AIIM

Association for Information and Image Management

1100 Wayne Avenue, Suite 1100
Silver Spring, Maryland 20910
301/587-8202

Centimeter

Inches

MANUFACTURED TO AIIM STANDARDS
BY APPLIED IMAGE, INC.

1 of 1

AUTOMATIC CALIBRATION OF LASER RANGE CAMERAS USING ARBITRARY PLANAR SURFACES*

James E. Baker
Center for Engineering Systems Advanced Research
Oak Ridge National Laboratory
P.O. Box 2008
Oak Ridge, TN 37831-6364

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

"The submitted manuscript has been authored by a contractor of the U.S. Government under contract DE-AC05-84OR21400. Accordingly, the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes."

To be published in the Proceedings of SPIE's Orlando'94 Conference, Orlando, Florida,
April 4-8, 1994

* Research sponsored by the Engineering Research Program, Office of Basic Energy Sciences, under contract DE-AC05-84OR21400 with Martin Marietta Energy Systems, Inc., and by the Naval Air Warfare Center Aircraft Division Lakehurst under Interagency Agreement No. 2072-E123-A1 with the U.S. Department of Energy.

MASTER

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

Automatic calibration of laser range cameras using arbitrary planar surfaces

James E. Baker

Oak Ridge National Laboratory, P.O. Box 2008,
Bldg. 6025, MS-6364, Oak Ridge, TN 37831-6364

ABSTRACT

Laser Range Cameras (LRCs) are powerful tools for many robotic/computer perception activities. They can provide accurate range images and perfectly registered reflectance images of the target scene, useful for constructing reliably detailed 3-D world maps and target characterizations. An LRC's output is an array of distances obtained by scanning a laser over the scene. To accurately interpret this data, the angular definition of each pixel, i.e., the 3-D direction corresponding to each distance measurement, must be known. This angular definition is a function of the camera's intrinsic design and unique implementation characteristics, e.g., actual mirror positions, axes of rotation, angular velocities, etc. Typically, the range data is converted to Cartesian coordinates by calibration-parameterized, non-linear transformation equations. Unfortunately, typical LRC calibration techniques are manual, intensive, and inaccurate. Common techniques involve imaging carefully orchestrated artificial targets and manually measuring actual distances and relative angles to infer the correct calibration parameter values.

This paper presents an automated method which uses Genetic Algorithms to search for calibration parameter values and possible transformation equations which combine to maximize the planarity of user-specified sub-regions of the image(s). This method permits calibration to be based on an arbitrary plane, without precise knowledge of the LRC's mechanical precision, intrinsic design, or its relative positioning to the target. Furthermore, this method permits rapid, remote, and on-line recalibration — important capabilities for many robotic systems. Empirical validation of this system has been performed using two different LRC systems and has led to significant improvement in image accuracy while reducing the calibration time by orders of magnitude.

1. INTRODUCTION

Laser Range Cameras (LRCs) are designed to produce a "depth image" of the scene rather than the typical ambient-light reflectance image, e.g., like those produced by CCDs. A "depth image" is an array of pixels whose grey-scale values correspond to actual distances to each target point. LRC imaging is typically performed by scanning a laser beam in a two dimensional grid pattern and measuring the returned (reflected) signal to determine the target distances. This array of distances can then be converted into a 3-dimensional representation of the scene. However, to *accurately* make such a conversion, one must know the *precise* angular definition of each pixel, i.e., the 3-D direction corresponding to each distance measurement. Without this information, converting a depth image to an accurate world map is impossible. Unfortunately, LRCs provide only the measured distances and their relative position within the scan, e.g., row and column number, the user must determine the angular definitions and hence the spatial transformations appropriate for their particular camera model *and* unit.

We will refer to a particular set of transformation equations and corresponding parameter values which convert an LRC pixel (**row, col, range**) into a 3-dimensional Cartesian coordinate (**x, y, z**) as an LRC's "mapping". The LRC mapping which accurately converts a particular camera's images into 3-dimensional space is defined by the camera's design model, i.e., how the mirrors combine to optically provide the two dimensional scan, and specific camera parameters, e.g., precise field-of-view, initial angular position, mirror pivot positions relative to scan sweep, accuracy of mirrors' relative axes of rotation, etc. The camera's model can be determined by analyzing its design schematics or examining the actual mechanism. From this analysis, the basic spatial transformation equations can be derived, see subsection 3.2. The specific camera parameters are more difficult to determine and require either precision measurements of the individual camera components during normal operation or an empirical analysis of the camera's performance.

