vehicle Velocity

24. DISCUSSION AND CONCLUSIONS

24.1 Relative Force and Relative Strength

The concept of relative strength is a useful metric for examining the effects of the external forces on the body, and dividing a performance spectrum up into a set of sub-spectra, each of which is dominated by a particular force. The relative strength is an expression of the relative *magnitudes* of the forces, thus avoiding the problem of comparison at or near mechanical equilibrium. In this fashion, relative strength is a quite different comparison from relative force. The relative strength

compares a force to a set of other forces. However, the relative force compares a force to the change in state of the body. When attempting to determine whether particular disturbances dominate the problem, the concept of relative strength is more useful. For the NavLab II analyses it seems apparent that gravitational load and rolling friction dominate the problem. Aerodynamic drag is such a small force in comparison that it does not seem to be necessary to compensate for it.

24.2 Relative Impulse and Relative Work

The concepts of relative impulse and relative work are useful metrics for examining the *significance* of an external force.¹ Significance is an odd word to use in mechanics - what I mean here is that intuitively, a hill (gravity load) that is 10 miles long probably has more of an effect on a vehicle than a similar hill that is only 10 feet long. And either hill is probably more significant to a slow moving vehicle than to one that is already travelling along at 100 m.p.h. The concept of relative momentum is therefore useful for distinguishing the cruise control problem from the autonomous speed control problem in rugged, unstructured terrain. It was shown that at low speeds the relative impulse of a hill approaches a singularity.

1. Relative work differs from relative impulse mainly in that it is a scalar concept.

CHAPTER VI - CONTROL FORCE MODELS

The linear control algorithm experiments performed early in this work indicated that the generation of propulsive force through the control of the engine throttle plate is a non-linear process. This chapter produces a model of the engine and powertrain that is used later in the production of integrated motion control algorithms. A second model of the braking dynamics is also produced.

25. PROPULSION FORCE

A conventional vehicle generates a propulsive force through the combined actions of an engine and a powertrain. The engine is a device which supports a chemical reaction to produce rotary motion. The powertrain is a device which acts as a mechanical transformer, to reduce rotary velocities in favor of torque.

25.1 An Engine Cycle

Most engines operate on the four-stroke cycle. The cycle is named for the four strokes made by a single piston during a complete cycle. During the *intake stroke*, an intake valve in the head of the cylinder opens and the piston draws down, pulling fresh air into the cylinder. At the bottom of the stroke, the valve closes; the piston rises during the *compression stroke*, causing a sudden rise of temperature with the sudden compression. At or near the top of the compression stroke, fuel is injected and ignited. The combustion gases expand and force the piston back down during the *power stroke*, it is during this stroke that the piston is doing useful work. When the piston reaches the bottom of the power stroke, an exhaust valve is opened; during the *exhaust stroke*, the spent gases of combustion are forced from the cylinder. When the piston reaches the top of the stroke, the exhaust valve closes, the intake valve opens and the cycle begins again.

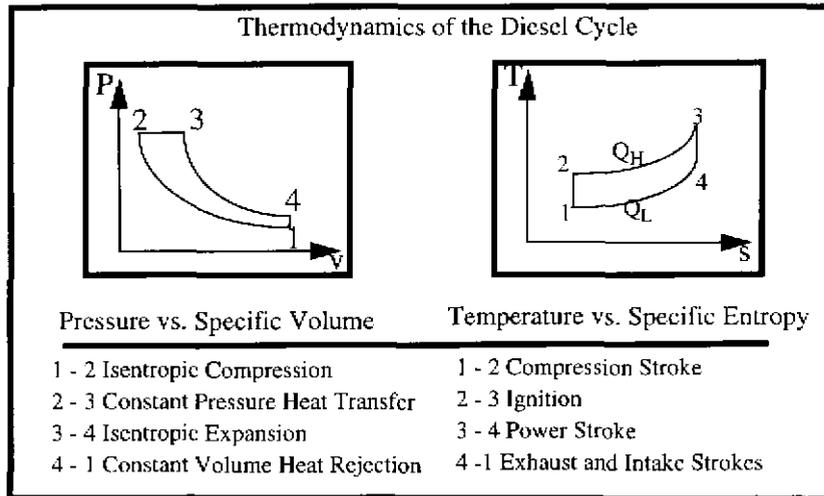
25.2 The Diesel Engine and Diesel Cycle

The Diesel engine is a compression ignition engine, meaning that the compression of the combustibles alone is sufficient to cause ignition, there is no need for an electric spark. The thermodynamics of the Diesel cycle generally follow the actions of the engine cycle. The first diagram below is a graph of pressure vs. specific volume¹. The engine cycle begins with the piston fully retracted and the air within the cylinder at ambient pressure at (1). This corresponds to the end of the intake stroke. From (1) to (2) the piston compresses the air fuel mixture; note that the specific volume decreases and the pressure increases. From (2) to (3) the pressure remains constant and the specific volume increases - this corresponds to the instant at which the fuel ignites. The motion from (3) to (4) corresponds to the power stroke, and (4) to (1) corresponds to both the exhaust and subsequent intake stroke.²

1. Thermodynamic quantities are calculated on a per mass basis; specific volume is volume per unit mass, the inverse of density. The calculations in this section will find the work per mass unit of working fluid.

2. The thermodynamic cycle is actually open ended, the final section is an idealized closure of the open cycle.

The work performed during a heat engine cycle is obtained from the accompanying temperature vs. entropy diagram. The net work of the engine is the area enclosed by the cycle; it is the difference between the heat addition Q_H and the heat loss Q_L .



The net work of the engine can be written as follows.

$$w_{net} = q_H - q_L$$

$$q_H = C_p (T_3 - T_2)$$

$$q_L = C_v (T_4 - T_1)$$

Note that as the ambient temperature T_1 is reduced, q_L decreases, causing the net work performed by the cycle to increase.

25.3 Control of the Internal Combustion Engine

One controls the rate at which an engine produces work by controlling the flow of the combustibles into the combustion chamber. From thermodynamics, we know that net work (w_{net}) produced by a heat cycle is given on a *specific* or *per mass* basis. The power produced by a heat engine is found by multiplying the cycle's net work by the mass flow rate of the combustibles:

$$P = \dot{m} w_{net}$$

$$\left[\frac{kJ}{s} \right] = \left[\frac{kg}{s} \right] \left[\frac{kJ}{kg} \right]$$

Thus *the power of the engine may be controlled by controlling the mass flow of combustibles*. It is the function of the throttle to allow the driver of a vehicle to control this mass flow; note however, that there is no direct mapping from throttle position to either torque or engine RPM. This point will be central to the development of the thesis.

The mass flow rate of combustibles through the engine is a function of three events: mass flow through the throttle port; mass flow through the intake manifold; and mass flow into the cylinders. As noted by (Cho and Hedrick[3]), the relationship among these three process is described by the following rate equation:

$$\dot{m}_a = \dot{m}_i - \dot{m}_o$$

which states that the rate of mass accumulation in the intake manifold is equal to the difference between the mass flowing in from the throttle (\dot{m}_i) and the mass flowing out to the cylinders (\dot{m}_o). The rate of mass flow out to the cylinder may be described by the following equation:

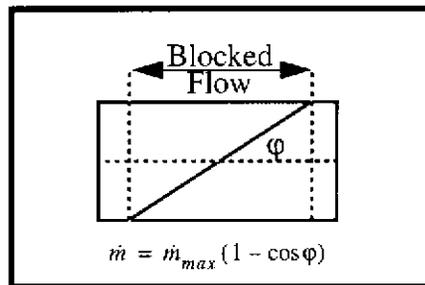
$$\dot{m}_o = c_1 \eta_{vol} \dot{m}_a \omega_e$$

which shows that the flow rate into the cylinders is a function of engine RPM, and the amount of mass in the intake manifold, scaled by a constant c_1 and a volumetric charging efficiency. The volumetric efficiency itself is a rather complicated expression that is a function of engine RPM's and the mass state of the intake manifold.

Rather than attempting to explicitly model these terms, we will opt for a simplifying assumption. For the purposes of this thesis, we will assume that the intake manifold remains in a *steady state*, such that no mass is either accumulated or removed. Under this condition, the rate relationship given above yields:

$$\dot{m}_i = \dot{m}_o$$

The validity of making such an assumption will be supported by the demonstrated performance of the servo-controls algorithm in later chapters. The mass flow into the intake manifold, through the throttle port may be estimated for the rotational valve as follows. One may also note that the mapping from valve position to power is nonlinear for rotational valves due to the valve projection geometry.



The force caused by the expansion of the combustibles acts through the slider crank mechanism to produce a torque on the crankshaft. The equation of rotational dynamics can be written by

considering the crankshaft and other rotational components to be an inertia I_c . The engine torque

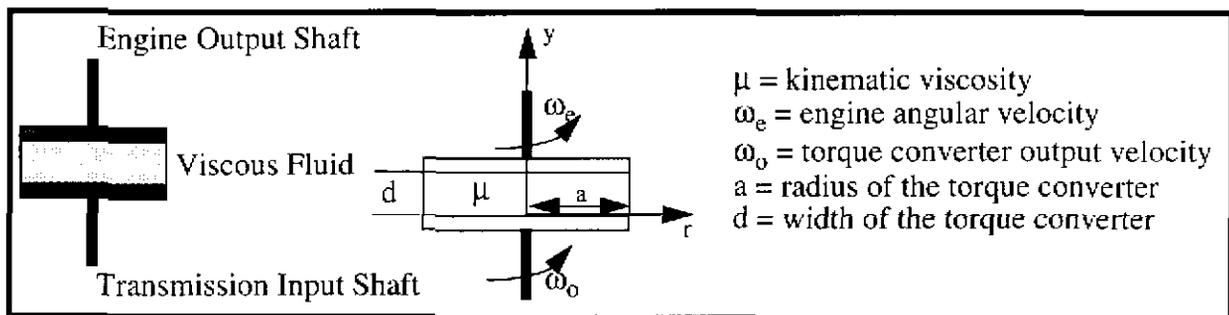
$$\Sigma \Gamma = I_c \dot{\omega}_e$$

$$\Gamma_e - \Gamma_l - \Gamma_o = I_c \dot{\omega}_e$$

from combustion (Γ_e) is opposed by the load torque (Γ_l) and the output shaft torque (Γ_o). The load torque is caused by internal engine frictions, engine accessory loads (such as air conditioner compressors, alternators, power steering pumps, etc.) and other miscellaneous engine loads. The output shaft load is caused by the load of the component being driven by the engine's output shaft. The rotational velocity of the engine (ω_e) (which is more commonly expressed in RPM¹) is thus dependent upon the time history of engine loads.

26. TORQUE CONVERTER

A transmission is a mechanical device that takes as its input a rotary load, and produces as its output a rotary load, scaled in torque and angular velocity. Conventional vehicles are equipped with transmissions that have several gear sets, allowing the driver to vary the input to output ratio. Manual transmissions couple the engine output to the transmission input using friction clutches. Automatic transmissions couple using a fluid clutch, also called a torque converter. A torque converter has two internal fans, called the rotor and the stator. These two fans face each other in a sealed chamber filled with a viscous fluid. The rotor's shaft is attached to the engine output shaft; the stator shaft is the input to the transmission. The engine turns the rotor, which causes the viscous fluid to flow with the rotor; the rotation of the viscous fluid in turn causes the stator to turn. In this manner, the engine torque is transferred to the transmission through the fluid; there is no mechanical coupling. A simple model of torque conversion comes from ideal planar rotational Couette flow. In the diagram below, the stator and the rotor are modelled as identical flat plates, facing one another, with a viscous fluid trapped between them.



1. RPM is Revolutions per minute.

Assuming a linear velocity distribution along the y direction of the fluid, the rotational velocity component u_{Θ} can be expressed as follows:

$$u_{\Theta} = -r \left(\omega_e \frac{y}{d} + \omega_o \frac{(d-y)}{d} \right) = -r \left((\omega_e - \omega_o) \frac{y}{d} + \omega_o \right)$$

The specific rate of transfer of Θ momentum in the y direction, or the flux density of Θ momentum in the y direction can be calculated as follows.

$$j_{y\Theta} = -\mu \frac{du_{\Theta}}{dy}$$

$$\frac{du_{\Theta}}{dy} = -\frac{r}{d} (\omega_e - \omega_o)$$

For a Newtonian fluid, the shear stress, $\tau_{y\Theta}$, is equal to the rate of momentum transfer $j_{y\Theta}$. The torque on the output shaft, T_o , caused by the momentum transfer is found by integrating the shear stress across the area of the output plate.

$$T_o = \int_0^a (r \tau_{y\Theta} 2\pi r) dr$$

Substituting for $\tau_{y\Theta}$ and carrying out the integration yields:

$$T_o = K\mu (\omega_e - \omega_o)$$

where K is a geometric constant.

27. THE COMBINED POWERTRAIN EQUATIONS

The propulsive force F_p is found from the torque converter equation. The output torque, applied through a radius R, produces a force F_p :

$$F_p = T_o R = KR\mu (\omega_e - \omega_o) = K_2 (\omega_e - \omega_o)$$

where K_2 is a combined constant. The output angular velocity ω_o is proportional to the vehicle speed, allowing the equation to be rewritten as:

$$F_p = (K_2 \omega_e - K_3 V)$$

The engine angular velocity is the state variable of the engine dynamic system, and may be found as follows. The engine is a power output device; engine power is equal to the product of engine torque and engine angular velocity. Engine power is also proportional, from combustion considerations, to the mass flow rate of the combustants. The mass flow rate for the HMMWV can be approximated by the given valve flow equation.

$$\begin{aligned} P &= \Gamma_e \omega_e = I_e \dot{\omega}_e \omega_e \\ P &= \dot{m} w_{net} \\ \dot{m} &= \dot{m}_{max} (1 - \cos \varphi) \end{aligned}$$

Setting the two power equations equal to one another yields a differential equation for engine angular velocity.

$$I_e \dot{\omega}_e \omega_e = \dot{m} w_{net}$$

Substitution for the mass flow rate relates ω_e to the input variable φ , which is the valve rotation.

$$I_e \dot{\omega}_e \omega_e = \dot{m}_{max} (1 - \cos \varphi) w_{net}$$

Isolating the change of state, yields the following equation, where K_e is a combining constant.

$$\dot{\omega}_e = \frac{K_e}{I_e \omega_e} (1 - \cos \varphi)$$

RPM is proportional to ω_e , thus the equation may be rewritten as follows, noting that the numerical value of K_e changes accordingly.

$$R\dot{P}M = \frac{K_e}{I_e (RPM)} (1 - \cos \varphi)$$

Similarly the output force equation may be modified to RPM:

$$F_p = (K_{RPM} RPM - K_v V)$$

28. TOTAL DYNAMICAL SYSTEM AND COMPENSATION

In this section we construct a total system model by combining the vehicle model with the powertrain model. (These models were developed in Chapters IV and V, respectively). The powertrain dynamical model describes the vehicle's propulsion force, which is one term in the larger vehicle system model. The production of a total system model allows us to revisit the

disturbance forces relevant to the problem and note the type of compensation that may be performed to counter each.

28.1 Vehicle Model

The following vector matrix vehicle model was derived in Chapter IV. Each column vector represents a force acting upon the rigid body model of the vehicle. The columns, in order from left to right, represent:

- Gravity Force
- Aerodynamic Drag
- Rolling Resistance
- Terrain or Suspension Reaction Forces
- Apparent Forces

$$ma = \begin{bmatrix} c\theta s\theta mg \\ -s\theta mg \\ -c\theta c\phi mg \end{bmatrix} + \begin{bmatrix} 0 \\ -\frac{\rho V^2}{2} \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ -fmg \\ 0 \end{bmatrix} + \begin{bmatrix} f_{sx} \\ f_{sy} \\ f_{sz} \end{bmatrix} + m \left(\begin{bmatrix} \dot{\omega}_y r_z - r_y \dot{\omega}_z \\ -\dot{\omega}_x r_z + r_x \dot{\omega}_z \\ \dot{\omega}_x r_y - r_x \dot{\omega}_y \end{bmatrix} + \begin{bmatrix} (-\omega_x r_z + r_x \omega_z) r_z - (\omega_x r_y - r_x \omega_y) r_y \\ -(\omega_y r_z - r_y \omega_z) r_z + (\omega_x r_y - r_x \omega_y) r_x \\ (\omega_y r_z - r_y \omega_z) r_y - (-\omega_x r_z + r_x \omega_z) r_x \end{bmatrix} \right)$$

Powertrain control is the problem of regulating or servoing speed along the y-axis, which isolates the following single dimensional equation from the matrix:

$$ma = -s\theta mg + \frac{-\rho V^2}{2} - fmg + f_{sy} + m(-\dot{\omega}_x r_z + r_x \dot{\omega}_z - (\omega_y r_z - r_y \omega_z) r_z + (\omega_x r_y - r_x \omega_y) r_x)$$

Now, consider the terms on the right hand side of the equation from left to right.

• $-s\theta mg$

This term is the projection of the gravity vector along the longitudinal axis of the vehicle. From our relative force analysis it was shown that the magnitude of this force is significant throughout the expected range of operation. Because a model of this disturbance exists, *feedforward compensation* can be performed.

• $-\rho V^2 / 2$:

This term represents the aerodynamic drag force acting along the longitudinal axis of the vehicle. From the relative force analysis, it was shown that the magnitude of this force is relatively small, and essentially constant throughout the expected range of operation. As a small constant force, its effects can be eliminated through the use of *feedback control*.

• $-fmg$:

This term represents the loss of energy due to soil compaction, expressed as a force. Similarly, the magnitude of this force is small. Within the scope of this work it is assumed to be essentially constant, though in practice this may or may not be true. It is assumed that this force may be compensated using *feedback control*.

- $\omega \times r$ projection:

This extended term represents the apparent force generated by the motion of the vehicle in an inertial frame. For the low speeds and reasonably smooth trajectories over which the vehicle is assumed to operate, these forces remain low. It is interesting to note that this projection represents the coupling between steering control and powertrain control. It is assumed that this force may be compensated using *feedback control*.

28.2 Powertrain model

The powertrain model describes the generation of the vehicle's propulsion force; it is encapsulated in two equations. The first is a differential equation relating the engine state (expressed in RPM's) to the angle of the throttle plate, ϕ . The second is an algebraic equation relating the propulsive force to a proportional difference between the engine state and the vehicle state. This equation models the dynamics of the torque converter.

$$R\dot{P}M = \frac{K_e}{I_e(RPM)} (1 - \cos\phi) - \frac{K_\Gamma \Gamma_e}{I_e}$$

$$F_p = (K_{RPM}RPM - K_v V)$$

28.3 System Model

The union of the powertrain model with the previously derived vehicle model yields a total system model. First, combine the propulsive force model with the vehicle force equations by adding a propulsive force F_p , as shown:

$$ma = -s\theta mg + \frac{-\rho V^2}{2} - fmg + F_p + F_{apparent}$$

Substituting the powertrain force model yields a new equation, derived below:

$$m\dot{V} = -s\theta mg + \frac{-\rho V^2}{2} - fmg + F_p + F_{apparent}$$

$$F_p = (K_{RPM}RPM - K_v V)$$

$$m\dot{V} = -s\theta mg + \frac{-\rho V^2}{2} - fmg + K_{RPM}RPM - K_v V + F_{apparent}$$

The apparent forces, terrain friction force and aerodynamic force are grouped as F_{dist} :

$$m\dot{V} = -s\theta mg + K_{RPM}RPM - K_v V + F_{dist}$$

Together, the RPM state equation and this new vehicle equation form the total system equations:

$$m\dot{V} = -K_v V + K_{RPM}RPM + F_{dist} - s\theta mg$$

$$\dot{RPM} = \frac{K_e}{I_e(RPM)}(1 - \cos\varphi) - \frac{K_\Gamma \Gamma_e}{I_e}$$

To simplify our treatment of the problem, we make the following substitution of $MF = 1 - \cos\varphi$:

$$\dot{RPM} = \frac{K_e(MF)}{I_e(RPM)} - \frac{K_\Gamma \Gamma_e}{I_e}$$

where MF represents mass flow.

29. STATE SPACE ANALYSIS

State space control theory offers a number of analytical techniques based upon linear systems theory. In this section, some of these techniques are applied to help characterize the system. The system model indicates a single input / two state system. The two states are coupled, as shown in the torque converter equation.

29.1 Achieving State Space Form

The first term in the RPM differential equation contains an inverse state - $(1 / RPM)$, which makes this state equation nonlinear. To achieve state space form, this equation is linearized about an operating point RPM_0 , using a Taylor's series expansion.

$$\frac{1}{RPM} = \frac{RPM}{RPM^2}$$

$$\frac{1}{RPM} = RPM \left\{ \frac{1}{RPM_0^2} - \frac{2}{RPM_0^3}(RPM - RPM_0) + \frac{3}{RPM_0^4}(RPM - RPM_0)^2 - \dots \right\}$$

Now, replace the Taylor's series expansion with an approximation, P , of the series, comprising the first n terms.

$$\frac{1}{RPM} = RPM(\hat{P})$$

Recall the state space form, and write our current model in those terms:

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

$$\begin{bmatrix} \dot{V} \\ RPM \end{bmatrix} = \begin{bmatrix} -\frac{K_v}{m} & \frac{K_{RPM}}{m} \\ 0 & \frac{K_e \hat{P}}{I_e} \end{bmatrix} \begin{bmatrix} V \\ RPM \end{bmatrix} + \begin{bmatrix} 0 \\ -\frac{K_e}{I_e RPM_0} \end{bmatrix} [\cos \phi]$$

29.2 Controllability Analysis

The controllability matrix M_c describes the reachable state space of the system, and for this system is defined to be:

$$M_c = [B \ AB]$$

From the state-space model above, we may calculate the controllability matrix for this system:

$$B = \begin{bmatrix} 0 \\ -W \end{bmatrix} \quad AB = \begin{bmatrix} \left(-\frac{K_{RPM}}{m} \right) W \\ -\frac{1}{RPM_0} W^2 \end{bmatrix}$$

where $W = K_e / (I_e RPM_0)$.

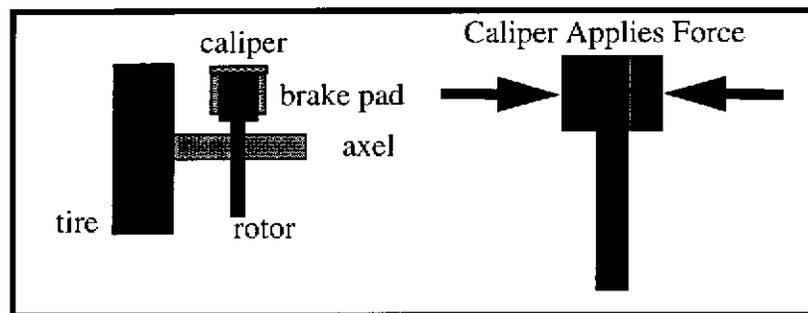
$$M_c = \begin{bmatrix} 0 & \left(-\frac{K_{RPM}}{m} \right) W \\ -W & -\frac{1}{RPM_0} W^2 \end{bmatrix}$$

Note that M_c is of full rank, and so the state space is fully reachable. However, the state space is only fully reachable when both columns of M_c are used, due to the 0 value in the (1,1) position. This indicates that only by rotating the input vector through the transition matrix (the physical system), can we achieve arbitrary placement in state space.

The system, while fully controllable, is coupled. *It is thus physically possible to achieve control, but the system directs the manner in which we perform this control.* Specifically, in order to place the system's velocity we must simultaneously place the system's RPM. This makes some intuitive sense based solely on the number of inputs and states of the system. This is a single input, two state system; in other words, the single input is somehow "smeared" across the two states. *In order to ultimately control the velocity of the vehicle, we must respect and compensate for its coupling to the engine RPM.*

30. BRAKING FORCE

The braking system transforms some of the kinetic energy of the vehicle into heat energy, reducing the vehicle speed in the process. When travelling downgrade, the braking system must also transform the gained potential energy into heat energy. In conventional vehicles, this is most often accomplished through the generation of frictional forces.¹ To generate the braking friction, a metal rotor or drum is attached to the wheel axle. Braking pads are placed against the surface of the rotor, and a force is applied to the pads by the calipers. This force produces a sliding friction, which dumps kinetic energy from the vehicle to the surroundings. The braking force is produced when the driver presses his or her foot against the brake pedal. The pedal is the input to a hydraulic system that multiplies the applied foot force and applies that force to the pad.



30.1 Force Model

Braking systems are specifically designed to provide a linear increase in applied force, and thus a linear increase in deceleration for increased pedal force and pedal travel. (Limpert[20]) Given this relationship, a model relating deceleration to pedal travel may be written as follows:

$$F_b = \frac{\mu N}{x_{max}} x = K_b x$$

1. There are exceptions, such as the Jacobs brake found on most large trucks and tractor trailers.

Application of proportional control to vehicle braking under this force law in the absence of other externally applied forces results in the following equation:

$$\dot{V} = -\frac{F_b}{m} = -\frac{K_b}{m}x$$

Application of the proportional control law to this system results in the following closed loop law:

$$x = -K_p V_e$$

$$\dot{V} = -\frac{F_b}{m} = \frac{K_b K_p}{m} V_e$$

Which may be written as a closed loop transfer function:

$$V = \frac{K_b K_p}{m + K_b K_p} V_d$$

This transfer function matches well with the braking performance of the NavLab II operating on the open terrain of the slag heap flats. In a subsequent chapter, it will be shown that linear braking control can be extended to 3-D terrain in a straightforward manner. We thus conclude here that *improved throttle compensation is the key to achieving good speed control*.

