An Analysis of Requirements for Rough Terrain Autonomous Mobility

Alonzo Kelly and Anthony Stentz
Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213
alonzo@ri.cmu.edu
axs@ri.cmu.edu

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, throughput, and the resolution and accuracy of sensors and computations. Several nondimensional expressions emerge which characterize requirements in canonical form.

The automatic generation of dense geometric models for autonomously navigating vehicles is a computationally expensive process. Using first principles, it is possible to quantify the relationship between the raw throughput required of the perception system and the maximum safely achievable speed of the vehicle. We show that terrain mapping perception is of polynomial complexity in the response distance. To the degree that geometric perception consumes time, it also degrades real-time response characteristics. Given this relationship, several strategies of adaptive geometric perception arise which are practical for autonomous vehicles.

Keywords: mobile robots, autonomous vehicles, rough terrain mobility, terrain mapping, obstacle avoidance, goal-seeking, trajectory generation, requirements analysis

Table of Contents

1 Introduction

This paper is concerned with the requirements that must be satisfied by an autonomous vehicle which operates safely in its environment. A typical autonomous vehicle has been fitted with low level vehicle-specific control loops to enable computer control of propulsion, steering, and brakes. Some position estimation system is typically incorporated to determine position. At least one perception sensor is needed to enable it to perceive its environment.

For most of the purposes of this paper, the perception sensor can be any imaging sensor measuring range or intensity in any electromagnetic band of frequencies. Indeed, while we will mention sensor field of view, a line scanned sensor can be considered a special case for which all of the results apply.

The need for high throughput perception algorithms has been acknowledged for some time in the field of autonomous vehicle navigation. Yet, the evidence for this need has not been based on any underlying theory. The problem addressed in this paper is the lack of such a theory.

This paper proposes a rudimentary theory of obstacle avoidance and uses it to quantify the complexity of terrain mapping perception in autonomous vehicles under a set of assumptions that render the problem tractable.

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**, and will appear in a glossary at the end of the paper for easy reference.

2.1.2 Nomenclature

In this paper, 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 the direction of travel.
- y - **downrange**, in the groundplane, along the direction of travel.
- 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.

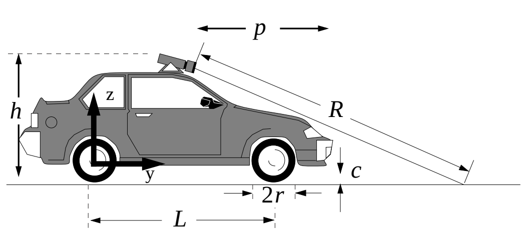

Figure 1 Important Dimensions

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 prevelance 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 prevalance of terrain self occulsions.

2.4 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**.

2.4.1 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.

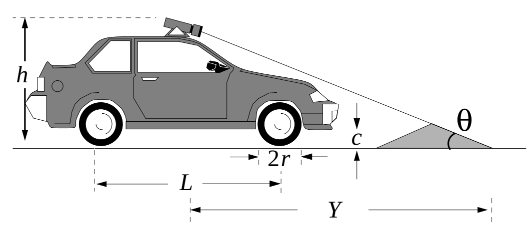

Figure 2 Hill Occlusion

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

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

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**. 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.

2.4.2 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.

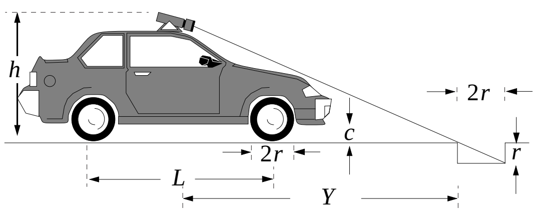

Figure 3 Hole Occlusion

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:

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

This will be called the **hole occlusion rule**. While properly placed scanlines could indeed 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.

2.5 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.

2.5.1 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.

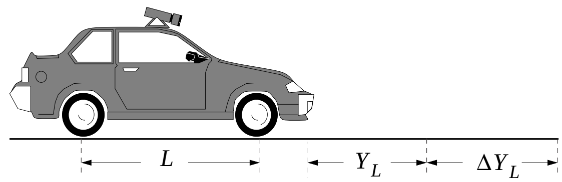

Figure 4 Response

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

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

2.5.2 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.

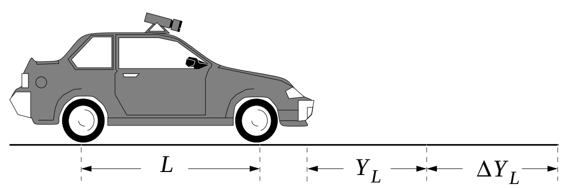

Figure 5 Throuhput

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 less than 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.

2.5.3 Acuity

Obstacles cannot be avoided unless the system can reliably detect them. Reliability in obstacle detection is at least a question of the spatial resolution of the sensor pixel footprint. However, a larger vehicle requires a larger obstacle to challenge it, so it is natural to normalize the spatial resolution of the sensor by a characteristic vehicle dimension.

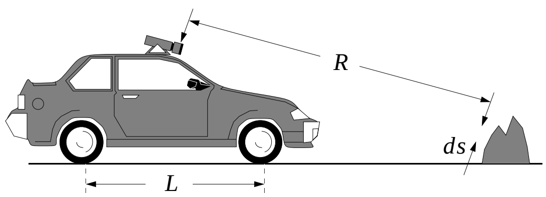

Figure 6 Acuity

The **acuity ratio** will be defined as:

$$\rho_{acuity} = ds/L$$

In order to resolve a difference in the size of an environmental feature that is as small as the vehicle dimension chosen, the acuity ratio must be kept, by the sampling theorem, below one-half.

2.5.4 Fidelity

Obstacles cannot be avoided unless the system can locate them sufficiently accurately with respect to itself and execute an avoidance trajectory sufficiently accurately. In this context, “sufficiently accurately” depends on the size of the vehicle and the spacing between obstacles in some average, worst-case, or other useful sense.

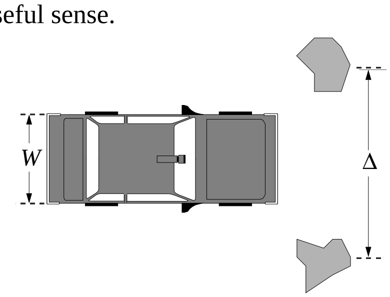

Figure 7 Fidelity

The **fidelity ratio** will be defined as:

$$\rho_{fidelity} = ds/(\Delta - W)$$

where ds is the error between the intended and actual paths of the vehicle. This quantity depends on the accuracy of the perception sensor used to locate the vehicle relative to obstacles, the position estimation system, and the command following controls.

The margin for error available when driving exactly between two separated obstacles is half the difference between the obstacle spacing and the vehicle dimension aligned between them. That is, the fidelity ratio must be kept below one-half.

2.6 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 t o measure the position of an object only once and assume that it stays put while the vehicle moves around it.

2.7 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.

2.7.1 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.

2.7.2 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.

2.7.3 Acuity Problem

The **acuity problem** is that of guaranteeing detection of obstacles. It is often the case that sensor intrinsic angular or range resolution is inadequate for a given lookahead distance but other subproblems can be identified as well:

- **sampling problem:** Unfavorable variation in the size, density, or shape of sensor pixels due to terrain shape, sensor mounting configuration, and radiometric considerations.
- **motion distortion problem:** Distortion of images due to the motion of the vehicle during image acquisition.

2.7.4 Fidelity problem

The **fidelity problem** is that of guaranteeing adequate fidelity of models and measurements. Several subproblems can also be identified:

- **sensitivity problem:** Extreme sensitivity of changes in one quantity to small changes in another.
- **registration problem:** Inability to match redundant measurements of the environment due to errors in the measurements.
- **command following problem:** Inability of the vehicle control systems to cause the vehicle to execute its commands sufficiently well.
- **stability problem:** Instability of obstacle avoidance and/or goal seeking due to the use of insufficiently accurate models.

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 that they should be modeled.

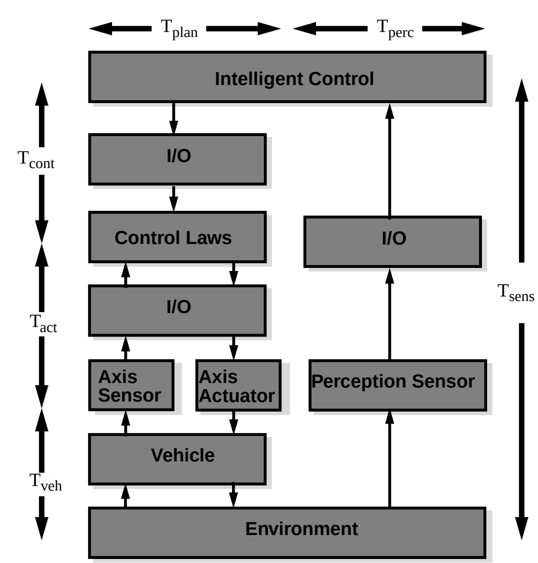

Figure 8 Response Time Elements

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.

