



(19) **United States**

(12) **Patent Application Publication**
Dimsdale et al.

(10) **Pub. No.: US 2003/0001835 A1**

(43) **Pub. Date: Jan. 2, 2003**

(54) **INTEGRATED SYSTEM FOR QUICKLY AND ACCURATELY IMAGING AND MODELING THREE-DIMENSIONAL OBJECTS**

PCT/US97/06793, filed on Apr. 24, 1997, which is a continuation-in-part of application No. 08/638,961, filed on Apr. 24, 1996, now Pat. No. 5,988,862.

(76) Inventors: **Jerry Dimsdale**, Berkeley, CA (US);
Jonathan Apollo Kung, Berkeley, CA (US)

Publication Classification

(51) **Int. Cl.⁷ G06T 15/00**

(52) **U.S. Cl. 345/419**

Correspondence Address:
STALLMAN & POLLOCK LLP
Suite 290
121 Spear Street
San Francisco, CA 94105 (US)

(57) **ABSTRACT**


An integrated system generates a model of a three-dimensional object. A scanning laser device scans the three-dimensional object and generates a point cloud. The points of the point cloud each indicate a location of a corresponding point on a surface of the object. A first model is generated, responsive to the point cloud, that generates a first model representing constituent geometric shapes of the object. A data file is generated, responsive to the first model, that can be inputted to a computer-aided design system.

(21) Appl. No.: **10/079,007**

(22) Filed: **Feb. 20, 2002**

Related U.S. Application Data

(60) Division of application No. 09/177,777, filed on Oct. 23, 1998, which is a division of application No.