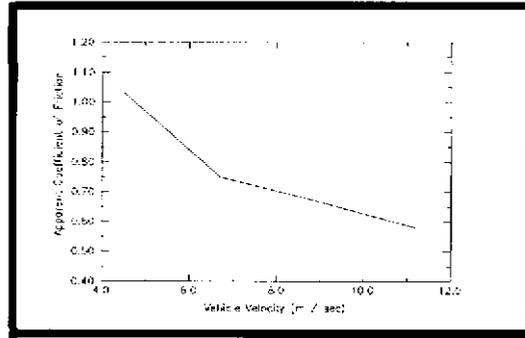
30.2 Friction Limitations

In the Coulomb model of sliding friction, $F = \mu N$. Frictional force is generated in proportion to the normal force applied at the surface. Equating the energy expended over the braking distance Δx , to the kinetic energy dissipated over that distance, yields the following relationship, from which an expression for μ may be derived.

$$\mu mg \Delta x = \frac{1}{2} m v^2$$

$$\mu = \frac{v^2}{2g \Delta x}$$

A simple experiment with the HMMWV yields the following relationship between the apparent coefficient of friction and the vehicle velocity. The apparent coefficient of friction is that calculated from the previous expression, assuming that all energy is depleted from a single friction force of constant coefficient. The experimental results do not match the previous expression because in the experiment, the wheels locked up for some portion of the stopping distance - the vehicle takes longer to stop when its wheels are sliding than it does if they are rolling. The purpose of this experiment was to show the importance of preventing wheel lock-up in practice. Although the braking problem is a force control problem, it is insufficient to consider only the frictional forces applied at the braking surface. One may also need to consider the tire / terrain interaction forces to



ensure that the wheels do not lock. An appropriate solution to this problem would be to install an anti-lock braking system, thus eliminating, or at least reducing, the concern over tire skid. For the purposes of this thesis, terrain interaction phenomena have been largely ignored, and we will not further concern ourselves with friction limitations.

Part III

An Autonomous Speed Control System

The central work of this thesis is the production of an autonomous speed control system. **Part III** describes the design, construction, implementation and results of an autonomous speed control system, beginning with *LITHIA*, a model based speed controller. *LITHIA* closes the throttle and brake loops to provide stable, asymptotic achievement of the commanded speed. *LITHIA* is shown to be extensible to a model based optimal control speed planner through the feedforward solution of a constrained optimization metric. Finally, *SPECTRE*, a frequency space motion evaluator is presented as a computationally efficient method for estimating the response of the passive dynamic system for use in speed and / or motion planning.

CHAPTER VII - LITHIA: A MODEL BASED CONTROLLER

Throttle compensation has been shown to be the principal challenge to attaining autonomous speed control of conventional vehicles operating over unstructured 3-D terrain. This chapter will detail the development of LITHIA - a robust throttle compensation control algorithm; the extension of the brake linear controller to 3-D operations; and the consideration of appropriate switching algorithms for trading control between throttle and brake compensation.

LITHIA is an integrated speed controller whose algorithm is based upon a model of the disturbance forces acting upon the vehicle, and of the dynamics internal to the generation of the vehicle's propulsion force. Using this model, a feedback / feedforward proportional servo controller is designed. LITHIA's construction is both simple and powerful - it models only the basic dynamics necessary to appropriately reject force disturbances and compensate for internal powertrain dynamics. The design of the LITHIA algorithm is driven by both the prior work and the noted responses of the linear control experiments. These experiences motivate the generation of a control algorithm that is capable of:

- Rejecting force disturbances.
- Performing servo control.
- Maintaining algorithmic stability over a wide range of vehicle speeds and terrain conditions.

Analyses performed in this thesis indicate that there are two basic problems which the control algorithm must address. First, the state coupling effect shown in the state space analysis must be overcome. Second, the disturbance forces, especially the gravity and rolling resistance forces, must be compensated. These two issues characterize the construction requirements for stable, robust velocity control. Two generations of LITHIA are presented in this chapter. The first generation, five-gain algorithm is designed directly from the models developed in preceding chapters. Using the analysis methods of classical control, the first generation LITHIA algorithm is re-designed, and reduced in form to an improved two-gain, second generation algorithm.

31. THE LITHIA CONTROL ALGORITHM - FIRST GENERATION

The first generation LITHIA control algorithm was designed to back-solve through the coupling in the transition matrix to determine the appropriate throttle body position for the current physical conditions. It simultaneously performs feedforward to compensate for measured gravity disturbances. Consider the LITHIA algorithm in three separate steps. First, given the desired velocity, compute the RPM that the engine must turn in order to achieve that velocity. This desired RPM is calculated as follows:

$$\begin{aligned}
 Vel_e &= Vel_d - Vel \\
 RPM_d &= M3 (Vel_d) + M4 (Vel) + M5 (Load)
 \end{aligned}$$

Second, given the newly calculated desired RPM, compute the mass flow required to achieve that RPM:

$$RPM_e = RPM_d - RPM$$

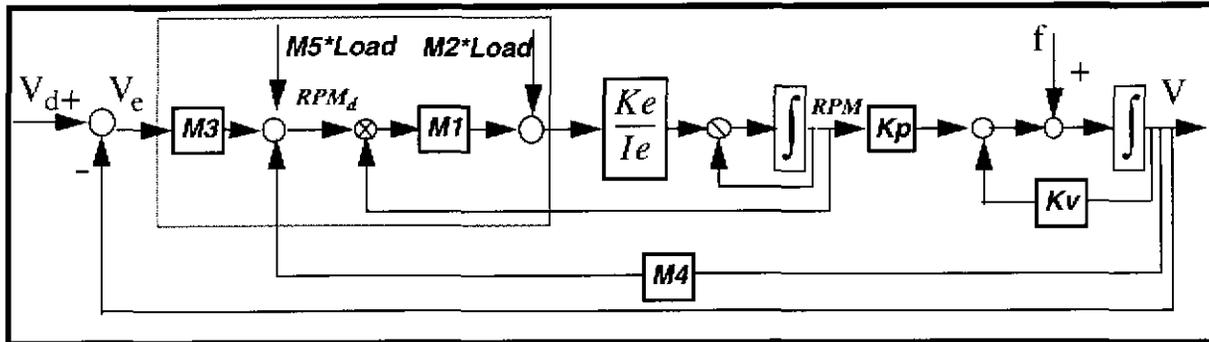
$$MassFlow_d = M1(RPM)(RPM_e) + M2(Load)$$

Finally, compute the throttle-plate angle necessary to achieve the required mass flow:

$$\phi = \arccos(1 - MassFlow_d)$$

31.1 LITHIA Throttle Compensation Algorithm Block Diagram

The following block diagram is an extension of linear control diagrams that illustrates the LITHIA algorithm. It was necessary to add multiply (circles with cross) and divide (circles with slash) to



the normal summing circles as the LITHIA algorithm is a non-linear compensator.

31.2 Algorithm Design Justification

The LITHIA algorithm, as designed above, achieves three things. First, it compensates for the coupling between the RPM state and the velocity state by explicitly calculating a desired RPM based on the velocity error. RPM alone is then controlled by altering the mass flow of fuel through the engine. Second, it compensates for the substantial gravitational and rolling friction disturbances through a feedforward term in both the RPM loop and the velocity loop.

The terms in the desired RPM equation arose in a somewhat *ad hoc* fashion based on the powertrain model:

$$m\dot{V} = -K_v V + K_{RPM} RPM + F_{dist} - s\theta mg$$

$$R\dot{P}M = \frac{K_e}{I_e(RPM)} MF - \frac{K_r \Gamma_e}{I_e}$$

The RPM desired control law is produced by solving for solving for RPM in the velocity differential equation. The velocity derivative term is replaced with the velocity error, and the constant parameters replaced with gains. The mass flow desired equation is similarly derived by solving for the mass flow in the second equation above.

31.3 Closed Loop Response Analysis

The closed loop response of the algorithm can be generated from the signal block diagram, or explicitly from the algorithm and powertrain model. Combining these functions yields the following closed loop response:

$$V = \frac{K_1}{s^2 + sK_v + K_2} V_d + \frac{s + m_1 m_5 K + \frac{m_2 K}{RPM}}{s^2 + sK_v + K_2} Load - \frac{Km_1}{s^2 + sK_v + K_2} RPM$$

where:

$$K = \frac{K_e K_R}{I} \quad K_1 = m_1 m_3 K \quad K_2 = -m_1 K (m_4 - m_3)$$

Note that the characteristic equation in either case is second order. The commanded velocity transfer function is no longer a function of the RPM state, indicating that it has been decoupled. The load transfer function retains one term in the numerator of the transfer function which is a function of RPM. The final term passes the RPM state to the output.

31.4 Comparative PID Closed Loop Analysis

Now that we have a model of the powertrain dynamics, a closed loop transfer function for the PID controller can be written as well. Combining the PID control law with the powertrain plant model yields the following equation, which cannot be solved in closed form due to the cosine term:

$$sV = \left[\frac{K}{RPM} \left(1 - \cos \left(K_p + sK_D + \frac{K_I}{s} (V_e) \right) \right) + (s + K) Load \right] \frac{1}{s + K_v}$$

Under the small angle assumption, we replace the cosine with the first two terms of its Taylor series expansion, yielding the following closed loop transfer function:

$$V = \frac{\frac{K}{RPM} \left[K_p + sK_D + \frac{K_I}{s} \right]}{s^2 + sK_v + \frac{K}{RPM} \left[K_p + sK_D + \frac{K_I}{s} \right]} V_d + \frac{s + K}{s^2 + sK_v + \frac{K}{RPM} \left[K_p + sK_D + \frac{K_I}{s} \right]} Load$$

A comparison of this transfer function with the LITHIA transfer function yields two key differences. First, the presence of the $1/RPM$ term in the PID's characteristic equation, indicating that the dynamic response of the system is still a function of the engine state. Second, the zero of the PID's Load function is fixed by the vehicle dynamics, there is no opportunity to move it through gain adjustment.

32. EXPERIMENTATION - FIRST GENERATION LITHIA

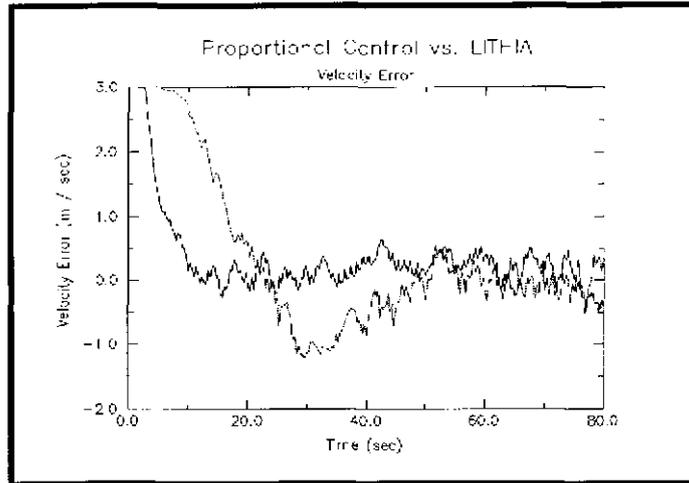
The LITHIA algorithm was tested on Flagstaff Hill using the proportional control algorithm as a basis of comparison, and also in rugged unstructured terrain. The purpose of these experiments was to quantify LITHIA's improvements in dynamic bandwidth, command following rise-time, and disturbance rejection. For the purposes of these experiments, we have isolated the throttle compensation algorithm, and a human operates the brakes - this necessity arises to avoid the so-called switching problem, which will be discussed later.

Two basic sets of experiments were performed. In the first, LITHIA operated on Flagstaff Hill, on the same course used for the earlier linear experiments, to enable easy comparison with proportional control techniques. For the second set of tests, LITHIA operated on a large industrial slag heap near the Homestead High Level Bridge on Pittsburgh's Monongahela River. The slag heap encompasses approximately one square kilometer, and has three basic terrain types. The *flats* are packed slag and earth, with low frequency undulations. The *plateau* rises above the flats by a height of about 100 feet. The slag of the plateau is much less densely packed - one can dig through this slag with a bare hand. Finally, there is a small section in which researchers built an obstacle course of meter tall mounds. The natural terrain in this area has higher frequency terrain. The whole of the slag heap is accessed by a small network of roads of varying pitch. LITHIA was tested over each of these areas, and across all of the roads. It is already known from prior experience that the proportional control algorithm cannot operate with algorithmic stability on terrain other than the flats. The plateau and the obstacle area have never been surmounted.

32.1 Flagstaff Hill Experiments

The LITHIA algorithm performed well against the linear controllers in the Flagstaff Hill test. LITHIA provided faster command attainment, with zero overshoot. The following graph in many

ways sums up the fundamental differences between the LITHIA algorithm and the linear approaches:



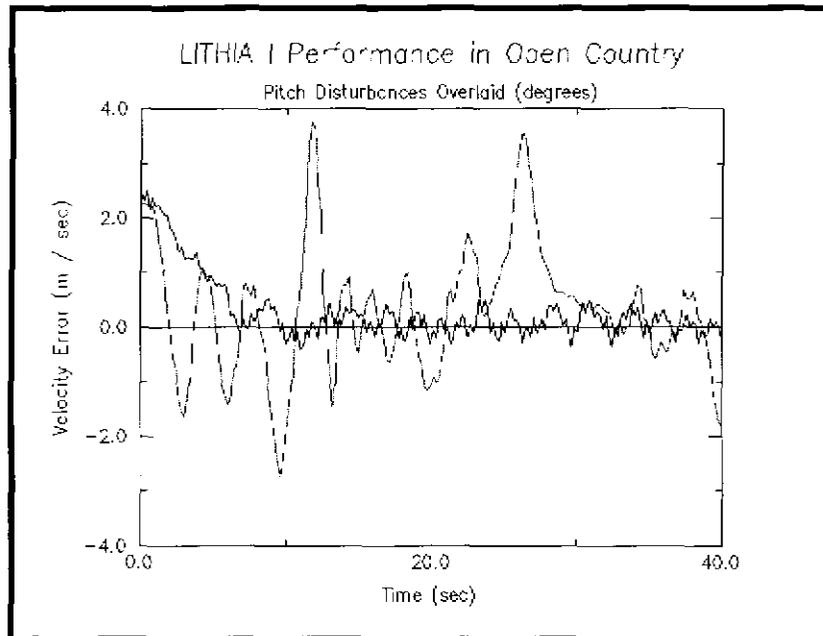
This graph illustrates the performance of the best of the proportional control runs (as judged by minimal overshoot, and minimal oscillation) against the LITHIA algorithm. Note that the proportional controller overshoots, then takes approximately 15 seconds to reattain the zero crossing. The LITHIA algorithm attains and maintains the zero crossing without overshooting. To compare the two algorithms, consider the following metrics:

Metric	Proportional	LITHIA
Time to 95% commanded velocity.	22 seconds	11 seconds
% commanded velocity, max overshoot.	39%	9%
Integral square error	398.11	140.6
Normalized ISE	1.327	0.469
Integral error	100.05	72.20
Normalized IE	0.333	0.24

It is important to note that the signal to noise ratio of the velocity measurement device is approximately 10%; the overshoot metric given for LITHIA is thus a direct measurement of noise. The error metrics given above include the standard IE and ISE measurements. Normalization of each is done by the number of samples taken across the run - these normalizations give some indication of the command following capabilities of the algorithms. The normalized ISE is essentially an average squared error, while the normalized IE gives some indication of the noise of the system - random variations tend to cancel in the NIE, but not in the NISE.

32.2 Slag Heap Experimentation

The first generation algorithm was tested at the slag heap site described earlier. A set of extended runs, over terrains of varying pitch, were performed to gauge the general performance abilities of the LITHIA speed control algorithm. Specifically, the purpose of the test was to determine whether the algorithm was capable of maintaining stability over undulating flat terrain, mounting significant slopes, and continuing to maintain stability. For these runs, brake actuation was again performed by the driver of the vehicle, so as to isolate the software test to the throttle algorithm. The driver was to hold the speed of the vehicle to a nominal 3 m/sec, braking when he would apply the brakes in normal operations. In the first test, the truck is driven over an undulating section of essentially flat terrain. The following graphic plots LITHIA's velocity error as a function of time, with the terrain pitch signal overlaid:

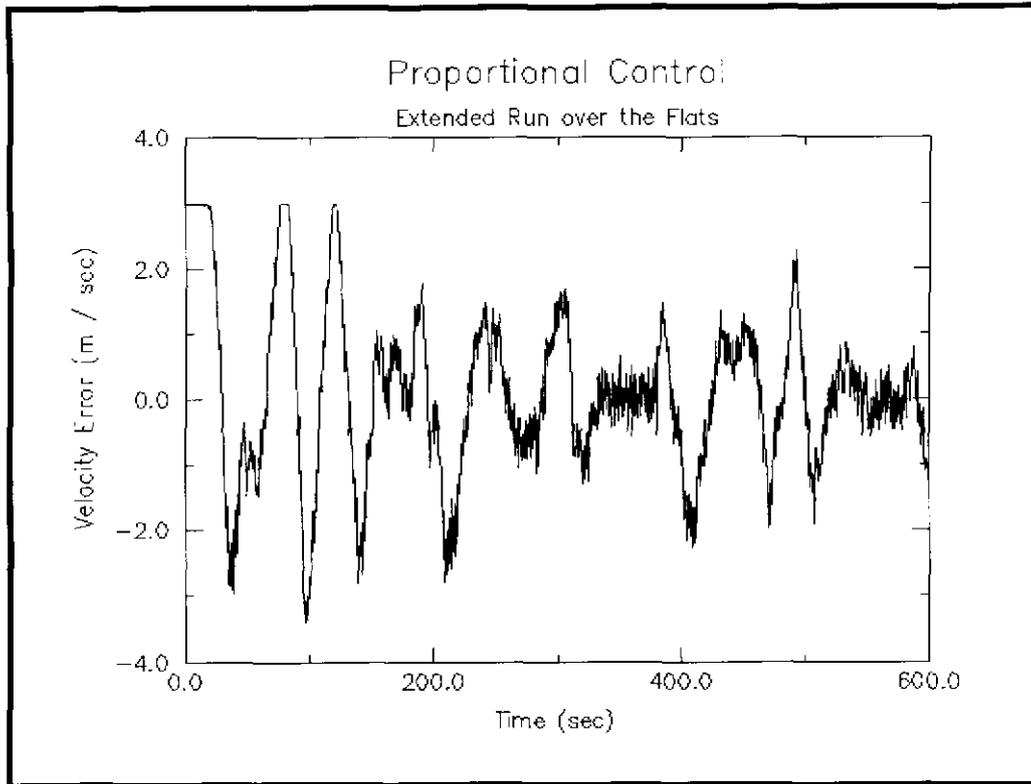


The performance of the controller can be quantified using standard control error metrics. Below, the values for the normalized IE, normalized ISE, and the RMS error are tabulated.

Normalized Sum of Error	Normalized Sum of Squared Error	RMS Error
0.294224	0.443181	0.665718

As a general comparison, the proportional algorithm was run over an extensive course, approximately one kilometer long, across the same area of the flats. The following graphic demonstrates that the basic behavior of the traditional approach is fundamentally different, displaying the same oscillatory characteristic seen in the Schenely Park tests. Note that the time scale of the graphic is compressed, compared to the former graphic; some 600 seconds, as opposed to the previous 40 second duration. Note that the proportional controller takes some 10 seconds

before it begins to move, and that twice during the run it encounters sufficient slope to bring it to a complete halt.

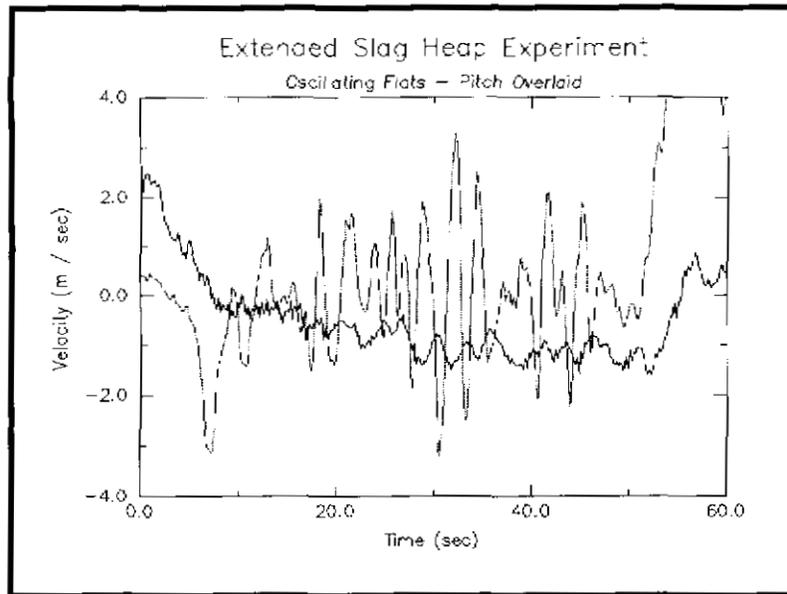


The application of standard controls tests yields the following metrics for the proportional controller:

Normalized Sum of Error	Normalized Sum of Squared Error	RMS Error
0.502128	4.437786	2.106605

The next series of graphics are produced from a single 15 minute LITHIA run in which the vehicle traversed over 2 kilometers. The graphics are divided according to terrain types encountered along the run. The vehicle travelled across the undulating flats of the main slag heap, mounted the steep access road to the top of the central plateau, turned around on top of the plateau, descended to the main roadway, drove into the back flats area, turned around and returned to the slag heap entrance via the main access road. The following graphic depicts the first portion of the journey, in which the vehicle traverses the oscillating flats. Note that the pitch signal varies rather quickly between

+3.0 degrees of inclination. These changes are sudden, and drastic enough at 3.0 m/sec to jostle the driver about the compartment.



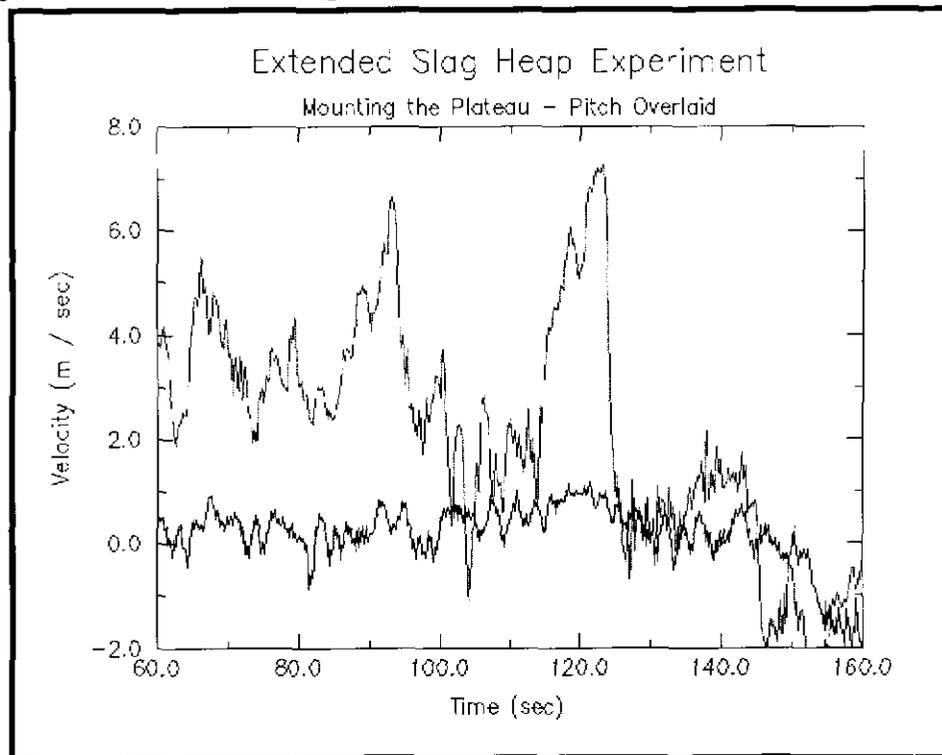
The performance of this run may be quantified using the standard controls metrics.

Normalized Sum of Error	Normalized Sum of Squared Error	RMS Error
-0.494652	1.017155	1.008541

It is interesting to note that although the throttle brings the vehicle quickly to the zero error axis, the overall positive inclination of the slope causes the vehicle to maintain a slight steady state error. In order to remove this error the driver would have had to apply the vehicle brakes. The driver based the braking decision on the speedometer read-out, and his own sense of appropriate speeds, and did not feel that the brakes should be applied. In general, we will see that the controlled speed of the vehicle is at a higher resolution than that of the human driver.

In the next graphic, the vehicle transitions from the oscillating flats area to the main access road. The inclination changes suddenly throughout, yet the LITHIA algorithm maintains tight control over the vehicle speed. It is interesting to note that there is no detectable lag in transition; this is in contrast with traditional PI control in which integrator wind-up takes some time to advance the throttle to re-attain the commanded speed. This particular maneuver is one that could not previously be safely performed due to the lack of PI stability under such drastic load changes. Near the end of the graphic's time range, at about 120 seconds, the pitch returns to level as the truck attains the top of the plateau - note that the vehicle's speed is maintained, and does not overshoot,

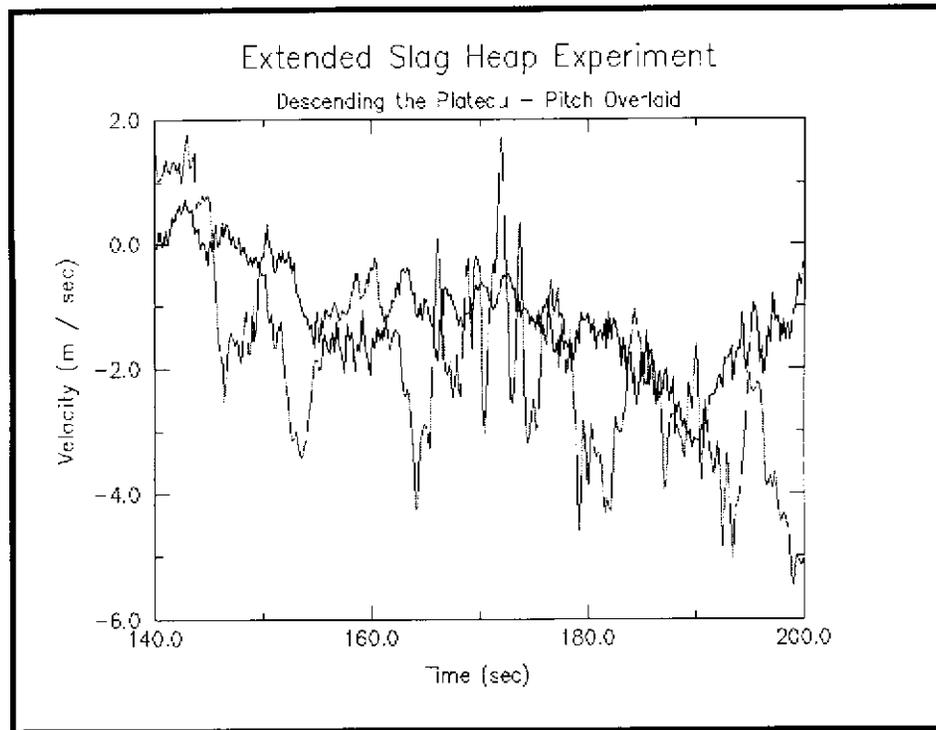
despite the sudden reduction in disturbance force. At 140 seconds, the vehicle has turned about and is beginning its descent on the same path.



