Submitted to the 1997 IEEE
International Conference on Robotics and Automation
Albuquerque, New Mexico - April 1997

Analysis of Requirements for High Speed Rough Terrain Autonomous Mobility Part II: Resolution and Accuracy

Alonzo Kelly and Anthony Stentz

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

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

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 the purposes of this paper, the perception sensor can be any imaging sensor measuring range or intensity in any electromagnetic band of frequencies. This paper proposes aspects of a rudimentary theory of obstacle avoidance and uses it to quantify some of the requirements placed upon autonomous systems that are derived from the need to ensure safety

2 Guaranteed Safety

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

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

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

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

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

2.1 Preliminaries

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

2.1.1 Lexical Conventions

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

2.1.2 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.

2.1.3 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.

2.2 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.

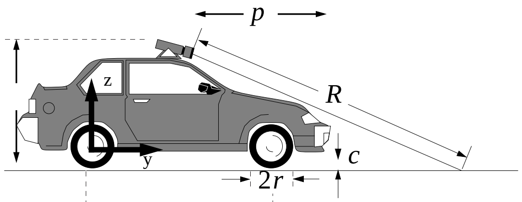

Figure 1 Important Dimensions

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 Nondimensional Safety Requirements

One way to characterize scale is to choose a characteristic vehicle dimension to represent its size. We will sometimes use the wheelbase L, the width W, or the wheel radius r to characterize scale. In this way, results will be expressable in scale-independent terms.

2.4.1 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.

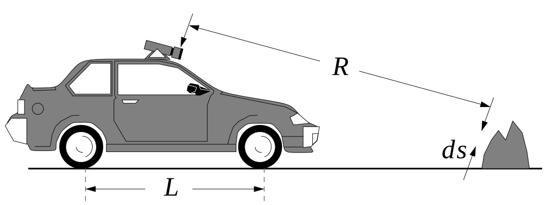

Figure 2 Acuity

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.

2.4.2 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.

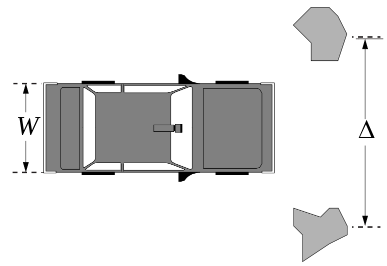

Figure 3 Fidelity

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.

2.5 Standard Assumptions

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

- **small incidence angle assumption**: the assumption that the perception ratio is small. When adopted, allows us to equate the range to a point on the ground to its groundplane projection with a minimal relative error equal to the square of the perception ratio.
- **point vehicle assumption**: the assumption that the finite extent of the vehicle can be ignored in the analysis. When adopted, allows us to ignore the extension of the vehicle nose in front of the perception sensor, for example.
- **low latency assumption**: the assumption that the delays associated with passing energy or information through an element of the system can be ignored. When adopted, allows us to ignore actuator dynamics, for example.
- **flat terrain assumption:** the assumption that the terrain is at least locally flat at the scale of the sensory lookahead distance. When adopted, allows us to simplify many aspects of the analysis.
- **smooth terrain assumption:** the assumption that the terrain does not contain any high spatial frequencies. When adopted, allows to assume reasonable limits on the need to resolve small hazards in the environment.
- **stationary environment assumption**: the assumption that the environment is rigid. When adopted, allows us to measure the position of an object only once and assume that it stays put while the vehicle moves around it.

2.6 Standard Problems

Given the description of the problem outlined above, a set of natural subproblems emerge when vehicle subsystems do not meet the underlying requirements of the problem of autonomous mobility. Many of the following subproblems will be subsequently elaborated in more detail.

2.6.1 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.

2.6.2 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 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.

3.1 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**.

3.2 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.

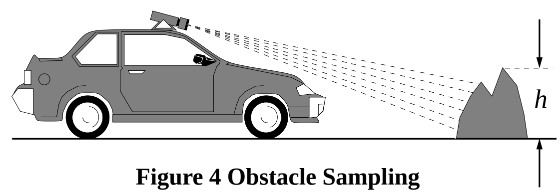

Figure 4 Obstacle Sampling

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.

