DEVELOPMENT OF TECHNIQUES REQUIRED FOR THE APPLICATION OF A LASER TO THREE DIMENSIONAL VISUAL SENSING

NAGW-1333

by

Arthur M. Ryan and Lester A. Gerhardt

Rensselaer Polytechnic Institute
Electrical, Computer, and Systems Engineering Department
Troy, New York 12180-3590

August 1991

CIRSSE REPORT #101

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The ongoing vision research at the Center for Intelligent Robotic Systems for Space Exploration (CIRSSE) is directed toward identifying and addressing the relevant issues involved in applying visual sensing to space assembly tasks. A considerable amount of effort has been devoted to passive sensing techniques such as using multiple cameras to identify objects in a scene. To compliment the capabilities of the passive visual system in the CIRSSE robotics testbed, research is being conducted in active sensing techniques. This report is a description of the research associated with the testbed's laser scanner and its application as an active sensing device. The report is comprised of five major topics. First is a brief description of the CIRSSE visual system and a summary of the active sensing sensing research that has been conducted up to this point. Second, some of the methods currently used to calibrate CIRSSE's laser scanner are described as well as an appraisal of the effectiveness of these methods. Third, is a discussion of how the laser scanner can be employed in concert with a camera to provide a three dimensional point estimation capability. Fourth, there is a description of methods that can be used to detect the presence of the laser beam in a cluttered camera image. Finally, there will be a summary of the current state of this research and a description of research planned for the future.

1.1 Description of CIRSSE Testbed

The CIRSSE robotic testbed is designed to support research in robotic assembly tasks for space applications. The centerpiece of the testbed is a pair of PUMA robots each of which is mounted on a movable cart. Both carts

are mounted on a track; thereby permitting the robots to operate over a large work volume. The testbed is also equipped with a multi-camera vision system which provides a three dimensional visual sensing capability. The vision system is equipped with five cameras (two mounted above the testbed, two mounted on the wrist of one PUMA robot and one camera mounted on the second PUMA robot). The vision system is also equipped with a laser scanner (also mounted above the testbed) manufactured by General Scanning [1]. This laser scanner provides the capability to direct a laser beam by deflecting it with a pair of mirrors that can be controlled by computer. The entire vision system is controlled by an image processing system manufactured by Datacube.

The CIRSSE vision system provides the capability to support research in both passive and active techniques. A passive technique is one that uses available light sources (i.e. general illumination) while an active technique employs the projection of some externally supplied prestructured light [2].

The laser scanner provides a unique active sensing capability that complements and enhances the passive sensing capabilities of the CIRSSE vision system. Such an enhancement is possible by virtue of the laser's ability to inject a pre-defined signature into the workspace that a camera can subsequently detect. A laser beam has specific direction, intensity, wavelength, and predictable behavior when it reflects off of objects. Further, because the laser is an active sensor, it can operate under certain lighting and environmental conditions in which a strictly passive system would be unable to function.

For the purposes of system calibration and three dimensional point estimation, the CIRSSE testbed has a well defined world origin point. This point is located at the center of the track utilized by the robot carts. The exact location and orientation of this point is described in a CIRSSE technical memorandum [3]. This report will refer to the testbed's world point as either the *world origin* or the *world coordinate system.*

1.2 History and Development of Laser Research Conducted at CIRSSE

Research involving the laser scanner began shortly after the vision system was installed in the CIRSSE testbed. Initially, the laser scanner was used to generate discrete or continuous patterns such as grid lines or predefined

shapes. This pattern generating capability is possible since the laser scanner is equipped with a shutter that can block the laser beam when blanking is required between points. At this early phase of the research, the laser could be controlled only in terms of its scanning mirrors, Initially, there was no ability to control the laser in terms of cartesian space.

The next phase of the research was directed towards using a camera as a feedback mechanism to direct the laser at a specific pixel coordinate. The laser was directed in the field of view of a camera and the camera was used to identify the reflection of the laser beam off objects in the workspace (this reflection is colloquially referred to as the *laser spot*). Control of the laser was still in terms of the angles of the scanning mirrors and these angles were repeatedly adjusted by the computer until the laser spot was centered onto the desired pixel coordinates. The process of locating the laser spot in the camera image was premised on the assumption that the laser spot was the brightest object in the image. Hence, this laser and camera configuration operated best under subdued lighting conditions. This work highlighted several issues that served as the basis for the current laser research:

- The CIRSSE vision was capable of using the laser and cameras in a well coordinated manner.
- Using a camera to detect the laser spot can be a valuable sensing capability since the laser could illuminate objects that the camera might otherwise be unable to distinguish.
- The current camera feedback method assumed that the laser spot is the brightest object in the image. Methods must be developed that circumvent this assumption to permit a laser and camera to operate in a wider variety of environmental conditions.
- If the laser could be calibrated to the same coordinate system as the camera, it would then be possible to directly place the laser at a world point and use a calibrated camera to confirm the proper placement of the laser. Further, because the laser would be calibrated, it would be unnecessary to repeatedly direct the laser to settle on a desired point.

With these ideas in mind the laser research was directed toward accomplishing three major tasks. **First**, calibrate the laser to a level of accuracy comparable to that of the passive multiple camera system. **Second**, adapt

the point estimation techniques used by the multiple camera system for use with a calibrated laser and camera. Third, create techniques that permit a camera to detect a laser spot in a cluttered environment under normal lighting conditions. The result of these three goals would be an active sensing capability that can complement the existing multiple camera system, and provide avenues for future research in active three dimensional ranging and structured light.

Chapter 2

Calibration of a Laser Scanner

A laser is useful in 3-D visual sensing because it provides an active sensing capability. The laser emits a beam of light that a camera can detect as it reflects off of objects within the camera's field of view. An active sensing configuration, such as a laser and camera, can enhance the reliability and flexibility of a vision system since it can generate structured light and "ground truth". In some applications, the laser does not have to be calibrated[2], but calibration is necessary to fully utilize the capabilities of the laser.

The phrase "calibrated laser" is misleading in that the laser itself is not calibrated. It is usually incorporated in an assembly that can direct the laser beam in some well-defined manner. When the laser is calibrated, the entire assembly is actually calibrated. One useful laser assembly is a laser scanner[1] which uses mirrors to deflect the laser beam in a controlled manner.

Calibrating a laser scanner is similar to calibrating a camera in that both devices have intrinsic and extrinsic parameters. For a laser scanner, intrinsic parameters include the distance between the scanner's mirrors and the relationships between the mirror's rotation and the voltage applied to their rotational mechanisms (galvanometers). A laser scanner's extrinsic parameters describe the pose of the laser scanner with respect to some coordinate frame. To better understand the process of laser scanner calibration, it is instructive to examine the internal arrangement of the device and identify the mathematical relationships that describe its operation.

2.1 The Mathematical Model of a Laser Scanner

The internal arrangement of the laser scanner is depicted in Figure 2.1. The device consists of a laser and two scanning mirrors. Each mirror is connected to a galvanometer that rotates the mirror as a function of a control voltage applied to it. The mirrors are configured such that their rotational axes are mutually orthogonal. The rotational axis of the θ_x mirror is parallel with the z axis of the laser scanner, and the rotational axis of the θ_y mirror is parallel to the x axis of the laser scanner.

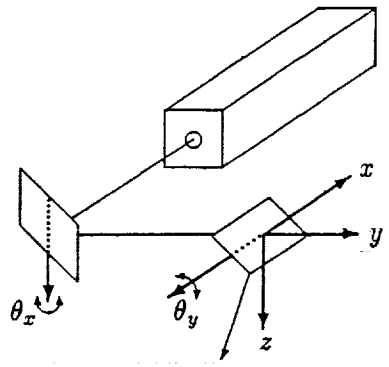

Figure 2.1: Internal Arrangement of a Laser Scanner

Prior research with laser scanners, [4], highlighted the problems associated with creating mathematical relationships between the deflection of the scanning mirrors and the vector of the outgoing laser ray in closed form. If the laser and mirrors are placed at arbitrary locations and orientations, these mathematical relationships become intractable. This is due, in part, to the difficulty in determining the values of some of the necessary parameters[4]. To alleviate these problems, two constraints on the placement of the laser and mirrors must be established. First, the beam emitted by the laser must be parallel to the rotational axis of the θ_y mirror. Second, the laser beam must intersect each mirror at a point along its rotational axis. These constraints are reasonable and practical when one considers that a laser scanner can be assembled with high precision using current manufacturing technology. These constraints reduce the mathematical relationships between the scanning mirrors and the outgoing laser ray to two simple relationships which are described later in more detail.

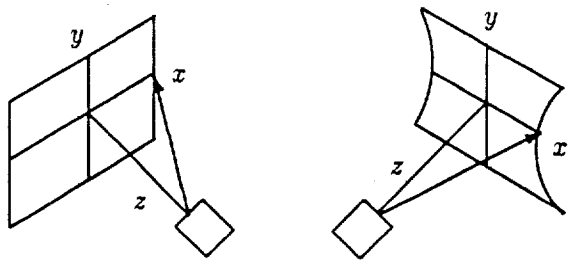

With Pincushion Compensation Without Pincushion Compensation

Figure 2.2: Effect of Pincushion Distortion

When the laser beam is scanned onto a plane normal to the laser's z axis and at some fixed distance from the laser's origin, it is possible to determine the (x, y) coordinates of the laser spot as a function of the mirror's angular deflection. In this situation, the y coordinate of the laser spot is strictly a function of θ_y and the distance to the plane as shown below:

$$y = z \tan \theta_y \tag{2.1}$$

where z is the distance from the laser's origin (L) to the plane normal to the laser's z axis.

