Calibration and Control for Range Imaging in Mobile Robot Navigation

Ole H. Dørum
Norwegian Institute of Technology
Division of Computer Systems and Telematics
Trondheim, Norway

Adam Hoover
Dept. Computer Science and Engineering
University of South Florida
Tampa, FL 33612 USA

Judson P. Jones
Oak Ridge National Laboratory
Oak Ridge, TN 37831-6364 USA

Abstract

This paper addresses some issues in the development of sensor-based systems for mobile robot navigation which use range imaging sensors as the primary source for geometric information about the environment. In particular, we describe a model of scanning laser range cameras which takes into account the properties of the mechanical system responsible for image formation and a calibration procedure which yields improved accuracy over previous models. In addition, we describe an algorithm which takes the limitations of these sensors into account in path planning and path execution. In particular, range imaging sensors are characterized by a limited field of view and a standoff distance - a minimum distance nearer than which surfaces cannot be sensed. These limitations can be addressed by enriching the concept of configuration space to include information about what can be sensed from a given configuration, and using this information to guide path planning and path following.

1 Introduction

This paper addresses two issues in the development of sensor-based systems for mobile robot navigation which use scanning laser range cameras as the primary sensor. Our motivation comes from the vision of autonomous robots replacing human beings in applications where hazardous environments pose a threat to life or health. It is generally agreed that the technology necessary to realize this vision will require sensor-based systems for navigation and object recognition, as well as a host of other capabilities too numerous to mention. Earlier we reported on complete systems designed for point-to-point navigation in a priori unknown environments using laser range cameras [5] and using multiple sensors [1]. These systems were constructed around a simple representation of the environment, an occupancy map [2] and [3], which can be computed reliably by sensors and which is also a useful form for navigation. An occupancy map represents space as an ordered set of elementary areas (pixels) or volumes (voxels), each of which can assume the binary states "empty" or "occupied". This simple representation has a variety of attractive properties for our purposes. Here we mention only two. First, occupancy maps are easy to construct given the data provided by active ranging sensors such as sonars or imaging laser radars, provided that these sensors are accurately calibrated. Second, reasonably efficient motion planning algorithms are available which use this representation as their input.

As these systems developed, we had the opportunity to improve the performance of some of the constituent components. In particular, laser range imaging systems have recently become available which have greatly improved resolution. To take advantage of this requires a more accurate camera model and a more sophisticated calibration procedure. In the first section of this paper, we present such a model and a calibration procedure for it. In the second part of the paper we present an improved set of heuristics for guiding a navigating robot through an a priori unknown environment using a sensor with a limited field of view and a minimum standoff distance. We have found that incorporating these limitations into path planning

is necessary in order to avoid planning motions into regions where the sensors cannot disambiguate unknown territory.

2 Camera Calibration

The purpose of the camera calibration is to determine the transformation equations which map each pixel $(row, col, range)$ of the range image into Cartesian coordinates (x, y, z). This is accomplished by finding the camera model which describes the path of the laser beam from the laser diode's emission point to the target[1] it strikes as a function of the row and column number of the image. Camera transforms for Laser Range Finders (LRFs) have generally been treated in a simplified manner [7],[8]. The *Orthogonal Axis model* and the *Azimuth-Elevation model* [7] are instances of mapping from spherical coordinates to Cartesian coordinates and differ only in how the angles are defined. These models assume that the rays that scan the scene eminate radially from a virtual focus. Due to the scanning mechanism used in laser range cameras these assumptions do not hold. Furthermore, these models also assume that range value 0 maps to the origin of the coordinate system. We have therefore developed a model which does not require these assumptions and a calibration procedure to establish the parameters of the model.

2.1 Electro-Mechanical Operation

The laser source of the Perceptron's 5000 LASAR Scanner is a 100mW peak output (50mW average) Continuous Wave laser diode which has a wavelength of 835nm. The laser is amplitude modulated at a frequency of about 17.75MHz. The range is determined by measuring the phase difference between the outgoing laser beam and the incoming laser light being reflected off the target. Since the scanner cannot detect whether the phase difference is beyond 2π or not, the wrap-around interval of the range, known as the *ambiguity interval*, is equal to the half wavelength of the AM frequency, 8.45m. The phase difference is discretized into 12 bits yielding range measurements precise to 2.0mm. There are two moving mirrors which enables the camera to scan the scene (see Fig 1). Initially, the laser beam it hits the *rotating mirror*, a three-faceted mirror with each facet sloping backwards. This mirror performs the sweeping motion of each raster scan. The laser gets reflected upward where it hits the *nodding mirror*. The nodding mirror deflects the upward motion of the laser beam into the scene. It is the position of