The performance of the controls algorithm during this maneuver may be quantified as:

Normalized Sum of Error	Normalized Sum of Squared Error	RMS Error
0.332859	0.233575	0.483296

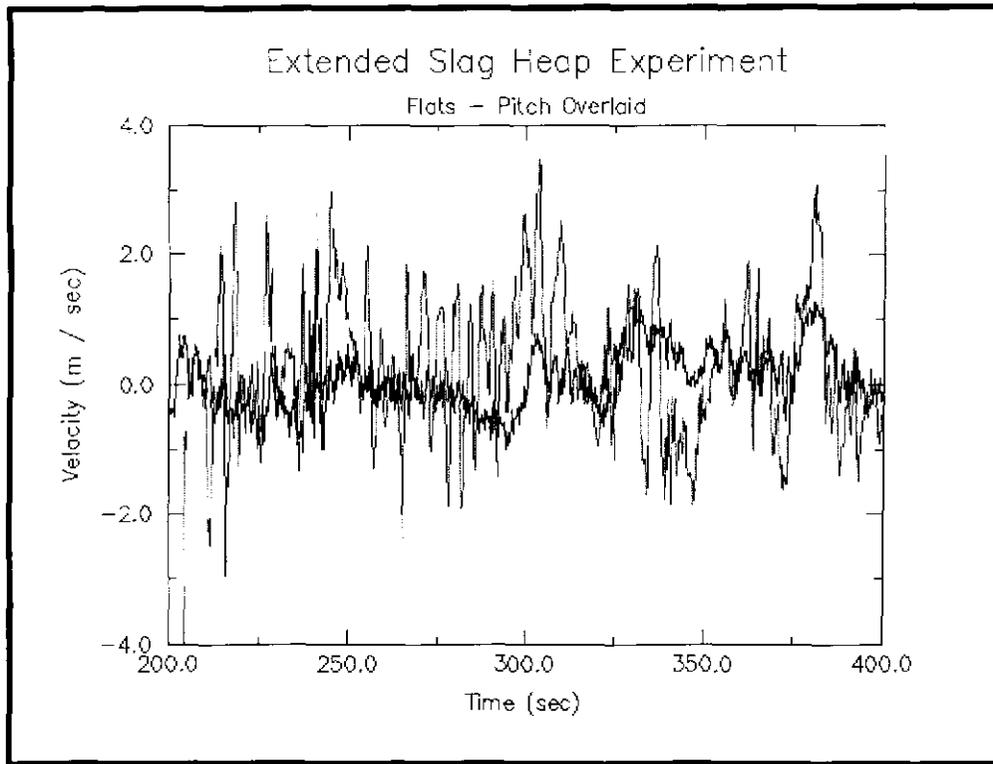
In the next sequence, the vehicle descends the plateau, with the driver applying the brakes. Note that the overshoot grows slowly during descent, until the driver appears to notice and take corrective action at approximately the 190 second mark.



The performance of the driver in the descent phase may be quantified using the following error metrics:

Normalized Sum of Error	Normalized Sum of Squared Error	RMS Error
-1.410273	2.458631	1.568002

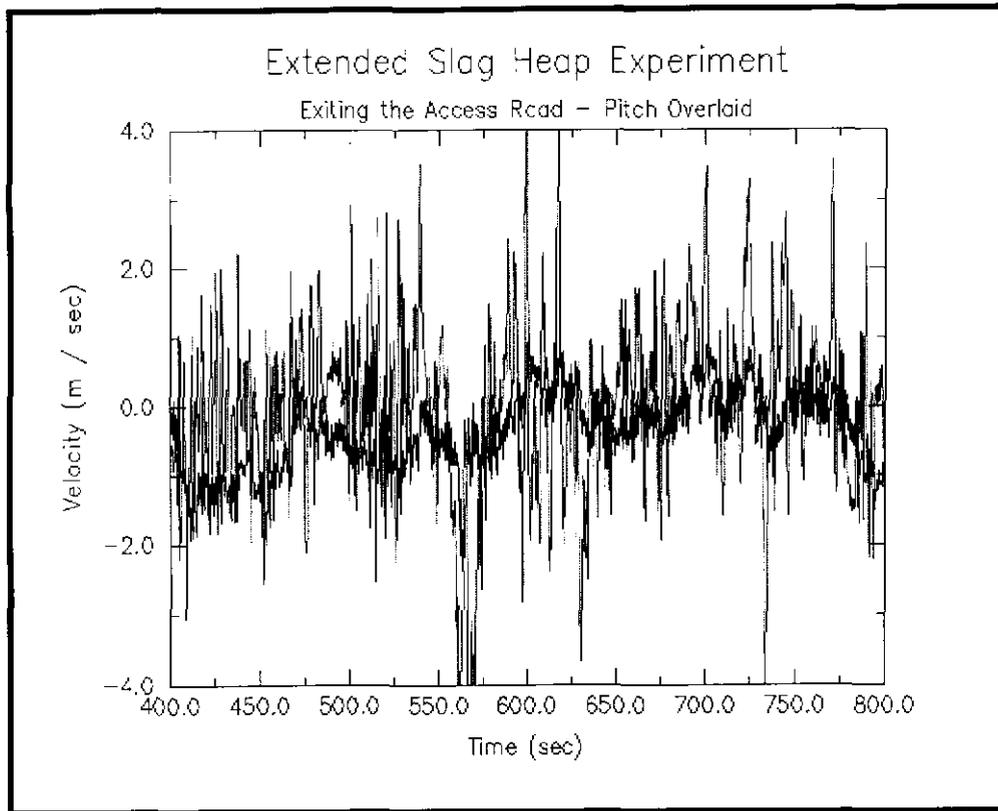
In the next portion of the run, the vehicle regains the main access road, and circles about through a relatively flat back lot, with more severe undulations. A particularly large undulation occurs at the 375 second mark, lasting for about 10 seconds, which sets the vehicle speed back. This speed obstacle indicated the need for a slightly higher gain on the feedforward terms. We should note that it is more difficult to perform the feedforward tuning for such short duration events, thus although the vehicle performs well in climbing longer segments of steep terrain, shorter impulse type disturbances are more difficult to reject.



The performance of this segment may be quantified as follows:

Normalized Sum of Error	Normalized Sum of Squared Error	RMS Error
0.073585	0.207186	.455177

The final segment takes the vehicle from the back lot to the entrance to the slag heap via the main access road. This section of road is filled with large, deep potholes, whose breadth approaches the wheelbase of the HMMWV.



Normalized Sum of Error	Normalized Sum of Squared Error	RMS Error
-0.351791	0.442026	0.664850

Overall the performance of the vehicle over the entire run is:

Normalized Sum of Error	Normalized Sum of Squared Error	RMS Error
-0.245294	0.529038	0.727350

33. DECONSTRUCTING THE ALGORITHM

Experimentation with the LITHIA algorithm on Flagstaff Hill led to a better understanding of how each term in the algorithm affects system performance. The LITHIA algorithm contains five gains, labelled M_i :

$$\begin{aligned} RPM_d &= M3 (Vel_d) + M4 (Vel) + M5 (Load) \\ MassFlow_d &= M1 (RPM) (RPM_e) + M2 (Load) \end{aligned}$$

Each of these equations takes a slightly augmented form of the standard servo control equation - a feedback term, plus a feedforward term. Gains M_2 and M_5 are the feedforward gains used to reject load disturbances. Gains M_3 and M_1 are feedback gains on velocity and RPM errors. Note the "extra" RPM in the M_1 term. This RPM derives from the physics of IC engines, and is an important property to be modelled. It is this property that prevents the mass flow equation from taking on the standard servo control form. The M_4 term similarly prevents the RPM equation from taking on the standard form. It is important to note that there is some overlap in terms, as was shown in the closed loop response equation. Final performance of the system is dependent upon the products of gain pairs, rather than on the values of strict individual gains:

$$V = \frac{K_1}{s^2 + sK_v + K_2} V_d + \frac{s + m_1 m_5 K + \frac{m_2 K}{RPM}}{s^2 + sK_v + K_2} Load - \frac{K m_1}{s^2 + sK_v + K_2} RPM$$

where:

$$K = \frac{K_e K_R}{I} \quad K_1 = m_1 m_3 K \quad K_2 = -m_1 K (m_4 - m_3)$$

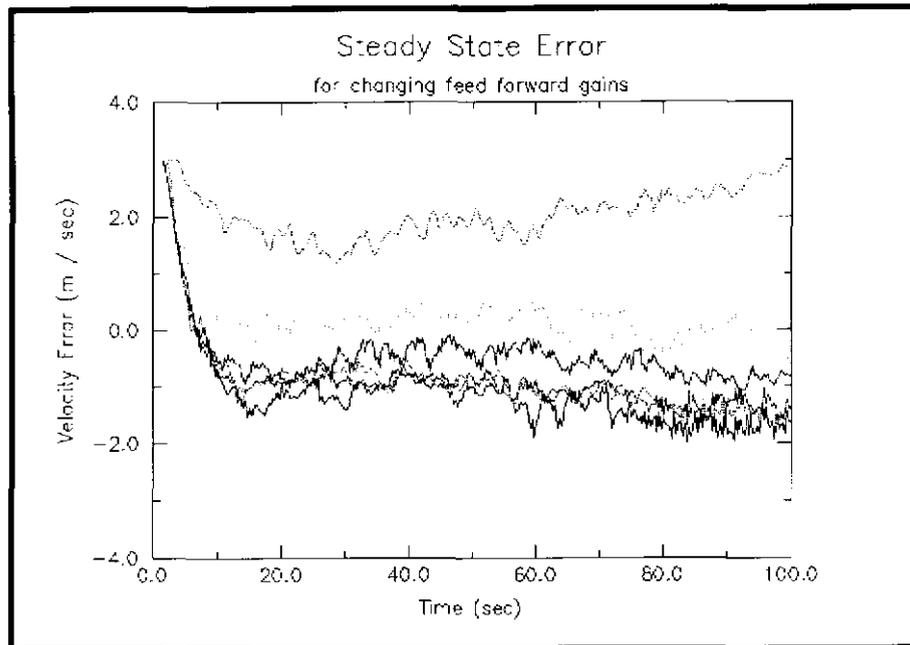
Fortunately, LITHIA has the advantage of an additional constraint in that its terms make physical sense to the user - thus, the gains $M_3 - M_5$ may be altered until the RPM rate requested by the controller makes some physical sense under the circumstances. Once reasonable RPM rates are being requested, one can move on to adjust the Mass Flow rate requests by adjusting gains M_1 and M_2 . The job is made even easier by recalling the feedforward / feedback performance split. Feedback tends to alter transient response and to cause oscillation in the steady state, while feedforward principally affects the steady state value. These effects are summarized in the following table:

	RPM	Velocity
Transient	M3	M1
Steady State	M5	M2

The final gain, M_4 , is found to be an unnecessary additional degree of freedom, originally intended to compensate for possible torque converter lag. The reason for its redundancy will be discussed in the LITHIA redesign section of this chapter.

33.1 Steady State Performance

As noted above, from servo control theory, we expect that the steady state performance of the algorithm will be altered by changes to the feedforward gains. An experiment performed on Flagstaff Hill helped to support this expectation. The following plot illustrates the change in the vehicle's hill-climbing performance as a function of the vehicle's feedforward gain.



33.2 Tuning the LITHIA Algorithm

One of the principal drawbacks of the gain scheduling or lookup table approaches is the amount of time (or number of experiments) required to construct the gain table. The LITHIA algorithm is very easy to tune, requiring about a half hour of effort, using the following method, which is based on our understanding of the gain's effects on the vehicle's performance. It is important to very slowly increase the gain values in order to prevent damage to the engine from opening the throttle too wide under low-load conditions. The tuning process occurs in three steps:

- Set initial values to the feedforward gains M_2 and M_5 to enable motion on a flat surface.
- Tune the feedback gains M_1 and M_3 to produce reasonable transient response.
- Tune the feedforward gains M_2 and M_5 to adjust the steady state response.

33.2.1 Initial Tuning of the Feedforward Gains

We begin on a reasonably level surface with all gains except M_1 set to zero. There are two goals to this portion of the tuning process. The first is to produce some basic motion to enable the tuning

process. The second is to set initial values to the feedforward gains, and an initial estimate of the rolling resistance portion of the load function; the level surface prevents the $\sin\Theta$ term from effecting the calculations. Note that M_1 requires a small non-negative value in order to pass the effects of the RPM calculation on to the mass flow calculation. The value of M_1 should be picked only after setting initial values to M_2 and M_5 , and noting the value of the requested RPM such that the order of magnitude of M_1 prevents the throttle plate from opening too wide. Let's take a look at how this process effects the transfer function:

$$V = \frac{K_1}{s^2 + sK_v + K_2} V_d + \frac{s + m_1 m_3 K + \frac{m_2 K}{RPM}}{s^2 + sK_v + K_2} Load - \frac{Km_1}{s^2 + sK_v + K_2} RPM$$

$$K = \frac{K_e K_R}{J} \quad K_1 = m_1 m_3 K \quad K_2 = -m_1 K (m_4 - m_3)$$

It is important to remember that most of the terms in this transfer function are found in *pairs*. Only the final term in the numerator of the *Load* transfer function contains a single gain. M_1 plays an especially prominent role as it is represented in every pair, this is the reason that a small non negative value is required of M_1 . Adjusting the values of the feedforward gains M_2 and M_5 affects only the numerators of each transfer function, which equates this part of the tuning process to *an adjustment of the steady state gain of each term*. Note also that if the *Load* function itself contains an estimate of the rolling resistance, then this constant must be estimated. Let's now look at the effects of this same process in the algorithm as expressed in the time domain:

$$RPM_d = M3 (Vel_d) + M4 (Vel) + M5 (Load)$$

$$MassFlow_d = M1 (RPM) (RPM_e) + M2 (Load)$$

Note again the role of the gain M_1 . This is the gain which passes the results of the RPM calculation on to affect the required mass flow, and thus the throttle plate position. Here again, M_1 must have a non-zero value in order for the effects of changing any gain $M_3 - M_5$ to be seen in the output.

In this first pass of tuning, we will be adjusting the values of the feedforward gains M_2 and M_5 . Adjusting M_5 will alter the value of RPM's that the algorithm requests. The RPM output of the algorithm is observed to determine whether M_5 is giving us a reasonable request; for example, 100,000 RPM would be far too large and 100 RPM would be too low. For the diesel engine on the NavLab II, values between 800 and 3000 are expected. Adjusting the value of M_2 has the effect of altering the requested mass flow directly. During the tuning process I find it useful to write out individually the values of the two terms in the mass flow calculations so that I can independently verify their reasonableness.

33.2.2 Tuning the Feedback Gains

Having achieved motion, we now adjust the feedback gains M_1 and M_3 to improve the transient response of the vehicle. Note that increasing the M_1 proportional gain will require that the

feedforward gain be simultaneously lowered to maintain reasonable steady state performance. We can see this effect in the time domain expression of the algorithm:

$$\begin{aligned} RPM_d &= M3 (Vel_d) + M4 (Vel) + M5 (Load) \\ MassFlow_d &= M1 (RPM) (RPM_e) + M2 (Load) \end{aligned}$$

Altering M_1 has the effect of altering all of the RPM expression's terms as seen in the output. Thus to maintain the same effect of the M_5 term in the output, M_5 must be lowered in proportion to M_1 's increase. This reasoning applies to the M_3 and M_4 terms as well. Similarly, raising the value of M_1 may require that the value of M_2 be lowered to effect better steady state output. This issue is addressed in the next section. The pairs coupling effect just discussed may be seen in the transfer function as well:

$$\begin{aligned} V &= \frac{K_1}{s^2 + sK_v + K_2} V_d + \frac{s + m_1 m_5 K + \frac{m_2 K}{RPM}}{s^2 + sK_v + K_2} Load - \frac{Km_1}{s^2 + sK_v + K_2} RPM \\ K &= \frac{K_e K_R}{I} \quad K_1 = m_1 m_3 K \quad K_2 = -m_1 K (m_4 - m_3) \end{aligned}$$

The effect of altering M_1 is to move the zeros that were placed by the M_2 and M_5 terms. If we think of the tuning process as a *pole placement process*, then we desire to place poles with M_1 and M_3 , while not altering the steady state response. In order to maintain the same steady state response, M_2 and M_5 must be correspondingly adjusted, as we will do in the next section.

33.2.3 Tuning the Feedforward Gains

The process of placing the poles in the last section may have altered the steady state gain of each transfer function, and in any case their *original placement was not intended to be terribly precise*. In this section we readjust M_2 and M_5 while operating on inclined surfaces in order to excite the $\sin\Theta$ term. This process is otherwise identical to their original adjustment. It is necessary to perform this process in three steps because in practice the feedforward terms (or equivalently, the disturbance forces) dominate the problem. If we were to eliminate the first step and simply adjust the feedback gains high enough to produce the initial motion, then the response would be highly erratic, as the poles would be driven into the left-half plane by the increase in magnitude of the K_2 term of the characteristic equation.

Driving the vehicle up an inclined slope repeatedly is a good way to fine tune this gain, as shown in a previous graph. It may be necessary to reduce the estimate of rolling resistance as the feedforward gain is increased - remember that the feedforward gain is amplifying at least two parameters, one of which is sinusoidal.

34. REDESIGNING LITHIA USING CLASSICAL CONTROL TECHNIQUES

The implementation of the first-generation LITHIA control algorithm pointed the way toward three improvements in the algorithm. First, the M_4 term was empirically noted to have little effect on the performance of the algorithm. In this section, a classical controls analysis characterizes its role in the system's closed loop response. Second, the structure of the first-generation LITHIA control technique achieved near decoupling of the RPM state from the closed loop transfer function. The second generation LITHIA algorithm eliminates this coupling, achieving full decoupling of velocity and RPM states. Finally, the pairs-coupling of gain terms seems indicative of a simpler formulation of the algorithm. The second-generation LITHIA algorithm eliminates pair-wise gains, reducing the number of gains in the process to two.

34.1 Eliminating M_4

It was empirically determined that the gain M_4 could be set to zero with no appreciable loss of performance. We can see why this is the case by first considering M_4 's role in the transfer function:

$$V = \frac{K_1}{s^2 + sK_v + K_2} V_d + \frac{s + m_1 m_3 K + \frac{m_2 K}{RPM}}{s^2 + sK_v + K_2} Load - \frac{K m_1}{s^2 + sK_v + K_2} RPM$$

$$K = \frac{K_e K_R}{I} \quad K_1 = m_1 m_3 K \quad K_2 = -m_1 K (m_4 - m_3)$$

M_4 appears only in the K_2 term, which affects the placement of the system poles. Note that the value of K_2 may be arbitrarily specified even if M_4 is held to zero. M_4 offers an additional, though seemingly unnecessary degree of freedom. M_4 's main effect is noted in the steady state response. The steady-state response of an s-space system is produced by noting that:

$$\lim_{t \rightarrow \infty} f(t) = \lim_{s \rightarrow 0} s f(s)$$

Taking a step input for the desired velocity, the steady state response of the velocity transfer function becomes:

$$V_{SS} = \frac{K_1}{K_2} = \frac{m_1 m_3 K}{-m_1 K (m_4 - m_3)} = \frac{m_3}{(m_3 - m_4)}$$

$$K = \frac{K_e K_R}{I} \quad K_1 = m_1 m_3 K \quad K_2 = -m_1 K (m_4 - m_3)$$

M_4 has the ability to shift the steady state response from unity, without offering any appreciable benefit. Should M_4 equal zero, then the desired unity steady state response emerges.

34.2 Achieving Complete Decoupling

It is preferable to attain a transfer function in which the RPM state does not appear. There are two terms which currently prevent this - the RPM term in the Load transfer function, and the existence of the entire *RPM* transfer function itself. The *RPM* transfer function can be traced to the use of the RPM_e term in the Mass Flow equation:

$$\begin{aligned} RPM_d &= M3 (Vel_d) + M4 (Vel) + M5 (Load) \\ MassFlow_d &= M1 (RPM) (RPM_e) + M2 (Load) \end{aligned}$$

Had this term been instead RPM_d , then the entire *RPM* transfer function would not arise. We have also noted the existence of a single RPM coupled term in the numerator of the *Load* transfer function. This RPM coupling could have been eliminated if, instead the M_2 term had been multiplied by the RPM.¹ If the algorithm is adjusted in this fashion, complete decoupling results.

34.3 Eliminating the Pair Gain Structure - Second Generation LITHIA

The first-generation LITHIA control algorithm represents an improvement in ease of tuning over gain scheduling or lookup table methods. LITHIA requires about a half an hour of experimentation to properly adjust its five gains. A good portion of this time is consumed because the gains occur in pair, which requires an additional adjustment round. It would be much more straightforward to tune the algorithm if the gains did not appear pair-wise. In this section, the second-generation LITHIA algorithm is derived, in a non pair-wise expression. The M_2 term has been multiplied by RPM to produce complete decoupling, and the M_4 term has been eliminated. We begin with a modified version of the first-generation LITHIA algorithm which reflects these changes to the M_2 and M_4 terms.

$$\begin{aligned} RPM_d &= M3 (Vel_e) + M5 (Load) \\ MassFlow_d &= [M1 (RPM_d) + M2 (Load)] (RPM) \end{aligned}$$

Substitution of the RPM_d term into the RPM_e term in the mass flow equation yields:

$$MassFlow_d = [M1 (M3 (Vel_e) + M5 (Load)) + M2 (Load)] (RPM)$$

1. This was actually a mathematical error on my part. It should have originally been $M_2 * RPM$, but I did not multiply the Load function by RPM when solving for MF.

Collection of terms yields:

$$MassFlow_d = [M1M3 (Vel_e) + (M1M5 + M2) (Load)] (RPM)$$

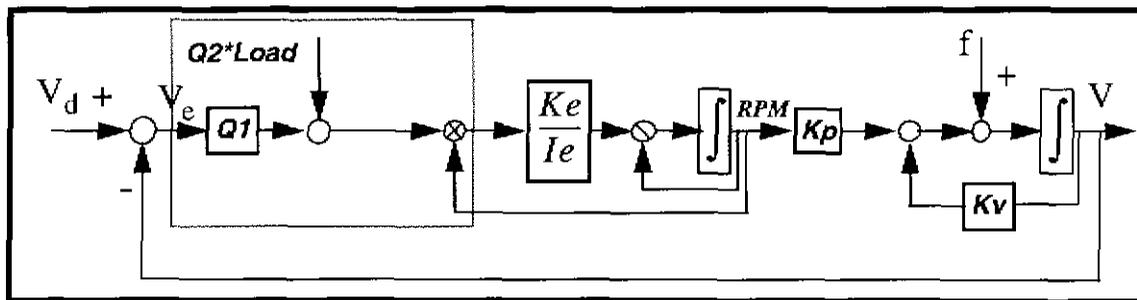
Now, the collected gain terms may be replaced with a new set of two gains, Q_i :

$$MassFlow_d = [Q1 (Vel_e) + Q2 (Load)] (RPM)$$

This algorithm represents the second generation form of the LITHIA control algorithm. It offers a more compact representation, with fewer gains to tune. The gains are clearly split into a feedback term (Q_1), and a feedforward, disturbance rejection term (Q_2).

34.4 Second Generation LITHIA - Closed Loop Response

The second generation LITHIA algorithm block diagram and corresponding closed loop transfer function can be produced from the time domain algorithm:



The corresponding closed loop transfer function is found to be:

$$V = \frac{KQ_1}{s^2 + sK_v + KQ_1} V_d + \frac{s + K(1 + Q_2)}{s^2 + sK_v + KQ_1} Load$$

$$K = \frac{K_e K_R}{I}$$

A steady state analysis, with step inputs for both the commanded velocity and the load, produces the following response:

$$V_{SS} = 1 + \frac{(1 + Q_2)}{Q_1}$$

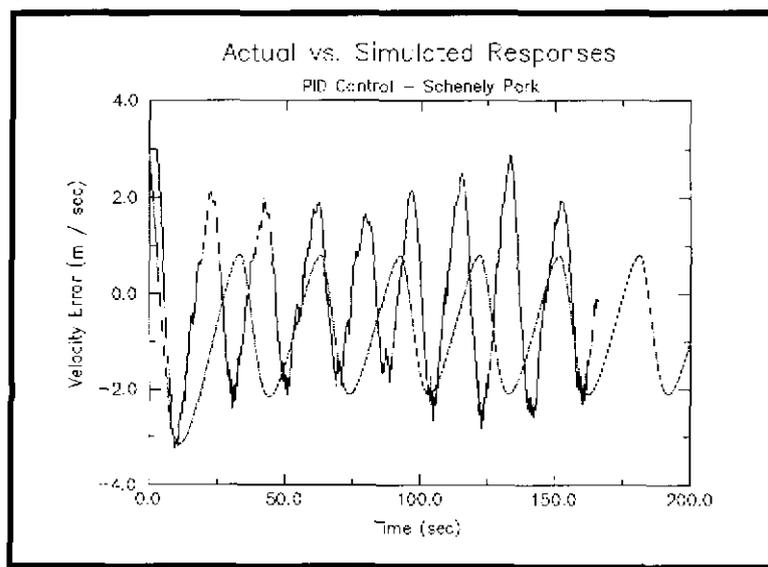
The steady state response to the commanded velocity is unity, while the steady state response to load inputs is a ratio of both the feedforward and feedback gains.