If the x coordinate of the spot is assumed to be independent of θ_y, then the laser exhibits *pincushion distortion* (as depicted in Figure 2.2). In reality, the x coordinate of the spot is a function of both θ_x and θ_y. This interdependency is due to the fact that the laser beam strikes the θ_y mirror after it is deflected by the θ_x mirror (see Figure 2.1). The expression for the x coordinate of the laser spot is

$$x = (z \sec \theta_y + e) \tan \theta_x \tag{2.2}$$

where z is the distance from L as in (2.1), and e is the distance between the two scanning mirrors. Note that the $z \sec \theta_y$ term increases as $|\theta_y|$ increases. Hence, the displacement of the laser spot "flares out" away from the origin in the x direction as the spot moves away from the origin in the y direction. Equations (2.1) and (2.2) provide the necessary relationships to direct the laser spot to any 3-D point defined with respect to the laser coordinate frame without pincushion distortion.

The mathematical relationships presented above describe the direction of the laser ray. Hence, while it is possible to determine the value of the

scanning mirrors angles given a three dimensional point, the mathematical relationships are not closed since a given set of mirror angles does not relate to a unique three dimensional point. This is an inherent property of the laser scanner since the mirrors control the direction of the laser beam, and there are an infinite number of three dimensional points that are colinear with this beam. Closed form solutions are only possible if one of the axes (usually the Z axis) of the three dimensional point is fixed.

2.2 Calibration of the Intrinsic Parameters

The distance between the two mirrors (e) can be obtained by direct measurement. The degree to which errors in this measurement will affect the accuracy of the scanner depends on the environment in which it will be used. If the laser scanner is situated at a large z distance from the workspace, then the effects of error in the measurement of e will be reduced. This is the case in the CIRSSE Testbed, since the distance between the scanning mirrors is 5 mm and the z distance to the workspace is typically 2000 mm. If the laser scanner is used in situations where z is small, then the distance between the scanning mirrors should be determined analytically. This can be done in conjunction with determining the laser's extrinsic parameters (see Section 2.3).

Calibrating the scanning mirrors is critical to proper operation of the laser scanner. Each mirror is rotated with a galvanometer, which transforms a control voltage into an angular rotation of the mirror. Zero volts is assumed to correspond to a mirror angle of zero degrees (e.g., the laser beam is assumed to coincide with the laser's z axis when both galvanometers have zero input). The θ_y mirror is calibrated by directing the beam onto a plane at a fixed z distance from the laser with respect to the laser's origin (L). With θ_x fixed at zero, the θ_y mirror is rotated with a fixed voltage and the amount of y displacement on the plane is recorded. Using this information and (2.1) it is possible to determine θ_y. Assuming the relationship between voltage and mirror rotation is linear, mirror rotation is determined by dividing θ_y by the voltage applied to the mirror. To confirm the linearity of the galvanometers, the y mirror should be displaced to several different positions and the relationship should be verified to not change within measurement error. The θ_x mirror is calibrated in the same manner except that θ_y is fixed at zero (so there will be no pincushion distortion) and (2.2) is used to determine θ_x.

2.3 Calibration of the Extrinsic Parameters using an LSE Method

The extrinsic parameters of a laser scanner can be obtained using an LSE method as follows. Direct the laser at a set of 3-D points and record the scanning mirror angles at each point. These points and associated mirror angles can then be used to solve an overdetermined system of linear equations to obtain the laser's extrinsic parameters. The method described in this section is analogous to an LSE approach proposed by Roger Tsai[5] for the calibration of cameras, except that the terms of the linear equations are different for a laser.

The extrinsic parameters of the laser scanner consist of the rotation and translation of the laser coordinate frame with respect to some other fixed coordinate frame. This rotation and translation should ultimately be represented as a 4×4 homogeneous transform (of the type defined by Craig[6] of the form

$$ {}^a_bT = \begin{bmatrix} r_1 & r_2 & r_3 & t_x \\ r_4 & r_5 & r_6 & t_y \\ r_7 & r_8 & r_9 & t_z \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{2.3} $$

The homogeneous transform ${}^a_bT$ is primarily composed of a 3×3 rotation matrix and a 3×1 translation vector which define the orientation and position of frame b with respect to frame a. To calibrate the laser we need to find ${}^l_wT$, which is the transformation from the laser coordinate frame to a desired world coordinate frame. What is required is a mathematical relationship that will determine these parameters given a set of points defined in the world coordinate frame and a set of corresponding scanning mirror angles. A point $\vec{P}_w$ defined in the world coordinate frame is transformed to the laser coordinate frame using (2.4) to produce $\vec{P}_l$ ($\vec{P}_l$ and $\vec{P}_w$ are 3×1 vectors):

$$ \vec{P}_l = \begin{bmatrix} x_l \\ y_l \\ z_l \end{bmatrix} = {}^l_wT\vec{P}_w = {}^l_wT \begin{bmatrix} x_w \\ y_w \\ z_w \end{bmatrix} = \begin{bmatrix} r_1x_w + r_2y_w + r_3z_w + t_x \\ r_4x_w + r_5y_w + r_6z_w + t_y \\ r_7x_w + r_8y_w + r_9z_w + t_z \end{bmatrix} \tag{2.4} $$

Also, x_l and y_l can be expressed in terms of the scanning mirrors using (2.1) and (2.2) with (2.4) as shown below:

$$ (z_l \sec\theta_y + e)\tan\theta_x = r_1x_w + r_2y_w + r_3z_w + t_x \tag{2.5} $$

$$ z_l \tan\theta_y = r_4x_w + r_5y_w + r_6z_w + t_y \tag{2.6} $$

Substituting the expression for z_w in (2.4) into (2.5) and (2.6) results in:

$$\begin{aligned}((r_7x_w + r_8y_w + r_9z_w + t_z)\sec\theta_y + e)\tan\theta_x &= \\ r_1x_w + r_2y_w + r_3z_w + t_x & \qquad (2.7)\\ (r_7x_w + r_8y_w + r_9z_w + t_z)\tan\theta_y &= \\ r_4x_w + r_5y_w + r_6z_w + t_y & \qquad (2.8)\end{aligned}$$

Simplifying (2.7) and (2.8) and then dividing through by t_z ($t_z \neq 0$):

$$\begin{aligned}&\frac{\tan\theta_x}{\cos\theta_y}x_w\frac{r_7}{t_z} + \frac{\tan\theta_x}{\cos\theta_y}y_w\frac{r_8}{t_z} + \\ &\frac{\tan\theta_x}{\cos\theta_y}z_w\frac{r_9}{t_z} + \frac{\tan\theta_x}{\cos\theta_y} + \tan\theta_x\frac{e}{t_z} = x_w\frac{r_1}{t_z} + y_w\frac{r_2}{t_z} + z_w\frac{r_3}{t_z} + \frac{t_x}{t_z} \qquad (2.9)\\ &x_w\tan\theta_y\frac{r_7}{t_z} + y_w\tan\theta_y\frac{r_8}{t_z} + \\ &z_w\tan\theta_y\frac{r_9}{t_z} + \tan\theta_y = x_w\frac{r1}{t_z} + y_w\frac{r_2}{t_z} + z_w\frac{r_3}{t_z} + \frac{t_y}{t_z} \qquad (2.10)\end{aligned}$$

Equations (2.9) and (2.10) can be expressed in the form $Ax = b$, where A is $2n \times 12$, b is $2n \times 1$, n is the number of data points collected, and x is a 12×1 vector of unknowns. Note that this system of equations not only determines the laser's extrinsic parameters, but also the distance between the scanning mirrors (e). The final form (for $t_z \neq 0$) is presented below:

$$\begin{bmatrix} \tan\theta_{x_1} & -x_{w_1} & -y_{w_1} & -z_{w_1} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -x_{w_1} & -y_{w_1} & -z_{w_1} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \tan\theta_{x_n} & -x_{w_n} & -y_{w_n} & -z_{w_n} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -x_{w_n} & -y_{w_n} & -z_{w_n} \end{bmatrix}$$

$$\begin{bmatrix} \frac{\tan\theta_{x_1}}{\cos\theta_{y_1}} x_{w_1} & \frac{\tan\theta_{x_1}}{\cos\theta_{y_1}} y_{w_1} & \frac{\tan\theta_{x_1}}{\cos\theta_{y_1}} z_{w_1} & -1 & 0 \\ x_{w_1}\tan\theta_{y_1} & y_{w_1}\tan\theta_{y_1} & z_{w_1}\tan\theta_{y_1} & 0 & -1 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \frac{\tan\theta_{x_n}}{\cos\theta_{y_n}} x_{w_n} & \frac{\tan\theta_{x_n}}{\cos\theta_{y_n}} y_{w_n} & \frac{\tan\theta_{x_n}}{\cos\theta_{y_n}} z_{w_n} & -1 & 0 \\ x_{w_n}\tan\theta_{y_n} & y_{w_n}\tan\theta_{y_n} & z_{w_n}\tan\theta_{y_n} & 0 & -1 \end{bmatrix} - \begin{bmatrix} e \\ r_1 \\ r_2 \\ r_3 \\ r_4 \\ r_5 \\ r_6 \\ r_7 \\ r_8 \\ r_9 \\ t_x \\ t_y \end{bmatrix} \frac{1}{t_z} = \begin{bmatrix} \frac{\tan\theta_{y_1}}{\cos\theta_{x_1}} \\ \tan\theta_{y_1} \\ \vdots \\ \frac{\tan\theta_{y_n}}{\cos\theta_{x_n}} \\ \tan\theta_{y_n} \end{bmatrix} \tag{2.11}$$

After solving (2.11) using singular value decomposition [7], $|t_z|$ can be found from

$$|t_z| = \frac{1}{\sqrt{\left(\frac{r_7}{t_z}\right)^2 + \left(\frac{r_8}{t_z}\right)^2 + \left(\frac{r_9}{t_z}\right)^2}} \tag{2.12}$$

Once $|t_z|$ is determined, it is a straightforward process to obtain e, r_1, through r_9, t_x and t_y. With several hundred data points, this method produces reliable results provided that the data points are measured accurately (i.e., measurement error is $\leq$ 1 mm). The major flaw in this calibration method is that it treats the twelve parameters as being independent (which is obviously incorrect) and thereby fails to meet the constraints inherent in the rotation matrix.