3.2 Maneuverability

3.2.1 Canonical Maneuvers

The **actuation space** of a vehicle is spanned by a command vector whose elements may be steering, throttle, brake, or perhaps individual wheel velocities. Each of these elements is a time-continuous function and many vehicles have nontrivial dynamics. Hence, 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.

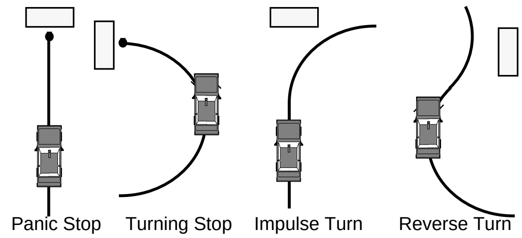

Figure 9 Canonical Maneuvers

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:

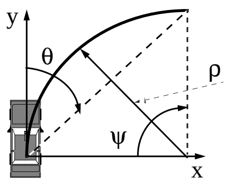

$$\theta = \psi/2$$

Figure 10 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.

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, on conventional (Ackerman) steered vehicles, the 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. In short, the reaction distance is larger for turning than for braking.

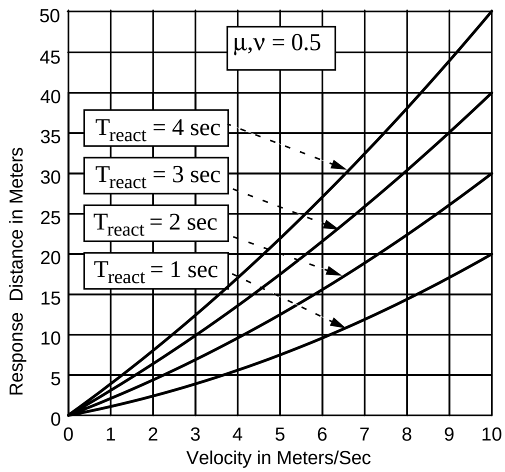

Figure 11 Generic Response Distance

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:

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

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:

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

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.

Clearly, the response angle grows roughly linearly while the turn radius is limited by the mechanism.

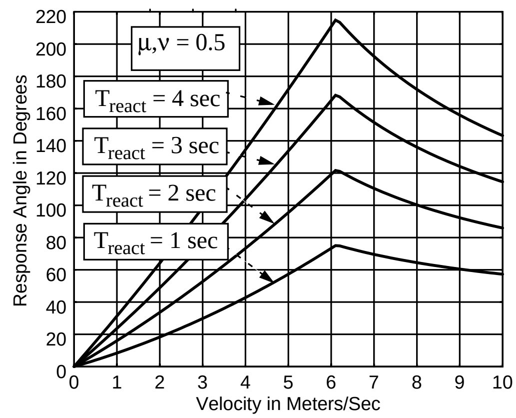

Figure 12 Response Angle

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 the ratio of these two. 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 11. 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.

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

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 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:.

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

Where T_{turn} is the **turning reaction time** and $\dot{t}$ is the associated **turning 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 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.

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

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 horizontal field of view does decrease with velocity beyond some point, contemporary sensors generally do not generate the field of view necessary to image all reachable terrain. Consider the following figure in which the vehicle is initially turning to the left. If the steering wheel turns at constant speed, 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**.

This consideration argues for mechanisms which physically point the narrow field of view sensor. In the right of the figure, it is clear that sensor panning is an effective approach in this situation.

Nonetheless, it is important to distinguish the width of the field of view from its angular position. While pointing may help, there is also a minimum width that covers all reachable terrain. In the figure below, the

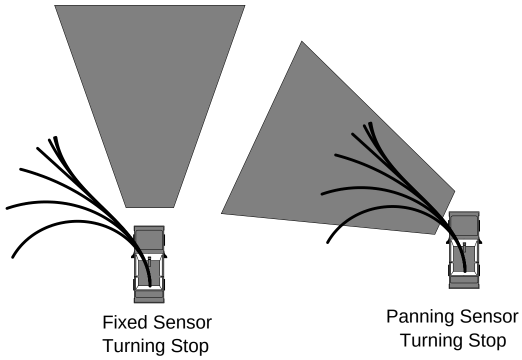

Figure 13 Tunnel Vision - Initial Curvature

initial curvature is zero. In this case, the field of view is not wide enough regardless of where it is pointed.

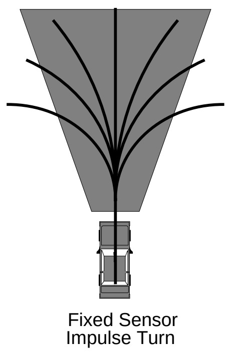

Figure 14 Tunnel Vision - No Initial Curvature

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.

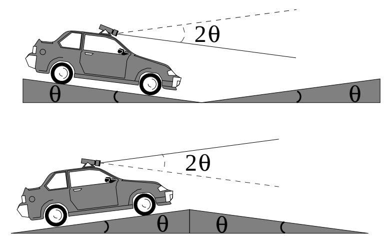

Figure 15 Rough Terrain Vertical Field of View

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

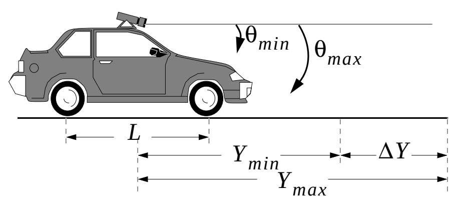

$$VFOV = \theta_{max} - \theta_{min} = \frac{h}{Y_{max}} - \frac{h}{Y_{min}} = \frac{h\Delta Y}{Y_{max}Y_{min}}$$

Figure 16 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:

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

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 **throughput problem**. Any fixed VFOV is a compromise between these two considerations of computing less than necessary or more than is feasible.

5 Acuity

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 a **flat terrain assumption** so it is not entirely correct in rough terrain. Nonetheless, it is a useful theoretical approximation.

5.1 Acuity Limits

The size of a spatial feature that presents an obstacle to a vehicle has both an upper and a lower useful limit. The largest feature of interest is one the size of the vehicle wheelbase because this is the lowest resolution that still allows the vehicle pitch angle to be predicted. At resolutions below this, the entire vehicle is smaller than the sensor resolution and vehicle pitch cannot be resolved. This lower useful limit on acuity will be called **minimum acuity**. Based on earlier comments on acuity, we can express this limit in terms of the wheelbase as follows:

$$\rho_{acuity} = \frac{1}{2} = ds/L \qquad ds = L/2$$

Another important form of obstacle is one which could collide with or trap a tire at operating velocity such as a pothole or step. The ability to resolve a spatial feature on the order of the size of a wheel radius is needed to ensure that a wheel does not fall in a hole or drive over a step which would cause damage. This upper limit on acuity will be called **maximum acuity**. Based on earlier comments, we can express this limit in terms of the wheel radius as follows:

$$\rho_{acuity} = \frac{1}{2} = ds/r \qquad ds = r/2$$

While it can be argued that the smallest feature of interest is one the size of a nail, this leads to results that are impossible to achieve so we will assume that such 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 resolution limit. This will be called the **smooth terrain assumption**.

5.2 Obstacle Sampling Factor

Formally, resolution is the smallest difference that a system can resolve. Thus, the acuity problem is to reliably distinguish a spatial feature of a given size from one somewhat smaller. The choice of what is to be considered somewhat smaller is arbitrary but it relates directly to reliability of obstacle detection.

Consider the following figure in which an obstacle appears in the field of view of a sensor. The obstacle is of height h. We will define a one-dimensional obstacle **sampling factor** $n/2$ as half the number of pixels that intersect the obstacle in a particular direction.

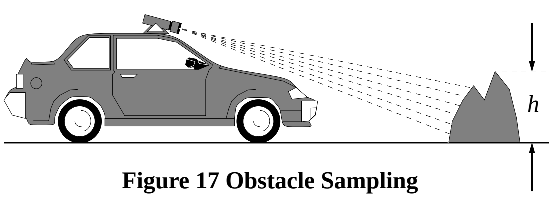

Figure 17 Obstacle Sampling

The spatial resolution of the system is governed by the sampling theorem. The sensor can distinguish a difference δ in obstacle size no smaller than:

$$\delta = 2ds = 2\frac{h}{n}$$

Thus, the sampling theorem is just satisfied for a given feature size when the sampling factor is unity. One measure of reliability in obstacle detection is the frequency of false positives and false negatives and both of these measures can be expected to improve as the sensor spatial resolution exceeds the amount required by the sampling theorem, or equivalently, as the sampling factor increases.