3.3 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:

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

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.

3.4 Sampling Problem

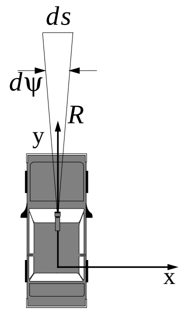

$$x = R\sin\psi$$
$$ds = Rd\psi$$
$$dx = ds = Rd\psi$$
$$d\psi = dx/R$$

Figure 5 Crossrange Resolution

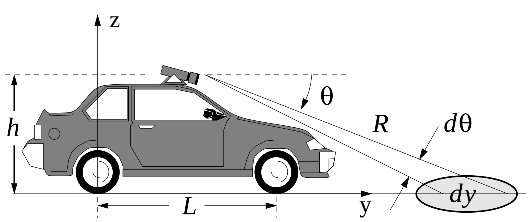

$$R = \frac{h}{\sin\theta}$$

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

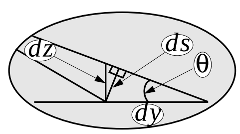

$$z = h - R\sin\theta = 0$$

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

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

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

Figure 6 Downrange Resolution

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**.

3.4.1 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.

3.4.2 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.

3.5 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.

3.5.1 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:

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.

3.5.2 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.

3.5.3 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.

3.6 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.

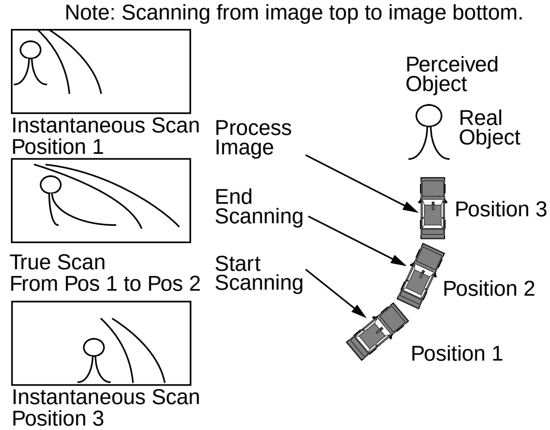

Figure 7 Motion Distortion Problem

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.

4 Fidelity

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

4.1 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.

4.1.1 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.

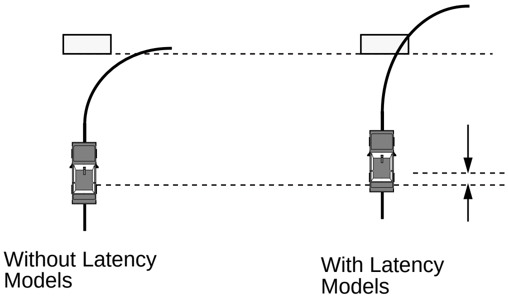

Figure 8 Unmodelled Latency Problem

4.1.2 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:

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

4.1.3 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.

4.2 Ackerman Steering Kinematics

4.2.1 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.

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

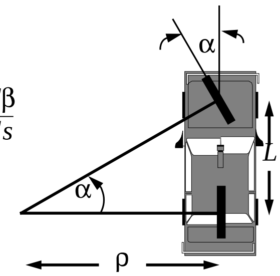

Figure 9 Bicycle Model

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:

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

4.2.2 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:

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

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

$$\psi(t) = \psi_0 + \int_0^t V(t)\kappa(t)dt$$

4.2.3 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.

4.3 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.

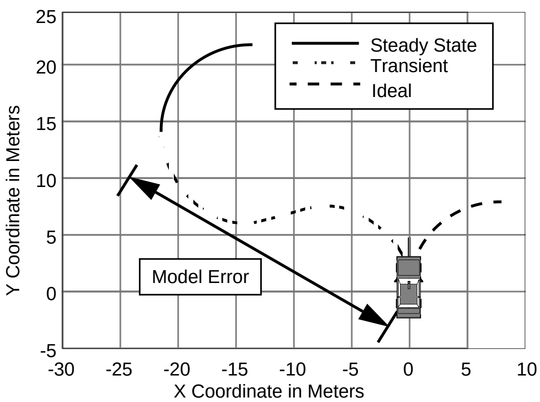

Figure 10 Transience in the Reverse Turn

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.