34.5 Tuning the Second Generation LITHIA algorithm

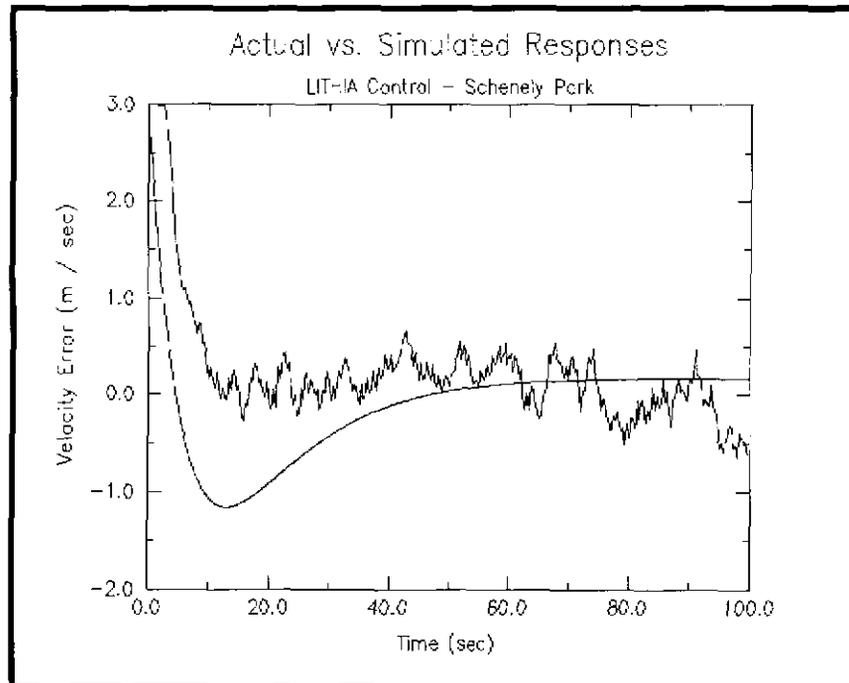
The second generation LITHIA algorithm is considerably easier to tune. There are only two gains, whose effects map easily onto observable physical effects. No special procedure is needed in this case, one may simply begin with small gains for each term. Adjust the feedforward gain to produce improved load rejection by moving the Load transfer function zero. Adjust the feedback gain to produce better command following by moving the location of the poles of the characteristic equation. Good performance may be achieved in less than fifteen minutes of effort. The second generation algorithm, because of its compact form and clean transfer functions, permits *a priori* calculation of optimal gains, given that the parameters of the transfer function have been identified.

34.6 Powertrain Simulator Experiments

The second generation algorithm has been tested using a powertrain simulator based on the models generated in this thesis, with an improved torque converter model developed in this chapter. The simulator is of sufficient resolution to enable qualitative comparisons of the gross dynamic system response to be drawn among the various algorithms. In general, I find the simulator to be a useful tool, but not an adequate replacement for in-vehicle runs. The following plots compare simulated vs. actual responses of the HMMWV driving the Schenely Park Course. The simulated response assumes a constant 5 degree course inclination, whereas the inclination of the actual run varies in either direction by an additional two degrees. The simplification is necessary due to discrepancies between the measured pitch rate and the pitch rate to which the vehicle actually responds. The following plot compares the simulated and actual responses for the PID controller.



The next plot compares the simulated and actual responses for the LITHIA controller:



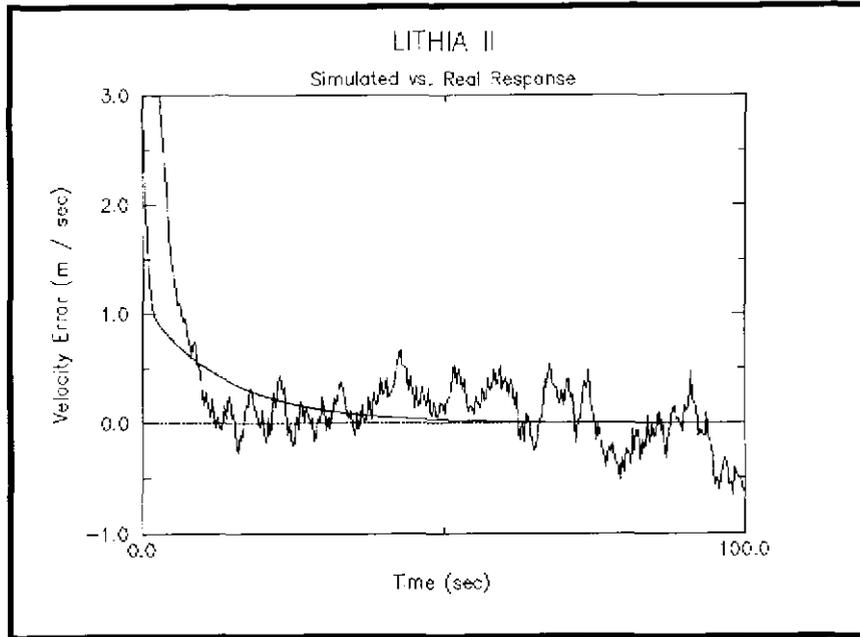
These simulations indicate that while the gross response of the real vehicle is captured by the model, additional unmodelled dynamics appear to be at work. There are two specific phenomena which appear to be important. In the PID control loop, it is apparent that the frequency of oscillation of the simulated vehicle is different from that of the real vehicle - this frequency is correlates well with the models of internal engine frictions and external vehicle drags. This point will become important later when considering possible future improvements to the LITHIA controller. The second phenomena is the magnitude of the overshoot in the LITHIA simulation. In practice, I have witnessed no such overshoots. Experimentation with tuning demonstrated that this phenomena appears to be tied to the assumed form of the powertrain equations, specifically the torque converter equation. As noted earlier, the planar rotational Couette flow model is primitive. (Cho and Hedrik[3]) proposed the use of quadratic functions to model the torque converter; such a model takes the form:

$$T = U_1 (RPM)^2 + U_2 (RPM) (Vel) + U_3 (Vel)^2$$

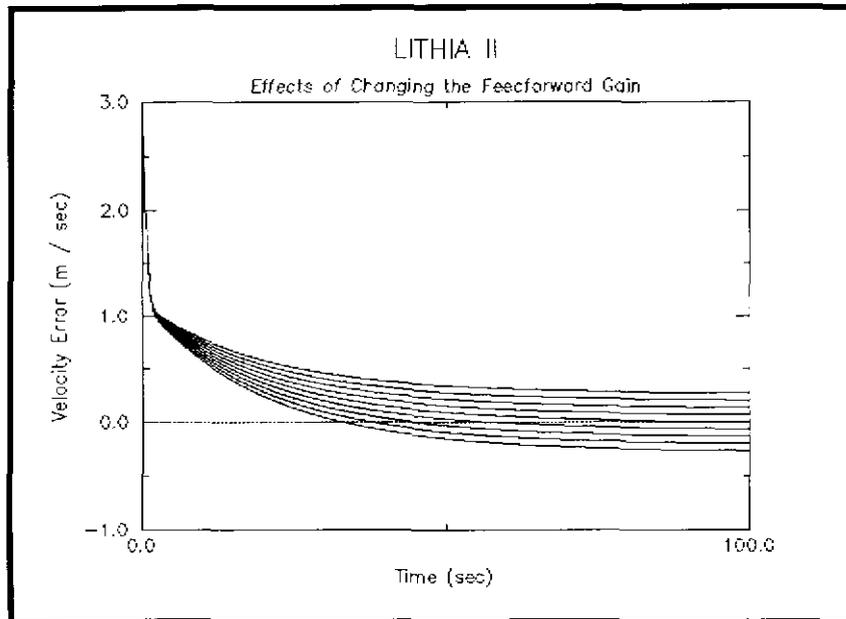
where U_i is constant.

The effect of moving to the quadratic model is two-fold. First, the troublesome overshoot is eliminated. Second, the rise time of the velocity can be made to correspond much more closely to

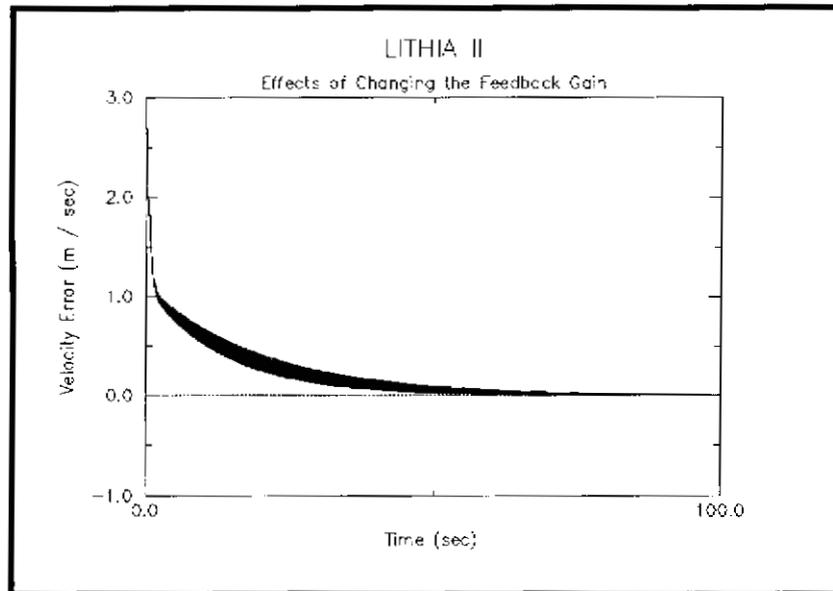
reality. The following graphic illustrates the improved performance of the simulator, as compared to the same real data taken from Flagstaff Hill:



It is interesting to note that the second generation LITHIA algorithm produces a clear division of effects between the feedforward and feedback gains during tuning. The following graphic illustrates the change in steady state that follows from the alteration of the feedforward gain:

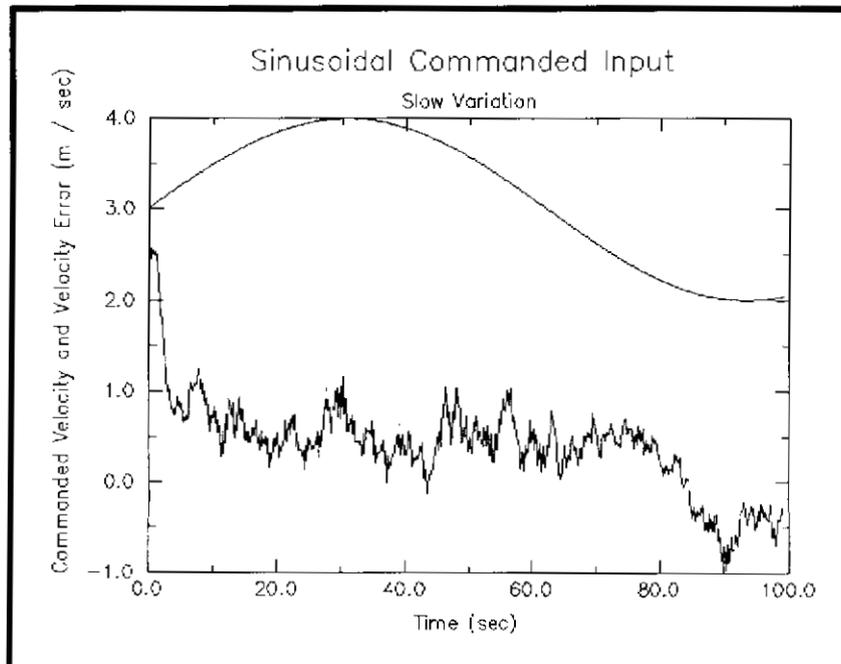


The next graphic demonstrates the change in transient response that follows from altering the feedback gain:

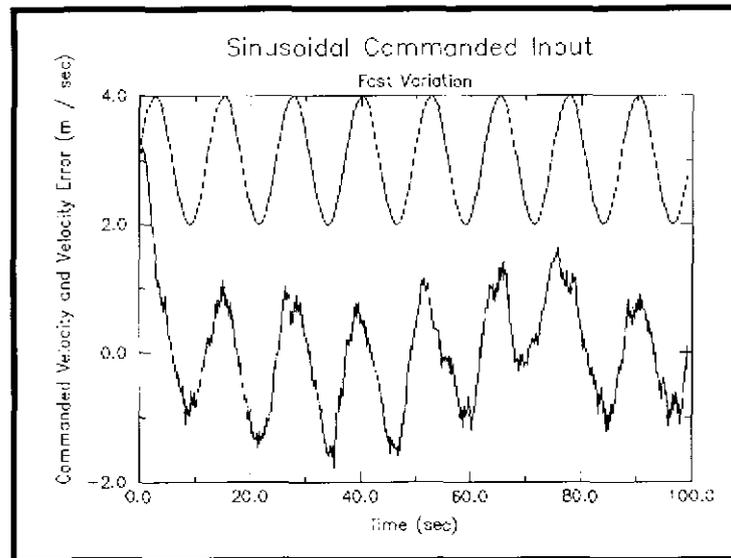


34.7 LITHIA II Experimentation

The LITHIA II algorithm demonstrates functional performance that is nearly identical to LITHIA I. The principal distinction between the two algorithms is mathematical - LITHIA II being a more compact representation. Consider the following graphs which demonstrate the ability of the algorithm to follow a sinusoidal input while simultaneously rejecting disturbances on the flats area of the slag heap:



In the first example, LITHIA is able to maintain velocity; however, it holds a general overshoot due to inadequate braking on the driver's part. In the second example, LITHIA is not able to maintain velocity, as the frequency of the commanded signal exceeds the system's physical response bandwidth.



35. DISTINGUISHING LITHIA FROM SO-CALLED “LINEAR TECHNIQUES”

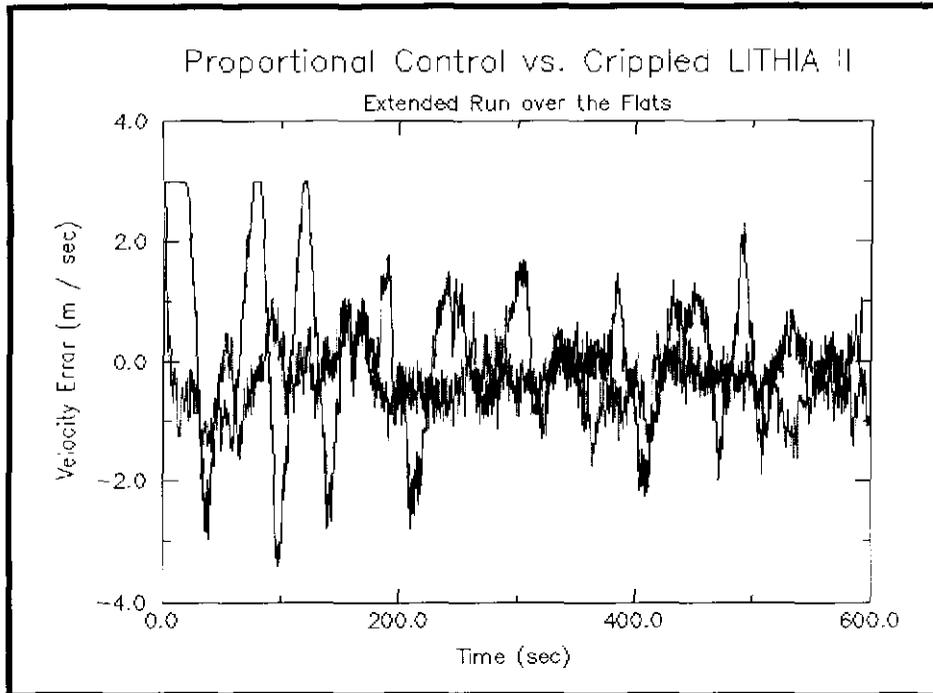
The inability of early proportional control laws to properly control the powertrain system of the NavLab II is not grounded in any failing on the part of the control law itself, but rather on the manner in which it was applied. It should be clear by now that the LITHIA controller is essentially a proportional controller; however, LITHIA is applied in the power or energy space, rather than the traditional velocity / force space. There are two practical requirements forced by the adoption of this perspective. First, that we properly distinguish the control variable from the actuated variable; and second, that we command this control variable in a manner mathematically compatible with its physical form.

35.1 Control vs. Actuated Variables

It is important that one distinguish between the axis over which one may exert physical control (Actuated Variable), and the input parameter which controls the system (Control Variable). In this application, LITHIA makes a distinction that previous control systems did not - LITHIA servos vehicle speed by actuating the throttle to control the mass flow of the engine. Previous control systems attempt to servo speed by controlling and actuating the same axis - the throttle plate. We can demonstrate the importance of this distinction through an experiment, in which we take the LITHIA II algorithm and restrict its input such that it resembles Murphy's gain scheduling method.

$$\text{MassFlow} = Q_1 V_e + Q_2 \text{Load}$$

The RPM feedback is held to be a constant signal, and zero pitch angle is returned to the controller, the *Load* function estimates only the rolling friction. Now, we run the crippled LITHIA II algorithm over the flats area, along with the proportional controller and note the form of their responses.



This comparison demonstrates that *by distinguishing between the actuation of the throttle plate and the control of the mass flow, a fundamental change in behavioral form is achieved*. This experiment implies that, had Murphy's control method included the throttle port geometry transform, reasonable performance would have been achieved without the need for gain scheduling. It further implies, but cannot prove, that the values in Murphy's table somehow encode this port transform.

35.2 Compatibility between Forcing Function and Physical System

A second distinction between LITHIA and previous algorithms is that it forces a compatibility between the form of the forcing function and the linear differential equation description of the system. This compatibility enforcement results in the application of a servo controls algorithm in the power / energy space, not the velocity / force space of the vehicle. The differential equation of the powertrain dynamics must be written in the *power space* in order to produce a *linear* ordinary differential equation appropriate to the application of a *linear* control law.

Control theory predicates the existence of a physical system, which can often be described by a linear ordinary differential equation of the following form:

$$a_0 \frac{d^n y}{dt^n} + a_1 \frac{d^{(n-1)} y}{dt^{(n-1)}} + a_2 \frac{d^{(n-2)} y}{dt^{(n-2)}} + \dots + a_n y = f(z)$$

where $f(z)$ is called the forcing function, and z is called the control law. We seek to produce a control law z that alters the dynamics of the system in a preferred manner. The value of z maps onto the forcing function, producing the forced response of the system; thus, the effects of control are represented within the particular solution to the differential equation.

35.3 The Particular Solution

The particular solution to the linear differential equation is obtained from the homogeneous solution using a technique called the *variation of parameters*. The homogeneous solution of the differential equation may be written in the form:

$$y_H = \sum_{k=1}^n c_k u_k(t)$$

The variation of parameters technique assumes that the constant parameters c_k in the above expression are themselves functions, and that the particular solution may be written as:

$$y_P = \sum_{k=1}^n C_k(t) u_k(t)$$

where the derivatives of the functions C_k are produced through the application of Cramer's rule, with the common denominator determinant being the Wronskian of the functions u_k . Thus, for example, the derivative of C_1 may be produced as follows:

$$C_1' = \frac{\begin{vmatrix} 0 & u_2 & u_3 & \dots & u_n \\ 0 & u_2' & u_3' & \dots & u_n' \\ 0 & u_2'' & u_3'' & \dots & u_n'' \\ \dots & \dots & \dots & \dots & \dots \\ f & u_2^{(n)} & u_3^{(n)} & \dots & u_n^{(n)} \end{vmatrix}}{\begin{vmatrix} u_1 & u_2 & u_3 & \dots & u_n \\ u_2 & u_2' & u_3' & \dots & u_n' \\ u_3 & u_2'' & u_3'' & \dots & u_n'' \\ \dots & \dots & \dots & \dots & \dots \\ u_1^{(n)} & u_2^{(n)} & u_3^{(n)} & \dots & u_n^{(n)} \end{vmatrix}} = \frac{\begin{vmatrix} 0 & u_2 & u_3 & \dots & u_n \\ 0 & u_2' & u_3' & \dots & u_n' \\ 0 & u_2'' & u_3'' & \dots & u_n'' \\ \dots & \dots & \dots & \dots & \dots \\ f & u_2^{(n)} & u_3^{(n)} & \dots & u_n^{(n)} \end{vmatrix}}{W[u_1, u_2, \dots, u_n]}$$

from which the function C_1 is found through integration:

$$C_1 = \int C_1' dt$$

From this solution, the required unit-wise compatibility between the forcing function $f(x)$ and the particular solution $y_H(x)$ may be found. First, we denote the units of the variable x by $[x]$. Thus, if x is a position, then $[x]$ may be meters. From the determinant expression given above, we note that the units of C_1' will be reported as:

$$[C_1'] = \frac{[f(t)] [u_i(t)]^{(n-1)}}{[u_i(t)]^n} = \frac{[f(t)]}{[u_i(t)]}$$

Upon integration, this quantity will be multiplied by time:

$$[C_i] = \frac{[f(t)] [t]}{[u_i(t)]}$$

The form of the particular solution gives a unit-based relationship between the output y_H and the solution pairs $C_i u_i$:

$$[y_H] = [C_i] [u_i] = \frac{[f(t)] [t]}{[u_i(t)]} [u_i] = [f(t)] [t]$$

This relationship expresses the required compatibility between the system output and the system forcing function, namely that *the forcing function be an expression that is one time derivative removed from the output*. In the particular case of conventional vehicles, it is clear that the system output may not be chosen to be velocity, given that *the forcing function maps onto power - the system output must be an energy-based variable*.

35.4 The Force Space Solution

The incompatibility of the traditional system description is apparent when attempting to produce the particular solution to the following system model of a conventional vehicle. Assume the following first order system description:

$$\frac{dV}{dt} + aV = f(z)$$

The particular solution to this equation may be produced in closed form using the integrating factor method, yielding:

$$V = e^{-at} \int e^{at} f(z) dt$$

The function $f(z)$ must map the control law z onto force (or acceleration) to produce a compatible solution. In the case of conventional vehicles, the forcing function $f(z)$ maps the control law z onto power. Recall from the modelling chapters that the engine is an exothermic chemical reactor; changes in throttle position map onto changes in reactant rates, and thus changes in power production. Thus, *this system model is incompatible with its forcing function.*

35.5 The Energy Space Solution

Choose now the kinetic energy, E_k , of the system as the independent variable y :

$$\frac{dE_k}{dt} + bE_k = f(z)$$

Now, rewrite the kinetic energy in terms of the velocity:

$$\frac{d}{dt} \left(\frac{1}{2} m V^2 \right) + b \left(\frac{1}{2} m V^2 \right) = f(z)$$

This equation can be reduced in form such that its left-hand side resembles the original equation:

$$mV \frac{dV}{dt} + b \left(\frac{1}{2} m V^2 \right) = f(z)$$

$$\frac{dV}{dt} + \frac{b}{2} V = \frac{f(z)}{mV}$$

Since $f(z)$ maps onto power, the product of force and velocity, the V in the new forcing function's denominator will cancel. We can rename this new forcing function $g(z)$, and write a new differential equation:

$$\frac{dV}{dt} + \frac{b}{2} V = g(z)$$

Yielding the following solution:

$$V = \frac{2}{b} \int e^{-\frac{b}{2}t} g(z) dt$$

In this case, the function $g(z)$ produces a force or acceleration, which, when integrated yields a velocity. This system equation is compatible with its forcing function. **It is this point which drives the design of the LITHIA control laws.** In either case, the so-called decoupling step “removes” the RPM term by enforcing the compatibility of the forcing function with the chosen system model. This can be seen most readily in the second-generation LITHIA control law:

$$MassFlow_d = [Q1 (Vel_e) + Q2 (Load)] (RPM)$$

The terms inside the brackets produce an estimate of the force or torque required to servo to the desired velocity. This estimate is multiplied by the current RPM to move it into the power space of the engine. This step closes the loop from velocity feedback, to torque estimate, to mass flow command, to torque production, to velocity production in a compatible manner.

35.6 Underlying Physics

Normally the constituent equation of the forcing function is a force law; for example, electromotive force or hydraulic force. These laws may be directly substituted into Newton's Second Law as externally applied forces. In this case, the constituent equation of the engine is the First Law of Thermodynamics, an energy equation - or as expressed here in mass flow terms - a power equation. Newton's Second Law and the First Law of Thermodynamics may be written compatibly only in the energy / power / work domain, having integrated Newton's Second Law appropriately. It is this domain which offers the linear control space to our problem.

36. BRAKE SERVO CONTROL

We noted in an earlier chapter that traditional proportional control, as applied to vehicle braking, worked well when driving about on reasonably level terrain. It is a straightforward matter to extend the original proportional brake controller to 3-D terrain by considering the additional loads applied to the vehicle during traversal:

$$\dot{V} = -\frac{F_b}{m} + g \sin \Theta = -\frac{K_b}{m} x + g \sin \Theta$$

and compensating through the addition of a feedforward term:

$$x = -K_p V_e - K_f \sin \Theta$$

which yields the following closed loop control law:

$$\dot{V} = -\frac{K_b}{m} (K_p V_e - K_f \sin \Theta) + g \sin \Theta$$

$$V = \frac{K_b K_p}{sm + K_b K_p} V_d + \frac{K_b K_f + g}{sm + K_b K_p} \sin \Theta$$

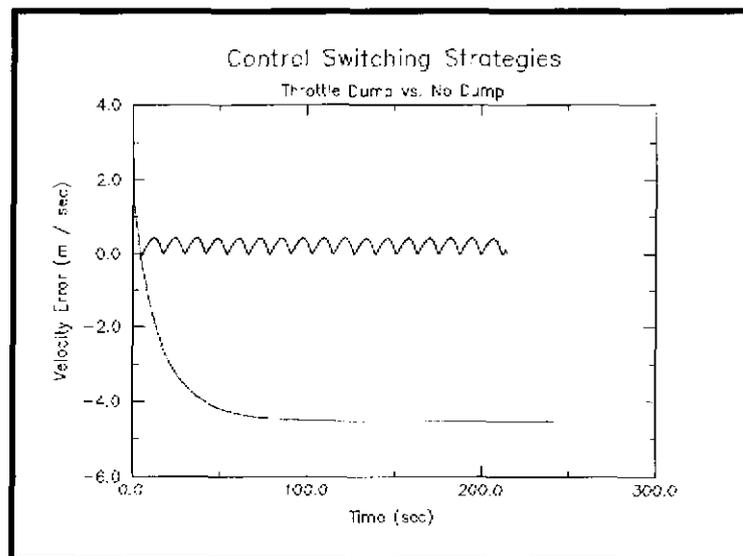
It is clear from this model that we may select $K_f = -g/K_b$ to produce complete disturbance rejection.