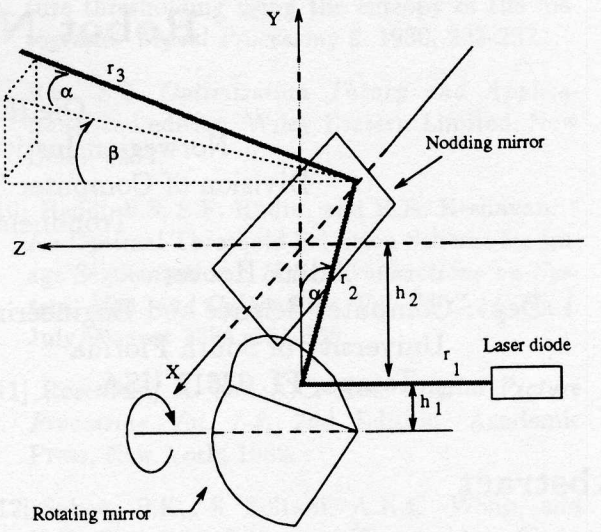

Figure 1: *This figure illustrates the scanning geometry of the camera. The x-axis of the camera's right-handed coordinate system is embedded in the nodding mirror axis with the y-axis at the center of the mirror axis pointing straight up and the z-axis pointing straight into the scene. The laser is emitted by the laser diode and travels parallel to the z-axis the distance r_1 before it hits the rotating mirror. The sloping angle of each facet of the mirror reflects the laser upward parallel to (or in) the xy-plane. In that plane, the angle of rotation α of the rotating mirror causes the laser to deviate the shown angle α from its otherwise vertical path. It is this sweeping motion which scans each raster. In its upward path r_2 between the mirrors, the laser strikes the nodding mirror at the point $p_n = (x_0, y_0, z_0)$ close to (or on) the x-axis where it gets reflected out into the scene where it travels the distance r_3 from the mirror until it hits the target.*

the nodding mirror which determines which raster of the image is currently being scanned.

2.2 Mathematical Model

In order to describe the mechanical operation of the scanner it is necessary to refer to its local coordinate system. We have chosen our right-handed local coordinate system (see Fig. 1) to have its origin at the center of the nodding mirror axis with the x-axis embedded in the mirror axis, positive y-axis pointing straight up, and the positive z-axis point-

[1] Target is defined as the point where the laser hits an object in the scene.

Symbol	Description
H	Horizontal field of view.
V	Vertical field of view.
H_0	Horizontal offset angle.
V_0	Vertical offset angle.
α	Angular position of rotating mirror (=0 when the normal of a facet lies in the yz-plane).
β	Angle between xz-plane and plane going through the x-axis containing the laser beam r_3.
$\gamma(45°)$	Slope of the facets of the rotating mirror with respect to the z-axis.
$\theta(45°)$	Angle of the nodding mirror with respect to the z-axis when $\beta = 0$.
$h_1(3.0cm)$	Distance (y) between rotating mirror axis and the parallel laser beam r_1.
$h_2(5.5cm)$	Distance (y) between nodding mirror axis and the laser beam r_1.
l	Length of r_1 when $\alpha = 0$.
δ	Length of one range unit in [cm].
R	True length of laser beam from laser diode to target.
r	Range value returned by the scanner.
r_0	Standoff distance; length of laser beam to the point on the beam where $r = 0$.
r_1	Length of laser path from laser diode to where it strikes the nodding mirror.
r_2	Length of laser path between rotating mirror and nodding mirror.
r_3	Length of laser path from where it leaves the nodding mirror until it strikes the target.
p_r	Point of intersection of the laser beam with the rotating mirror.
p_n	Point of intersection of the laser beam with the nodding mirror.
x_0, y_0, z_0	Coordinates of p_n.
t_i	(x, y, z) coordinates of calibration target i.

