

A Partial Analysis of the High Speed Autonomous Navigation Problem

Alonzo Kelly

CMU-RI-TR-94-16

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

May 2, 1994

©1994 Carnegie Mellon University

This research was sponsored by ARPA under contracts “Perception for Outdoor Navigation” (contract number DACA76-89-C-0014, monitored by the US Army Topographic Engineering Center) and “Unmanned Ground Vehicle System” (contract number DAAE07-90-C-RO59, monitored by TACOM).

The views and conclusions expressed in this document are those of the author and should not be interpreted as representing the official policies, either express or implied, of the US government.

Abstract

This report examines the general problem of high speed autonomous navigation from range image data as it applies to both stereo and lidar sensing systems.

In order to intelligently guarantee its own safety, a high speed vehicle must be able to resolve the smallest obstacle that can present a hazard, process sensory data at a rate commensurate with its speed, respond fast enough to avoid obstacles, and maintain a sufficiently accurate model of the world to enable it to make correct decisions.

These dimensions of the problem are analysed in a nondimensional manner and the implications of satisfying all requirements simultaneously are investigated. In this analysis, it is shown that to adopt a policy of guaranteed vehicle safety is to adopt a computational complexity of $O([TV]^N)$ for range image processing where T is the vehicle reaction time and V is the velocity.

This result implies that increased vehicle speed will require nonlinear growth in computational bandwidth. Further, it identifies the fundamental tradeoff of finite computing resources as one of speed for either resolution or reliability.

The conclusions of this report are the theoretical justification for the adaptive, real-time controller design of the RANGER cross country navigator.

List of Figures

Figure 1 - Response Ratio	8
Figure 2 - Throughput Ratio	10
Figure 3 - Acuity Ratios	12
Figure 4 - Fidelity Ratios	14
Figure 5 - Nondimensional Vehicle	16
Figure 6 - System Control Loops	18
Figure 7 - Empirical Braking Distance	21
Figure 8 - Stopping Distance	22
Figure 9 - Steering Limits	26
Figure 10 - Stopping Region	27
Figure 11 - Impulse Turning Trajectories	28
Figure 12 - Impulse Turning Distance	29
Figure 13 - Reverse Turning Trajectories	30
Figure 14 - Lookahead Zones	34
Figure 15 - Planner Lookahead	35
Figure 16 - Vertical Field of View	37
Figure 17 - Minimum Range	40
Figure 18 - Maximum Range	41
Figure 19 - Horizontal Field of View	43
Figure 20 - Tunnel Vision Problem	45
Figure 21 - Vertical Field of View	46
Figure 22 - Occlusions	46
Figure 23 - Nondimensional Vertical Field of View	47
Figure 24 - Stabilization Problem	49
Figure 25 - Hill Occlusion	51
Figure 26 - Hole Occlusion	52
Figure 27 -Lateral Occlusion Problem	53
Figure 28 - Practical Sampling	56
Figure 29 - Perception Ratio	57
Figure 30 - Differential Imaging Kinematics	59
Figure 31 - Typical Scanning Pattern	68
Figure 32 - Typical Spot Pattern	68
Figure 33 - Scanning Density	69
Figure 34 - Scanning Density	70
Figure 35 - Imaging Density	71
Figure 36 - Geometric Efficiency	72
Figure 37 - Latency Problem	76
Figure 38 - Transience in the Reverse Turn	78
Figure 39 - Transient Steering Response	79
Figure 40 - Feedforward	81
Figure 41 - Bicycle Model	87
Figure 42 - Constant Speed Reverse Turn	89
Figure 43 - Constant Speed Partial Reverse Turn	90

Figure 44 - Ackerman Steer Configuration Space	91
Figure 45 - Throughput Problem	93
Figure 46 - Throughput for Constant Flux, Constant Scan	97
Figure 47 - Throughput for Adaptive Sweep, Constant Scan	100
Figure 48 - Throughput for Adaptive Sweep, Adaptive Scan	102
Figure 49 - Area Consumption	103
Figure 50 - Throughput for Adaptive Sweep, Uniform Scan	104
Figure 51 - Throughput for All Algorithms	105
Figure 52 - Computational Spiral Effect	107
Figure 53 - Small Incidence Angle Assumption	112

List of Symbols

a	acceleration	α	steer angle
\bar{b}, \bar{b}_d	(dynamic) braking coefficient	$\dot{\alpha}$	steer angle rate
\bar{b}_k	kinematic braking coefficient	δ	map resolution
c	undercarriage clearance	Δ	obstacle spacing
f_{cells}	map cell throughput	η_S	perceptual software efficiency
f_{pixels}	sensor throughput	η_G	geometric efficiency
f_{images}	frame rate	κ	curvature
f_{cpu}	CPU speed	μ	coefficient of friction
f_{comm}	communications bandwidth	ν	coefficient of lateral acceleration
g	acceleration due to gravity	Ψ	pixel azimuth angle, vehicle yaw
h	sensor height	$\dot{\Psi}$	vehicle yaw rate
\bar{h}	perception ratio	Ψ_L	planner lookahead angle
L	vehicle wheelbase	ρ	radius of curvature
M	vehicle mass	ρ_K	kinematic steering limit
r	wheel radius	ρ_D	dynamic steering limit
R	range	ρ_{cyc}	throughput ratio
R_L	planner lookahead distance	ρ_{dx}	fidelity ratio
ΔR_L	incremental lookahead distance	ρ_{dy}	minimum acuity ratio
s	arc length	ρ_{dz}	maximum acuity ratio
s_B	stopping distance	σ	area density
s_T	turning stopping distance	σ_I	imaging density
s_{IT}	impulse turning distance	σ_S	scanning density
\bar{t}_k, \bar{t}	(kinematic) turning coefficient	τ_B	braking reaction time
\bar{t}_d	dynamic turning coefficient	τ_T	turning reaction time
T	time	θ	beam elevation angle
T_{lat}	frame buffer latency	$\dot{\theta}$	vertical sweep rate
T_{cyc}	software cycle time	Ψ	sensor flux
T_{act}	actuator delay		
T_{react}	system reaction time		
T_{images}	frame period		
V	vehicle speed		
x	crossrange coordinate		
y	downrange coordinate		
z	vertical coordinate		

Table of Contents

1	Introduction	1
1.1	Commentary	1
1.2	Acknowledgments	2
1.3	The Nature of High Speed Autonomy	3
1.4	Structural Problems of High Speed Autonomy	4
1.5	Summary of Conclusions	5
PART I: Elements		6
<hr/>		
1	Guaranteed Safety	6
1.1	The Four Dimensions	7
1.2	Safety	7
1.2.1	Requirement - Guaranteed Safety	7
1.2.2	Mitigating Assumption - Safe Terrain Assumption	7
1.3	Response	8
1.3.1	Requirement - Guaranteed Response	8
1.3.2	Nondimensional Requirement - Response Ratio	8
1.3.3	Design Rules - Response Adapted Lookahead and Speed	9
1.3.4	Algorithmic Solutions - Response Adaptive Lookahead and Speed	9
1.3.5	Mitigating Assumptions - Low Latency Assumption, Wide Depth of Field Assumption	9
1.3.6	Related Subproblems - Latency Problem and Myopia Problem	9
1.4	Throughput	10
1.4.1	Requirement - Guaranteed Throughput	10
1.4.2	Nondimensional Requirement - Throughput Ratio	10
1.4.3	Design Rules - Throughput Adapted Sweep and Speed	11
1.4.4	Algorithmic Solutions - Throughput Adaptive Sweep and Speed	11
1.4.5	Mitigating Assumption - High Throughput Assumption	11
1.4.6	Related Subproblems - Stabilization Problem, Tunnel Vision Problem	11
1.5	Acuity	12
1.5.1	Requirement - Guaranteed Detection	12
1.5.2	Nondimensional Requirement - Acuity Ratios	12
1.5.3	Design Rules - Acuity Rules	12
1.5.4	Algorithmic Solutions - Acuity Adaptive Scan and Planning	13

1.5.5 Mitigating Assumptions - Uniform Scan and Terrain Smoothness	13
1.5.6 Related Subproblems - Sampling Problem	13
1.6 Fidelity	14
1.6.1 Requirement - Guaranteed Localization	14
1.6.2 Nondimensional Requirement - Fidelity Ratios	14
1.6.3 Design Rules - Fidelity Rules	14
1.6.4 Algorithmic Solutions - Fidelity Adaptive Planning	15
1.6.5 Mitigating Assumption - Benign Terrain Assumption, Low Dynamics Assumption	15
1.6.6 Related Subproblems - Sensitivity Problem, Image Registration Problem	15
2 Configuration	16
2.1 Summary of Configuration Nondimensionals	17
PART II: Response	18
1 Reaction Time	18
1.1 System Reaction Time	18
1.2 Braking Reaction Time	19
1.3 Turning Reaction Time	19
1.4 Nondimensional Response	20
2 Maneuverability	21
2.1 Braking	21
2.2 Panic Stop Maneuver	22
2.3 Nondimensional Braking	23
2.4 Braking Regimes	24
2.5 Turning	25
2.6 Steering Limits	26
2.7 Turning Stop Maneuver	27
2.8 Impulse Turn Maneuver	28
2.9 Impulse Turning Distance	29
2.10 Reverse Turn	30
2.11 Nondimensional Turning	31
2.12 Turning Regimes	31

2.13	Turning Stop Maneuver	32
2.14	Impulse Turn Maneuver	33
2.15	Impulse Turning Regimes	33
3	Lookahead	34
3.1	Adaptive Regard	34
3.2	Pointing Rules	35
3.3	Adaptive Lookahead	36
3.4	Nondimensional Lookahead	36
3.5	Nondimensional Pointing Rules	37
3.6	Summary of the Response Nondimensionals	38
PART III: Throughput		39
1	Depth of Field	39
1.1	Minimum Sensor Range	40
1.2	Maximum Sensor Range	41
1.3	Nondimensional Maximum and Minimum Range	42
1.4	Myopia Problem	42
2	Horizontal Field of View (HFOV).....	43
2.1	Nondimensional Horizontal Field of View	44
2.2	Tunnel Vision Problem	45
3	Vertical Field of View (VFOV)	46
3.1	Nondimensional Vertical Field of View	47
4	Sweep Rate	48
4.1	Stabilization Problem	49
4.2	Nondimensional Sweep Rate	50
4.3	Adaptive Sweep	50
5	Occlusion	51
5.1	Hill Occlusion	51

5.2	Hole Occlusion	52
5.3	Occlusion Problem and Unknown Hazard Assumption	53
5.4	Lateral Occlusion	53
6	Perceptual Bandwidth.....	54
6.1	Sensor Flux	54
6.2	Sensor Throughput	54
6.3	Sweep Rate	54
6.4	Processor Load	54
6.5	Perceptual Software Efficiency	55
6.6	Computational Bandwidth	55
6.7	Communications Bandwidth	55
PART IV: Acuity		56
1	Acuity	56
1.1	Sampling Theorem	56
1.2	Terrain Smoothness Assumption	57
1.3	Impact of Imaging Geometry on Acuity	57
1.4	Nomenclature	58
1.5	Sampling Problem	58
1.6	Differential Imaging Kinematics	59
1.7	Pixel Footprint Area and Density Nonuniformity	60
1.8	Pixel Footprint Aspect Ratio	61
1.9	Minimum Sensor Acuity in Image Space	62
1.10	Maximum Sensor Acuity in Image Space	63
1.11	Kinematic Maximum Range and the Myopia Problem	64
1.12	Maximum Angular Resolution	64
1.13	Acuity Problem	64
2	Positioning Bandwidth	65
2.1	Heading and Positioning Bandwidth	65
2.2	Attitude and Positioning Bandwidth	66

2.3	Motion Distortion Problem	67
3	Geometric Efficiency	68
3.1	Imaging Geometry	68
3.2	Scanning Density	69
3.3	Imaging Density	71
3.4	Geometric Efficiency	72
3.5	Adaptive Scan	73
3.6	Acuity Nondimensionals	73
3.7	Summary of the Acuity Nondimensionals	74
PART V: Fidelity		75

1	Delays	75
1.1	Latency Problem	76
1.2	Minimum Significant Delay	76
1.3	Dynamic Systems	77
1.4	Characteristic Times and Low Latency Assumption	77
1.5	Transience in Turning	78
1.6	Heading Response	79
1.7	Nondimensional Transient Turning	80
2	Dynamics Feedforward.....	81
2.1	Dynamics Feedforward	81
2.2	Characteristic Times and Loop Bandwidth	82
2.3	Impact on Planner/Controller Hierarchy	82
2.4	Impact on Trajectory Generator/ Trajectory Tracker Hierarchy	82
3	Positioning Fidelity.....	83
3.1	Absolute Attitude Accuracy Requirement	83
3.2	Rigid Terrain Assumption	83
3.3	Rigid Suspension Assumption	83
3.4	Image Registration Problem	84

3.5	Linear Relative Accuracy Requirement	84
3.6	Angular Relative Accuracy Requirement	84
4	Perceptual Fidelity	85
4.1	Incidence Sensitivity Problem	85
4.2	Attitude Sensitivity Problem	85
PART VI: Results		86
<hr/>		
1	Effect of Response on Fidelity - Rationale for Feedforward Approach	86
1.1	The Bicycle Model	87
1.2	The Fresnel Integrals	88
1.3	Dynamics of the Constant Speed Reverse Turn	89
1.4	The Clothoid Generation Problem	90
1.5	Configuration Space	91
1.6	A Real Time Control View of High Speed Autonomy	92
2	Effect of Acuity and Response on Throughput - Rationale for Adaptive Perception	93
2.1	Throughput Problem	93
2.2	The Illusion	94
2.3	Adaptive Perception	94
2.4	Assumptions of the Analysis	95
2.5	Common Throughput Expression	95
2.6	Basic Mechanism	96
2.7	Complexity of Constant Flux Range Image Processing	97
2.8	Complexity of Adaptive Sweep Range Image Processing	98
2.9	Complexity of Adaptive Sweep, Adaptive Scan Range Image Processing	101
2.10	Complexity of Adaptive Sweep, Uniform Scan Image Processing	103
2.11	The Fundamental Speed/Resolution Trade-off	105
3	Effect of Throughput on Response - Rationale for A Real Time Approach	107
3.1	Computational Spiral Effect	107

4	A Strategy For High Speed Autonomous Navigation	108
4.1	Basic Strategy	108
4.1.1	Real Time Systems Analysis and Design	108
4.1.2	Adaptive Sensors and Adaptive Perception	108
4.1.3	A Deliberative Approach to Autonomy	108
4.1.4	Key Assumptions	108
4.1.5	Vehicle Configuration	108
4.2	Real Time Systems Analysis and Design	109
4.2.1	Hardware Platform	109
4.2.2	Real Time Software	109
4.3	Adaptive Sensors	109
4.4	Adaptive Perception	110
4.5	Deliberative Approach	110
4.6	Key Assumptions	110
4.6.1	Continuity Assumption	110
4.6.2	Terrain Smoothness Assumption	110
4.6.3	Benign Terrain Assumption	111
4.6.4	Obstacle Sparsity Assumption	112
4.7	Small Incidence Angle Assumption	112
4.8	Vehicle Configuration	113
4.8.1	Mechanical Design	113
4.8.2	Sensor Design	113
5	References	114

1. Introduction

At the speeds required of the next generation of autonomous vehicles, it becomes necessary to explicitly address the need for the vehicle to react to hazards in real time. As speeds increase, computing and sensory hardware and software must be re-evaluated to remove the bottlenecks that were acceptable at the lower speeds encountered in earlier research.

Higher speeds require both looking further ahead and reacting faster. Hence, autonomous navigation must mature to the point where significantly more computation is performed in less time, or some new approach must be embraced.

This report presents a rudimentary theory of high speed autonomous navigation. The analyses performed include real-time analysis, dimensional analysis and complexity analysis. The report provides background support for the design of the real-time adaptive controller called RANGER which is described in [16].

1.1 Commentary

Requirements analysis is a matter of basic doctrine in the engineering design method. The goal of the technique is to understand the *nature of the problem* being solved. In this process, one makes every attempt to abstract the problem to a level where the assumptions of specific solutions are avoided. For example, an autonomous vehicle needs to see what is out there and do something about it. These requirements are intrinsic to the problem and require no assumptions about how the environment is perceived or how the vehicle goes about responding. This report attempts to take such a view of autonomous vehicles.

Dimensional analysis is a powerful technique that can be used to aid the designer in forming an understanding of the basic issues in the design of an engineered system. The technique is used in its simplest form by checking the units in an equation. There is, however, much more that can be done with this simple idea.

There exists a theorem which has been used for decades in fluid mechanics, and aeronautical engineering called the **Buckingham Pi Theorem**. The rudiments of this theorem are that any group of equations describing physical variables of interest can *always* be reduced to canonical relationships in a minimum number of nondimensional variables which codify the essence of the dependencies in a scale independent manner. These nondimensional variables, or **Pi products**, distill the dependencies relating the physical quantities involved to the fundamental relationships that explain the deep issues.

An engineering design can be a difficult process to follow because there are so many different variables which interact in complex ways. Use of nondimensionals permits reduction of the dimension of a design space to manageable proportions and, at the same time, ensures consistency of the analysis.

1.2 Acknowledgments

Barry Brumitt and the author investigated the impact of steering delays on the planning and path generation problems respectively prior to this work. Based on this work, and the precedent of FASTNAV, the idea of generalizing steering dynamics feedforward to a complete state space model ultimately emerged. Barry did initial work on the extent of the dynamically feasible subset of C space. The C space planner constructed by Barry Brumitt and Tony Stentz provided the initial impetus for the use of a feedforward simulator.

The notion of relating pixel size to subtended length on a vertical surface has been around for some time on the ARPA UGV program. Its origin is unknown to the author. Wherever it came from, it is the key element which permits an analysis of computational complexity.

R Coulter first investigated the braking performance of the CMU HMMWV in a quantitative manner. This lead naturally to the question of quantifying steering performance.

Omead Amidi was probably the first person to suggest that a tightly coupled real-time implementation of the high level perceive-think-act loop would be ultimately necessary at high speed.

Dong Hun Shin and Sanjiv Singh were the first to recognize the significance of steering delay and the first to implement a feedforward solution on the FASTNAV vehicle some five years ago. Shin implemented the first pure pursuit path tracker for the HMMWV vehicle to be used in off road work. Shin identified the difficulty of the clothoid generation problem with the Fresnel integral. Dean Pomerleau later provided a second precedent for steering feedforward in the ALVINN system.

Many of the ideas contained here can be considered to be a natural evolution of the ideas of Tony Stentz which formed the basis of the first full geometry off road navigator at CMU since the system implemented earlier by Martial Hebert. This software system was implemented by Barry Brumitt, R Coulter, Al Kelly, Bill Burky and George Mueller under the direction of Tony Stentz.

Tony first pointed out the difficulty of treating negative obstacles and started the development of adaptive regard by relating lookahead to the vehicle stopping distance. Tony also pointed out the basis for kinematic requirements on the vertical field of view.

Martial Hebert pointed out the sampling problem, the attitude sensitivity problem, and the significance of the pixel aspect ratio on throughput. Many conversations with Martial suggested the need for adaptive perception.

1.3 The Nature of High Speed Autonomy

It will be shown that the *high speed autonomy problem is a control problem* (as opposed to a planning problem) because the vehicle configuration space is degenerate at high speed and the mapping from configuration space to actuation space is not defined over something like 97% of the extent of configuration space. Search based C space planners based on many fine AI algorithms are brittle and waste resources because the clothoid generation problem is impossible in practical terms and most vehicle configurations in C space are not feasible when vehicle dynamics are considered. Therefore, an optimal system inverts the AI hierarchy and considers path feasibility before obstacle avoidance. In optimization terms, such a system considers the constraints before the utility function, and in doing so, becomes a controller.

It will be shown that *high speed autonomy is a real time control problem* because response time requirements are stringent and throughput and response time depend on each other to an alarming degree. Contemporary sensor technology coupled with the latencies of sluggish massive vehicles and distributed control schemes limit vehicle speeds severely. Computational resource limitations are increasingly severe at even moderate speeds and a real time approach is indicated which minimizes response time and maximizes sensor lookahead. Through the adoption of a managed minimum reaction time strategy, an optimal system becomes a real-time system.

It will be shown that *high speed autonomy is an adaptive autonomy problem* because the computational complexity of range image perception is severe when measured against contemporary general purpose computing hardware and non adaptive image processing techniques are used. Indeed, up to 98% of the information provided by contemporary sensors is redundant. The report introduces a novel adaptive perception algorithm which effectively eliminates the historically significant throughput problem of high speed navigation at contemporary speeds.

It will be shown that *high speed autonomy is a feedforward control problem* because algorithmic stability of path tracking and obstacle avoidance can only be achieved by high fidelity models of vehicle actuator dynamics.

It will be shown that *high speed autonomy is a state space control problem* because the only valid model of a high speed vehicle is a coupled nonlinear multidimensional differential equation. An optimal system models this equation explicitly, and in doing so becomes a state space controller.

1.4 Structural Problems of High Speed Autonomy

The report will identify many structural problems of high speed autonomy which relate to suboptimal design of the vehicle itself. Many of the limitations of contemporary high speed navigators arise from underlying hardware limitations. This can be comically illustrated with an overly exaggerated analogy. Take a typical race car driver. Normally, such people would be expected to have excellent vision, reflexes and an ability to make decisions quickly under pressure. Consider replacing this driver with another less capable one which is roughly analogous to today's autonomous systems.

The substitute driver's vision is blurred to reflect the fact that contemporary sensor pixels are an order of magnitude too large to resolve small, yet significant, obstacles. Replace the driver's excellent human brain with a 10 Mflop processor to reflect the limited throughput of today's computers. Alternately, consider that the substitute driver was not very bright to begin with. Give the half blind driver a little too much to drink in order to slow the reflexes to be equivalent to the reflexes of a distributed computer system with limited power actuators and significant sensor dwell and latency.

Next, in order to reflect the limited field of view of sensors, strap a tube to the drivers face or alternately give him a severe case of tunnel vision. Also, because many sensors have relatively limited depth of field, give the driver a severe case of myopia to go along with his blurred, tunnel vision. Sensors today normally cannot be physically pointed, so it is necessary to strap the drivers head into a fixed orientation, and point the head just over the vehicle hood instead of straight ahead.

Of course, the driver is not placed in his own backyard but is blindfolded, transported, and dropped thousands of miles away with no map and no idea of the gross layout of the place, and in some cases, without so much as a compass to guide him. Our driver also has a limited cognitive ability to formulate a strategic plan to cover or traverse a large area, and coupled with his limited ability to remember where he has been before, is prone to wander aimlessly about.

Finally, place a brick on the gas pedal to achieve constant speed of about 20 mph and demand that the driver drive, not on the road where there is a nicely painted line to follow and no stationary obstacles, but across country on rough terrain where stationary obstacles abound, tipover is a real constant possibility, tires can be easily destroyed, holes and ponds exist which can swallow the whole vehicle, mud and snow and ice can cause traction failure, and many dense forested areas and rocky regions are completely untraversable to begin with.

So by analogy a drunk, stiff-necked, nearsighted driver, who is not very bright, with intrinsically blurred tunnel vision, staring just over the vehicle hood, is trying to drive 20 mph across country with no idea where he is and no idea what is over the next rise. This is the contemporary high speed navigation problem.

It will be shown that the fundamental safety requirement cannot be met with contemporary technology. At first glance, the problem appears impossible but the requirements analysis demonstrates that this is an illusion which is generated from a non systems-oriented approach. Indeed, the problem is not as difficult as it appears if computational resources are managed optimally and this bodes well for research progress in the medium term.

1.5 Summary of Conclusions

The major conclusions of the report are restated here for easy reference.

- The policy of guaranteed vehicle safety implies a computational complexity of $O([TV]^N)$ for range image processing where T is the vehicle reaction time and V is the velocity. This result implies that increased vehicle speed will require nonlinear growth in computational bandwidth. Further, it identifies the fundamental tradeoff of finite computing resources as one of speed for either resolution or reliability.
- Vehicle maneuverability and sensor field of view are intimately related and both change substantially over the speed regimes of current autonomous vehicle research.
- Analogous braking and turning regimes can be defined which identify points where vehicle maneuverability changes in important ways.
- The clothoid generation problem for the high speed Ackerman vehicle is a problem in nonlinear underdetermined differential equations which is impossible to solve in practical terms.
- The planning “state space” of the high speed Ackerman vehicle is degenerate so the AI hierarchical view of the high speed autonomy problem is unsound. The problem is a predictive control problem. A feedforward state space control approach is indicated.
- The low latency assumption is fundamentally wrong at surprisingly moderate speeds and kinematic planners will always be unstable and unreliable above some speed threshold.
- A consistent application of the small incidence angle assumption is a key element in the solution of the throughput problem. The throughput problem can be completely eliminated at contemporary speeds. Adaptive perception techniques which computationally stabilize the vertical field of view provide the best of both worlds. They provide the high throughput necessary for high speed motion and the wide field of view necessary for rough terrain.
- Most contemporary vehicle testbeds have a tunnel vision problem. Horizontal field of view required for rough terrain work exceeds 120 degrees.
- Most contemporary vehicle testbeds have a myopia problem brought about by poor angular resolution. Angular resolution must be increased to about 1 mrad for high speed work.
- Most contemporary vehicle testbeds have a stabilization problem. Guaranteed visibility of terrain requires 60 degrees of vertical field of view.
- A vehicle specifically designed for high speed autonomy would have a very high sensor and very large wheels.
- A sensor specifically designed for high speed rough terrain autonomy would have very wide horizontal field of view and significantly nonsquare pixels.

PART I:Elements

In order to intelligently guarantee its own safety, a high speed vehicle must be able to resolve the smallest obstacle that can present a hazard, process sensory data at a rate commensurate with its speed, respond fast enough to avoid obstacles, and maintain a sufficiently accurate model of the world to enable it to make correct decisions. This section develops these requirements into a simple set of nondimensional requirements, and a related set of assumptions and design rules.

1. Guaranteed Safety

The fundamental requirements of fast cross country navigation are that of *robust, high speed*, navigation over *rough terrain*. Of these, the robustness requirement is considered paramount. It is cast in terms of vehicle safety, and is further analyzed along the four dimensions of **timing**¹, **speed**, **resolution**, and **accuracy**.

There is some speed beyond which the problem of high speed autonomy becomes one of intelligent real time control. In such a problem, a few basic failure modes are typically encountered:

- the system cannot react fast enough to respond to environmental events
- the system cannot process all of the data presented to it in the time allotted
- the system cannot recognize environmental events when they occur
- the system cannot decide on the proper course of action for a specific event

These four requirements are intrinsic to autonomy at any speed on any type of terrain but they are more difficult to meet at high speed on rough terrain.