5.3 Differential Imaging Kinematics

The relationships between pixel angular width and its projections onto three orthogonal axes are approximated below *for flat terrain*. In the crossrange direction the following figure applies:

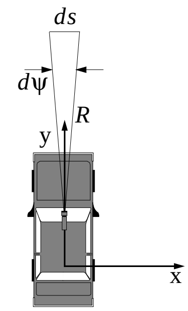

$$x = R\sin\psi$$
$$ds = Rd\psi$$
$$dx = ds = Rd\psi$$
$$d\psi = dx/R$$

Figure 18 Crossrange Resolution

In the downrange and vertical directions, the following figure applies:

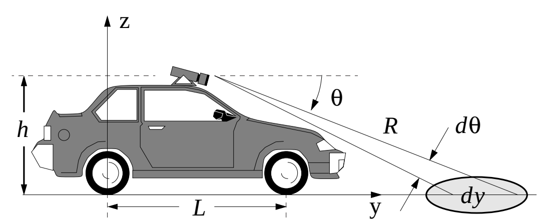

$$R = \frac{h}{\sin\theta}$$

$$y = R\cos\theta = h/\tan\theta$$

$$z = h - R\sin\theta = 0$$

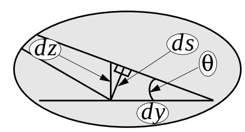

$$dy = \frac{ds}{\sin\theta} = \frac{hd\theta}{\sin^2\theta} \qquad ds = Rd\theta = \left(\frac{h}{\sin\theta}\right)d\theta$$

$$dz = \frac{ds}{\cos\theta} = \frac{hd\theta}{\cos\theta\sin\theta}$$

$$d\theta = \frac{dy\sin^2\theta}{h} = \frac{dz\cos\theta\sin\theta}{h}$$

Figure 19 Downrange Resolution

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

$$dy = \frac{ds}{\sin\theta} \approx \frac{Rd\theta}{(h/R)} \qquad dx = dz = Rd\theta$$

These approximations will be called the resolution transforms and used extensively throughout the rest of this section.

5.4 Sampling Problem

It is clear from the previous expressions that the size of pixels projected on the ground plane varies linearly with range in the crossrange and vertical directions while it varies quadratically in the downrange direction. The **perception ratio** has appeared in the denominator of the dy expression and because it is normally smaller than unity, the downrange pixel projection is normally largest.

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 orders of magnitude, and it varies with both position and direction. Significant variation in groundplane resolution can cause undersampling at far ranges and oversampling close to the vehicle. This problem will be called the **sampling problem**.

5.4.1 Pixel Footprint Area and Density Nonuniformity

Multiplying the above expressions:

$$dxdy = Rd\theta\frac{Rd\theta}{(h/R)} = \frac{R^2 d\theta^2}{(h/R)}$$

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

5.4.2 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 is given by the perception ratio. This result expresses the variation of pixel size with direction.

5.5 Acuity Limits in Image Space

This section develops expressions for sensor angular resolution requirments based on vehicle dimensions and sensory lookahead. For reasons of simplicity, we will define sensor angular resolution in this section as the smallest difference in sensor pixel azimuth and elevation that can be resolved. It is important to distinguish this definition from the angle subtended by the smallest obstacle that can be resolved. The quantum of motion or measurement of pixel angle may not be related to the angle subtended by a pixel in the case of a laser rangefinder.

5.5.1 Minimum Acuity

When R >> h, the downrange projection of a pixel significantly exceeds the crossrange projection. Consider what happens when the downrange spacing between pixels begins to approach the size of the vehicle itself.

The ability to resolve vehicle pitch angle from terrain data 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. At resolutions below this level, sensor data contains no useful information at all.

Equating downrange resolution to one-half the wheelbase and substituting the resolution transforms

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

Rewriting gives the following relationship that relates two key nondimensional variables and relates the vehicle shape and lookahead distance to the required sensor angular resolution:

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

The lowest useful resolution occurs when the product of the **normalized wheelbase** and the **perception ratio** equals one-half the angular resolution of the sensor. This is an image space expression of the **minimum sensor acuity rule**. Any of the variables can be considered to be absolutely limited by the others in the expression.

5.5.2 Maximum Acuity

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 two pixels on a vertical surface, equal in height to the wheel radius, at any given range.

Equating vertical resolution to one-half the wheel radius and substituting the resolution transforms

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

Rewriting gives the following relationship that relates the vehicle shape and lookahead distance to the required sensor angular resolution:

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

The highest useful resolution occurs when the ratio of wheel radius to range equals one-half the angular resolution of the sensor. This is an image space expression of the **maximum sensor acuity rule**. Again, any of the variables can be considered to be absolutely limited by the others in the expression.

5.5.3 Relative Importance of Acuity Limits

Notice that the minimum rule is quadratic in 1/R, whereas the maximum rule is linear. Both constraints are equal when:

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

At long ranges, the minimum acuity limit actually dominates the maximum limit. Solving the minimum acuity expression for range gives an expression for the maximum useful range of a sensor:

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

The condition that this range is small compared to that required by response considerations has been called the **myopia 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 terrain 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.

5.6 Motion Distortion Problem

By the time an image is received by the perception system, the vehicle may have moved a considerable distance since the image was acquired. So, the processing of the geometry in the image must account for the exact position of the vehicle when the image was taken. Further, some sensors such as scanning laser rangefinders may require significant time to scan the laser beam over the environment. In the worst case, there must be a distinct vehicle pose associated with each pixel in a ladar image. If this motion distortion is not corrected, the terrain maps computed from images will be grossly in error.

The worst case is a high angular velocity turn as indicated in the figure below. Suppose the input latency of a range image is 0.5 secs, that rangefinder scanning takes a further 0.5 secs, and that the vehicle is travelling at 6 mph and turning sharply, so its angular velocity is 1 rad/sec. If this motion is not accounted for, all of the following effects will occur:

- objects will be smeared by $30°$ in the image
- objects will be shifted by $30°$ in their percieved location
- the range to an object will also be overestimated by the distance travelled in 1 second.

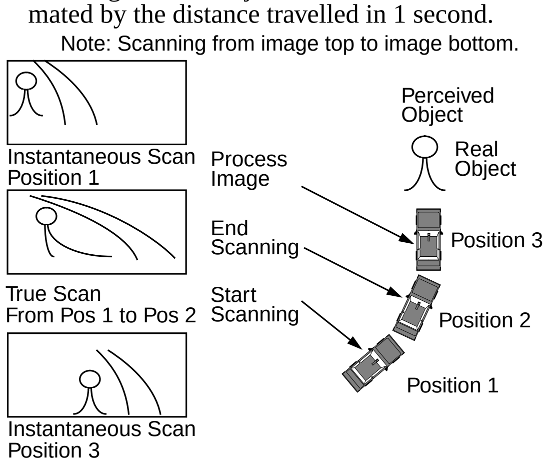

Figure 20 Motion Distortion Problem

This distortion of range images can be removed by maintaining a history of vehicle poses sampled at regular intervals for the last few minutes of execution and searching this list for the precise vehicle position at which each range pixel was measured.

6 Fidelity

This section investigates the manner in which the accuracy of models of vehicle maneuverability determine the ability of a vehicle to operate robustly.

6.1 Modeling Dynamics and Delays

In the context of high-speed motion, the time it takes to pass information into and out of the system becomes a significant factor. Any delays in time which are not modeled are ultimately reflected as errors between both:

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

Time delays, also called **latencies**, may arise in general from several sources - all of which occur in a contemporary autonomous system:

- **sensor dwell latency** is the time it really takes for a measurement to be acquired even though it is often a nominally instantaneous process.
- **communication latency** is the time it takes to pass information between system processes and processors.
- **processing latency** is the time it takes for an algorithm to transform its inputs into its outputs.
- **plant dynamics latency** is the delay that arises in physical systems because they are governed by differential equations.

Feedback controllers often cannot significantly reduce the raw delay associated with response of actuators and the vehicle body. While delays affect response directly, they also affect the ability of the system to localize obstacles correctly if they are not modeled in perceptual processing. This section investigates these matters in the context of high-speed motion.

6.1.1 Latency Problem

Unmodeled latencies in both sensors and actuators can cause the vehicle to both underestimate the distance to an obstacle and underestimate the distance required to react. This behavior is indicated in the following figure. When latencies are modeled, the system is aware of its closer proximity to the obstacle and its reduced ability to turn sharply. In the following scenario, it should choose an alternative obstacle avoidance trajectory to avoid collision.

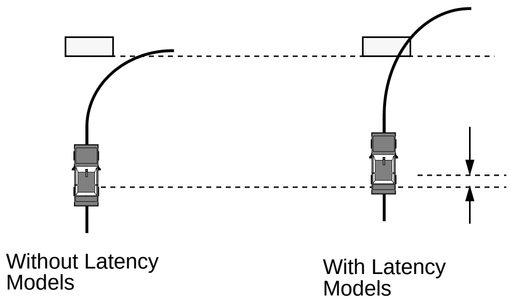

Figure 21 Unmodelled Latency Problem

6.1.2 Minimum Significant Delay 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. In the case of dynamic systems characterized by a differential equation, the **time constant** τ is a related concept.

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 relatively small. This **low latency assumption** is not correct for high-speed autonomy above some speed.