Fig. 1

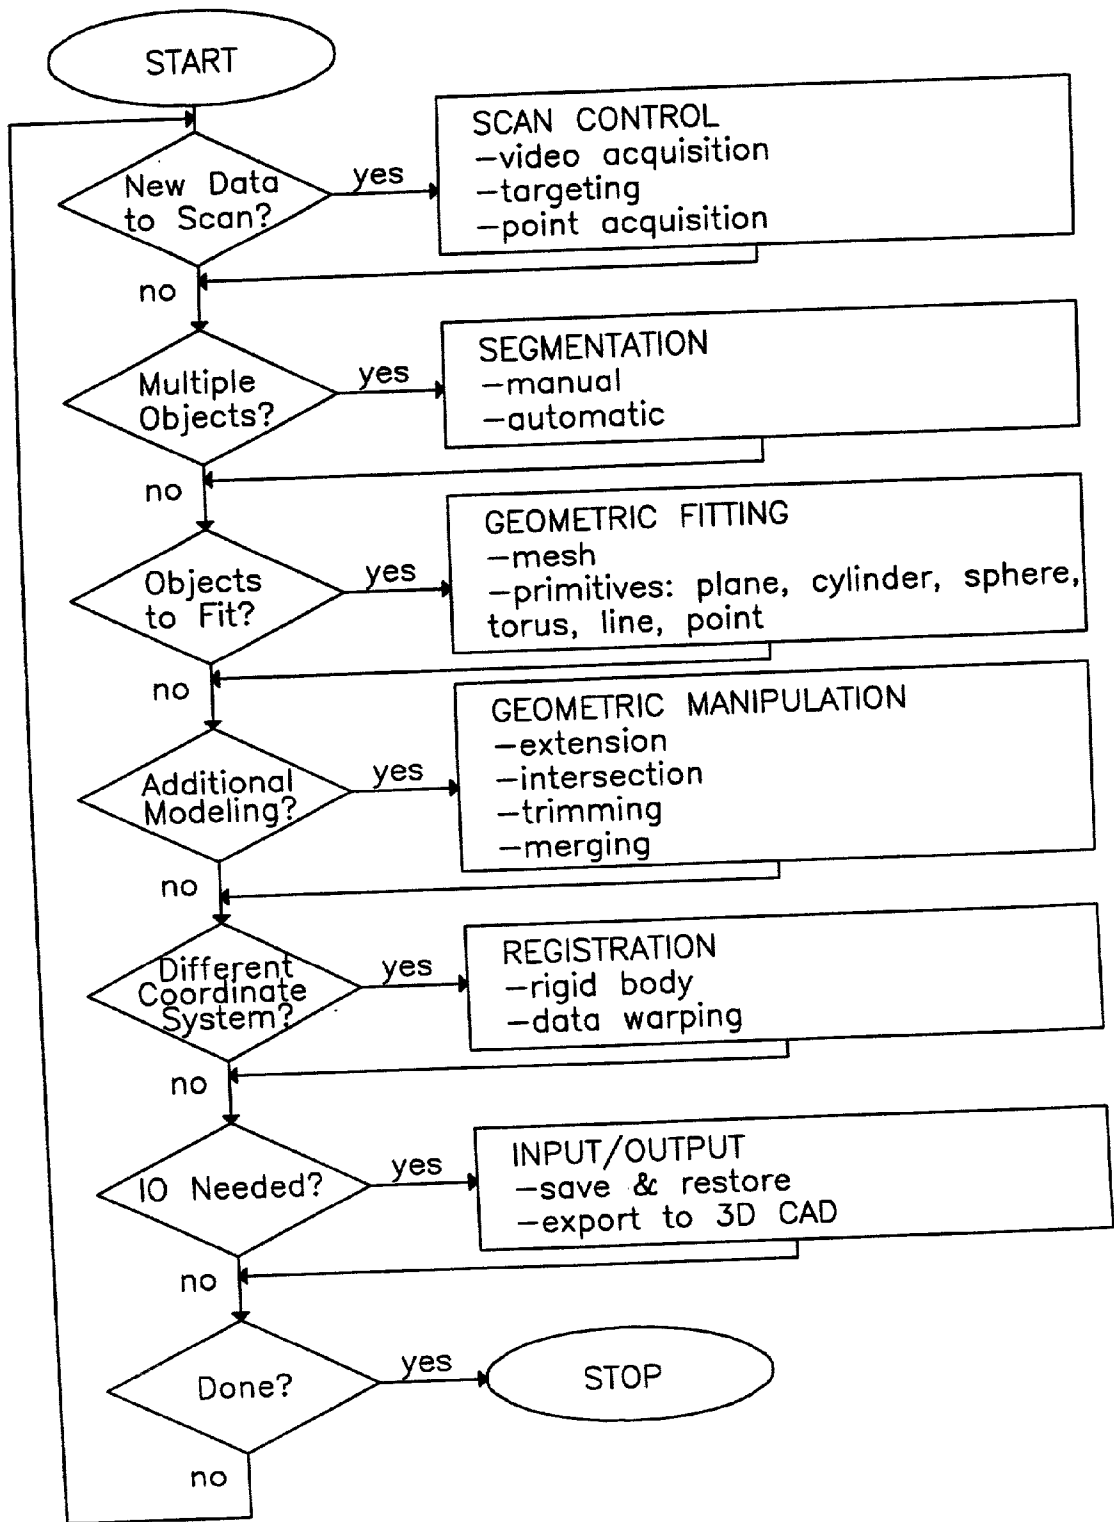


Fig. 1A




Fig. 2




Fig. 3A




Fig. 3B




Fig. 4A




Fig. 4B




Fig. 5




Fig. 6




Fig. 6A




Fig. 7




Fig. 7A




Fig. 8




Fig. 8A




Fig. 9

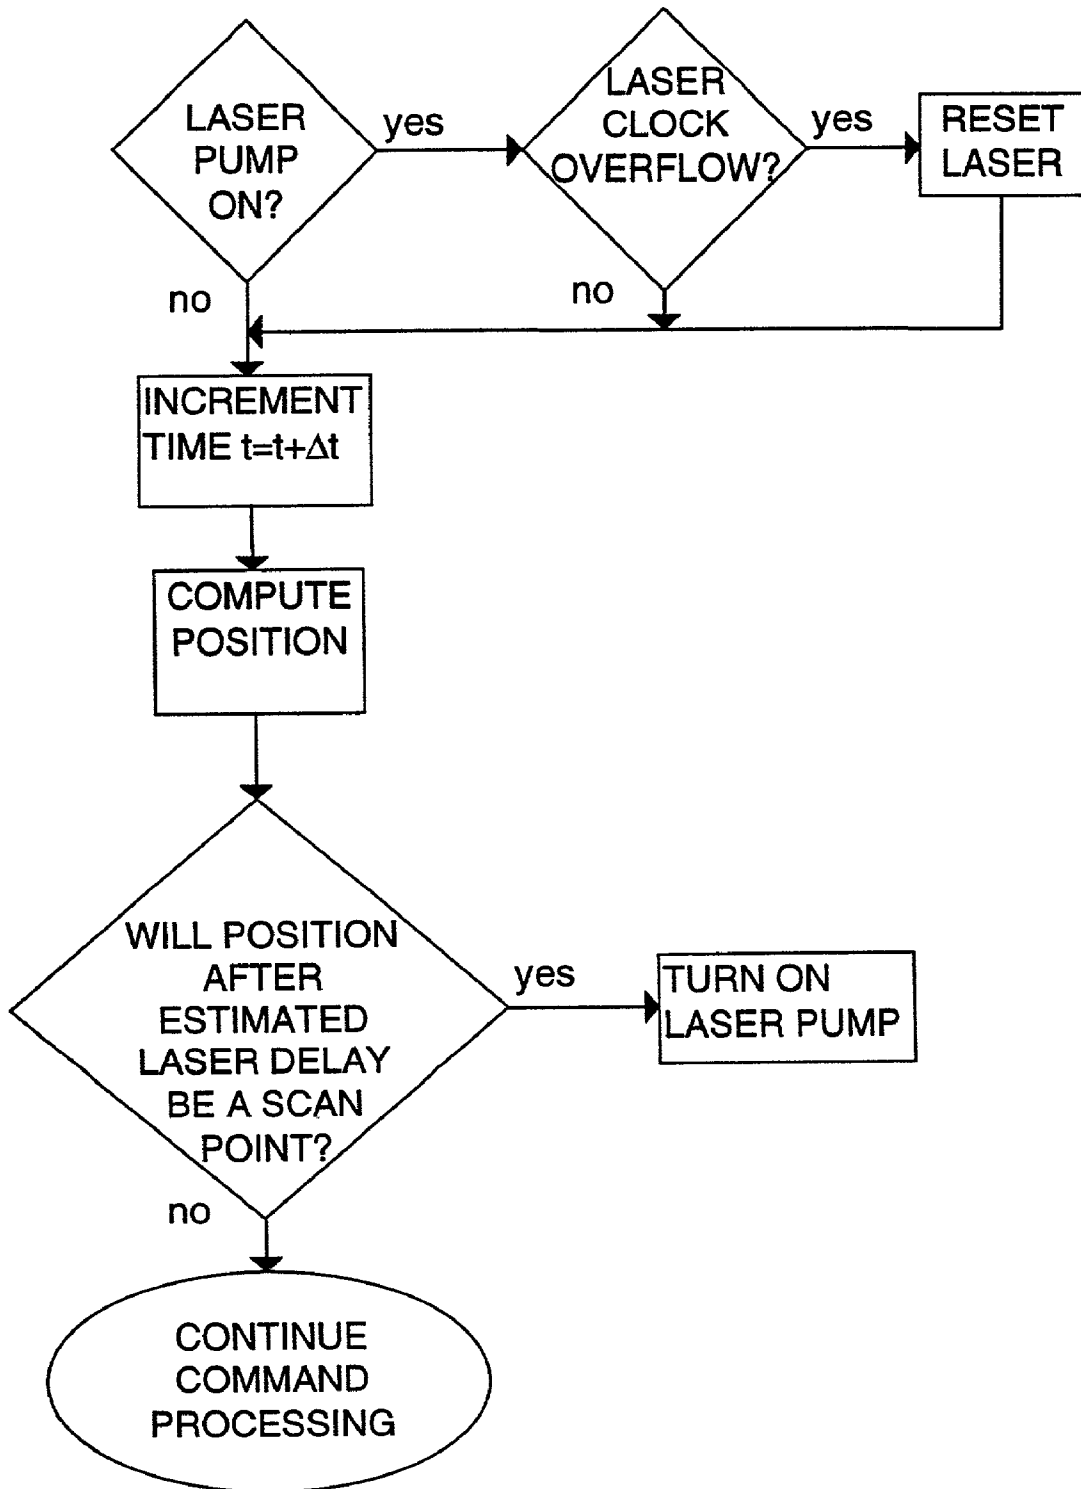


Fig. 10




Fig. 11A




Fig. 11B




Fig. 12




Fig. 13




Fig. 14




Fig. 15




Fig. 16A




Fig. 16B




Fig. 17A




Fig. 17B




Fig. 18




Fig. 19




Fig. 20




Fig. 21




Fig. 22




Fig. 23




Fig. 24

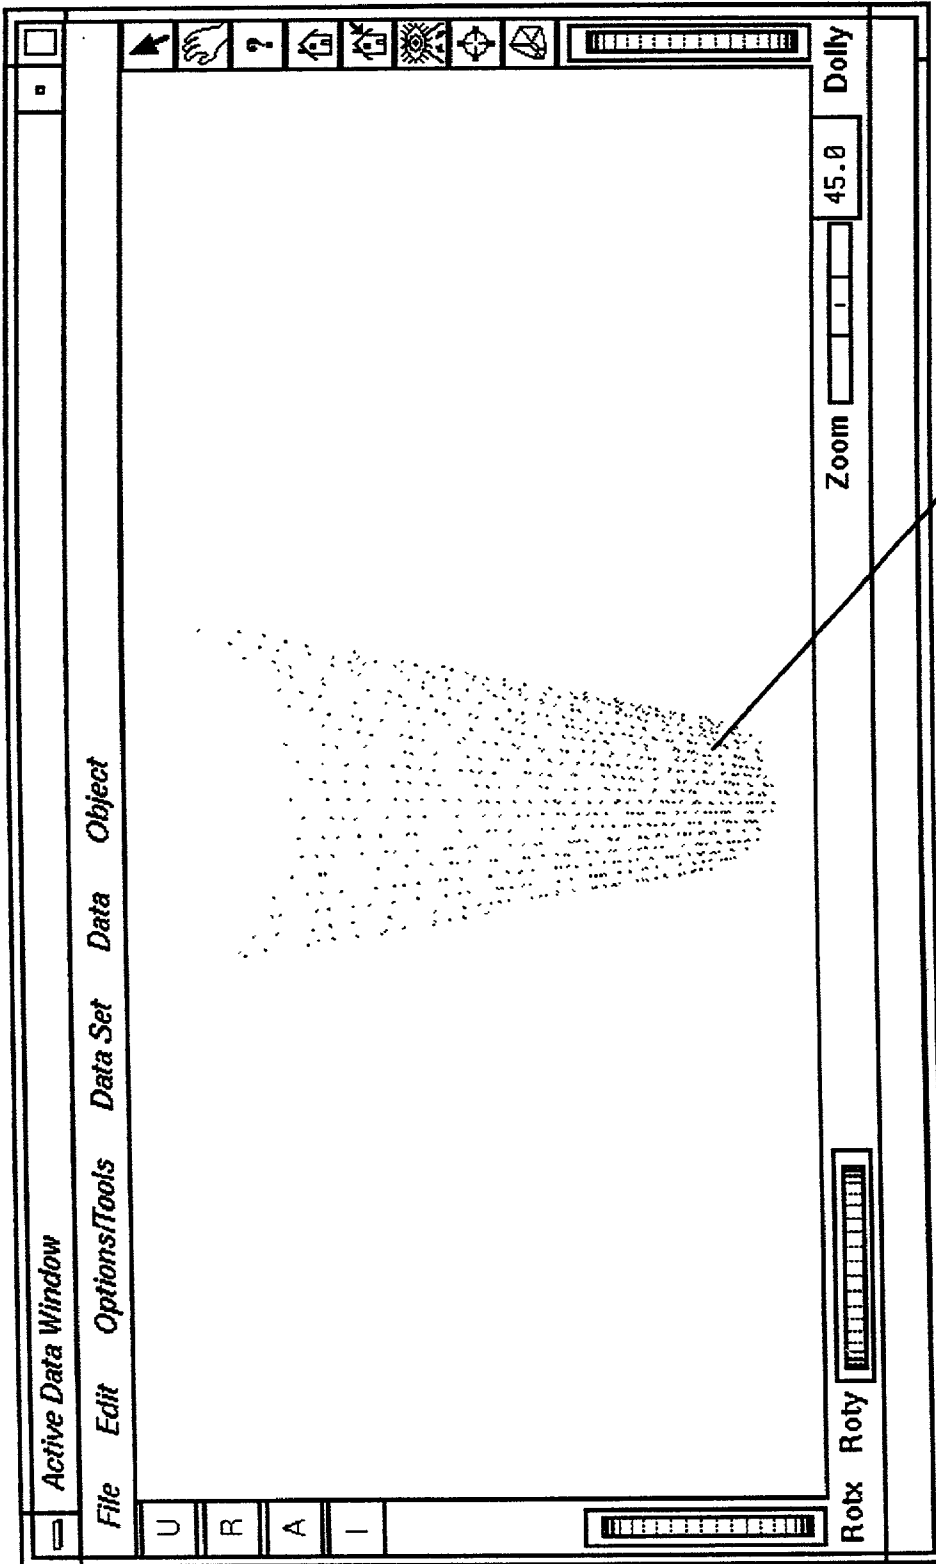


Fig. 25


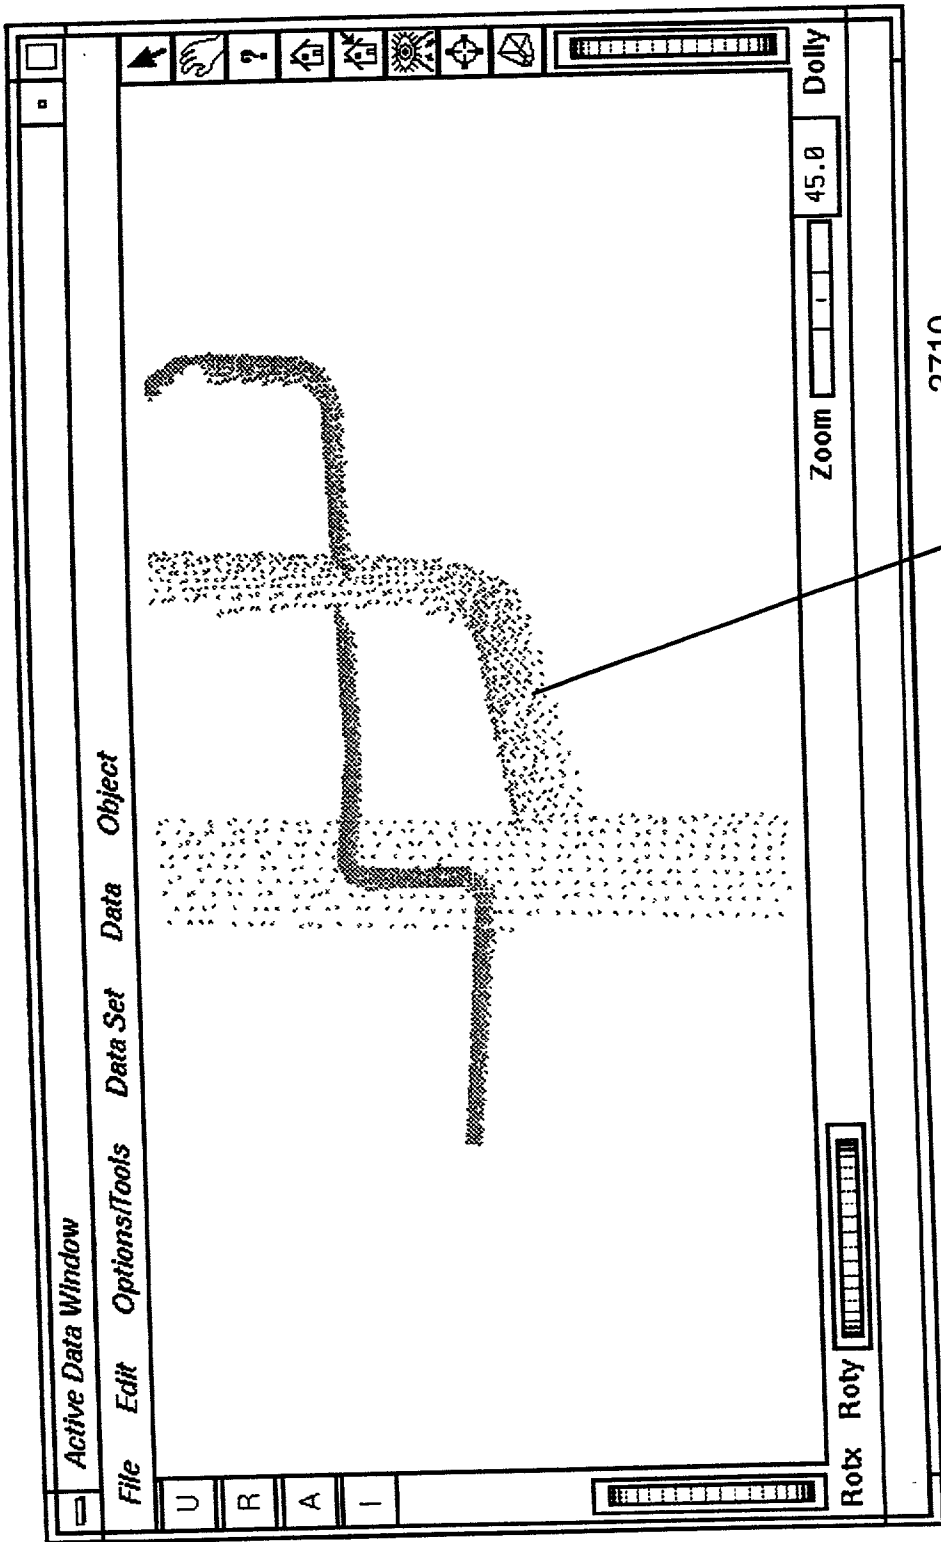


Fig. 26



2710

Fig. 27




Fig. 28




Fig. 29




Fig. 30




Fig. 31




Fig. 32




Fig. 33




Fig. 34




Fig. 35




Fig. 36




Fig. 37




Fig. 38




Fig. 39




Fig. 40

INTEGRATED SYSTEM FOR QUICKLY AND ACCURATELY IMAGING AND MODELING THREE-DIMENSIONAL OBJECTS

TECHNICAL FIELD