Typically, LRCs are calibrated "manually"¹ by imaging carefully designed targets and/or physically measuring absolute distances and angles, and inferring the corresponding parameter settings. Such a method is extremely time-consuming, error prone, and imprecise. Although many factors may adversely effect the quality/applicability of one's calibration, e.g., camera "wear and tear" or subsequent minor repairs, changing ambient light conditions, varying target distances and reflectivity, etc., recalibration can be considered prohibitively expensive and online recalibration impossible. Thus, one is left with an expensive camera system which may never reach its imaging potential for even static environments.

This paper presents a fast, automatic, "self-calibration" method which requires neither carefully designed targets nor a detailed analysis of the camera model. (This method is applicable to any type of scanning LRC, thus eliminating the need to

analyze/characterize the internal operation of a particular camera.) Section 2 describes the method's basic approach. Section 3 describes its implementation, including a detailed overview of its global search method, Genetic Algorithms. Section 4 presents an empirical analysis of the system by comparing its results with those obtained by more traditional methods, and Section 5 summarizes the work to date.

2. APPROACH

The basic premise on which this calibration system is based is that an LRC's mapping can be accurately and automatically inferred by evaluating the planar quality of its output when imaging a surface known to be planar. Note that there is a subtle, yet significant, distinction between relying on a "known planar surface" versus a "surface known to be planar". While both approaches are a type of template-based pattern matching, the former approach requires precise knowledge of the target's relative position and orientation with respect to the camera and its internal components. This requirement places significant demands and limitations on the calibration process, including increased LRC down-time for calibration, increased calibration turn-around time, and a reduced remote calibration capability. The latter approach requires simply evaluating the planar quality of the data, independent of its position and orientation.

The basic design of this calibration system can be outlined by the following six steps (which will be numerically referenced throughout the remainder of this paper): 1) a set of range pixels from a single LRC image, which are known to target a single plane, are chosen as a calibration training set; 2) a single candidate mapping is defined for evaluation; 3) the candidate mapping is used to convert the training set of range pixels into 3-dimensional Cartesian coordinates; 4) the overall planar quality of these Cartesian points is quantified for use as a measure of the quality of this mapping; 5) based on this and previous such evaluations, a new mapping is defined for evaluation; 6) the process cycles back to step #3, until some stopping criterion is met. (Note that a number of separate planes (training sets) can be simultaneously used for calibration by modifying step #4 to appropriately merge the multiple results into a single quantity.)

This adaptive, results-based approach eliminates the need for exacting measurements describing the camera's internal hardware and/or the relative positional relationship between the camera and the target points. Furthermore, this approach converts the LRC calibration problem into a straightforward, albeit complex, global search for the best mapping for a given LRC and is well suited to available adaptive global search techniques, e.g., Genetic Algorithms.²

This approach has a number of advantages, among them are:

-) *minimal down-time* — By "down-time", we refer to the period during which the LRC is unavailable for normal operation. This calibration system's dependance upon arbitrary planes within the scene minimizes down-time by requiring, at most, the additional capture of a single LRC image for calibration analysis.

-) *rapid turnaround time* — By "turnaround time" we refer to the total delay between initiating calibration and obtaining a new mapping. Rapid turnaround time enables the system to support greater calibration analysis and potential utilization. e.g., understanding the effect of ambient light on the calibration parameters could facilitate the adaptive modification of the mapping in order to compensate for changing environmental conditions. Manual calibration approaches can often take days or even weeks to perform, as they require precise physical measurements of distances between coordinate target points and internal camera positions, e.g., the actual distance to given pixels' target points from their reflectance points on the first incident mirror. The turnaround time for our system is a function of available processing speed and the training set's size. When employing a single processor analyzing a large training set (60K+ pixels), we achieved a turnaround time on the order of an hour, see subsection 4. However, the turnaround times could be reduced by orders of magnitude with parallel processing or approximate evaluation.³

-) *minimal hardware dependance* — Just as hardware implementation eccentricities, e.g., mirror rotational velocities, are a part of the overall system for which the calibration parameters are being optimized, the basic camera model can also be considered a parameter space to be searched. The correct parameter values combined with the correct camera model equations should yield a superior mapping, and hence, the global search need not be restricted to equation coefficients, but can be extended to include the form of the equation as well. This capability virtually eliminates the need to analyze and characterize specific LRC hardware — emphasis is correctly placed on its output, not its particular implementation.

-) *minimal scene orchestration* — By relying on arbitrary planes in the scene, this system does not require carefully designed and positioned targets or external determination of specific target points' relative position to the camera — a very tedious, time-consuming, and error prone process.

-) *automatic parameter optimization* — By automating the parameter optimization process one increases the speed, objectivity, and "user-friendliness" of the calibration and enables online, less labor intensive, and remote calibration.

-) *online operability* — By "online" we refer to the ability to perform (re)calibration while the LRC remains in normal operation. This capability is of significant importance in critical or real-time applications, and would permit adaptive recalibration for some dynamic environments and/or utilizations.