Table 1: *Symbols used in the camera model. The parenthesis hold the nominal value for those symbols for which it is known.*

ing perpendicularly out of the camera. The angular position of the mirrors is a linear function of the row and column number of the image (image dimension: $n+1$ rows and $m+1$ columns):

$$\alpha = H_0 + H\left(\frac{\frac{m}{2} - col}{m}\right) \text{ for } 0 \leq col \leq m \quad (1)$$

$$\beta = V_0 + V\left(\frac{\frac{n}{2} - row}{n}\right) \text{ for } 0 \leq row \leq n \quad (2)$$

When $\alpha = 0$ the rotating mirror projects the laser beam straight up along the y-axis and when $\beta = 0$ the laser beam is projected out of the scanner in the xz-plane (along the z-axis if both angles are 0). The point of intersection p_r of the laser beam with the rotating mirror moves parallel to the z-axis as the mirror rotates. It is furthest away from the laser diode when $\alpha = 0$ at $z = 0$. The distance p_r moves is

$$dz = -h_1\frac{(1 - \cos\alpha)}{\tan\gamma}$$

Since the slope of each facet is $\gamma = 45°$ the reflected beam always lies in a plane parallel to the xy-plane. The angle between the laser beam and the yz-plane is α so that the direction of the laser beam is $[\sin\alpha\ ,\ \cos\alpha\ ,\ 0]$. Since the angle of the nodding mirror with respect to the z-axis is $\theta = 45°$ when $\beta = 0$ we get for an arbitrary β that the mirrors angle is $\theta + \frac{\beta}{2}$. The point p_n where the laser intersects the nodding mirror is therefore:

$$\begin{aligned} x_0 &= (h_2 + y_0)\tan\alpha \\ y_0 &= dz\tan(\theta + \tfrac{\beta}{2}) \\ z_0 &= dz \end{aligned} \quad (3)$$

Since the nodding mirror rotates about the x-axis it can only change the y and z component of the laser beam. We get that the direction of the laser leaving the mirror is $[\sin\alpha\ ,\ \cos\alpha\sin\beta\ ,\ \cos\alpha\cos\beta]$.

To find the length r_3 of the laser from the mirror to the target we first need to determine the length of the laser path up to the point where it leaves the nodding mirror, namely $r_1 + r_2$. The total length R of the beam can also be expressed as the range value r returned by the camera plus the standoff distance r_0. We get the following relationship:

$$R = r + r_0 = r_1 + r_2 + r_3 \quad (4)$$

$$\Rightarrow r_3 = r + r_0 - (r_1 + r_2) \quad (5)$$

where the path length from the laser diode to the rotating mirror is $r_1 = (l + dz)/\delta$ and the path length from the rotating mirror to the nodding mirror is $r_2 = \sqrt{x_0^2 + (h_2 + y_0)^2}/\delta$. We can now set up the

complete transformation equations:

$$\begin{aligned} x &= x_0 + r_3 \sin\alpha \\ y &= y_0 + r_3 \cos\alpha \sin\beta \\ z &= z_0 + r_3 \cos\alpha \cos\beta \end{aligned} \quad (6)$$

The transformation equations can broken down into two parts. The first part is the offset term (x_0, y_0, z_0) which describes where the laser leaves the nodding mirror. The second part is simply a mapping from spherical to Cartesian coordinates given by the direction of the laser (the trigonometric terms) and the length of the laser r_3 from the nodding mirror to the target. Notice that what distinguishes this model from [7], [8] is the *offset term* which translates the coordinate system's origin so that it intersects the laser. This enables a standard spherical to Cartesian mapping to be used on the part of the laser beam for which it applies, namely the *external path* (r_3).

2.3 Determining the Parameters of the Model

The process of determining the parameters of the scanner is decoupled so that the parameters can be solved for analytically avoiding the need for searching the parameter space. The parameters given by the geometry of the mirrors are: $h_1, h_2, l, \theta, \gamma$, the remaining parameters: $H, H_0, V, V_0, r_0, \delta$ are determined experimentally (see Table 1).