1. Bolded nonitalic text contains keywords that repeat in the document and appear in the index.

1.1 The Four Dimensions

Thus, there are four aspects to the problem of ensuring that a vehicle remains safe. The first is a question of **reaction time** (timing), of whether the system is fast enough to react once it perceives a hazard. The second is a question of **throughput** (speed), of whether the system is fast enough to perceive everything needed. These questions are conceptually independent but become related when considered along with resolution. *These first two dimensions limit the speed of a system.*

It is possible to construct a system with excellent reaction time that cannot supply the throughput necessary by simply reducing the field of view of a sensor to one, or a few pixels. Such a system could respond provided it happened to see a hazard. Conversely, it is possible to design a system with sufficient throughput, but poor reaction time by processing range data out to the horizon. Such a system could see, but would be unable to respond. *There is a spectrum of how often to see how much*, and there is a single point on this spectrum which optimizes performance against both requirements.

The third is a question of **acuity**² (resolution), of whether the system has high enough resolution perceptual equipment to enable it recognize important aspects of the environment. The fourth is a question of **fidelity** (accuracy), of whether the system has sufficiently accurate models to permit it to make correct decisions. *These last two dimensions limit the reliability of a system.*

1.2 Safety

1.2.1 Requirement - Guaranteed Safety

Together, response, throughput, acuity and fidelity requirements must be continuously met in order to guarantee vehicle safety. This will be called the policy of **guaranteed safety**. An adaptive system can directly implement the policy of guaranteed safety by reasoning in real time about the four dimensions of safety discussed above and by adapting its perception and planning subsystems to comply directly with the need for safety.

1.2.2 Mitigating Assumption - Safe Terrain Assumption

Although it may be obvious, it is important to mention that many systems can profitably reject the guaranteed safety requirement in favor of an implicit assumption that the terrain is safe. This is the **safe terrain assumption**. Human automobile drivers routinely make this assumption when driving on highways, and they do so both out of necessity and with some success. There are often situations when a driver would be unable to react to a stationary obstacle which suddenly appeared around a bend in the highway. This key assumption permits high speed driving.

2. The report is replete with invented expressions used as a brevity device. Consider these expressions to be an aspect of notation like a mathematical symbol or an acronym.

1.3 Response


A vehicle must be able to react to an obstacle that it sees before it is reached. Therefore, the relationship between the distance that the vehicle looks ahead R_L , the reaction time T_{react} , and the speed V is central.

1.3.1 Requirement - Guaranteed Response

Following real-time systems terminology, the policy of guaranteeing a timely response to important sensory input events will be called the policy of **guaranteed response**. The problem of achieving guaranteed response will be called the **response problem**.

1.3.2 Nondimensional Requirement - Response Ratio

If it takes T_{react} seconds to react, then the **response ratio**³ is the ratio of the distance travelled during the reaction time to the distance at which the obstacle is perceived:



To say that a vehicle must avoid obstacles is equivalent to saying that the response ratio must never exceed one. If the highest possible practical speed has been achieved, the only way to improve on this (while maintaining the ratio below 1) is to either decrease the reaction time or increase the lookahead distance. For a given vehicle with fixed latencies, and a given sensor with fixed maximum range, there exists some speed that cannot be safely exceeded without risking an encounter with hazards that were *seen too late to avoid them*. Adaptive systems can deliberately increase lookahead or reduce speed on the basis of the response ratio.

There are many related issues. There are practical limits on decreasing the reaction time due to the physical response of the vehicle given power limitations on actuators, and the limited speed of the data processing hardware. There are also practical limits on increasing the lookahead because the maximum useful range of a sensor is often limited by image occlusions, limited accuracy, limited resolution, power limitations, or safety considerations, and most of these are extremely aggravated by shallow pixel incidence angles. Different reaction times apply to different obstacle avoidance maneuvers. Turning typically requires more time than braking, for instance. It is the response ratio which is the central concern, not the reaction time itself. A vehicle with half the speed which also responds half as quickly is equivalent. None of this matters until the sensor lookahead is considered.

3. The ratio of a velocity-time product to a distance is a central nondimensional variable in this kind of analysis.

1.3.3 Design Rules - Response Adapted Lookahead and Speed

Any element of the response ratio can be considered to be absolutely limited by the other two. The **response adapted lookahead rule** expresses how the planner lookahead must adapt to the state of the vehicle (speed, curvature) and its ability to respond (braking, turning) in order to guarantee that the response ratio remains less than unity. It can be written as:

$$R_L \geq VT_{\text{react}}$$

Notice that the product VT_{react} is a kind of characteristic vehicle distance which encodes the ability to respond expressed as a distance. An autonomous system must always scan for hazards *beyond* this characteristic distance.

The **response adapted speed rule** expresses how the vehicle speed must adapt to the sensor range and the ability of the vehicle to respond. It can be written as:

$$V \leq \frac{R_L}{T_{\text{react}}}$$

The ratio R_L/T_{react} is a kind of velocity which encodes the “speed” at which hazards present themselves to the vehicle. A more in-depth reaction time analysis reveals that various system latencies, including vehicle dynamics, cause reaction times that can be larger than might be expected, and therefore, maximum speeds are lower than might be expected.

1.3.4 Algorithmic Solutions - Response Adaptive Lookahead and Speed

The idea of **response adaptive lookahead** is to always ensure that the vehicle has time to react to any obstacle it may encounter at the current speed. The idea of **response adaptive speed** is to ensure that the vehicle speed always remains below the critical speed determined by the sensor maximum range. These measures are important because latencies are large and uncontrollable and speed cannot be changed instantaneously and is influenced to a great extent by the slope of the terrain.

1.3.5 Mitigating Assumptions - Low Latency Assumption, Wide Depth of Field Assumption

It may be possible under certain circumstances to simply ignore the issue of whether or not the system can respond quickly enough to avoid hazardous situations without explicitly considering it. Of course, this amounts to an assumption that response is instantaneous relative to any particular situation. This is the **low latency assumption**. An assumption with equivalent consequences is the assumption that sensor useful range is sufficiently large. This could be called a **wide depth of field assumption**. Vehicles which execute start-stop motions and slow speed vehicles can normally make these assumptions.

1.3.6 Related Subproblems - Latency Problem and Myopia Problem

The most central concern of guaranteed response is often the overall latency of the system. A **latency problem** exists when system latencies are too large for any particular goal speed and sensor maximum range. If the sensor maximum range is too small compared to the latencies and the speed, then a **myopia problem** exists.

1.4 Throughput


There is also a basic requirement that the system acquire geometric information as fast as it is consumed by driving over it. Therefore, the relationship between the projection of an image on the ground plane ΔR , the system cycle time T_{cyc} , and the speed V is also central⁴.

1.4.1 Requirement - Guaranteed Throughput

The policy of guaranteeing sufficient throughput will be called the policy of **guaranteed throughput**. The problem of achieving guaranteed throughput will be called the **throughput problem**.

1.4.2 Nondimensional Requirement - Throughput Ratio

The **throughput ratio** is the ratio of the distance travelled during the cycle time to the projection of an image on the groundplane:



To say that a vehicle must not drive over unknown terrain is equivalent to saying that the throughput ratio must never exceed one. If the highest possible practical speed has been achieved, the only way to improve on this (while maintaining the ratio below 1) is to either decrease the cycle time or increase the image projection. For a given vehicle with fixed sensor and computational throughput, there exists some speed that cannot be safely exceeded without risking an *encounter with unseen hazards*. Adaptive systems can deliberately reduce speed or increase the sensor field of view on the basis of the throughput ratio.

There are many related issues. There are practical limits on decreasing the cycle time due to the computer speed and the sensor throughput. There are also practical limits on increasing the image projection due to pitching of the vehicle and the finite angular field of view of the sensor. Also, the maximum range is limited by many concerns, and the minimum range is limited by the height of the sensor and the extension of the vehicle nose in front of it. It is the throughput ratio which is the central concern, not the throughput itself. A sensor with half the field of view which generates twice the frame rate is equivalent. None of this matters until the velocity is considered.

4. Notice that response determines the minimum sensor range whereas throughput determines the vertical field of view. The two together specify a focus of attention which is the real issue at any point in time. This idea will be central later. Notice also that response determines reaction time whereas throughput determines cycle time. These two dimensions of real-time analysis are almost always important.

1.4.3 Design Rules - Throughput Adapted Sweep and Speed

Any element of the throughput ratio can be considered to be absolutely limited by the other two. The **throughput adapted sweep** rule expresses how the sensor field of view must be adapted based on the system cycle time and the vehicle speed. It can be written as:

$$\Delta R \geq VT_{cyc}$$

The product VT_{cyc} is a kind of characteristic vehicle distance, the distance travelled per cycle, which encodes the basic throughput necessary. The **throughput adapted speed** rule expresses how the vehicle speed must be adapted based on the sensor field of view and the system cycle time. It can be expressed as follows:

$$V \leq \Delta R / T_{cyc}$$

The ratio $\Delta R / T_{cyc}$ is a characteristic speed which encodes the throughput necessary in terms of geometry per second. A more in-depth throughput analysis reveals that the efficiency with which traditional sensors generate geometry is unnecessarily low and that unprecedented throughput is possible by simply modifying the sensor geometry and optics without increasing its fundamental throughput⁵.

1.4.4 Algorithmic Solutions - Throughput Adaptive Sweep and Speed

The idea of **throughput adaptive sweep** is to always ensure that the vehicle acquires new environmental information as fast as it is consumed by driving over it. The idea of **throughput adaptive speed** is to ensure that the vehicle speed remains below a critical speed given by the cycle time and the sensor field of view. These measures are important because they can allow a system to achieve unprecedented speeds by computing only the minimum amount of information necessary to ensure safety.

1.4.5 Mitigating Assumption - High Throughput Assumption

It may be possible under certain circumstances to simply ignore the issue of whether or not the system can measure the environment fast enough. This amounts to an assumption that the computers are fast enough to process everything in an image without significantly affecting safety. This is the **high throughput assumption**.

1.4.6 Related Subproblems - Stabilization Problem, Tunnel Vision Problem

On rough terrain, it is possible that either the shape of the terrain in the image or the shape of the terrain upon which the vehicle moves will cause rapid motion of the sensor vertical sweep unless something is done about it. A **stabilization problem** exists when this motion can cause holes between images. Further, another aspect of sensor requirements is that they image all reachable terrain. When the field of view is too narrow to achieve this, a **tunnel vision problem** exists.

5. This is important for two reasons, laser rangefinder throughput is limited fundamentally by the ability of the laser diode to shed generated heat coupled with the need to maintain reasonable signal to noise ratios. Stereo triangulation throughput is limited by the processing speed of the computer used. In both cases, there is much that can be achieved through judicious design of the sensor.

1.5 Acuity


There is a basic requirement that the system be able to resolve the smallest obstacle that can present a hazard to the vehicle at operating velocity. This resolution requirement applies to both the sensor and the computations performed on the data. Clearly, the larger the vehicle, the larger the hazard necessary to challenge it, so it is to be expected that the resolution requirement will be dependent on the vehicle dimensions.

1.5.1 Requirement - Guaranteed Detection

In order to guarantee detection of hazards, some acuity requirement must be maintained at all times. This will be called the policy of **guaranteed detection**. The problem of maintaining adequate acuity will be called the **acuity problem**.

1.5.2 Nondimensional Requirement - Acuity Ratios

Considering the projection of a single range pixel onto the horizontal and onto a vertical surface, two very useful expressions can be formed:



called the **acuity ratios**. To say that a vehicle must resolve hazards reliably is equivalent to saying that the acuity ratios do not exceed one half, as will be shown below.

1.5.3 Design Rules - Acuity Rules

In order to ensure that the vehicle pitch can be computed at all, Nyquist's sampling theorem tells us that the relevant acuity ratio must not exceed one half. This can be written as:

$$dy \leq \frac{L}{2}$$

This will be called the **minimum acuity rule**, for below this resolution, a sensor measures nothing useful at all. Conversely, the resolution of a wheel step hazard will require that a pixel subtend no more than half of the height of the smallest obstacle. The smallest obstacle which presents a hazard is on the order of the wheel radius. Choosing to require two pixels on this surface can be written as:

$$dz \leq \frac{r}{2}$$

This will be called the **maximum acuity rule** for pixel sizes smaller than this are excessively small⁶. There is some range at which the acuity requirements change relative severity, so in the most general case, both must be met simultaneously.

1.5.4 Algorithmic Solutions - Acuity Adaptive Scan and Planning

The idea of **acuity adaptive scan** is to actively modify sensor resolution in order to ensure adequate acuity over the entire field of view of the sensor. This measure is important because traditional sensors, such as the ERIM laser rangefinder, admit two orders of magnitude of variation in the density of pixels on the groundplane.

The idea of **acuity adaptive planning** is to accept that sensory information is finite in resolution and to perform planning computations at a consistent resolution. For example, vehicle pitch evaluated by moving the vehicle over the map cannot change until a distance of one map cell is moved, so there is no point in wasting computer cycles trying to extract higher resolution information. This would be equivalent graphically to using many small steps to integrate a function over an interval for which it is constant.

1.5.5 Mitigating Assumptions - Uniform Scan and Terrain Smoothness

In the most general case, pathological hazards may exist which are impossible to resolve for both sensor resolution and computational throughput reasons. A nail, for instance, could feasibly exist on the vehicle trajectory. Therefore it is necessary to assume that pathological cases do not exist. This will be called the **terrain smoothness assumption**. It must always be adopted to some degree, and it is adopted implicitly when a terrain map of any finite resolution is used in planning computations. This assumption is also important because the projection of a pixel onto the groundplane depends on the shape of the terrain surface itself⁷. To assume that pixel resolution is inherently adequate is to adopt a **uniform scan assumption**. One way to achieve almost uniform scan is to mount a sensor directly over the terrain of interest⁸.

1.5.6 Related Subproblems - Sampling Problem

The true relationship between the angular resolution of any particular sensor and the linear resolution of the measurements it provides is affected primarily by the height at which the sensor is mounted and the shape of the terrain. In practice, it can vary by orders of magnitude over the field of view. This will be called the **sampling problem**.

Acuity has a temporal element as well as the spatial elements mentioned above. It is necessary to sample vehicle position fast enough to correctly localize range pixels. The distortion of an image that is caused by motion of the sensor is called the **motion distortion problem**.

6. Read excessively small as information theory. In practice, it is a good idea to require more than 2 pixels on the surface.

7. From radiometry, the projected area of a surface varies with the cosine of the angle of incidence.

8. This, of course can only be done for limited excursions or with a flying vehicle. Neither option is practical here, but both have limited domains of usefulness to other problems.

1.6 Fidelity


There is a basic requirement that the system be able to locate the vehicle with respect to hazards in the environment to within some limit of error⁹. This accuracy requirement applies to both the range sensor, the position sensor, and the computations performed on the data. Clearly, the smaller the vehicle and the larger the distance between hazards, the lower this required accuracy is, so it is to be expected that the accuracy requirement will be dependent on the vehicle dimensions as well as the density of hazards.

1.6.1 Requirement - Guaranteed Localization

In order to guarantee localization of hazards, some fidelity requirement must be maintained at all times. This will be called the policy of **guaranteed localization**. The problem of maintaining adequate fidelity will be called the **fidelity problem**.

1.6.2 Nondimensional Requirement - Fidelity Ratios

Considering the groundplane projection of the vehicle and two hazard areas, two very useful expressions can be formed:



Called the **fidelity ratios**. In the figure Δ is the minimum distance between hazards, W is the vehicle dimension aligned with the line between the hazards, dx is the maximum allowed error, and δ is the terrain **map resolution**.

1.6.3 Design Rules - Fidelity Rules

In order to ensure that the vehicle does not collide with the hazards, position error of the vehicle *relative to the hazards*, must not exceed the maximum, or equivalently the ratio must not exceed one.

$$\frac{dx}{(\Delta - W)} \leq 1$$

This will be called the **minimum fidelity rule**, for below this accuracy, collision is guaranteed.

9. This is deliberately stated in relative terms. If both the hazard and the vehicle position are off by exactly the same error, system viability is not affected. Therefore, some aspects of the calibration problem are not important.

Conversely, there is little point in measuring geometry to superb accuracy when it will be reduced to the intrinsic map resolution before it is used.

$$\frac{dx}{\delta} \leq 1$$

This will be called the **maximum fidelity rule** for accuracies better than this are excessive.

1.6.4 Algorithmic Solutions - Fidelity Adaptive Planning

When predicting the trajectory of the vehicle forward in time, the idea of **fidelity adaptive planning** is to accept that system accuracy is finite and to conservatively avoid obstacles so that hazards are guaranteed to be avoided despite the levels of uncertainty that exist. This measure is important because, in dense obstacle environments, the correct response to an overly dense collection of hazards is to avoid them as a unit instead of trying to drive between them.

1.6.5 Mitigating Assumption - Benign Terrain Assumption, Low Dynamics Assumption

It is clear that a vehicle cannot navigate between two hazards that are closer together than the size of the vehicle and in practice, it may be necessary to accept suboptimal localization from a poor sensor. To do so is to fundamentally assume that hazards are sparse in the environment. This will be called the **obstacle sparsity assumption** which is a special case of the **benign terrain assumption**. A special extreme case of the benign terrain assumption is the **flat world assumption**. In situations where dynamics can be neglected, a **low dynamics assumption** is being adopted.

1.6.6 Related Subproblems - Sensitivity Problem, Image Registration Problem

There is often a high degree of sensitivity of particular parameters to changes in another. For example, the localization of a range pixel is very sensitive to errors in the range measurement or in the angular measurement of the position of the ray through the pixel with respect to the navigation coordinate system. When sensitivity becomes an important consideration, a **sensitivity problem** exists. Terrain map fidelity is sensitive to many factors.

Also, guaranteed detection implies that an autonomous system is concerned mostly with predictions of vehicle position relatively far into the future. When nonlinear differential equations are involved in this prediction, there is the potential for extreme sensitivity to exist¹⁰. The solution to the Fresnel equations for the vehicle position several seconds into the future, for example, is very sensitive to dynamic model miscalibration errors.

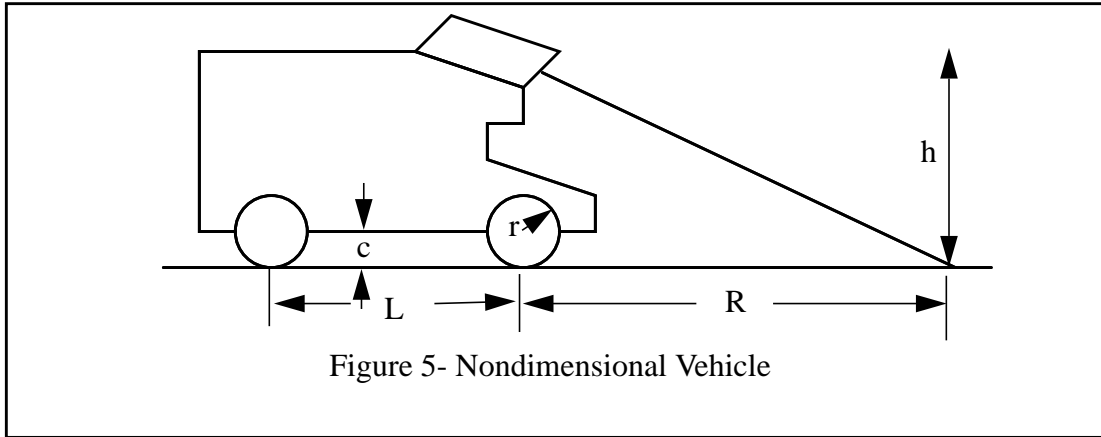
There are many other dimensions to the fidelity problem. The **image registration problem** arises when the relative accuracy of the entire system geometric model is insufficient to ensure that redundant measurements of the same geometry agree. Sometimes, especially in control and estimation applications, situations can arise where there are not enough sensors to measure all quantities of interest. This problem is a kind of extreme fidelity problem known as the **observability problem**.

10. When the sensitive parameters are the initial conditions, this is, of course, the notion of chaos in nonlinear dynamical systems. Vehicle position is indeed quite sensitive to the initial curvature if it is not accounted for at all.

2. Configuration

The shapes of most terrestrial vehicles vary little even while their sizes vary alot. Most are about twice as long as they are high and about as wide as they are high. This is especially true of vehicles designed for roads. Simple models of tipover stability dictate that the height should be reduced as much as possible and, all other things being equal, length and width should be equal.

From the point of view of impact on a perception system, a few dimensions are most important. Consider the dimensions indicated in the following figure.



Let the largest important dimension of the vehicle, the **wheelbase**, be called L . Define the **normalized wheelbase** to be:

$$\bar{L} = \frac{L}{R}$$

The normalized wheelbase relates each of the two following sets of three variables. Define the **perception ratio**, **normalized wheel radius**, and **normalized undercarriage clearance** to be:

$$\bar{h} = \left(\frac{h}{R}\right) \quad \bar{r} = \left(\frac{r}{R}\right) \quad \bar{c} = \left(\frac{c}{R}\right)$$

These variables will play a key role in rules which relate perception system requirements to the vehicle itself. By borrowing terminology from wing theory, the vehicle shape can be expressed in terms of aspect ratios. Define the **longitudinal aspect ratio**, **wheel fraction**, and **undercarriage tangent** to be:

$$A_L = \left(\frac{h}{L}\right) \quad A_r = \left(\frac{r}{L}\right) \quad A_c = \left(\frac{c}{L}\right)$$

These measure vehicle shape. The first measures how oblong the profile is. The second measures the overall roundness of the traction system and is one of many factors affecting terrainability. The third is another measure of terrainability and is a key element in uncovering the basic reason for occlusions in a perception system. Future results will refer to the normalized vehicle frequently. Of the seven variables presented, only four are independent.

2.1 Summary of Configuration Nondimensionals

The following table summarizes the vehicle shape nondimensionals.

Table 1: Vehicle Shape Nondimensionals

Symbol	Name	Expression	Symbol	Name	Expression
\bar{L}	Normalized Wheelbase	$\frac{L}{R}$			
\bar{h}	Perception Ratio	$\frac{h}{R}$	A_L	Longitudinal Aspect Ratio	$\frac{h}{L}$
\bar{r}	Normalized Wheel Radius	$\frac{r}{R}$	A_r	Wheel Fraction	$\frac{r}{L}$
\bar{c}	Normalized Undercarriage Clearance	$\frac{c}{R}$	A_c	Undercarriage Tangent	$\frac{c}{L}$

PART II:Response

As was shown earlier, the **response ratio** relates the ability of the vehicle to react to its speed and its sensory lookahead. This section analyses these aspects of vehicle performance for typical vehicles.

1. Reaction Time

The **system reaction time** is the time period between the instant that an object appears in the field of view of the sensor and the instant of time when the vehicle can be considered to have reacted to it. This aspect of performance depends on both software and hardware components.

1.1 System Reaction Time

The system reaction time can be computed by tracing the flow of information through the basic system processing loops. These can be understood from the following flow diagram:

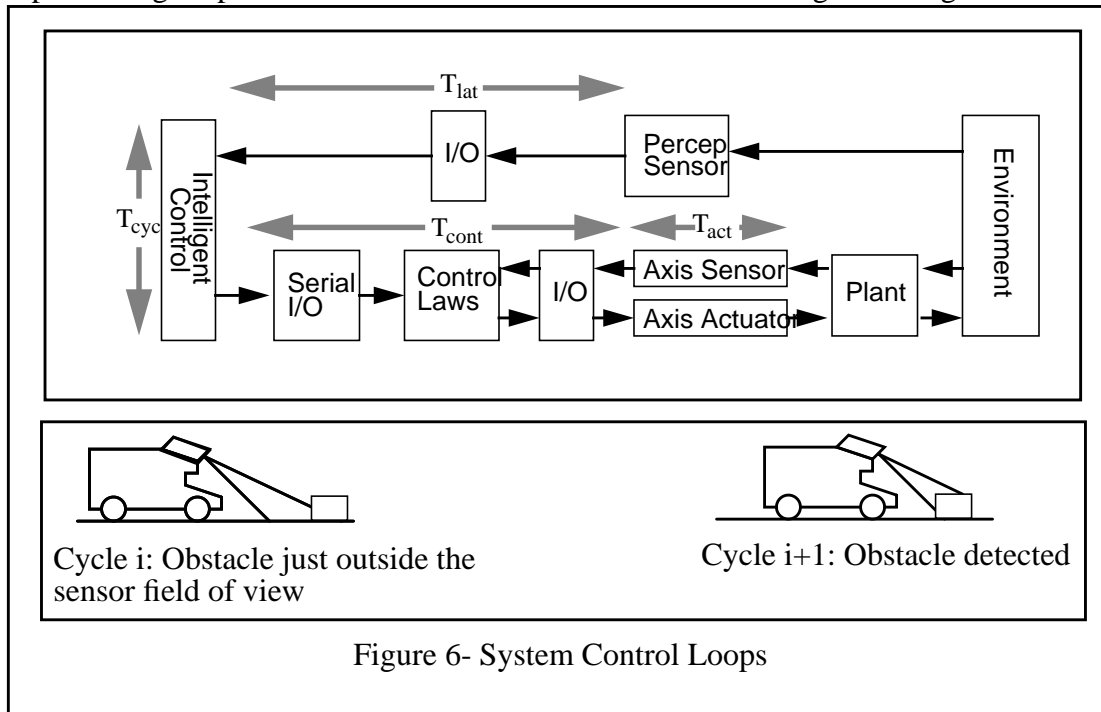


Figure 6- System Control Loops

Consider that a clock is started the instant that an obstacle appears in the field of view of the sensor. After the frame buffer latency T_{lat} has elapsed, the obstacle appears in a new image in the frame buffer. Unless software is synchronized with the sensor, in the worst case, software has just started a new cycle immediately before this occurs. After T_{cyc} more seconds have elapsed, software starts a new cycle on the latest image. This image may be the first, second or some other image since the obstacle was first seen, but the image does contain the obstacle.

When this image is processed in cycle $i+1$, T_{cyc} more seconds will elapse before the planner has decided to stop and issues the brake command. Hence $T_{lat} + 2T_{cyc}$ seconds of time have elapsed before the system has decided to react by braking the vehicle. This analysis assumes perfect obstacle detection. Next, the communication link to the control computer incorporates a delay called T_{cont} . Finally, after the controller receives the brake command, a delay of T_{act} applies before the mass of the actuator moves far enough to be considered to have responded to the amplifier drive current. Thus the worst case time it takes for the system to react to a situation is given by:

$$T_{react} = T_{lat} + 2T_{cyc} + T_{cont} + T_{act}$$

The coefficient of 2 arises from lack of synchronization between the sensor and the perception software. It is a worst case assumption. In reality, the coefficient of the software cycle time can be considered to vary randomly between 1 and 2 unless the cycle time is precisely constant.

Reaction time may be different for different actuators. In general, the steering and brake actuators may incorporate different delays. For this reason, two different reaction times are defined. The **braking reaction time** is called τ_B , and the **turning reaction time** is called τ_T .

1.2 Braking Reaction Time

According to the definition used here, the braking reaction time has elapsed after the brakes are fully engaged. That is, the time during which the vehicle decelerates is not included in this time. Therefore, an expression for the braking reaction time is:

$$\tau_B = T_{lat} + 2T_{cyc} + T_{cont} + T_{act}$$

where T_{act} is the small amount of time required for the brake actuator to move.