•) *remote operability* — Since relative target to camera position measurements are not necessary with this system, a complete calibration can be performed remotely and, after image acquisition, offline. This can be of significant benefit for applications in remote or hazardous environments.

3.0 IMPLEMENTATION

As described in the previous section, this calibration system approach requires two major subsystems: a planar evaluation function which will quantify the planar quality of a given set of 3-dimensionally placed points (see step #4, above); and a global search method to efficiently and automatically reveal superior mappings, as defined by the planar evaluation function (see step #5, above). This section outlines these subsystems and defines the mappings' search space.

3.1 Planar evaluation function

The planar evaluation function is responsible for quantifying the planar quality of a set of 3-dimensional Cartesian points. This function will be used to infer the quality/accuracy of a candidate mapping, i.e., a mapping's ability to accurately convert the LRC's range pixels into 3-dimensional space will be inferred by its handling of target points known to be interplanar.

The planar evaluation function can be subdivided into three tasks:

1) *to find the plane which "best fits" the Cartesian points* — Finding the "best fit" plane is somewhat complicated by the fact that we equally distrust all three coordinates — calibration errors are no more likely to impact a pixel's range value than its angular definition. Given this situational guidance, a least squares fit plane fitting algorithm, for data equally uncertain in all coordinates, was developed by D. B. Reister, et. al.⁴ and is used for this task.

2) *to determine each point's error with respect to the "best fit" plane* — Our uniform distrust of each coordinate also guided our error determination method. We chose to use the point's orthogonal distance to the plane, rather than, for example, the distance along the laser ray, since it placed the fewest assumptions on the relative quality of the coordinates.

3) *to accumulate the points' errors into a single measure of the training set's planarity* — There are many possible ways to resolve this task, e.g., sum of magnitudes, standard deviation, maximum magnitude, etc. We somewhat arbitrarily chose to use the sum of the squares. This method is appealing since the "square" term concentrates attention on the worst points while the "summation" ensures a fine resolution, i.e., each pixel substantively contributes to the outcome.

3.2 Mapping — mathematical model

A mapping converts an LRC pixel (**row**, **col**, **range**) into a 3-dimensional Cartesian coordinate (**x**, **y**, **z**). We perform this conversion in two steps, first from an LRC pixel (**row**, **col**, **range**) to a spherical type of coordinate (α , β , **radius**) then from this triplet to a Cartesian coordinate (**x**, **y**, **z**). For clarity, we will henceforth refer to the first conversion as the calibration and to the second as the transformation.

3.2.1 Calibration conversion

For the initial system, we will make two simplifying assumptions: 1) α and β are linear functions of **row** and **col**, hence, $\alpha = A + B \cdot \text{row} + C \cdot \text{col}$ and $\beta = D + E \cdot \text{col} + F \cdot \text{row}$, for some unknown constants **A**, **B**, **C**, **D**, **E**, and **F**; and 2) the differential range value is accurate, i.e., the LRC can accurately measure distance with respect to some unknown origin, hence, **radius** = **G** + **range**, for some unknown constant **G**. If appropriate, these assumptions will be relaxed in future research.

While these parameterized equations could be searched by GAs, they are unnecessarily inefficient and span enormous, "trivial" local optima, i.e., equations which result in effectively zero or one dimensional objects rather than the fundamentally 2 dimensional objects of the training set. These shortcomings can be resolved by judiciously restricting the parameters' search space to regions likely to contain the LRC's true mapping. In this initial implementation, the calibration conversion is based on a general solution for scanning LRCs and the search is restricted to values near those theoretically implied.

A scanning LRC samples according to some grid pattern, as defined by the dynamics of its internal mirrors. We will approximate the grid elements as parallelograms, rather than trapezoids or arbitrary quadrangles. These more complex shapes are considered inconsistent with the realities motivating LRC design choices. LRCs are designed for general applicability and ease-of-use (rectangular grids) and mechanical simplicity and rapid image acquisition, therefore continuous or continuous-step motion mirror movement (parallelogram grids). Admittedly, more complex grid shapes may result from subtle implementation inaccuracies, e.g., multisided, rotating mirrors with non-parallel surfaces, but at this stage, these effects are assumed to be negligible by comparison.

If the LRC scan painted a perfect rectangular grid pattern, then the angular calibration equations could be written as:

$$\begin{aligned}\alpha &= \text{fovV} / 2 - \text{fovV} * \text{Row} / (\text{ResV} - 1) \\ \beta &= \text{fovH} / 2 - \text{fovH} * \text{Col} / (\text{ResH} - 1)\end{aligned}$$

where **fovV** = vertical field of view
ResV = vertical resolution, i.e., number of rows
Row = row of interest, will range from 0 to **ResV-1**
fovH = horizontal field of view
ResH = horizontal resolution, i.e., number of columns
Col = column of interest, will range from 0 to **ResH-1**

and the coordinate system is oriented so that the radial axis (0, 0, **radius**) passes through the center of the camera and of the image.