[0001] The present invention relates generally to systems that document the geometry and other attributes of objects in three dimensions and, specifically, to a system that employs a scanning lidar (range finding laser) to quickly and accurately sense the position in three-dimensional space of selected points on the surface of an object to generate a point cloud which represents the sensed positions of the selected points; that recognizes geometric shapes represented by groups of points in the point cloud, and that generates a model that represents these geometric shapes. The model may be transformed into a further model usable by computer-aided design (CAD) tools, including conventional CAD tools.

BACKGROUND

[0002] Mapping the geometry (shape, dimensions and location) and other attributes (e.g., color, texture and reflectance intensity) of complex real objects (whether small components such as small mechanical parts or large objects such as buildings and sites) has conventionally been a tedious and time consuming process. That is, such measurement have traditionally been performed manually. In addition, transforming these measurements into drawings or computer models required manual drafting or input into a CAD system for the production of the drawing or computer models.

[0003] Recently innovations have endeavored to simplify this process, but all have fallen short of achieving full integration, automation, precision, speed and range. For example, in the building industry, mapping a structure conventionally requires three basic steps:

- [0004] 1. Field data gathering
- [0005] 2. Data reduction and preparation
- [0006] 3. Drafting and CAD

[0007] The field data gathering step is performed by a team of surveyors who manually measure and record dimensions of pertinent components of the structure such as