1.3 Turning Reaction Time

In the case of turning, the turning reaction time has elapsed after the steering mechanism reaches the commanded curvature. The time for which the vehicle moves at this curvature is not included. Without loss of generality, let the steering mechanism move through an angle $\Delta\alpha$ and let its maximum velocity be $\dot{\alpha}_{max}$. Then the actuator delay is given by:

$$T_{act} = \frac{\Delta\alpha}{\dot{\alpha}_{max}}$$

which can be as much as 3 seconds under some circumstances.

One of the most important aspects of the **latency problem** is the large value of the turning reaction time. For this reason, systems which attempt continuous motion in a dense obstacle field require excellent turning response characteristics.

1.4 Nondimensional Response

The only physical dimension represented in the reaction time equation is time, so it is not very interesting from the point of view of dimensional analysis. A dimensional analysis can be accomplished by simple division. Consider:

$$\frac{2T_{\text{cyc}}}{T_{\text{react}}} + \frac{T_{\text{lat}} + T_{\text{cont}} + T_{\text{act}}}{T_{\text{react}}} = 1$$

which states the intuitively obvious result that it is the relative ratios of the software and hardware components of the system reaction time which are important. If software efficiency approaches the point where it accounts for only a small fraction of the total system reaction time, attempted performance improvements should concentrate on hardware and vice versa.

2. Maneuverability


In addition to its ability to react, the ability of a vehicle to maneuver to avoid obstacles or otherwise ensure safety is important. This section investigates the manner in which *computational reaction time and mechanical maneuverability together* determine the ability of a vehicle to avoid obstacles. Four special obstacle avoidance maneuvers will be defined. These maneuvers are the **panic stop**, the **turning stop**, the **impulse turn**, and the **reverse turn**. These maneuvers will be defined in an ideal sense for a point robot with actuators which respond instantaneously after a command is received. Actuator dynamics will be considered in a later section.

2.1 Braking

Once a command to stop is issued to the vehicle brakes, a certain amount of distance is travelled in order to exhaust the vehicle kinetic energy through friction between the wheels and the terrain. This is the **braking distance**. Simple analysis suggests that this distance is quadratic in vehicle velocity V . Let the vehicle mass be M , then its kinetic energy is $MV^2/2$. This energy must be removed through the work done by the frictional force. If the dynamic coefficient of friction is μ , and the distance travelled is s , and g is the local acceleration due to gravity, then this work is μMgs . Equating these gives:

$$\frac{MV^2}{2} = \mu Mgs \Rightarrow V^2 = 2\mu gs$$

An experiment was conducted¹¹ which verifies this quadratic result. Fitting the data to the model generates a coefficient of friction of from 0.5 to 1.0. The data are presented below:



11. This graph is courtesy of R. Coulter of CMU based on an experiment conducted in the HMMWV.

2.2 Panic Stop Maneuver

Consider that the vehicle approaches an obstacle which it cannot avoid. That is, its only alternative is to stop as quickly as it can. This will be called a **panic stop maneuver**. The **stopping distance** is the distance travelled by the vehicle between the instant an object appears in the sensor field of view and the instant when the vehicle comes to a complete stop. If the vehicle velocity is V , then the **planning distance** is the distance travelled before the brake is applied, or $[\tau_B] V$. The **braking distance** is the distance actually travelled while the brake is on. If the two are added, the following expression is obtained for the distance travelled by the vehicle:

$$s_B = \tau_B V + \left[\frac{V^2}{2\mu g} \right]$$

This relationship is plotted below for three values of reaction time using a conservative coefficient of friction of 0.5. Notice that stopping distance is quadratic in speed with a coefficient of about 0.1.




Figure 7- Stopping Distance

2.3 Nondimensional Braking

A rudimentary dimensional analysis is performed by asking at what velocity of the vehicle do the planning distance and the braking distance become equal. This happens when the ratio of the two distances is 1. The dimensions in the equation are distance and time. Hence there are two pi products. There are many ways to choose these because pi products are not unique. Later results will benefit from the following choice of variables.

The first is the **kinematic braking coefficient**:

$$\bar{b}_k \equiv \frac{T_{\text{react}} V}{s}$$

which is the proportion of the stopping distance which is consumed in simply deciding to brake.

The second is the **dynamic braking coefficient**:

$$\bar{b}_d \equiv \frac{V}{2\mu g T_{\text{react}}}$$

which is the ratio of the braking distance to the planning distance.

Rewritten, the stopping distance expression takes the following forms:

$$\bar{b}_k = \frac{1}{1 + \bar{b}_d} \qquad \bar{b}_d = \frac{1}{\bar{b}_k} - 1$$

These expressions capture everything in Figure 7 in nondimensional form. Since both Pi products must satisfy this equation, either can be derived from the other. The dynamic braking coefficient will be considered to be fundamental and referred to simply as the **braking coefficient** \bar{b} . When the braking coefficient exceeds, say 0.1, it becomes important for a planning system to reason about the dynamics of braking.

Real Numbers

For the HMMWV at speeds of 5 meters/sec, with a 2 second reaction time, the braking coefficient is 0.25.

2.4 Braking Regimes

The braking coefficient identifies two key regimes of operation for autonomous vehicles. After substituting into the original expression, some algebra gives:

$$\frac{s_B}{\tau_B V} = 1 + \frac{V}{2\mu g \tau_B}$$
$$s_B = \tau_B V [1 + \bar{b}]$$

When the braking coefficient is 1.0, braking distance and planning distance are equal. At this point, stopping distance enters a regime of quadratic growth with initial velocity. Based on the braking coefficient, two regimes of operation can be defined. In the **kinematic braking regime** it is much less than 1.0. In the **dynamic braking regime** it is much greater than 1.0.

In the kinematic regime, stopping distance is linear in both initial velocity and reaction time. This is the reason why the curves in Figure 7 are basically linear even though they are given by a quadratic equation. In the dynamic regime, stopping distance is quadratic in initial velocity and independent of reaction time. In order to achieve a given target speed, *the path planning horizon can be reduced significantly if the system reaction time is reduced to a minimum.*

Real Numbers

For the HMMWV, with a 2 second reaction time, the braking coefficient is 1 at a velocity of 20 meters/second. Hence, the system will operate in the kinematic regime throughout the range of speeds considered in this report. Based on this single fact, many of the results which will follow can now be predicted based on intuition. Further, this explains why Figure 7 is basically linear although it is given by a quadratic equation.

2.5 Turning


Consider a **constant curvature turn**. Let the vehicle yaw be given by ψ , the velocity be given by V , the curvature be given by κ , and the radius of curvature be given by ρ . In a turning maneuver, the instantaneous vehicle **yaw rate** is given simply by the chain rule of differentiation:

$$\dot{\psi} = \frac{d\psi}{ds} \frac{ds}{dt} = \kappa V = \frac{V}{\rho}$$

If the time spent in the turn is T , the yaw of the vehicle after the turn is given by:

$$\psi = \frac{s_T}{\rho} = \frac{TV}{\rho}$$

Where s_T is the **turning distance**. When a constant curvature turn is executed, there is a simple relationship between the yaw of the vehicle and the angle subtended at the start point by the stop point. This is given below:



The range from the start point to the endpoint is given by:

$$R = \sqrt{x^2 + y^2} = \rho \sqrt{2} \sqrt{1 - \cos\psi}$$

The angle from the start point to the endpoint is given by:

$$\theta = \text{atan}\left(\frac{y}{x}\right) = \text{atan}\left(\frac{1 - \cos\psi}{\sin\psi}\right) = \text{atan}\left(\frac{\sin\psi/2}{\cos\psi/2}\right) = \psi/2$$

Which is a very useful result for determining the angular width of the region which is reachable in a given time at a given speed. **The angular width, subtended at the start point, of the region reachable by the vehicle in a turn, is the yaw of the turn itself.**

If the vehicle decides to brake while in a turn, the **planning angle** is the angle travelled before the brakes are actuated. The **braking angle** is the angle travelled while the brakes are actuated before coming to a stop. The **stopping angle** is the sum of these two. In a manner analogous to the coefficient of friction, a **coefficient of lateral acceleration** can be defined thus:

$$v = \frac{a_{\max}}{g} = \frac{V^2}{\rho g} = \frac{\kappa V^2}{g}$$

which is simply the lateral acceleration expressed in g's. The lateral acceleration at all points on a turning trajectory must be limited for safety reasons.

2.6 Steering Limits

There exists a maximum steering angle for any vehicle speed which will cause the vehicle to enter an unsafe state. This will be called the **dynamic steering limit** ρ_D . The radius of curvature of a vehicle may also be kinematically limited by the steering mechanism. This is the **kinematic steering limit** ρ_K . To determine the dynamic limits, let the maximum safe lateral acceleration be $0.5 g^{12}$. Let κ be the path curvature, V be the velocity, and ρ be the radius of curvature. Then, the velocity and radius of curvature are related by:

$$V^2 \kappa = \frac{V^2}{\rho} < v g \Rightarrow \rho_D = \frac{V^2}{v g}$$

This relationship is plotted below for all vehicle speeds up to 10 m/s using a kinematic steering limit of 7.5 meters.




Figure 8- Steering Limits

Hence, minimum radius of curvature is quadratic in speed with a coefficient of about 0.2. The minimum radius of curvature can be expressed by the function:

$$\rho_{\min} = \max(\rho_K, \rho_D) = \max\left(\rho_K, \frac{V^2}{v g}\right)$$

The discontinuous derivative of this function will show up several times later.

12. This analysis is a gross approximation which is sufficient for the purpose it is used.

2.7 Turning Stop Maneuver

The **turning stop maneuver** occurs when the vehicle is executing a constant curvature turn and detects an obstacle in its path that is avoided by braking. The **stopping region** is the region of space in front of the vehicle which includes all possible braking trajectories. Clearly, this region is a function of the initial velocity and the initial curvature. It can be quantified by considering all initial curvatures at all speeds, and simulating forward along all trajectories until the vehicle comes to a stop. The stopping regions for two different speeds are plotted below.

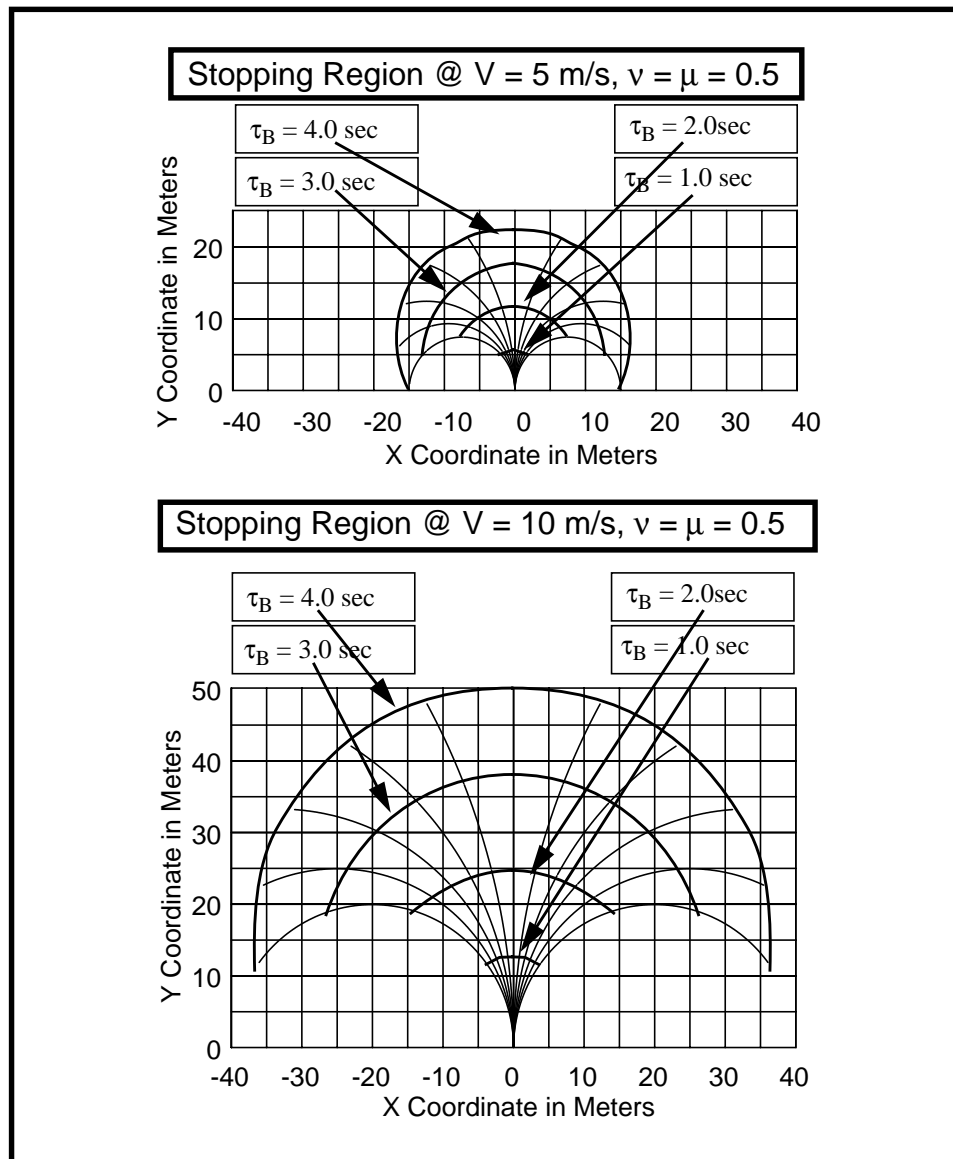


Figure 9- Stopping Region

2.8 Impulse Turn Maneuver

Consider the case when the vehicle is executing a linear trajectory and decides to avoid an obstacle by executing a turn and not by hitting the brakes. This will be called an **impulse turn maneuver**. The vehicle will travel a distance given by the velocity and the reaction time before the steering mechanism is actuated. Then, the sharpest turn is determined by the kinematic and dynamic steering limits. Idealized turning trajectories for a point robot based upon a lateral acceleration limit of 0.5 g are plotted below¹³.




Figure 10- Impulse Turning Trajectories

13. This analysis is a gross conservative approximation for several reasons. The dynamics of the steering mechanism are neglected even though it takes several seconds for a standard steering mechanism to reach full deflection. For this reason, reaction times applied to steering must be increased relative to braking. Also, the finite width of the vehicle must be considered. This can be achieved by further increasing the reaction time.

2.9 Impulse Turning Distance

In the impulse turn maneuver, notice the V-shaped region between the opposite turning trajectories. The width of this region is the width of the largest obstacle directly in the path of the vehicle that can be avoided¹⁴. This width becomes infinite after 90 degrees of turn.

For the impulse turn maneuver, the **planning distance** is the time taken for the steering mechanism to be actuated. The **turning distance** is the distance *along the original trajectory* consumed in turning. The **impulse turning distance** is the sum of these two.

An impulse turn of particular interest is the 90 degree impulse turn. This maneuver is required to avoid a large object without braking. The impulse turning distance is given by:

$$s_{IT} = \tau_T V + \rho_{\min}$$

Which can be written as:

$$s_{IT} = \tau_T V + \max(\rho_K, \frac{V^2}{vg})$$

The impulse turning distance is plotted below:




Figure 11- Impulse Turning Distance

14. Remember that a point robot was assumed, so this V shaped region should be moved back a few meters to account for this or the turning reaction time should be increased. On the other hand, point obstacles require substantially less turning than full deflection in order to avoid them which reduces the effective impulse turning distance. A detailed analysis is too detailed for the purposes of the report. The point is that large obstacles really do occur regularly, so this condition must drive sensor maximum range. Ideally, the offset of the vehicle nose from the sensor must be accounted for as well as the effect of vehicle motion on the scanning pattern.

2.10 Reverse Turn

Consider a situation in which the vehicle is already engaged in a sharp turn in one direction and decides to reverse curvature, and proceed in the opposite direction. This will be called a **reverse turn**. The vehicle will travel along a trajectory given by the velocity, the initial curvature, and the reaction time before the steering mechanism is actuated. Then, for an idealized trajectory, the curvature will instantaneously switch. Any dynamics in the steering response are accounted for in the turning reaction time. Idealized turning trajectories for a point robot based upon a lateral acceleration limit of 0.5 g are plotted below:




Figure 12- Reverse Turning Trajectories

It is important to recognize that the vehicle will continue on its original curved trajectory in this model until the turning reaction time expires. This is why these trajectories deviate so significantly from the trajectory that would be followed if the reaction time were zero.

2.11 Nondimensional Turning

The angle through which the vehicle turns in any constant curvature trajectory is given by:

$$\psi = \left(\frac{s}{\rho} \right)$$

If the vehicle were to execute a constant velocity turn for a time period of the reaction time at some radius of curvature, it would turn through an angle given by the arc length over the radius of curvature. Since it will be useful later, the inverse of this angle will be called the **kinematic turning coefficient** or simply the **turning coefficient**:

$$\bar{t}_k = \bar{t} = \frac{\rho}{V T_{\text{react}}}$$

This quantity can be defined for any velocity and any radius curve. It is the inverse of the amount that the vehicle turns over a period of one reaction time. Now a quantity analogous to the dynamic braking coefficient can be defined. Let the **dynamic turning coefficient** be given by:

$$\bar{t}_d = \frac{V^2}{vg\rho}$$

It can be defined for any velocity and any radius curve. This quantity measures, for any curve and velocity, the proximity of the vehicle to its lateral acceleration limit. As such, it is a measure of safety of a turn.

2.12 Turning Regimes

If the variation in the angular width of the stopping region is plotted against speed, it exhibits a maximum when the dynamic turning coefficient first reaches one. Beyond this speed, the vehicle tends to become less omnidirectional as the dynamics of turning come into play. Hence there are turning regimes which are analogous to the braking regimes. In the **kinematic turning regime**, the dynamic turning coefficient is less than one and turns of the minimum kinematic radius are safe. In the **dynamic turning regime**, the dynamic turning coefficient is greater than or equal to one.

In the dynamic regime, the effect of increased speed is to **reduce** the angular width of the stopping region. The vehicle approaches kinematic omnidirectionality when the turning coefficient significantly exceeds one. It will be shown that, at any speed, the horizontal field of view required is given by the minimum value of the turning coefficient for any safe trajectory.

2.13 Turning Stop Maneuver

Suppose a vehicle is required to stop to avoid an obstacle while it is executing a turning trajectory. The angle that it will turn through before it stops is again given by:

$$\psi = \frac{s}{\rho}$$

Substituting for the stopping distance in terms of the braking coefficient:

$$\psi = \frac{\tau_B V [1 + \bar{b}]}{\rho} = \frac{[1 + \bar{b}]}{\bar{t}}$$

which is a pleasingly simple result.

By analogy to the linear braking trajectory, three angles can be defined. The **planning angle** is the angle through which the vehicle turns while deciding to stop. The planning angle for a turning stop maneuver is the inverse of the turning coefficient.

The **braking angle** is the angle through which it turns while braking. This is given by the ratio of the braking coefficient and turning coefficients. The **stopping angle** is the sum of these two. For the turning stop trajectory, the ratio of the braking angle to the planning angle is the braking coefficient just as it was in the linear braking case.

Further, for a turning stop, the braking and turning coefficients are related as given below:

$$\bar{t}_d = 2 \bar{b} \left(\frac{\mu}{v} \right) \qquad \bar{b}_k = \frac{1}{\bar{t}} \left(\frac{\rho}{s} \right)$$

It has been shown that only one braking coefficient is independent. From this result, it is clear that there is only a single second independent variable - the ratio of stopping distance to radius of curvature. This is, of course, the **stopping angle** which was just defined.

For a turning stop trajectory, only the braking and turning coefficients are independent. All others can be derived from them. When the stopping distance is much greater than the minimum radius of curvature, the turning coefficient is significantly less than one. Under these conditions, the vehicle can be considered to be kinematically omnidirectional at the resolution of the stopping distance.

2.14 Impulse Turn Maneuver

The impulse turning distance was defined as the planning distance plus the minimum radius for the tightest possible turn. This is given by:

$$s_{IT} = \tau_T V + \rho_{\min} = \tau_T V \left[1 + \frac{\rho_{\min}}{\tau_T V} \right] = \tau_T V [1 + \bar{t}]$$

2.15 Impulse Turning Regimes

This result is a perfect analogy to the stopping distance relationship. When the turning coefficient is greater than one, the second term dominates and growth is quadratic. More precisely, the nature of the growth of the impulse turning distance then depends on the dynamic turning coefficient. The different regimes are illustrated in Figure 16. At sufficient speeds, the quadratic growth of radius of curvature dominates the linear growth of the planning distance. This occurs when

$$\rho_{\min} = \tau_T V = \frac{V^2}{vg} \Rightarrow V = \tau_T vg$$

Real Numbers

On the HMMWV, for a 2 second turning reaction time, this occurs at 10 m/s. For a 4 second reaction time, it occurs at 20 m/s. Unless reaction times can be reduced below 2 seconds (and this is not likely) the impulse turning distance can be considered linear in velocity.

Notice that the turning coefficient plays the same role for impulse turning that the braking coefficient plays for the panic stop. The dynamic turning coefficient plays a similar role for continuous turning.


3. Lookahead

The final element of the **response ratio** is the lookahead distance. The highest level requirement on perception is to image all of the terrain that the vehicle still has an option of traversing at any particular point in time. Ideally, a range sensor should provide geometric information *beyond* the entire region of terrain that can be reached by the vehicle before coming to a stop.

3.1 Adaptive Regard

From the perspective of turning and its impact on the horizontal field of view, it is clear from earlier sections that there are situations during turns when the entire left half or right half of the field of view need not be processed because the vehicle cannot go there. Further, from the perspective of braking and its impact on the vertical field of view, it is clear that there is some distance roughly given by the reaction time times the speed within which the vehicle cannot stop.

It is useful to partition the space in front of the vehicle into different zones based upon the different implications of an obstacle in each zone. Consider the following figure:




Regions are defined for each particular possible trajectory and for each type of avoidance maneuver. The **dead zone** is the region within the swath of the planned vehicle trajectory which the vehicle is committed to travelling. Should an obstacle ever enter this region, a collision is unavoidable. The **detection zone** is the region along the trajectory which can still be avoided. **Free zones** are areas in neither of the other two. These are regions where the vehicle cannot go because of its steering and braking limitations.

Since an obstacle in the dead zone cannot be avoided, there is no reason to allocate precious range pixels there. Rather, obstacles must be detected in the detection zone before they ever reach this region. As a minimum, a range sensor must allocate range pixels over the region formed by the union of the detection zones of all possible trajectories for all possible avoidance maneuvers. The dead zone grows as vehicle speed increases. This is the mechanism which drives the sensor field of view away from the vehicle with increased velocity.

In the context of planning these notions are important and the dead and free zones are a very large fraction of the total. There is no point in wasting perceptual and planning cycles in discovering the geometry of regions in these categories because *there is no useful decision that the planner can make*. This notion will be called **adaptive regard**. The principle of adaptive regard is to process geometry only in the detection zone which is to say that *a system must adapt in real time to both its speed and its curvature*.

3.2 Pointing Rules

In order to detect obstacles beyond the dead zone, a range sensor must supply range pixels in the region beyond the dead zone, called the detection zone. The width of the detection region used is called the **incremental lookahead distance** and the angular width beyond the dead zone is called the **incremental lookahead angle**. The planner lookahead required for a turning stop maneuver is indicated in the following figure:



The lookahead distance must be chosen such that obstacles can be detected beyond the dead zone of all possible avoidance maneuvers. There are two considerations involved in choosing this distance.

- In order to compute a reliable prediction of vehicle pitch, it is necessary for the sensor field of view to extend for at least one vehicle wheelbase beyond the dead zone.
- Reliable small obstacle detection may require that the sensor field of view be such that several images fall on an obstacle before it leaves the detection zone. On this basis, the planner lookahead should exceed the dead zone by several times the product of the velocity and the sensor frame period.

In this report, the vehicle wheelbase will be chosen as the incremental lookahead distance.

3.3 Adaptive Lookahead

Adaptive regard is a notion defined for path planning purposes. It restricts the region that an obstacle detector considers to the detection zone for some set of avoidance maneuvers. A direct implementation of the pointing rules for the purposes of guaranteeing response will be called **adaptive lookahead**. Adaptive lookahead is a notion defined for perception purposes. It restricts the focus of attention of perception to the detection zone. Practically, this mechanism computes the minimum range or maximum range based on the obstacle avoidance maneuver in question. In some implementations, the distinction between these two devices is irrelevant.

3.4 Nondimensional Lookahead

For the purpose of providing information to the path planner, the perception system must look beyond the dead zone. The amount of lookahead is called the **incremental lookahead distance** ΔR_L . Similarly, the maximum sensor horizontal field of view is greater than the stopping angle by an amount called the **incremental lookahead angle** $\Delta \psi_L$. Based on these, two convenient nondimensionals can be defined.

The **normalized incremental lookahead distance** is the lookahead distance normalized by planning distance.

$$\Delta \bar{R}_L = \frac{\Delta R_L}{V T_{\text{react}}}$$

It is related to the **throughput ratio** defined earlier. The **lookahead angle** is already nondimensional. It is the incremental lookahead distance divided by the instantaneous radius of curvature


$$\Delta \bar{\psi}_L = \frac{\Delta R_L}{\rho}$$

Let the ratio of maximum to minimum range be called the **range ratio**:

$$\bar{R} = \frac{R_{\text{max}}}{R_{\text{min}}}$$

3.5 Nondimensional Pointing Rules

The pointing rules relate the sensor field of view to the vehicle maneuverability at any speed. The vertical field of view is shown below:



The pointing rules that will be used to calculate the field of regard for planning purposes can be expressed as follows:

$R_{\min} = s_B$	stopping distance
$R_{\max} = s_{IT} + L$	impulse turning distance + lookahead
$\psi_{\max} = \frac{s_T + L}{\rho_{\min}}$	stopping angle + lookahead

3.6 Summary of the Response Nondimensionals

All of the nondimensionals given in this section are related to the response of the vehicle. The following table gives everything that will be useful later:

Table 3: Response Nondimensionals

Symbol	Name	Expression	Symbol	Name	Expression
μ	coefficient of friction	numeric	ν	coefficient of lateral acceleration	numeric
\bar{b}_k	kinematic braking coefficient	$\frac{VT_{\text{react}}}{s}$	\bar{t}_k, \bar{t}	(kinematic) turning coefficient	$\frac{\rho}{VT_{\text{react}}}$
\bar{b}_d, \bar{b}	(dynamic) braking coefficient	$\frac{v}{2\mu g T_{\text{react}}}$	\bar{t}_d	dynamic turning coefficient	$\frac{v^2}{vg\rho}$
$\Delta\bar{R}_L$	normalized incremental lookahead distance	$\frac{\Delta R_L}{VT_{\text{react}}}$	$\Delta\bar{\psi}_L$	incremental lookahead angle	$\frac{\Delta R_L}{\rho}$
\bar{R}	range ratio	$\frac{R_{\text{max}}}{R_{\text{min}}}$			

PART III:Throughput

As was shown earlier, the **throughput ratio** relates the ability of the vehicle to process information at a sufficient rate to its speed and its incremental sensory lookahead. This section analyses these aspects of vehicle performance for typical vehicles.