However, if the LRC scans the rows and/or columns in a continuous motion, rather than the discrete step motion required to paint a perfect rectangular grid, then additional terms are necessary. In the following angular calibration equations, the last terms have been added to account for continuous motion effects of both the row and column based mirrors:

$$\begin{aligned}\alpha &= \text{fovV} / 2 - \text{fovV} * \text{Row} / (\text{ResV} - 1) - \text{fovV} * \text{Col} / (\text{ResV} * (\text{ResV} - 1)) \\ \beta &= \text{fovH} / 2 - \text{fovH} * \text{Col} / (\text{ResH} - 1) - \text{fovH} * \text{Row} / (\text{ResH} * (\text{ResH} - 1))\end{aligned}$$

Thus, between any two adjacent rows, α is uniformly interpolated according to the value of the column, and similarly, between any two adjacent columns, β is uniformly interpolated according to the value of the row. By generalizing these calibration equations with parameter coefficients, we make them applicable to any LRC which scans a parallelogram grid pattern:

$$\begin{aligned}\alpha &= A / 2 - (A+B) * \text{Row} / (\text{ResV} - 1) - (A-C) * \text{Col} / (\text{ResH} * (\text{ResV} - 1)) \\ \beta &= D / 2 - (D+E) * \text{Col} / (\text{ResH} - 1) - (D-F) * \text{Row} / (\text{ResV} * (\text{ResH} - 1)) \\ \text{radius} &= G + \text{range}\end{aligned}$$

A+B and **D+E** represent the respective field-of-views, **C** and **F** indicate the rectangular quality of the LRC's scan, while the magnitudes of **B** and **E** indicate the image's symmetry with respect to the (0, 0, **radius**) axis. The trivial solutions described above, i.e, conversions which result in effectively zero or one dimensional objects, occur whenever the magnitudes of **(A+B)** and/or **(D+E)** approach zero. By appropriately constraining each parameter's ranges, one can increase the efficiency of the search without jeopardizing its potential, e.g., **A** and **D** should be limited to values near their respective fields-of-view while **B** and **E** should be limited to values near zero. In general, only one of the two angular calibration equations may be effected by continuous motion of its respective mirror, hence, **C** and **F** should be allowed to search between near zero (corresponding to full effect) and near their respective field-of-view (corresponding to no effect).

3.2.2 Transformation conversion

We referred to the (α , β , **radius**) coordinate system as a "spherical type" since the reference frame of its angular coordinates is hardware implementation dependent and, hence, unknown. Either or both of these angles may be either with respect to the original coordinate frame or to the transformed coordinate frame, e.g., α may be with respect to the original axes or with respect to the axes obtained by performing the β transformation. Thus, four distinct transformation equation sets are possible, depending entirely upon the internal hardware design of the LRC's mirrors' motions. We will refer to these equation sets by "type number", **T**:

Type **T=0**:

$$\begin{aligned}x &= \sin(\alpha) \\ y &= \sin(\beta) \\ z &= \sqrt{\cos^2(\alpha) - \sin^2(\beta)}\end{aligned}$$

Type **T=1**:

$$\begin{aligned}x &= \sin(\alpha) \\ y &= \cos(\alpha) \sin(\beta) \\ z &= \cos(\alpha) \cos(\beta)\end{aligned}$$

Type T=2:

$$\begin{aligned}x &= \sin(\alpha)\cos(\beta) \\ y &= \sin(\beta) \\ z &= \cos(\alpha)\cos(\beta)\end{aligned}$$

Type T=3:

$$\begin{aligned}x &= \sin(\alpha)\cos(\beta) \\ y &= \cos(\alpha)\sin(\beta) \\ z &= \cos(\alpha)\cos(\beta)\end{aligned}$$

The transformation equation set corresponding to a particular LRC can be determined by analyzing the internal mirror configuration and its dynamics. Alternatively, the GA can optimize an additional parameter which would indicate the transformation equation set to be used when evaluating a candidate solution. This latter approach adds significant flexibility since it furthers the "black box" approach to sensor analysis.