2.3.1 The Standoff Distance r_0 and Range Unit δ

To determine the standoff distance r_0 and the length of a range unit δ we take multiple images of a planar surface at several different z values (we chose intervals of 30-50cm). The center of the plane is always aligned with the center (Field Of View) FOV pixel ($\alpha = \beta = 0$) and perpendicular to that ray (in most cases it will be very close to the z-axis). For each z position of the plane we took 4 images and measured the z distance from the origin of the camera's local coordinate system to the plane. At the pixel of interest we placed a 7×7 *median mask* to deal with noise in the data. The most consistent median range value over each of the four images was chosen as the correct range value. The next step is to plot range vs. distance [2](see Fig. 2). Outliers are determined by either visual inspection of a manual line fit to the data or by performing a *Least Median of Squares* line fit. Once the outliers are removed, a standard *Least Squares* line fit is used on the inliers. The equation of the fitted line expresses the range as a function of distance:

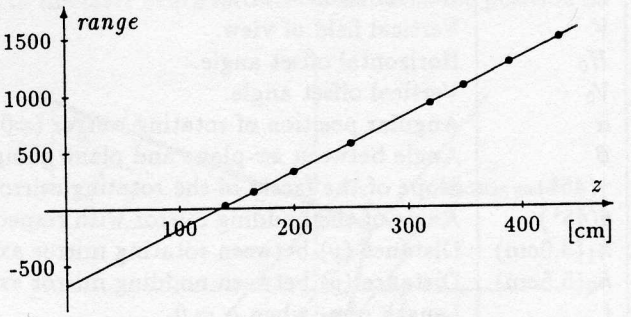

Figure 2: *Range plotted vs. distance for center FOV pixel.*

$$range(z) = \frac{1}{\delta} z - r_0 = 4.844\, z - 655.95$$

yielding the desired parameter values $r_0 = 656.0$ range units and $\delta = 0.20644$cm.

2.3.2 Field Of View Parameters H, V, H_0, V_0

In order to find these parameters it is necessary to determine the horizontal and vertical slopes α and β of r_3 for a set of pixels in the image. For this purpose we construct a calibration scene in which we place calibration targets. The coordinates of the calibration targets t_i are measured to within 1.0mm using standard surveying equipment and triangulation methods. The row and column pixel position of each calibration target is currently determined by manual inspection of the intensity image returned by the scanner along with each range image. This step will in the future be replaced by an image processing step yielding sub-pixel accuracy. At this point we have for each calibration target in the image both its (x, y, z) position with respect to the local coordinate system of the scanner and its (sub)-pixel position in the image. Notice that the range values of the calibration targets are not used in the step. To determine the approximate direction of the laser beams striking the calibration targets, we find the approximate *virtual* focus for the yz-component and the xz-component of the rays separately. The yz-components converge close to the x-axis while the xz-components converge at a narrow vertical band in the yz plane behind the nodding mirror at $z \approx -h_2$. This step although approximate yields angles accurate to about $\pm 0.1°$ for calibration targets at least 4m away from the scanner.

[2]Since the z distance is measured from the origin, we either have to use *true* path length: $z_{true} = z + h_2 + l$, or we can simply set $l = -h_2$ and use the z-distance as measured.

For each calibration target t_i we can find its approximate angle β_{ap}:

$$\beta_{ap} = \arctan\left(\frac{y}{z}\right)$$

From the approximate focus of the rays xz-component behind the mirror we can find α_{ap}:

$$\alpha_{ap} = \arctan\left(\frac{x\cos\beta_{ap}}{z+h_2}\right)$$

From Eq. 3 using the approximate expressions for α and β for each calibration point we can now compute the approximate point $p_n = (x_0, y_0, z_0)$ where each ray leaves the nodding mirror.

The true angles can now be expressed as:

$$\beta = \arctan\left(\frac{y-y_0}{z-z_0}\right) \quad (7)$$

$$\alpha = \arctan\left(\frac{(x-x_0)\cos\beta}{z-z_0}\right) \quad (8)$$

Since we for each calibration target i have determined its (sub)-pixel position (row_i, col_i) we have the following information about each calibration target t_i : $\{row_i, col_i, \alpha_i, \beta_i\}$. For each calibration target we can set up the following equations to solve for the FOV parameters:

$$\alpha_i = H_0 + H\left(\frac{\frac{m}{2}-col_i}{m}\right) \quad (9)$$

$$\beta_i = V_0 + V\left(\frac{\frac{n}{2}-row_i}{n}\right) \quad (10)$$

We can for N calibration targets express the equations in matrix form $\mathbf{a} = \mathbf{A}\,\mathbf{h}$ and $\mathbf{b} = \mathbf{B}\,\mathbf{v}$ where row i of matrix $\mathbf{A}$ is $[1 \quad \left(\frac{\frac{m}{2}-col_i}{m}\right)]$ and where row i of matrix $\mathbf{B}$ is $[1 \quad \left(\frac{\frac{n}{2}-row_i}{n}\right)]$. We solve for the unknown FOV parameters $\mathbf{h}, \mathbf{v}$ by finding the *pseudoinverse* [3] $\mathbf{A}^+$ of $\mathbf{A}$ and $\mathbf{B}^+$ of $\mathbf{B}$. The Least Squares solution becomes

$$\mathbf{h} = \mathbf{A}^+\mathbf{b} \quad (11)$$

$$\mathbf{v} = \mathbf{B}^+\mathbf{a} \quad (12)$$

where $\mathbf{h} = [H_0 \;\; H]^T$,$\mathbf{v} = [V_0 \;\; V]^T$ and where $\mathbf{a} = [\cdots\alpha_i\cdots]^T$,$\mathbf{b} = [\cdots\beta_i\cdots]^T$.

2.3.3 Conclusion

The presented procedure has been applied to a Perceptron 5000 LASAR Scanner. Initially we applied the simplified spherical camera model which yielded warps in the mapping of the transformation equations. Consequently, surfaces which should have been planar were warped and angles between surfaces were distorted as well. This led us to the conclusion that such a simple model was not accurate enough. The operation of the scanner was investigated in detail to improve the model. That effort resulting in the present camera calibration procedure (yielding the following parameters for the scanners current programmed mode of operation: $H = 61.03°, V = 37.68°, H_0 = -3.89°, V_0 = -1.05°, r_0 = 656.0$ range units and $\delta = 0.20644$cm). Although precise experiments for verifying the quality of the model and its parameters have not yet been performed, preliminary results indicate correct surface shape is obtained along with true angles between surfaces.

[3] Since $\mathbf{A}$ and $\mathbf{B}$ are rectangular $N \times 2$ matrices.

3 Incorporating Sensor Constraints in Motion Planning

The problem of motion planning in in a partially known environments presents several challenges for real robotic systems, including constraints on the possible motions of the robot, bounds on robot velocity and acceleration and, most importantly here, sensor constraints. Figure 3 illustrates the field of

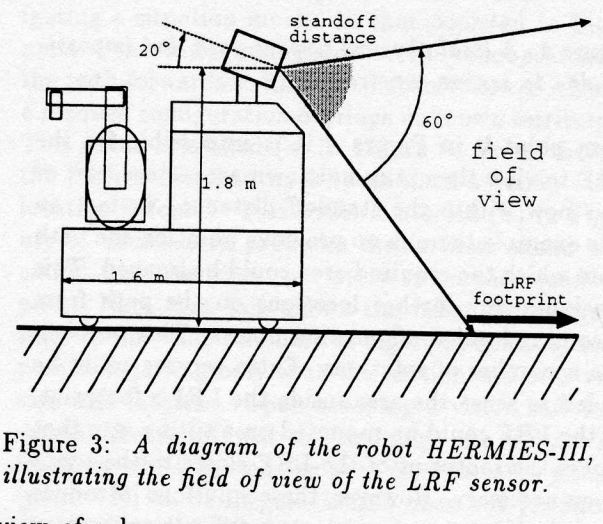

Figure 3: *A diagram of the robot HERMIES-III, illustrating the field of view of the LRF sensor.*

view of a laser range camera, as mounted on the robot HERMIES-III [4]. Notice the area between the footprint of the sensor and the front of the robot. It is impossible for the range camera to sense obstacles in this area. Although the geometry of this area is a consequence of the construction of this particular robot, active and passive ranging sensors, such as sonar, stereo vision, and laser range finders all have a characteristic minimum *standoff distance*. Objects closer than the standoff distance cannot be

sensed due to sensor-specific limitations. In planning the motions of mobile robots, it is desirable to take these and other sensor limitations into account prior to moving the robot. To illustrate the effects of ignoring these limitations, consider Figure 4, an overhead view of HERMIES-III navigating through a partially known environment. At point (a), the robot is approaching a corner. The areas marked as *obstacle* and *empty* have been previously imaged by the LRF and are stored in a global map. The area marked as *unknown* has not been imaged yet. Since unknown territory may be unoccupied, the path marked in Figure 4 may be physically traversable by the robot. To verify whether the unknown area is freespace, the robot must travel along the path to a point where its view is unobstructed. However,