walls, ceilings, beams, columns, doors, windows, fixtures, pipes, conduits and equipment. The surveyors attempt to determine the geometry of the components as well as the relative location of the components in the structure. The surveyors recorded the data in a field notebook. The field-collected data is then organized and reduced to tables and organized sketches, and a CAD operator or drafter utilizes these tables to generate final drawings or models.

[0008] This process is labor intensive, time consuming, and error prone. In addition, using traditional surveying methods, the number of points which can actually be measured is very limited, due to the high cost of acquiring each point in terms of time and effort. Furthermore, if it is desired to acquire color, texture and other attribute information, additional field notes must be taken (e.g., still photographs and video).

[0009] Recently, the field step has been somewhat automated by using a laser ranging device built into or mounted

on an electronic theodolite. Precision reflection targets (retro reflectors) are placed at the locations of the object for which measurements are desired. Then, the laser ranging device obtains a precise measurement of the distance between the instrument and the target, which the theodolite provides an accurate indication of the horizontal and vertical angle offsets to the point relative to a given coordinate system. The distance and angle data are either recorded automatically on a magnetic device connected to the instrument or are reduced within the instrument to Cartesian coordinates relative to the instrument axes. This procedure is then repeated as many times as necessary to map a desired number of points of the object. The collected coordinates data can then be plotted directly on a CAD system.

