In 1997 IEEE
International Conference on Robotics and Automation
Albuquerque, New Mexico - April 1997

Analysis of Requirements for High Speed Rough Terrain Autonomous Mobility Part 1: Throughput and Response

Alonzo Kelly and Anthony Stentz

Field Robotics Center, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213-3890
email: alonzo@ri.cmu.edu, axs@ri.cmu.edu
www: http://www.frc.ri.cmu.edu/~alonzo, http://www.frc.ri.cmu.edu/~axs

Abstract

A basic requirement of autonomous vehicles is that of guaranteeing the safety of the vehicle by avoiding hazardous situations. This paper analyses this requirement in general terms of real-time response, and throughput. Several nondimensional expressions emerge which characterize requirements in canonical form.

1 Introduction

The need for high throughput perception algorithms has been acknowledged for some time [1][2][3][7][8][9][10][11] in the field of autonomous vehicle navigation. Yet, evidence for this need has not been based on any underlying theory. This paper is concerned with the requirements that must be satisfied by an autonomous vehicle which operates safely in its environment.

For the purposes of this paper, the perception sensor can be any imaging sensor measuring range or intensity in any electromagnetic band of frequencies. We will propose aspects of a rudimentary theory of obstacle avoidance and use it to quantify some of the requirements placed upon autonomous systems that are derived from the need to ensure safety. It is one of three related papers [4], [5] in these proceedings.

2 Guaranteed Safety

Any vehicle which attempts to navigate autonomously in the presence of unknown obstacles must exhibit performance that satisfies a basic set of requirements. At the highest level, if the system is to survive on its own, the vehicle control system must implement a **policy of guaranteed safety**.

It may be possible in simple environments to make the default assumption that the terrain is navigable in the absence of direct evidence to the contrary. The *weak form* of the policy is optimistic. It requires that the vehicle guarantee, to the best of its ability, that collisions with *identified* obstacles will be avoided. The system must prove an area is not safe before not traversing it. An example of such an environment is a flat floor indoor setting.

In more complex environments, it is necessary to make the default assumption that the terrain is not navigable in the absence of direct evidence to the contrary. In its *strong form*, the policy is pessimistic. It requires that a vehicle not enter terrain that it has not both perceived and understood. The system must prove that an area is safe before traversing it. An example of such an environment is a rough terrain outdoor environment.

This requirement to guarantee safety can be further broken down into four other requirements on performance and functionality expressed in terms of timing, speed, resolution, and accuracy. In order to survive on its own, an autonomous vehicle must implement the four policies of:

- **guaranteed response**: It must respond fast enough to avoid an obstacle once it is perceived.
- **guaranteed throughput**: It must update its model of the environment at a rate commensurate with its speed.
- **guaranteed detection**: It must incorporate high enough resolution sensors and computations to enable it to detect the smallest event or feature that can present a hazard.
- **guaranteed localization**: It must incorporate sufficiently high fidelity models of itself and the environment to enable it to make correct decisions and execute them sufficiently accurately.

2.1 Preliminaries

A nondimensional expression of the above policies provides the most compact expression of the relationships between speed, reaction time, and other system performance parameters. Results will be expressed in a scale-independent form when this is possible. Before developing such expressions, a brief background discussion is in order.

2.1.1 Lexical Conventions

The paper will introduce many new terms as a device to foster brevity and precision. New terms will be defined in their first appearance in the text. They will generally be highlighted **thus**.

2.1.2 Nomenclature

The words **response** and **reaction** will be distinguished for reasons of notational convenience. Generally, response will refer to the entire autonomous system including the vehicle, and reaction will refer to the computational and control aspects of the system only. Finally, the term **maneuver** will apply to the vehicle physical response only.

For example, if the vehicle applies the brakes, the time it took to decide to brake is the reaction time, the time spent stopping is the maneuver time, and the sum of these is the response time.

$$T_{response} = T_{react} + T_{maneuver}$$

The **instantaneous field of view** will be defined as the angular width of a pixel.

2.1.3 Coordinate Conventions

The angular coordinates of a pixel will be expressed in terms of horizontal angle or **azimuth** ψ, and vertical angle or **elevation** θ. Three orthogonal axes are considered to be oriented along the vehicle body axes of symmetry. Generally, we will arbitrarily choose z up, y forward, and x to the right:

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

2.1.4 Notation

We will carefully distinguish range, R measured in 3D from a range sensor, and the projection of range Y onto the groundplane. Generally, both will be measured forward from the sensor unless otherwise noted.

2.1.5 Acronyms

The following acronyms will be employed:

- •**VFOV** - vertical field of view
- •**HFOV** - horizontal field of view
- •**IFOV** - instantaneous field of view
- •**HIFOV** - horizontal instantaneous field of view
- •**VIFOV** - vertical instantaneous field of view.