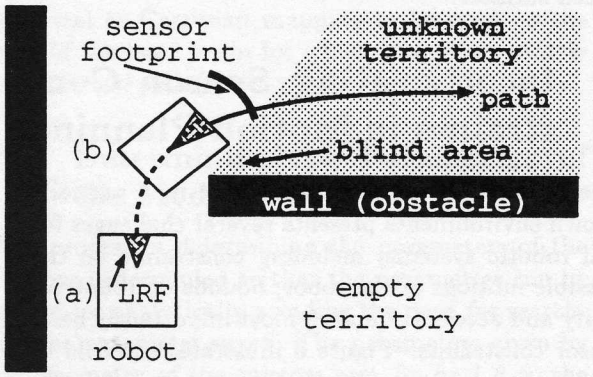

Figure 4: *A legal physical Cspace path, but impossible due to sensor constraints.*

from point b in Figure 4 it is impossible for the LRF to view the entire unknown area, since part of it is now within the standoff distance. In fact, in this example there is no previous point on the path from which the required area could be imaged. This precludes any further locations on the path from becoming legal configuration points. This problem has a number of solutions. Extra sensors could be added to sense the area inside the LRF's footprint, or the LRF could be mounted on a tilting arm that moves the footprint of the LRF closer to the robot when necessary. However, these solutions introduce extra hardware costs and pose difficult calibration problems. Alternatively, we could increase the apparent size of the robot to encompass its minimum sensing distance while planning motion. However, this unnecessarily restricts the motion capability of the robot. Finally, it is possible to keep the robot at least its minimum sensing distance from any unknown territory. However, as the example in Figure 4 shows, there are cases where this does not suffice.

We have instead taken the approach of explicitly incorporating sensor constraints in the motion planning problem. To do this we use a version of configuration space. The next section explains the expanded Cspace and the following demonstrates its success in application to HERMIES-III.

3.1 Searching Cspace with Sensor Constraints

The projection of obstacles into the configuration space (Cspace) of a robot defines the positions and orientations which a robot can occupy without collision. This information is enough for motion planning in a known environment, or when ideal sensing is possible. However, when motion planning in a partially known environment in which real sensing is required to discover new legal configurations, the limitations of the sensors must be considered. Conceptually, this adds a property to the Cspace representation of the robot. A complete description of this property would define exactly what area or volume the sensors can measure from the given physical configuration, and what areas and volumes are occluded by known obstacles. However, an algorithm for searching configuration space given a complete description of these properties would be computationally burdensome.

Instead, we have developed a a few, somewhat less precise concepts which describe a robot's sensing capabilities, both for reasons of generality and practicality. For the purposes of discussion, we will illustrate these concepts using the HERMIES-III robot as an example. Figure 5 illustrates a taxonomy of configurations we use in navigating HERMIES-III.

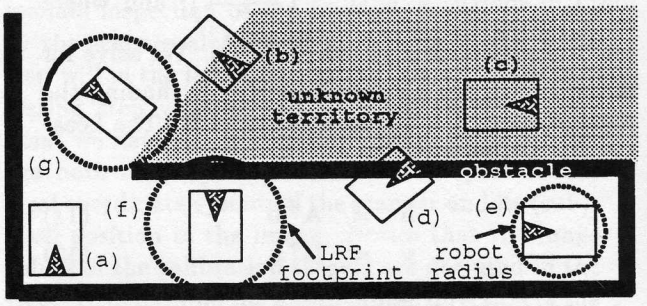

Figure 5: *Configurations are classified according to what can and cannot be sensed.*

The robot at position (d) is in a collision configuration. The configurations illustrated at positions (b) and (c) are classified as non-collision configurations for the purpose of path planning, and collision configurations in path following. Configuration (e)