Let a time delay of Δt occur which is not modeled 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 either a range pixel or the vehicle to an accuracy of δ, the **minimum significant delay** occurs when the fidelity ratio is unity, or when:

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

6.1.3 Normalized Time Constant

Motion planners operating on a mission level may find it convenient to abstract away the dynamics of the problem for reasons of efficiency or irrelevance. However, obstacle avoidance must be aware that a steering actuator may not reach its commanded position before an obstacle is reached because this will dramatically affect the trajectory followed. This spectrum can be formalized roughly with a quantity called the **normalized time constant**:

$$\bar{\tau} = \frac{\tau}{T_{look}} = \frac{T_{act}}{T_{look}}$$

where T_{look} is the **temporal planning horizon** or the amount of time the system component is looking ahead in its deliberations.

When the normalized time constant is small, dynamics are not important but when it approaches or exceeds unity, dynamics are a central issue.

6.2 Ackerman Steering Kinematics

6.2.1 Bicycle Model

It is useful to approximate the kinematics of the Ackerman 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.

$$\frac{1}{\rho} = \frac{\tan\alpha}{L} = \frac{d\beta}{ds}$$

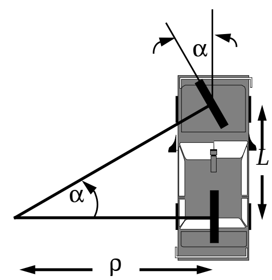

Figure 22 Bicycle Model

Rotation rate is obtained from the speed V as:

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

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

$$\alpha = \operatorname{atan}\left(\frac{L\dot{\beta}}{V}\right) = \operatorname{atan}(\kappa L)$$

When the dependence on time of inputs and outputs is represented explicitly, this steering mechanism is modeled by a coupled nonlinear differential equation thus:

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

6.2.2 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 controlled. 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 **dead reckoning** in navigation. For example, the integral and differential equations which map A-space to C-space in a flat 2D world are given below:

$$\frac{dx(t)}{dt} = V(t)\cos\psi(t) \qquad x(t) = x_0 + \int_0^t V(t)\cos(\psi(t))dt$$

$$\frac{dy(t)}{dt} = V(t)\sin\psi(t) \qquad y(t) = y_0 + \int_0^t V(t)\sin(\psi(t))dt$$

$$\frac{d\psi(t)}{dt} = V(t)\kappa(t) \qquad \psi(t) = \psi_0 + \int_0^t V(t)\kappa(t)dt$$

6.2.3 Nonholonomic Constraint

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 length and a is a constant. These curves are linear equations for curvature in the arc length parameter and are known as the **clothoids**. 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.

6.3 Rough Ackerman Steering Dynamics

The following sections consider the latencies associated with a typical Ackerman steering column. When such 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.

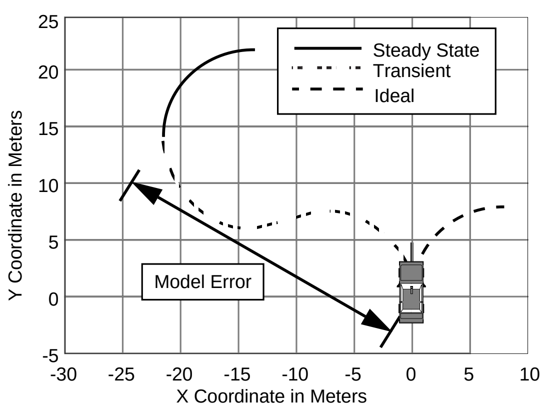

Figure 23 Transience in the Reverse Turn

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.

6.3.1 Heading Response

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

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

where ψ is vehicle heading, κ is curvature, and T_{act} is the time required for the actuator to reach commanded deflection. 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:

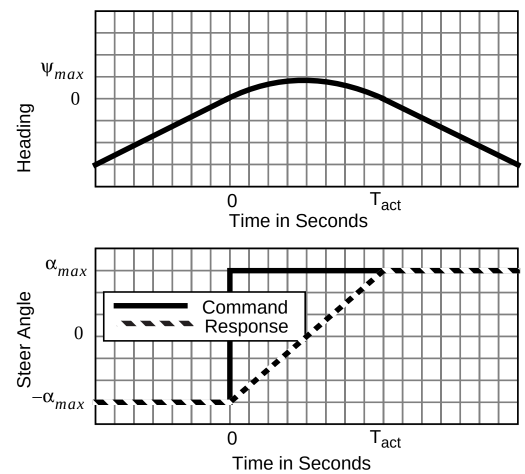

Figure 24 Transient Steering Response

6.3.2 Nondimensional Transient Turning

If $\Delta\alpha$ is the commanded change in steer angle, and α_{max} is the maximum rate of change of steer angle, the actuator reaction time for a reverse turn is given by:

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

The temporal horizon of obstacle avoidance is the time required to turn through an angle $\Delta\psi$ at constant curvature

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

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

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

This nondimensional is a particular instance of the **normalized time constant**. It provides 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. It can easily exceed unity for a conventional automobile.

6.3.3 Command Following Problem

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:

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

Thus, the modeling error associated with an ideal model of steering is equal to the reaction distance of the steering actuator.

To cast this result in terms of the fidelity ratio, consider the minimum fidelity ratio for an acceptable model error on the order of the wheel radius. Let this be called the **turning fidelity ratio**:

$$\rho_t = \frac{dx}{(r-W)} = \frac{T_{act}V}{(r-W)}$$

This number must be significantly less than unity to allow ignoring dynamics. It is often on the order of 10.

6.4 Exact Ackerman Steering Dynamics

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 modeled 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.

6.4.1 Dependence of Steering Response on Speed

The limited rate of change of curvature for an Ackerman steer vehicle is an important modeling 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. It should be clear from this analysis that obstacle avoidance must account somehow 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.

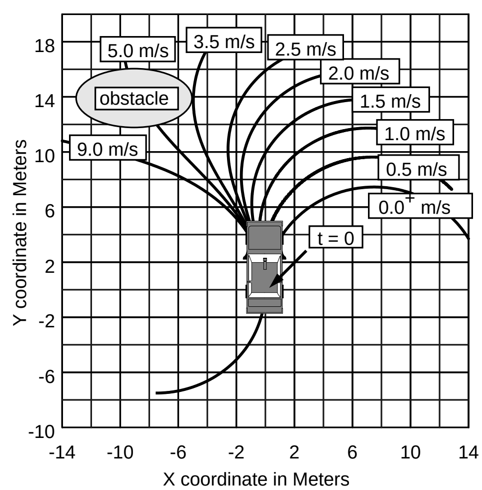

Figure 25 Constant Curvature Reverse Turn

6.4.2 Stability Problem

Feedforward of dynamics can be necessary for stable control. In the above figure, 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 issued although a slight turn was commanded. Acceptable control is not possible without knowledge of these dynamics.

6.4.3 Exact Response of Steering at Constant Speed

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 issued at a speed of 5 m/s. 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*. 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.

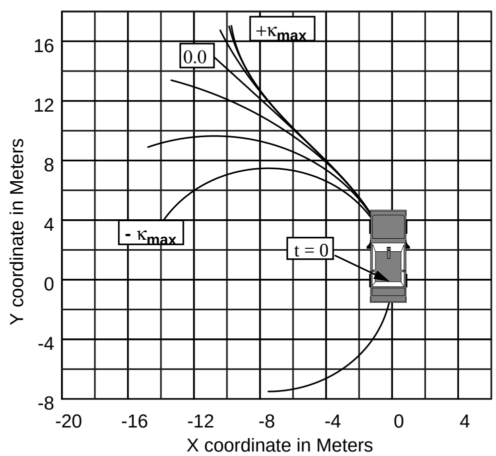

Figure 26 Constant Speed Reverse Turn

One valid model of this system is a coupled system of nonlinear differential equations.

7 Interactions

The satisfaction of the guaranteed safety policy requires that response, throughput, acuity, and fidelity requirements all be met simultaneously. They cannot, however, be treated individually because they are all interrelated. For example:

- Throughput depends on resolution if the same area of ground is to be covered per unit time.
- Resolution depends on range through the projective geometry of the sensor.
- Sensor range depends on speed if lookahead is modulated to match or exceed response distance.

Resolving the above relations leads to the conclusion that throughput depends at least on speed. This section will quantify this relationship between throughput and speed based on the preceding analysis. Indeed, we will show that the complexity of terrain mapping perception is polynomial in the reaction distance.

Guaranteed safety implies that throughput is proportional to a high power of velocity because:

- Maximum range increases quadratically with speed because response distance does.
- Pixel size decreases quadratically with maximum range for constant groundplane resolution..
- Throughput increases quadratically with pixel size assuming a fixed sensor field of view.

Naive analysis will suggest that the problem of high speed navigation is difficult because the necessary throughput approaches impractical levels. On the other hand, if one computes the rate at which a vehicle covers ground area as it moves, any reasonable spatial resolution for sampling this area leads to a perceptual throughput that is trivial to meet. This difference has many sources which will be presented in this section.