3.3 Global search — Genetic Algorithms

The success of the overall calibration system is largely dependant on the quality and effectiveness of its global search subsystem, (see step #5, above). For even a relatively coarse mapping search, the search space has significant dimensionality (8+), is very large (2^{64+}), and is represented by noisy sensor data. Depending on the size of the training set, a single mapping evaluation (steps #3 and #4, above) may take a few seconds. Hence, to ensure a turnaround time of less than even two hours would permit the evaluation of only a few thousand mappings — or less than one trillionth of a percent of the search space. (Note that these numbers represent minimum values used in my empirical analysis and are a function of the quality and precision desired.) Thus, the search strategy must be highly efficient and applicable to noisy, nonclosed-form equations that are either highly dimensional or, if one considers the concatenated representation as a single dimension, highly multi-modal, nondifferentiable, and discontinuous. These limitations led us to choose Genetic Algorithms (GAs)² for our global search method.

GAs are extremely powerful adaptive global search techniques and have been shown to perform well for many types of functions including those exhibiting very difficult characteristics, e.g., discontinuities, non-differentiability, multi-modality, high dimensionality, huge search spaces and noise. GAs require no specific, a priori function information; all that is required is the form of a candidate solution, i.e., the number of parameters to be optimized and the desired level of precision (number of bits), and a comparative performance measure for candidate solutions. This makes GAs well suited to our application; with GAs, we do not need to understand the potentially complex interrelationships between the calibration parameters and transformation equations being optimized, we must simply define their existence and use the previously described planar evaluation function as the comparative performance measure.

Conceptually, GAs are a simplified simulation of the natural genetic model. As such, GAs simulate a population of individuals evolving over multiple generations: individuals are specified by a series of genes (bits) that can be independently inherited; reproduction is accomplished by a crossover operation which forms offspring from genetic material of two parents; and, an individual's reproduction frequency is based on his performance in the environment (evaluation function). Thus, trait encoding, sexual reproduction and "survival of the fittest" propagation are all simulated; and the average performance of the population tends to improve over successive generations. When some user-specified stopping criterion is met, the best individual produced is taken as the GA's solution for function optimization.

Clearly, GAs can not guarantee discovery of the optimal solution, but they have proven themselves as powerful global search techniques⁵ capable of simultaneously searching extensive regions of the parameter space. The basic GA framework used in this research was provided by the GENESIS⁶ GA package with modifications suggested by Ref. 7. However, the quality and efficiency of the search is highly dependant upon the user defined, evaluation function's ability to correctly guide the search.

In this application, an individual in the GA's search space is defined as a set of values for each of the eight parameters: **A, B, C, D, E, F, G, and T**, described in subsection 3.2. The GA's evaluation function will consist of: using the individual's parameter values to define a mapping according to the mathematical model provided in subsection 3.2, converting the training set of LRC pixels to three dimensional Cartesian coordinates, and evaluating the resulting points according to the planar evaluation function described in subsection 3.1. In the following section, we present an empirical analysis/validation of this calibration method on two different LRC systems.

4. EMPIRICAL ANALYSIS

The LRC calibration system presented in the preceding sections was empirically validated on two distinct LRC systems: a "3-D Laser Imaging System" developed by the Advanced Intelligent Machines Division of Odetics⁸ and a "5000 LASAR Series Scanner" by Perceptron, Inc.⁹ Since the Odetics LRC has been manually calibrated,¹⁰ we will perform a head-to-head comparison with our automated calibration system. The results of the automated calibration will also be presented for the Perceptron LRC, though no comparative calibration is available since that system has only recently been acquired.

Most of the figures in this section present either reflectance images or planar error plots. Typically, LRCs output both reflectance and depth images of their target scene. In a reflectance image, the pixel's color indicates the intensity of the reflected laser light and, as such, appears very similar to a standard CCD image. Because of this, reflectance images will be used in this section to more easily present and describe the target scenes.

In the planar error plots, the color indicates the amount of orthogonal error of each pixel with respect to some reference plane. All pixels having an error within the range $[-1, +1]$ LRC range units are linearly grey scaled from white to black — all other pixels are colored white. Hence, zero error corresponds to a medium grey. The reference plane is the "best fit" plane⁴ to points which are known to correspond to some planar surface of interest, e.g., the floor or a wall. All of the points within the intersection of the rows and columns marked by brackets at the edges of the images were used to define the "best fit" plane. Note that by using the data points themselves to define the reference plane, we remove all necessity to determine the camera's relative position and orientation with respect to the surface of interest.

4.1 Odetics LRC

The Odetics LRC produces both depth and reflectance images of the target scene. Each image consists of 128×128 pixels and corresponds to a claimed field-of-view of $60^\circ \times 60^\circ$. For automatic calibration purposes, we acquired an image consisting of approximately 80% unobstructed floor for use as the calibration plane. These floor pixels were accumulated as (**row, col, range**) triplets and stored as a training set. The GA was permitted to evaluate 5000 individuals, where an individual is a set of parameter values which defines a candidate mapping. The best individual found during that search, according to the planar evaluation function, is considered the GA's "solution". Note that no claim is made as to this being the optimum set of parameter values or that any of the values accurately define any single component of the LRC being calibrated. Rather, what is shown is that with this set of parameter values, the corresponding mapping is relatively effective in correctly converting the training set into 3-D points having a high degree of planarity.