2.2 Nondimensional Configuration

Certain vehicle dimensions that will be generally important in the analysis are summarized in the following figure. One distinguished point on the vehicle body will be designated the vehicle control point. The position of this point and the orientation of the associated coordinate system is used to designate the pose of the vehicle.

The wheelbase is L, and the wheel radius is r. The height of the sensor above the groundplane is designated h and its offset rear of the vehicle nose is p. The height of the undercarriage above the groundplane is c. Range measured from the sensor is designated R.

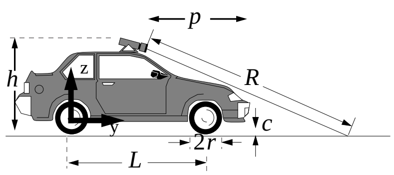

Figure 1: Important Dimensions

2.3 Key Nondimensionals

Certain nondimensional variables that encode relevant aspects of the vehicle geometry will be employed later in the paper.

- •L/R : **normalized wheelbase**, the ratio of wheelbase to measured range, encodes the size of the vehicle relative to its sensory lookahead, relates to requirements on sensor angular resolution.
- •h/R : **perception ratio**, the ratio of sensor height to measured range, encodes the sensor height relative to vehicle sensory lookahead, encodes angle of incidence of range pixels with the terrain, relates to requirements on sensor angular resolution, pixel footprint aspect ratio, and prevalence of terrain self occlusions.
- •c/L : **undercarriage tangent**, the ratio of undercarriage clearance to wheelbase, encodes body clearance aspects of terrainability in scale independent terms, relates to the prevalence of terrain self occlusions.

2.4 Nondimensional Safety Requirements

One way to characterize scale is to choose a characteristic vehicle dimension to represent its size. Let the wheelbase L be chosen for this purpose here. Let V represent vehicle speed and let T represent an interval of time. One nondimensional quantity that will concern us is the ratio of a velocity-time product to a distance. This generic nondimensional can be expressed as:

$$\sigma = VT/L$$

If $T_{response}$ represents the time required to respond to an obstacle, the product of speed V and this response time will be called a **response distance**. This distance can be defined for any particular obstacle avoidance maneuver or class of maneuvers. If we normalize this distance by the wheelbase, a nondimensional is created which expresses response distance in scale-independent terms. Thus, the **normalized response distance** is:

$$\sigma_{response} = VT_{response}/L$$

This number encodes the capacity of a vehicle to respond relative to its own size. If the number is large, it implies that vehicle maneuverability is low in scale-independent terms.

2.4.1 Response

Obstacles cannot be avoided unless the vehicle can react fast enough. The response distance can never be allowed to exceed the sensory lookahead distance Y_L.

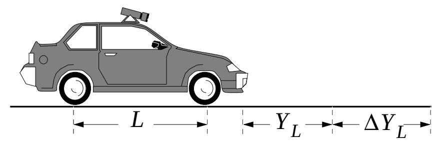

Figure 2: Response

Thus, the **response ratio** must be continuously kept below unity:

$$\rho_{response} = VT_{response}/Y_L$$

2.4.2 Throughput

Obstacles also cannot be avoided unless the vehicle sees them. The vehicle must see all terrain that it will, or can, traverse. Without loss of generality, let a sensor capture one image every T_{cyc} seconds. Let the sensor field of view project onto a distance ΔY_L on the groundplane.

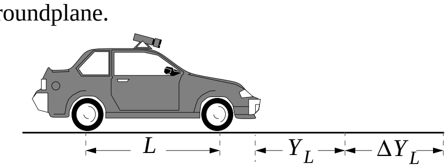

Figure 3: Throughput

To see all obstacles, there can be no gaps in the groundplane coverage of the sensor so the distance moved per frame cannot exceed the groundplane projection. Thus, the **throughput ratio** must be continuously kept below unity:

$$\rho_{throughput} = VT_{cyc}/\Delta Y_L$$

Notice that for both response and throughput ratios, we can fix any one quantity in the ratios and generate an adaptive rule that encodes how the remaining two quantities depend on each other when safety is guaranteed.

2.5 Standard Assumptions

Certain assumptions will be important either because they must be adopted, or because they simplify analysis. These assumptions are not always necessary, justified, or even correct, but we will employ them when they are:

- •**small incidence angle assumption**: the assumption that the perception ratio is small. When adopted, allows us to equate the range to a point on the ground to its groundplane projection with a minimal relative error equal to the square of the perception ratio.
- •**point vehicle assumption**: the assumption that the finite extent of the vehicle can be ignored in the analysis. When adopted, allows us to ignore the extension of the vehicle nose in front of the perception sensor, for example.
- •**low latency assumption**: the assumption that the delays associated with passing energy or information through an element of the system can be ignored. When adopted, allows us to ignore actuator dynamics, for example.
- •**flat terrain assumption:** the assumption that the terrain is at least locally flat at the scale of the sensory lookahead distance. When adopted, allows us to simplify many aspects of the analysis.
- •**smooth terrain assumption:** the assumption that the terrain does not contain any high spatial frequencies. When adopted,

allows to assume reasonable limits on the need to resolve small hazards in the environment.