7.1 Assumptions of the Analysis

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

We will be interested in terrain mapping perception algorithms [reference] which associate a single unique elevation with each point in a sampled representation of the groundplane. The most important assumption of the analysis is the **stationary environment assumption** because this permits us to percieve a point on the groundplane only once and avoid dealing with moving obstacles. While points in the environment will certainly move relative to the vehicle, they will be assumed not be moving relative to each other.

Based on earlier analysis of field of view, the HFOV will be fixed because it must be chosen based on the worst case speed which is often below the maximum vehicle speed. We will employ the following assumptions:

- Horizontal field of view is fixed at 80°, 120°, 170°, and 215° for each increasing reaction time respectively, based on earlier analysis.
- Sensor frame rate is set to 2 Hz because this is a typical value for a laser rangefinder with a wide VFOV.
- Minimum acuity will be used because this is actually the most stringent requirement beyond

some range.

The estimates that are produced are underestimates for many reasons, including the following additional assumptions, all of which lower the required throughput:

- The graphs estimate perception geometric transform processing only. Planning, position estimation, and control are not included at all.
- The processor load is assumed to be 50 flops per pixel when experience suggests that many times this is required in a practical system.
- 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 - being based on a turning maneuver.
- The maximum range that is chosen is based on the response distance. Actually it is the minimum range which should be set to the response distance.
- The point vehicle assumption is used to avoid dealing with the offset of the vehicle nose from the sensor.
- We use an obstacle sampling factor of 1. A practical factor is perhaps between 3 and 10. This implies that the results must be multiplied by the square of a practical sampling factor.

7.2 Common Throughput Expression

It will be necessary to quantify the number of operations performed per unit time in terms of the number of sensor pixels processed times the number of operations used per pixel. This section develops a basic expression which is then modified based on further assumptions.

7.2.1 Sweep Rate

The product of the vertical field of view and the frame rate is measures the effective angular velocity of the sensor in the vertical direction, and is known as the **sweep rate**:

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

where $VFOV$ is the vertical field of view and f_{images} is the frame rate.

7.2.2 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:

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

where $HFOV$ is the horizontal field of view. Note that the sensor flux is the two dimensional analog of the sweep rate. Not surprizingly, the two are related by:

$$\Psi = HFOV \times \dot{\theta}$$

and the sensor flux has units of angular flux - solid angle per unit time.

7.2.3 Sensor Throughput

The number of range pixels generated per unit time by a sensor will be called the **sensor throughput** f_{pixels}. If the field of view is fixed and pixels are square, the sensor throughput is given by:

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

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**.

7.2.4 Processor Load

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

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

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

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

7.2.5 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:

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

$$f_{cpu} = \sigma_P \frac{HFOV \times IFOV \times f_{images}}{(IFOV)^2}$$

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

$$f_{cpu} = \sigma_P \frac{HFOV \times VFOV \times f_{images}}{IFOV_H IFOV_v}$$

7.3 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 the number of pixels generated per unit time times the cost of processing one pixel.

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

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

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:

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

We will invoke the point robot assumption and eliminate p and the small incidence angle assumption to equate Y to R :

$$R = T_{react}V[1+\bar{b}]$$

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 assumption, the stopping distance is the product of speed and reaction time - the reaction distance.

The resulting complexity estimate represents the minimum computational throughput necessary in order to meet guaranteed response, throughput, and detection simultaneously. Any system which cannot supply this throughput must either:

- reduce resolution and violate guaranteed detection.
- reduce field of view and violate guaranteed throughput.
- reduce lookahead and violate guaranteed response.

7.4 Constant Flux

A real sensor usually has a fixed field of view and fixed frame rate, so the sensor flux Ψ is constant. It is straightforward to compute the throughput required to keep up with the sensor. Throughput under guaranteed detection is obtained by substituting the acuity expression into the basic throughput expression:

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

Substituting the stopping distance for range gives:

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

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

$$f_{cpu} \sim \sigma_P O([T_{react}V]^4)$$

The following graph indicates the variation of throughput with speed when square pixel size is chosen to satisfy the minimum acuity resolution requirement at the maximum range. The processing rates required are substantial - even under our liberal assumptions.

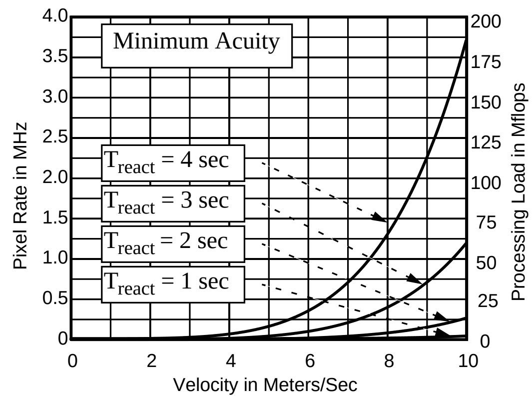

Figure 27 Throughput for Constant Flux

7.5 Adaptive Sweep

Adaptive sweep will be defined as the process of modulating the sweep rate of the sensor to generate an imaging density of unity and thereby barely satisfy guaranteed throughput. This does not compromise guaranteed response and it leads to significant reduction in throughput.

The basic throughput expression under guaranteed detection is:

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

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

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

The complexity expression is now:

$$f_{cpu} = \sigma_P\left(\frac{4R^4}{(Lh)^2}\right)VFOV \times HFOV \times f_{images}$$

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

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

Guaranteed throughput is implemented by setting $\rho_{throughput} = 1$. Also, assuming every image is processed:

$$T_{cyc} = 1/f_{images}$$

Which give the sweep rate as:

$$VFOV \times f_{images} = h\frac{V}{Y^2}$$

This gives:

$$f_{cpu} = \sigma_P\left(\frac{4R^4}{(Lh)^2}\right)h\frac{V}{Y^2}HFOV$$

Cancelling an R^2 and a Y^2 and substituting the stopping distance for R gives:

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

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

$$f_{cpu} \sim \sigma_P O([T_{react}V]^2[V])$$

which is less than the previous result by the factor $T^2_{react}V$.

This result leads to the conclusion that adjusting the vertical field of view based on vehicle speed can significantly reduce computational throughput requirements.

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.

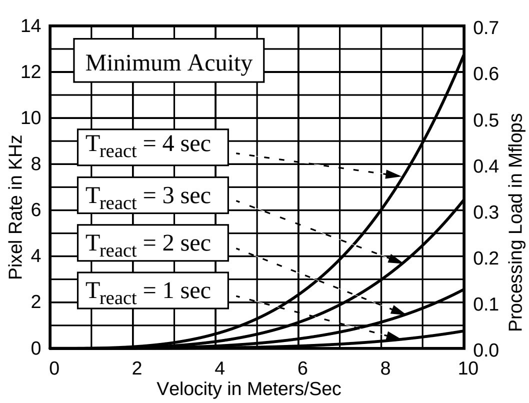

Figure 28 Throughput for Adaptive Sweep

7.6 Adaptive Sweep/Scan

Adaptive scan will be defined as the process of modulating the aspect ratio of image pixels in order to achieve roughly the same crossrange and downrange resolution on the groundplane. This process does not compromise guaranteed detection and it leads to further reduction in throughput.

Recall that the cpu load required to process all sensory data is given by:

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

In practical adaptive scan, the pixel aspect ratio is adjusted to be a constant over the field of view and equal to the perception ratio. The vertical and horizontal instantaneous field of view then have different expressions at minimum acuity:

$$IFOV_V = \frac{1}{2}\frac{Lh}{R^2}$$

$$IFOV_H = \frac{1}{2}\frac{Lh}{R^2}\left(\frac{R}{h}\right) = \frac{1}{2}\frac{L}{R}$$

where the horizontal image resolution was multiplied by the factor R/h in order to implement adaptive scan.

The throughput required to process all sensory data is then given by:

$$f_{cpu} = \sigma_P\left(\frac{4R^3}{L^2h}\right)VFOV \times HFOV \times f_{images}$$

As before, the sweep rate under guaranteed throughput is:

$$VFOV \times f_{images} = h\frac{V}{Y^2}$$

This gives:

$$f_{cpu} = \sigma_P\left(\frac{4R^3}{L^2h}\right)h\frac{V}{Y^2}HFOV$$

Cancelling an R^2 and a Y^2 and substituting the stopping distance for R gives:

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

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

$$f_{cpu} \sim \sigma_P O(T_{react}V^2)$$

which is less than the previous result by the factor $T_{react}V$. 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.

These two simple adaptive techniques have reduced

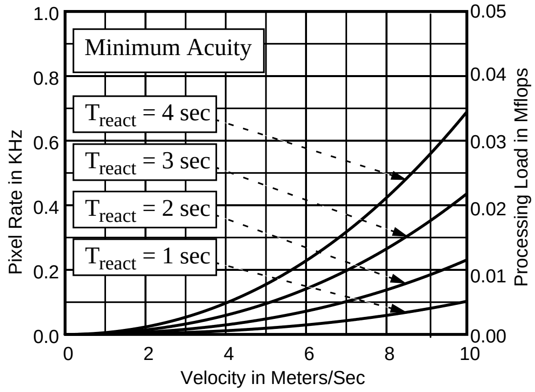

Figure 29 Throughput for Adaptive Scan

the throughput requirements by 4 orders of magnitude over constant flux at speeds of 20 mph and reaction times of 4 seconds.