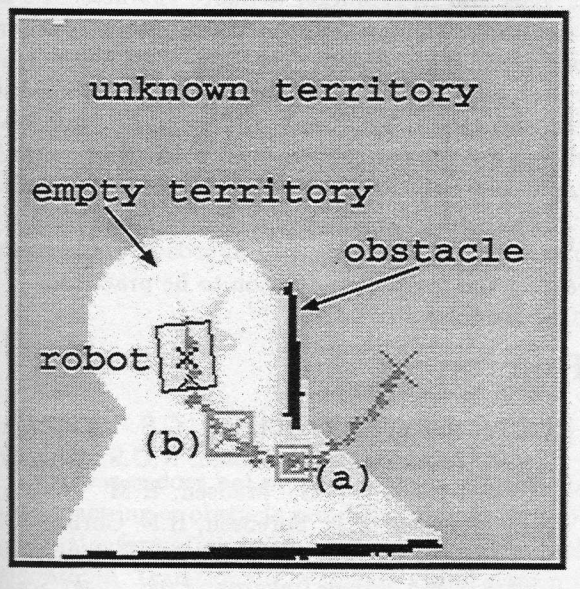

Figure 6: *All points up to (a) are physical safe configurations. Backtracking from (a), the farthest ahead vision safe location is (b). HERMIES-III can safely move to point (b), but no further.*

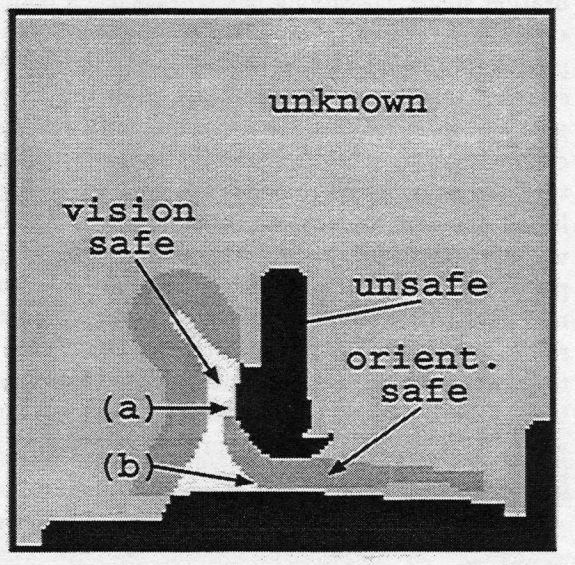

Figure 7: *Four levels of safety are depicted in the map: unsafe (collision), physical safe, vision safe and unknown. HERMIES-III is stuck at point (a), and will use point (b) as a temporary off-path goal.*

shows a location in which the robot would not collide with obstacles in any orientation. We call such a location *physical safe*. This definition allows us to search a reduced Cspace (by one physical dimension) for motion planning. It also inherently implements the optimality criteria of trying to keep the robot some minimum distance from obstacles. Finally, configurations (f) and (g) of Figure 5 illustrate our notion of the sensing limitations. For a location to be *vision safe* it (1) must be physical safe and (2) must have no unknown territory within the sensor footprint radius, *unless that territory is occluded by an obstacle*. Under this definition, location (f) is vision safe, but location (g) is not.

Motion planning is done from a location which is vision safe. This allows the robot to plan a path in any direction from the location. We are assured that any adjacent location can be discovered to be either a collision configuration or a legal configuration. We have implemented a motion planner based upon [6]. The algorithm has been extended to incorporate the two optimality criteria described above. Note that these are desired or favored, not essential, criteria. The minimum condition which must be satisfied is simply a non-collision configuration.

3.2 Navigating with Sensor Constraints

Figure 6 shows an example of HERMIES-III navigating a situation much like that depicted in Figure 4. Prior to beginning motion, the robot searches the path forwards to find the first instance of where a planned configuration overlaps unknown territory. This point is labeled (a) in Figure 6. From there, the path is searched backwards until a vision safe location is found. This point is labeled (b) in Figure 6. This search produces the farthest ahead location that the robot can move to on the planned path. Note that this allows the robot to traverse single-point-wide corridors of legal physical Cspace points, requiring only that the motion end at a vision safe location. Upon reaching a vision safe location, the robot will try to image the unknown territory which caused the next configuration point (point (a)) to be illegal. If this does not produce any further vision safe locations, it tries to image the unknown territory preventing further path locations (between (a) and (b)) from being vision safe. If a vision safe location can be created further along the path, motion resumes. Otherwise, if the imaged territory is found to create a collision on the path, the robot plans a new path.