37. THE SWITCHING PROBLEM

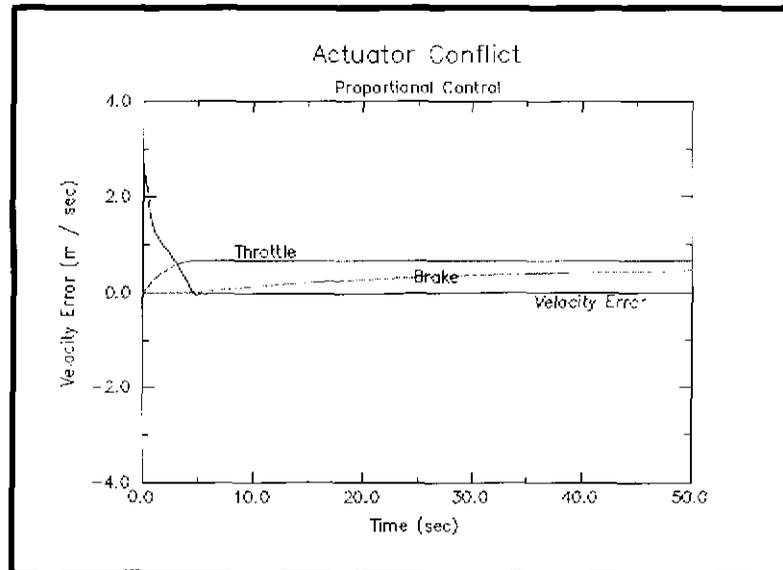
Velocity control of conventional vehicles requires the coordinated control of two *unilateral* control axes - the *throttle* and the *brake*. Switching between these two control axes requires switching between two control laws, and in so doing, setting the end and initial conditions for each control law during each transition. Using the powertrain simulator, a set of switching experiments were performed. In these of experiments the switching law considers the velocity error:

if ($V > 0$) then (ACCELERATE)
 if ($V \leq 0$) then (BRAKE)

Both the proportional algorithm and the second generation LITHIA algorithm are tested in a constant uphill slope scenario in under two sets of end conditions: *Throttle Dump* - in which the throttle position is set to zero before switching to the brake algorithm; and *No Dump* - in which the last remaining throttle position is kept before switching to the brake algorithm. The following graphic illustrates the test results for the proportional controller:

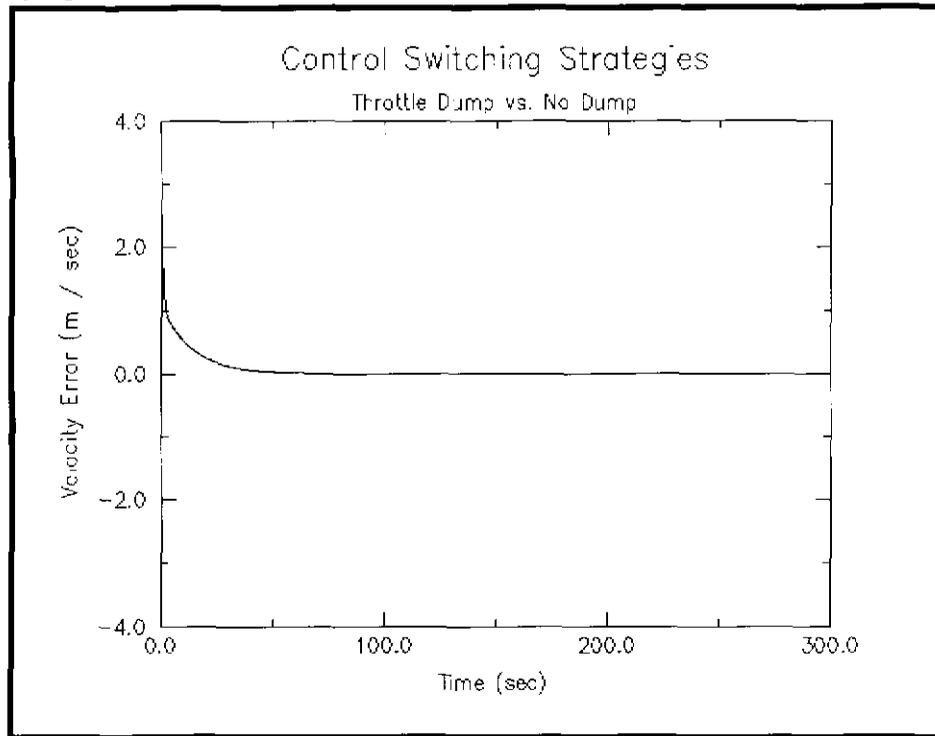


The throttle dump end conditions produce a limit-cycle behavior, in which the brakes are never actuated. At the zero crossing, the proportional algorithm is exited for one cycle, in which the throttle position is reset to zero; the vehicle slows while the throttle opens enough to accelerate the vehicle back to the zero crossing, where the cycle is repeated. The no-dump end condition is more dramatic, in that the throttle retains the partially open position that it had at the zero crossing. The brakes are applied, but in this simulation are unable to slow the vehicle to the zero crossing; the partially open throttle and the fully applied brakes reach an equilibrium at some overshoot speed. In the following graphic, under the simulation of stronger brakes, a zero velocity error is held, but under the efforts of combined throttle and brake. The throttle and brake actuator positions are superimposed on the graph:



The switching problem is troublesome, in that it requires the introduction of discontinuities in motion control algorithms that are typically designed to be continuous. The behavior observed in the proportional controller is related to the algorithm's convergence behavior - in a previous chapter it was pointed out that the proportional algorithm is guaranteed to overshoot because of its zero crossing behavior - this behavior leads to the switching problem because switching is guaranteed to occur at mid-points of the actuator's range. Switching problems can be alleviated if the algorithm's convergent behavior forces axis control switches to occur when the actuator returns to its zero position, as in the LITHIA algorithm. In the graphics below, no switching problem is

observed, and in this case the brake is never actuated, because the throttle servos the velocity of the vehicle asymptotically to the zero error.



38. OPTIMIZATION

LITHIA's achievement of stable control now allows us to consider optimizing the use of control against other system parameters such as stability and fuel consumption. Optimal control admits generally those problems which seek to minimize a cost function or performance index, of the form:

$$J = F(x(t_f), t_f) + \int_0^{t_f} L(x(t), u(t), t) dt$$

where $F(x(t_f), t_f)$ describes the terminal performance of the action, and the integral describes the transitional performance of the action. Such a performance index may be used either in the *design process* or during the control process. In the design process one use J to *select gains* which provide *optimal performance for the execution of a specific action*. The dynamic bandwidth over which the performance remains near optimal is a function of the system sensitivity. A sensitivity analysis is often performed with the design analyses to guarantee a robust optimality margin. One may also use the performance index to *design an optimal control law*. Two standard approaches to this solution include the dynamic programming principle, and the use of Pontryagin's minimum principle.

38.1 Linear Quadratic Optimal Control

The LQ performance index considers the state error and the magnitude of the control signal in the production of an optimal control law - linear quadratic control.¹ The performance index is given in continuous form by:

$$J = \int_{t_i}^{t_f} [x^T(t) Q_{1c} x(t) + 2x^T(t) Q_{12c} u(t) + u^T(t) Q_{2c} u(t)] dt$$

The performance index is given in discrete form by:

$$J(k) = x^T(kh) Q_{1x}(kh) + 2x^T(kh) Q_{12} u(kh) + u^T(kh) Q_{2u}(kh)$$

where:

$$Q_1 = \int_{kh}^{(kh+h)} \Phi^T(s, kh) Q_{1c} \Phi(s, kh) ds$$

$$Q_{12} = \int_{kh}^{(kh+h)} \Phi^T(s, kh) [Q_{1c} \Gamma(s, kh) + Q_{12c}] ds$$

$$Q_2 = \int_{kh}^{(kh+h)} [\Gamma^T(s, kh) Q_{1c} \Gamma(s, kh) + 2\Gamma^T(s, kh) Q_{12c} + Q_{2c}] ds$$

The LQ control rule is given to be:

$$u(k) = -L(k)x(k)$$

where $L(k)$ is defined to be:

$$L(k) = \left(Q_2 + \Gamma^T S(k+1) \Gamma \right)^{-1} \Gamma^T S(k+1) \Phi$$

1. The following derivation is taken from Astrom and Wittenmark, chapter 11 - Optimal Design Methods: State-Space Approach

and $S(k)$ is defined to be:

$$S(k) = [\Phi - \Gamma L(k)]^T S(k+1) [\Phi - \Gamma L(k)] + Q_1 + L^T(k) Q_2 L(k)$$

and is called the discrete time Riccati Equation.

38.2 Application of LQ to the NavLab II

It is possible, but not quite straightforward, to apply the LQ method, as formulated, to conventional vehicles. The first problem arises with the linearity of the powertrain model. As noted in the LITHIA section, the constituent equations are not linear in velocity, but in energy. This forces us to adopt an energy based set of state equations. Beginning with the system equations, rewrite each such that it is a power equation:

$$P = [MF + (Load) (RPM)] \frac{K_e}{I_e}$$

$$P = mV \frac{dV}{dt} = mK_{RPM} RPMV + mK_V V^2 + mK_L LoadV$$

These equations can be combined to form a single power equation:

$$P = mK_{RPM} (RPM) (V) + mK_V V^2 + mK_L (Load) (V) + m \frac{K_e}{I_e} (Load) (RPM) + \frac{K_e}{I_e} MF$$

Which will be somewhat artificially grouped as follows:

$$\frac{dE}{dt} = E + \frac{K_e}{I_e} MF$$

$$E = mK_{RPM} (RPM) (V) + mK_V V^2 + mK_L (Load) (V) + m \frac{K_e}{I_e} (Load) (RPM)$$

This artificial grouping does not live up to the mathematics of the first order equation in that the load terms appearing in the Energy equation above give power, not energy. Strictly speaking they should appear as system inputs. The purpose here is to isolate the mass flow term such that the LQ control law may be applied. Given this formulation of the first order equation, the following performance index is derived:

$$J = \int_{t_i}^{t_f} (E^2 Q_1 + MF^2 Q_2) dt$$

This performance index provides optimizing information that is directly related to the velocity error of the vehicle, the applied loads, and the state of the engine. Should one wish to include other dynamic responses, such as the suspension response of the vehicle, then power / energy

formulations may be added to the energy terms above. Solution via the Riccati equations produces an optimal control law.

38.3 Limitations of the LQ Method

The LQ method was ultimately rejected, with some reservations, in favor of the feedforward optimal control or action space planning method. In making this design decision, I reasoned over several potential problems which I discuss here.

38.3.1 Role of Prediction

The LQ method is implemented by precomputing the feedback gain $L(k)$ from forward time N back to present time. In practice, the stationary controller achieved by extending N out in time is often used. The performance of this precomputed step has the potential advantage of allowing us to achieve optimal control without the necessity of performing a costly feedforward computation each cycle. The potential problem with this approach is that the fidelity of the model may not be sufficient to support reliable pre-computing. The models developed in this thesis were derived for the purpose of capturing gross dynamics and general physical relationships. I suspect that higher order models may be necessary to achieve robust control.

38.3.2 Control Signal Saturation

The LQ controller has no knowledge of the limitations of the physical system. The linear energy formulation does not admit of engine RPM limits, acceleration limits, throttle limits, etc. The inability to easily insert a saturation mechanism in the control formulation was seen to be a significant problem with the method.

38.3.3 Switching Unidirectional Control

The derived LQ controller does not recognize the throttle control signal as a unilateral controller. In order to practically implement such a system, some method of separating brake control from throttle control would need to be included; or some method of mapping energy requirements back onto throttle and brake commands.

38.3.4 Velocity / Energy Equivalency

It is unclear that performing LQ control on kinetic energy is equivalent to optimization on velocity. Specifically, the energy term in the energy differential equation includes energies of the vehicle as a body, as well as energies of the engine as a rotational system. Such a formulation appears to demand optimization across a larger system than intended.

39. FEEDFORWARD OPTIMAL CONTROL

Feedforward optimal control differs from classical LQ control in the manner in which it obtains a solution to the performance index minimization problem. Whereas LQ control offers a closed form solution generated through dynamic programming, the feedforward method obtains an approximate optimal solution through simulation. This method of solution generation is called the

feedforward solution, although it has recently been called action-space planning (Kelly and Singh[18]) in the mobile robotics community when contrasted with classical configuration space planning methods. The selection of an *open ended solution technique* brings two concerns to bear. First, the resolution of the search must be sufficiently fine that the optimal solution may be resolved, without being so fine that computation time is wasted; and second, the cycle time must be consonant with the total cycle time of the system.

39.1 Speed Profile Generation

An estimate of the powertrain response may be produced by forward simulating the closed-loop response of the engine and / or braking system for a given commanded velocity V_d . For the powertrain system, the mass flow may be commanded using any control law. In the following example, the second generation LITHIA algorithm is used. For the braking system, the extended proportional servo control algorithm is used.

39.1.1 LITHIA Forward Simulation:

$$MF = [Q_1 V_e + Q_2 Load] RPM$$

$$R\dot{P}M = \frac{K_e}{I_e(RPM)} MF + \frac{K_e}{I_e} Load$$

$$\dot{V} = U_1 (RPM)^2 + U_2 (RPM) (Vel) + U_3 (Vel)^2 + U_4 Load$$

39.1.2 Braking System Forward Simulation:

$$x = G_1 V_e + G_2 Load$$

$$\dot{V} = \frac{1}{m} [F_{Load} - Kx]$$

$$K = \mu \frac{N_{max}}{x_{max}}$$

39.2 Search

Search may be performed by choosing a set of commanded velocities V_d about a nominal desired velocity provided by a higher planning agent. The vehicle system employs two unilateral control algorithms, implying that a method for determining how best to search the two independent spaces is needed. It is not necessarily the case that when the higher planning agent requests a velocity higher than the current velocity that the throttle will be employed. For example, it may be the case that other factors in the optimization algorithm actually require that the vehicle be slowed in the next cycle for the sake of stability.

39.3 Performance Index

Each speed profile generated during the search is evaluated against a performance index, which takes the form:

$$J = \int_{t_i}^{t_f} [a_1 V_e^2 + a_2 T_{R95} + a_3 V_O^2 + a_4 V_U^2 + a_5 S + a_6 MF] dt$$

Where V_e = velocity error; T_{R95} is the time to 95% of the commanded velocity; V_O is the velocity overshoot; V_U is the undershoot of a preset lower limiting velocity; S is a stability metric returned by the SPECTRE system; and MF is the mass flow of fuel to the engine. The members of the performance index were chosen to represent the criteria by which the servo control algorithm performance is typically judged, the need for a stability assessment, and the desire to conserve fuel whenever possible. These relative importance of each member is quantified by its associated constant a_i .

40. SUMMARY

This chapter has developed a novel, energy-based approach, to the problem of servo-controlling the speed of a conventional vehicle. The algorithms developed have demonstrated an entirely different class of behavior from the traditional approach. The performance has been shown to exhibit good command following, and robust disturbance rejection over rugged, unstructured terrain. The method is further shown to be easily extensible to current, proven planning methodologies such as feedforward optimal control. The inherent link between speed planning and motion evaluation was noted in the formulation of the performance index. In the following chapter, we will consider a method of performing computationally efficient motion evaluation for both speed and motion planning.

CHAPTER VIII - SPECTRE: A FREQUENCY SPACE MOTION EVALUATION SYSTEM

Motion evaluation comprises part of the central loop of the motion planning algorithm. Its computational complexity has a significant impact on the cycle time of the entire autonomous navigation software system. To date, only *time domain motion evaluation techniques*, based on ODE system models have been implemented. The numerical solutions to ODE's are given by a set of techniques derived from the Euler method which require the explicit computation of function derivatives to attain the forward propagation of the system state. Such techniques are of sufficient computational cost to be prohibitive to current computational resource for all but the simplest of system models. The fundamental limitations of system cycle time offer sufficient motivation to search for alternate evaluation techniques that may prove computationally tractable in the near term and computationally advantageous in the long term when more generous computational resources are available.

Current methods of motion evaluation follow from time domain system models, represented by the Cayley Hamilton theory, in which future states are shown to be the convolution of inputs over the present state. Such a model, when numerically implemented, require a similar forward propagation of state to arrive at future states, which methods will be shown to be computationally intractable under current resource constraints. It is a fundamental principle of mathematics that convolution in the time space is equivalent to multiplication in the frequency space. We employ this principle to eliminate the on-line computational burden of forward propagation, replacing instead with an off-line calibration and discretization of a frequency space filter model.

41. MODELLING

The motion of a mechanical system is completely described by the system Lagrangian, which is a function of the generalized positions and generalized velocities of the system. Given the Lagrangian of the system and a set of initial conditions for each generalized position and velocity, the future states of the system may be described. The Lagrangian may be reduced to a set of equations of motion through the application of the *first variation* - also called *Lagrange's equation*. It is a phenomena of Nature that many processes exhibit linear behavior through large portions of their bandwidth and may be described in a linear system model, such as the classical state space form:

$$\dot{x}(t) = Ax(t) + Bu(t)$$

The solution of this equation offers a closed form function which describes future system states in terms of the current state and inputs to the system. We will consider two such solutions, the first being the classical solution found in the Cayley Hamilton theory, the second being obtained through the application of the LaPlace transform.

41.1 Cayley Hamilton Theory

Cayley Hamilton Theory supplies a classical time domain solution to the linear system model. For the sake of completeness, we will repeat this solution here. Given a linear system, represented in state space form:

$$\dot{x}(t) = Ax(t) + Bu(t)$$

The time domain solution to this differential equation is the sum of the homogeneous and particular solutions. The homogeneous solution is obtained by considering the unforced system

$$\dot{x}_h(t) = Ax_h(t); x_h(t) = x_0$$

The particular solution satisfies the forced system equation with zero initial conditions. One obtains the particular solution by beginning with a function $w(t)$:

$$w(t) = e^{-At}x_p(t)$$

Taking the time derivative of $w(t)$ yields:

$$\begin{aligned}\dot{w}(t) &= \frac{d}{dt}[e^{-At}]x_p(t) + e^{-At}\dot{x}_p(t) \\ \dot{w}(t) &= -Ae^{-At}x_p(t) + e^{-At}[Ax_p(t) + Bu(t)] = e^{-At}Bu(t)\end{aligned}$$

Solving through integration yields:

$$w(t) - w(0) = \int_0^t e^{-A\tau}Bu(\tau) d\tau$$

We may now solve for the particular solution $x_p(t)$:

$$\begin{aligned}x_p(t) &= e^{-At}w(t) \\ x_p(t) &= e^{-At}\left[\int_0^t e^{-A\tau}Bu(\tau) d\tau\right]\end{aligned}$$

Which yields:

$$x_p(t) = \int_0^t e^{-A(t-\tau)}Bu(\tau) d\tau$$

A simplified form, given by Cayley Hamilton theory, assumes that there is no need to rotate the input u into a new space:

$$y_{zerostate} = \int e^{-At} u(t) dt$$

The response of the system is found to be a *time integral convolution* of the system inputs over the state transition.

41.2 The System Transfer Function

An alternate solution to the matrix equation is produced through the application of the LaPlace transform, and is often called the frequency-space solution. This technique yields a transfer function in s -space, as opposed to the integral formulation produced by the time domain solution. Again, for completeness sake, we step through the derivation process, beginning with the LaPlace transform, which is defined to be:

$$f(s) = L\{F(t)\} = \int_0^{\infty} e^{-st} F(t) dt$$

Application of the LaPlace transform to the state space equation yields the following:

$$\begin{aligned} \dot{x}(t) &= Ax(t) + Bu(t) \\ sx(s) &= Ax(s) + Bu(s) \end{aligned}$$

Which can be solved for $x(s)$:

$$x(s) = (sI - A)^{-1} Bu(s)$$

Let us define a new variable G :

$$G = (sI - A)^{-1} B$$

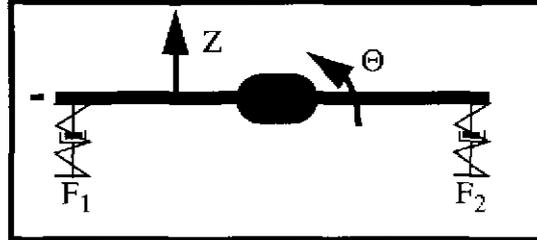
which simplifies the expression to:

$$x(s) = Gu(s)$$

where G may be considered to be the system transfer function.

42. DEVELOPMENT OF A SIMPLE SUSPENSION MODEL

In this section we will develop a very simple suspension model in both time and frequency domain, and perform a cost complexity analysis to determine how its evaluation effects the overall robustness of the system. Consider the following simplified model of the passive dynamics of the HMMWV. We consider the vehicle to be a point mass suspended from a massless beam, which is supported at either end by parallel spring-dashpot pairs:



Let the forces applied to the beam ends by the spring-dashpot pairs be labelled F_1 and F_2 . The positive sense of the two model dimensions, Z and Θ are noted in the diagram. We will denote the position of each of the spring-dashpot pair's lower end with p . This position is assumed to be fixed by contact with the ground. Applying N.S.L. in the vertical and rotational directions yields the following set of equations:

$$\begin{aligned} \Sigma F_z &= m\ddot{z} & \Sigma \Gamma &= I\ddot{\Theta} \\ -mg + F_1 + F_2 &= m\ddot{z} & F_2 l_2 - F_1 l_1 &= I\ddot{\Theta} \end{aligned}$$

Expanding the forces under the assumption of lumped parameter spring and dashpot models yields the following equation for the z -equation:

$$\begin{aligned} m\ddot{z} &= k(p_2 - z_2) + b(\dot{p}_2 - \dot{z}_2) + k(p_1 - z_1) + b(\dot{p}_1 - \dot{z}_1) - mg \\ m\ddot{z} + b(\dot{z}_2 + \dot{z}_1) + k(z_2 + z_1) &= -mg + b(\dot{p}_2 + \dot{p}_1) + k(p_2 + p_1) \end{aligned}$$

Expanding the forces under the assumption of lumped parameter spring and dashpot models yields the following equation for the Θ -equation:

$$\begin{aligned} I\ddot{\Theta} &= [k(p_2 - z_2) + b(\dot{p}_2 - \dot{z}_2)] l_2 - [k(p_1 - z_1) + b(\dot{p}_1 - \dot{z}_1)] l_1 \\ I\ddot{\Theta} + kz_2 l_2 - kz_1 l_1 + b\dot{z}_2 l_2 - b\dot{z}_1 l_1 &= kp_2 l_2 - kp_1 l_1 + b\dot{p}_2 l_2 - b\dot{p}_1 l_1 \end{aligned}$$

To further simplify the resulting Θ -equation, we will assume that $l_2 = l_1$.

$$I\ddot{\Theta} + bl(\dot{z}_2 - \dot{z}_1) + kl(z_2 - z_1) = bl(\dot{p}_2 - \dot{p}_1) + kl(p_2 - p_1)$$

which allows us to make the following trigonometric substitutes, where $L = 2l_1 = 2l_2$.

$$\begin{aligned} z_2 &= z + \frac{L}{2} \sin \Theta & \dot{z}_2 &= \dot{z} + \frac{L}{2} \cos \Theta \dot{\Theta} \\ z_1 &= z + \frac{L}{2} \sin \Theta & \dot{z}_1 &= \dot{z} - \frac{L}{2} \cos \Theta \dot{\Theta} \\ z_2 + z_1 &= 2z & \dot{z}_2 + \dot{z}_1 &= 2\dot{z} \\ z_2 - z_1 &= L \sin \Theta & \dot{z}_2 - \dot{z}_1 &= L \cos \Theta \dot{\Theta} \end{aligned}$$

Substitution in the Θ -equation yields the following:

$$I\ddot{\Theta} + \frac{bL^2}{2} \cos \Theta \dot{\Theta} + \frac{kL^2}{2} \sin \Theta = bl(\dot{p}_2 - \dot{p}_1) + kl(p_2 - p_1)$$

To further simplify the Θ -equation, the small angle assumption is employed:

$$I\ddot{\Theta} + \frac{bL^2}{2} \dot{\Theta} + \frac{kL^2}{2} \Theta = bl(\dot{p}_2 - \dot{p}_1) + kl(p_2 - p_1)$$

Substitution into the z-equation yields:

$$m\ddot{z} + 2b\dot{z} + 2kz = -mg + b(\dot{p}_2 + \dot{p}_1) + k(p_2 + p_1)$$

The resulting time domain model of the simple suspension model is thus given by the following equation set:

$$\begin{aligned} m\ddot{z} + b(\dot{z}_2 + \dot{z}_1) + k(z_2 + z_1) &= -mg + b(\dot{p}_2 + \dot{p}_1) + k(p_2 + p_1) \\ I\ddot{\Theta} + \frac{bL^2}{2} \dot{\Theta} + \frac{kL^2}{2} \Theta &= bl(\dot{p}_2 - \dot{p}_1) + kl(p_2 - p_1) \end{aligned}$$

Application of the LaPlace transform yields the following frequency domain representations of the simple suspension model:

$$z = \frac{-mg + (sb + k)(p_2 + p_1)}{s^2 m + 2sb + 2k}$$

$$\Theta = \frac{(sbl + kl)(p_2 + p_1)}{s^2 I + s \frac{bL^2}{2} + \frac{kL^2}{2}}$$

43. SOLUTION OF THE TIME DOMAIN PROBLEM

The time domain equations of motion are time dependent, ordinary differential equations specified as initial value problems. Numerical solutions are called for due to the discrete (non-functional) form of the input, which does not support direct solution of the integral. The numerical solution to such equations are obtained by first reducing the higher order differential equation to a set of coupled first order differential equations, which are then numerically integrated by replacing the derivative term with a first order difference. This principle is most easily seen in Euler's method, which may be considered the basis of all ODE solution techniques. Euler's method is given by the formula:

$$y_{n+1} = y_n + hf(x_n, y_n)$$

This formula, like all ODE solution techniques, calculates an approximation to the derivative of the function and advances the solution through a step of size h . Euler's method calculates the derivative estimate based solely on information from the beginning of the time step. Euler's method is impractically unstable, and of low order accuracy - $O(h^2)$. More advanced methods use information from throughout the time step h to estimate the function's derivative. There are three principal methods for the numerical solution of initial value problems:

1. Runge Kutta Methods
2. Richardson Extrapolation
3. Predictor Corrector Methods

Runge Kutta techniques are generally considered to be the "workhorse" techniques of numerical ODE integration. It operates by applying several Euler steps across the step and combining this information to match a Taylor's series expansion up to some chosen order - the most common being the fourth order Runge Kutta, which offers accuracy $O(h^5)$.

Richardson extrapolation techniques are based on the idea that one may, from several integration estimates, combine these estimates to a higher order estimate by considering the effect of a parameter of the integration, in this case the step size, and extrapolating the parameter to a limit - in this case, extrapolating the step size to zero. These techniques are limited in their application to smooth functions containing no singularities internal to the integration bounds. Richardson extrapolation is generally a more computationally efficient method to apply, but it is also potentially less stable, depending upon the form of the ODE. We will show that, although this technique is computationally more tractable, its output is not in a form useful to the motion evaluation problem.

The final technique, Predictor-Corrector is a middle of the road method that is typically dominated by Richardson extrapolation techniques when it comes to solving computationally expensive ODE problem where high precision is needed. Where such expense or precision is not present, Runge Kutta suffices, being more convenient to implement.