7.7 Adaptive Sweep, Uniform Scan

Uniform scan will be defined as the process of modulating pixel size and field of view in order to achieve perfect homogeneous and isotropic distribution of resolution on the groundplane. No existing sensor can provide this capability, but it is a useful theoretical approximation.

This analysis considers the fundamental acuity and throughput requirements of perception. As a minimum requirement, any sensor must generate geometry at a rate that is consistent with the rate at which the vehicle consumes geometry through its motion. Consider that the motion of the vehicle consumes a swath of geometry directly in front of it as shown below:

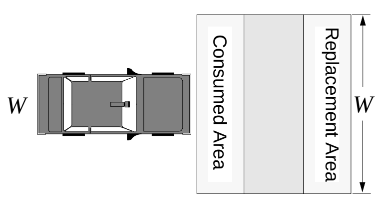

Figure 30 Area Consumption

In the simplest case, this consumed area must be replaced by adding new information to the map shown to the right. The area consumed per second, expressed in appropriate units, is the required absolute minimum throughput of a perception system under guaranteed throughput.

Let the width of the swath be W, the velocity of the vehicle be V, and the required spatial resolution be δ. This minimum rate is given by:

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

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 swath to twice the maximum range gives:

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

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

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

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

$$f_{cpu} \sim \sigma_P O(T_{react}V^2)$$

which is, in complexity terms, equal to the adaptive sweep, adaptive scan expression. There is a multiplicative constant difference of $2 \times HFOV$ 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 and the HFOV is fixed. This relationship is plotted below for minimum acuity spatial resolution of 3.3 meters versus vehicle velocity for various values of system reaction time.

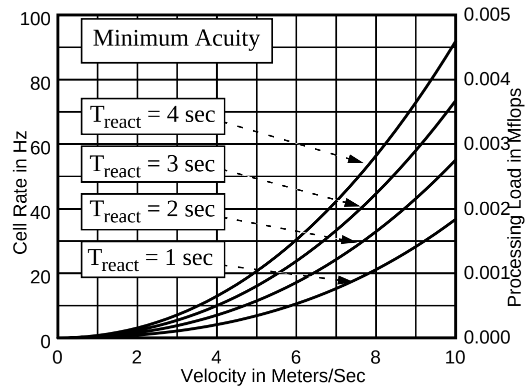

Figure 31 Throughput for Uniform Scan

7.8 Throughput for all Algorithms

Recall that the preceeding complexity estimates are all consistently based on a kinematic braking regime assumption. The true power of velocity is actually squared as speeds increase. Identical assumptions of minimum acuity, 4 second reaction time, and 10 meter / second speed, have led to the following throughput estimates for different image processing algorithms:

Table 1: Throughput Estimates

Algorithm	Estimate	Complexity
constant flux	250 Mflops	$O(T_{react}^4 V^4)$

Table 1: Throughput Estimates

Algorithm	Estimate	Complexity
adaptive sweep	0.7 Mflops	$O(T_{react}^2 V^3)$
adaptive scan	0.035 Mflops	$O(T_{react} V^2)$
ideal	0.0045 Mflops	$O(T_{react} V^2)$

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

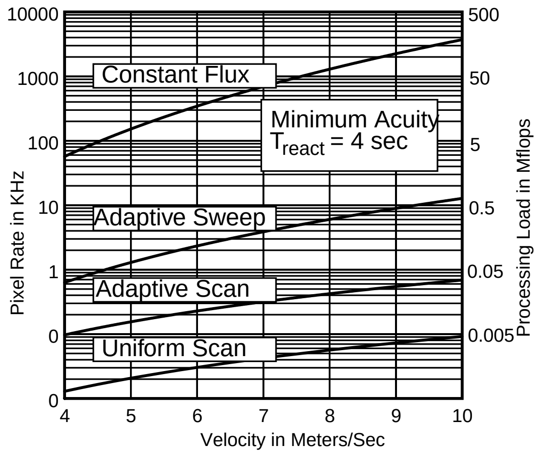

Figure 32 Throughput for All Algorithms

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

$$f_{cpu} \sim \sigma_P O([T_{react}V]^N [V]^M)$$

There are a few ways to read this result. If throughput is fixed, then speed is inversely proportional to reaction time. If speed is fixed, throughput required grows with the nth power of reaction time. If reaction time is fixed, throughput grows with the (n+m)th power of speed. In general, the complexity of terrain mapping perception is polynomial in the vehicle reaction distance.

8 Conclusions

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

8.1 Sensor Geometry

One very important distinction of high-speed autonomous mobility is the fact that sensor height is typically an order of magnitude smaller than the vehicle response distance. This observation has many implications relating to the prevalence of occlusions in images and the complexity of image processing algorithms.

8.2 Obstacle Avoidance

From the perspective of reliability in obstacle detection and avoidance, it is important to recognise that the planning horizon of obstacle avoidance (reaction time) is roughly equal to the characteristic time (time constant) of the actuators, so the system operates almost entirely in the transient regime. This leads to the conclusion that the absence of dynamic models of response will lead to unreliability in obstacle avoidance. Specifically, “arc” based models of Ackerman steering will be unreliable at even moderate speeds.

8.3 Goal Seeking

In the particular case of steering delays, the raw trajectory error associated with higher speeds implies that stability problems will emerge with control algorithms that do not account for the delay.

8.4 Trajectory Generation

From a trajectory generation and planning perspective, it seems advisable not to attempt the C-space to A-space transform in any form such as the generation of clothoids if another method can be found. Feedforward, for example, is one alternative that generates the C space curve from the A space curve with little algorithmic difficulty at the level of trajectory generation.

8.5 Adaptive Sweep Perception

As speeds increase, the redundant measurement of the same geometry that happens when images overlap on the groundplane becomes more of an efficiency concern because the imaging density increases without bound. We have proposed a technique called adaptive sweep which deliberately modulates the vertical field of view and shown two orders of magnitude reduction in required perceptual throughput.

8.6 Adaptive Scan Perception

Another technique which improves computational requirements is the modulation of the range pixel aspect ratio in order to precisely meet the groundplane resolution requirement. We have shown an order of magnitude reduction in required perceptual throughput under certain operating conditions.

8.7 Complexity of Terrain Mapping

In general, the complexity of terrain mapping perception is polynomial in the vehicle reaction distance.

8.8 Fundamental Tradeoff

This complexity result quantifies the perceptual throughput problem of autonomous mobility and identifies the **fundamental trade-off** associated with the use of finite computing resources. This trade-off is one of resolution for speed, or equivalently, reliability for speed. Computing resources establish a limit on vehicle performance which can be expressed as either high speed and low reliability or vice versa.

9 References

1. M. G. Bekker, “Off-the-Road Locomotion”, The University of Michigan Press, 1960.
2. M. G. Bekker, “The Theory of Land Locomotion”, The University of Michigan Press, 1956.
3. R. Bhatt, L. Venetsky, D. Gaw, D. Lowing, A. Meystel, “A Real-Time Pilot for an Autonomous Robot”, Proceedings of IEEE Conference on Intelligent Control, 1987.
4. B. Blochl, “System Requirements and architectures for vision-guided autonomous robots”, in Proceedings of the 12th European meeting on Cybernetics and Systems Research, Vienna, 5-8 April, 1994.
5. B. Brumitt, R. C. Coulter, A. Stent, “Dynamic Trajectory Planning for a Cross-Country Navigator”, Proceedings of the SPIE Conference on Mobile Robots, 1992.
6. 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.
7. 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.
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. E. D. Dickmanns, “Dynamic Computer Vision for Mobile Robot Control”, Proceedings of the 19th International Symposium and Exposition on Robots, pp. 314-27.
10. J. C. Dixon, “Linear and Non-Linear Steady State Vehicle Handling”, Proceedings of Instn. Mechanical Engineers, Vol. 202, No. D3, 1988.
11. R. T. Dunlay, D. G. Morgenthaler, “Obstacle Avoidance on Roadways Using Range Data”, Proceedings of SPIE Conference on Mobile Robots, 1986.
12. D. Feng, S. Singh, B. Krogh, “Implementation of Dynamic Obstacle Avoidance on the CMU Navlab”, Proceedings of IEEE Conference on Systems Engineering, August, 1990.
13. U. Franke, H. Fritz, S. Mehring, “Long Distance Driving with the Daimler-Benz Autonomous Vehicle VITA”, PROMETHEUS Workshop, Grenoble, December, 1991.
14. J. Gowdy, A. Stentz, and M. Hebert, “Hierarchical Terrain Representation for Off-Road Navigation”, In Proc SPIE Mobile Robots, 1990.
15. P. Grandjean and L. Matthies, “Perception Control for obstacle detection by a cross-country rover”, in Proceedings of the !EEE International Conference on Robotics and Automation, May 1993.
16. M. Hebert and E. Krotkov, “Imaging Laser Radars: How Good Are They”, IROS 91, November 91.
17. 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.
18. 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.
19. B. Hotz, Z. Zhang, P. Fua, “Incremental Construction of Local DEM for an Autonomous Planetary Rover”, Workshop on Computer Vision for Space Applications - Antibes, Sept. 22-24, 1993.
20. D. Keirsey, D. Payton, J. Rosenblatt, “Autonomous Navigation in Cross-Country Terrain”, proceedings of Image Understanding Workshop, 1988.
21. A. Kelly, A. Stentz, M. Hebert, “Terrain Map Building for Fast Navigation on Rough Terrain”, Proceedings of the SPIE Conference on Mobile Robots, 1992.
22. E. Krotkov, and M. Hebert, “Mapping and Positioning for a Prototype Lunar Rover”, proceedings of the International Conference on Robotics and Automation, May 1995.
23. In So Kweon, “Modelling Rugged Terrain by Mobile Robots with Multiple Sensors”, CMU Ph.D. Thesis, 1990.
24. L. Matthies, “Stereo Vision for Planetary Rovers”, International Journal of Computer Vision, 8:1, 71-91, 1992.
25. L. Matthies, and P. Grandjean, “Stochasitc performance modelling and evaluation of obstacle detectability with imaging range sensors.” IEEE Transactions on Robotics and Automation, Vol 10, Number 6, December 1994, pp 783-791
26. L. Matthies, A. Kelly, T. Litwin, “Obstacle Detection for Unmanned Ground Vehicles: A Progress Report,” to appear in the Proceedings of the IEEE International Symposium on Intelligent Vehicles.
27. K. Olin, and D. Tseng, “Autonomous Cross-Country Navigation”, IEEE Expert, August 1991, pp. 16-30.
28. 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.