[0010] Unfortunately, the plot is of little practical use since it does not indicate the object geometry. Moreover, because of the requirement for retro reflectors which must be manually placed, and because of the relatively long time per reading required by the laser range finder, the gathering of sufficient points to describe most objects is very labor intensive, time consuming and error prone.

[0011] Another known field gathering data process employs stereo photography and aerial photogrammetry. That is, stereoscopic images are taken of the objects and the resulting stereo photographs are registered either manually or using computerized techniques to reproduce the relative location of the camera picture plane location at the time each photograph was taken. The data reduction and preparation step is performed manually by a specially trained operator. Specifically, with the aid of specially mounted stereoscopic viewing lenses, the operator digitizes the coordinates of a sufficient number of points to allow the definition of the objects using the stereo photographs. Again, the digitized data is input into a CAD system or is manually drawn on paper.

SUMMARY

[0012] The present invention is an integrated system for generating a model of a three-dimensional object. A scanning laser device scans the three-dimensional object and generates a point cloud. The points of the point cloud each indicate a location of a corresponding point on a surface of the object. A first model is generated, responsive to the point cloud, representing constituent geometric shapes of the object. A data file is generated, responsive to the first model, that can be inputted to a computer-aided design system.

[0013] The subject invention further includes a method of controlling the timing of output pulse from a laser for use in a device which requires scanning of the laser output, wherein each output pulse is generated in response to a pump pulse comprising the steps of: monitoring the time delay between the initiation of the pump pulses and the subsequent generation of the associated output pulses; predicting the time delay between the initiation of next pump pulse and the associated output pulse based on the monitored time delays and; initiating the next pump pulse at a time selected to insure the output pulse is-generated at a time to permit proper positioning of the laser output during the scan of the beam.

[0014] The present invention further includes a method of manually separating from a plurality of clouds of points, representing three-dimensional features in a scene, a subset

of the points that represents a desired feature in the scene, the method comprising: selecting all the point clouds that include at least some data points representing the desired feature; and changing a view of the clouds and drawing a polygonal lasso to refine a selected subset of points to be included in a point sub-cloud and repeating the refining as many times as required to obtain the desired sub-cloud.

[0015] The present invention further includes a method for automatically segmenting a scan field of a scene into subsets of points that represent different surfaces in the scene, comprising the steps of: separating the scan field into a depth grid that includes depth information for scanned points of surfaces in the scene and a normal grid that includes an estimate of a normal to scanned points of the surfaces; convolving the depth information of the depth grid to generate a depth rating image whose values represent a gradient of depth change from one scanned point to another scanned point in the scene; convolving the components of the normal grid to generate a scalar value for each component for each point of the normal grid; for each point of the normal grid, determining from the scalar values for the components of that particular point a gradient of the normal at that point, wherein the gradients determined for the points of the normal grid collectively constitute a normal rating image; converting the depth rating image to a binary depth image using a recursive thresholding technique; converting the normal rating image to a binary normal image using a recursive thresholding technique; combining the binary depth image and the binary normal image to determine a single edge image; and grouping subsets of non-edge points as belonging to corresponding surfaces of the scene.

[0016] The method can further include the steps of determining the type of geometric primitive that would best first each group of points; fitting the geometric primitive to the data points; and intersecting adjacent planar regions in the scene.

[0017] The subject matter further includes a method for fitting a point cloud representing a corner, comprising: determining a fit of three planes to the points of the point cloud and creating the planes for a model; determining the three lines at the intersection of pairs of planes and creating the lines for the model; and determining the vertex point at the intersection of the three planes and creating a vertex point for the model.

[0018] The subject invention further includes a method for modeling a three-dimensional scene, comprising: generating a plurality of points that each represent a point on a surface of the scene; determining a best fit of a cylinder for a group of the points using surface normal estimates and global error minimization.

[0019] The subject invention further includes a method for modeling a three-dimensional scene, comprising: generating a plurality of points that each represent a point on a surface of the scene; determining a best fit of a cylinder for a group of the points using a quadric surface fit and global error minimization.

[0020] The subject invention further includes a method for modeling a three-dimensional scene, comprising: generating a plurality of points that each represent a point on a surface of the scene; determining a best fit of a sphere for a group of the points using a quadric surface fit and global error minimization.