The GA-based solution implies a field-of-view of approximately $66.1^\circ \times 63.5^\circ$ with the image being approximately $8.5^\circ \times 5.0^\circ$ off center. The magnitudes of the **C** and **G** parameters indicate little impact due to this factor. This may be a real effect, caused by rapid raster acquisition relative to the motion of the nodding mirror, or coincidental, resulting from the relative insignificance of these terms in the presence of noise and limited search time. Even though the correct mapping for the Odetics LRC is known to require the Type #2 transformation equations (see subsection 3.2.2), based on a theoretical analysis of the Odetics hardware,¹⁰ this parameter was also searched and the GA correctly agreed with the Type #2 assignment.

The quality of the GA-based and manual calibrations will be compared on three images:

- 1) a scene containing a significant wall segment — A wall was chosen as the first comparison since the manually derived parameter values were based on shapes roughly lying in a vertical plane, and hence, this surface orientation should favor the manual mapping and, in so doing, provide the most difficult challenge for the GA's solution, (see Fig. 1).
- 2) a scene containing a significant floor segment — Like the wall for the manually derived parameter values, the floor should favor the GA's solution, since the GA trained on a plane lying roughly in this orientation, (see Fig. 2).
- 3) a relatively cluttered scene — Given a sufficiently accurate calibration, a relatively small section of a plane should be sufficient to correctly identify a large, complex and even disjoint planar surface. The back wall of Fig. 5 provides such a surface for comparison and quality assessment.

Figure 3 compares the planar quality of the manually derived mapping¹⁰ with that of the GA's solution on the three piece wall segment shown at the top of Fig. 1. While the results look very similar, the GA-based calibration is slightly better — more of the wall is colored and the overall color is closer to zero error medium grey. The circular rings are assumed to be caused by the discretization accuracy (nominally $\sim 36\text{mm/LRC range unit}$) of the Odetics LRC and are an unavoidable product of imaging a plane with a discrete (digital) sensor which measures distances in a spherical coordinate system. Despite the GA's solution's superiority, the overall quality of the manually derived mapping on vertical planes is evident.

Figs. 1 and 2. Odetics LRC reflectance images with regions being: 1) WALL - three partition coplanar wall at top center, note trapezoidal shape; and 2) FLOOR - lower ~90% of image, receding to back walls.

Fig. 3a and b. Orthogonal distance of pixels from back wall plane of Fig. 1. Only pixels within one range unit are colored. Three-dimensional position of pixels is determined by: a) manually derived calibration, and b) automatic calibration.

Figure 4 compares the planar quality of the manually derived mapping with that of the GA's solution on the large floor segment shown in Fig. 2. There are several items of note in this figure:

- 1) The GA-based mapping is significantly better — nearly all floor pixels having an error of less than one range unit.
- 2) The GA conducted an effective parameter search for the calibration model currently being used. Our current mathematical model assumes α and β are linear functions of row and col, (see subsection 3.2.1). These assumptions place limitations on the types of errors that our parameter search can redress, i.e., calibration errors which effect all rows (columns) in a similar fashion. In Fig. 4a, the manual mapping results in each row of the floor being concave, a condition caused by a larger than appropriate horizontal field-of-view, but in Fig. 4b, errors which this mathematical model can resolve are all but eliminated.

- 3) The planar quality of the manually derived mapping is significantly dependant on the orientation of the planar surface, compare Fig. 3a with Fig. 4a. This problem reveals the difficulty of deriving robust global solutions from some manual calibration methods, due to both the very limited number of points feasible and the correspondingly high impact of erroneous measurements. While the automated system presented by this paper can simultaneously optimize multiple planar surfaces in order to assure orientation robustness, the system's solution quality using only a single plane, (compare Fig. 3b with Fig. 4b), demonstrates the advantage of the automatic optimization of a large, easily obtained training set — in this case, the single calibration image provided a calibration plane of nearly 12,000 points, without the need for any manual measurements.
- 4) The two spikes appearing in the middle of the images in Fig. 4 are present in all images taken with the Odetics LRC and are assumed to be the result of specular reflections of the laser light off of the LRC's glass cover plate. The small spike approximately 60% of the way down the left edge of the image, see Fig. 4b, was an object initially thought to be outside the field-of-view. The pixels corresponding to these three spikes were not used in the training set. However, the calibration quality of the GA-based mapping is such that all three spikes are clearly discernable.
- 5) The multiple nature of the interference pattern evident on both of the images cannot be explained by discretization precision. These patterns appear centered along all planar surfaces and, hence, are not simply the result of mechanical, e.g., rotational, effects of the moving mirrors. Their explanation remains, at present, unsatisfactorily resolved.