44. TIME COMPLEXITY ANALYSIS

The relative costs of performing motion evaluation in the time and frequency domains can be estimated without the necessity of resorting to experimental computation. It is often the case that the run time cost of basic algorithms may be significantly reduced through the application of secondary techniques such as adaptive step-size and error correction algorithms. In the case of real time motion evaluation, the brevity of the integration bounds is severe - only about 100 - 500 steps are taken along the observed path, thus there is little room for computational improvement over the cost of the basic algorithm. In this section, we will compute explicitly the estimated cost of the fourth order Runge-Kutta technique as applied to a typical run-time cycle of an autonomous navigation system. The models given above will be used in the example.

44.1 Runge Kutta Solution

Using the suspension model given above, we will now construct a fourth-order Runge-Kutta numerical integrator in two steps - in the first step we will explicitly calculate the costs of estimation for the z-model. In the second step, for the sake of brevity, we will not derive the costs, only present their final values. Recall the model:

$$m\ddot{z} + 2b\dot{z} + 2kz = -mg + b(\dot{p}_2 + \dot{p}_1) + k(p_2 + p_1)$$

Let us first convert the second order model into a set of first order equations:

$$\begin{aligned} \dot{y} &= -mg + b(\dot{p}_2 + \dot{p}_1) + k(p_2 + p_1) - 2by - 2kz \\ \dot{z} &= y \end{aligned}$$

For the sake of simplicity, let us assume that the inputs $(p_2 + p_1)$ and $(dp_2/dt + dp_1/dt)$ are provided in tabulated form and can be accessed at essentially zero cost. The fourth order Runge Kutta equations are given by:

$$\begin{aligned} y_{n+1} &= y_n + \left[\frac{1}{6} (k_1 + 2k_2 + 2k_3 + k_4) \right] h \\ z_{n+1} &= z_n + \left[\frac{1}{6} (l_1 + 2l_2 + 2l_3 + l_4) \right] h \end{aligned}$$

where:

$$\begin{aligned}
 k_1 &= f(y_i, z_i) \\
 k_2 &= f\left(y_i + \frac{h}{2}, z_i + \frac{hk_1}{2}\right) \\
 k_3 &= f\left(y_i + \frac{h}{2}, z_i + \frac{hk_2}{2}\right) \\
 k_4 &= f(y_i + h, z_i + hk_3) \\
 f(y, z) &= -mg + b(\dot{p}_2 + \dot{p}_1) + k(p_2 + p_1) - 2by - 2kz
 \end{aligned}$$

and:

$$\begin{aligned}
 k_1 &= f(z_i, y_i) \\
 k_2 &= f\left(z_i + \frac{h}{2}, y_i + \frac{hk_1}{2}\right) \\
 k_3 &= f\left(z_i + \frac{h}{2}, y_i + \frac{hk_2}{2}\right) \\
 k_4 &= f(z_i + h, y_i + hk_3) \\
 f(z, y) &= y
 \end{aligned}$$

Assuming that the cost of multiplications and additions are equal, we can evaluate the floating point operation cost of calculating the four terms k_i in the Runge-Kutta equations:

	Calculate Parameters	Calculate Function
Calculate k_1	0	10
Calculate k_2	5	10
Calculate k_3	5	10
Calculate k_4	3	10
Total = 55	13	40

Add to this the number of operations needed to combine the y-equation terms once:

	Operations Count
Sum Inner Block	5
Multiply by 1/6	1
Multiply by h	1
Add y_i	1
Total	8

bringing the total for the y-equation computation to 61. We may now do the same for the z-equation:

	Calculate Parameters	Calculate Function
Calculate k_1	0	0
Calculate k_2	5	0
Calculate k_3	5	0
Calculate k_4	3	0
Total = 9	13	0

and for combining:

	Operations Count
Sum Inner Block	5
Multiply by 1/6	1
Multiply by h	1
Add z_i	1
Total	8

which totals 21 for the z -equation. The total for one cycle of the z suspension model through the Runge Kutta implementation is equal to $61+21 = 82$ operations. Through a similar process we may arrive at the cost of implementing the Θ model, which we recall to be:

$$I\ddot{\Theta} + \frac{bL^2}{2}\dot{\Theta} + \frac{kL^2}{2}\Theta = bl(\dot{p}_2 - \dot{p}_1) + kl(p_2 - p_1)$$

which can be similarly converted to a pair of first order equations:

$$\begin{aligned}\dot{w} &= -mg + b(\dot{p}_2 + \dot{p}_1) + k(p_2 + p_1) - \frac{bL^2}{2}w - \frac{kL^2}{2}z \\ \dot{\Theta} &= w\end{aligned}$$

The cost of implementing fourth-order Runge Kutta to this model is found to be 79 operations, for a total of 161 operations for the system model. Let's now estimate the cost for running through a typical motion evaluation cycle. Let us assume that we are to check 10 paths, 20 meters long, 200 points per path.

$$CycleCost = \left(10 \frac{paths}{cycle}\right) \left(200 \frac{points}{path}\right) \left(161 \frac{operations}{point}\right) = 322,000 \frac{operations}{cycle}$$

for a machine that operates at 4MFLOPS, this cycle cost translates into 0.08 sec = 80 msec. This cost represents approximately 80% of the total cycle of a system operating at 10Hz. Remember, however, that this cost is suffered for the *simplest possible approximation* to the suspension response. More sophisticated models will likely consume the entire cycle time of the mobile robot. This is the fundamental near-term driver for exploring alternate formulations. In general, a more computationally efficient algorithm permits greater flexibility in the software design process - FLOP's freed from motion evaluation can then be used to improve the system resolution, or to improve the system cycle time and hence its reaction time.

44.2 Richardson Extrapolation Technique - Bulirsch-Stoer Method

The Richardson Extrapolation technique offers a more sophisticated approach to numerical integration of linear ODE's. The technique is based upon the same fundamental idea as Romberg integration, which is called *Richardson's deferred approach to the limit*. The idea is that the final answer of a numerical calculation, such as numerical integration, is in itself a function of many parameters, such as the step size. If one is able to fit a function to these stepsize / answer pairs, then one could drive the stepsize parameter to the zero limit, achieving a theoretically precise answer. The Bulirsch-Stoer Method uses a rational function to perform the zero limit extrapolation, and the modified midpoint method to provide the extrapolation data. Let's first consider the cost of

performing a midpoint evaluation. The modified midpoint method advances $y(x)$ from x to a point $(x + H)$ by taking n substeps each of size h :

$$h = \frac{H}{n}$$

The advancement cycle is given by the following set of equations in which we define a set of dummy variables z_i to carry forward the estimate. We first define the initial variables z_0 and z_1 :

$$\begin{aligned} z_0 &\equiv y(x) \\ z_1 &= z_0 + hf(x, z_0) \end{aligned}$$

The function is then advanced for $m = 1 \dots (n-1)$ using a center differencing estimate:

$$z_{m+1} = z_{m-1} + 2hf(x + mh, z_m)$$

The new value at $y(x+H)$ is then given by:

$$y(x + H) = \frac{1}{2} [z_n + z_{n-1} + hf(x + H, z_n)]$$

The Bulirsch-Stoer Method assembles a set of data by applying the modified midpoint method for sequentially finer step sizes. After each iteration, a rational function is fit to the data set and extrapolated to the zero stepsize limit. From the previous example, we know that the cost of evaluating the z -function is 10 FLOP's and the cost of evaluating the Θ -function is also 10 FLOP's. We can deduce the cost of performing a single iteration of the suspension model by considering each equation in the modified midpoint method, beginning with the initialization:

$$\begin{aligned} z_0 &\equiv y(x) \\ z_1 &= z_0 + hf(x, z_0) \end{aligned}$$

in which 1 function evaluation is performed, plus one addition and one multiplication, for a total of 12 FLOP's. Now consider the center differencing equation:

$$z_{m+1} = z_{m-1} + 2hf(x + mh, z_m)$$

which requires one function evaluation, plus two additions and three multiplications for a total of 15 FLOP's. The center differencing equation is iterated n times, where n is supplied by the Bulirsch-Stoer sequence $n = 2, 4, 6, 8, 12 \dots 2_{n-2}$. The total cost of the modified midpoint method is therefore:

$$MMM_{Cost} = 12 + 15n$$

The second operation in the Bulirsch-Stoer method is the fitting of the rational function and the subsequent extrapolation to the zero stepsize limit. For the sake of this comparative study, we will ignore this cost, and estimate the Richardson Extrapolation cost's lower limit based on the modified midpoint method's forward propagation cost. To form an estimate of the lower limit of the cycle time cost, let us assume that only four cycles $n = 2, 4, 6, 8$ are required which yield a cost of $42 + 72 + 102 + 132 = 348$ operations for both the z-equation and Q-equation. The total cost for the suspension is thus 696 operations per path.

$$\text{CycleCost} = \left(10 \frac{\text{paths}}{\text{cycle}} \right) \left(696 \frac{\text{operations}}{\text{path}} \right) = 6960 \frac{\text{operations}}{\text{cycle}}$$

The computational savings of this method is impressive; however, at this point we will point out the cost at which this savings comes, as shown in the extrapolation procedure. The sequential application of the modified midpoint operation produces a set of estimates of the value $y(x + H)$, tabulated for the number of substeps n taken to traverse H . The extrapolation procedure fits a rational function to this sequence, thus refining the local estimate of $y(x + H)$. Motion evaluation necessarily requires the examination of the path, not merely the path endpoints. Thus the integration process must achieve a resolution on an order much smaller than that of H , which requires integration on the stepsize order of the Runge-Kutta method. Thus, although the cost savings of the Richardson Extrapolation technique is tempting, it does not achieve the fundamental objective of ODE integration for motion evaluation.

44.3 Output Evaluation of Time Domain Solutions to the ODE

The forward propagation of the system model provides an estimate of the future states under that set of inputs. The evaluation of this output reduces this state estimate to a metric of goodness. Evaluation may require some further calculation of integrals or derivatives of the state, but typically reduces to a set of threshold comparison operations on state values such as velocity and acceleration. In the particular case of suspension evaluation, the principle metrics of concern are the acceleration, velocity, and travel of the two degrees of freedom.

44.4 Comments

Time domain motion evaluation techniques require three steps: input data preparation, ODE solution, and output evaluation. For the purposes of estimating computational cost the input data preparation cost is assumed to be zero, we rationalize this stance by noting that in most navigation systems a terrain map amenable to our purposes is available. We also note that the time domain motion evaluation techniques required a good state estimate everywhere along the path, as state itself is the final metric of comparison. This requirement acts as a cost driver to the motion evaluation system by forcing the integration stepsize to approximate the evaluation resolution along the path. The cost of the Runge Kutta method has been shown to be prohibitive for even the simplest suspension model.

45. SPECTRE

The SPECTRE system seeks to eliminate entirely the on-line or real-time computation of motion response, leaving only data preparation and output evaluation processes. We achieve this goal by moving to the frequency space or filter formulation which makes use of the system transfer function as the motion evaluation model. The fundamental principal at work here is that convolution (and therefore integration) in the time domain is equal to multiplication in the frequency domain. In order to make use of this method, the SPECTRE system must pay an additional cost in projecting the range data into the frequency domain, for which the Fast Fourier Transform method is employed.

45.1 The Filter Model

The SPECTRE system makes use of a frequency domain model of the system, which may be called a filter model. The filter model shows a system transfer function which filters the input pair ($p_2 + p_1$) onto the system output: z & Θ :

$$z = \frac{-mg + (sb + k)(p_2 + p_1)}{s^2 m + 2sb + 2k}$$

$$\Theta = \frac{(sbl + kl)(p_2 + p_1)}{s^2 I + s \frac{bL^2}{2} + \frac{kL^2}{2}}$$

Filters have the property of taking in a frequency of known magnitude and producing a change in magnitude and an accompanying phase shift. This information may be expressed in either vector form (usually in polar coordinates), or as the scalar power. The graphical display of the vector formulation produces the Nyquist and Bode plots, while the power formulation produces the power spectral density plot. The SPECTRE system makes use of the scalar power formulation to perform thresholding operations in frequency space. Recall the general filter formulation of the system model:

$$x(s) = Gu(s)$$

The fundamental idea behind SPECTRE is that, given this filter formulation of the model, we may discretize the input and output space into a set of buckets. The filter scales each input bucket i , to produce a power in output bucket i . In principle, the gains between input and output buckets may be computed through extensive off line testing of instrumented systems, or extensive modelling and exacting numerical analysis of system models. This produces a discrete set of n input / output gains:

$$x_i = G_i u_i$$

Having arrived at a set of gains between input and output buckets, we may assign to each input u_i such threshold that prevents the output bucket x_i from exceeding an imposed limit. Such a formulation requires that the system inputs be transformed into the frequency space, for which purpose we employ the Fast Fourier Transform.

45.2 The Fast Fourier Transform

The Fast Fourier Transform algorithm is a method of computing the discrete fourier transform in reduced time. It is $O(N\log_2N)$ as opposed to the ordinary discrete fourier transform's $O(N^2)$. The algorithm achieves this cost savings by employing a binary partitioning scheme that sequentially reduces an input set of N members into two sets of length $N / 2$. These two sets may then each be reduced to two more sets, and so on. The N members are split by separating the even components into one set and the odd components into the other. Having reduced an input set of N members into N individual sets, the Fourier transform of each individual member (the identity function) is taken. Recombination of these sets, with the appropriate frequency shift completes the FFT algorithm.

45.3 Cost of the FFT

Recombination of the partitions requires \log_2N cycles, in each cycle N complex additions and $N / 2$ complex multiplications are required. The cost of the FFT algorithm is thus found to be:

$$FFT_{cost} = \frac{3}{2}N\log_2N$$

The number of data points available across a path is on the order of 200. For simplicity of calculation, let's assume 256 data points for the FFT. The computational cost per path is thus 3072 FLOP's. For the assumed 10 paths per cycle, we produce a total cost of 30,720 FLOP's.

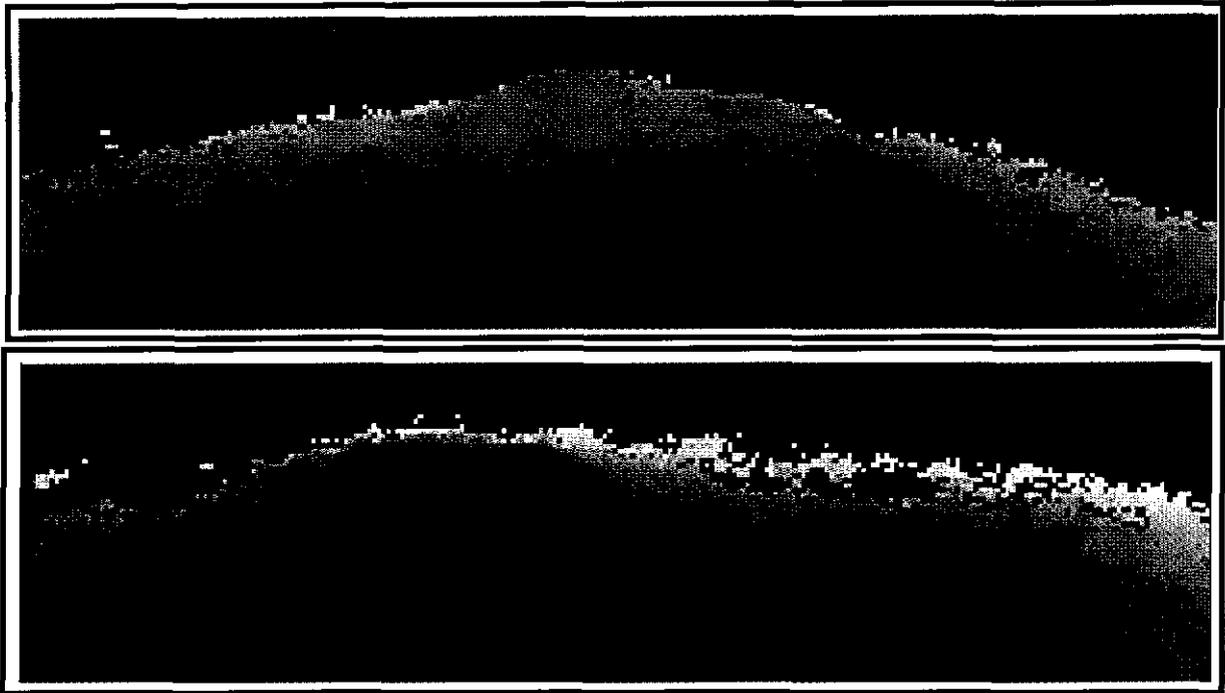
45.4 Cost Savings over Time Domain Approaches

The cost savings of FFT application over ODE integration is tremendous - approximately one order of magnitude. The cost of applying the FFT algorithm is approximately 10% of the cost of Runge-Kutta integration. The real advantage of the frequency domain approach is that *accuracy improvements do not require additional computation time, since the FFT is applied only to the input.*

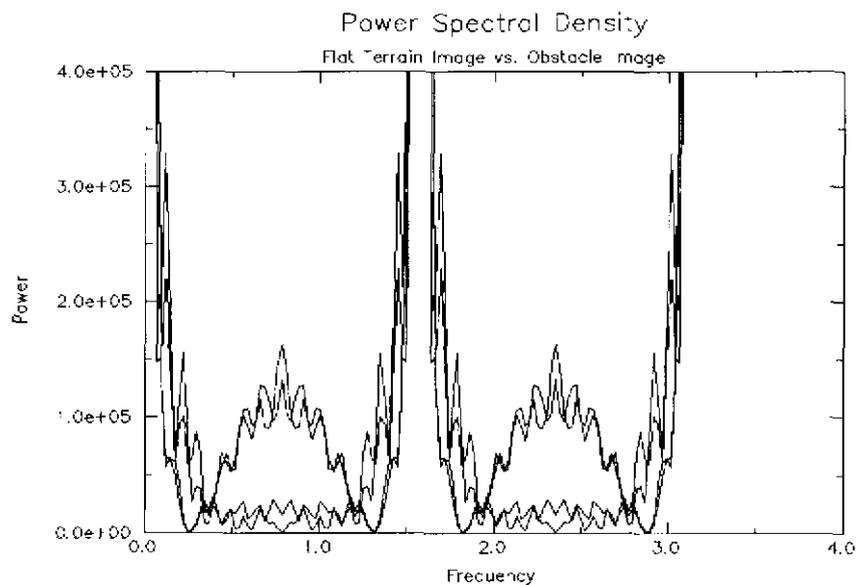
46. EXPERIMENTATION

A set of laser range images were taken at the slag heap site to determine what obstacles might look like, as seen in the frequency space. The FFT algorithm was applied to each column of the range image, with the resulting frequency curve converted to a power spectrum. The following two images are of the slag heap site - the first shows an essentially flat patch of terrain, one that is easily

traversable. The second image has a large mound that stretches across approximately the second quarter of the image - this mound is significant enough to be an obstacle to the HMMWV:



The following graphic superimposes four power spectra from four separate images. Two images are of the same patch of flat terrain, and the second pair of images capture the same obstacle-mound. The obstacle is easily identified by its large mid-frequency content, as seen below:



47. SUMMARY

Motion evaluation cost is one of the driving costs of an autonomous navigation system. The reaction time of the system as a whole is thus closely tied to speed of the motion evaluator, while the robustness of the system is a function of the quality of the evaluation. This chapter has described a new methodology for path evaluation that is mathematically equivalent to the traditional method, but at a fraction of the computational cost. The method is noted to require no increase in computational cost as model complexity is increased.

Part IV

Conclusions

CHAPTER IX - SUMMARY

This thesis addresses the autonomous speed control problem as applied to a large and important class of mobile robots - conventional vehicles. Speed control algorithms applied to conventional vehicles have historically lacked robustness in rugged natural terrain deployment. Early algorithms are characterized by performance sensitivity, poor stability, poor command following, and general tuning difficulties. An analysis of prior work in the field and a set of initial experiments indicated that unmodelled dynamics were the likely barrier to future performance improvements.

Modelling of the gross dynamics of the system enabled the discovery and description of the two phenomena principally responsible for prior poor performance. The magnitudes of the disturbance forces are large when compared to both vehicle momentum and engine output torque, and the system forcing function is incompatible with the use of a proportional controller about velocity. The second phenomena is explained by the physical properties of the system; specifically, the governing principle of the system forcing function is the First Law of Thermodynamics, an energy based law. A linear differential system equation may therefore only be formed in energy space, owing to the necessities of producing the particular solution to the differential equation. This result is non-intuitive as most forcing functions are force-based, resulting in the traditional velocity space linear differential equation. Unfortunately, the resulting equations, when transferred to velocity space, retain two state terms, yielding a coupled system equation.

Two basic goals were set for a new class of controllers - complete state decoupling and enhanced disturbance rejection. Two generations of LITHIA speed controllers were developed based on the engine and powertrain dynamical models developed in this thesis. The first generation LITHIA controller was formulated using direct model inversion. Head to head experimentation demonstrated that the LITHIA algorithm attains greater dynamic bandwidth (directly attributable to disturbance rejection) and faster time to commanded velocity than traditional approaches, with zero overshoot (attributable to proper model formulation). Classical control techniques, such as the use of s-space for design, were employed to improve the first generation LITHIA algorithm. The second generation LITHIA algorithm improves upon the first by achieving full state decoupling while reducing the number of gains from five to two.

The thesis demonstrated the extensibility of the LITHIA algorithm to optimal control planning, which trades off command following performance against fuel consumption and system stability. The classical LQ control techniques are first explored and shown to be applicable to the problem if formulated in energy space. There is an unfortunate lack of clarity as to the proper equivalence between optimization in energy space and optimization for stability, thus a feedforward planning method was instead implemented.

Optimal control planning systems require a real-time motion evaluation algorithm capable of analyzing the response of the vehicle's passive dynamical systems to predict the system stability margin. This thesis offered a fresh perspective on this problem by formulating a solution in the frequency space, rather than applying the Cayley-Hamilton time domain integral solution implemented by all previous systems. SPECTRE offers the unique advantage of filter

precomputation, reducing the evaluation problem to a thresholding operation on frequency inputs obtained through the application of the FFT algorithm to terrain images.

Modelling proved to be a key core technology to this thesis. The use of low fidelity models that qualified the gross dynamical interactions of the system allowed the designer to create a simple, but powerful *integrated speed control algorithm*. The creation of the SPECTRE system was brought about by a generally unstated precept of control theory - namely the convention of employing time domain and frequency domain mathematics and methods interchangeably while taking note of their individual strengths. SPECTRE is nothing more than the frequency domain equivalent of the time domain motion evaluator, yet it offers several distinct advantages over traditional methods.

48. RESULTS AND CONTRIBUTIONS

This thesis offers the first comprehensive look at autonomous speed control as applied to conventional vehicles - a more difficult problem than those previously solved by virtue of its complicated powertrain and suspension dynamics. The analyses and designs produced in this work result in a simple control algorithm that enables stable, wide-bandwidth powertrain control over rugged, unstructured 3-D terrain. These results, combined with the model-based approach of control theory, constitute a clear counter-example to the notion that the dynamics of combustion and torque conversion cannot be adequately modelled to effect good speed control - the LITHIA system provides good command following (as good as, or better than, a human) using *low-fidelity models*. As a result, one general contribution of the work is the partial *establishment of the physical theory underlying the autonomous control of conventional vehicle powertrains*. While I am convinced that the physical phenomena identified and modelled in this thesis constitute the foundation for the production of robust, stable conventional vehicle speed control algorithms, the physical theory is bound to expand with increasingly complex environments. Clearly, mechanics of soil deformation and other issues of terramechanics must soon be considered to provide further advancements in natural environments. None-the-less, the foundation created in this thesis is fundamental to such future work.

This thesis has shown that both classical control (s-space) design theory and state-space analysis techniques are powerful tools in the production and improvement of closed loop powertrain control laws. I feel that these results are particularly important as the community was at somewhat of a loss to marry this problem to basic control theory even a year ago. This match is due mostly to the identification of the proper linear space, and a willingness to interpret the linear energy regime in non-linear torque and velocity spaces when necessary. The opportunities taken resulted in the production of the simplified, second generation LITHIA algorithm. With respect to optimal control, classical LQ optimal control has been shown to be applicable to the autonomous speed control problem when formulated in the energy domain. Other practical design issues forced the extension of action space planning (or feedforward optimal control) to this non-spacial domain. Finally, an alternate, frequency space, formulation of the motion evaluation problem has been developed in which the significant advantage of filter pre-computation allows the implementation of fine resolution evaluation through input thresholding. The method has been shown to be an

identical mathematical solution to the traditional ODE integration approach, with the significant advantage of multiplication of input rather than convolution.

For purposes of clarity, I have restated the results and contributions of this thesis in as succinct a manner as possible. I believe that the following results and contributions sections properly characterize the technologies and distinguish them from prior art in the field.

48.1 Results

This thesis has produced *a novel energy-based servo-controls algorithm* which has demonstrated superior command following and disturbance rejection performance over all previously known algorithms. This energy or power-based algorithm has enabled robust, stable speed control over rugged, unstructured 3-D terrain for extensive distances while maintaining asymptotic convergence. The asymptotic behavior of the energy-based algorithm is not merely a performance improvement over the traditional methods, but is shown to be an entirely new class of response behavior, nearly identical to the response of standard second order systems. Experimentation demonstrates speed control performance that is as good as, or better than, that attainable by a human driving the same terrain.

This thesis has also developed *an optimal speed planner*, based on the feedforward method of motion planning. Speed planning is shown to be analogous to motion planning, with sufficient similarity to allow the combination of speed and motion planning into a single optimal planner. No specific claims are made of this work, other than the demonstration that the autonomous speed control problem may be formulated in a manner entirely analogous to the traditional autonomous navigation problem. Such a formulation leads in a straightforward manner to the eventual integration of autonomous speed control and autonomous navigation systems.

Finally, this thesis has produced *a novel frequency space motion evaluation methodology* that eliminates the need for expensive real time ODE integration. Extensive cost comparisons are provided which demonstrated a minimum of one order of magnitude difference in run time cost. The most compelling attribute of this new methodology is its constant run-time cost in the face of increasing model complexity and / or improved model resolution. The mathematics of the method are developed and a set of frequency space images of obstacles are obtained from laser range finder images.

48.2 Principal Contributions

There are two classes of contributions made in this thesis - fundamental and demonstrative. Fundamental contributions are both novel and unique in the field, and have a full grounding in either physical or mathematical theory. Demonstrative contributions lack such grounding, but offer instead a demonstration of either the applicability of a previously untried method, or the importance of an unnoticed phenomena. This thesis offers two fundamental contributions, and two demonstrative contributions.