2.4 Direct Geometric Method for Calibrating a Laser Scanner

The LSE approach described in the previous section requires a large number of accurate data points to generate an accurate solution. A solution typically occurs when the values generated by the LSE method are stable as the number of data points increases. Usually, two to three hundred points are required to generate a stable solution. There are situations where it is impractical to collect a large number of data points. However, it is possible to measure the laser scanner's extrinsic parameters directly since the laser emits a beam of light that can be measured with respect to a reference point. The method presented in this section treats laser scanner calibration as a

geometrical problem in which the laser scanner's Euler angles and location with respect to the world coordinate frame are measured directly. It should be emphasized that the objective of this calibration method is identical to that of the LSE method: to determine the homogeneous transform from the laser coordinate frame to some world coordinate frame.

The laser coordinate frame (L) is located at the center of the θ_y mirror and is oriented as in Figure 2.1. The pose of the world coordinate frame (W) is arbitrary. To assist in calibrating the laser, an intermediate coordinate frame F is defined whose origin is located directly below the center of the laser scanner's aperture. The frame F is located by suspending a plumb line from the laser scanner to some fixed plane (the floor is used in the CIRSSE system). This plumb line constitutes the z axis of the F coordinate frame. Note that while F's origin is on the floor, F's x and y axes are not necessarily coplanar with the floor. The location of F is selected in this manner to simplify the measurement of the laser scanner's orientation. The calibration procedure involves measuring twelve parameters, which are depicted in Figures 2.3 and 2.4 and defined as:

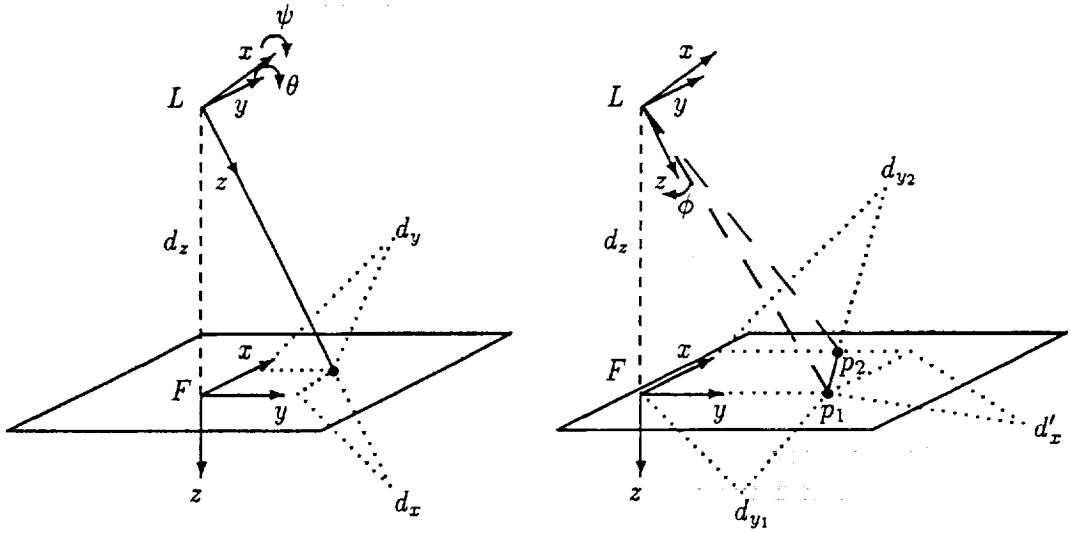

Figure 2.3: Determination of a Laser Scanner's Euler Angles. (Note: F is not necessarily coplanar with the floor).

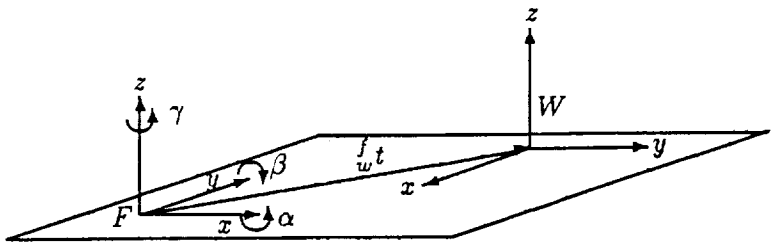

Figure 2.4: Transformation Parameters from Frame F to Frame W. (Note: F is not necessarily coplanar with the floor).

d_z The distance from F to L along F's z axis.

d_x The x coordinate of the point where the undeflected laser beam intersects the floor, measured with respect to the origin of F.

d_y The y coordinate of the point where the undeflected laser beam intersects the floor, measured with respect to the origin of F.

d_{y_1} The y coordinate of the point p_1 on the floor, measured with respect to the origin of F.

d_{y_2} The y coordinate of the point p_2 on the floor, measured with respect to the origin of F.

d'_x The length of the projection of the line segment joining p_1 and p_2 onto the vector formed by projecting the x axis of frame F onto the plane of the floor.

ρ The rotation about F's x axis from F's xy plane to the floor.

δ The rotation about F's y axis from F's xy plane to the floor.

${}^f_w t$ The translation vector from F to W.

α The pitch angle about F's x axis from frame F to frame W

β The yaw angle about F's y axis from frame F to frame W

γ The roll angle about F's z axis from frame F to frame W

The above twelve parameters provide all the information required to determine the transformation from the laser coordinate frame to the world coordinate frame: ${}^l_w T$. The calculations are broken into two steps:

1. Determining ${}^{l}_{f}T$

2. Determining ${}^{f}_{w}T$

Once these transforms are known, it will be possible to determine ${}^{l}_{w}T$ from

$$ {}^{l}_{w}T = {}^{l}_{f}T \cdot {}^{f}_{w}T \tag{2.13} $$

${}^{l}_{f}R$ can be determined by deriving the orientation of frame F with respect to the laser in terms of the Euler angles pitch, yaw and roll (ψ,θ,ϕ) about the laser's x, y, and z axes, respectively. Figure 2.3 shows how the Euler angles are measured independently of one another. The laser's scanning mirrors are set to zero and the coordinates of the point where the laser beam intersects the floor are measured (d_x, d_y). It is assumed that any offset of the scanning mirrors from zero is negligible compared to the magnitude of the Euler angles being measured. Since the mirrors are assumed to be in their undeflected state, the beam is coincident with the laser's z axis, and the point (d_x, d_y) is invariant to rotations about L's z axis. We will therefore assume that the laser coordinate frame is rotated about z by an amount ϕ (to be determined) so that the projection of the laser's y axis onto F's xy plane is coincident with F's y axis. Such an orientation implies that d_x is only a function of θ and d_y is only a function of ψ.

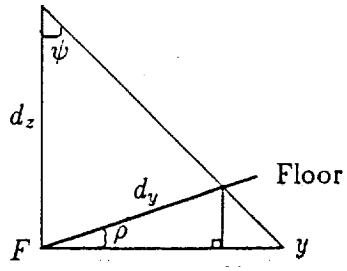

Figure 2.5: Relationship Between the Floor and F's xy Plane

When calculating ψ and θ, it cannot be assumed that the floor correctly defines the xy plane of F. Indeed, since d_z was measured with a plumb bob, F's z axis is aligned with earth's local gravity vector, but the floor may not be orthogonal to this vector. By using a level it is possible to determine the angles ρ and δ about F's x and y axes, respectively, between the plane of

the floor and the xy plane of F. The situation is depicted in Figure 2.5 for the Euler angle ψ. The formulas for determining ψ and θ can be derived directly from the figure and are presented below:

$$\psi = \tan^{-1}\left[\frac{d_y \cos\rho + d_y \sin\rho \tan\rho}{d_z}\right] \tag{2.14}$$

$$\theta = \tan^{-1}\left[\frac{d_x \cos\delta + d_x \sin\delta \tan\delta}{d_z}\right] \tag{2.15}$$

At this point it is necessary to determine ϕ so that the laser's y axis can be aligned with F's y axis. The θ_x mirror is repeatedly rotated by an arbitrary amount while the θ_y mirror is set to zero, resulting in a line segment traced on the floor. The slope of this line determines ϕ and is found by measuring two arbitrary points (p_1 and p_2) on the line segment (this yields the parameters d_{y_1}, d_{y_2} and d'_x).

The primary issue with determining ϕ in this manner is ensuring the slope of the line segment to be a function only of ϕ and not of other variables, such as ψ, θ, or θ_y (which causes pincushion distortion). Because we have assumed $\theta_y = 0$, the slope of the line segment will not be altered by pincushion distortion. The Euler angle θ affects the laser's x and z components of the beam direction. A change in the z direction will not affect the slope of the line since the line ultimately will lie in F's xy plane. Further, distortion in x will tend to move the line segment by some constant value, leaving the slope unchanged. The Euler angle ψ affects the laser's y and z components of the beam direction. The distortion in z will not affect the measurements for the same reasons stated for θ. The distortion in y due to ψ consists of a constant translation of the line segment along F's y axis. The slope of the line is a function of the relative change in y from points p_1 to p_2, so the effects of ψ will not change the slope since p_1 and p_2 will be translated in y by the same amount. Hence, the slope of the line segment is a function only of ϕ.

Since the measurements used to determine ϕ were taken from the floor, they will have to be corrected for the effects of ρ and δ for the same reasons that the corrections were necessary for determining ψ and θ. The final equation for ϕ is presented below:

$$\phi = -tan^{-1}\left[\frac{(d_{y_2} \cos\rho + d_{y_2} \sin\rho \tan\rho) - (d_{y_2} \cos\rho + d_{y_2} \sin\rho \tan\rho)}{d'_x \cos\delta + d'_x \sin\delta \tan\delta}\right] \tag{2.16}$$

At this point the Euler angles ψ, θ, ϕ have been determined, and it is now possible to create the rotation matrix ${}^{l}_{f}R$ from the laser to the F coordinate frame. Each Euler angle changes four parameters in the rotational matrix. The most straightforward approach is to determine separate rotational matrices for each Euler angle and then multiply them together to obtain ${}^{l}_{f}R$. The individual matrices for R(X,ψ), R(Y,θ) and R(Z,ϕ) are:

$$R(X,\psi) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\psi & -\sin\psi \\ 0 & \sin\psi & \cos\psi \end{bmatrix} \tag{2.17}$$

$$R(Y,\theta) = \begin{bmatrix} \cos\theta & 0 & \sin\theta \\ 0 & 1 & 0 \\ -\sin\theta & 0 & \cos\theta \end{bmatrix} \tag{2.18}$$

$$R(Z,\phi) = \begin{bmatrix} \cos\phi & -\sin\phi & 0 \\ \sin\phi & \cos\phi & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{2.19}$$

$${}^{l}_{f}R(\phi,\theta,\psi) = R(X,\psi)R(Y,\theta)R(Z,\phi) \tag{2.20}$$

As is apparent from (2.20), ${}^{l}_{f}R$ is obtained by combining the rotation matrices in the order roll, yaw, pitch. This ordering is essential for this calibration procedure. Applying the roll rotation first will align the projections of the laser's y axis to the F coordinate frame's y axis. This condition was assumed when the Euler angles ψ and θ were determined.