- **stationary environment assumption**: the assumption that the environment is rigid. When adopted, allows us to measure the position of an object only once and assume that it stays put while the vehicle moves around it.

2.6 Standard Problems

Given the description of the problem outlined above, a set of natural subproblems emerge when one component or another of each ratio does not meet the underlying requirements for fixed values of the other quantities of interest. Many of the following subproblems will be subsequently elaborated in more detail.

2.6.1 Response Problem

The **response problem** is the problem of guaranteeing timely response to external stimuli. Related subproblems include:

- **myopia problem:** The sensor lookahead is too short for a given speed and response time.
- **latency problem:** The response time is too large for a given speed and sensory lookahead.

2.6.2 Throughout Problem

The **throughput problem** is the problem of guaranteeing adequate sensory and processing throughput. It is often the case that raw computing power is insufficient to satisfy this requirement at adequate resolution, but other subproblems can be identified as well:

- **stabilization problem:** Attitude changes of the sensor cause gaps in the sensor coverage.
- **tunnel vision problem:** The sensor field of view is too small for a given vehicle speed and maneuverability, and a given terrain roughness.
- **occlusion problem:** The position of the sensor combined with the roughness of the terrain cause self occlusion of the terrain.

3 Response

This section investigates the manner in which computational reaction time and mechanical maneuverability together determine the ability of a vehicle to avoid obstacles. Analysis of response requires an analysis of the time and space required to react to external events. Up to this point, we have considered that the vehicle travelled at constant speed while executing some undefined obstacle avoidance trajectory.

In a practical model of response, we must consider such matters as the variation of speed with time, the precise trajectory followed including any relevant vehicle dynamics, and the spatial extent of both the vehicle and the obstacle(s). This section considers these matters in detail.

3.1 Response Time

A precise definition of response time requires a precise definition of two discrete events. The first is the event to which the vehicle must respond and the second is the completion of the response trajectory - however it is defined.

It is useful to think about response time in terms of a perceive-think-act loop which models the overall vehicle control and planning system. For the present purpose, we will define the **system response time** as the time period between the instant that an obstacle appears in the field of view of a perception sensor and the instant that the vehicle is considered to have completed execution of the associated avoidance trajectory. This time includes:

- T_{sens} : sensing the environment
- T_{perc} : perceiving what the sensor data means
- T_{plan} : deciding what to do
- T_{cont} : commanding actuators
- T_{act} : actuator response time
- T_{veh} : operating on the vehicle and environment

The following figure presents a potential configuration where one computer is used for intelligent control, and another is used for servo control. Also, several input/output operations are indicated because their delays are significant enough to model.

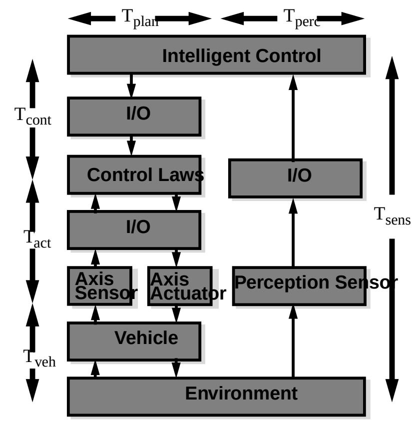

Figure 4: Response Time Elements

The total system response time is therefore:

$$T_{resp} = T_{sens} + T_{perc} + T_{plan} + T_{cont} + T_{act} + T_{veh}$$

It may be useful at times to distinguish the hardware and software components of the response time in order to assess where to make improvements. Thus,

$$T_{sw} = T_{perc} + T_{plan} + T_{cont}$$

$$T_{hw} = T_{sens} + T_{act} + T_{veh}$$

It may also be useful to distinguish the time before actuator response from the time after. The former is the **reaction time** and the latter, the **maneuver time**. Thus,

$$T_{react} = T_{sens} + T_{perc} + T_{plan} + T_{cont} + T_{act}$$

$$T_{maneuver} = T_{veh}$$

The distance travelled during the reaction time tends to be linear in initial velocity while the distance travelled during the maneuver time tends to be quadratic in it.

3.2 Maneuverability

3.2.1 Canonical Maneuvers

Precise analysis of vehicle maneuverability requires solution of the equations of vehicle dynamics under time-varying inputs while accounting for terrain-following loads. For the purpose of the paper, we will often resort to simplified canonical obstacle avoidance trajectories in order to avoid this complexity.

Four special trajectories are defined for a point robot under an assumption of instantaneous and complete response of actuators to their commands:

- **panic stop**: The vehicle is traveling at constant speed in a straight line, decides to fully apply the brakes, and skids or slows to a complete stop.
- **turning stop**: The vehicle is travelling at constant speed along a constant curvature arc, decides to fully apply the brakes, and skids or slows along the original arc to a complete stop.