1. Depth of Field

This section and the following two sections investigate the relationship between the maneuverability of the vehicle and the sensor field of view. The sensor field of view is considered to be a solid cone of rectangular cross section which can be described by the horizontal field of view, vertical field of view, minimum range, and maximum range.

1.1 Minimum Sensor Range

The minimum range required of a sensor as a function of velocity is given by the closest point of any detection zone for any trajectory and any avoidance maneuver. Of the four avoidance maneuvers considered, the turning stop is the one which remains closest to the vehicle start point. The following figure was generated by tracing the endpoint of all turning stop trajectories, and computing the minimum range to each endpoint.

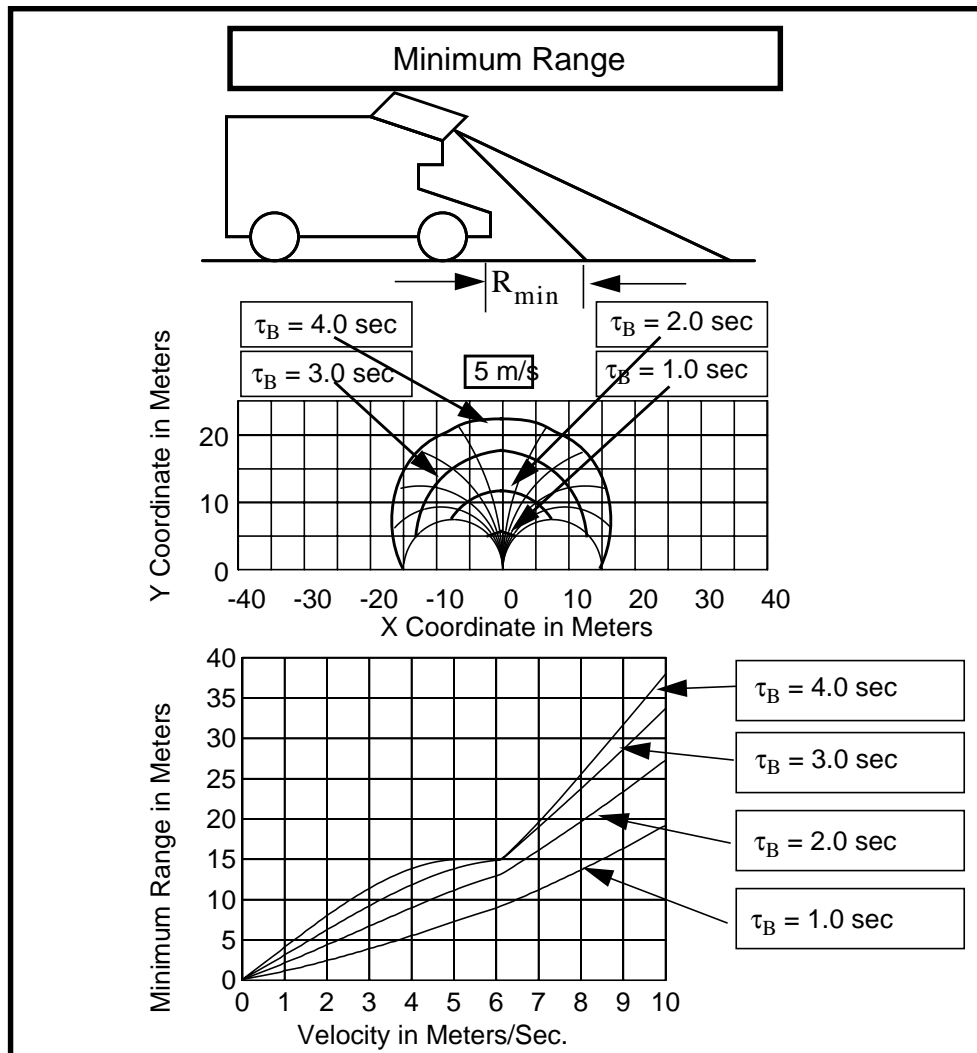


Figure 16- Minimum Range

The peculiar knee in the higher curves arises from the discontinuous derivative of minimum radius of curvature with velocity. Notice that at the knee, the minimum is twice the kinematic minimum radius. By comparing this figure with Figure 7, it is clear that the stopping distance overestimates the minimum range slightly, but it is a good approximation.

The minimum range is a loose requirement. There is no cost to response incurred by increasing the minimum range. It can be increased up to the point where throughput is barely guaranteed.

1.2 Maximum Sensor Range

Turning is preferred over braking for obstacle avoidance, so it is reasonable to let the impulse turning maneuver drive the specification of maximum range even though this maneuver requires much more space. The sensor maximum range is given by the sum of the impulse turning distance and the planner lookahead distance. The maximum range derived from this condition is given in the figure below¹⁵. Again, the discontinuity arises because of the discontinuity in the derivative of the radius of curvature.

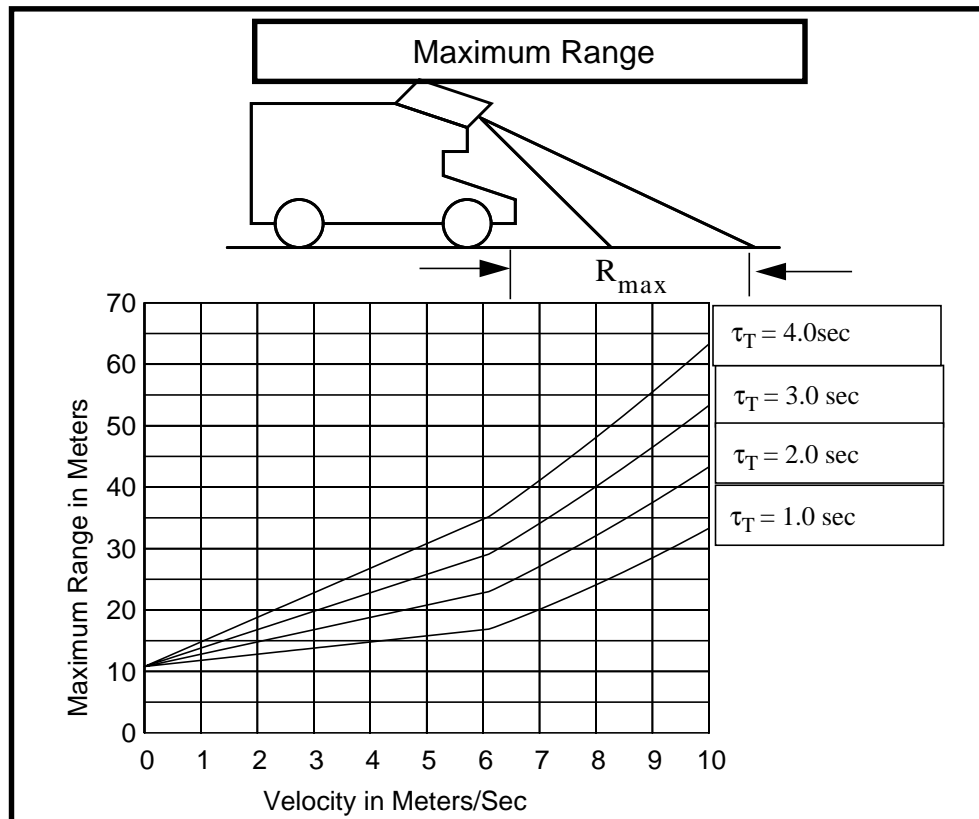


Figure 17- Maximum Range

For a single scanline sensor, the line must be placed at or beyond the maximum range computed above. Such sensors rely on very good obstacle detection.

Real Numbers

The maximum range of the ERIM scanner is 20 meters. Reaction times below 3 seconds for full deflection turns are probably not realistic. On the basis of this curve, the ERIM sensor supports no more than 3 m/s speeds while avoiding wide obstacles. This estimate incorporates no safety margin.

15. This graph is based on a very ambitious lateral acceleration limit and a small lookahead of the wheelbase. For this reason, a practical maximum range may be far larger than indicated by the graph. No numbers are available for realistic lateral acceleration limits.

1.3 Nondimensional Maximum and Minimum Range

According to the pointing rules, the sensor minimum range is given by the stopping distance.

$$R_{\min} = s_B = \tau_B V [1 + \bar{b}]$$

Also, the sensor maximum range is given by the impulse turning distance plus the lookahead distance. This gives.

$$R_{\max} = s_{IT} + L = \tau_T V [1 + \bar{t}] + \Delta R_L$$

Dividing gives an expression for the range ratio in terms of the vehicle maneuverability:

$$\frac{R_{\max}}{R_{\min}} = \frac{\tau_B V [1 + \bar{b}]}{\tau_T V [1 + \bar{t} + \Delta \bar{R}_L]} = \frac{\tau_B [1 + \bar{b}]}{\tau_T [1 + \bar{t} + \Delta \bar{R}_L]}$$

1.4 Myopia Problem

The obstacle resolving power of contemporary stereo and lidar systems is poor at the high ranges necessary to resolve obstacles at speed and *a system must attempt to succeed with a poor idea of what is out there*. The effective maximum range of a contemporary environmental perception sensor is limited by any of three concerns. First, angular resolution is often poor. Second, signal to noise ratios degrade with increased range and with shallow incidence angles. Third, occlusion is aggravated by the inherently shallow incidence of range pixels which arises from long range measurement.

2. Horizontal Field of View (HFOV)

The horizontal field of view will be determined by the turning stop maneuver. It can be argued that the **reverse turn** is a worse case, but this would drive the field of view to unreasonable size. It is reasonable to assume that stopping is the only option used when executing a tight turn.

The angle subtended at the vehicle by the dead zone is the maximum angle subtended at the vehicle by the stopping region for all vehicle speeds, since, in the worst case, the vehicle executes a continuous turn into new terrain. Recalling the 5 m/s stopping region, the angles can be read from the graphs of the stopping region for each value of system reaction time. The horizontal field of view is determined by the planner lookahead which is added to the angle subtended by the dead zone. The HFOV is derived by adding a lookahead distance of the vehicle wheelbase to the braking trajectory. In this way, a planner could detect and respond to a roll or pitch obstacle by braking.





Figure 18- Horizontal Field of View

The reaction times applicable to the turning stop are braking reaction times. Notice that the maximum is achieved at the fastest speed for which the smallest radius of curvature arc can be executed. From the previous graphs, this is 6.0 m/s. Beyond, this curvature, the required HFOV actually *decreases* because turns of the minimum radius are no longer possible.

A HFOV beyond 180 degrees seems counterintuitive. This is because the analysis assumes that information from previous images cannot be counted upon to provide information. The validity of this assumption depends on the frame rate, the scanning pattern, and the vehicle speed. More importantly, it depends on whether there are any occlusions which the sensor must see around. To adopt a HFOV smaller than the worst case requirement is to artificially reduce the maneuverability of the vehicle. This may be acceptable in some cases.

2.1 Nondimensional Horizontal Field of View

In dimensional analysis, ratios of important lengths often turn out to be important angles. As speeds increase beyond the maximum speed for which turns of the minimum radius of curvature are safe, curvature must be decreased to avoid excessive lateral acceleration. Consider the following figure:



This relationship gives the variation of the required HFOV with velocity for any vehicle. This is a very pleasing result. It states the intuition that both the hardness of the turn and the stopping distance should determine the result. In the kinematic braking regime, the result is linearly increasing with velocity. In the dynamic braking regime, it decreases quadratically. So the formula is consistent with the graphs in Figure 18.

Real Numbers

For the HMMWV, executing a minimum radius of curvature turn, with a reaction time of 2 secs, and 6 meters/sec speed, the formula gives:

$$\text{HFOV} = \left[\frac{T_{\text{react}} V}{\rho} \right] \left[1 + \frac{V}{2 T_{\text{react}} \mu g} \right] + \frac{L}{\rho} = \frac{2V}{7.5} \left[1 + \frac{V}{20} \right] + \frac{3.3}{7.5} = 2.52 \text{rads}$$

which is, from Figure 18, exactly correct.

Real Numbers

The maximum in HFOV is reached at 6 m/s for the HMMWV. Clearly, this must be the velocity at which the dynamic turning coefficient is one.

$$\bar{t}_d = \frac{V^2}{v g \rho} = \frac{6 \times 6}{\frac{1}{2} \times 10 \times 7.5} = \frac{72}{75}$$

The curious growth of HFOV followed by a decrease with velocity is unique to non omnidirectional vehicles. The minimum radius of curvature is zero for truly omnidirectional vehicles. They operate solely in the dynamic turning regime and the angular width of the stopping region is determined solely by the lateral acceleration limit.

2.2 Tunnel Vision Problem

Contemporary sensors have horizontal image sizes which are far too small to permit aggressive obstacle avoidance maneuvers on rough terrain. The analysis suggests that a HFOV of 120 degrees is a reasonable engineering guesstimate. Yet, typical stereo field of view is 40 degrees and the best rangefinder has a HFOV of 80 degrees. Consider the following figure in which the vehicle is executing a reverse turn. Turning dynamics imply that the entire region that the vehicle can reach is contained within the set of curves shown. Aggressive maneuvers cause collision simply because **a contemporary system cannot look where it is going**. In fact, in the reverse turn indicated below, there is no overlap at all between the field of view and the region that the vehicle is committed to travelling.






Figure 19- Tunnel Vision Problem

3. Vertical Field of View (VFOV)

The major *kinematic* system requirement which influences the vertical field of view is the pitch angle induced in the vehicle body by the most challenging, yet navigable, terrain. Later sections will develop *dynamic* requirements on the VFOV. Let the highest achievable pitch angle be 15° . Then the following figure illustrates the two cases which determine the vertical field of view required to ensure that the vehicle is able to see up an approaching hill or past a hill that it is cresting.



Therefore, *the vertical field of view required is four times the maximum pitch of the body*. It would also be useful to extend the lower half of the field of view to allow the vehicle to see behind occlusions as shown below:

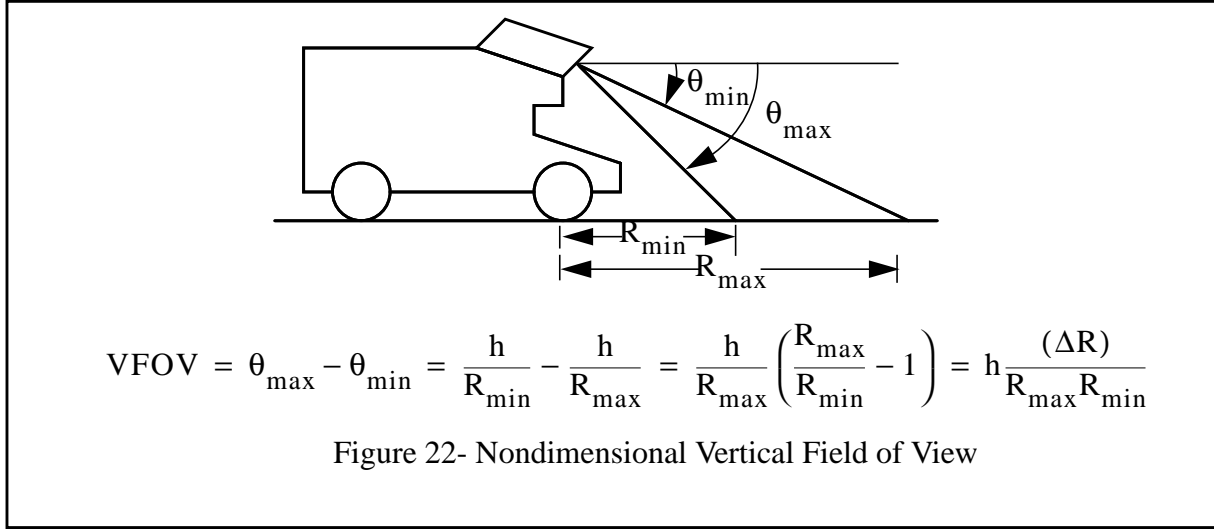


An extreme disadvantage of measuring range pixels above the horizontal is that most of the time they measure nothing. These pixels come into play only when the vehicle is executing challenging maneuvers. For this reason, it is advantageous to reduce the angle above the horizon as much as possible.

Another reason to reduce the VFOV above the horizon is that the left case in Figure 20 can be detected by a reduction in the highest measured range, so mechanisms to point the field of view or slow the vehicle are viable alternatives. In the right case, the highest measured range will likewise decrease, but slowing the vehicle does not deal with the problem, so sensor pointing is the only solution. Hence, it is better to allocate margin on the lower half of the vertical field of view.

3.1 Nondimensional Vertical Field of View

It is useful to assume that the height of the sensor is always less than the range measured. This is the **small incidence angle assumption**. The vertical field of view can then be expressed as follows:



When the sensor height is significantly less than the minimum range (and therefore the stopping distance), the vertical field of view can be expressed in terms of minimum value of the perception ratio and the range ratio.

$$\text{VFOV} = h \frac{(R_{\max} - R_{\min})}{R_{\max} R_{\min}}$$

$$\text{VFOV} = \bar{h}_{\min} (\bar{R} - 1) = \bar{h}_{\min} \bar{\Delta R}$$


This relationship gives the variation in required VFOV with speed for any vehicle. Substituting the pointing rules permits elimination of the range ratio and expressing the result in terms of the maneuverability coefficients:

$$\text{VFOV} = \frac{h}{R_{\min}} - \frac{h}{R_{\max}} = h \left[\frac{1}{\tau_B V [1 + \bar{b}]} + \frac{1}{\tau_T V [1 + \bar{t} + \bar{R}_L]} \right]$$

Therefore, in kinematic regimes, the VFOV decreases linearly with velocity. In dynamic regimes, it **decreases quadratically with velocity**. This result will be of extreme importance later.

4. Sweep Rate

The sweep rate of a sensor is defined in image space as the VFOV generated per unit time. It may represent the physical motion of the elevation mirror in a rangefinder or the product of the VFOV and the frame rate for a camera. The sweep rate required of a perception sensor is related to the velocity of the vehicle under guaranteed throughput. If $\dot{\theta}$ is the sensor sweep rate, the requirement on it can be approximated as follows:



$$\rho_{cyc} = \frac{VT_{cyc}}{\Delta R} \approx \frac{VT_{cyc}}{\dot{\theta} RT_{cyc}}$$

Thus, by insisting that the throughput ratio never exceed one, the equivalent requirement is that the sweep rate always exceed:

$$\dot{\theta} \geq \frac{V}{R}$$

which is pleasingly simple. The ratio of vehicle speed to sensor range is an effective angular velocity requirement which the sensor must meet. This gives the minimum sweep rate, or equivalently the minimum vertical field of view for some frame rate which guarantees that there are no holes in the coverage of the sensor. This is the linear velocity component of the **sweep rate rule**.

Equivalently, on rough terrain, the vehicle may pitch as a result of terrain following loads, and in the worst case, these motions add to the sweep rate requirement. If $\dot{\theta}_{max}$ is the maximum pitch rate of the vehicle, then the sweep rate rule becomes simply:

$$\dot{\theta} = \dot{\theta}_{max} + \frac{V}{R}$$

which relates the sensor field of view to both the velocity and the vehicle natural frequency. Thus rough terrain affects throughput to the degree that vehicle natural frequency may affect the pitch of the body.

4.1 Stabilization Problem

Notice that the linear component of the rule benefits from higher speeds whereas the angular component suffers. For the vertical field of view, the process of decreasing it in order to reduce throughput requirements will eventually lead to the creation of a **stabilization problem**. The following figure indicates that rough terrain can create severe stabilization constraints for a narrow VFOV sensor. However, it will also be shown that *a wide VFOV can be difficult to achieve* because it aggravates the **throughput problem**.




Figure 23- Stabilization Problem

4.2 Nondimensional Sweep Rate

The **imaging density** σ_I is the average number of images that fall on any patch of terrain. It is also the inverse of the **throughput ratio**. The expression of the sweep rate rule can be written in terms of the imaging density as:

$$\frac{\Delta R}{T_{\text{cyc}}} = \Delta R f_{\text{images}} = \sigma_I V$$

Reusing the result for vertical field of view from the pointing rules:

$$\frac{\text{VFOV} (R_{\text{max}} R_{\text{min}})}{h} f_{\text{images}} = \sigma_I V$$

Which can be written in terms of nondimensionals as:

$$\frac{\text{VFOV} \times R_{\text{max}}}{\left(\frac{h}{R_{\text{max}}}\right) \left(\frac{R_{\text{max}}}{R_{\text{min}}}\right)} f_{\text{images}} = \sigma_I V$$

The **sweep rate** is therefore given by:

$$\dot{\theta} = \text{VFOV} \times f_{\text{images}} = \frac{\sigma_I V h}{R_{\text{max}} R_{\text{min}}}$$

So, for a fixed angular field of view, the rule says that a constant imaging density is achieved by modulating the sweep rate by roughly the inverse of the square of the range. For any speed, there is some minimum sweep rate of the beam. Kinematic arguments were given earlier which constrain the vertical field of view based on vehicle dimensions, maneuverability, and the terrain. This rule specifies a *dynamic constraint* as well.

4.3 Adaptive Sweep

There is no fundamental requirement that an autonomous system have an arbitrarily large vertical field of view. The line scanner has been an important existence proof of this fact. It is possible to employ adaptive techniques that either physically point a narrow VFOV sensor or which computationally stabilize a wide VFOV sensor by processing a small portion of each image. This idea will be called **adaptive sweep**.

In an earlier section, **adaptive lookahead** was proposed as a mechanism for moving the position of the projection of the vertical field of view on the groundplane in order to guarantee response. While this mechanism addresses the position of the sweep, adaptive sweep addresses the width of the sweep. These two aspects of the sweep are related through the velocity.


5. Occlusion

The field of view analysis conducted so far is based on a **flat world assumption** and it is not entirely correct on rough terrain. This section investigates the relationship between vehicle configuration and the prevalence of terrain self occlusions. It turns out that merely mounting a sensor on the roof of a vehicle implies, for typical geometry, that terrain self occlusions are inevitable, and that holes cannot be detected until it is too late to react to them.

These are two aspects of the **occlusion problem**. Any system which does not deal with occlusions or which attempts to deal with hole obstacles may be designed suboptimally. There is an additional problem that the sampling problem is itself aggravated by rough terrain. There is no getting around the fact that an optical sensor cannot see through hills so a rough terrain system must be designed to live gracefully with this problem.

5.1 Hill Occlusion

A hill can also be called a **positive obstacle**. Ideally, a sensor could see behind a navigable hill at the maximum sensor range. The required sensor height can be derived from this requirement. The highest terrain gradient which is just small enough to avoid body collision is determined by the vehicle undercarriage tangent as shown below.



In order for occlusions of navigable terrain to be completely eliminated, the following condition must be met:

$$\left(\frac{h}{R}\right) = \left(\frac{c}{L/2}\right)$$


So, for complete avoidance of occlusion of navigable terrain, the ratio of sensor height to maximum range must equal or exceed half the undercarriage tangent. This will be called the **hill occlusion rule**. In order to satisfy the hill occlusion rule, a sensor must be mounted at a fantastic height which is not realistically achievable without taking very special measures. The hill occlusion rule is almost always violated. The **perception ratio** h/R can easily exceed the undercarriage tangent by a factor of three or four. Hence, *occlusions of navigable terrain are common when the terrain is rough*.

Real Numbers

For the HMMWV, the undercarriage tangent is about 1/3. For a maximum range of 50 m, this condition requires a sensor height of 16.6 m.

5.2 Hole Occlusion

A hole can also be called a **negative obstacle**. Such obstacles are particularly problematic to an autonomous vehicle. Consider a hole which is roughly the same diameter as a wheel and which is as deep as a wheel radius. Such a hole is roughly the smallest size which presents a hazard.



In order to detect that the hole was deep enough to present a hazard, the vehicle would have to wait until the hole was close enough to satisfy:

$$\frac{h}{R} = \frac{r}{2r} = \frac{1}{2}$$

This will be called the **hole occlusion rule**. While it may seem that properly placed high depression scanlines are all that is required to detect the hole, this is not the case. Recall from the pointing rules that obstacles inside the stopping distance cannot be avoided at all. It is often the case that to allocate range pixels for hole detection is to waste throughput on obstacles that cannot be avoided anyway. Hole detection *requires a separate higher speed obstacle avoidance channel*.

Holes generate range shadows at the leading edge when the map resolution is sufficiently high. Alternatively, a vehicle could adopt a policy of avoiding such shadows specifically when they are in the path of the wheels.

Real Numbers

On the HMMWV, this rule requires that the minimum range be about 5 meters. From Figure 16, it is clear that, depending on the braking reaction time, this may limit speeds unreasonably. Later results will show that the cost to throughput is extreme when the lower scanline is depressed to ranges close to the sensor height.

5.3 Occlusion Problem and Unknown Hazard Assumption

One of the fundamental problems of rough terrain navigation at any speed is the occlusion problem and little can be done about it. The occlusion problem can be mitigated somewhat by noticing that, most of the time, regions which are occluded are occluded by hazards. To accept occlusion then is to assume that occlusions are hazards - that is, to consider large unknown regions to be unnavigable by default¹⁶. This will be called the **unknown hazard assumption**. **Guaranteed safety** requires that this assumption be made.

A further mechanism for dealing with the problem is to limit the roughness of the terrain that will be navigated, or at least to design a system to assume some limit. This will be called the **benign terrain assumption**.

5.4 Lateral Occlusion

A special case of the occlusion problem arises when a large portion of an already narrow HFOV is occluded by a large object as shown in the following figure. Oftentimes, the system cannot see a clear path ahead even though there is one. This problem, the **lateral occlusion problem**, aggravates the tunnel vision problem and it is aggravated by increased system cycle time and steering dynamics.




Figure 26-Lateral Occlusion Problem

16. This assumption is far more necessary than it may appear. Without it, the system will happily drive straight off a cliff.

6. Perceptual Bandwidth

Another aspect of throughput is the raw data rate required of communications electronics and the computation required for a single image pixel. This section investigates these perceptual bandwidth requirements.

6.1 Sensor Flux

The **sensor flux** Ψ represents the solid angle subtended by the field of view generated per unit time. It can be written as:

$$\Psi = \text{HFOV} \times \text{IFOV} \times f_{\text{images}}$$

6.2 Sensor Throughput

The number of range pixels generated per unit time by a sensor will be called the **sensor throughput** f_{pixels} . Existing sensors can be characterized as completely nonadaptive since neither the field of view nor the angular resolution change with time. For such sensors, the sensor throughput is given by:

$$f_{\text{pixels}} = \frac{\Psi}{(\text{IFOV})^2}$$

The IFOV is the angular resolution of the sensor. A sensor for which Ψ is constant is called **constant flux**, and one for which the IFOV is constant is called **constant scan**.