In the example illustrated, no unknown territory can be be imaged on the new path either. As a final action, the robot exhaustively searches config-

uration space, classifying configurations. Figure 7 shows a *safe map*, completely enumerating which points in the world are vision safe. The robot searches this map for a vision safe point that is *off-path* and from which the elusive unknown territory can be imaged. If multiple candidates are available, then the point which provides the greatest triangulation between the current robot location and the unknown territory to be imaged is chosen. This point, labeled (b) in Figure 7, becomes a temporary goal. After arriving at the temporary goal, the robot images the unknown territory that caused it to take this action and then goes back to the original task. Figure 8 shows the conclusion to this example where HERMIES-III is in the process of reaching its goal.

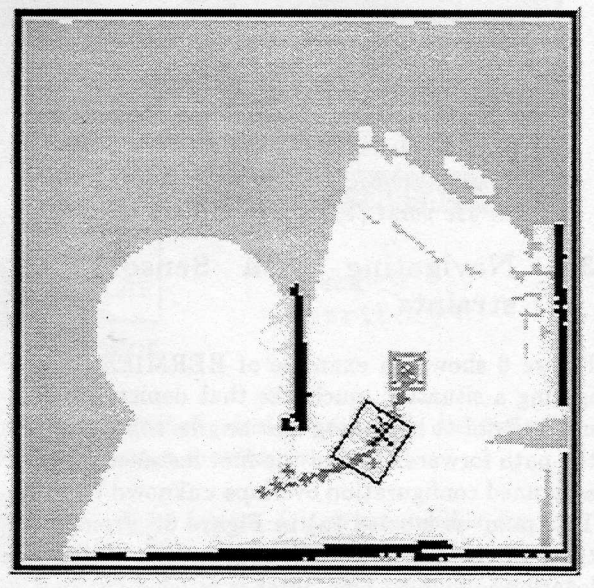

Figure 8: *Finally, HERMIES-III is able to image the troublesome unknown territory. The robot is then able to continue navigating towards its original goal.*

4 Conclusion

In this paper we have presented a new model for scanning range imaging sensors, and a method for calibrating the model. The present model differs from previous models in that explicit consideration is taken of the non-lens-like geometry of the image formation mechanism, which in practice means that rays do not converge to a focus. At the level of resolution of the Perceptron camera, we believe it is necessary to model this mechanism explicitly in order to take full advantage of the camera. We have also described a method by which the limited field of view and the standoff distance characteristic of laser range cameras can be taken into account in path planning and path execution. This method involves scoring configurations on the basis of what can and cannot be sensed, and using this information to guide a search for a legal path. This method could probably be made more accurate by explicitly increasing the number of dimensions in configuration space; however, this approach would likely require too much computation to be practical.

References

[1] J.P. Jones, O.H. Dørum, C.S. Andersen, S.B. Jacobsen, M.S. Jensen, N.O.S. Kirkeby, S. Kristensen, C.B. Madsen, H.M. Nielsen, E. Sorensen, J.J. Sorensen, H.I. Christensen (1993). Experiments in Mobile Robot Navigation and Range Imaging. *Proc. Eight Scandanavian Conference on Image Analysis*, May 25-28, 1993, Tromsø, Norway 371-387.

[2] A. Elfes, Sonar-based real-world mapping and navigation. *IEEE J. Robotics and Automation* (RA-3) 249-265, 1987.

[3] A. Elfes, Using occupancy grids for mobile robot perception and navigation. *Computer* (22:6) 46–58, 1989.

[4] C. R. Weisbin, B. L. Burks, J. R. Einstein, R. R. Feezell, W. W. Manges, D. H. Thompson. HERMIES III: A step toward autonomous mobility, manipulation and perception. *Robotica* (8) 7–12, 1990.

[5] C.S. Andersen, C.B. Madsen, J.J. Sorensen, N.O.S. Kirkeby, J.P. Jones, H.I. Christensen, Navigation using range images on a mobile robot. *Robotics and Autonomous Systems* (10) 147–160, 1992.

[6] J. Barraquand and J. C. Latombe, Robot Motion Planning: A Distributed Representation Approach. *Stanford University Technical Report*, 1989.

[7] P. J. Besl and R. C. Jain, Segmentation Through Variable-Order Surface Fitting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol 10, No. 2, 167–192, March 1988.

[8] M. Hebert and E. Krotkov, 3D measurements from imaging laser radars: how good are they?. *Image and Vision Computing*, vol 10. no 3, 170–178, April 1992.