- **impulse turn**: The vehicle is travelling at constant speed in a straight line, decides to turn at a given radius, and issues the turn command.
- **reverse turn**: The vehicle is travelling at constant speed at the minimum safe turn radius in one direction and issues a command to reverse curvature to the minimum safe radius in the other direction.

These canonical maneuvers are indicated in the following figure.

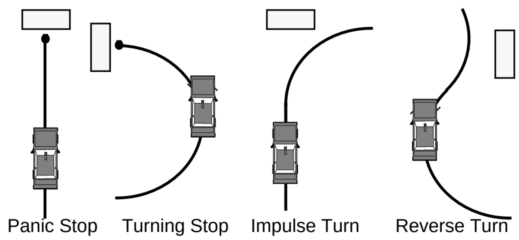

Figure 5: Canonical Maneuvers

3.2.2 Braking

Consider the trajectory followed if a full braking command is issued while travelling at constant speed in a straight line. Let μ be the coefficient of sliding friction, V be the initial velocity, and g be the acceleration due to gravity. Equating the initial kinetic energy to the work done by friction leads to an expression for the **braking distance**:

$$s_{brake} = \frac{V^2}{2\mu g}$$

3.2.3 Turning Radius Limits

Consider the trajectory followed if the vehicle turns at the turn radius which generates the highest safe lateral acceleration. In order to force an analogy with the coefficient of friction for braking, let ν be one-half the maximum permissible lateral acceleration expressed in g's, called the **coefficient of lateral acceleration**. Let V be the initial velocity, and g be the acceleration due to gravity. The **minimum dynamic turn radius** occurs at maximum lateral acceleration and is given by:

$$\rho_{dyn} = \frac{V^2}{2\nu g}$$

Note that many steering mechanisms, including the traditional Ackerman-steered automobile mechanism impose a **minimum kinematic turn radius**, ρ_{kin}. For such vehicles, the operative lower limit on the turn radius is the maximum of these two:

$$\rho_{min} = max(\rho_{dyn}, \rho_{kin})$$

3.2.4 Turning Angle

Consider the trajectory followed if the vehicle executes 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. For a constant curvature turn, the angle subtended at the start point, of the region reachable by the vehicle in a turn, is the yaw of the turn itself as shown below:

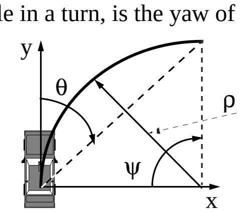

$$\theta = \psi / 2$$

Figure 6: Turning Angle

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, the **turning angle**, is given by:

$$\psi = \frac{s_{turn}}{\rho} = \frac{TV}{\rho}$$

Where s_{turn} is the **turning distance**.

3.3 Response Distance

We have nominally defined the **response distance** as follows:

$$s_{response} = T_{response}V$$

but in a more realistic situation, velocity is not constant throughout a particular trajectory. However, this definition can be retained if the **response velocity** is defined as the ratio of the response distance to the response time.

3.3.1 Panic Stop

Consider a **panic stop** obstacle avoidance trajectory. There is a period of time before the brakes are applied and a period of time after. Before the brakes are applied, the intelligent controller is processing images and deciding on a course of action. For constant velocity, this **reaction distance** is clearly:

$$s_{react} = T_{brake}V$$

The total response distance is the sum of the reaction distance and the braking distance. It expresses the distance travelled from the point where the obstacle first appeared to when the vehicle stops. Thus, for a **panic stop**:

$$s_{response} = T_{brake}V + \frac{V^2}{2\mu g}$$

3.3.2 Impulse Turn

Consider an **impulse turn** obstacle avoidance trajectory. In the worst case, the obstacle spans the entire sensor horizontal field of view and a turn of 90° is required to avoid hitting it. For such a turn, the corresponding distance moved along the original direction of travel is equal to one turn radius. This will be called the **impulse turning distance**.

If we consider the full system reaction time, then there is also a period of time, and associated distance travelled, when the steering has not yet been engaged while the intelligent controller is processing images and deciding on a course of action. For constant velocity, this **reaction distance** is clearly:

$$s_{react} = T_{turn}V$$

The total reaction distance is the sum of these two. It expresses the distance travelled from the point where the obstacle first appeared to when the vehicle completes a 90° turn. Thus, for an **impulse turn** a form analogous to the panic stop is obtained:

$$s_{response} = T_{turn}V + \frac{V^2}{2\nu g}$$

3.3.3 Response Distance

It is possible to define, for the panic stop and impulse turn maneuvers, a general form of the response distance:

$$s_{response} = T_{react}V + \frac{V^2}{2\mu g}$$

where the quantity T_{react} is understood to not include the time spent with the brakes on or turning at the minimum radius. Henceforth, we will write μ to represent the friction or lateral acceleration coefficient as the case requires. The first term can be called the **reaction distance** and the second is the **maneuver distance**.