6.3 Sweep Rate

The product of the vertical field of view and the frame rate is a measure of the angular velocity of the beam, and is known as the **sweep rate**:

$$\dot{\theta} = \text{VFOV} \times f_{\text{images}}$$

6.4 Processor Load

In a simple case, a range pixel must have its coordinates converted at least three times in order to place it into world coordinates. The complete transform is developed in [11]. First, the ray is converted to a sensor fixed cartesian system. This operation can be implemented via lookup tables. Each pixel is looked up and multiplied by the range to get the RHS vector of the transform. This costs 6 flops.

Next, the first RPY matrix converts from the sensor frame to the body frame. This matrix is usually fixed so its elements can be computed at initialization time. The matrix multiplication costs 12 flops for a total of 18 so far.

Finally, the elements of the second RPY matrix depend on the vehicle pose, so they must be computed in real time. Assume the trig functions are implemented via lookup table and hence cost 1 flop each, to give 6 flops. The matrix terms cost about 14 flops. The multiplication costs 12 flops. Thus the left matrix multiplication costs an additional 32 flops. The total processing of a range pixel is therefore 50 flops.

More generally, it is useful to define the **processor load** σ_P as the number of flops necessary to process a single range pixel.

$$\sigma_P = \frac{\text{flops}}{\text{pixel}}$$

Real Numbers

While the analysis suggests a value of 50 for this, actual data for an entire navigator is ten times as large. Specifically, the ERIM pulse rate is 16 KHz and navigation consumes 8 Mflops. This gives a value of 500.

Thus, the relationship between processing load and sensor throughput is:

$$f_{\text{cpu}} = f_{\text{pixels}} \times \sigma_P$$

6.5 Perceptual Software Efficiency

The inverse of the processor load is called the **perceptual software efficiency** η_S .

$$\eta_S = \frac{1}{\sigma_P}$$

6.6 Computational Bandwidth

The computational bandwidth is the number of flops required of a processor per unit time. If *the geometric transforms of perception are the only aspect of the system considered*, this quantity is related to the sensor bandwidth by the processor load:

$$f_{\text{cpu}} = \frac{f_{\text{pixels}}}{\eta_S} = \frac{1}{\eta_S} \frac{\Psi}{(\text{IFOV})^2} = \frac{1}{\eta_S} \frac{\text{HFOV} \times \text{IFOV} \times f_{\text{images}}}{(\text{IFOV})^2}$$

6.7 Communications Bandwidth

Define the **communications load** σ_C as the width of a range pixel in bits. Assuming that beam angular position information is implied by position in the scan, the bandwidth required is:

$$f_{\text{comm}} = f_{\text{pixels}} \times \sigma_C$$

At 10 bits per pixel, the required bandwidth in MHz is 1/5 the required processing rate in Mflops. Hence high bandwidth communications between the sensor and the computing engine is also a fundamental necessity of high speed range image based navigation.

PART IV:Acuity

As was shown earlier, the **acuity ratios** relate the configuration of the vehicle to the resolution required of the sensor in order to resolve obstacles. This section analyses these aspects of vehicle performance for typical vehicles.

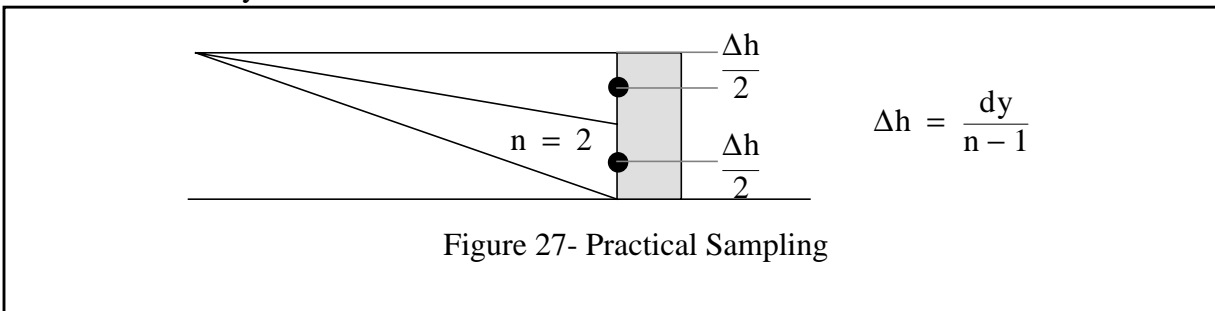
1. Acuity

The ability of a vehicle to resolve obstacles in order ensure safety is important. This section investigates the manner in which **vehicle configuration and sensor resolution together** determine the ability of a sensor to resolve obstacles. The following analysis is based on the **flat world assumption** and is not entirely correct in rough terrain. Nonetheless, it is a useful theoretical approximation.

1.1 Sampling Theorem

The acuity ratios were stated in the introduction according to a strict application of the **sampling theorem**. According to this theorem, proper recovery of a signal from its samples requires a sampling frequency at least twice that of the highest frequency of the original signal. This is a theoretical requirement and it must be interpreted carefully. Formally, **resolution** is *the smallest difference that a system can resolve*, and *not* the smallest thing that a pixel can cover.

Taking a wheel collision hazard as an example, a system with a 6 inch wheel radius requires a 3 inch vertical projection of a pixel. Such a system can only resolve a 3 inch step from a 6 inch step. Further, a 6 inch step appears as a 3 inch step if the center of the pixel is used in calculations. As shown in the figure the size of a step obstacle is underestimated by Δh which is, in turn, a larger fraction of the true height as the pixel size increases. The resolution dy is related to the height underestimate Δh by:



A practical sampling theorem requires “n” times the signal frequency for some number n greater than 2. A practical acuity requirement will require an **oversampling factor** of $\frac{n}{2}$. An oversampling factor of 3 will allow discrimination of a 5 inch step from a 6 inch step, for example. As a theoretical analysis, the report will continue to use an oversampling factor of 1.

1.2 Terrain Smoothness Assumption

The smallest feature of interest to a planning algorithm is the smallest feature which can cause the vehicle to be in an unsafe configuration. Three sizes of hazard can be identified.

From the point of view of tipover, the motion of the vehicle over the terrain amounts to a process whereby high spatial frequencies present in the terrain are filtered. The vehicle longitudinal and transverse wheelbase determine the vehicle pitch and roll. Hence, from the point of view of tipover obstacles, cells a fraction of the size of the vehicle seem sufficiently small.


Another important form of obstacle is one which presents a gradient sufficient to collide with or trap a tire at operating velocity. Potholes and steps fall into this category. Based on this reasoning, cells on the order of the size of a wheel radius are needed to ensure that a wheel does not fall in a hole or drive over a step which would cause damage.

A man made fixture such as a fence post or a natural feature such as a tree may feasibly exist in the path of the vehicle. Such features represent the extremes of terrain gradient and present the vehicle with both wheel collision and front bumper collision hazards. In the worst case, a nail in a board could feasibly exist in the vehicle's path.

Cell sizes on the order of a nail diameter are, as will be shown, infeasible for both throughput and sensor resolution reasons. It is therefore *necessary to assume that pathological cases do not exist*. A practical system must always assume that there are no man-made or natural hazards that are smaller than some practical limit. This is the **terrain smoothness assumption**.

1.3 Impact of Imaging Geometry on Acuity

Nothing could be more natural than to bolt the environmental sensor to the front of the vehicle - especially when long excursions are intended. For the present purpose, it can be considered a requirement to use such a sensor geometry. However, one important aspect of this approach to high speed autonomy is that the ratio, called the **perception ratio**¹⁷ \bar{h} :



is always a number much less than 1.0. This number shows up in many places in dimensional analysis. It can be shown that this geometric limitation of the sensor height is the root of some of the fundamental technological problems of autonomous navigation.

Real Numbers

For the HMMWV the sensor height is about 2.7 meters and up to 30 meters of lookahead is required. Hence, for this vehicle, the perception ratio is 0.1.

17. The ratio of two important lengths is often an important angle in dimensional analysis. Here, the perception ratio is the angle of incidence of the beam with the terrain.

1.4 Nomenclature

The **instantaneous field of view** is the angular width of the laser beam¹⁸ or the angular width of a camera pixel. This is also known as the **beam dispersion** for laser rangefinders. The angular coordinates of a pixel are often expressed in terms of horizontal sweep or **azimuth** ψ , and vertical sweep or **elevation** θ . When elevation is measured down from the horizon, it is called **depression**. The relationships between the beam width and its projections onto three orthogonal axes are given below. In the analysis, three orthogonal axes are considered to be oriented along the vehicle body axes of symmetry:

- x - **crossrange**, in the groundplane, normal to the direction of travel
- y - **downrange**, in the groundplane, along the direction of travel
- z - **vertical**, normal to the groundplane

1.5 Sampling Problem

The differential mapping from image space onto cartesian space is both nonlinear, and a function of the terrain geometry. The density of pixels on the groundplane can vary by three orders of magnitude, and it varies with both position and direction. Hence, the shape and density of pixels in traditional sensors is not optimal.

Significant variation in groundplane resolution can cause **undersampling** at far ranges and **oversampling** close to the vehicle. If resolution is chosen appropriately at the maximum range, pixel density at the minimum range can be extreme and cause significant waste of computational bandwidth if all pixels are processed. This is a subproblem of the **acuity problem** which has been called the **sampling problem**.


The traditional solution to the sampling problem is to use **interpolation** to fill in small unknown regions. There is a limit to how much interpolation can be performed. Fundamentally, the limit used is another aspect of the **terrain smoothness assumption**.

18. Actually, the IFOV is the angle subtended by the **receive** optics, but this distinction will be ignored here.

1.6 Differential Imaging Kinematics

Note that the following differential relationships come from projective geometry and not from differentiating the coordinate transforms. More precise relationships are available from the Jacobian of the coordinate transformation, but this depends on the precise scanning pattern chosen. In order to avoid dependence on the scanning pattern, the following approximations are used which are good enough for the purpose of the discussion here.

Approximations for the transformations from image space to the groundplane are easy to write as follows:



Consider the following approximations to these when elevation spacing $d\theta$ equals azimuth spacing $d\psi$ as is almost always the case:

$$dy = \frac{ds}{\sin \theta} \approx \frac{R d\theta}{\left(\frac{h}{R} \right)} \quad dx = dz = R d\theta$$

These approximations will be used extensively throughout the report.

1.7 Pixel Footprint Area and Density Nonuniformity

Multiplying the above expressions:

$$dx dy = R d\theta \frac{R d\theta}{\left(\frac{h}{R}\right)} = \frac{R^2 d\theta^2}{\left(\frac{h}{R}\right)}$$

Hence, *the area of a pixel when projected onto the ground plane is a constant times the cube of the range*. Due to the projection onto the groundplane¹⁹, it is increased by the inverse of the perception ratio over what would be expected based on the area of an expanding wavefront. This is the *variation of pixel size with position*.

It is typical to allocate the sensor vertical field of view such that $R_{\min} \approx 2h$. Under this configuration, the density variation over the entire field of view is given by:

$$\left(\frac{R_{\max}}{R_{\min}}\right)^3 = \frac{1}{8} \left(\frac{R_{\max}}{h}\right)^3$$

The **range ratio** is an important variable because it determines the degree to which groundplane resolution is wasted in the current generation of sensors.

Real Numbers

For the HMMWV using the first ambiguity interval of the ERIM sensor, this density variation is 50. So if data were of appropriate density at the extremes of the field of view, it would be 50 times too dense near the vehicle.

19. Again from radiometry. The pixel solid angle is constant with range, so an R^2 growth would be expected. However, the projection onto the groundplane by close to 90 degrees provides the extra R .

1.8 Pixel Footprint Aspect Ratio

Dividing the above expressions:

$$\frac{dx}{dy} = \frac{dz}{dy} = \left(\frac{h}{R} \right)$$

Hence, *the pixel footprint aspect ratio at any range for which $R \gg h$ is given by the perception ratio*. The assumption $R \gg h$ comes from the approximations to the resolution transforms. These formula are extremely good approximations whenever R is larger than $2h$. For most vehicles, this condition is always true²⁰. This is the *variation of pixel size with direction*.

The impact of this is that, since the perception ratio is often close to $1/10$, terrain is often oversampled in the crossrange direction by an order of magnitude with respect to the downrange direction. Practically, then, only one *column* in ten is needed to ensure adequate coverage of the environment if square pixels are used and downrange resolution is adequate. More importantly, a sensor designed specifically for high speed autonomy would have significantly non square pixels. This arises fundamentally because stopping distance exceeds sensor height by a large factor.

Real Numbers

On the HMMWV, at the 5 meter minimum range given by 30 degrees depression, it is 2:1. At 27 meters, this is 10:1, at 50 meters, it is just under 20:1.

20. Perhaps the most important distinction of high speed autonomy is the fact that sensor height is an order of magnitude smaller than the vehicle braking distance. This fact has many implications.

1.9 Minimum Sensor Acuity in Image Space

The nonuniformity of pixel spacing is quite extreme when the perception ratio is close to 0.1 and the scanline depression angle approaches 30 degrees. Consider now what happens when the spacing between pixels begins to approach the size of the vehicle itself. At far ranges dy is much larger than dx . So it is dy which will first approach L .

Real Numbers

This is no mere abstraction. At 50 meters range, for the ERIM sensor, the spot spacing is an incredible 10 meters or three times the size of the vehicle.

Consider that since computation of the vehicle pitch angle depends on having two different elevations under the front and rear wheels, the pixel spacing dy must be no larger than one half the wheelbase for this to be practical. Beyond the range at which this occurs, sensor data contains no useful information at all, unless data from several images can be registered accurately enough to contain useful information.

Equating dy to one half the wheelbase:

$$dy = \frac{L}{2} = R d\theta / \left(\frac{h}{R}\right)$$

Rewriting gives a very elegant expression:

$$\left(\frac{L}{R}\right) \left(\frac{h}{R}\right) = 2 d\theta$$

This is a very interesting relationship which says that at some resolution, the sensor measures nothing useful, or equivalently, the holes in the map are almost large enough to swallow the whole vehicle. This resolution occurs when the product of the normalized wheelbase and the perception ratio equals one half the angular resolution of the sensor. This is another expression of the **minimum sensor acuity rule**. Any of the variables can be considered to be absolutely limited by the others in the expression. The equation relates two key nondimensional variables and connects the vehicle shape to the required sensor angular resolution.

Two conclusions can be immediately drawn from this relationship. First, the IFOV itself limits the maximum range and therefore the maximum speed of the vehicle. Second, for any given IFOV, there exists a vehicle speed beyond which sufficiently accurate map registration becomes *essential* to reliable hazard detection.

1.10 Maximum Sensor Acuity in Image Space

It is possible to formulate a similar rule by considering the much more stringent requirements of resolving a wheel collision hazard at the maximum range. In order to resolve a wheel collision hazard, spatial resolution in the vertical direction must be sufficient to land, say, 2 pixels on a vertical surface at any given range. A collision hazard for a given wheel radius r is a step of the same order as the wheel radius. Therefore, resolving a collision hazard requires:

$$dz = \frac{r}{2}$$

which can be rewritten in terms of the approximations to the forward resolution transforms thus:

$$Rd\theta = \frac{r}{2}$$

Solving for the pixel angular resolution:

$$d\theta = \frac{1}{2} \left(\frac{r}{R} \right)$$

This is another expression of the **maximum sensor acuity rule**²¹. Pixel sizes below this limit are excessive. Pixel angular resolution must be no larger than one half the normalized wheel radius in order to resolve a collision hazard at a given range. This can be rewritten in several ways:

$$d\theta_{\min} = \frac{1}{2} \left(\frac{r}{R} \right) = \frac{1}{2} \left(\frac{r}{h} \right) \left(\frac{h}{R} \right) = \frac{1}{2} \left(\frac{r}{L} \right) \left(\frac{L}{h} \right) \left(\frac{h}{R} \right)$$

Notice that maximum acuity is related to wheel size and lookahead. Since lookahead increases as the wheelbase increases, maximum acuity is related to the **wheel fraction**. On this basis, an all wheel vehicle configuration (where the wheels are large compared to the body size) requires the lowest perceptual resolutions. It is often necessary to violate this rule and adopt assumptions about the terrain in order to lower throughput requirements to a practical level. Notice that the minimum rule is quadratic in $1/R$, whereas the maximum rule is linear. Both constraints are equal when:

$$R = \frac{Lh}{r}$$

Real Numbers

For the HMMWV, using a wheelbase of 3.3 meters, a sensor height of 2.7 meters, and a wheel radius of 0.475 meters, this range is 18.8 meters. This implies that minimum acuity is the more stringent requirement beyond this range.

21. Other acuity rules can be defined by relating any resolution requirement to some vehicle dimension, and substituting the resolution transforms. For instance, if downrange resolution must be less than a wheel radius, set $dy = r$. Then a rule is generated which is similar to the minimum acuity rule, but more stringent.

1.11 Kinematic Maximum Range and the Myopia Problem

The minimum sensor acuity rule gives a kinematic limit on maximum sensor range which says that at this range, the holes in the map are too large. This limit is surprisingly close. Solving for R:

$$R = \sqrt{\frac{hL}{2d\theta}}$$

Real Numbers

Using the HMMW vehicle wheelbase of 3.3 m, the sensor height of 2.7 m, and the ERIM beam dispersion of 10 mrad, this gives a remarkable 21.1 meters. The conclusion is that the information at the extremes of the first ambiguity interval of the ERIM sensor is barely useful unless some mechanism is available to register subsequent images accurately enough in pitch.

1.12 Maximum Angular Resolution

The minimum sensor acuity rule also gives a limit on sensor angular resolution. For a given nondimensional specification of the vehicle, the largest feasible sensor angular resolution is given by:

$$d\theta_{\max} = \frac{1}{2} \left(\frac{L}{R} \right) \left(\frac{h}{R} \right)$$

The minimum acuity rule takes any of these forms:

$$\left(\frac{L}{R} \right) \left(\frac{h}{R} \right) = 2d\theta_{\max} \quad \left(\frac{L}{R} \right)^2 \left(\frac{h}{L} \right) = 2d\theta_{\max} \quad \left(\frac{L}{h} \right) \left(\frac{h}{R} \right)^2 = 2d\theta_{\max}$$

and it can be rewritten in terms of any of the other nondimensionals describing the vehicle shape.

Notice that all of these results are *independent of h or R itself*, and *independent of the scanning pattern* of the sensor, except that uniform pixel angular spacing is assumed. All of these results are intrinsic when $h \ll R$. The issue which causes nonuniform sampling is the perception ratio: it has very little to do with the sensor itself.

1.13 Acuity Problem

For contemporary vehicles, the myopia problem and the acuity problem are linked because poor angular resolution is the typical limit on the useful range of a sensor. The above analysis is based on the **flat world assumption**. On rough terrain, there is no practical way to guarantee adequate acuity over the field of view because there will always be situations where pixels have glancing incidence to the terrain.

2. Positioning Bandwidth

Another aspect of the acuity requirement is the rate at which the vehicle position is sampled. This section investigates the interaction between vehicle speed and the bandwidth required of position estimation. Although the position of the ray through each pixel is known in body coordinates and can be accounted for, the entire image field of view sweeps over the terrain as the vehicle moves. It is necessary to localize a range pixel to within the map resolution in order to meet the maximum fidelity requirement. The smearing of the environmental model which arises from incorrect models of vehicle motion during the digitization process is called the **motion distortion problem**. The following analysis can be conducted for translational motion as well.

2.1 Heading and Positioning Bandwidth

For a turn of radius ρ , and a velocity of V , the yaw rate of the vehicle is given by:

$$\dot{\psi} = \kappa V = \frac{V}{\rho}$$

If the range pixel returns a value of R , then the velocity of the pixel due solely to this vehicle motion is given by:

$$V_{\text{pixel}} = \frac{RV}{\rho}$$

At this speed, the pixel moves through a distance equal to the map resolution over a time period of:

$$\Delta t = \frac{\delta}{V_{\text{pixel}}} = \frac{\delta \rho}{RV}$$

Therefore, the vehicle yaw must be sampled at a rate of:

$$f_{\text{heading}} = \frac{RV}{\delta \rho}$$

If this positioning bandwidth is not met, the terrain map will smear unacceptably and hazards will be localized incorrectly due to poor resolution in the computations. In practice, interpolation can be used to supply this position estimate provided the sampling of the underlying heading signal satisfies the sampling theorem. This requirement implies that *every scanline* in a rangefinder image must use an independent vehicle pose estimate.

Real Numbers

Using the HMMW with a 30 meter lookahead, a velocity of 4.5 m/s, a turn radius of 7.5 meters, and a map resolution of 1/3 meter, this gives a frequency of 60 Hz.

2.2 Attitude and Positioning Bandwidth

In the case of attitude, the vehicle pitch will be considered. Let the vehicle pitch as it encounters a patch of rough terrain. The front wheel will lift by a small amount dz and the pitch of the vehicle will change by a small amount given by:

$$d\theta = \frac{dz}{L}$$

The pitch rate is therefore related to the terrain gradient in the forward direction as follows:

$$\dot{\theta} = \frac{1}{L} \frac{dz}{dt} = \frac{1}{L} \frac{dz}{dy} \frac{dy}{dt} = \left(\frac{dz}{dy} \right) \frac{V}{L}$$

For a terrain gradient of 1, a velocity of 3 m/s and a wheelbase of 3 meters, this figure is 1 rad/sec. This is unrealistically high because the suspension will attenuate this figure by a large factor. Let the maximum *instantaneous* pitch rate be called $\dot{\theta}_{\max}$. Recall from the acuity analysis earlier that the differential relationship between the downrange position of a pixel and the elevation angle of the pixel is as follows:

$$dy_{\text{pixel}} = \frac{R^2}{h} d\theta$$

Therefore, the velocity of the pixel in the downrange direction due solely to this vehicle motion is:

$$V_{\text{pixel}} = \frac{dy_{\text{pixel}}}{dt} = \frac{dy_{\text{pixel}}}{d\theta} \frac{d\theta}{dt} = \frac{R^2}{h} \dot{\theta}_{\max}$$

At this speed, the pixel moves through a distance equal to the map resolution over a time period of:

$$\Delta t = \frac{\delta}{V_{\text{pixel}}} = \frac{\delta h}{R^2 \dot{\theta}_{\max}}$$

Therefore, the vehicle pitch must be sampled at a rate of:

$$f_{\text{attitude}} = \frac{R^2 \dot{\theta}_{\max}}{\delta h}$$

If this positioning bandwidth is not met, the terrain map will smear unacceptably and hazards will be localized incorrectly due to poor resolution in the computations. In practice, interpolation can be used to supply this position estimate provided the sampling of the underlying attitude signal satisfies the sampling theorem. This requirement often implies that *every pixel* in a rangefinder image must use an independent vehicle pose estimate.

Real Numbers

Using the HMMW with a 30 meter lookahead, a maximum instantaneous pitch rate of 1/8 rad/sec, a sensor height of 2.7 meters, and a map resolution of 1/3 meter, this gives a frequency of 150 Hz.

2.3 Motion Distortion Problem

Although the last analysis made some assumptions about the instantaneous vehicle pitch rate, the assumption is not too important in practice. In the most general case, the motion distortion problem has no solution because it also depends on the terrain gradient. As a vehicle crests a hill, for example, there will come a time when a single pixel will sweep several meters forward in a small fraction of a second. This can be seen, for example, in Figure 23.

3. Geometric Efficiency

This section investigates the impact of typical imaging geometry on the efficiency with which a system can sense its environment at appropriate resolution.

3.1 Imaging Geometry

For typical imaging geometry, the density of range pixels on the groundplane varies significantly with both position and direction. Consider the following graph of the groundplane positions of each pixel from a range image for a typical rangefinder. A 10 m grid is superimposed for reference purposes and every fifth pixel is shown in azimuth to avoid clutter. The density variation is more extreme than can be shown on paper. The variation of density with range is obvious




Figure 30- Typical Scanning Pattern

The spot pattern is obtained by intersecting the ground plane with each pixel. Spots are shown below for the same sensor at one tenth density in elevation and one seventh density in azimuth. The variation in aspect ratio, as well as size, is apparent.




Figure 31- Typical Spot Pattern

3.2 Scanning Density

In a single range image, the average number of range pixels that fall into a particular map cell is not necessarily one. A quantity called the **scanning density** σ_s can be defined which is the average number of range pixels per map cell over the entire field of view of a single image. The scanning density depends on the scanning pattern of the sensor as well as its position and orientation on the vehicle.

Real Numbers

For the ERIM scanner mounted on the HMMWV at 2.7 meters height and 16.5 degrees tilt, using a map resolution of 1/6 m, the scanning density is 2.0. However, this figure is the average value so it does not reflect the considerable variation in density over the field of view. Recall from an earlier section that the maximum variation in density is the cube of the range ratio - or about 50.

The scanning density is purely a geometric matter. It is easy to compute as follows:

$$\sigma_s = \frac{\text{pixels}}{\text{cells}} = \left(\frac{\text{HFOV} \times \text{VFOV}}{\text{IFOV} \times \text{IFOV}} \right) / (\text{AREA} / \delta^2) \quad \text{AREA} = \frac{\text{HFOV}}{2\pi} \pi [R_{\max}^2 - R_{\min}^2]$$

$$\sigma_s = \frac{\text{HFOV} \times \text{VFOV}}{\text{AREA}} \times \frac{\delta^2}{\text{IFOV}^2}$$

$$\sigma_s = \frac{2 \times \text{VFOV}}{[R_{\max}^2 - R_{\min}^2]} \times \frac{\delta^2}{\text{IFOV}^2} = \frac{2h \frac{(R_{\max} - R_{\min})}{R_{\max} R_{\min}}}{[R_{\max} - R_{\min}] [R_{\max} + R_{\min}]} \times \frac{\delta^2}{\text{IFOV}^2}$$

$$\sigma_s = \frac{2h}{(R_{\max} R_{\min}) [R_{\max} + R_{\min}]} \times \frac{\delta^2}{\text{IFOV}^2}$$

$$\sigma_s = 2 \left(\frac{h}{R_{\max}} \right) \left(\frac{1}{\frac{R_{\max}}{R_{\min}} + 1} \right) \left(\frac{\delta}{R_{\min}} \right)^2 \left(\frac{1}{\text{IFOV}} \right)^2$$

$$\sigma_s = 2 \bar{h}_{\min} \left(\frac{1}{R + 1} \right) \left(\frac{\delta}{R_{\min}} \right)^2 \left(\frac{1}{\text{IFOV}} \right)^2$$




Figure 32- Scanning Density

For a given map resolution and IFOV, the scanning density depends only on the perception ratio and the range ratio and a new nondimensional ratio of map resolution to minimum range. Let this new variable be called the **normalized map resolution**.

$$\bar{\delta} = \frac{\delta}{R}$$

The significance of the extra variable is that it is necessary to specify where the field of view is pointed in order to compute the area that it covers.

The scanning density is plotted below for a sensor height of 2.7 m, a map resolution of 1/6 m, and various values of the pixel size. The growth with IFOV is quadratic and the growth with decreased minimum range is also quadratic. Pixel sizes were chosen in the range necessary to identify step obstacles at the maximum range²². Clearly, for any choice of parameters, there is a minimum range below which scanning density begins to increase very rapidly and significant computational bandwidth is wasted.