Fig. 4a and b. Orthogonal distance of pixels from floor plane of Fig. 1. Only pixels within one range unit are colored. Three-dimensional position of pixels is determined by: a) manually derived calibration, and b) automatic calibration.

The previous two comparisons demonstrate the quality and robustness of the GA's solution for the Odetics LRC calibration problem. This final comparison was designed to demonstrate both its quality and the potential usefulness of more accurate calibration. Assume the task is to isolate (segment) the pixels corresponding to some planar surface, e.g., a wall. The inherent planarity of the surface may be critical for accurate segmentation, but exploitation of this intrinsic property is entirely dependant upon calibration accuracy. Assume that, based on some preprocessing algorithms, a subset of the surface's pixels can be confidently identified. Given this subset and an accurate calibration, we should be able to identify most of the remaining pixels according to their planar error. Note that some of the pixels may not be identified due to noise or systemic errors, even with an optimal calibration parameter set.

For this comparison, we will attempt to identify those pixels targeting the back wall of Fig. 5a based on their orthogonal error with respect to the plane defined by the relatively small subset indicated, (see the margin indicators). For clarity, Fig. 5b presents a "cartoon" version of the data of Fig. 5a with the primary "objects" isolated: the floor; a large box in the center of the room; a moveable wall partition behind this box with a crate to its left and two writing desks next to that; two regions of the back wall, divided by a supporting pillar; a tool chest on the right with a wall-mounted shelf and a coil of electrical cables above it; the ceiling in the top center of the image and a heating duct in the top left; a segment of the left wall and a triple faceted, corner support structure.

Fig. 5a and b. Scene from which back wall pixels will be identified based on coplanarity with indicated regions: a) Odetics LRC reflectance image, and b) "cartoon" version, used solely for clarity of explanation.

Figure 6 compares the planar error resulting from the manually derived mapping with that resulting from the GA's solution. Both of the calibrations are sufficiently accurate to correctly disassociate most of the objects from the back wall. However, the manually derived mapping also disassociated many of the back wall pixels. Note the GA's solution's superior handling of pixels above the tool chest, above the wall partition, and above the crate and desks. Furthermore, the manual calibration sufficiently warped the back wall pixels that one of the corner supports erroneously fit the plane much better than nearby wall pixels. While the GA-based mapping resulted in some potentially fallacious planar associations, the result was significantly better and of sufficient quality to be useful for pixel characterization and subsequent object isolation in a system like that described in.¹¹

Fig. 6a and b. Orthogonal distance of pixels from back wall plane of Fig. 5. Only pixels within one range unit are colored. Three-dimensional position of pixels is determined by: a) manually derived calibration, and b) automatic calibration.

4.2 Perceptron LRC

Many of the imaging specifications are parameter controlled on the Perceptron LRC. For this calibration analysis, we chose parameters which most closely match ongoing project needs, including a 256 x 256 pixel depth image with a designated field-of-view of 60°H x 40°V. For automatic calibration purposes, we acquired an image consisting of approximately 95% unobstructed wall for use as the calibration plane. These wall pixels were accumulated as (row, col, range) triplets and stored as a training set. The GA was permitted to evaluate 5000 individuals, and the best individual found during that search, according to the planar evaluation function, is considered the GA's solution. As with the Odetics calibration, no claim is made as to this being the optimum set of parameter values.

The GA-based solution implies a field-of-view of approximately 69.8° x 41.7° with the image being approximately 1.2° x 2.3° off center. Figure 7 presents the planar error resulting from the GA's solution. In this figure, pixels having an error within the range [-7, +7] LRC range units are linearly grey scaled from white to black, rather than the [-1, +1] range displayed in the Odetics images in Figs. 3 through 6. All other pixels are still colored white. Several properties should be noted in Fig. 7:

- 1) The band down the right side of the image is apparently due to the Perceptron's errant scanning of the interior of its case. This band occurs on all 256 x 256 images and was not used to define the "best fit" plane.
- 2) Overall, the Perceptron's data appears very noisy. This significantly undermines the Perceptron's greater range measuring precision, i.e., the Odetics returns either 8 or 9 bit range values depending on the user's specifications while the Perceptron returns a 12 bit range value. However, if the three least significant bits are below the noise threshold, then the Perceptron is also effectively operating at no more than 9 bits.
- 3) There is a clear "ring" pattern of planar error in Fig. 7. Unlike the rings of Fig. 3 this pattern is not caused by discretization precision (these rings are sinusoidal, not discrete step functions) and unlike Figs. 4 and 6 this pattern seems to always be centered with respect to the image, not targeted planes. These rings are believed to be caused by the motions of the mirrors and their effect on the laser beam's internal flight path length during imaging. This may explain their sinusoidal nature and ovalar shape, as well as their absence in the Odetics images, due to its 8 bit resolution and internal hardware design differences.
- 4) Despite applying a linear model to what is clearly a non-linear problem, the GA's search was largely successful; linear error effects appear to be resolved.
- 5) The magnitude of the "ring" error may have a significant impact on the quality and robustness of the automated calibration. To properly calibrate this LRC, we must increase the complexity of the mathematical model from its current linear assumptions. This work is ongoing.