At this point, ${}^{l}_{f}R$ has been determined, but to obtain ${}^{l}_{f}T$ it is necessary to determine the translational component ${}^{l}_{f}t$ from the laser's origin to F with respect to the laser. Given the configuration of the CIRSSE testbed and the available measuring equipment, it is difficult to directly measure this value with any degree of accuracy. However, it is not necessary to directly measure this value. The translation vector ${}^{f}_{l}t$ from F to the laser can be determined since F is located directly below the aperture of the laser, and d_z is known. The resulting value of ${}^{f}_{l}t$ is $[0, 0, d_z]^T$.

At this point ${}^{f}_{l}T$ can be determined by combining the inverse of ${}^{l}_{f}R$ (for a rotation matrix, $R^{-1} = R^T$) with the translation vector ${}^{f}_{l}t$. It is then possible to obtain ${}^{l}_{f}T$ by taking the inverse of ${}^{f}_{l}T$. This relationship is

defined as

$$ {}^{l}_{f}T = \begin{bmatrix} & {}^{l}_{f}R^{T} & & {}^{f}_{l}t \\ 0 & 0 & 0 & 1 \end{bmatrix}^{-1} \tag{2.21} $$

With ${}^{l}_{f}T$ defined, the next step is to determine ${}^{f}_{w}T$. W should be chosen so that the Euler angles α, β, and γ can be easily measured. If W is chosen such that its z axis is plumb (as in the case with F), then the Euler angles from F to W can be readily determined. By applying (2.17) through (2.20), it is possible to determine ${}^{f}_{w}R$, and this matrix can be combined with ${}^{f}_{w}t$ (which is one of the twelve calibration parameters) to obtain ${}^{f}_{w}T$. At this point, ${}^{l}_{w}T$ can be determined by matrix multiplication:

$$ {}^{l}_{w}T = {}^{l}_{f}T \cdot {}^{f}_{w}T \tag{2.22} $$

This concludes the calibration of the laser's extrinsic parameters. It is now possible to transform points defined in the world coordinate frame into points defined in the laser coordinate frame. Additionally, points in the laser frame can be transformed into the world frame using the inverse of ${}^{l}_{w}T$.

2.5 Appraisal of Calibration Method Performance

When presented with two alternatives for calibrating the laser's extrinsic parameters, the question arises as to which method is best. The answer depends on how and where the laser scanner will be used. The transform obtained using the direct geometrical method typically results in errors of less than 0.5%. This accuracy has consistently been obtained in the CIRSSE testbed where this method is currently implemented in software. Typical values for the calibration parameters are presented in table 2.1 This level of accuracy is sufficient for many visual sensing tasks. Further, the geometric approach achieves its results in a simple systematic manner. Hence, calibrating the laser using this method requires less effort than the LSE method.

Since the LSE method generates a solution based on a large set of data points, measurement errors among individual points should have less affect on the LSE solution. This is in contrast to the geometric method which uses a small number of measurements to obtain its results, and, hence, these few

Parameter	Value
d_x	5.0 mm
d_y	101.1 mm
d_z	2597.0 mm
ρ	0.0 radians
δ	0.0 radians
d_{y_1}	-10.0 mm
d_{y_2}	0.0 mm
dx'	946.0 mm
α	0.0 radians
β	π radians
γ	0.0 radians
${}^f_w t_x$	1409.9 mm
${}^f_w t_y$	887.65 mm
${}^f_w t_z$	0.0 mm
e	5.0 mm

Table 2.1: Typical Calibration Parameter Values for Direct Geometric Method

points must be more accurately measured. The LSE method has been simulated in software and tested with simulated sets of data points containing differing degrees of error. The results of these simulations indicate that *if* it is possible to collect a large number of points with high accuracy, the LSE could produce more accurate results than the direct geometric method. At the time this research was conducted, the CIRSSE testbed had no means to collect a large number of highly accurate data points, but the LSE method has the potential to be highly effective when the testbed acquires the necessary data collection capability.

Chapter 3

Point Estimation with a Calibrated Laser and Camera

Once the laser scanner is calibrated to a world coordinate frame it is possible to use it in concert with a calibrated camera to perform three dimensional sensing. Three dimensional sensing with a camera and a laser is different from strictly passive methods such as dual cameras. In a dual camera system, features identified in one camera image are correlated to similar features in the other camera image. The pixel coordinates of these features are then used to determine the three dimensional point of the object corresponding to the feature in the images.

Three dimensional sensing with a laser and camera, however, employs a slightly different approach. The laser directs its beam into the field of view of the camera and the camera image is scanned for the reflection of the laser beam off of an object in the image. The pixel coordinates of the laser beam's reflection (usually referred to as the laser spot) are correlated to the mirror angles of the laser scanner to obtain the three dimensional point of the object in the workspace.

Methods for estimating three dimensional points using dual cameras have been developed by Repko, Sood, and Kelly [8] and Noseworthy[9]. One method solves an overdetermined set of equations to obtain a lease squared estimate of the three dimensional point, while a second method calculates two three dimensional rays projecting from the image planes of the cameras,

and estimates the coordinates of the corresponding point by determining the midpoint of the common normal of the rays.

These two methods for point estimation can be readily adapted for use with a calibrated camera and laser. In order to do this, the mathematical relationships contributed by one camera are substituted with mathematical relationships for the laser expressed in the same form as the camera. Hence, in order to prove that these point estimation methods will work with a laser and a camera, it is only necessary to show that the laser's mathematical model can be expressed in such a way as to be compatible with each method.

3.1 Point-Estimation Using Least Squared Error

The overdetermined system of linear equations approach for dual cameras[9] can readily be adapted to a camera and a laser scanner. What is needed is a relationship for the laser between the scanning mirror angles and the coordinates of a three dimensional point ($\vec{P}_w$) defined in the world frame. Such a relationship was derived in Section 2.3 as part of the LSE method for calibrating the laser (see equations (2.7) and (2.8)). Ultimately, the equations contributed by the laser will be included in a system of linear equations of the form $Ax = b$ where x is a three row by one column vector representing the x, y, and z coordinates of the estimated three dimensional point.

In the context of laser calibration, the unknown variables were the distance between the scanning mirrors (e) and the rotation and translation components of the transformation from the laser coordinate frame to the world coordinate frame. Since the laser is assumed to be calibrated at this point, all these values are known. Additionally, the angles of the scanning mirrors are known. What is not known are the x, y, and z coordinates of the world point. The terms in (2.7) and (2.8) can be regrouped into a form that is more suitable for point estimation:

$$\left(\frac{\tan\theta_x}{\cos\theta_y}r_7 - r_1\right)x_w + \left(\frac{\tan\theta_x}{\cos\theta_y}r_8 - r_2\right)y_w + \left(\frac{\tan\theta_x}{\cos\theta_y}r_9 - r_3\right)z_w = t_x - t_z\frac{\tan\theta_x}{\cos\theta_y} - e\tan\theta_x \quad (3.1)$$

$$(r_7\tan\theta_y - r_4)x_w + (r_8\tan\theta_y - r_5)y_w + (r_9\tan\theta_y - r_6)z_w = t_y - t_z\tan\theta_y \quad (3.2)$$

Therefore, the laser contributes two equations with the same three unknowns as the camera equations (assuming that the laser and camera are calibrated to the same world coordinate frame). These two equations can be combine with the two equations contributed by the camera[9]. For sake of brevity, the full derivation of the camera's equations will not be given here (see Noseworthy(1991)[9]), but the final result is given by equations (3.3) and (3.4).

$$x_w(r_7x_u - fr_1) + y_w(r_8x_u - fr_3) + z_w(r_9x_u - fr_3) = ft_x - t_zx_u \quad (3.3)$$

$$x_w(r_7y_u - fr_4) + y_w(r_8y_u - fr_5) + z_w(r_9y_u - fr_6) = ft_y - t_zy_u \quad (3.4)$$

Where x_u and y_u are camera pixel coordinates, f is the focal length of the camera, and r_1 - r_9, t_x, t_y and t_z are elements of the homogeneous transformation ${}^c_wT$ from the camera's coordinate space to the world coordinate space. Equations (3.3), (3.4), (3.1), and (3.2) can be expressed as a system of linear equations of the form $Ax = b$:

$$\begin{bmatrix} (r_7x_u - fr_1) & (r_8x_u - fr_3) & (r_9x_u - fr_3) \\ (r_7y_u - rr_4) & (r_8y_u - fr_5) & (r_9y_u - fr_6) \\ \left(\frac{\tan\theta_x}{\cos\theta_y}r_7 - r_1\right) & \left(\frac{\tan\theta_x}{\cos\theta_y}r_8 - r_2\right) & \left(\frac{\tan\theta_x}{\cos\theta_y}r_9 - r_3\right) \\ (r_7\tan\theta_y - r_4) & (r_8\tan\theta_y - r_5) & (r_9\tan\theta_y - r_6) \end{bmatrix} \cdot \begin{bmatrix} x_w \\ y_w \\ z_w \end{bmatrix} = \begin{bmatrix} (ft_x - t_zx_u) \\ (ft_y - t_zy_u) \\ \left(t_x - t_z\frac{\tan\theta_x}{\cos\theta_y} - e\tan\theta_x\right) \\ (t_y - t_z\tan\theta_y) \end{bmatrix} \quad (3.5)$$

The system of equations presented in (3.5) can be solved using singular value decomposition[7]. Care must be taken, however, in interpreting the results of the LSE solution[9]. The LSE approach solves for x by minimizing the expression: $(b - Ax)^T(b - Ax)$ where $x = \left(A^TA\right)^{-1} \cdot A^Tb$ where the quantity $\left(A^TA\right)^{-1}$ is the pseudo-inverse of A. Minimizing $(b - Ax)^T(b - Ax)$ does not mean that the error between the actual point

and the estimated point has been minimized. Since this is "implicit" estimation, the error is minimized, and by so doing, it is assumed that the parameter of interest (x) is optimized in the process.