Figure 33- Scanning Density

22. The situation is far worse if the beam dispersion is chosen to keep the spot size within the map resolution.

3.3 Imaging Density

It may be necessary to image each map cell more than once for several reasons:

- to ensure robust obstacle detection
- to track moving objects
- to track stationary landmarks for position estimation purposes

A quantity called the **imaging density** σ_I can be defined which is the number of images that include a map cell in the field of view, averaged over all map cells. This quantity is independent of the scanning density by definition.

Real Numbers

For the ERIM scanner mounted on the HMMWV at 2.7 meters height and 16.5 degrees tilt, at 5 m/s speed, the imaging density is about 6. So the vehicle sees an obstacle 6 times before it leaves the field of view of the sensor completely.

The imaging density depends on vehicle speed as well as the sensor configuration. A simple expression for it is available by considering linear trajectories of a given velocity. The vertical field of view of the sensor projects onto the groundplane to give a certain length in the downrange direction. Over a time period of the sensor frame period T_{image} (the inverse of the frame rate), the vehicle moves a distance given by its velocity. The governing relationship is given below:

$$\sigma_I = \frac{\Delta R}{V T_{\text{image}}}$$

but this is just the inverse of the **throughput ratio** as would be expected. The imaging density is plotted below versus speed for a sensor with a 0.4 sec frame period for some values of the vertical field of view.




Figure 34- Imaging Density

3.4 Geometric Efficiency

The scanning density and imaging density together give the average total number of image pixels that fall on a patch of terrain over several images. Putting these factors together, a **geometric efficiency** η_G can be defined as follows:

$$\eta_G = \frac{1}{\sigma_S \sigma_I}$$

In order to demonstrate the importance of the geometric efficiency, it is plotted below for the HMMWV using the ERIM rangefinder as currently configured. The configuration is maximum range 20 m, minimum range 4.5 m, beam dispersion 10 mrad, map resolution 1/3 m, sensor height 2.7 m:




Figure 35- Geometric Efficiency

Thus, at an operating speed of 3 meters/second²³, the **geometric efficiency** of the ERIM is a mere 2.5%. Equivalently, *roughly 98 pixels out of 100 are waste, providing no new information.*

23. This is the typical operating speed of contemporary off-road vehicles.

3.5 Adaptive Scan

Two techniques are available to alleviate poor geometric efficiency. **Adaptive sweep** was proposed earlier as a mechanism for minimizing the imaging density. Typical geometry implies that the area density of range pixels becomes more or less constant when this is done because the range ratio is fairly small. However, the shape of those uniformly distributed pixels is still very elongated. The mechanism for alleviating this problem, associated with the scanning density, is to elongate the pixel in the reverse direction in the image. This will be called **adaptive scan**. For contemporary sensors, adaptive scan can be implemented by filtering or subsampling the image in the azimuth direction. Ideally, a sensor would have nonsquare hardware pixels.

3.6 Acuity Nondimensionals

From the point of view of dimensional analysis, it is not so much the height of the sensor as the ratio of that height to the maximum sensor range that matters. This quantity was defined earlier as the **perception ratio**:

$$\bar{h} \equiv \frac{h}{R}$$

The quantity gives quick rules of thumb for computing laser spot aspect ratio and area. Let the ratio of maximum to minimum range be called the **range ratio**:

$$\bar{R} = \frac{R_{\max}}{R_{\min}}$$

Similarly, the **normalized range difference** is:

$$\overline{\Delta R} = \frac{R_{\max} - R_{\min}}{R_{\min}} = \frac{\Delta R}{R_{\min}}$$

These are obviously related by:

$$\overline{\Delta R} = \bar{R} + 1 \qquad \bar{R} = \overline{\Delta R} - 1$$

Recall that the quantity

$$\eta_G = \frac{1}{\sigma_S \sigma_I}$$

is called the **geometric efficiency** and it expresses the average number of range pixels in a map cell when it is consumed by driving over it. The **scanning density** σ_S is the average density of pixels per map cell over a single entire image. The **imaging density** σ_I is the average number of images that fall on a patch of terrain.

3.7 Summary of the Acuity Nondimensionals

The identified pi products for acuity are summarized below:

Table 4: Perception Nondimensionals

Symbol	Name	Expression	Symbol	Name	Expression
\bar{h}	Perception Ratio	$\frac{h}{R}$	\bar{R}	Range Ratio	$\frac{R_{\max}}{R_{\min}}$
σ_I	Imaging Density	numeric	Scanning Density	σ_S	numeric
$\overline{\Delta R}$	Normalized Range Difference	$\frac{\Delta R}{R_{\min}}$			

PART V:Fidelity

As was shown earlier, the **fidelity ratios** relate the configuration of the vehicle, the accuracy of its geometric models, and the density of hazards. This section analyses these aspects of vehicle performance for typical vehicles.

1. Delays

A software system is, of course, a disembodied entity which relies on sensors and effectors to connect it to, and permit it to influence its environment in robotics applications. In the context of high speed motion, the time it takes to pass information into and out of the system becomes a significant factor. As a moving entity, any delays in time which are not accounted for are ultimately reflected as errors between both:

- what is sensed and reality, or
- what is commanded and reality

Time delays may arise in general from:

- sensor dwell
- communication
- processing
- plant dynamics

and all of these types of delay occur in a contemporary autonomous system.

While delays affect response directly, they also affect the ability of the system to localize obstacles correctly. Delays themselves are not so much a problem as are *unmodelled delays*. Hence, one of the impacts of high speed on the fidelity requirement is the need for high fidelity temporal models. This implies several things:

- significant events must be time stamped
- significant delays must be modelled
- i/o which is not correlated in time must be buffered

This section investigates these matters in the context of high speed motion.

1.1 Latency Problem

Unmodelled latencies in both sensors and actuators cause the vehicle to both underestimate the distance to an obstacle and underestimate the distance required to react. This is indicated in the following figure.




Figure 36- Latency Problem

1.2 Minimum Significant Delay

Let a time delay of Δt occur which is not modelled by the system. If the vehicle travels at a speed V then the distance travelled is, naturally, $V\Delta t$. In order to guarantee correct localization of a range pixel, the **maximum fidelity** requirement must be met. The impact of any delay on the fidelity requirement can be expressed in terms of the **maximum fidelity ratio** thus:

$$\rho_{dx} = \frac{V\Delta t}{\delta}$$

and thus the **minimum significant delay** occurs when the ratio is one, or when:

$$\Delta t = \frac{\delta}{V}$$

Real Numbers

At a mere 5 m/s, using a map resolution of 1/3 meter, this gives less than one tenth of a second whereas the delays of real systems typically exceed this by a large factor.

1.3 Dynamic Systems

A matter similar to delay is the time constant of dynamic systems. Oftentimes, a control algorithm actuates a derivative of the variable of interest and there is often a limit on the magnitude of the derivative that can be commanded. The net effect is a significant delay in the transfer function. In general, any system which commands a derivative of the controlled variable satisfies a differential equation at least as complex as:

$$\tau \frac{dy_{\text{response}}}{dt} + y_{\text{response}} = y_{\text{command}}$$

which is the equation of a **filter** and where τ is called the **time constant**.

Typically, the position and attitude of the vehicle are the variables of interest. For the Ackerman steer vehicle, the input commands are speed (the derivative of arc length) and curvature (the derivative of heading wrt arc length) so *the Ackerman steering system is fundamentally a filter*.

1.4 Characteristic Times and Low Latency Assumption

The **characteristic time** of any element is the total delay, whatever its source, which relates the input to the associated correct, steady state output. The total characteristic time of all information processing elements, hardware or software, and all energy transformation elements is the quantity which matters, so it is not correct to discount delays individually. To assume that delays are irrelevant is to assume that the characteristic time is zero. This **low latency assumption** is not correct for high speed autonomy above some speed.

1.5 Transience in Turning

When a vehicle executes a reverse turn, the actuator response can be divided into a transient portion and a steady state portion as shown in the following figure. During the transient portion the steering mechanism is moving to its commanded position at a constant rate. This portion of the curve in the groundplane is a **clothoid**. During the steady state portion, the curvature is constant, and the curve is a circular arc.




Figure 37- Transience in the Reverse Turn

1.6 Heading Response

If the mechanism actuates curvature more or less directly, then the heading response curve is the direct integral of the steering mechanism position at constant velocity because yaw rate is given by:

$$\dot{\psi} = \kappa V \quad \psi(t) = \psi_0 + V \int_0^{T_{act}} \kappa(t) dt$$

This implies that the heading will grow quadratically, reach a maximum and descend back to zero exactly as the steering mechanism reaches its goal because the area under the curvature signal is zero as shown below:




Figure 38- Transient Steering Response

1.7 Nondimensional Transient Turning

Recall that the actuator reaction time is given by:

$$T_{\text{act}} = \frac{\Delta\alpha}{\dot{\alpha}_{\text{max}}} = \frac{2\alpha_{\text{max}}}{\dot{\alpha}_{\text{max}}}$$

Real Numbers

For the HMMWV, it takes 3 seconds to turn the front wheels through their 60 degrees of travel.

The time required to turn through an angle $\Delta\psi$ at constant curvature is:

$$T_{\text{turn}} = \frac{\Delta\psi}{\dot{\psi}_{\text{max}}} = \frac{\Delta\psi}{\kappa_{\text{max}} V} = \frac{\Delta\psi \rho_{\text{min}}}{V}$$

Thus, a **transient turning coefficient** can be defined as the ratio of these two:

$$\bar{t}_t = \frac{2\alpha_{\text{max}}}{\dot{\alpha}_{\text{max}}} \frac{\Delta\psi}{\psi} = \frac{2\alpha_{\text{max}} V}{\dot{\alpha}_{\text{max}} \Delta\psi \rho_{\text{min}}} = \frac{T_{\text{act}} V}{\Delta\psi \rho_{\text{min}}}$$

This nondimensional is a measure of the importance of turning dynamics in a sharp turn. When it exceeds, say 0.1, it becomes important to explicitly consider turning dynamics. Note that the number increases for smaller constant curvature turns.

Real Numbers

For the HMMWV, this number is 1.27 at 5 m/s for a 90 degree turn. Hence, it is imperative to consider turning dynamics.

Another important aspect of the high curvature turn at speed is the raw error involved in assuming instantaneous response from the steering actuators. The difference between the two models is illustrated in the previous figure. The length of this vector can be approximated by.

$$s_{\text{error}} = T_{\text{act}} V$$

Real Numbers

For the HMMWV, the model error of the ideal arc response is 15 meters at 5 m/s. Hence, the use of “arc” models of steering are fundamentally wrong at surprisingly moderate speeds.

2. Dynamics Feedforward

For some actuators, such as steering, response cannot be modelled as a kinematic function of the input command. These systems require differential equation models. The solution of such equations for the purpose of predicting response is called feedforward. Feedforward is investigated in this section.

2.1 Dynamics Feedforward

Delays affect feedback control systems in a special way. In effect, a delay amounts to a pathological disturbance, active during transients, which always ensures that a system fails to react instantaneously to its inputs. In high speed autonomy, the issue has special significance. The control of vehicle speed and curvature for traditional steering mechanisms causes massive errors in feedback controllers. In the face of such dynamics, two general conclusions can be made.

- the system fundamentally cannot follow commands given to it above some frequency
- error signals formed directly on the controlled variable will cause control loop instability

The traditional technique used control highly latent systems is **feedforward**. The philosophy of feedforward is to:

- accept that the system cannot be configured as a high fidelity command follower
- form an error signal over *system response* and not commands by predicting the response to commands and comparing this to the feedback. In general, a feedforward loop would look like the following block diagram:




Figure 39- Feedforward

where the system box may include lower levels of control as well as the real system.

2.2 Characteristic Times and Loop Bandwidth

While large delays cause problems that require feedforward, they are somewhat beneficial with respect to reducing the bandwidth requirements of control loops. Nyquist's sampling theorem of course says that sampling frequencies above twice the highest frequency in the sampled signal provide no new information. In the context of the numerical solution of a differential equation, or in the context of a digital feedback control law, there is little to be gained by employing a high bandwidth servo to control a highly latent system, or at least, there are rapidly diminishing returns available as bandwidth increases.

The system will respond no faster than its time constant, related to the actuator power limit, predetermines. Increasingly higher loop bandwidth reaches an eventual point of diminishing returns in a throughput limited system where the cycles can be put to better use elsewhere²⁴.

2.3 Impact on Planner/Controller Hierarchy

The overall impact of these issues on the problem is that the hierarchical view of planners commanding paths from controllers through some high level interface will not work above some critical speed. The system will be unstable, and command following fidelity will fail to meet the fidelity requirements of the problem. Collisions will inevitably occur. In high speed work, the emphasis must shift from the traditional AI hierarchy view, to the tightly coupled real time control view.

2.4 Impact on Trajectory Generator/ Trajectory Tracker Hierarchy

A high bandwidth path tracking algorithm is not justified if the bandwidth significantly exceeds the natural frequency of the system. Even if feedforward is introduced into a tracker, the fidelity of the feedforward transfer function is far more important from a fidelity standpoint than the loop bandwidth.

The steering column and the body dynamics of a large vehicle will attenuate any high frequency control inputs that are given to them. Therefore, it is not necessary to control these with high bandwidth servos, and any attempt to do so is a waste of computing resources.

24. Consider the solution to the first order filter equation presented earlier, for time steps below about 10% of the time constant, the solutions are indistinguishable from a relative error standpoint.

3. Positioning Fidelity

In order to correctly localize hazards, the entire system geometric model must be sufficiently accurate. This implies that the vehicle location sensors and the perception sensors have accuracy requirements. This section investigates the first of these.

3.1 Absolute Attitude Accuracy Requirement

In order to compute static stability, a system fundamentally needs to know the projection of the body vertical axis onto the gravity vector, and this generates the need for:

- a sensor to measure 3D attitude at time t
- an ***absolute angular accuracy*** requirement on that sensor
- a 3D model of vehicle attitude and terrain following which allows the system to predict the presence of the hazard along candidate trajectories

A rough terrain system must both measure in 3D and think in 3D. As a rough rule of thumb, an absolute accuracy on pitch and roll indications which is an order of magnitude less than the minimum angle of rotation which brings the center of gravity outside the horizontal support polygon of the wheels is necessary. This arises again from the Nyquist sampling theorem and a requirement to distinguish hazard from non hazard as an angular proximity measurement. It can be called the **attitude acuity rule**.

3.2 Rigid Terrain Assumption

In order to predict tipover before it occurs, the system requires a model which maps the vehicle position in the terrain map onto an attitude. This requires some assumptions about the ability of the terrain to sustain compression and shear loads as well as a model of the vehicle suspension (which may be a trivial, rigid model). If the terrain is assumed to always provide the support loads required, without deformation, a **rigid terrain assumption** is being adopted.

3.3 Rigid Suspension Assumption

The suspension model allows the system to associate a unique vehicle attitude which each position, and the computation of any volumetric intersection between the vehicle body and the terrain. If it is assumed not to deflect at all, a **rigid suspension assumption** is being adopted.


3.4 Image Registration Problem

A well known problem in outdoor range image based navigation is a special case of the **fidelity problem** known as the **image registration problem**. A fidelity requirement on the measurement of environmental geometry is that the relative error in elevation for the same spot on the ground between two consecutive images not exceed the acuity requirement. Otherwise, artifacts of systematic sensor error will appear as phantom obstacles in the terrain map. There are many potential sources for such systematic errors including:

- systematic range sensor errors
- range sensor miscalibration²⁵
- vehicle kinematic miscalibration
- systematic position and attitude sensor errors
- position and attitude sensor miscalibration

3.5 Linear Relative Accuracy Requirement

From the point of view of the impact of the fidelity requirement on linear position estimates, the relative accuracy required is related to the vehicle vertical error excursion between images as shown in the following figure:



Thus, it requires only 12% **relative** position accuracy to avoid an artifact that is within the maximum acuity limit at 5 m/s.

Hence the **image registration problem** ought not to exist at all (on this basis) unless there are gross systematic perception errors, overall miscalibration errors, or some inherent geometric sensitivity.

3.6 Angular Relative Accuracy Requirement

The relative angular accuracy follows directly from the acuity limit itself. The relative angular accuracy required of a position estimation system is on the order of the angle subtended by a pixel. These last two requirements are well within the performance specifications of available hardware.

25. Cartesian bias is the only form of miscalibration which can be ignored in a relative error model. Even small angular or range errors generate nonconstant errors in elevation over the sensor field of view.

4. Perceptual Fidelity

Correct localization of hazards also depends on the accuracy of the range image itself. This section investigates the effect of imaging geometry on the accuracy of a sensor.

4.1 Incidence Sensitivity Problem

The small incidence angle associated with large lookahead distance creates a severe sensitivity problem in the **localization** of obstacles. Consider again the downrange projective differential imaging kinematics expressions.

$$dy = \frac{R^2 d\theta}{h} \quad \therefore \frac{dy}{d\theta} = \frac{R^2}{h}$$

Therefore range resolution grows with the square of range for typical imaging geometry. This fact is related to the minimum acuity requirement. The footprint of a range pixel is elongated in the downrange direction, and the returned range value will have a large random component when the range itself is large.

Real Numbers

For the HMMWV, a 30 meter range pixel of 10 mrad width is elongated to 3.3 meters. Range measurement noise can be expected to be some fraction of this.

4.2 Attitude Sensitivity Problem

Notice that the vertical error excursion of a pixel due to a small error in vehicle pitch indications is linear in the measured range because:

$$dz = R d\theta$$

Unless the relative angular attitude accuracy is sufficiently high and the positioner bandwidth is sufficiently high, phantom obstacles will be generated. Another source of these errors is relative motion between the vehicle body and the sensor. Sensor shock mounts should be designed to remove high frequencies, particularly in pitch, from this relative motion.

PART VI:Results

The fundamental requirements of timing, speed, resolution, and accuracy are largely independent in concept but they become related as soon as it is insisted that they be met *simultaneously*. This section will investigate the interactions of requirements at a more fundamental level. In doing so, it will establish the need for:

- a feedforward control approach to trajectory generation and tracking
- an adaptive approach to environmental perception
- a real-time approach to obstacle detection


1. Effect of Response on Fidelity - Rationale for Feedforward Approach

An earlier section presented an analysis of the relative importance of computational reaction time and vehicle maneuverability on the **response ratio**. In that analysis, actuators were considered to respond instantaneously and perfectly to an input command - after some time delay had elapsed. While this is a useful theoretical approximation, and while it is a good model of braking, the same is not true of turning. Steering dynamics can only be modelled correctly by a differential equation. This section presents an accurate steering model for an Ackerman steer vehicle.

While this section is written specifically for the Ackerman steer vehicle, many of the conclusions apply in general because high speeds and rollover hazards limit the curvatures that a vehicle can safely sustain. There are three principle problems associated with the Ackerman steering mechanism. The first is steering dynamics, a **latency problem**, and the second is the nonintegrability of constraints, a **nonholonomy problem**. The third is the fact that the conversion from steer angle to curvature is *nonlinear*. A further problem is that attitude rate is *coupled* to both speed and steer angle. This section discusses these issues and their impact.

1.1 The Bicycle Model

It is useful to approximate the kinematics of the steering mechanism by assuming that the two front wheels turn slightly differentially so that the instantaneous center of rotation can be determined purely by kinematic means. This amounts to assuming that the steering mechanism is the same as that of a bicycle. Let the angular velocity vector directed along the body z axis be called $\dot{\beta}$. Using the bicycle model approximation, the path curvature κ , radius of curvature ρ , and steer angle α are related by the wheelbase L .



Where $\tan \alpha$ denotes the tangent of α . Rotation rate is obtained from the speed V as:

$$\dot{\beta} = \frac{d\beta}{ds} \frac{ds}{dt} = \kappa V = \frac{V \tan \alpha}{L}$$

The steer angle α is an indirect measurement of the ratio of $\dot{\beta}$ to velocity through the measurement function:

$$\alpha = \tan^{-1} \left(\frac{L \dot{\beta}}{V} \right) = \tan^{-1} (\kappa L)$$

Although it is common to think of these equations in kinematic terms, this is only possible when the dependence on time is avoided. In fact, this steering mechanism is modelled by a coupled nonlinear differential equation thus:

$$\frac{d\beta(t)}{dt} = \frac{1}{L} \tan [\alpha(t)] \frac{ds}{dt} = \kappa(t) \frac{ds}{dt}$$

1.2 The Fresnel Integrals

The **actuation space** (A space) of a typical automobile is the space of curvature and speed since these are the variables that are directly controlled²⁶. The **configuration space** (C space) on the other hand is comprised of (x, y, heading) or perhaps more degrees of freedom in cartesian 3D. The mapping from A space to C space is the well known **Fresnel Integrals** which are also the equations of **deduced reckoning** in navigation. For example, the following equations map A space to C space in a flat 2D world:

$$\begin{aligned}
 x(t) &= x_0 + \int_0^t V(t) \cos(\psi(t)) dt & \frac{dx(t)}{dt} &= V(t) \cos \psi(t) \\
 y(t) &= y_0 + \int_0^t V(t) \sin(\psi(t)) dt & \frac{dy(t)}{dt} &= V(t) \sin \psi(t) \\
 \psi(t) &= \psi_0 + \int_0^t V(t) \kappa(t) dt & \frac{d\psi(t)}{dt} &= V(t) \kappa(t)
 \end{aligned}$$

The inverse mapping is that of determining curvature $\kappa(t)$ and speed $V(t)$ from the C space curve. Notice that C space is three dimensional while A space is two dimensional. Not only is the problem of computing this mapping a nonlinear differential equation, but it is underdetermined or **nonholonomic**. This is a difficult problem to solve and, from a mathematics standpoint, there is no guarantee that a solution exists at all.

Practical approaches to the C space to A space mapping problem often involve the generation of curves of the form:

$$\kappa(s) = \kappa_0 + as$$

where s is arc length²⁷. These curves are linear equations for curvature in the arc length parameter and are known as the **clothoids**. There is a growing body of literature on the generation and execution of clothoid curves in indoor, non-omnidirectional autonomous vehicle applications. The generation of clothoids can be computationally expensive. Their generation can also be unreliable if the algorithm attempts to respect practical limits on the curvature or its derivatives.

26. The choice of what is to be considered the controlled variable depends on the level of abstraction. At some boundary in the system, a speed is asked for and an error signal on speed is formed. Lower level loops may form errors on other quantities.

27. Mathematically, the clothoid generation process is equivalent to the power series solution of a differential equation. The idea is to assume a power series solution, substitute it into the original differential equation and then solve the simpler recurrence relation which results. This is, for example, how many higher transcendental functions (the Bessel function, for example) are defined.

1.3 Dynamics of the Constant Speed Reverse Turn

The limited rate of change of curvature for an Ackerman steer vehicle is an important modelling matter at even moderate speeds. A numerical feedforward solution to the dead reckoning equations was implemented in order to assess the realistic response of an automobile to steering commands. It was used to generate the following analysis. The maneuver is a reverse turn. The following figure gives the trajectory executed by the vehicle at various speeds for a 3 second actuator delay.

For a vehicle speed of 5 m/s, a kinematic steering model would predict that an immediate turn to the right is required to avoid the obstacle. However, the actual response of the vehicle to this command would cause a direct head-on collision. Any planner must explicitly account for steering dynamics - even at low speeds, in order to robustly avoid obstacles.

There are two fundamental reasons for this behavior. First, steering control is *control of the derivative of heading*, and any limits in the response of the derivative give rise to errors that are integrated over time. Second, curvature is an arc length derivative, not a time derivative. Hence the heading and speed relationships are *coupled differential equations*. The net result is that *the trajectory followed depends heavily on the speed*.




Figure 41- Constant Speed Reverse Turn

Feedforward of dynamics is necessary for stable control. If the vehicle decided to turn slightly right at 5 m/s speed, position feedback would indicate that the vehicle was not turning right. Any feedback control law which attempted to follow the ideal commanded arc would continue to increase the turn command while the steering servo tries to turn right. This overcompensation will eventually lead to the maximum turn command being issued²⁸ although a slight turn was commanded. Acceptable control is not possible without knowledge of these dynamics.

28. This is a major reason for the poor reliability of kinematic arc-based planners at higher speeds.

1.4 The Clothoid Generation Problem

The previous graph investigated the variability of the response to a steering command at various speeds. Consider now the response at a single speed to a number of steering commands. Again using the reverse turn at $t = 0$, the response curves for a number of curvature commands are as shown in the figure below:

The vehicle cannot turn right at all until it has travelled a considerable distance. Further, a configuration space planner which placed curve control points in the right half plane would consistently fail to generate the clothoid necessary, if it attempted to model the steering dynamics, *because the vehicle fundamentally cannot execute such a curve*²⁹. If the clothoid generator did not model such limits, the error would show up as instability and ultimate failure of the lower levels of control to track the path. The x-y region bounded by the curves is the entire region that the vehicle can reach.




Figure 42- Constant Speed Partial Reverse Turn

The only valid model of this system at even moderate speed is a coupled system of nonlinear differential equations. From the perspective of trajectory generation, it is advisable not to attempt the C space to A space transform in any form - particularly the generation of clothoids³⁰.

29. Thus, it is no accident that kinematic planners get “confused” as soon as a wide turn is executed. The incorrect low latency assumption in its particular form as a kinematic steering model is the reason for this behavior. The behavior will not arise if sharp turns are not attempted.

30. Of course, the fundamental issue is that of attenuation. Clothoid generation is feasible at low speed and for moderate turns at any speed.

1.5 Configuration Space

While it is difficult to compute the shapes of regions in configuration space in closed form, it is relatively easy to write a computer program to enumerate all possibilities and fill in boxes in a discrete grid which represents C space at reduced resolution. The three dimensional C space for an Ackerman steer vehicle for a 4.5 m/s impulse turn was generated by this technique.

The results are plotted below in heading slices of $1/16$ of a revolution. Symmetry generates mirror images along the heading axis, so two slices are plotted on each graph. The maneuver is an impulse turn from zero curvature to the maximum issued at $t=0$. A dot at a particular point (x,y) in any graph indicates that the heading of the slice is obtainable at that position. There are 16 slices in total of which 6 are completely empty (i.e the vehicle cannot turn around completely in 20 meters). The total percent occupancy of C space is the ratio of the total occupied cells to the total number of cells. This can be computed from the figure to be 3.1%.

So 97% of the C space of the vehicle is empty if the limited maneuverability of the vehicle is modelled. The maneuverability is limited by both the nonzero minimum turn radius and the steering actuator response. The impulse turn is the best case. If the initial curvature was nonzero, the percentage occupancy is even less.

The occupancy of C space does not account for higher level dynamics. There are severe constraints on the ability to “connect the dots” in these graphs, so the total proportion of feasible C space *paths* is far lower.