48.2.1 Linearity of the Energy Domain

This thesis makes a fundamental contribution by proving that linear servo control techniques are appropriate to the throttle control portion of the autonomous speed control problem when the system equations of motion are formulated in the system's *energy space*, rather than the traditional velocity space. This phenomena is shown to arise from the necessity of enforcing a differential equation compatibility between the First Law of Thermodynamics (the governing principal of the system forcing function) and Newton's Second Law (the governing principal of the characteristic system equation). This has the effect of:

1. Allowing the full application of linear control techniques.
2. Allowing the formulation of a control law that produces a standard second order transfer function. Such a formulation allows the use of fundamental design techniques such as pole placement.

48.2.2 Frequency Space Obstacle Detection

This thesis makes a second fundamental contribution by developing an alternate, frequency-space, solution to the motion evaluation problem. This new method is shown to be of constant order in model complexity, whereas the time space formulation is guaranteed to be of non-constant order. In even the simplest suspension modelling case, the method is shown to produce a one order of magnitude improvement in run time.

48.2.3 Optimal Control Application

This thesis demonstrates that the classical LQ optimal control method can be applied to the autonomous speed control problem when formulated in the linear energy space rather than velocity space. Questions of the appropriateness of such an application arise, which leads to the ultimate rejection of the method in favor of the so-called action space planning, or feedforward optimal control method. Such a formulation demonstrates the compatibility of integrating autonomous speed control and autonomous navigation problems in a single optimal control planner.

48.2.4 Disturbance Reduction Functions

This thesis demonstrates that the reduction of the magnitude of the disturbance forces through mechanical means is tightly tied to the thermodynamic efficiency of the retro-fit design, as expressed by the Clausius statement of the Second Law of Thermodynamics. This work was ancillary to the main theme of the thesis, and so is relegated to an appendix.

49. TECHNICAL PERSPECTIVES

This thesis was produced through the application of a traditional engineering approach to problem solving. It was my intention to allow the engineering process to force the physics and the mathematics to form a solution, rather than to produce and test solutions which I invented. To a large extent this method worked well; the LITHIA algorithms were produced by-and-large through a straightforward model inversion followed by s-space refinement techniques. The frequency space work was also produced by such a method, in its case I back-chained the mathematics of

motion evaluation towards its source, until I found the fork in the road. In truth, I had had ideas of using such an alternate approach years earlier while performing my master's degree work, but the formal acceptance of the method as a true alternate to ODE integration was produced through such mathematical chaining.

There were a few ironies in this thesis which I have mused on without any general conclusion. First, I find it interesting to note that the analysis work of the thesis followed from a decision to produce force-based models of the vehicle, yet the final engine controller takes an energy-based form. I thought it particularly important to follow what I called a force-centric approach. It is hard to remember my original reasoning, but I find it likely that I followed the general recipe for producing control algorithms, the first step generally being the generation of a set of equations of motion. I considered different methods (Newtonian and Lagrangian methods) of arriving at the same goal - the production of the *differential equations of motion*. I think that this intermediate goal might not have been so blindly chosen if I had instead begun with a description of the power plant, rather than the system under actuation. The importance of energy-based methods arose only after the importance of control algorithm compatibility with the First Law of Thermodynamics was recognized.

I find that this work progressed at a faster rate once the problem of speed control was sub-divided into the problems of brake control and throttle control. As was shown in the thesis, the underlying physics are fundamentally different, requiring the production of different classes of control laws for the two. The inter-weaving of the two into a single control algorithm has been performed using a switching algorithm, with appropriate attention paid to the end conditions. I feel that more work should be done testing the switching algorithm under a variety of terrain conditions. There was insufficient time to do so in the course of this work, and for the purposes of algorithmic comparisons, it was inappropriate to test fully integrated systems. Having demonstrated robust capabilities, and reasonable integration strategies, it is appropriate to produce a second generation prototype system for extensive field testing.

49.1 A Speed Control Algorithm Comparison

In the prior work and linear experimentation chapters, a set of three previous control algorithms were analyzed or empirically tested. These control laws, which represent the prior state of the art in control theory, were tantalizingly close to an appropriate form. In this section, I compare each law to the second generation LITHIA control law produced in this thesis:

$$MassFlow_d = [Q1 (Vel_e) + Q2 (Load)] (RPM)$$

Recall the lookup table approach of Murphy, which took the form:

$$\phi = C + K_1 V_e + K_2 \sin\theta$$

the second two terms of the control law are an estimate of the required torque. If this estimate had been multiplied by the engine RPM, a power law would have been produced, though the critical

throttle valve transform is still required - we have previously noted its importance. I have not performed a rigorous comparison, but the values forwarded from Karl Murphy rise as a function of the commanded speed, which would be consistent with the required increase in RPM. Recall now the second, PID-style approach examine in the preliminary section of this thesis:

$$\varphi = K_P V_e + K_I \int V_e + K_D \dot{V}_e$$

This algorithm requires the addition of the feedforward term, as well as the engine RPM multiplication, and the throttle valve geometry transform.

It is significant to note that the mathematical alteration of these algorithms to realize much higher performance is almost insignificant. The production of robust behavior is in no way dependent upon complicated real-time modelling or computation - merely (and importantly) it is entirely dependent upon the proper identification of the underlying physical mechanics and proper design model formulation. The simplicity of the resulting control algorithm supports the notion that the gross physics of powertrain dynamics dominate the problem.

49.2 Performance Improvements through Increased Model Fidelity

This thesis demonstrated that speed control performance, as good as, or better than a human operator driving the same terrain could be attained through the use of simplified engine and powertrain models. It would be interesting to see whether any significant performance improvements can be made by improving the fidelity of the models on which the control algorithm is based. It has been suggested to me that some simplifying assumptions made in the control modelling section, such as the steady state intake manifold assumption, may be inappropriate and incompatible with the attainment of good control. The thesis has shown such significant performance improvements over the prior art that this assumption seems to have been validated. None-the-less, it is interesting to consider how the future performance of the algorithm might be improved through more rigorous modelling activities.

I believe that the load estimate is of critical importance to the performance of the algorithm, and so I find it most likely that performance improvements can be made by either improving the terramechanics models themselves, or by providing a feedback or feedforward measurement of terramechanical properties. I have previously noted that the difference in rolling friction between packed and unpacked slag was significant enough to cause reasonable steady state errors. Adaptation of the load seems necessary for future all-terrain navigation. Of course, the general assumption of non-yielding terrains made in the course of the thesis must also be confronted to produce autonomy in general, non-homogeneous, yielding terrains. Other model improvements, such as an improved mapping of engine and powertrain functions, while useful for improved simulation, do not seem likely to produce significant changes in control performance. For these reasons, I consider terramechanics to be generally the next most significant topic to be addressed in autonomous speed control for conventional operations in unstructured terrain.

49.3 Pitch Measurements

The pitch measurement is the most critical measurement made in the LITHIA control law. Yet, it is very difficult to accurately measure. There are two fundamental problems with accurate measurements - the first is the physical location of the device. Pitch sensors are mounted on the moving vehicle, with the entire suspension between their measurement axis and the plane which they actually should be measuring. The dynamic response of the suspension is thus mixed with the measurement signal, and at sufficiently low frequency that the vehicle can respond to this error. In the tests performed over the course of the thesis work, such responses occurred only when the HMMWV mounted a small high-frequency object, such as a log or a cement block, which produced a lasting low frequency suspension response. Although in this particular set of work the response did not significantly affect the performance of the vehicle, in general this dynamical transform exists between the measurement device and the slope of interest and should be considered. The use of non-contact sensors such as laser rangefinders to estimate the slope does not eliminate this problem as the rangefinder measurements themselves must be referenced through a pitch sensor. This problem is entirely fundamental in its nature - it is an instantiation of Einstein's elevator, the thought problem in general relativity in which Einstein demonstrated the equivalence of gravitation and acceleration. This is one of the basic problems that designers of high fidelity inertial navigation units face. This problem will only be exacerbated as we move to more rugged terrains where suspension response becomes more exaggerated. One possible solution to the problem is to filter the pitch measurements against a frequency spectrum of the terrain, taken by an imaging sensor such as a laser range-finder. Such a process might allow us to separate suspension response frequencies from spatial frequencies in the terrain. A second possibility, though computationally expensive, is the inclusion of a roll (or in this case, pitch) prediction model, akin to those used in naval ships to attain high accuracy projectile targeting. This seems to be rather the same method as the frequency methods previously mentioned.

49.4 Frequency Space Motion Evaluation

This thesis addressed the fundamental mathematics of motion evaluation within its scope; however, the full benefits of frequency-space evaluation have not yet been measured. To do so requires a modelling effort beyond the scope of this thesis, and the current resource of the university. The brute force method of performance quantification would be to tune the evaluator against high performance filter models produced through exhaustive modelling and testing of the target vehicles. I am rather opposed to such a method, strictly on the grounds of its abhorrent cost, for applications where the domain of operation can be adequately described, or for those domains where limit-approaching performance is not required. For autonomous navigation over unstructured, unknown terrain in performance critical arenas, such as space exploration, the cost of extensive modelling is reasonable compared to the cost of failure. However, for many other applications, threshold tuning through experimentation, or through on-line learning or adaptation may suffice.

50. FUTURE RESEARCH AND WORK

The work addressed in this thesis has advanced to the point where engineering and commercialization are appropriate. Many of the fundamental problems of autonomous speed control have been overcome, notably the production of stable, asymptotic servo-control. The problem has generally been grounded to established controls theory indicating that further work is most likely to result from engineering improvements to the base technologies. I have stated my belief that the next research front in speed control is opposed by the problems of terramechanics. The work in this thesis assumed rigid soils, effectively allowing us to ignore the traction energy transferal problem. Limitations in soil stability limit the effectiveness of current techniques in the presence of mud or yielding soils.

I find it likely that the next generation of speed controllers will build upon the current work (one cannot escape the physical limitations of conventional vehicles addressed in this thesis) by adding further feedback in the form of torque measurements or torque observers at the contact points. Such feedback could enable either low level reactive-type torque limitation schemes or the implementation of higher level plans that perhaps shift to higher (low torque) gears. Advances in autonomous speed planning also call for improvements in either basic sensor implementation or sensor interpretation algorithms, either of which enables the recognition or measurement of viscous, limited trafficability soils before encounter.

I would like to further pursue the formulation of the LQ problem. Perhaps there is a way to satisfactorily equate the energy based perspective with the velocity based perspective. I'm not sure that the pursuit of this goal will bear any relevant practical result. I believe that the formulation presented in this thesis is sufficiently broad that in practice it encompasses as a subset the LQ result; however, there is a certain elegance that the given solution lacks and that a usable LQ solution may possess.

I would like to combine the autonomous speed controller with a similar autonomous navigation system (the steering problem) to produce a single optimal controller. I believe that there is an opportunity to do just that by combining the RANGER system (Kelly[16]) with LITHIA. The proposed RANGER-LITHIA system would then use LITHIA as an *intermediary integrated control* device and SPECTRE as a predictive tool to perform optimal motion planning in the two dimensional steering- propulsion space.

51. GENERAL PERSPECTIVES

I wanted to produce a systems work that demonstrated that the autonomous speed control problem exists as a complementary problem to the autonomous navigation problem, and that its solution is a necessary part of the achievement of full autonomy in conventional vehicles. I also wanted to demonstrate that the problem is fundamentally an *engineering* problem, and that when it is approached from an engineering perspective a reasonable, straightforward solution is achieved. The connection to control theory and physical mechanics which was forced throughout the thesis had the effect of grounding the work to physical theory - the continuity of this connection made all

of the difference between this work and its prior art. In short, I credit the use of an engineering perspective with the advancement of this work above all other factors.

Conventional vehicles represent the most significant population of work vehicles in the planet. They are engaged in work tasks in nearly every field from agriculture to mining to forestry to military service, usually in the performance of basic utility and haulage service. I believe that the legacy of this work will be the production of practical commercial autonomous systems for use in such outdoor applications where conventional vehicles perform routine, mundane, basic haulage tasks.

Part V

Appendices

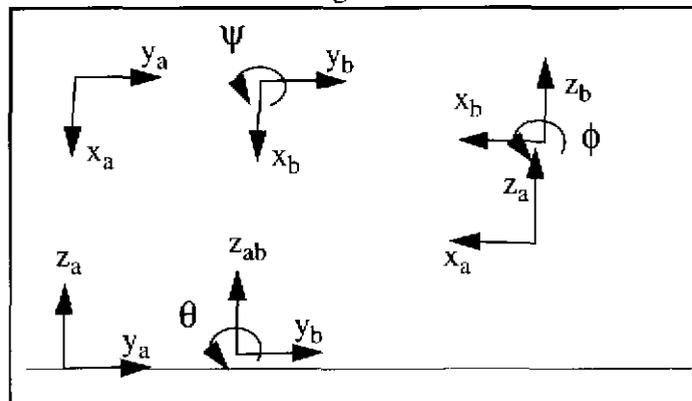
APPENDIX I - KINEMATIC TRANSFORMS

52. DEFINITION OF THE KINEMATIC TRANSFORMATIONS¹

In this section, the kinematic transformations that are necessary to move from a coordinate system attached to the vehicle frame and a coordinate system situated in a fixed inertial frame are presented. Note that these transformations are *kinematic* (mathematical) and are to be distinguished from the dynamic laws that will transform physical measurements taken in one frame to their equivalent in another. Let two general frames be defined called 'a' and 'b' and consider the moving axis operations which transform the fixed frame 'a' into coincidence with the moving frame 'b'. *In order*, these are:

- translate along the (x,y,z) axes of frame 'a' by (u,v,w) until its origin coincides with that of frame 'b'
- rotate about the new z axis by an angle Ψ called **yaw**
- rotate about the new x axis by an angle Θ called **pitch**
- rotate about the new y axis by an angle ϕ called **roll**

Angles are measured counterclockwise positive according to the right hand rule. These operations are indicated below for the case of transforming the fixed frame into the vehicle frame.



1. This chapter originally appeared in a technical report entitled "Essential Kinematics for Autonomous Vehicles," by Alonzo Kelly. CMU-RI-TR-94-14. The text of this chapter has been altered to suit the purposes of this report.

The forward kinematic transform that represents this sequence of operations is, according to Kelly's rules for forward kinematics:

$$\begin{aligned}
 T_{moving}^{fixed} &= Trans(u, v, w) Rotz(\psi) Rotx(\theta) Roty(\phi) \\
 T_{moving}^{fixed} &= \begin{bmatrix} 1 & 0 & 0 & u \\ 0 & 1 & 0 & v \\ 0 & 0 & 1 & w \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} c\psi & -s\psi & 0 & 0 \\ s\psi & c\psi & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & c\theta & -s\theta & 0 \\ 0 & s\theta & c\theta & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} c\phi & 0 & s\phi & 0 \\ 0 & 1 & 0 & 0 \\ -s\phi & 0 & c\phi & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \\
 T_{moving}^{fixed} &= \begin{bmatrix} (c\psi c\phi - s\psi s\theta s\phi) & -s\psi c\theta & (c\psi s\phi + s\psi s\theta c\phi) & u \\ (s\psi c\phi + c\psi s\theta s\phi) & c\psi c\theta & (s\psi s\phi - c\psi s\theta c\phi) & v \\ -c\theta s\phi & s\theta & c\theta c\phi & w \\ 0 & 0 & 0 & 1 \end{bmatrix}
 \end{aligned}$$

In the case of applied forces, translation is not of interest. The four by four matrix can thus be collapsed to a three by three rotation matrix:

$$T_{moving}^{fixed} = \begin{bmatrix} (c\psi c\phi - s\psi s\theta s\phi) & -s\psi c\theta & (c\psi s\phi + s\psi s\theta c\phi) \\ (s\psi c\phi + c\psi s\theta s\phi) & c\psi c\theta & (s\psi s\phi - c\psi s\theta c\phi) \\ -c\theta s\phi & s\theta & c\theta c\phi \end{bmatrix}$$

The inverse transformation, from the moving frame 'b' to the fixed frame 'a' is formed by taking the transpose of this matrix:

$$T_{fixed}^{moving} = \begin{bmatrix} (c\psi c\phi - s\psi s\theta s\phi) & (s\psi c\phi + c\psi s\theta s\phi) & -c\theta s\phi \\ -s\psi c\theta & c\psi c\theta & s\theta \\ (c\psi s\phi + s\psi s\theta c\phi) & (s\psi s\phi - c\psi s\theta c\phi) & c\theta c\phi \end{bmatrix}$$

52.0.1 The ω - Euler Angle Relationship

The roll, pitch, and yaw angles are, as defined, measured about moving axes. Therefore, they are a sequence of **Euler angles**, specifically, the z-x-y sequence¹. The Euler angle definition of vehicle attitude has the disadvantage that the roll, pitch, and yaw angles are not the quantities that are actually indicated by strapped down vehicle mounted sensors such as gyros. The relationship between the rates of the Euler angles and the angular velocity vector is nonlinear. The angles are measured neither about the body axes nor about the navigation frame axes. It is important to know the exact relationship between the two because it provides the basis for determining vehicle attitude from angular rate measurements.

1. The sequence depends on the convention for assigning the directions of the linear axes.

In order to determine the angular velocity, consider that the total angular velocity is the sum of three components, each measured about one of the intermediate axes in the chain of rotations which bring the navigation frame into coincidence with the body frame. Using the fundamental transforms, each of the three rotation rates are transformed into the body frame by the remaining rotations in the sequence to give the result in the body frame.

$$\omega^b = \begin{bmatrix} 0 \\ \dot{\phi} \\ 0 \end{bmatrix} + \text{rot}(y, \phi) \begin{bmatrix} \dot{\theta} \\ 0 \\ 0 \end{bmatrix} + \text{rot}(y, \phi) \text{rot}(x, \theta) \begin{bmatrix} 0 \\ 0 \\ \dot{\psi} \end{bmatrix}$$

$$\omega^b = \begin{bmatrix} \omega_x \\ \omega_y \\ \omega_z \end{bmatrix} = \begin{bmatrix} c\phi\dot{\theta} - s\phi c\theta\dot{\psi} \\ \dot{\phi} + s\theta\dot{\psi} \\ s\phi\dot{\theta} + c\phi c\theta\dot{\psi} \end{bmatrix} = \begin{bmatrix} c\phi & 0 & -s\phi c\theta \\ 0 & 1 & s\theta \\ s\phi & 0 & c\phi c\theta \end{bmatrix} \begin{bmatrix} \dot{\theta} \\ \dot{\phi} \\ \dot{\psi} \end{bmatrix}$$

This result gives the vehicle angular velocity expressed in the body frame in terms of the Euler angle rates. Notice that when the vehicle is level the x and y components are zero and the z component is just the yaw rate as expected.

This relationship is also very useful in its inverted form. One can verify by substitution that:

$$\begin{bmatrix} \dot{\theta} \\ \dot{\phi} \\ \dot{\psi} \end{bmatrix} = \begin{bmatrix} \omega_x c\phi + \omega_z s\phi \\ \omega_y - t\theta[\omega_z c\phi - \omega_x s\phi] \\ [\omega_z c\phi - \omega_x s\phi]/c\theta \end{bmatrix} = \begin{bmatrix} c\phi & 0 & s\phi \\ t\theta s\phi & 1 & -t\theta c\phi \\ -\frac{s\phi}{c\theta} & 0 & \frac{c\phi}{c\theta} \end{bmatrix} \begin{bmatrix} \omega_x \\ \omega_y \\ \omega_z \end{bmatrix}$$

because $\omega_z c\phi - \omega_x s\phi = c\theta\dot{\psi}$

APPENDIX II - DISTURBANCE REDUCTION

It is an unwritten rule of control system design that one should consider how the dynamics of the system might be improved through *mechanical means* before implementing a control strategy. It is often the case that simple mechanical additions or improvements to the system can reduce system non-linearities, improve system response time, or improve the quality of system feedback. In the case of the autonomous speed control problem, the most significant disturbances arise due to the gravity forces acting upon the vehicle. The relative strength metric derived in the disturbance analysis section will be expanded to show that disturbance reduction is equivalent to an improvement in the propulsive force to weight ratio, or equivalently, the power to weight ratio of the vehicle.

For the purposes of our application to the NavLab II, it is difficult to increase the available propulsion force, therefore we choose to decrease vehicle weight through a configuration redesign. In the course of the redesign some weight was reduced through the elimination of unnecessary components and the replacement of components with lighter weight alternatives. However, a significant portion of the total weight shed was attributable to small changes made in the heat transfer efficiency of the on-board computing's active cooling system. The strong relationship between *thermal efficiency and system weight is similar to problems* encountered in the design of other mobile systems, including automobiles and satellites, in which a thermal partition must be maintained at an energy cost to the system - classically stated in the Clausius statement of the Second Law of Thermodynamics. The attendant weight cost is attributable to the required size up of the power production system.

In this chapter we produce a set of cost functions that allow the system designer to consider the effects of component addition to the physical speed control performance of the machine, as described by the system power to weight ratio. We will show that the relationship between system weight and system power requirements is non-linear and possessive of a step like behavior properly described by a floor function. Finally we will show that disturbance reduction is most strongly linked to the *thermal efficiency of the active cooling system*.

53. RELATIVE ANALYSIS - THE CASE FOR WEIGHT REDUCTION

The case for performing mass reduction is based on the same relative analysis techniques presented in the modelling chapters. The basis of the question is whether the reduction of mass improves the physical ability of the vehicle's propulsion system (independent of any speed control algorithms) to reject externally applied disturbance forces. We expand the relative strength analysis to include a powertrain propulsion force term F_p :

$$S_r = |F_p| + |F_g| + |F_d| + |F_r|$$

The relative strength of propulsion is then found to be:

$$S_p = \frac{F_p}{S_r} = \frac{F_p}{mg(f + |\sin \Theta|) + \frac{1}{2}\rho V^2 C_d A + F_p}$$

We can see from this relationship that as m , the mass of the vehicle, increases the relative strength of propulsion S_p decreases. Reduction of mass is thus one method for improving the relative strength of propulsion.

54. COST FUNCTIONS AND PERFORMANCE METRICS FOR RECONFIGURATION

The performance of a mobile machine can be quantified by the ratio of the generated locomotive power to the weight of the machine - this ratio is called the **power to weight ratio**. The power to weight ratio is a good first order approximation to acceleration, and therefore to the physical ability of the machine to attain velocity and reject disturbances. In this section, the power to weight ratio of a reconfigured machine will be derived as a function of the power to weight ratio of the original machine altered by the cost of autonomy.

54.1 Cost Functions

The cost of autonomy can be predicted during the design phase using two cost functions, one for power and a second for weight. The **total power cost** is the sum of the power cost of those components powered from the vehicle's main engine and those components powered from an auxiliary on-board power source.

$$A_p = A_p^{Loc} + A_p^{Aux}$$

The **locomotive power cost** is equal to the summation of the power loads of the components that are to be powered by the vehicle's main engine, pro-rated by the efficiency of the machine's power generation system.

$$A_p^{Loc} = \left(\frac{1}{\eta_{Loc}} \right) \sum P_i^{Loc}$$

The **auxiliary power cost** is equal to the summation of the power loads of the individual components that require auxiliary power, pro-rated by the efficiency of the on-board auxiliary power generation system:

$$A_p^{Aux} = \left(\frac{1}{\eta_{Aux}} \right) \sum P_i^{Aux}$$

The total weight cost is comprised of two components, the first is the sum of the weights of the individual components which require power. The second is the weight of the auxiliary power generation system, which is found by sizing the system against the auxiliary power requirement. Because power generation components are available in discrete sizes, the weight of the power generation system is the weight of the smallest system capable of supplying the auxiliary power. This relationship is represented by the *floor function*.

$$A_w = \sum W_i + \left[A_p^{Aux} \right]$$

Consider the power cost of an individual component i , in the power cost summation above. Each component contributes its power cost to the total power cost; however, the power cost of some components is a function of the power costs of other components. If raising the powerdraw of component A affects the powerdraw of component B, then component A exhibits **power cost duality**. An example of such a component is a computing card. As the powerdraw of the computing card increases, so does its need for air conditioning. Raising the powerdraw of the card may require, in turn, raising the powerdraw of the air conditioner. The overall effect on the system is more than just the card's powerdraw increase.

54.2 Configuration Power to Weight Ratio

The capability of the configuration¹ is defined by the configuration power to weight ratio, which is found as follows. The configuration power C_p is the power generation capability of the original machine, M_p , depleted by the locomotive power cost:

$$C_p = M_p - A_p^{Loc}$$

The configuration weight is the sum of the original machine's weight M_w and the weight of the autonomy hardware.

$$C_w = M_w + A_w$$

The final configuration power to weight ratio is thus found to be:

$$\Pi = \frac{C_p}{C_w} = \frac{M_p - A_p^{Loc}}{M_w + A_w}$$

54.3 Configuration Efficiencies

As a basis of comparison among concept configurations, two efficiency metrics are introduced, the *Coefficient of Power Performance* and the *Coefficient of Weight Performance*. The **coefficient of**

1. The *configuration* is the machine after alteration for autonomy.

power performance describes the distribution of power toward cognitive activity and support. A ratio of zero would require no power to support cognition and would be the most efficient possible configuration. A ratio of 1 would require all of the robot's power just to think, and would be incapable of doing any work.

$$\beta = \frac{A_p^{Loc} + A_p^{Aux}}{M_p + A_p^{Aux}}$$

The **coefficient of weight performance** describes the distribution of weight toward cognitive activity and support. A ratio of zero would indicate a massless cognition system, while a ratio of one would indicate an immobile system - a brain without a body.

$$\Gamma = \frac{A_w}{C_w}$$

54.4 Component Addition and Π Deration

Consider an expanded form of the configuration power to weight ratio equation.

$$\Pi = \frac{M_p - \left(\left(\frac{1}{\eta} \right) \sum P_i^{Loc} \right)}{M_w + \sum W_i + \left[\left(\frac{1}{\eta} \right) \sum P_i^{Aux} \right]}$$

The form of this equation provides some insights into the nature of Π deration as components are added to the configuration. First, note that in the numerator, the configuration power is derated discretely, in an amount proportional to the powerdraw of the individual vehicle-powered components. Second, note that, in the denominator, the weight of the configuration is increased discretely by the weight of the individual components. Finally, note that, in the denominator, the presence of the floor function indicates that the addition of an auxiliary powered component *may or may not* further increase the weight of the vehicle configuration.

To summarize, adding a component can have the following effect:

If powered by the vehicle:

- Reduces the configuration power.
- Adds its own weight to the configuration weight.

If the component is powered by an auxiliary source:

- Adds its own weight to the configuration weight.