3.2 Point Estimation Using Midpoint of Common Normal

Another method for estimating three dimensional points was developed by Noseworthy[9]. This method calculates a ray projecting from the camera's image plane into the three dimensional environment. A brief summary of this work is presented below followed by a description of how this method can be adapted for use with a laser

A three dimensional point is calculated by determining the midpoint to the common normal of the rays calculated from two different cameras. the ray for each camera is expressed as a linear parametric equation of the form:

$$\vec{r}_c = s_c \vec{d}_c + \vec{O}_c \tag{3.6}$$

Where $\vec{O}_c$ is the 3×1 vector describing the location of the origin of the camera with respect to the world coordinate system, and $\vec{d}_c$ is a 3×1 vector describing the direction of the ray projected from the camera's image plane with respect to the world coordinate system. These two terms can be further expressed as:

$$\vec{d}_c = {}^w_c R \begin{bmatrix} x_u \\ y_u \\ f_c \end{bmatrix} s_c \tag{3.7}$$

$$\vec{O}_c = {}^w_c R {}^w_c t \tag{3.8}$$

${}^w_l R$ and ${}^w_l t$ can be obtained from the inverse of the homogeneous transform ${}^c_w T$, and x_u and y_u are camera pixel coordinates. Once these parametric equations have been determined it is possible to calculate the unit direction, in world coordinates, of the common normal to the rays from two cameras as:

$$\hat{n} = \frac{\vec{d}_{c_1} \times \vec{d}_{c_2}}{| \vec{d}_{c_1} \times \vec{d}_{c_2} |}$$
$$| \vec{d}_{c_1} \times \vec{d}_{c_2} | \neq 0 \tag{3.9}$$

The shortest distance between $\vec{r}_{c_1}$ and $\vec{r}_{c_2}$, l, can be determined by projecting $\vec{O}_{c_1} - \vec{O}_{c_2}$ in the $\hat{n}$ direction. Mr. Noseworthy points out that $\vec{r}_{c_1}$ and $\vec{r}_{c_2}$ are assumed to be skew (i.e. $| \vec{d}_{c_1} \times \vec{d}_{c_2} | \neq 0$). l is determined using the following expression:

$$l = \frac{\left(\vec{O}_{c_1} - \vec{O}_{c_2}\right) \cdot \left[\vec{d}_{c_1} \times \vec{d}_{c_2}\right]}{| \vec{d}_{c_1} \times \vec{d}_{c_2} |} \tag{3.10}$$

Finally, a 3×1 vector, $\vec{m}$, representing the coordinates of the midpoint to the common normal of $\vec{r}_{c_1}$ and $\vec{r}_c c_2$ is determined by:

$$\vec{m} = \left(s_1 \vec{d}_{c_1} + \vec{O}_{c_1}\right) - \left(s_1 \vec{d}_{c_1} + \vec{O}_{c_1}\right) = l\hat{n} \tag{3.11}$$

The midpoint to the common normal method can also be used with a calibrated camera and laser scanner. To do this, it is necessary to derive a parametric equation for the 3-D ray of the laser beam. This equation is expressed as

$$\vec{r}_l = s_l \vec{d}_l + \vec{O}_l \tag{3.12}$$

where $\vec{O}_l$ is the origin of the laser beam in terms of the world coordinate frame, $\vec{d}_l$ is the direction vector of the laser ray and s_l is a parameter.

Equation (3.12) is of the same form as (3.6). To use the midpoint to the common normal method with a calibrated laser, it is necessary to derive expressions for $\vec{O}_l$, $\vec{d}_l$ and s_l. Once these values are determined, the mathematical relationships for the midpoint of the common normal for two cameras will also work for a calibrated camera and laser.

$\vec{O}_l$ can be determined based on the values of the intrinsic and extrinsic parameters of the laser scanner. Specifically, the required parameters are the distance between the scanning mirrors, and the transformation ${}^l_w T$ from the laser coordinate frame to the world coordinate frame. The laser calibration assumes the laser's origin to be at the center of the θ_y mirror. The coordinates of the laser's origin with respect to the world coordinate frame

can be derived by extracting the translational component of the inverse of the ${}^{l}_{w}T$. This translation vector is defined as ${}^{w}_{l}t$.

The origin of the laser has one additional component. Recalling the arrangement of the laser scanner described in Section 2.1, the θ_x mirror deflects the laser beam along the rotational axis of the θ_y mirror. Since the origin of the laser scanner is defined to be on the θ_y mirror, the origin of the laser ray is translated along the x axis of the laser scanner (this is the rotational axis of the θ_y mirror) by the rotation of the θ_x mirror. The translation of the laser's origin as a function of θ_x with respect to the laser's coordinate frame can be expressed as a 3×1 vector:

$$\vec{t}_{\theta_x} = \begin{bmatrix} e \tan \theta_x \\ 0 \\ 0 \end{bmatrix} \tag{3.13}$$

The term $e \tan \theta_x$ from (2.2) defines the x coordinate of the laser beam with respect to the laser's coordinate frame given a set of mirror angles and a specific z coordinate. The vector $\vec{t}_{\theta_x}$ is defined in terms of the laser coordinate frame; hence it must be transformed into a translation with respect to the world coordinate frame. This is accomplished by multiplying $\vec{t}_{\theta_x}$ by the rotation matrix ${}^{w}_{l}R$ contained in the inverse of the homogeneous transform ${}^{l}_{w}T$. This yields a new translation vector defined in the world coordinate frame:

$$\vec{t_{w_{\theta_x}}} = {}^{w}_{l}R \cdot \vec{t_{\theta_x}} \tag{3.14}$$

Therefore, the final value for $\vec{O}_l$ can be expressed as

$$\vec{O}_l = {}^{w}_{l}t + \vec{t}_{w_{\theta_x}} \tag{3.15}$$

The next step is to determine $\vec{d}_l$ from (2.2) and (2.1). Since the expression for $\vec{O}_l$ already compensates for the translation of the laser's origin due to rotation of the θ_x mirror, the $e \tan \theta_x$ term in (2.2) can be removed. Hence, the direction of the laser ray with respect to the laser coordinate frame can be expressed as a 3×1 vector of the form

$$\vec{d}_l = \begin{bmatrix} z \sec \theta_y \tan \theta_x \\ z \tan \theta_y \\ z \end{bmatrix} \tag{3.16}$$

The direction vector is currently defined with respect to the laser coordinate frame. In order to use it in determining the midpoint of the common normal, it will have to be transformed with respect to the world coordinate frame. This can be done simply by multiplying $\vec{d_l}$ by the rotation matrix ${}^{l}_{w}R$. Further, the z term in $\vec{d_l}$ can be factored out and used as the variable parameter s_l. Therefore, the parametric equation for the 3-D laser ray can be expressed as

$$\vec{r_l} = {}^{w}_{l}R\vec{d_l}z \cdot + \vec{O_l} \tag{3.17}$$

Equation (3.17) can be used in place of the parametric equation for the second camera to determine the midpoint of the common normal for a camera and a laser scanner. The remaining mathematical expressions for the midpoint of the common normal calculation remain valid.

3.3 Appraisal of Point Estimation Methods

As mentioned previously, the LSE point estimation method generates a solution by attempting to optimize the value of x. Because this method attempts to determine an optimized solution, it can accommodate minor errors in the camera and laser calibration parameters. This property of the LSE method can be useful in that minor calibration errors will not necessarily result in poor point estimates.

The midpoint to the common normal method attempts to model the exact behavior of the laser and camera. Specifically, it projects rays from the laser and the camera into space based on the mirror angles of the laser scanner and the pixel coordinates on the camera's image plane. If all the calibration parameters for the laser and camera are perfectly accurate, the rays should intersect, but in reality, the rays do not intersect due to errors in calibration, and hence, the midpoint of the common normal to these two rays is used as the estimate of the three dimensional point.

The decision as to which method to use for point estimation depends on the nature of the application in which the method will be used. Since the LSE method is more robust in terms of accommodating calibration errors it may be useful in situations where the accuracy of calibration parameters is questionable. However, it is not clear at what point calibration errors will severely affect the performance of the LSE method.

While the midpoint to the common normal method is less tolerant of calibration errors, it does have one significant quality. Since this method is

a direct representation of the geometry of the point estimation scenario, it may be possible to use this method as a means of predicting how calibration errors will affect performance. If performance can be predicted ahead of time it might be possible to modify the point estimates to account for calibration error. Future research in this area will be necessary to better understand the appropriate application of these two point estimation methods.

Chapter 4

Locating a Laser Spot in a Camera Image

To use a calibrated laser scanner in concert with a camera, it is necessary to be able to locate the laser spot in the camera image. This is a simple problem if one can guarantee that the laser spot is the brightest region in the image. However, such an assumption restricts the utility of a calibrated laser by placing illumination constraints on the image. If techniques can be employed to locate the laser spot in the presence of "noise" (e.g. pixels of similar intensity), then a calibrated laser can be used in a wider variety of situations. The method developed by the author to locate the laser spot in a noisy image is a heuristic approach whereby regions in the image are successively eliminated based on a set of criterion tests.