Fig. 7. Orthogonal distance of pixels from a plane, using a Perceptron LRC's image of a wall. Only pixels within seven range units are colored. Three-dimensional position of pixels is determined by automatic calibration.

5. SUMMARY

The automated calibration system presented in this paper has proven to be robust, undemanding, and relatively fast. Its robustness has been demonstrated by its successful application to two distinctly different LRCs and its applicability without any predesigned calibration objects or direct physical measurements. However, to properly calibrate some existing LRCs, e.g., Perceptron's 5000 LASAR Series Scanner, the underlying mathematical model must be extended to cover the non-linear effects of the mirrors' motions. It is undemanding by its "point and shoot" black box nature, i.e., after designating a planar region of any image, the user need have no further interaction with the system. Its speed is apparent in its minimal down-time — requiring the acquisition of a single image, and its rapid turnaround time — requiring on the order of an hour using common computing hardware, compared with the days or weeks required for many conventional (manual) calibration methods.

This paper presents an ongoing research effort. At present, research is focusing on expanding the mathematical model to include LRC appropriate non-linear terms. This addition will facilitate calibration of cameras like the Perceptron without decreasing the system's general applicability. Future research will include online recalibration for LRCs in slowly changing environments, e.g., due to changing weather conditions in outdoor applications.

6. ACKNOWLEDGMENTS

I wish to acknowledge the significant support provided by: 1) David Reister in his coordinated development of a method for obtaining the least squares fit plane to uncertain data; 2) Ole Doerum for his ongoing assistance in modeling LRCs and interfacing with the Perceptron LRC; and 3) Judson Jones for his assistance in interfacing with the Odetics LRC.

Research sponsored by the Engineering Research Program, Office of Basic Energy Sciences, under contract DE-AC05-84OR2100 with Martin Marietta Energy Systems, Inc., and by the Naval Air Warfare Center Aircraft Division Lakehurst under Interagency Agreement No. 2072E123A1 with the U.S. Department of Energy.

7. REFERENCES

1. Lenz, R. K. and R. Y. Tsai, "Techniques for Calibration of the Scale Factor and Image Center for High Accuracy 3-D Machine Vision Metrology," *IEEE Trans. Pattern Analysis and Machine Intelligence*, 10(5), pp. 713-720 (1988).
2. Holland, J. H., *Adaptation in Natural and Artificial Systems*, Univ. of Michigan Press, Ann Arbor, MI, 1975.
3. Grefenstette, J. J. and Fitzpatrick, J. M., "Genetic Search with Approximate Function Evaluations", *Proc. Intern. Conf. on Genetic Algorithms and Their Applications*, Pittsburgh, PA, 112-120 (1985).
4. Reister, D. B. and M. D. Morris, "A Method for Obtaining the Least Squares Fit of a Hyperplane to Uncertain Data", ORNL/TM-12686 (1994).
5. Booker, L. B., ed., *Proc. of the Fourth Intern. Conf. on Genetic Algorithms*, Morgan Kaufmann Publ., Inc., San Mateo, CA, (1991).
6. Grefenstette, J. J., "A User's Guide to GENESIS", Version 5.0, direct distribution: gref@aic.nrl.navy.mil (1990).
7. Baker, J. E., "An Analysis of the Effects of Selection in Genetic Algorithms", Ph.D. thesis, Computer Science Dept., Vanderbilt Univ., Nashville, TN (1989).
8. *3-D Laser Imaging System User's Guide*, Odetics, Anaheim, CA (1990).
9. *LASAR Hardware Manual*, Part Number #009-0064, Revision B, Perceptron, Inc. (1993).
10. Storjohann, K. and E. L. Saltzen, "Laser Range Camera Calibration", ORNL/TM-11723 (1991).
11. Baker, J. E., "Image Accuracy and Representational Enhancement Through Low-Level, Multi-Sensor Integration Techniques", *Proc. Sensor Fusion and Aerospace Applications, SPIE Proc.* 1956, pp. 170-186 (1993).

DATE

FILMED

6 / 27 / 94

END