29. A. M. Shkel and V. J. Lumelsky, “The joggers problem: accounting for body dynamics in real-time motion planning”, in Proceedings of the 1995 International Conference on Intelligent Robots and Systems, Pittsburgh, PA, 5-9 Aug, 1995.
30. S. Singh, et. al, “FastNav: A System for Fast Navigation”, Robotics Institute Technical Report CMU-RI-TR-91-20, Carnegie Mellon University, 1991.
31. A. Stentz, “Optimal and Efficient Path Planning for Unknown and Dynamic Environments”, CMU Tech report CMU-RI-TR-93-20.
32. J. Spofford and B. Gothard, “Stopping distance analysis of ladar and stereo for unmanned ground vehicles.”, Proceedings of the International Society of Optical Engineering, Mobile Robots IX, Boston Mass, 2-4 Nov, 1994.

Alonzo Kelly is a project scientist at the Robotics Institute of Carnegie Mellon University. He recieved his B. A. Sc. in Aerospace engineering from University of Toronto in 1984, his B. Sc. in computer science from York University in 1990, and his Masters and PhD degrees in robotics from Carnegie Mellon University in 1994 and 1996 respectively. His research interests include perception, planning, control, simulation, and operator interfaces for indoor and outdoor mobile robots.

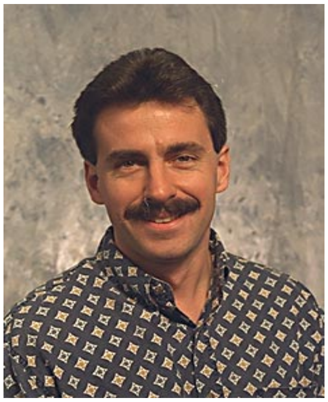

Anthony Stentz is a senior research scientist at the Robotics Institute of Carnegie Mellon University. He received his Ph.D. in computer science from Carnegie-Mellon University in 1989. He received his M.S. in computer science from CMU in 1984 and his B.S. in physics from Xavier University of Ohio in 1982. His research interests include mobile robots, path planning, computer vision, system architecture, and artificial intelligence in the context of fieldworthy robotic systems.

10 Appendix A - List of Symbols

10.1 Lowercase Alphabetics

$\bar{b}$ ....braking coefficient

c ....undercarriage clearance

f ....frequency

f_{cells} ....map cell throughput

f_{cpu} ....computational bandwidth

f_{images} ....frame rate

f_{pixels} ....sensor throughput

g ....acceleration due to gravity

h ....sensor height

$n/2$ ....sampling factor

p ....sensor / vehicle nose offset

r ....wheel radius

s ....arc length, distance travelled

s_{brake} ....braking distance

s_{turn} ....turning distance

s_{react} ....reaction distance

$s_{response}$ ....response distance

s_{error} ....error distance

$\bar{s}$ ....maneuver coefficient

t ....time

$\bar{t}$ ....turning coefficient

$\bar{t}_t$ ....transient turning coefficient

x ....crossrange coordinate

x_0 ....initial crossrange coordinate

y ....downrange coordinate

y_0 ....initial downrange coordinate

z ....vertical coordinate

z_0 ....initial vertical coordinate

10.2 Alphabetics

L ....vehicle wheelbase

R ....range

R_{max} ....maximum range
R_{min} ....minimum range
R_L ....lookahead distance
T ....time, time interval
T_{act} ....actuator delay
T_{brake} ....braking reaction time
T_{cont} ....control reaction time
T_{cyc} ....software cycle time, frame period
$T_{maneuver}$ .....maneuver time
T_{hw} ....hardware reaction time
T_{look} ....temporal planning horizon
T_{perc} ....perception reaction time
T_{plan} ....planning reaction time
T_{react} ....system reaction time
$T_{response}$ ....system response time
T_{sens} ....sensing reaction time
T_{sw} ....processing reaction time
T_{turn} ....turning reaction time
T_{veh} ....vehicle reaction time
V ....vehicle speed
W ....vehicle width, swath width
Y ....groundplane projected range
Y_{min} ....min groundplane projected range
Y_{max} ....max groundplane projected range
Y_L ....lookahead groundplane projected range

10.3 Greek Alphabetics

α ....steer angle
α_{max} ....maximum steer angle
$\dot{\alpha}$ ....steer angle rate
$\dot{\alpha}_{max}$ ....maximum steer angle rate
β ....yaw
$\dot{\beta}$ ....angular velocity (z component)
δ ....spatial resolution
Δ ....obstacle spacing
κ ....curvature
κ_0 ....initial curvature
μ ....coefficient of friction
ν ....coefficient of lateral acceleration
ψ ....yaw, pixel azimuth, vehicle yaw
ψ_0 ....initial yaw
$\dot{\psi}$ ....vehicle yaw rate
ψ_L ....planner lookahead angle
$\psi_{response}$ ......response angle
ρ ....radius of curvature
ρ_{acuity} ....acuity ratio
ρ_{cyc} ....throughput ratio
$\rho_{fidelity}$ ....fidelity ratio
ρ_{dyn} ....minimum dynamic turn radius
ρ_{kin} ....minimum kinematic turn radius
ρ_{min} ....minimum turn radius
ρ_t ....turning fidelity ratio
$\rho_{response}$ ....response ratio
σ ....area density, generic nondimensional
σ_I ....imaging density
σ_p ....processor load
$\sigma_{response}$ ....normalized reponse distance
τ ....time constant
$\bar{\tau}$ ....normalized time constant
θ ....pitch, elevation
θ_{max} ....maximum depression
θ_{max} ....maximum allowable body pitch
θ_{min} ....minimum depression
$\dot{\theta}$ ....vertical sweep rate
ϕ ....roll
Ψ ....sensor flux

10.4 Increments and Differentials

dx, Δx ....crossrange incremental distance
dy, Δy ....downrange incremental distance
dz, Δz ....vertical incremental distance
ds ....general incremental distance
Δh ....height increment, resolution
Δt ....time delay
ΔR, ΔR_L ....incremental lookahead distance
ΔY, ΔY_L ....groundplane projected incr. lookahead dist.
$\Delta\alpha$ ....steer angle increment
$d\theta$, $\Delta\theta$ ....pitch/elevation increment or error
$d\psi$, $\Delta\psi$ ....yaw/azimuth increment or error
$\Delta\psi_L$ ....incremental lookahead angle