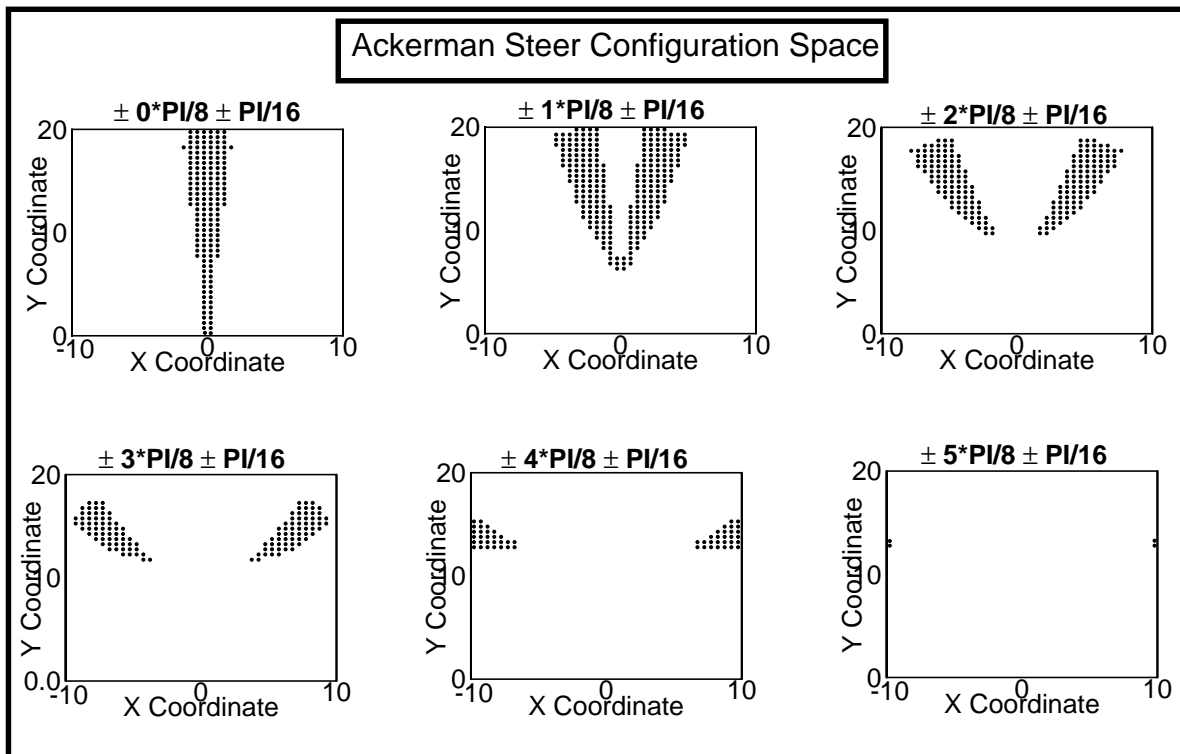


Figure 43- Ackerman Steer Configuration Space

1.6 A Real Time Control View of High Speed Autonomy

In general, the path planning problem of autonomous navigation can be cast into one of searching some space of alternatives for those paths which meet two constraints:

- they must be safe from dangerous hazards
- they must be reachable by the vehicle

On serial computers, these two constraints must inevitably be applied in the order stated or the reverse order. However, while the local result is the same, the order of application of the constraints can have a large impact on both the efficiency of the computation and the overall robustness of the system. The *hierarchical view* of the robot navigation problem is to check for collision first and give responsibility for execution of the path to a control algorithm placed lower in the hierarchy. The *control systems view* of the problem is to generate feasible paths first and then check for collisions. Both views are equally valid and have their domains of applicability. However, the hierarchical view is not optimal in high speed autonomy.

It is not efficient to plan in C space because too many solutions will be generated that do not satisfy the actuator dynamics constraint. C space is almost completely *degenerate* in the heading dimension³¹.

The search space for planning purposes is degenerate because heading is practically not an option at all and vehicle position is confined to a narrow cone. Thus *the configuration space for an Ackerman steer vehicle is degenerate above about 5 meters/sec speeds* when typical steering column response is considered. There is effectively no state space to be searched and search based planners would accomplish nothing useful.

Explicit enumeration of the few alternatives that exist is the efficient method of planning paths. The planning problem degenerates into a simple decision process, and quite often, even though some regions in front of the vehicle may be clear of obstacles, the vehicle cannot turn fast enough to reach them and the only alternative is to stop.

An efficient method of planning and evaluating paths would be based on, not candidate paths, but candidate command signals *expressed in actuation space* (in terms of curvature and speed as a function of time). Such a system would use a high fidelity feedforward simulator to solve the differential equations of motion in order to determine the exact response to these candidate commands. This strategy has the following advantages:

- The paths generated *meet the mobility constraints of the vehicle by construction*, so the difficult and often impossible problem of conversion from C space to A space is completely avoided. Instead, the reverse process of dead reckoning is used in the simulator. This is a kind of feedforward.
- The paths coarsely span the entire region of space that the vehicle can reach, so no real alternatives are missed within the fidelity requirement.

31. It is not surprising that this has occurred. Most research on the AI approach to path planning has been geometric in character - the implicit assumption being that dynamics could be safely ignored. Indeed they can be at slow speed.

2. Effect of Acuity and Response on Throughput - Rationale for Adaptive Perception

An autonomous vehicle must satisfy all component requirements of guaranteed safety simultaneously and these requirements are all interrelated. For example, sensor range depends on speed through the response requirement and resolution depends on range through the acuity requirement and throughput depends on resolution, so throughput depends on speed. This section will quantify this relationship between throughput and speed.

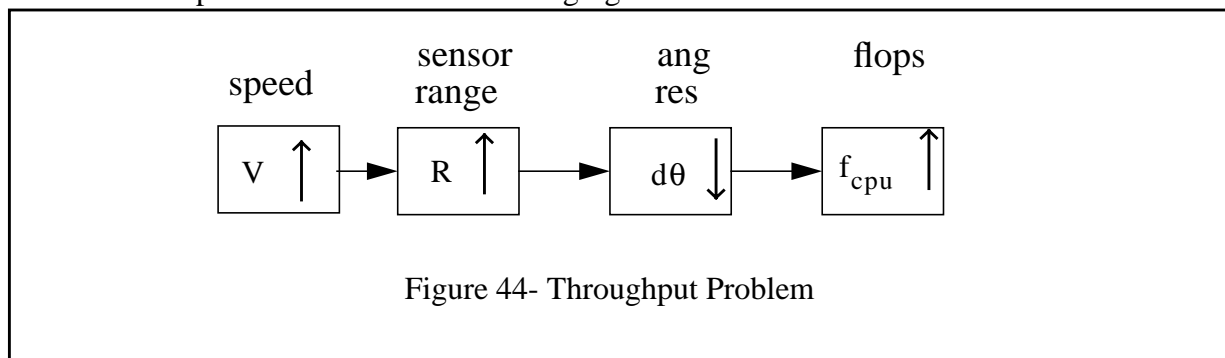
2.1 Throughput Problem

The throughput required to process an image depends on the number of pixels in the image. The number of pixels depends on the field of view and resolution. Resolution depends on the acuity requirement which implies it depends on range. The response requirement implies that range depends on speed so that resolution depends on speed. The throughput requirement also implies that field of view depends on speed. So ultimately, throughput can be expressed solely in terms of speed.

With an analysis of response and acuity it is possible to analyze the computational complexity of perception. In intuitive terms, guaranteed response implies that throughput is proportional to a high power of velocity because:

- Maximum range increases quadratically with speed (because braking distance does)
- Pixel size increases quadratically with maximum range (in order to resolve obstacles)
- Throughput increases quadratically with pixel size (assuming fixed field of view)

This relationship is indicated in the following figure:



When throughput is limited, this relationship gives rise to a trade-off between speed and resolution. Naive analysis suggests that the problem of high speed navigation is nearly impossible, because the necessary throughput is impractical. This will be called the **throughput problem**.

2.2 The Illusion

From an image processing perspective, the throughput problem appears to be impossible. Consider that contemporary rangefinders are 10 mrad resolution and many researchers believe that 1 mrad or so is needed to resolve obstacles. A ten fold increase in resolution is a hundredfold increase in pixels and a hundredfold increase in required throughput. Today, it is not possible to process 10 mrad images fast enough on a 10 Mflop processor. Therefore, if resolution were increased tenfold, it would be impossible to process 1 mrad images on a 1 Gflop processor. Brute force is not the elegant way to solve this problem.

On the other hand, the raw requirement is the throughput requirement, and this is trivial to meet. Consider that a 5 m/s vehicle covers about 6 map cells between images at 2 Hz, so there is band in the image about six pixels wide which would supply exactly the needed steady state throughput.

This section will show that *the throughput problem is an illusion* which arises from an image processing view of the problem and that simple adaptive techniques can solve it completely at contemporary sensor resolutions.

2.3 Adaptive Perception

It is possible to solve the throughput problem while simultaneously guaranteeing safety as efficiently as is theoretically possible by employing four principle mechanisms.

- **Adaptive Lookahead** is a mechanism for guaranteeing that the vehicle can respond to any hazards that it may encounter at any speed.
- **Adaptive Sweep** is a mechanism for guaranteeing barely adequate throughput and the fastest possible reaction time. In this way, speed is maximized.
- **Adaptive Scan** is a mechanism for ensuring barely adequate resolution that is as constant as possible over the field of view. In this way, speed is maximized without compromising robustness of the system.
- **Adaptive Regard** is a mechanism for ensuring that the system minimizes the spatial extent of the region it perceives based on vehicle maneuverability so that speed is maximized without compromising safety.

Together, these mechanisms can increase the efficiency (measured in terms of range pixel throughput) of a system by four orders of magnitude at 20 mph while simultaneously making it considerably more robust.

2.4 Assumptions of the Analysis

The following subsections will analyze the throughput problem in terms of the design of a fixed vehicle which is optimized for some maximum speed. The pixel size is permitted to change with speed, so the graphs represent the variation of system designs versus speed and not the throughput requirement for a single design as it drives faster.

The estimates that are produced are underestimates for many reasons.

- They are based on an oversampling factor of 1. A practical factor is at least 3. This implies that the results must be multiplied by the square of 6, or 36.
- Minimum acuity will be used because this is actually the most stringent requirement beyond some range.
- The maximum range that is chosen is based on the stopping distance. Actually it is the minimum range which should be set to the stopping distance, but the approximation is useful.
- Braking is chosen as the obstacle avoidance maneuver. This is viable for a system which stops when a hazard is detected. However, when a vehicle turns to avoid obstacles, sensor lookahead must exceed the stopping distance by a large factor.
- The processor load is assumed to be 50 flops per pixel when experience suggests that up to ten times this is required in a practical system.
- The graphs estimate perception geometric transform processing only. Planning, position estimation, and control are not included at all.
- Horizontal field of view is fixed at 80°, 120°, 170°, and 215° for each increasing reaction time respectively, based on earlier analysis.
- Frame rate is set to 2 Hz.

2.5 Common Throughput Expression

Recall from previous analysis that the throughput required from the computer for perception processing can be written as:

$$f_{\text{cpu}} = \frac{f_{\text{pixels}}}{\eta_s} = \frac{1}{\eta_s} \frac{\Psi}{(\text{IFOV})^2}$$

This can be written in terms of field of view and frame rate as follows:

$$f_{\text{cpu}} = \frac{1}{\eta_s} \frac{\text{HFOV} \times \text{VFOV} \times f_{\text{images}}}{(\text{IFOV})^2}$$

When it is necessary to employ a nonsquare pixel size, the horizontal and vertical pixel dimensions can be differentiated as follows:

$$f_{\text{cpu}} = \frac{1}{\eta_s} \frac{\text{HFOV} \times \text{VFOV} \times f_{\text{images}}}{\text{IFOV}_H \text{IFOV}_V}$$

2.6 Basic Mechanism

The basic mechanism for generating a complexity estimate is as follows:

- choose an angular resolution that is consistent with the need to resolve obstacles at the maximum range (guaranteed detection)
- choose a maximum range consistent with the need to stop if necessary (guaranteed response)
- choose a fixed field of view and frame rate (because sensors are designed that way)
- throughput is then simply the number of pixels generated per second times the cost of processing one pixel

Guaranteed detection is enforced by substituting for the IFOV from the minimum acuity rule developed earlier:

$$\text{IFOV} = \frac{1}{2} \frac{Lh}{R^2}$$

Guaranteed response is enforced by substituting for the maximum range based on the expression derived in an earlier section for the stopping distance in terms of the braking coefficient:

$$R = s_B = \tau_B V [1 + \bar{b}]$$

Complexity is estimated by noting that the braking coefficient does not approach 1 for the speed regimes of current research, so it can be neglected. Under this kinematic braking regime assumption, the stopping distance is the product of speed and reaction time. This is a characteristic vehicle distance - the distance required to stop. All complexity results will be polynomial in this distance.

2.7 Complexity of Constant Flux Range Image Processing

It turns out that, because a real sensor usually has a fixed field of view and fixed frame rate, it is possible to compute, not the underlying requirement, but the complexity of processing all of the data that the sensor generates if it is designed to properly resolve obstacles. This is the historical view of the problem - that of processing images in order to resolve obstacles. The whole image is evaluated and geometry and perhaps an obstacle map is later passed to a planner which determines an appropriate response.

The basic throughput expression under guaranteed detection is.:

$$f_{\text{cpu}} = \frac{f_{\text{pixels}}}{\eta_S} = \frac{1}{\eta_S} \left(\frac{4R^4}{(Lh)^2} \right) \Psi$$

Substituting the stopping distance gives:

$$f_{\text{cpu}} = \frac{f_{\text{pixels}}}{\eta_S} = \frac{1}{\eta_S} \left[\frac{4 (T_{\text{react}} V [1 + \bar{b}])^4}{(Lh)^2} \right] \Psi$$

In the kinematic braking regime, the following result for the computational complexity is obtained:

$$f_{\text{cpu}} \sim \frac{1}{\eta_S} O([T_{\text{react}} V]^4)$$

The following graph indicates the variation of throughput with speed when square pixel size is chosen to satisfy the maximum acuity resolution requirement at the maximum range. This paints a bleak picture. However, later sections will show that this graph arises from a design decision and is not intrinsic.




Figure 45- Throughput for Constant Flux, Constant Scan

2.8 Complexity of Adaptive Sweep Range Image Processing

By analogy to the last section, it is possible to compute the processing requirements associated with guaranteed response for a constant flux sensor if only the requirement for guaranteed throughput is met. The basic idea is similar to the previous section except that the sensor vertical field of view is chosen so that it barely meets the guaranteed throughput requirement. This does not compromise guaranteed response and it leads to significant improvement.

The basic throughput expression under guaranteed detection is:

$$f_{\text{cpu}} = \frac{f_{\text{pixels}}}{\eta_s} = \frac{1}{\eta_s} \left(\frac{4R^4}{(Lh)^2} \right) \Psi$$

The sensor flux is, again, the solid angle measured per unit time. Thus:

$$\Psi = \text{VFOV} \times \text{HFOV} \times f_{\text{images}}$$

where the product of the vertical field of view and the frame rate is a measure of the angular velocity of the beam, and is known as the **sweep rate**:

$$\dot{\theta} = \text{VFOV} \times f_{\text{images}}$$

An expression which relates the vertical field of view to its projection on the groundplane was given earlier as:

$$\text{VFOV} = \theta_{\text{max}} - \theta_{\text{min}} = \frac{h}{R_{\text{min}}} - \frac{h}{R_{\text{max}}} = \frac{h}{R_{\text{max}}} \left(\frac{R_{\text{max}}}{R_{\text{min}}} - 1 \right) = h \frac{(\Delta R)}{R_{\text{max}} R_{\text{min}}}$$

Recall that the throughput ratio is given by:

$$\rho_{\text{cyc}} = \frac{VT_{\text{cyc}}}{\Delta R}$$

Using this, eliminate all ranges but the maximum from the expression for VFOV, giving:

$$\text{VFOV} = \left(\frac{h}{R_{\text{max}}} \right) \frac{VT_{\text{cyc}} / (R_{\text{max}} \rho_{\text{cyc}})}{(1 - VT_{\text{cyc}} / (R_{\text{max}} \rho_{\text{cyc}}))}$$

Guaranteed throughput is implemented by setting $\rho_{\text{cyc}} = 1$. Defining, the nondimensional:

$$\rho_R = \frac{VT_{\text{cyc}}}{R_{\text{max}}}$$

This gives:

$$VFOV = (h/R_{\max}) \left(\frac{\rho_R/\rho_{\text{cyc}}}{\rho_R/\rho_{\text{cyc}} - 1} \right) = (h/R_{\max}) \left(\frac{\rho_R}{\rho_R - 1} \right)$$

which is an angle considerably smaller than that used in the previous graph. The complexity expression is now:

$$f_{\text{cpu}} = \frac{f_{\text{pixels}}}{\eta_S} = \frac{1}{\eta_S} \left(\frac{4R^4}{(Lh)^2} \right) (h/R_{\max}) \left(\frac{\rho_R}{\rho_R - 1} \right) (HFOV \times f_{\text{images}})$$

using the fact that:

$$\frac{1}{x-1} \sim -1 - x - x^2 - \dots$$

there results:

$$f_{\text{cpu}} = \frac{f_{\text{pixels}}}{\eta_S} \sim \frac{1}{\eta_S} \left(\frac{4R^4}{(Lh)^2} \right) \left(\frac{h}{R_{\max}} \right) \left(\frac{VT_{\text{cyc}}}{R_{\max}} \right) (HFOV \times f_{\text{images}})$$

which is, assuming every image is processed:

$$f_{\text{cpu}} = \frac{f_{\text{pixels}}}{\eta_S} \sim \frac{1}{\eta_S} \left(\frac{4R^4}{(Lh)^2} \right) \left(\frac{h}{R_{\max}} \right) \left(\frac{V}{R_{\max}} \right) (HFOV)$$

Substituting the stopping distance gives:

$$f_{\text{cpu}} = \frac{1}{\eta_S} \left[\frac{4 (T_{\text{react}} V [1 + \bar{b}])^2}{L^2 h} \right] (V) (HFOV)$$

In the kinematic braking regime, the following result for the computational complexity is obtained:

$$f_{\text{cpu}} \sim \frac{1}{\eta_S} O([T_{\text{react}} V]^2 [V])$$

which is less than the previous result by the factor $T_{\text{react}}^2 V$ or roughly two orders of magnitude.

This result leads to the conclusion that, in general, if the vertical field of view were adjusted for speed, computer requirements would be reduced significantly. This idea will be called **adaptive sweep** because it arises fundamentally from the direct relationship between throughput and the sensor sweep rate. An alternate name is **computational stabilization** which borrows terminology from strapdown inertial navigation.

The following graph indicates the variation of throughput with speed when the vertical field of view is computed from the above expressions and square pixel size is chosen to satisfy the

resolution requirement at the maximum range.




Figure 46- Throughput for Adaptive Sweep, Constant Scan

This graph demonstrates that the problems of the previous graph were an illusion. In fact, if the vertical field of view is chosen such that it satisfies the guaranteed throughput requirement, the throughput requirements are 300 times lower than constant flux at 10 meters/second speed.

Thus, while the logic of decreasing pixel size for higher speeds is inescapable, equivalent logic leads to a reduced vertical field of view requirement. Adaptive sweep is the key to keeping the navigation problem tractable because it is worth two orders of magnitude improvement in efficiency.

The key to this complexity difference is the redundant measurement of the same geometry that happens when images overlap on the groundplane. The average number of images that fall on a given patch of terrain is called the **imaging density** σ_I .

2.9 Complexity of Adaptive Sweep, Adaptive Scan Range Image Processing

The complexity analysis of this section gives the best complexity that can be practically achieved using a conventional fixed field of view sensor where the pixel shape is optimized in order to match the groundplane resolution requirement.

Although adaptive sweep is a significant improvement in complexity, it addresses only the optimum size of the image. It does not address the fact that the density of pixels on the groundplane and the aspect ratio of a pixel are grossly suboptimal. If the sensor pixel size and shape were optimal, then the throughput would be exactly equal to the theoretical minimum possible throughput and another two orders of magnitude improvement would be possible. While there are limits to what can be done, the idea of dynamically adjusting pixel size and shape, called **adaptive scan**, is an important one.

It is difficult to construct a system with perfect density on the ground. However, it is easy to subsample an image in azimuth to remove the effects of poor aspect ratio. Further, guaranteed response already requires that high depression scanlines (where the density is too high) be avoided. These two approaches can effectively eliminate the effects of poor geometric efficiency. Recall that the cpu load required to process all sensory data is given by:

$$f_{\text{cpu}} \sim f_{\text{pixels}} = \frac{\Psi}{(\text{IFOV})^2}$$

In practical adaptive scan, the pixel aspect ratio is adjusted to be equal to the perception ratio, and this gives the opportunity to cancel one of the ranges. The vertical and horizontal instantaneous field of view now have different expressions at minimum acuity:

$$\text{IFOV}_V = \frac{1}{2} \frac{Lh}{R^2} \quad \text{IFOV}_H = \frac{1}{2} \frac{Lh}{R^2} \left(\frac{R}{h} \right) = \frac{1}{2} \frac{L}{R}$$

thus, the throughput required to process all sensory data is given by:

$$f_{\text{cpu}} = \frac{f_{\text{pixels}}}{\eta_S} = \frac{1}{\eta_S} \left(\frac{4R^3}{L^2h} \right) \Psi$$

As before, the vertical field of view necessary for guaranteed throughput is:

$$\text{VFOV} = (h/R_{\text{max}}) \left(\frac{\rho_R/\rho_{\text{cyc}}}{\rho_R/\rho_{\text{cyc}} - 1} \right) = (h/R_{\text{max}}) \left(\frac{\rho_R}{\rho_R - 1} \right)$$

The complexity expression is now:

$$f_{\text{cpu}} = \frac{f_{\text{pixels}}}{\eta_S} = \frac{1}{\eta_S} \left(\frac{4R^3}{L^2h} \right) (h/R_{\text{max}}) \left(\frac{\rho_R}{\rho_R - 1} \right) (\text{HFOV} \times f_{\text{images}})$$

which is, assuming every image is processed:

$$f_{\text{cpu}} = \frac{f_{\text{pixels}}}{\eta_S} \sim \frac{1}{\eta_S} \left(\frac{4R^3}{L^2 h} \right) \left(\frac{h}{R_{\text{max}}} \right) \left(\frac{V}{R_{\text{max}}} \right) (\text{HFOV})$$

Substituting the stopping distance gives:

$$f_{\text{cpu}} = \frac{1}{\eta_S} \left[\frac{4 (T_{\text{react}} V [1 + \bar{b}])}{L^2} \right] (V) (\text{HFOV})$$

In the kinematic braking regime, the following result for the computational complexity is obtained:

$$f_{\text{cpu}} \sim \frac{1}{\eta_S} O(T_{\text{react}} V^2)$$

which is less than the previous result by the factor $T_{\text{react}} V$ or roughly one order of magnitude. The following graph indicates the variation of throughput with speed when vertical field of view is computed from the above expressions, nonsquare pixel size is chosen to satisfy the resolution requirement at the maximum range, and system cycle time is set to the frame rate.




Figure 47- Throughput for Adaptive Sweep, Adaptive Scan

It is clear from this graph that a problem which at first appeared impossible, requiring supercomputers, is in fact *theoretically trivial*. In fact, by using simple adaptive techniques, the throughput requirements are 20 times lower at 10 meters/second speed than adaptive sweep and *5,700 times lower than constant flux at 20 mph*.

2.10 Complexity of Adaptive Sweep, Uniform Scan Image Processing

The following considers the fundamental acuity and throughput requirements of perception in terms of a terrain map but it does not necessarily imply that a map is explicitly formed. As a minimum requirement, any sensor must generate map cells at a rate that is consistent with the rate at which the vehicle consumes them through motion. For now, consider that the motion of the vehicle consumes a swath of map cells directly in front of it as shown below:




Figure 48- Area Consumption

In the simplest case, this consumed area must be replaced by adding new information to the map shown to the right. In general, the new information need not be at the end of the map, but could be anywhere where an unknown cell exists. Also, although the information immediately to the left and right of the vehicle is not useful, the system did not know this at the time it was measured. Further, guaranteed response requires that the vehicle measure geometry a long way off, so, for now, there appears to be no solution to this problem³².

The area consumed per second, expressed in map pixels, is the required absolute minimum throughput of a perception system. This quantity is the minimum rate at which new geometric information must be generated, regardless of the scanning pattern of any sensor, or the vehicle must either

- drive over unknown terrain and violate guaranteed throughput
- accept inadequate resolution information and violate guaranteed detection

Let the width of the map be W , the velocity of the vehicle be V , and the map resolution (unit length per cell) be δ . This minimum rate is given by:

$$f_{\text{cells}} = \frac{WV}{\delta^2}$$

In previous sections it was shown that, under guaranteed response, the maximum range can be determined from the stopping distance. Let L be the vehicle wheelbase. Setting the width of the

32. Later, it will be shown that something, at least, can be done.

map to twice the maximum range gives:

$$f_{\text{cells}} = \frac{2s_B V}{L^2} = \frac{2T_{\text{react}} V^2 [1 + \bar{b}]}{L^2}$$

Putting all of these results together, gives the following expression for the processing load:

$$f_{\text{cpu}} = \frac{1}{\eta_s} f_{\text{cells}} = \frac{1}{\eta_s} \frac{2T_{\text{react}} V^2 [1 + \bar{b}]}{L^2}$$

In the kinematic braking regime, the following result for the computational complexity is obtained:

$$f_{\text{cpu}} \sim \frac{1}{\eta_s} O(T_{\text{react}} V^2)$$

which is, in complexity terms, ***equal to the adaptive sweep, adaptive scan expression***. This is the fundamental complexity of the problem³³. There is a multiplicative constant difference of $2 \times \text{HFOV}$ between this minimum requirement and the adaptive sweep, adaptive scan expression because the whole image is processed at the same nonsquare pixel resolution in adaptive scan.

This relationship is plotted below for minimum acuity map resolution of 3.3 meters versus vehicle velocity for various values of system reaction time. Again, for consistency reasons, 50 flops per cell are assumed. The analysis does not assume a uniform scan on the ground. It only assumes that the average density of the scanning pattern is one. This average is accumulated up to the time the vehicle drives past a particular point in the map.




Figure 49- Throughput for Adaptive Sweep, Uniform Scan

33. This is also why autonomy based on line scanners is feasible and so very efficient. Effectively, adaptive sweep and adaptive regard convert an imaging sensor into an ideal computationally stabilized line scanner.

2.11 The Fundamental Speed/Resolution Trade-off

Recall that the complexity estimates are all consistently based on a kinematic braking regime assumption. The true power of velocity is actually squared as speeds increase. Identical resolution assumptions have lead to the following throughput estimates for different image processing algorithms:

Table 5: Throughput Estimates

Algorithm	Estimate at Minimum Acuity, 4 second Reaction Time, and 10 m/s speed	Complexity
constant flux	250 Mflops	$O(T_{\text{react}}^4 V^4)$
adaptive sweep	0.7 Mflops	$O(T_{\text{react}}^2 V^3)$
adaptive sweep, scan	0.035 Mflops	$O(T_{\text{react}} V^2)$
ideal	0.0045 Mflops	$O(T_{\text{react}} V^2)$

The actual data for all 4 second reaction time curves is plotted below on a logarithmic vertical scale.