- Increase the auxiliary power load, which MAY force a power and / or weight size up of the power generation system.

If the component exhibits *cost duality* it also:

- MAY force a power and / or weight size up in a second component.

54.5 Cost Duality and Hidden Floor Functions

If a component exhibits cost duality, it may force a power and / or weight size up in a second component. Cost duality exists when a second component must be sized to meet the specifications of a group of subcomponents. The given example in this work has been air conditioning. Because the second component is being sized, its power draw and weight *are also floor functions*. In the expanded Π equation, *the individual weights W_i and power draws P_i may themselves be floor functions*, lending an increasingly complex behavior to the configuration process. Because the floor function exhibits a step like behavior at its threshold, the significant of changes in Π due to component addition are often more a matter of how close the configuration is to the threshold than on the component itself.

55. THE POWER WEIGHT SPIRAL

The power-weight spiral is an effect commonly noted by systems integrators. Occasionally, the addition of component *A* results in the need to upsize one or more additional components. It is also possible that the upsizing of these additional components will, in turn, cause the performance of component *A* to be insufficient, requiring it to be upsized, which creates positive feedback. There are two fundamental types of power-weight spirals. The first is truly a spiral, and is called the *Positive Feedback Power-Weight Spiral*. The second is not truly a spiral, in the sense that it is not a closed loop response, but merely a chain reaction - this reaction is called the *Chain Reaction Power-Weight Spiral*.

55.1 Positive Feedback Power-Weight Spiral - the Power Weight Density

A positive feedback power-weight spiral occurs as a result of internal coupling in the Π equation. Consider the following simplified version of the Π equation.

$$\Pi = \frac{M_p - A_p^{Loc}}{M_w + A_w}$$

Assume now that the value of Π is too low for the needs of the mobile robot; the designer wishes to increase Π . Further assume that he or she chooses to increase Π by increasing M_p .¹ Because there is no such thing as a massless power source, *increasing M_p necessarily affects M_w as well*. M_p and

1. Note that increasing the auxiliary power can only reduce Π , as it increases weight without supplying additional *locomotive* power.

M_w are functionally coupled. The change in Π for a given change in M_p can be calculated as follows:

$$\frac{\partial \Pi}{\partial M_p} = (M_p - A_p^{Loc}) \frac{\partial}{\partial M_p} \left((M_w + A_w)^{-1} \right) + \left(\frac{\partial}{\partial M_p} (M_p - A_p^{Loc}) \right) (M_w + A_w)^{-1}$$

Taking the partial derivatives and noting that there is no functional dependence of A_w or A_p on M_p :

$$\frac{\partial \Pi}{\partial M_p} = \frac{(M_w + A_w) - \frac{\partial M_w}{\partial M_p} (M_p - A_p^{Loc})}{(M_w + A_w)^2}$$

The condition for increasing Π is found when this partial derivative is greater than zero:¹

$$(M_w + A_w) - \frac{\partial M_w}{\partial M_p} (M_p - A_p^{Loc}) > 0$$

Noting that $(M_w + A_w)$ and $(M_p - A_p)$ are both always positive quantities, the partial derivative may be isolated without loss of generality:

$$\frac{\partial M_w}{\partial M_p} < \frac{(M_w + A_w)}{(M_p - A_p^{Loc})} = \frac{1}{\Pi}$$

This equation may be inverted to yield the following condition for increasing Π , with two caveats that will be explored in a moment.

$$\frac{\partial M_p}{\partial M_w} > \Pi$$

This condition may be interpreted as follows. The ratio of the *change in power production* to the *change in weight* must be *greater than the current configuration power to weight ratio* in order for the addition of the component to increase Π . An intuitive way to think about this relationship is through a density analogy. Π is the **power to weight density** of the system. Increasing Π increases the system density. A density can only be increased by adding something that is *more dense* to the system. The ratio of change in power to change in weight is the *density of a differential quantity*. The condition requires the density of the differential to be greater than the density of the system.

1. I have omitted some algebraic steps here by noting that the denominator is always positive; I can thus consider only the numerator in the inequality.

The *positive feedback power weight spiral* occurs when a power producing component that does not meet the condition is added to the configuration. Instead of raising Π , as expected, the additional power producing element derates Π . Positive feedback can occur if the designers mistakenly continue to add more power producing elements - effectively worsening the situation.

55.2 Two Caveats to the Power Weight Density Condition

When the differential quantity was inverted to yield the last equation, there was *no consideration* given to the possibility that the differential itself could be a negative number, or undefined. It could be the case that the replacement of a power source with a new power source (such as the replacement of a gasoline engine with a jet-engine) could result in a net power increase and a weight *decrease*. The case where the partial derivative is strictly negative is therefore an equally acceptable condition, *as long as the power is increasing*. In the case where the ratio is undefined, due to a zero change in weight, the condition is governed by the sign of the change in power. A positive power change raises Π , a negative power change lowers Π . The following conditions also lead to an increase in Π .

- Negative Ratio: due to increasing power, decrease in weight:

$$\left(\frac{\partial M_p}{\partial M_w} < 0 \right) \wedge (\partial M_p > 0)$$

- Undefined Ratio: due to zero weight change, with increase in power.

$$(\partial M_w = 0) \wedge (\partial M_p > 0)$$

55.3 Chain Reaction Power-Weight Spiral - the Floor Function

A chain-reaction type power-weight spiral occurs when one or more cost duality linked components are near their thresholds when a component is added. What appears to be a closed loop effect is actually the result of exceeding several component thresholds one after the other. It is *significant to note* that small changes in variables (power draw) can result in large changes in Π (power to weight ratio). Note also that *the effect (and thus the spiral) is reversible*. If the requirements of a functional design can be reduced by a small amount, it is possible to increase the power to weight ratio by a large amount. Intuitively, this means edging the design under the limits of the next smaller series (thus lighter) of power components. This point will turn out to be key to the application example in this chapter.

56. APPLICATION TO THE NAVLAB II MOBILE ROBOT

56.1 M_p and M_w

Changes to the power production of the vehicle's engine and the original weight of the vehicle were ruled out as too costly for the following reasons. First and foremost, there is little or no mobile robotics research content in redesigning the vehicle itself. Secondly, both the dollar and time costs were expected to be high due to the mismatch between the organization and the task - the university is not organized to perform custom automotive redesign. Costs for such a venture are expected to be high due to the steep learning curve and possible required capital equipment investments.

56.2 A_p and A_w

Changes to the engine power draw were ruled out because the original configuration drew no power from the vehicle engine, except for three small actuators. The power draw of these actuators is minimized for the application and practically insignificant to the engine. Coulter and Mueller[6] identified in their configuration study likely sources of weight reduction in the autonomy hardware. As a result of process of elimination (there was no less costly alternative) and due to the high likelihood of a performance improvement, A_w was chosen as the single component for redesign study.¹

56.3 Coefficients of Performance

The coefficient of power performance is calculated as follows:

$$\beta = \frac{A_p^{Loc} + A_p^{Aux}}{M_p + A_p^{Aux}} = \frac{0 + 4100W}{111,900W + 4100W} = 0.03534$$

This indicates that about 3.5% of the power of the entire vehicle is dedicated to cognition and cognitive support. The coefficient of weight performance is calculated as follows:

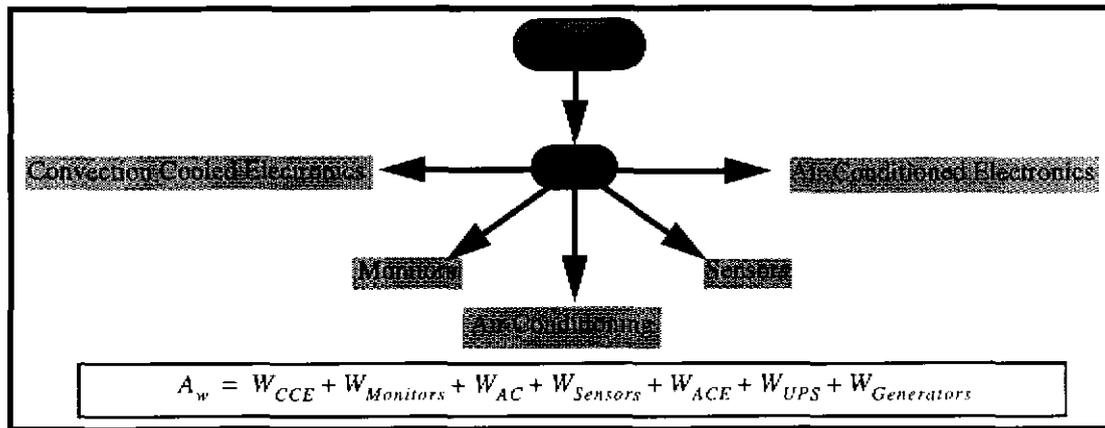
$$\Gamma = \frac{A_w}{C_w} = \frac{2500lbf}{10,200lbf} = 0.2451$$

This indicates that approximately 25% of the total vehicle weight is dedicated to cognitive support.

1. Note that in general, more than one component can be chosen for redesign. The selection of a single component is specific to this application.

56.4 The A_w Cost Function

Consider the following schematic of the autonomy hardware, representing the flow of electrical power from the generator, through the UPS, to the electrical components. Autonomy components are subdivided into five primary groups. The A_w cost function is written as the summation of the weight cost functions of each of these five primary groups.



(Coulter[5]) demonstrates that several of these weight functions are floor functions and nested floor functions, whose parameters are electronics devices requiring active cooling. The imposition of a thermal barrier requires additional power; inefficient barriers require even more power, unnecessarily sizing up power production and distribution components. Function backchaining indicated that small changes in thermodynamic efficiency could lead to large changes in system weight.

56.5 Reconfiguration

The study by (Coulter and Mueller[6]) provided for systematic weight reduction. A reduction in weight (310 pounds) was made by removing a monitor and an unneeded passenger work station. Another reduction (390 pounds) was made by replacing a large inertial pan and tilt platform with a smaller pan and tilt platform. The remainder of the weight loss came from a reconfiguration of the power generation and distribution system through the redistribution of components requiring active cooling, and through improvements to the thermal barrier separating the cold compartment from the rest of the vehicle (Coulter[5]).

The following table illustrates the power and weight savings breakdown:

Current Design Weight and Power Savings by Sub-System

Sub-System	Weight Loss	Percentage of Total Weight Loss	Power Draw Reduction	Percentage of Total Power Draw Reduction
Air Conditioning	100 lbf	7.14%	600 W	63.0%
Power	600 lbf	42.86%	~0	0
Staget	390 lbf	27.86%	336 W	35.0%
Structure	310 lbf	22.14%	0	0
Totals	1400 lbf	100%	936 W	98.0%

The effects of the reconfiguration are summarized below. *Aut. Power* stands for the power draw of the on-board autonomy hardware. *Aut. weight* is the weight of the autonomy hardware. *Vehicle power* is the power rating of the vehicle's engine. *Weight* is the total system weight.

Current Design Power and Weight Change

	Original Total	Modified Total	Percent Change
Aut. Power	4100 W	3150 W	-23.17%
Aut. Weight	2500 lbf	1100 lbf	-56.0%
Vehicle Power	150 HP	150 HP	0%
Weight	10,200 lbf	8800 lbf	-13.73%
Π	0.015	0.017	+12.0%

57. CONCEPT CONFIGURATION DESIGNS

The reconfiguration design that was performed on the NavLab II has, inherently, a number of assumptions about the mission of the vehicle, and its support requirements. Broadly stated, the NavLab series of mobile robots were designed as software development tools - a mobile NAVigational LABoratory for mobile robotics research. As such, this required the emplacement of seats, monitors, computing stations and other equipment for the researcher's use. The current design still reflects the philosophy of supporting research at the expense of performance. In this

section, we will consider a series of conceptual configurations that undo some of the major assumptions in this philosophy. The performance cost of each assumption can then be calculated in several ways. The cost of the assumption in terms of the change in Π from the original vehicle will be considered, as well as the change in Π from the last concept configuration in the series.

57.1 Concept #1: Off-load Passengers and Data Displays

Consider removing from the NavLab II all of the equipment that supports the human researcher's software development needs. Such equipment includes monitors, keyboards, seats, and non-essential video and electronics equipment. The following weight and power draw reductions, and changes in Π ensue:

Component Power and Weight Reductions

	Weight	Power
Original	1100 lbf	3150 W
Monitors	-157 lbf	-371 W
Seats and Structure	-150 lbf	-0
Video Switchers & VCR	-46 lbf	-75 W
Resize UPS	-50 lbf	-0
Resize Generators	-30 lbf	-0
Total	667 lbf	2704 W

Power and Weight Change - Concept #1

	Modified Total	Concept #1 Total	Percent Change
Aut. Power	3150 W	2704 W	-14.16%
Aut. Weight	1100 lbf	667 lbf	-39.36%
Vehicle Power	150 HP	150 HP	0%
Weight	8800 lbf	8367 lbf	-4.92%
Π	0.0170	0.0179	+5.3%

From Table 5, it is seen that a performance increase of 5% over the new (here called the modified) configuration results directly from removing the human from the environment.

57.2 Concept #2 Eliminate Air Conditioning

The current configuration does not require air conditioning to cool the computing, if adequate air circulation is maintained. The computing is capable of running on 100° days, as long as the heat is properly ducted away from the electronics. Air conditioning is required only when the computing environment must be sealed from the ambient environment. This is not necessary, as ambient air could be adequately filtered for cooling purposes. The following tables consider the savings incurred by eliminating the air conditioner.

Air Conditioning Power and Weight Reductions

	Weight	Power
Original	667 lbf	2704 W
Eliminate Air Conditioning	-130 lbf	-1200 W
Resize UPS	-120 lbf	-0
Resize Generators	-68 lbf	-0
Total	349 lbf	1504W

Power and Weight Change - Concept #2

	Modified Total	Concept #2 Total	Percent Change
Aut. Power	3150 W	1504 W	-52.25%
Aut. Weight	1100 lbf	349lbf	-68.27%
Vehicle Power	150 HP	150 HP	0%
Weight	8800 lbf	8049 lbf	-8.53%
Π	0.0170	0.0186	+9.41%

The new savings over the modified design is drastic. Over 50% of the power draw, and nearly 70% of the configuration weight is comprised of air conditioning and human support equipment that was eliminated in the previous concept design.

57.3 Concept #3: Production Computing and Sensing

The sensing equipment aboard the NavLab II is research grade, which means that it is larger, heavier and most power hungry than would be seen on a production line system. The following power and weight improvements are based on the latest, compatible laser scanning devices and video cameras.

Production Computing and Sensing Configuration

	Power	Weight
Computing	290 W	28 lbf
Sensing	126 W	20 lbf
UPS	0	50 lbf
Generator	0	51
Total	416 W	149 lbf

Power and Weight Change - Concept #3

	Modified Total	Concept #3 Total	Percent Change
Aut. Power	3150 W	416 W	-86.79%
Aut. Weight	1100 lbf	149 lbf	-86.45%
Vehicle Power	150 HP	150 HP	0%
Weight	8800 lbf	7849 lbf	-10.81%
Π	0.0170	0.0191	+12.35%

57.4 Summary of Concepts

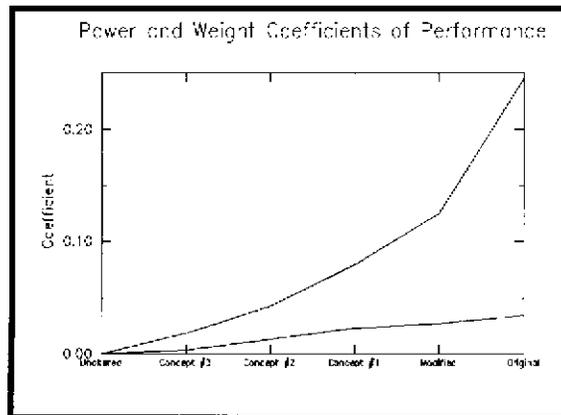
The following table compares the six design variation across a number of parameters. Π , $\Delta\Pi$ and the percent change in Π from the original state of the vehicle are all given. In addition, the power performance coefficient β , and the weight performance coefficient Γ , are all calculated for each design.

Design Performance Comparison

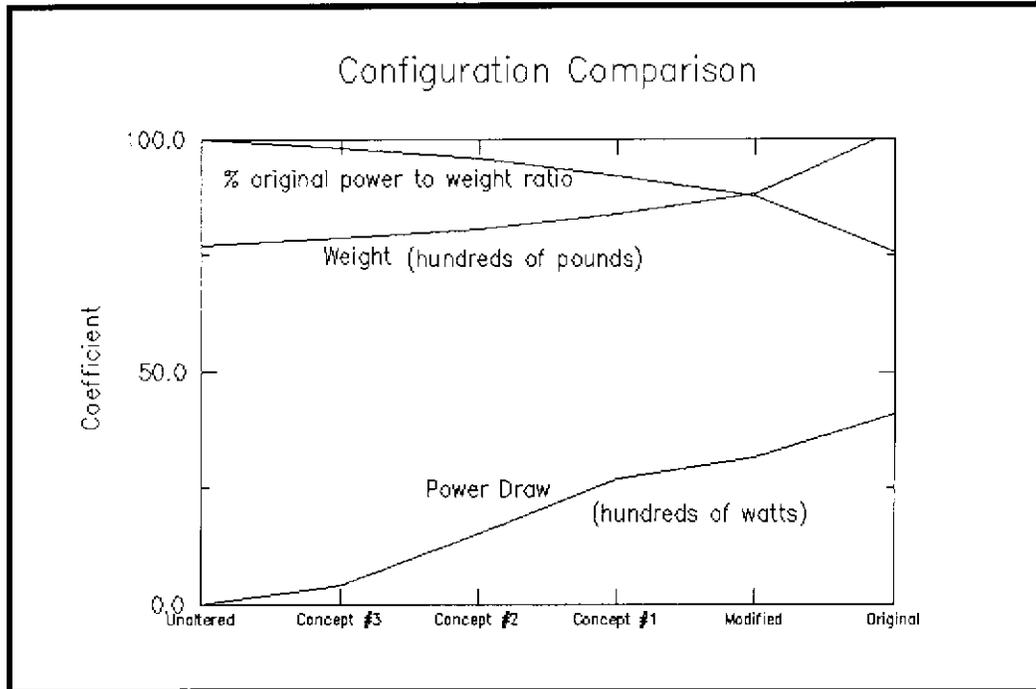
	Unaltered Vehicle	Concept #3	Concept #2	Concept #1	Modified Design	Original Design
Aut. Power	0	416 W	1504 W	2704 W	3150 W	4100 W
Aut. Weight	0	149 lbf	349 lbf	667 lbf	1100 lbf	2500 lbf
Vehicle Power	150 HP	150 HP	150 HP	150 HP	150 HP	150 HP
Total Weight	7700 lbf	7849 lbf	8049 lbf	8367 lbf	8800 lbf	10,200 lbf
Π	0.01948	0.01911	0.01864	0.01793	0.01705	0.01471
$\Delta\Pi_{\text{last}}$	0	-0.00037	-0.00047	-0.00071	-0.00088	-0.00234
$\%\Delta\Pi_{\text{org}}$	100%	98.10%	95.69%	92.04%	87.53%	75.51%
β	0	0.00370	0.01326	0.02359	0.02738	0.03534
Γ	0	0.01898	0.04336	0.07972	0.12500	0.24510

57.4.1 Graphical Comparisons

The relationships among power, weight and system performance are more readily seen in the following graphs. The β and Γ plots below show that the weight increase does not follow that same rate as the power increase. Notice that the slopes of the curves do not correlate. Sharp increases in weight coefficient do not necessarily follow from sharp increases in power coefficient.



The following graphs compare the system performance, Π against the total system weight and the total system power draw. Note that the performance drop follows precisely from the weight increase¹, however neither Π nor the system weight follow the change in power draw.



57.4.2 Comments

These graphs illustrate that, in this case, the relationship between the system weight and the system power is clearly non-linear. Changes in weight increase in a manner prescribed by the floor functions. Large weight changes correlate to resizing of power components and the elimination of supporting structure. The largest change, between the modified and original designs, corresponds to the doubling of the number of required power components, because physical upsizing of a single component was no longer practical.

1. This is a special case, as no power is being drawn from the engine, the power to weight change is only dependent upon the system weight change.

58. REFERENCES

- [1] Amidi, O. "Integrated Mobile Robot Control.", Masters Thesis, Dept. Of Electrical and Computer Engineering, CMU, May, 1990.
- [2] Mike Blackwell, "The Uranus Mobile Robot" Robotics Institute Technical Report, CMU-RI-TR-91-06.
- [3] D. Cho and J.K. Hedrick, "Automotive Powertrain Modelling for Control," Transactions of the ASME, Vol. 111, December 1989.
- [4] R. Craig Coulter, "Path Admissibility based on the Physical Mechanics of Vehicle / Terrain Interaction," Master's Thesis, Department of Mechanical Engineering, Carnegie Mellon University, January, 1992.
- [5] R. Craig Coulter, "A Systems Engineering Approach to Electro-Mechanical Reconfiguration for High Mobility Autonomy," Robotics Institute Technical Report CMU-RI-TR-94-27.
- [6] R. Craig Coulter and George Mueller, "A Reconfiguration Study of the NavLab II Mobile Robot," Robotics Institute Technical Report CMU-RI-TR-94-01.
- [7] Daily et al. "Autonomous Cross-Country Navigation with the ALV," Proceedings of the 1988 IEEE International Conference on Robotics and Automation.
- [8] E.D. Dickmanns and V. Graefe, "Dynamic Monocular Machine Vision," Machine Vision and Applications, Vol. 1, 1988.
- [9] E.D. Dickmanns and V. Graefe, "Applications of Dynamic Monocular Machine Vision," Machine Vision and Applications, Vol. 1, 1988.
- [10] Dickmanns and Zapp "A Curvature-based Scheme for Improving Road Vehicle Guidance by Computer Vision," SPIE's Cambridge Symposium on Optical and Optoelectric Engineering, October 1986.
- [11] Kevin Dowling, et al. "NAVLAB: An Autonomous Navigation Testbed," Robotics Institute Technical Report CMU-RI-TR-87-24.
- [12] Alberto Elfes, "Occupancy Grids: A Stochastic Representation for Active Robot Perception," Proceedings of the Sixth Conference on Uncertainty in AI, July 1990.
- [13] Dai Feng and Bruce Krogh, "Satisficing Feedback Strategies for Local Navigation of Autonomous Mobile Robots," IEEE Transactions on Systems, Man and Cybernetics, Vol. 20, No. 6, November / December 1990.
- [14] J.S. Freeman et al. "The Iowa Driving Simulator: An Implementation and Application Overview," SAE Vehicle Computer Applications - Vehicle Systems and Driving Simulation; SAE SP-1080; International Congress and Exposition, Detroit, MI, Feb. / Mar. 1995.
- [15] Thomas Gillespie, "Fundamentals of Vehicle Dynamics" SAE Publications, 1992.
- [16] Alonzo Kelly, "An Intelligent Predictive Controller for Autonomous Vehicles," Robotics Institute Technical Report CMU-RI-TR-94-20.
- [17] Alonzo Kelly "A Feedforward Control Approach to the Local Navigation Problem for Autonomous Vehicles", CMU-RI-TR-94-17, Robotics Institute, CMU May 2, 1994.
- [18] Sanjiv Singh and Alonzo Kelly, "Robot Planning in the Space of Feasible Actions: Two Examples," Proc. Int'l Conference on Robotics and Automation, April 1996, Minneapolis.
- [19] Dirk Langer, Julio Rosenblatt and Martial Hebert, "A Behavior Based System for Off-Road Navigation," IEEE Transactions on Robotics and Automation, December 1994, Vol. 10, No. 12.
- [20] Rudolf Limpert, "Brake Design and Safety," SAE, Inc. ISBN 1-56091-261-8, 1992.
- [21] Hans P. Moravec, "The Stanford Cart and the CMU Rover," Proceedings of the IEEE, Vol 71, No. 7, July 1983.

-
- [22] P.F. Muir and C.P. Neuman, "Kinematic Modelling of Wheeled Mobile Robots," *Journal of Robotic Systems*, Vol. 4, No. 2, 1987
- [23] P.F. Muir and C.P. Neuman, "Dynamic Modeling of Multibody Robotic Mechanisms: Incorporating Closed-Chains, Friction, Higher-Pair Joints, and Unactuated and Unsensed Joints," *Proceedings of the IEEE International Conference on Robotics and Automation*, 1988.
- [24] N.J. Nilsson, "A Mobile Automaton: An Application of artificial Intelligence Techniques," *Proceedings of the International Joint Conference on Artificial Intelligence*, 1969.
- [25] David W. Payton, J. Kenneth Rosenblatt, David M. Keirsey, "Plan Guided Reaction," *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. 20, No. 6, November / December 1990.
- [26] Greg Podnar et al. "A Functional Vehicle for Autonomous Mobile Robot Research," *Robotics Institute Technical Report CMU-RI-TR-84-28*.
- [27] Dean A. Pomerleau, "Neural Network Perception for Mobile Robot Guidance," *Doctoral Thesis*, issued as *Computer Science Technical Report CMU-CS-92-115*
- [28] Dong Hun Shin "High Performance Tracking of Explicit Paths by Roadworthy Mobile Robots," *Doctoral Thesis*, Department of Civil Engineering, Carnegie Mellon University, May 1990.
- [29] Singh et al. "A System for Fast Navigation of Autonomous Vehicles," *Robotics Institute Technical Report*, CMU-RI-TR-91-20.
- [30] Anthony Stentz "The NAVLAB System for Mobile Robot Navigation.," *Doctoral Thesis*, Computer Science Department, CMU, 1990
- [31] Anthony Stentz et al. "An Autonomous System for Cross-Country Navigation," *Proceedings of SPIE Symposium on Mobile Robots*, 1992.
- [32] Anthony Stentz "Optimal and Efficient Path Panning for Unknown and Dynamic Environments," *Robotics Institute Technical Report*, CMU-RI-TR-93-20.
- [33] Charles Thorpe, Martial Hebert, Takeo Kanade, Steven Shafer "Vision and Navigation for the Carnegie-Mellon NavLab," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 10, No. 3, May 1988.
- [34] Whittaker, W. "Terragator - Terrestrial Navigator" *Robotics Institute Technical Report*, Carnegie Mellon 1984.
- [35] World Wide Web Site: <http://mrl.frc.ri.cmu.edu:8001/usr/nivek/html/MRL/mrl.shtml>. This web site contains an on-line history of the Mobile Robot Laboratory at Carnegie Mellon University.
-