The first step in locating the laser spot is to perform region growing over some selected area of the image. This results in a list of regions, their area, and their centroids. The laser spot should be one of the regions in this list. To isolate the laser spot it is necessary to eliminate all those regions that are not attributable to the laser. There are four different tests that can be applied to the region list to perform this elimination. Each test returns a list of the regions that passed the test. The laser spot should be the only region that passes all four tests.

The first test is to eliminate all regions that do not fall within a specified intensity range. Since the laser will appear as a small bright spot, it will be one of the brighter regions in the image. However, there is no guarantee that the laser spot will be among the brightest. Indeed, if the laser beam is illuminating a matte (low reflectivity) object, such as a piece of cloth, then

the intensity of the laser spot will be lower than if the beam was reflecting off a piece of metal. Additionally, specular reflections of ambient light off of high reflectivity objects can exhibit the same intensity as the laser spot.

A second test that can be applied to the region list is to eliminate regions that do not fall within a certain range of sizes. The laser spot typically occupies between two pixels and twenty pixels depending on the reflectance of the object the laser beam is striking. The tests for size and intensity can detect the presence of a laser spot in the image in most cases. Problems arise when there are other regions in the image that have the same size and intensity characteristics as the laser, such as specular reflections.

If after application of the intensity and size tests the region list still has more than one candidate region, two more tests can be applied to further reduce the list. One of these tests is to take the centroid coordinates of each region and the known scanning mirror angles, run them through a point estimation algorithm, and eliminate the regions that generate solutions that fall outside of the workspace. This method does assume that some *a priori* knowledge exists about the expected location of the laser spot. The more that is known about the expected location of the laser spot, the greater the chance of correctly identifying it. Usually, little *a priori* knowledge is required to locate the laser spot, since only those regions that lie along the laser ray will generate results that are reasonable.

A final test can be employed if all the previous tests have failed to return a unique solution for the laser spot. The laser beam can be moved and another image acquired. The new image is passed through the region growing algorithm just as the first image. If the scene is static, the only region that should have moved is the laser spot.

It is important to note that it is not necessary to use all four tests. If a subset of tests yields one region, then the remaining tests do not have to be run. Further, there are situations where it may be impossible to locate the laser spot. If the laser spot is within the bounds of a bright region, the camera may not be able to distinguish it. This problem is particularly acute if the camera's aperture is too wide, since bright regions could then saturate the camera's CCD element. The laser spot is also undetectable if it is physically occluded by an object in the workspace. It should be noted that these four tests do not necessarily have to be performed in the order stated above. Indeed, part of the evaluations presented later in this chapter address the question of an optimal ordering for these tests

4.1 Application of Region Growing Algorithm to a Camera Image

Once a camera image is acquired, a region of interest is selected whose boundaries are such that it encompasses the laser region. The pixels within the region of interest are grouped into regions of similar intensity. The algorithm employed in this research is similar to the blob coloring algorithm proposed by Ballard and Brown [10]. The specific heuristic algorithm used in this application is described in figure 4.1.

A pixel is considered part of a unique region if its intensity is similar (by T_{region}) to its top, left, or top-left neighbors. The algorithm generates a list of regions identified in the image. Each entry in the list contains data on the regions size (in pixels), maximum intensity, centroid, and equivalence to another region in the list. The concept of region equivalency deserves more explanation. Envision performing this algorithm on an image that contained a region that is shaped like the letter "U". As the pixels are scanned the top portion of the "U" would be identified as two distinct vertical regions. At the point where the pixels form the curve at the bottom of the "U", the algorithm will find that the two regions its has been growing are actually the same region. In this case, an equivalency pointer in one region is set to the value of the other region.

The number of regions that are generated by this algorithm depends on the values of T_{back} and T_{region}. T_{back} essentially dictates how much of the image is eligible for region growing; while T_{region} determines how much contrast is required between pixels before a new region is detected. While the specific values for these variables depends on lighting conditions and image complexity, the values used in the CIRSSE testbed for room lighting conditions are typically $\{150 \leq T_{back} \leq 190\}$, $\{8 \leq T_{region} \leq 10\}$.

After the region list is constructed, it is assumed that any of the regions could be the laser spot. The next step after region growing is to eliminate from consideration all those regions that are equivalent to other regions (i.e. those regions that do not have the equivalency pointer equal to itself). At this point, the equivalency list contains a list of N unique regions.

for each pixel **from** left to right **and** top to bottom **do**:

if pixel intensity > background threshold T_{back}

if pixel directly above current pixel is part of a region **and** intensity of current pixel is similar to pixel directly above to within a given threshold T_{region}

- Mark current pixel as belonging to the same region as the pixel directly above it

else if pixel to left is part of a region **and** intensity of current pixel is similar to pixel directly above to within a given threshold T_{region}

- Mark current pixel as belonging to the same region as the pixel directly above it

else if pixel to the top left is part of a region **and** intensity of current pixel is similar to pixel directly above to within a given threshold T_{region}

- Mark current pixel as belonging to the same region as the top-left pixel

else

- Current pixel is part of a new region

else

pixel is part of the background

/* Check for region equivalence */

if {pixel to left of current pixel is part of a region} **and** {pixel above is part of a different region} and {the intensities of the two regions are similar to within T_{region}}

- Region to left of current pixel is equivalent to the region above the current pixel.

end loop

Figure 4.1: Region Growing Algorithm

4.2 Region Selection Based on Intensity and Size

Regions can be eliminated based on intensity and size criterion. This process is straightforward in that each region in the list is checked to see if it falls within a range of intensity values or within a range of size values depending on which criteria is being used. To ensure efficient performance, the size and intensity tests should be performed only on those regions that have not been eliminated as possible laser regions.

The execution time of the intensity algorithm can be expressed as:

$$\begin{aligned} \Gamma(intensity) &= C_{valid}(N) + C_{intensity}(n_{l1}) \\ & \{0 \leq n_{l1} \leq N\} \end{aligned} \tag{4.1}$$

Where C_{valid} represents the execution time required to determine if a region is a possible laser region and $C_{intensity}$ represents the time required to determine if the region's intensity falls within the specified limits. Since both of these operations consist of if/then comparisons the computation times for these operations can be expressed as constants. n_{l1} represents the number of regions in the list that have not been eliminated as laser regions.

Similarly, the execution time to eliminate regions by size can be expressed as:

$$\begin{aligned} \Gamma(size) &= C_{valid}(N) + C_{size}(n_{l2}) \\ & \{0 \leq n_{l2} \leq N\} \end{aligned} \tag{4.2}$$

Where C_{size} is the execution time to determine if a region's size falls within specified limits, and this value is also a constant. In practice, C_{size} and $C_{intensity}$ are approximately equal, and hence, so are $\Gamma(intensity)$ and $\Gamma(size)$. The execution time analysis for these two algorithms is trivial, yet, as will be discussed later, the performance of these algorithms is critical in determining the order of execution for all the selection criteria.

4.3 Region Selection Based on Laser/Camera Triangulation

Another method for determining the laser region is to apply the three dimensional point estimation algorithms specified in chapter 3 to each region

in the region list. Assuming the deflection of the laser scanner's mirrors is known, and the laser and camera are calibrated, the centroid data from each region can be used to estimate the three dimensional location of each region. If there is some knowledge as to the three dimensional location of the laser spot, the three dimensional point estimates can be successively eliminated until only those points that are consistent with the expected value remain.

The primary issue at hand is how much knowledge of the three dimensional location of the laser spot is needed to yield a unique solution. In practice, the laser region must lie along the projection of the laser ray across the camera's image plane. If the centroid of a region deviates from this projected line, then the point estimation algorithm will be trying to triangulate two divergent rays.

In practice, the three dimensional point estimates for regions other than the laser spot become highly irregular and minimal knowledge of the laser spot location is needed to reduce the set of point estimates to a unique solution. For example, in the CIRSSE testbed, the world origin is located about 10cm above the floor with the Z axis directed up at the ceiling. If a point estimate yields a Z of -80cm, this implies that the laser spot is located somewhere in or under the concrete floor of the testbed, and such a condition is clearly impossible.

The implementation of the triangulation algorithm currently used by the author employs the LSE point estimation algorithm described in section 3.1. Each region that has not been eliminated as a possible laser region is passed through the point estimation algorithm and the region is either eliminated or accepted if the estimated point lies within a specific three dimensional volume. Typically, this volume is centered about an estimated position of the laser spot and is constrained to +/- (5-10cm) in each axis about this position.

The execution time of this algorithm can be expressed as:

$$\begin{aligned} \Gamma(triangle) &= C_{valid}(N) + C_{triangle} n_{l3} \\ &\quad \{0 \leq n_{l3} \leq N\} \end{aligned} \tag{4.3}$$

Where n_{l3} is the number of regions in the region list that are possible laser regions and $C_{triangle}$ is a constant representing the computation time to estimate a three dimensional point for an arbitrary region. The $C_{triangle}$

for each region in current region list **do**

if current region X has not been eliminated as a candidate laser region

for ith region $\{0 \leq i \leq N\}$ in previous region list **do**

if distance between centroid of region X and region i $\leq$ $MaxDistance$ **and** difference between size of region X and size of region i $\leq$ $MaxSizeDiff$ **and** difference between intensity of region X and intensity of region i $\leq$ $MaxIntensityDiff$

- Eliminate region X as a possible laser region

Figure 4.2: Algorithm for elimination of regions based on movement

term deserves more explanation. The LSE point estimation algorithm uses singular value decomposition to calculate the estimated point. The execution time of the algorithm is dependent on the size of the A and b matrices. The size of these matrices is dependent on the number of sensing devices used to estimate the three dimensional point. In the case of one laser and a camera, the A matrix is four rows by three columns wide and the b matrix is four rows by one column. Since the number and type of sensing devices *should* not change during the middle of the triangulation algorithm, the dimensions of the A and b matrices will not change. Hence, the execution time for the singular value decomposition algorithm will be the same for each estimated point, and this value can be expressed as a constant. It is also important to mention that the value of $C_{triangle}$ is *much* larger than either $C_{intensity}$ or C_{size}. Hence, while all three algorithms execute in $O(N)$ time, the triangulation algorithm requires greater time to execute that the size or intensity algorithms.