This relationship is plotted below for typical values of the coefficient of friction or lateral acceleration.

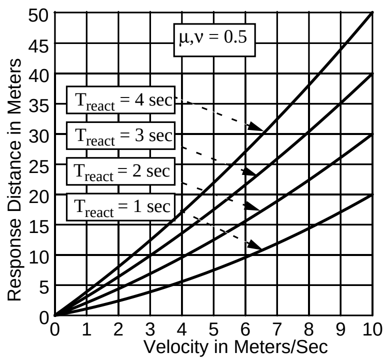

Figure 7: Generic Response Distance

In both cases, we have implicitly assumed that actuator transients can be neglected or absorbed into the reaction time.

All components of the reaction time except the actuator component can normally be considered equal for both braking and turning. However, the reaction distance is often larger for turning than for braking. On Ackerman steered vehicles, time required to complete movement of the steering actuator can often significantly exceed that required for braking. Also, the coefficient of lateral acceleration can be lower than the coefficient of friction because it is limited by the propensity to roll over in a turn.

3.4 Response Angle

3.4.1 Turning Stop

Consider a **turning stop** obstacle avoidance trajectory. For this maneuver, the angle through which the vehicle turns is governed by the braking response distance since the steering actuator does not move. If ρ is the radius of curvature, then the angle turned is:

$$\psi_{response} = \frac{s_{response}}{\rho} = \frac{T_{react}V + V^2/(2\mu g)}{\rho}$$

This will be called the **response angle**. By analogy, it is composed of the **reaction angle** and the **braking angle**.

3.4.2 Response Angle

It is possible to define, for the turning stop maneuver, a general form of the response angle:

$$\psi_{response} = \frac{s_{response}}{\rho}$$

where ρ is the radius of curvature of the turn and $s_{response}$ is the response distance. In the particular case of a turn at the minimum safe radius of curvature, we have:

$$\psi_{response} = \frac{T_{react}V + V^2/(2\mu g)}{max(V^2/(2\nu g), \rho_{kin})}$$

This relationship is plotted below for typical values of the coefficients of friction and lateral acceleration and a minimum kinematic turn radius of 7.5 meters.

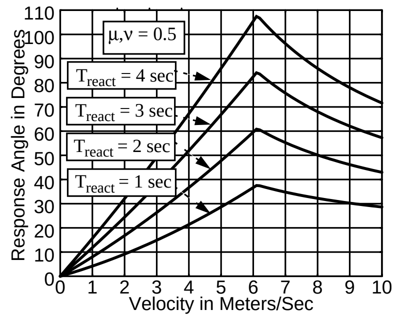

Figure 8: Response Angle

Clearly, the response angle grows roughly linearly while the turn radius is limited by the mechanism. Beyond some velocity (here 6 meters/sec.), the turn radius becomes limited by the lateral acceleration and the response angle decreases.

3.5 Nondimensional Response

One unique characteristic of a high-speed autonomous vehicle is the fact that it can spend as much or more time or distance deciding what to do as it takes to do it. The ratio of the reaction and maneuver distance is therefore a relative measure of how much precious spatial resources are used for each as the vehicle closes on an obstacle. Clearly, there is one unique speed where the reaction distance and the maneuver distance become equal. Let us define the **maneuver coefficient** $\bar{s}$ as their ratio. Thus:

$$\bar{s} = \left[\frac{V^2}{2\mu g}\right] / [T_{react}V] = \frac{V}{2\mu g T_{react}}$$

When this quantity is significantly less than one, the reaction distance dominates the maneuver, and the overall response distance is basically linear in initial velocity. This is the case for most of Figure 7. Note also, that the coefficient is also the ratio between the reaction angle and the braking angle.

For the turning stop, the limits on turn radius may be driven by either mechanical or dynamic concerns. Let us define the **turning coefficient** as the ratio of the kinematic and dynamic limits.

$$\bar{t} = \rho_{kin}/\rho_{dyn} = \rho_{kin}/\frac{V^2}{2\nu g}$$

If this ratio is less than unity, the response angle grows linearly with velocity. If it exceeds unity, the response angle decreases quadratically with velocity.

3.6 Response Regimes

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

$$s_{response} = T_{react}V[1 + \bar{s}]$$

Based on the response coefficient, two regimes of operation can be defined. In the **kinematic response regime** it is much less than unity. In the **dynamic response regime** it is much greater than unity. When the maneuver coefficient is unity, reaction distance and maneuver distance are equal. At this point, response distance enters a regime of quadratic growth with initial velocity. As speeds increase there comes a point where the system must explicitly reason about the “dynamics” of maneuvering in the sense that the maneuver distance is no longer an insignificant part of the overall response trajectory.