4.3.1 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:

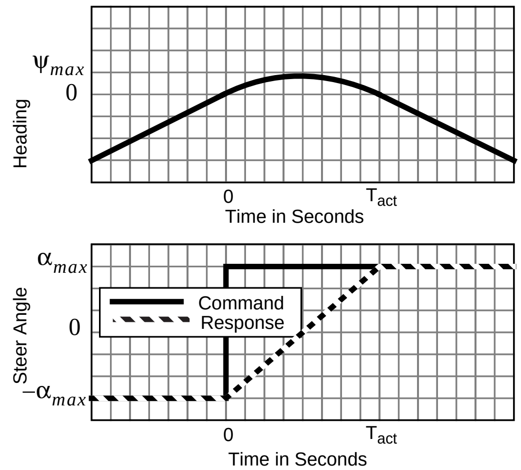

Figure 11 Transient Steering Response

4.3.2 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:

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

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.

4.3.3 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**:

This number must be significantly less than unity to allow

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

ignoring dynamics. It is often on the order of 10.

4.4 Exact Ackerman Steering Dynamics

While instantaneous response models of vehicle actuators is a useful theoretical approximation, and while it is a good model of braking, the same is not true of turning. Steering actuators exhibit dynamics that often must be modelled in practice. This section presents a reasonably 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.

4.4.1 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.

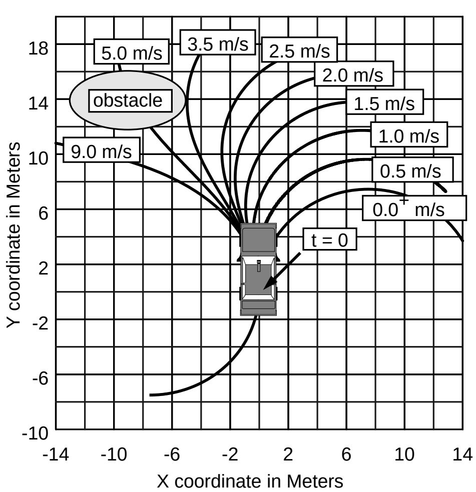

Figure 12 Constant Curvature Reverse Turn

4.4.2 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.

4.4.3 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.

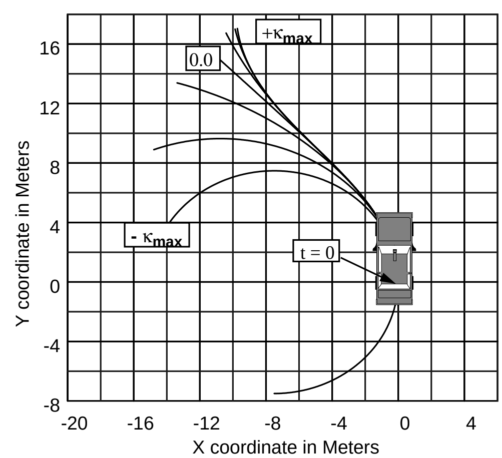

Figure 13 Constant Speed Reverse Turn

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

5 Conclusions

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

5.1 Sensor Mounting Geometry

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.

A primary difficulty associated with a low perception ratio is a severe groundplane sampling problem. However, certain simple approximations to fundamental requirements become available which are quite useful. Pixel aspect ratio is the perception ratio under these conditions and the area of a pixel footprint is cubic in range.

Elegant expressions for the angular resolution required of sensors also become available which depend directly and only on properties of the vehicle such as wheelbase, wheel radius, and (through sensory lookahed) the response distance.

Both the sampling problem and the motion distortion problem are expected to be severe in a typical situation.

5.2 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.

5.3 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.

5.4 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.

6 References

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

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

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

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

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

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

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

[8] K. E. Olin and David Tseng. “Autonomous Cross Country Navigation”, IEEE Expert, August 1991.

[9] 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

[10] In So Kweon, “Modelling Rugged Terrain by Mobile Robots with Multiple Sensors”, CMU PhD Thesis, 1990