4.4 Region Selection Based on Movement

Another method of determining the laser region is to acquire one image, move the laser, acquire a second image, and then eliminate all those regions that did not move. The algorithm used to determine if a region has moved is presented in figure 4.2

The algorithm in figure 4.2 will identify new regions in the current region list as regions that have moved. This is due to the fact that a new region in the current list cannot be correlated to a region in the previous list. This characteristic of the algorithm is neither a drawback nor an advantage as much as it is necessary to understand that the algorithm behaves in such a manner. The execution time of the algorithm can be expressed as:

$$\begin{aligned} \Gamma(movement) &= C_{valid} N_{current} + C_{movement} n_{l4} N_{previous} \\ &\quad \{0 \leq n_{l4} \leq N_{current}\} \end{aligned} \tag{4.4}$$

As one might expect, $N_{current}$ and $N_{previous}$ are the number of regions detected in the current image and the previous image respectively, and $C_{movement}$ is the execution time required to determine if a single region have moved from the previous frame. Assuming that $N_{current} \approx N_{previous}$ the algorithm in figure 4.2 executes as $O({N_{current}}^2)$ in the worst case.

4.5 Evaluation of Laser Region Identification Performance

A treatment of the issue of detecting a laser spot in a camera image would not be complete without a thorough evaluation of the performance of the algorithms under experimental conditions. The region selection algorithms were combined into a single program that directs the laser to specific three dimensional points and subsequently acquires images of the workspace for each point using a camera. This program was subjected to four different batteries of tests to determine the behavior of the laser spot selection algorithms to varying experimental conditions. The descriptions of the four test batteries are presented below:

- *Test Battery 1*: While maintaining constant lighting and region detection parameters, vary the complexity of the image by adding objects of differing size, and reflectance qualities. Low complexity images had few objects such cable, a few bits of metal and so forth, while more complex images contained everything in the low complexity images plus struts and unpainted metal nodes.

- *Test Battery 2*: While maintaining constant scene complexity and region detection parameters, vary the location and intensity of scene illumination.

Intensity	Min	210
	Max	255
Size	Min	4 pixels
	Max	50 pixels
Movement	MaxDistanceDiff	5 pixels
	MaxIntensityDiff	15
	MaxSizeDiff	10 pixels
Triangulation Tests 1,2,4	$\{-50 \leq X \leq 50\}$ mm $\{-50 \leq Y \leq 50\}$ mm $\{-50 \leq Z \leq 100\}$ mm	
Test 3	$\{-200 \leq X \leq 200\}$ mm $\{-200 \leq Y \leq 200\}$ mm $\{-350 \leq Z \leq 500\}$ mm	

Table 4.1: Parameters used for Test Batteries

- *Test Battery 3*: While maintaining constant scene complexity and illumination, place objects over a wide range of three dimensional locations in the workspace. The objective here is to determine how the triangulation algorithm performs when the valid three dimensional volume is set to encompass a large portion of the workspace.
- *Test Battery 4*: While maintaining constant scene complexity and illumination, vary the order in which the region selection algorithms are executed.

The parameters used for the region selection algorithms are summarized in table 4.1. The X, Y, and Z parameters for the triangulation algorithm define the valid three dimensional volume for region point estimates. The volume is defined with respect to the X, Y, and Z coordinates of the expected location of the laser spot. The expected location of the laser is determined by instructing the program to direct the laser beam at a specific world point. The three dimensional volume was changed for test battery three since the objective of these tests is to determine if the triangulation algorithm will work with a large valid volume.

The results of each test are encapsulated in eleven parameters defined as follows:

- *AvgRegion* The average number of regions detected by the region growing algorithm over all test trials.
- *AvgEquiv* The average number of regions eliminated as the laser region due to equivalency over all test trials.
- *AvgIntensity* The average number of regions eliminated as the laser region due to selection by intensity over all test trials.
- *AvgSize* The average number of regions eliminated as the laser region due to selection by size over all test trials.
- *AvgMovement* The average number of regions eliminated as the laser region due to selection by movement over all test trials.
- *AvgTriangle* The average number of regions eliminated as the laser region due to selection by triangulation over all test trials.
- *NoLaser* The number of test trials where no laser region was found.
- *Laser1* The number of test trials where one laser region was found.
- *Laser2* The number of test trials where two laser regions were found.
- *Laser3* The number of test trials where three laser regions were found.
- *LaserGT3* The number of test trials where more than three laser regions were found.

Each test in each test battery consisted of one hundred laser points. T_{back} and T_{region} were set to 190 and 10 respectively.

4.5.1 Analysis of Test Battery 1 Results

The first battery of tests were designed to study the behavior of the laser spot selection algorithms to images of varying complexity. The first test was conducted on a scene of low complexity that contained a few metal objects and a cable. The second test was conducted on a scene of higher complexity that included more multi-faceted metal objects and a few hand tools. Finally, the third test was conducted on a highly complex scene that included struts and nodes, grippers from the CIRSSE robot arms and a metal plate. The results of each test are presented in table 4.2.

Parameter	Test 1	Test 2	Test 3
AvgRegion	96	129	546
AvgEquiv	15	17	89
AvgIntensity	36	52	200
AvgSize	33	45	203
AvgMovement	10	12	49
AvgTriangle	0/1	0/1	2
NoLaser	3	4	8
Laser1	97	96	90
Laser2	0	0	2
Laser3	0	0	0
LaserGT3	0	0	0
Image Complexity	Low	Moderate	High

Table 4.2: Results of Test Battery One

The results of the three experiments indicates that elimination by size and by intensity had the greatest effect in reducing the number of regions in the image. Further, the triangulation algorithm was not required for the first two tests, which implies that size, intensity, and movement are sufficient criteria for identifying the laser in simple or moderately complex images. What is most significant about the test results is the number of times the laser spot was identified. The worst case results indicate that the selection algorithms produced a unique solution for the laser spot 90% of the time. What is even more interesting, is that there are few instances of identifying more than one region as the laser. Indeed, the results indicate that the selection algorithms either found a single laser spot or none at all.

The instances where the laser was not found can be attributed, in most cases, to effects of illumination or occlusion that prevented the camera from distinguishing the laser region. For example, it was observed in several instances that the laser spot was projected near an object such that the object occluded the laser spot from the camera's field of view. Further, there were regions in the image that registered as full intensity (i.e. 255) and when the laser was directed into these regions, the camera was unable to distinguish the laser spot since it was embedded in a region that saturated the CCD array. In short, failure to detect the laser region altogether is due mostly to the inherent limitations of the equipment used to perform the test and not to the performance of the algorithms.

4.5.2 Analysis of Test Battery 2 Results

The second test battery was designed to study the effects of illumination on the selection algorithms' performance. The first test in this set was performed on an image of moderate complexity (similar to test 2 in the first test battery) under normal room lighting conditions. This first test provided a baseline for gauging performance on other tests. The second test was conducted with the lights off. The third test was conducted with a single light source projected from one end of the scene (the top of the image) and oriented to place strong shadows on the objects in the scene. In the final test, a single light source was oriented not only to project strong shadows on the objects in the scene, but also to project reflection spots and halo effects into the camera's lens. The results of these tests are presented in table 4.3.

Parameter	Test 1	Test 2	Test 3	Test 4
AvgRegion	474	3	368	568
AvgEquiv	63	0/1	46	131
AvgIntensity	213	0/1	154	176
AvgSize	159	1	127	227
AvgMovement	31	0/1	37	28
AvgTriangle	5	0/1	1	3
NoLaser	5	1	5	15
Laser1	93	99	94	84
Laser2	2	0	1	1
Laser3	0	0	0	0
LaserGT3	0	0	0	0
Illumination	Normal	No lights	Shadow	Shadow & lens reflection

Table 4.3: Results of Test Battery Two

The results indicate that for the first three tests, the selection algorithms were largely resilient to changes in ambient light, in that the selection algorithms achieved a unique solution in more that 93% of the trials. What is particularly interesting is how performance degraded in test 4. It is not surprising that performance would degrade if light is directed in the camera, but what is interesting is that the pattern of performance is nearly identical to the results obtained in the most complex image of the first test battery. While the similarity in the exact numbers may be a coincidence, the pattern indicates that the performance of the selection algorithms degrades in a

Parameter	Test 1	Test 2
AvgRegion	499	189
AvgEquiv	103	41
AvgIntensity	208	74
AvgSize	132	52
AvgMovement	43	17
AvgTriangle	9	2
NoLaser	17	8
Laser1	68	79
Laser2	5	11
Laser3	4	2
LaserGT3	6	0

Table 4.4: Results of Test Battery Three

consistent manner. That is, the algorithms either achieve a unique solution for the laser spot or none at all. When test 4 was actually conducted, it was observed that the laser region was lost when the laser was directed into a region of the scene that was highly illuminated (to the point where the camera was saturated) or the laser was directed into a region of the image that contained a lens reflection. In both of these cases, the intensity of the image registered as 255 which is the maximum intensity value for the vision system, and therefore, the laser spot was visually indistinguishable from the surrounding image.

4.5.3 Analysis of Test Battery 3 Results

The third test battery was designed to test the effectiveness of the triangulation algorithm if the valid three dimensional volume was enlarged to cover a greater portion of the testbed. In this test, objects were placed in the workspace in such a way as to ensure that specular reflections and other noise were present over a wide three dimensional volume in the workspace. This arrangement was adopted to increase the chance that specular reflections would lie along the laser ray thereby increasing the probability that the selection algorithms would misidentify some of these regions as being attributable to the laser. The results for test battery three are presented in table 4.4. The two tests were virtually identical, although the arrangement of objects in the scene was altered between tests to provide different scenes of similar complexity.

There are two significant observations that can be made about these results. First, while enlarging the valid three dimensional volume does increase the probability that more than one region in an image will be identified as the laser spot, in the large majority of cases the selection algorithms either achieved a unique solution for the laser spot or could not find the laser region at all. In other words, the general behavior of the selection algorithms in this battery of tests is similar to the the other test batteries.