4 Throughput

This section investigates the manner in which computational cycle time, maneuverability, and sensor field of view determine the ability of a vehicle to measure the environment fast enough and comprehensively enough to avoid missing anything.

We will investigate the relationship between the maneuverability of the vehicle and the sensor **field of regard**. The sensor field of regard will be described by:

- **depth of field**: minimum and maximum range
- **field of view:** horizontal and vertical field of view

4.1 Depth of Field

There are many potential ways to determine requirements on sensor range. This section will propose one plausible way based on limits on response distance. For the sake of simplicity, we will work in terms of the distance y from the sensor measured in the groundplane, rather than distance R in the plane formed by a sensor scanline. Recall that the distance from the sensor to the nose of the vehicle is given by p.

4.1.1 Minimum Range

We could determine minimum required sensor range from the minimum response distance of any obstacle avoidance trajectory. This approach would be based on the argument that the vehicle is already committed to travel at least this far. In many cases, the **panic stop** is the trajectory that consumes the least space. Such an analysis would give a minimum useful range of:

$$Y_{min} = p + s_{brake} = p + T_{brake}V[1 + \bar{b}]$$

Where T_{brake} is the **braking reaction time** and $\bar{b}$ is the associated **braking maneuver coefficient** equal to the ratio of braking distance to reaction distance.

4.1.2 Maximum Range

Likewise, we could determine maximum range from the maximum response distance associated with any obstacle avoidance trajectory based on the argument that the vehicle cannot travel any further before another computational cycle of obstacle avoidance. We will consider the impulse turn to be the trajectory that consumes the most space. This would give a maximum useful range of:

$$Y_{max} = p + s_{turn} = p + T_{turn}V[1 + \bar{t}]$$

Where T_{turn} is the **turning reaction time** and $\bar{t}$ is the associated **turning maneuver coefficient** equal to the ratio of turning distance to reaction distance.

4.2 Horizontal Field of View

The horizontal field of view will be determined by the turning stop maneuver and hence by twice the response angle. The rationale for this choice is that when the vehicle is executing a turn, it will have just enough sensory lookahead to stop if an obstacle appears. Another important matter to consider is that a sensor normally cannot change its horizontal field of view dynamically, so it is necessary to allocate horizontal field of view for a range of velocities.

$$HFOV = 2 \times max_V[\psi_{response}]$$

This may mean that even though the field of view requirements reduce as speeds increase, a typical sensor cannot take advantage of it.

4.2.1 Tunnel Vision Problem

Although the required horizontal field of view does eventually decrease with velocity, contemporary sensors generally do not generate the field of view necessary to image all reachable terrain. Consider the following figure. If the steering wheel turns at constant speed, for the maneuvers indicated, the entire region that the vehicle can reach is contained within the set of curves shown. Each curve corresponds to an alternative steering angle.

It is often the case that a contemporary autonomous system cannot look where it is going. At times, there may be no overlap at all between the projection of the field of view on the groundplane, and the region that the vehicle is committed to travelling. This problem will be called the **tunnel vision problem**.

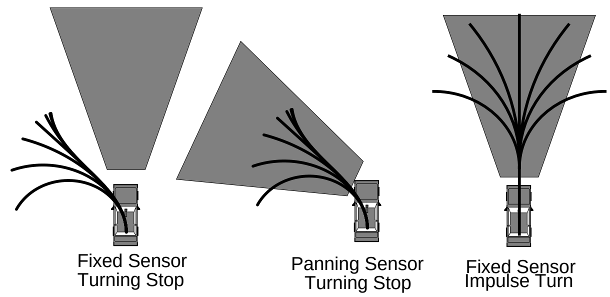

Figure 9: Tunnel Vision

In the center figure, it is clear that sensor panning can be an effective solution. Nonetheless, it is important to distinguish the width of the field of view from its direction. While pointing may help, there is also a minimum width that covers all reachable terrain. In the right figure, the initial curvature is zero, and the field of view is not wide enough regardless of where it is pointed.

For fixed sensors, overall latency severely complicates this problem. If the vehicle turns with angular velocity $\dot{\psi}$ and the horizontal field of view is small, it is not unusual for the vehicle to have driven completely off of the imaged terrain by the time that the data is processed. If the overall system reaction time is T_{react}, then by the time that a command reaches the hardware, the vehicle has turned through an angle:

$$\Delta\psi = \dot{\psi}T_{react}$$

This angle can easily exceed the available field of view.

4.3 Vertical Field of View

There are several potential mechanisms that might be used to determine requirements on the vertical field of view. The major kinematic requirement which influences the vertical field of view is the pitch angle induced in the vehicle body by the most challenging, yet navigable, terrain. On the other hand, we might choose the vertical field of view based on the overall sensory throughput required. Both options are considered below.

4.3.1 Worst Kinematic Case - Rough Terrain

Rough terrain considerations generate the worst case requirement on vertical field of view. Under the strong form of guaranteed safety, we can assume that there is no need to view terrain that cannot be traversed. Let the highest safe body pitch angle be θ. The following figure illustrates the two extreme 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.