Figure 50- Throughput for All Algorithms

The logic of decreasing pixel size for higher speeds is inescapable, but equivalent logic leads to a reduced vertical field of view requirement, so if the vertical field of view is not reduced, extreme throughput waste is being tolerated. Further, because pixel aspect ratio is extremely elongated at high ranges, the density of measurements in the crossrange direction is grossly suboptimal unless it is managed.

Notice that the complexity in either the above cases contains a constant times a power of the product $T_{\text{react}} V$. That is:

$$f_{\text{cpu}} \sim \frac{1}{\eta_s} O([T_{\text{react}} V]^N [V]^M)$$

This will be called the **fundamental trade-off** because it indicates that the trade-off of finite computing resources is one of reliability for speed. This is a basic trade-off of speed and resolution which always arises from a system throughput limit. Computing resources establish a limit on vehicle performance which can be expressed as either high speed and low reliability or vice versa.

There are a few ways to read the result. If throughput is fixed, then speed is inversely proportional to reaction time. If speed is fixed, throughput required is the n th power of reaction time. If reaction time is fixed, throughput is the $(n+m)$ th power of speed.

3. Effect of Throughput on Response - Rationale for A Real Time Approach

The previous sections have shown that the throughput required of a system is proportional to high powers of both speed and reaction time by substituting the dependence of maximum range on speed, and the dependence of pixel size on range into the throughput expression. This was called the **throughput problem**. The resulting throughput depended on both reaction time and speed.

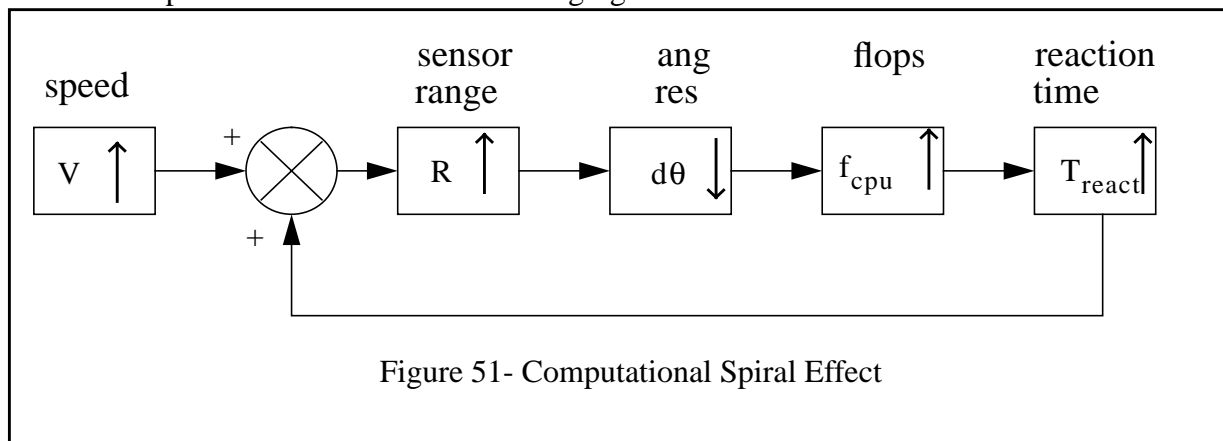
3.1 Computational Spiral Effect

However, there is still another dependence which has not been resolved. Reaction time itself depends on throughput. It increases as throughput increases for a fixed computer system. The intuitive logic behind this is:

- Driving faster, or reacting slower, requires looking farther ahead
- Looking farther ahead requires increased sensor angular resolution in order to resolve hazards at the lookahead distance.
- Increased resolution generates higher processing loads.
- Higher processing loads cause longer reaction times
- Longer reaction times require looking farther ahead

Therefore increased reaction time causes increased reaction time because computer cycle time is a component of reaction time. That is, as speeds increase, the computer has to do even more work per cycle than what would be expected based on the previous analysis. Effectively, ***reaction time itself is a function of throughput*** and therefore it is a function of velocity. This will be called the **computational spiral effect**.

The spiral affect arises fundamentally from the policy of guaranteed safety because it closes a relational loop. It is indicated in the following figure:



This relationship gives rise to a trade-off between throughput and reaction time. The analysis suggests that optimal real-time performance is achieved by reducing the field of view and increasing the frame rate in order to improve reaction time, but to ensure that guaranteed throughput is still maintained. In theoretical terms, ***the trade-off is managed by reducing throughput to a bare minimum and reducing reaction time as far as possible. This is the real time approach to the problem.***

4. A Strategy For High Speed Autonomous Navigation

The throughput curves indicate that supercomputer computational bandwidth is required when the constant flux sensor is used and system reaction time exceeds 1 second if a guaranteed safety policy is adopted. However, the earlier graphs and the logic which leads to them indicate practical policies which will render the throughput problem solvable. Optimal systems will operate robustly at or near the throughput performance limit. Systems based on today's technology will succeed through management of the spiral effect. The spiral effect is both a hardship and an advantage because it is reversible. While every increase in reaction time is costly, every decrease is equally advantageous.

4.1 Basic Strategy

The throughput performance limit represents a speed / reliability tradeoff. Hence, any techniques which increase speed, can be used instead to increase reliability. Among those techniques, the following have been identified as most important.

4.1.1 Real Time Systems Analysis and Design

A real time systems approach to the problem is important because the sensitivity of computational bandwidth to reaction time is extreme. Vehicle reaction time must be aggressively reduced as far as is possible.

4.1.2 Adaptive Sensors and Adaptive Perception

The constant flux sensor combined with straightforward image processing is a poor long term approach to the problem. Possibilities exist for sensor performance adaptability in very many areas. For nonadaptive sensors, much can still be done with adaptive perception software. Such a solution represents the best of both worlds because it permits wide field of view for rough terrain work without the traditional computational cost.

4.1.3 A Deliberative Approach to Autonomy

For the purpose of the report a deliberative system will be considered to be one which maximizes its memory of the state of the environment from previous sensor measurements.

4.1.4 Key Assumptions

It is clear that it is an easy matter to overdesign a system for its given application. A careful specification of the nature of the terrain to be traversed in terms of smoothness, obstacle sparsity, and benignness is a critical parameter affecting throughput. A degree of relaxation of the guaranteed safety policy is a key consideration.

4.1.5 Vehicle Configuration

If the guaranteed safety policy cannot be relaxed, hardening the vehicle against hazards, increasing wheel size, lifting the sensor as high as possible, and perhaps other such measures promise to improve matters. A conventional automobile is designed for people, not computers, and measures as simple as increasing the size of the wheels make a major difference. This is one reason why the throughput problem disappears for some planetary rovers.

4.2 Real Time Systems Analysis and Design

4.2.1 Hardware Platform

In a real time systems approach to the problem, the following techniques will decrease reaction time:

- Remove all sensor latency by providing a line by line interface with no buffering.
- Synchronize the sensor and the perception software.
- Utilize a dedicated pipelined geometry engine for real time coordinate transformation.
- Utilize multiprocessing by allocating perception and planning to separate processors.

4.2.2 Real Time Software

It is clear that software efficiency must be pushed to the limit. However, there are better ways to achieve this than to implement incomprehensible code. Some of the promising techniques include.

- Minimalist software configurations are called for. The fundamental tradeoff demands that the effect of each proposed software addition be scrutinized for its effect on reaction time. Those additions which do not satisfy a cost benefit test must not be implemented.
- An efficient terrain map data structure and access routines are mandatory. Sufficiently accurate position estimation or simple map matching techniques are required to make this possible.

4.3 Adaptive Sensors

Requirements for horizontal field of view, vertical field of view, maximum range, and beam dispersion all change radically with velocity. A fully adaptive sensor which dynamically changes all of the above parameters represents a major research effort. However, a rangefinder with a pointable vertical sweep is relatively easy to construct. For such a sensor, the following attributes are necessary:

- Pointing the sweep must be accomplished by moving the least mass, and that means the scanning mirrors. Heavy stabilized sensor heads with poor response are not optimal.
- The vertical field of view must always contain enough time correlated lines to permit map matching to recover the vehicle z excursion.
- The maximum and minimum depression angle must be independently controllable and controllable in real time through an application command interface to the sensor.
- A stabilization servo must be supplied for which the groundplane y coordinate rate is the controlled variable. That is, inertial or other angular stabilization is not optimal. This is because such stabilization causes extreme sweep rates when the vehicle is elevated or depressed by terrain geometry. Since this approach requires knowledge of the terrain, this loop must be closed in the application, not in the sensor electronics.

A pointable stereo or rangefinder sensor head is another way to improve matters, provided successful adaptive control strategies can be developed for it.

4.4 Adaptive Perception

Perhaps the most elegant solution to the throughput problem is the implementation of adaptive sweep, adaptive scan, adaptive lookahead and adaptive regard in software.

- Software subsampling of the range image will alleviate the problem of poor scanning efficiency.
- Software flux adaptation will avoid processing range pixels which add nothing useful to the map. This includes ignoring the lower scanlines and peripheral columns at higher speeds.
- It is a simple matter to compute the vehicle stopping or impulse turning distance in software as a function of velocity in order to constrain the search for useful data in an image.

4.5 Deliberative Approach

It is clear that high speed navigators cannot afford the luxury of viewing the entire world, taking one step, and then viewing the entire world all over again³⁴. Computers are not fast enough to process the data at the density required to resolve obstacles at the required distance from the vehicle.

This problem is a key design driver, and it is so serious that it is legitimate to let its solution drive the design decisions of otherwise unrelated systems. One way to view the trade-off of computational throughput is the question of whether a **deliberative** or **reactive** system is to be designed, or, at least where on the spectrum the design point should be. For, in a deliberative system, the use of “maps” allows memory of portions of the environment from cycle to cycle and obviates the need to view them all over again.

4.6 Key Assumptions

It is important to adopt as many assumptions as are feasible without compromising the nature of the mission. Extreme computational advantages are available when this is done. For research purposes, it is legitimate to engineer the test area to enforce consistency of the terrain with the assumptions.

4.6.1 Continuity Assumption

Many aspects of adaptive perception and feedforward control involve circular logic. For example, it is impossible to know where the vehicle will go until the terrain is known and it is impossible to know what image data to use until the vehicle trajectory is known. This is one of the traditional justifications for the image processing view. Luckily, there are often conservative approximations that can be made to break the circle. When this is not feasible, information from the last cycle of computations can be used as an approximation for the current cycle. This will be called the **continuity assumption**.

4.6.2 Terrain Smoothness Assumption

Some degree of **terrain smoothness assumption** must always be adopted. Use of a particular map resolution is tantamount to adopting this assumption. Without it, throughput is infinite. This assumption is critical to the successful operation of today’s generation of autonomous navigators. This, of course, also implies that systems will fail with some regularity when this assumption is

34. In general, all other things being equal, this implies that deliberative systems should be able to drive faster than reactive ones.

violated. A system cannot avoid features that it cannot see (at planning time), and it cannot see what it cannot represent. Random success with obstacle detection must be expected when obstacle feature sizes equal or are less than the map resolution.

4.6.3 Benign Terrain Assumption

There are three aspects of system performance which benefit from the **benign terrain assumption**

4.6.3.1 *Reduced VFOV*

Typically, on rough terrain, kinematic concerns dominate the choice of the VFOV. Thus, the benign terrain assumption is a key assumption which can significantly reduce throughput. Use of a programmable sweep is an alternative way of reducing the VFOV without adopting this assumption.

The body pitch consideration leads to extreme waste in benign terrain. In the kinematic braking regime, the vertical field of view relates directly to the imaging density and hence to the required throughput. Intuitively, this is the fundamental advantage of the single scanline sensor. However, the single scanline sensor makes it impossible to use the time honored technique of image matching to recover the vehicle z coordinate.

4.6.3.2 *Reduced Depression Angle, Increased Lookahead*

Extant constant angular resolution sensors can give rise to scanning densities within a single image which are extreme and attendant extreme growth in the throughput requirements. This happens when the vertical field of view is increased without sacrificing obstacle resolving power at high ranges.

The angle of a scanline below the horizon is called the **depression** angle. The *sole* justification for high depression angles is the need to see behind occlusions in rough terrain. Of course, systems which operate in image space are prevented from seeing proximal terrain, but this is only necessary to begin with because of inadequate angular resolution at high ranges. Scanning density grows with the third power of the minimum range, while angular resolution grows only with the first or second power of range. Hence, *it is cheaper as well as safer to detect obstacles at far range even if it requires higher angular resolution.*

4.6.3.3 *Unknown Obstacle Assumption*

A system which can faithfully adopt a policy that occlusions are obstacles can reduce the vertical field of view to a point where image matching is still feasible and enjoy significant computational advantages. The fact that humans manage to navigate nonetheless under similar, in fact, worse constraints is an indication that the problem of navigating in the face of occlusion is not insurmountable.

One factor which mitigates the problem of occlusion significantly is that it is often only obstacles which cause occlusion. A system which avoids occlusions also avoids obstacles. This will be called the **unknown obstacle assumption**. Unless the hill occlusion rule is satisfied, this assumption is not completely correct. If this assumption is adopted anyway (and it often has to be under a policy of guaranteed vehicle safety because some occlusions will never be revealed), it is only a practical strategy when alternate feasible vehicle trajectories exist. Hence, adopting this assumption

compells a system to also adopt the **obstacle sparsity assumption**.

4.6.4 Obstacle Sparsity Assumption

There are many costs associated with a low mounted sensor. The geometric limitation of the sensor height is the root of some of the fundamental technological problems of autonomous navigation. It gives rise to all of the following problems:

- prevalence of terrain self occlusions (stereo or lidar)
- very nonuniform sampling of terrain geometry (stereo or lidar)
- poor sensor signal to noise due to glancing incidence at high range (stereo or lidar)
- extreme growth of computational complexity with velocity

The nonuniform sampling varies not only with position but also with direction. For typical geometry, it is usual that the angular width of the beam projects to a much smaller linear distance on the vertical than it does in the groundplane. Hence, the **detection** of a vertical step obstacle is more precise than its **localization**.


In order to minimize throughput requirements, high speed systems are forced to adopt sensors of barely adequate angular resolution. Systems which roof-mount such sensors are compelled to assume that detection is more important than localization. This assumption works provided giving obstacles a wide berth is practical. Hence, use of a finite angular resolution roof mounted sensor forces the high speed system to adopt the **obstacle sparsity assumption** to some degree dependent on the ratio of vertical to horizontal resolution (this was shown earlier to be equal to the perception ratio). This is a second way in which limited height requires this assumption.

4.7 Small Incidence Angle Assumption

It turns out that the low perception ratio (h/R) which causes problems with respect to scanning density gives some payback with respect to adaptive regard and the sampling problem.

On the surface, adaptive regard appears to be impossible because the mapping from image space to world space is unknown until it is computed. Once a pixel is computed, it might as well be used. This logic leads to processing the entire image and it is doomed to failure as was shown earlier. The only remaining question is how to do it.

Luckily, the range measurement from the sensor to the environment is almost identical to its groundplane projection (because the angle involved is so shallow). Indeed, the relative error in assuming the two are equal is the square of the perception ratio:



Which is on the order of 1% for high speeds simply because the range is so large. The assumption

of a small perception ratio will be called the **small incidence angle assumption**. If the perception system attempts to process all geometry within a **range window**, the quest for the end of the range window will automatically walk right to the top of the image if necessary and, as a side effect, it will discover the height of a near vertical surface as long as it remains within the range window. This mechanism is far superior to simply processing a fixed subset of the vertical field of view (an **elevation angle window**) in part because of its performance on near vertical surfaces. Such surfaces would fall outside and elevation angle window and would not be completely processed.

The processing of pixels outside the range window is comprised of nothing more than reading their values and comparing them to the window. Further, because a terrain map already makes a 2-1/2 world assumption, the range values can be **consistently** assumed to be monotonic in elevation angle and this further reduces the required processing because some pixels need never be visited at all. This **monotone range assumption** also provides the basis for ambiguity removal in phase ambiguous sensors like AM rangefinders.

It was shown earlier that as the minimum range increases, the scanning density decreases quadratically and approaches one. Therefore, because adaptive regard discards high depression scanlines anyway, *the sampling problem is far less severe at high speed.*

4.8 Vehicle Configuration

The acuity rules and occlusion rules have uncovered some of the relationships between vehicle configuration and sensor requirements. A vehicle can make the job of perception easier.

4.8.1 Mechanical Design

Some techniques that have been uncovered include:

- Hardening the vehicle against those particular hazards which are difficult to detect. These include small steps and holes.
- Increasing wheel size to be large with respect to the wheelbase.
- Raising the sensor height as far as possible, perhaps even providing an actuated degree of freedom for this purpose.

4.8.2 Sensor Design

A ranging sensor that was specifically designed for high speed autonomy would have a very narrow vertical field of view and significantly nonsquare pixels. Further, stereo perception algorithms designed specifically for high speed autonomy can enjoy immediate efficiency increases of orders of magnitude if they implement adaptive sweep and adaptive scan directly in the stereo algorithm. Indeed, simply prefiltering the input images by a factor of 10 to 1 in azimuth provides an order of magnitude increase in stereo throughput. Also, the traditional square camera image can have a large fraction of the lower pixels eliminated outright. Depending on the lookahead distance another tenfold increase in throughput is available from this measure. In general, the throughput that is being wasted on redundant computations can instead be reallocated to improve resolution by an overall factor of 10.

5. References

- [1] O. Amidi, "Integrated Mobile Robot Control", Robotics Institute Technical Report CMU-RI-TR-90-17, Kelly, A., "A 3D State Space Formulation of a Navigation Kalman Filter for Autonomous Vehicles", CMU-RI-TR-XXX
- [2] M. G. Bekker, "The Theory of Land Locomotion", The University of Michigan Press, 1956.
- [3] M. G. Bekker, "Off-the-Road Locomotion", The University of Michigan Press, 1960.
- [4] B. Brumitt, R. C. Coulter, A. Stentz, "Dynamic Trajectory Planning for a Cross-Country Navigator", Proceedings of the SPIE Conference on Mobile Robots, 1992.
- [5] T. S. Chang, K. Qui, and J. J. Nitao, "An Obstacle Avoidance Algorithm for an Autonomous Land Vehicle", Proceedings of the 1986 SPIE Conference on Mobile Robots, pp. 117-123.
- [6] M. Daily et al. "Autonomous Cross Country Navigation with the ALV", Proceedings of the 1988 IEEE International Conference on Robotics and Automation, pp. 718-726.
- [7] E. D. Dickmanns, "Dynamic Computer Vision for Mobile Robot Control", Proceedings of the 19th International Symposium and Exposition on Robots, pp. 314-27.
- [8] E. D. Dickmanns, A. Zapp, "A Curvature-Based Scheme for Improving Road Vehicle Guidance by Computer Vision", Proceedings of the SPIE Conference on Mobile Robots, 1986.
- [9] D. Feng, S. Singh, B. Krogh, "Implementation of Dynamic Obstacle Avoidance on the CMU Navlab", Proceedings of IEEE Conference on Systems Engineering, August, 1990.
- [10] J. Gowdy, A. Stentz, and M. Hebert, "Hierarchical Terrain Representation for Off-Road Navigation", In Proc SPIE Mobile Robots 1990.
- [11] A. J. Kelly, "Essential Kinematics for Autonomous Vehicles", CMU Robotics Institute Technical Report CMU-RI-TR-94-14.
- [12] A. J. Kelly, "Modern Inertial and Satellite Navigation Systems", CMU Robotics Institute Technical Report CMU-RI-TR-94-15.
- [13] A. J. Kelly, "A Feedforward Control Approach to the Local Navigation Problem for Autonomous Vehicles", CMU Robotics Institute Technical Report CMU-RI-TR-94-17.
- [14] A. J. Kelly, "Adaptive Perception for Autonomous Vehicles", CMU Robotics Institute Technical Report CMU-RI-TR-94-18.
- [15] A. J. Kelly, "A 3D State Space Formulation of a Navigation Kalman Filter for Autonomous Vehicles", CMU Robotics Institute Technical Report CMU-RI-TR-94-19.
- [16] A. J. Kelly, "An Intelligent Predictive Controller for Autonomous Vehicles", CMU Robotics Institute Technical Report CMU-RI-TR-94-20.
- [17] A. J. Kelly, "Concept Design of A Scanning Laser Rangefinder for Autonomous Vehicles", CMU Robotics Institute Technical Report CMU-RI-TR-94-21.
- [18] D. Langer, J. K. Rosenblatt, M. Hebert, "A Reactive System For Off-Road Navigation", CMU Tech Report
- [19] M. Hebert and E. Krotkov, "Imaging Laser Radars: How Good Are They", IROS 91, November 91.
- [20] M. Hebert. "Building and Navigating Maps of Road Scenes Using an Active Sensor", In Proceedings IEEE conference on Robotics & Automation, 1989; pp.36-1142.
- [21] M. Hebert, T. Kanade, and I. Kweon. "3-D Vision Techniques for Autonomous Vehicles", Technical Report CMU-RI-TR-88-12, The Robotics Institute, Carnegie Mellon University, 1988
- [22] B.K Horn and J. G. Harris, "Rigid Body Motion from Range Image Sequences", Image Understanding, Vol 53, No 1, January 1991, pp 1-13
- [23] Hoffman, R., Krotkov, E., "Terrain Mapping for Outdoor Robots: Robust Perception for Walking in the Grass", Submitted to IEEE International Conference on Robotics and Automation, 1993.

- [24] In So Kweon, "Modelling Rugged Terrain by Mobile Robots with Multiple Sensors", CMU PhD Thesis, 1990
- [25] Olin, K. and Tseng, D. "Autonomous Cross Country Navigation", IEEE Expert, August 1991, pp. 16-30.
- [26] Singh, S., et. al, "FastNav: A System for Fast Navigation," Robotics Institute Technical Report CMU-RI-TR-91-20, Carnegie Mellon University, 1991.
- [27] D. H. Shin, S. Singh and Wenfan Shi. "A partitioned Control Scheme for Mobile Robot Path Planning", Proceedings IEEE Conference on Systems Engineering, Dayton, Ohio, August 1991

Index

A

accuracy	6
actuation space	88
acuity	7
acuity adaptive planning	13
acuity adaptive scan	13
acuity problem	12, 58
acuity ratios	12, 56
adaptive lookahead	36, 50, 94
adaptive regard	35, 36, 94
adaptive scan	73, 94, 101
adaptive sweep	50, 73, 94, 99
attitude acuity rule	83
azimuth	58

B

beam dispersion	58
benign terrain assumption	15, 53, 111
braking angle	25, 32
braking coefficient	23
braking distance	21, 22
braking reaction time	19
Buckingham Pi theorem	1

C

characteristic time	77
clothoid	78, 88
coefficient of lateral acceleration	25
communications load	55
computational spiral effect	107
configuration space	88
constant flux	54
constant scan	54
continuity assumption	110
crossrange	58

D

dead zone	34
deduced reckoning	88
deliberative	110
depression	58, 111
detection zone	34
dimensional analysis	1
downrange	58
dynamic braking coefficient	23
dynamic braking regime	24

dynamic steering limit	26
dynamic turning coefficient	31
dynamic turning regime	31
E	
elevation	58
elevation angle window	113
F	
feedforward	81
fidelity	7, 75
fidelity adaptive planning	15
fidelity problem	14, 84
fidelity ratios	14
filter	77
flat world assumption	15, 51, 56, 64
free zones	34
Fresnel Integrals	88
fundamental tradeoff	106
G	
geometric efficiency	72, 73
guaranteed detection	12
guaranteed localization	14
guaranteed response	8
guaranteed safety	7
guaranteed throughput	10
H	
high throughput assumption	11
hill occlusion rule	51
hole occlusion rule	52
I	
image registration problem	15, 84
imaging density	50, 71, 73, 100
impulse turn	21
impulse turn maneuver	28
impulse turning distance	29
incremental lookahead distance	35
instantaneous field of view	58
interpolation	58
K	
kinematic braking coefficient	23
kinematic braking regime	24
kinematic steering limit	26
kinematic turning coefficient	31
kinematic turning regime	31
L	
latency problem	9, 19, 86

lateral occlusion problem	53
localization	85
longitudinal aspect ratio	16
lookahead angle	35, 36
lookahead distance	36
low dynamics assumption	15
low latency assumption	9, 77

M

map resolution	14
maximum acuity rule	13
maximum fidelity	76
maximum fidelity ratio	76
maximum fidelity rule	15
maximum sensor acuity rule	63
minimum acuity rule	12
minimum fidelity rule	14
minimum sensor acuity rule	62
minimum significant delay	76
monotone range assumption	113
motion distortion problem	13, 65
myopia problem	9, 64

N

negative obstacle	52
nonholonomic	88
nonholonomy problem	86
normalized incremental lookahead distance	36
normalized map resolution	70
normalized range difference	73
normalized undercarriage clearance	16
normalized wheel radius	16
normalized wheelbase	16

O

observability problem	15
obstacle sparsity assumption	15, 112
occlusion problem	51
oversampling	58
oversampling factor	56

P

panic stop	21
panic stop maneuver	22
perception ratio	16, 51, 57, 73
perceptual software efficiency	55
pi product	1
pixel processing load	55
planning angle	25, 32

planning distance	22, 29
positive obstacle	51
processing efficiency	55

R

range ratio	36, 60, 73
range window	113
reaction time	7
reactive	110
requirements analysis	1
resolution	6, 56
response adapted lookahead rule	9
response adapted speed	9
response adaptive lookahead	9
response adaptive speed	9
response problem	8
response ratio	8, 18, 34, 86
reverse turn	21, 30, 43
rigid suspension assumption	83
rigid terrain assumption	83

S

safe terrain assumption	7
sampling problem	13, 58
sampling theorem	56
scanning density	69, 73
sensitivity problem	15
sensor flux	54
small incidence angle assumption	47, 113
speed	6
stabilization problem	11
stopping angle	25, 32
stopping distance	22
stopping region	27
sweep rate	54, 98
sweep rate rule	48
system reaction time	18

T

terrain smoothness assumption	13, 57, 58, 110
throughput	7
throughput adapted speed	11
throughput adapted sweep	11
throughput adaptive speed	11
throughput adaptive sweep	11
throughput problem	10, 49, 93, 107
throughput ratio	10, 36, 39, 50, 71
time constant	77
timing	6

transient turning coefficient	80
tunnel vision problem	11
turning coefficient	31
turning distance	25, 29
turning reaction time	19
turning stop	21
turning stop maneuver	27
U	
undercarriage tangent	16
undersampling	58
uniform scan assumption	13
unknown hazard assumption	53
unknown obstacle assumption	111
V	
vertical	58
vertical sweep rate	50
W	
wheel fraction	16, 63
wheelbase	16
wide depth of field assumption	9
Y	
yaw rate	25