11 Appendix B - Glossary

Ackerman steering - a steering mechanism, typical of automobiles, where the two front wheels turn together.
actuation space - an abstract space consisting of all independent control inputs to a system. For an automobile, this

space is spanned by steering, brake, and throttle.
acuity - a synonym for resolution. Often used in the context of visual angular resolution.
acuity problem - the problem of maintaining adequate resolution in computations and sensing.
acuity ratios - nondimensional ratios of world model resolution and an appropriate dimension of the vehicle.
adaptive scan - an algorithm which attempts to make perceptual groundplane resolution homogeneous and isotropic.
adaptive sweep - an algorithm which projects a focus of attention on the groundplane into image space.
azimuth - the rotation of something about a vertical axis.
bicycle model - a model of a vehicle in terms of an analogous two-wheeled vehicle (bicycle).
braking angle - the angle travelled while the brakes are actuated before coming to a stop.
braking coefficient - the ratio of braking distance to reaction distance.
braking distance - the distance travelled while the brake is actuated before coming to a stop.
braking reaction time - the time it takes from image acquisition until the brakes are engaged.
characteristic time - a measure of the delay associated with a system from input to output.
clothoid - a linear polynomial for curvature expressed in terms of arc length. The trajectory corresponding to this polynomial.
coefficient of lateral acceleration - one-half the ratio of lateral acceleration and gravitational acceleration. Equal to one-half the lateral acceleration in units of g's.
command following problem - the problem of causing a servo-controlled device to follow its commands acceptably well.
communication latency - the time it takes to pass information between system processes and processors.
computational bandwidth - number of floating point operations (flops) per second.
configuration space - any abstract space of variables which completely determines the positions of all points on a vehicle or mechanism.
constant flux - a property of a sensor of scanning a fixed solid angle at a fixed rate.
constant scan - a property of a sensor of scanning at fixed angular resolution.
crossrange - the horizontal direction transverse to the sensor optical axis.
curvature - the derivative of heading (or vehicle yaw) with respect to distance travelled.
dead reckoning - the process of integrating certain equations which express position and heading in terms of curvature and either distance or time.
depth of field - here, the maximum and minum sensor range.
downrange - the horizontal direction aligned with the sensor optical axis.
dynamic braking regime - the speed regime for which the braking coefficient exceeds unity.
dynamic response regime - the speed regime for which the response coefficient exceeds unity.
elevation - the rotation of something about a horizontal axis
fidelity - a synonym for accuracy. Often used in the context of sound reproduction or dynamic simulation.
fidelity problem - the problem of maintaining adequate fidelity in computations and measurements.
fidelity ratio - the ratio of actual or predicted system error to some allowable threshold of error.
field of view - the region of space that a sensor can see expressed in terms of an angle.
field of regard - the entire region of space that a sensor can see in terms of both depth and angle.
flat terrain assumption - the assumption that the vehicle operates in extremely benign terrain.
Fresnel Integrals - the equations of dead reckoning which express position and heading in terms of curvature and either distance or time.
fundamental trade-off - the trade-off of speed for resolution that occurs in any system with a throughput limit.
guaranteed detection - the policy of ensuring adequate resolution in computations, sensing, and actuation.
guaranteed localization - the policy of ensuring adequate accuracy in computations, sensing, and actuation.
guaranteed response - the policy of ensuring adequate response time.
guaranteed safety - the policy of guaranteeing vehicle survival.
guaranteed throughput - the policy of ensuring adequate system throughput.
hardware reaction time - the total reaction time of all hardware components.
hill occlusion rule - a relationship between vehicle configuration parameters that must be satisfied to see behind a small hill.
hole occlusion rule - a relationship between vehicle configuration parameters that must be satisfied to see inside a small hole.
imaging density - a measure of the number of images which contain the same point in the environment.
impulse turn - a turn from zero curvature to the maximum.
impulse turning distance - the distance measured along the original straight trajectory consumed in an impulse turn.
instantaneous field of view - the solid angle subtended by a single range measurement.
kinematic braking regime - the speed regime for which the braking coefficient does not exceed unity.
kinematic response regime - the speed regime for which the response coefficient does not exceed unity.
latency - any type of delay in transforming inputs into outputs in a system.
latency problem - the problem of latencies that are too large for the current speed or sensory lookahead.
low latency assumption - the assumption that latencies are small enough to be neglected for any particular speed and sensory lookahead.
terrain mapping - the process of generating a map of the environment surrounding the vehicle.

maneuver - the maneuver dynamics aspect of response..
maneuver time - the time it takes to completely execute a commanded maneuver.
maneuver coefficient - the ratio of maneuver distance to reaction distance.
maneuver distance - the distance consumed in the execution of a maneuver.
maximum acuity - the highest resolution that should be needed.
maximum sensor acuity rule - a rule for determining sensor angular resolution.
minimum acuity - the minimum resolution that should be needed.
minimum sensor acuity rule - the condition that the minimum acuity ratio not exceed one-half.
minimum dynamic turn radius - the minimum turn radius imposed by a given speed and lateral acceleration limit.
minimum kinematic turn radius - the minimum turn radius, if any, imposed by the steering mechanism.
minimum significant delay - the smallest unmodelled delay that affects the fidelity of a model.
motion distortion problem - the distortion of the world model due to unmodelled delays or unmodelled motion.
myopia problem - the problem of inadequate sensory lookahead.
negative obstacle - holes and depressions in the terrain.
nonholonomic - the property of a differential constraint that it cannot be integrated.
normalized response distance - the response distance normalized by a characteristic vehicle dimension. A measure of the capacity of a vehicle to respond quickly.
normalized time constant - the ratio of a characteristic time to a temporal planning horizon.
normalized wheelbase - the ratio of wheelbase to range.
occlusion problem - the problem of missing parts in a world model because radiation reflects from the first reflecting surface only.
panic stop - the obstacle avoidance maneuver of slamming on the brakes.
perception ratio - the ratio of sensor height to measured range. Equal to the tangent of the range pixel incidence angle for flat terrain.
plant dynamics latency - the delay that arises in physical systems because they are governed by differential equations.
point vehicle assumption - the assumption that the spatial extent of the vehicle can be ignored so that it can be modelled as a point.
positive obstacle - a hill or rise in the terrain.
processing latency - the time it takes for an algorithm to transform its inputs into its outputs
processor load - the number of floating point operations used to process an average pixel.
range image - an image whose intensity values correspond to the range to the first reflecting surface in the environment.
reaction - the computer and sensory processing aspects of response.
reaction angle - the angle through which the vehicle turns while deciding on a course of action.
reaction distance - the product of a vehicle speed and some appropriate reaction time.
reaction time - the time it takes to decide on a course of action and issue the associated commands to the hardware.
registration problem - the problem that redundant measurements of the same geometry do not agree.
resolution - the smallest difference that a system can resolve.
response - the total response of the vehicle including the computer and sensory processing and the maneuver dynamics.
response angle - the angle through which a vehicle turns while responding.
response distance - the sum of reaction distance and maneuver distance.
response problem - the problem of maintaining adequate response time.
response ratio - the ratio of reaction distance to sensory lookahead.
response time - the time it takes a vehicle to respond to an external event. Equal to the sum of reaction time and maneuver time.
response velocity - the ratio of response distance to response time.
reverse turn - a turn from one curvature extreme to the other.
sampling problem - the problem of variation in the shape and size of range pixels when projected onto the ground plane.
sampling factor - the ratio of actual resolution to the minimum required.
sensitivity problem - a high degree of sensitivity of one parameter with respect to another.
sensor dwell latency - the time it takes for a measurement to be acquired.
sensor flux - the solid angle scanned per unit time by a sensor.
sensor throughput - the sensor measurement rate in range pixels per second.
small incidence angle assumption - the assumption that the perception ratio is small.
smooth terrain assumption - the assumption that the terrain is not rough.
stability problem - the problem of maintaining stability.
stabilization problem - the problem of stabilization of sensors.
stationary environment assumption - the assumption that the environment is a single rigid body that is stationary, that is, that there are no dynamic obstacles or changes in the terrain etc.
steady-state regime - the regime of operations where the system output is not changing in a material way.
sweep rate - the angular elevation rate of scanning of a range sensor.
sweep rate rule - a design rule for determining sweep rate for a given vehicle speed.
system response time - the total reaction time of the vehicle from image acquisition to executed response.

temporal planning horizon - the amount of time a system or subsystem looks ahead.

throughput - a measure of amount of information processed per unit time.

throughput problem - the problem of maintaining adequate throughput.

throughput ratio - the ratio of the area covered by the vehicle in one cycle to the area measured by the sensor in the same time.

time constant - the coefficient of the first derivative in a first order system.

transient regime - the regime of operations where the system output is changing in a material way.

transient turning coefficient - the normalized time constant associated with a turning maneuver.

tunnel vision problem - the problem of inadequate horizontal field of view.

turning angle - the angle through which the vehicle turns in in a turning maneuver.

turning coefficient - the ratio of turning distance to reaction distance.

turning distance - the distance travelled in a turning maneuver.

turning fidelity ratio - the fidelity ratio associated command following error in turning maneuvers.

turning reaction time - the time it takes to execute a turn maneuver.

turning stop - a panic stop maneuver issued while in a turn.

turning stop maneuver - see turning stop.

undercarriage tangent - the tangent of the angle from the center of the undercarriage to the bottom of a wheel.

uniform scan - a theoretical ability of a sensor to produce a scanning pattern that does not suffer from the sampling problem.

undersampling - the process of sampling below the Nyquist rate.

vertical - the direction aligned with local gravity.

wheelbase - the length of the vehicle measured from back to front wheels.

yaw - rotation about the vertical axis.