A second observation was made while studying the behavior of the selection algorithms as the experiments were conducted. In cases where the selection algorithm generated multiple solutions for the laser spot the regions in question were either in close proximity to the actual laser spot (to within a few millimeters) or were a significant distance away from the actual laser spot and situated along the projection of the laser ray through the image. These observations are not surprising when one considers that the laser ray and the projected ray from the camera to these regions are not significantly divergent. To eliminate regions that do not result in divergent rays, it is necessary to be able to make accurate estimates of where the laser spot is expected. However, as is apparent by these test results, even a rough estimate of the valid three dimensional volume results in reasonable performance.

4.5.4 Analysis of Test Battery 4 Results

The fourth test battery was designed to determine if changing the execution order of the selection algorithms resulted in a significant change in performance. The results of these tests are presented in table 4.5.

The results of these tests indicate that while the pattern of region elimination differs between the different orders of algorithm execution, the final results for identifying the laser spot are virtually identical across all the tests. The implications of this result is that the order of execution can be arranged to optimize the overall performance of the selection algorithms without sacrificing reliability.

To optimize performance of the selection algorithms it is necessary to examine the equations for execution time of each algorithm that were derived previously. The total execution time of the laser selection process is the sum of these individual equations as indicated below:

Parameter	Test 1	Test 2	Test 3
AvgRegion	362	363	364
AvgEquiv	38	38	38
AvgIntensity	115	6	6
AvgSize	174	286	15
AvgMovement	30	29	306
AvgTriangle	1	1	1
NoLaser	6	5	5
Laser1	92	94	93
Laser2	2	1	2
Laser3	0/1	0/1	0
LaserGT3	0/1	0/1	0
Order of Execution	Intensity Size Movement Triangle	Size Intensity Movement Triangle	Movement Size Intensity Triangle

Table 4.5: Results of Test Battery Four

$$\Gamma(total) = \Gamma(intensity) + \Gamma(size) + \Gamma(movement) + \Gamma(triangle) \quad (4.5)$$

$$\Gamma(intensity) = C_{valid}(N) + C_{intensity}(n_{l1}) \quad (4.6)$$

$$\Gamma(size) = C_{valid}(N) + C_{size}(n_{l2}) \quad (4.7)$$

$$\Gamma(movement) = C_{valid}N_{current} + C_{movement}n_{l4}N_{previous} \quad (4.8)$$

$$\Gamma(triangle) = C_{valid}(N) + C_{triangle}n_{l3} \quad (4.9)$$

where the expressions for $\Gamma(intensity)$, $\Gamma(size)$, $\Gamma(movement)$, and $\Gamma(triangle)$ are restated here for convenience. The total execution time is dependent on the total number of regions in the image (N) and the number of possible laser regions passed to the individual selection algorithms (n_{l1}, n_{l2}, n_{l3}, n_{l4}). Recall from the previous discussion that $\Gamma(movement)$ executes in $O(N^2)$ time in the worst case and $C_{triangle} \gg \{C_{intensity}, C_{size}\}$. Therefore, the best way to reduce overall execution time is to reduce the contribution of the $\Gamma(movement)$ and $C_{triangle}$ terms. This can be accomplished by keeping n_{l3} and n_{l4} small. In other words, use the size and intensity al-

gorithms to eliminate as many regions as possible before executing the more computationally complex movement and triangulation algorithms. Since the execution time for the intensity and size algorithms are roughly equivalent (see section 4.2), the order in which these algorithms are executed will have little effect on performance.

At this point, it is clear that the intensity and size algorithms should be executed first to reduce the number of possible laser regions that must be sent through the movement and triangulation algorithms. The next issue is whether the movement algorithm or the triangulation algorithm should be executed next. The results from all the tests indicates that $\{n_{l3} \approx n_{l4}\} \ll N$, and, hence, the movement algorithm's execution time will more closely approximate $O(N)$ instead of $O(N^2)$.

The execution time of the triangulation algorithm also approximates $O(N)$. It was mentioned previously that the value of $C_{triangle}$ is significantly larger than $C_{intensity}$ and C_{size}. In fact, $C_{triangle}$ is also much larger than than $C_{movement}$ which represents the execution time of a small number of comparisons to determine if two regions are equivalent. When n_{l3} is small, the value of the the execution time for the triangulation algorithm is comparable to that of the movement algorithm despite the fact that the movement algorithm theoretically should be less efficient. Therefore, the total execution time of the laser selection algorithms will be relatively constant regardless of whether movement is executed before triangulation or vice-versa.

4.5.5 Conclusions about Laser Selection Performance

The results presented in the previous sections have provided a plethora of information about the behavior of the laser selection algorithms, both individually and in concert with each other. From the results and the subsequent analysis it is possible to draw several conclusions about the performance of these algorithms.

1. The combination of all four selection algorithms locates the laser spot reliably under normal lighting conditions and moderate to high scene complexity.

2. Degradation of algorithm performance results in a decreased potential for locating the laser spot in the image as opposed to inadvertently selecting multiple regions as the laser spot.

3. Algorithm performance is most affected by occlusion of the laser spot in the workspace and saturation of the camera's CCD array due to aperture setting.

4. Laser selection algorithms operate most efficiently when the intensity and size selections are executed first followed by movement and triangulation.

There are other properties of the laser spot that may be useful in enhancing the discrimination of the spot in a camera image. First, the laser spot has a specific spectral wavelength and if the camera were fitted with a filter that is sensitive only to the wavelength of the laser, it would be easier to locate the laser spot. Additionally, the laser spot has an elliptical geometry. Therefore, if the region growing algorithm were modified to record more information about the geometry of each region such as the length of the region's perimeter and the region's moments of inertia, it might be possible to eleiminate regions that do not resemble a small ellipse.

Overall, the techniques discussed in this section for locating a laser spot in an image offer a reliable method for laser tracking under a variety of lighting and scene conditions. These methods will permit a calibrated laser and camera to operate under the same conditions as multiple camera configurations. Such a capability permits a laser and camera to not only be a useful active three dimensional sensing device in its own right, but it also provides the ability to verify results obtained by passive techniques.

Chapter 5

Conclusions and Future Research

The laser research has successfully fulfilled the major objectives described in section 1.2. Specifically, methods have been developed to calibrate the laser to a level of accuracy that is comparable to the multiple camera system. Second, the laser can operate in concert with a calibrated camera to perform three dimensional point estimation. Finally, techniques are now available to permit a camera to locate the laser spot under a variety of conditions.

With these capabilities in place, there are many opportunities to use the laser in future research. In the near term, research will be focusing on three major areas. First, the laser and a camera will be used to perform three dimensional ranging. Under this arrangement, the laser can be used to scan the workspace to detect objects or major features. For example, recalling the case studies performed in November of 1990, the experiments were performed on a work surface that was covered with a black cloth. If the partially completed triangle was not on the work surface, it would be impossible for the camera system to know where the work surface is located because a plain black image offers no features to identify and extract. However, a calibrated laser could scan over the work surface and by employing the point estimation algorithms described in chapter 3 it would be possible to determine the location and orientation of the work surface.

A second avenue of research to be examined in the near term is the use of the laser to highlight objects or features already identified by the multiple camera system. This capability could be incorporated in future vision research as a method of verifying the performance of the vision system,

and hence, improving its reliability. As part of this research, the laser could be integrated into the current CIRSSE testbed applications for assembly of structures.

Another area of research is to continue exploring the effects of lighting not only on the laser and camera configuration, but also on passive multiple camera configurations. This area of research should include both quantitaive and qualitative analysis of lighting and the development of methods to compensate for its effects.

Finally, effort will be directed toward creating and testing a set of mathematical models that can predict the effects of calibration error on system performance. Such models can serve several purposes. First, by predicting how the laser will behave under certain conditions, it may be possible to determine how the laser can be used in an optimal manner. Second, a mathematical model for the laser can be used to determine a level of reliability for the three dimensional data generated by a camera and laser. Finally, these models can shed light on the appropriate use of the point estimation alorithms as discussed in section 3.3

Bibliography

[1] K. Pelsue, "Precision, post-objective, two-axis, galvanometer scanning," in *High Speed READ/Write Techniques for Advanced Printing and Data Handling, SPIE Proceedings*, (Los Angeles, California), p. 70, January 1983.

[2] J. R. Noseworthy and L. A. Gerhardt, "Three dimensional vision - requirements and applications in a space environment," in *Symposium on Advances in Intelligent Systems, SPIE Proceedings*, (Boston, Massachusetts), p. 1387, November 1990.

[3] J. Watson, "Testbed kinematics and routines." CIRSSE Technical Memorandum, March 1991.

[4] W. J. Mistretta, "Analysis, enhancement, and application of a three dimensional inspection system in computer integrated manufacturing," Master's thesis, Rensselaer Polytechnic Institute, Electrical & Computer Systems Engineering Dept., Troy, NY 12180-3590, December 1989.

[5] R. Y. Tsai, "A versatile camera calibration technique for high-accuracy 3d machine vision metrology using off-the-shelf tv cameras and lenses," *IEEE Journal of Robotics and Automation*, vol. RA-3, pp. 323–344, August 1987.

[6] J. J. Craig, *Introduction to Robotics, Mechanics, & Control.* Reading, MA: Addison-Wesley Publishing Company, 1986.

[7] W. H. Press, B. R. Flannery, S. A. Taukolsky, and W. T. Vetterling, *Numerical Recipes in C: The Art of Scientific Computing.* Cambridge University Press, 1990.

[8] D. Sood, M. C. Repko, and R. B. Kelley, "An implementation of a camera technique to obtain three-dimensional (3d) vision information for simple assembly tasks," CIRSSE Report 76, Rensselaer Polytechnic Institute, Troy, NY, November 1990.

[9] J. R. Noseworthy, "Inaccuracies of three dimensional vision systems - theory and practice," Master's thesis, Rensselaer Polytechnic Institute, Electrical & Computer Systems Engineering Dept., Troy, NY 12180-3590, August 1991.

[10] D. H. Ballard and C. M. Brown, *Computer Vision.* Englewood Cliffs, NJ: Prentice-Hall Inc., 1982.