On this kinematic basis, the vertical field of view required is four

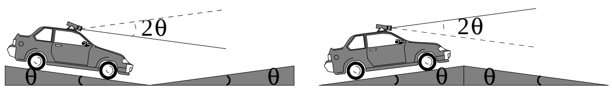

Figure 10: Rough Terrain Vertical Field of View

times the maximum pitch of the body.

$$VFOV = 4\theta$$

4.3.2 Best Kinematic Case - Flat Terrain

Under the **small incidence angle assumption**, the vertical field of view can be expressed in terms of the maximum and minimum range as follows:

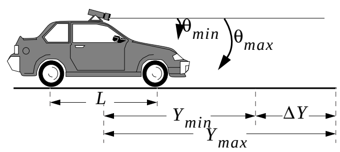

$$VFOV = \theta_{max} - \theta_{min}$$

$$VFOV = \frac{h}{Y_{max}} - \frac{h}{Y_{min}}$$

$$VFOV = \frac{h\Delta Y}{Y_{max}Y_{min}}$$

Figure 11: Flat Terrain Vertical Field of View

4.3.3 Best Dynamic Case - Flat Terrain

This section considers a dynamic basis for specifying the required vertical field of view in the sense that the result is dependent on velocity instead of angles. Accordingly, we will find it important to think in terms of a sensor measurement rate instead of the angular VOFV.

Let Y be the average of Y_{max} and Y_{min}. When ΔY is small compared to Y, we have:

$$Y_{max}Y_{min} \approx Y^2$$

If the average range is used to approximate the maximum and the minimum range, the required vertical field of view, from the above expression, is simply:

$$VFOV = h\frac{(\Delta Y)}{Y^2}$$

Recall that the **throughput ratio** is defined as the ratio of the distance travelled by the vehicle to the amount of terrain measured in the same unit of time:

$$\rho_{throughput} = VT_{cyc}/\Delta Y$$

Substitution yields the VFOV in terms of the throughput ratio:

$$VFOV = h\frac{VT_{cyc}}{\rho_{throughput}Y^2}$$

The **imaging density** σ_I will be defined as the average number of images that fall on any patch of terrain. It is the reciprocal of the **throughput ratio**:

$$\sigma_I = 1/\rho_{throughput}$$

The **sweep rate**, $\dot{\theta}$, of a sensor can be defined, in image space, as the vertical field of view (VFOV) generated per unit time. It has units of angular velocity. It may be related to the physical motion of the elevation mirror in a laser rangefinder or the product of the VFOV and the frame rate for a video camera. Rewriting the above, we have:

$$\dot{\theta} = \frac{VFOV}{T_{cyc}} = h\frac{V}{\rho_{throughput}Y^2} = \frac{\sigma_I hV}{Y^2}$$

Under **guaranteed throughput,** the throughput ratio is unity or lower, and the imaging density is correspondingly unity or higher, so the sweep rate must always exceed:

$$\dot{\theta} \ge \frac{hV}{Y^2}$$

We will call this relationship the linear velocity component of the **sweep rate rule**. If range Y is related to stopping distance, the sweep rate can be expressed solely in terms of reaction time, sensor height, and velocity - making it a function only of vehicle parameters and state.

4.3.4 Worst Dynamic Case - Rough Terrain

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 caused by terrain following loads, then the sweep rate rule becomes:

$$\dot{\theta} = \dot{\theta}_{max} + \frac{hV}{Y^2}$$

A final consideration in determining sweep rate and VFOV is the gradient of the terrain in front of the vehicle. The terrain gradient is unconstrained in general, and not usually known a priori. If a maximum terrain gradient can be specified, it can be used in the linear sweep rate expression. Such a maximum may be determined either from the a priori characteristics of the terrain, or from considering that terrain that is not navigable need only be imaged to the degree necessary to classify it as not navigable.

4.3.5 Stabilization Problem

Notice that the linear component of the sweep rate rule benefits from higher linear speeds whereas the angular component suffers from higher angular speeds. If the VFOV is either too small or too slowly adjustable to avoid holes in the coverage of the sensor, the situation will be known as the **stabilization problem**. This consideration, when it occurs, argues for a wider VFOV.

On the other hand, if all information in an image is processed, the required computational speed increases directly with VFOV. This has been called the perceptual **throughput problem**. Any fixed VFOV is a compromise between these two considerations of computing less than necessary or more than is feasible.

4.4 Occlusion

This section investigates the relationship between vehicle configuration and the prevalence of terrain self-occlusions. 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**.

4.4.1 Hill Occlusion

A hill can also be called a **positive obstacle**. Ideally, a sensor should see behind a navigable hill at the maximum sensor range. The necessary 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:

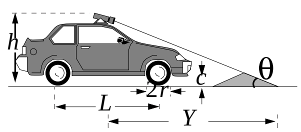

Figure 12: Hill Occlusion

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**. This rule is almost always violated because it is impractical to mount a sensor at the required height. The **perception ratio,** h/R, approximately h/Y, 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.

4.4.2 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.

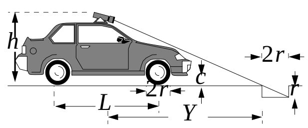

Figure 13: Hole Occlusion

Detecting that the hole is deep enough to present a hazard requires waiting until the hole was close enough to satisfy:

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

This will be called the **hole occlusion rule**. While properly placed scanlines could detect the hole, it is also the case that obstacles inside the stopping distance cannot be avoided at all.

Practical hole detection must be based on subtler cues than interior geometry for high speed vehicles. For example, holes generate range shadows beyond the leading edge.

5 Conclusions

Requirements analysis is an activity that attempts to study the problem rather than any particular solution. This paper has analyzed some of the requirements of high speed autonomous mobility in general terms and has supported the following conclusions about the nature of the problem.

5.1 Sensor Mounting Geometry

Sensor height is typically an order of magnitude smaller than the vehicle response distance. This relationship has many implications relating to the prevalence of occlusions in images and the complexity of image processing algorithms.

5.2 Response

For both the panic stop and the impulse turn trajectories defined earlier, the response distance consists of a linear and a quadratic velocity term. The linear term dominates at lower velocities and the quadratic term at higher velocities. Two regimes of operation, termed kinematic and dynamic are thus defined.

5.3 Field of View

Rational methods for specifying sensor field of view are available if we attempt to image all reachable terrain. However, by assessing the quantitative requirements it is clear that the tunnel vision, myopia, latency, stabilization, and occlusion problems are all severe in typical situations with typical hardware. Perhaps coincidently, requirements on field of view decrease quadratically with velocity after a certain point is reached. For horizontal field of view that point is reached when turns are limited by considerations of rollover (lateral acceleration). For vertical field of view, the point is reached when range significantly exceeds sensor height and the same length on the groundplane projects onto an ever smaller region in the image plane as range increases.

5.4 Obstacle Avoidance

From the perspective of reliability in obstacle detection and avoidance, it is important to recognize that the planning horizon of obstacle avoidance (i.e. the vehicle reaction time) is roughly equal to the time it takes for the steering actuator to reverse direction. Hence, the system operates almost entirely in the "transient" regime of continuously changing curvature when aggressively avoiding obstacles.

6 Acknowledgments

This research was sponsored by ARPA under contracts "Perception for Outdoor Navigation" (contract number DACA76-89-C0014, monitored by the US Army Topographic Engineering Center) and "Unmanned Ground Vehicle System" (contract number DAAE07-90-C-R059, monitored by TACOM).

7 References

[1] M. Daily et al., "Autonomous Cross Country Navigation with the ALV", In Proc. of the 1988 IEEE International Conference on Robotics and Automation, Philadelphia, Pa, April 1988; pp. 718-726.

[2] H. P. Moravec, "Obstacle Avoidance and Navigation in the Real World by a Seeing Robot Rover", Ph. D. Thesis, Stanford University, 1980.

[3] A. J. Kelly, "An Intelligent Predictive Control Approach to the High-Speed, Cross Country Autonomous Navigation Problem", Ph. D. thesis, Robotics Institute, Carnegie Mellon University, June 1995.

[4] A. Kelly, and A. Stentz, "Analysis of Requirements of Rough Terrain Autonomous Mobility: Part II: Resolution and Accuracy", IEEE International Conference on Robotics and Automation, Albuquerque, New Mexico, April 1997.

[5] A. Kelly, and A. Stentz, "Computational Complexity of Terrain Mapping Perception in Autonomous Mobility", IEEE International Conference on Robotics and Automation, Albuquerque, New Mexico, April 1997.

[6] L. Matthies, A. Kelly, T. Litwin, "Obstacle Detection for Unmanned Ground Vehicles: A Progress Report.", to appear in IEEE International Symposium on Intelligent Vehicles.

[7] L. Matthies, "Stereo Vision for Planetary Rovers", International Journal of Computer Vision, 8:1, 71-91, 1992.

[8] R. T. Dunlay and D. G. Morgenthaler, "Obstacle Detection and Avoidance from Range Data", In Proc. SPIE Mobile Robots Conference, Cambridge, Mass., 1986

[9] E.D. Dickmanns. "Dynamic computer vision for mobile robot control", Proceedings of the 19th International Symposium and Exposition on Robots, pp. 314-27;

[10] K. E. Olin and David Tseng. "Autonomous Cross Country Navigation", IEEE Expert, August 1991.

[11] 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

[12] In So Kweon, "Modelling Rugged Terrain by Mobile Robots with Multiple Sensors", CMU PhD Thesis, 1990