[0021] The subject invention further includes a method for modeling three-dimensional scene, comprising: generating a plurality of points that each represent a point on a surface of the scene; determining a best fit quadric surface for a group of points; and determining which geometric primitive of a plurality of the family described by the quadric surface best fits the group of points.

[0022] The subject invention further includes a method for merging two geometric primitives of the same type to form a single geometric primitive of the type, comprising: creating a new group of points by combining the points used to originally fit each of the two primitives; and fitting the new geometric primitive using any appropriate fitting technique and the newly generated point group with points from each of the original primitives.

[0023] The subject invention further includes a method of registering a first model, consisting of a plurality of points and geometric primitives and having a first coordinate system, with a second model, consisting of a plurality of points and geometric primitives and having a second coordinate system, comprising: identifying by a user common features of the first and second scenes; identifying a transformation between coordinate systems that is responsive to the identification; and transforming the objects of the second model so that they use the first coordinate system.

[0024] The subject invention further includes a method of warping, comprising: selecting one or more models represented by a plurality of point clouds and geometric primitives; specifying constraints on the locations of any number of points or geometric primitives; creating an artificial volume that surrounds the points and geometric primitives in each view and assigning mechanical material characteristics to the surrounding volume; computing a minimum energy configuration for the material in the surrounding volume in which the points or geometric primitives are embedded such that the configuration satisfies all applied constraints; and displacing the points and primitives in accordance with the computed minimum energy configuration of the surrounding volume of material. In the latter method, the constraints can be specified to eliminate closure errors.

[0025] The subject invention further includes an integrated system for generating a model of a three-dimensional scene, comprising: a scanning laser device that scans the three dimensional scene with pulsed laser beam wherein the pulses of light last less than 1 nanosecond with up to 0.2 μJ in each pulse and measures the time delay, with a resolution of 30 psec or less, between each emitted pulse and a corresponding pulse that is returned from a surface of the scene and wherein said scanning laser device further tracks and measures the angular orientation of the beam during the scan; and means for generating a point cloud based upon the measured time delays and angle measurements, the point cloud comprising a plurality of data points that each represents a location of a corresponding point on the surface.

[0026] The subject invention further includes a system for calibrating the measuring electronics in a device which requires monitoring the time of flight of the output pulses from a laser comprising: a single mode optical fiber with one end thereof positioned to receive the output pulses of the laser, said single mode optical fiber having a known length; a detector positioned at one of the ends of the fiber for monitoring when the pulses exit the fiber and generating a

signal in response thereto, said signal being passed through the measuring electronics; and a processor for calculating a theoretical length of the fiber based on the detection of the pulse exiting the fiber and comparing that -calculated length with known length of the fiber to calibrate the measuring electronics.

[0027] The optical fiber can include partial reflectors located at each end thereof so that for each laser pulse entering the fiber a train of pulses will exit the fiber and wherein said train of pulses are used to further calibrate the measuring electronics.

[0028] The system can further include delay measurement electronics and wherein the train of pulses have a fixed delay therebetween whereby the monitoring of the train of pulses can be used to calibrate the delay electronics.

[0029] The system can further include a means for varying the power of the pulses monitored by the detector and wherein said detector functions to generate a signal when the power of the detected light exceeds a predetermined threshold and wherein said processor functions to track the variation in the delay of the generation of the output signal by the detector as a function of the power of the output pulses, said processor further functioning to calibrate the measurement of the delay based on the measured power of successive pulses used for monitoring the time of flight.

[0030] The subject invention further includes an apparatus for obtaining position information about surface points of a three dimensional object comprising: a laser for generating an output beam; scanning system for moving the laser beam over the object; monitoring system for automatically measuring the range to the object based on the measurement of the reflection of the laser beam, said monitor system also tracking and measuring the angular position of the laser beam, said monitoring system having a positional accuracy for each point in three dimensional space equal to or better than six millimeters at one standard deviation over a range of up to 100 meters.

[0031] Each range measurement can be made in under 0.005 seconds. The laser can generate a pulsed output and the energy per pulse can be less than 0.2 micro joules and the average output power of the laser can be less than 1.0 milliwatts.

[0032] The subject invention further includes an apparatus for measuring the distance to an object comprising: a laser for generating a beam of output pulses; a monitoring system for measuring the distance to the object based on the reflection of the laser beam, said monitoring system having an accuracy equal to or better than 6 millimeters at one standard deviation over its entire range of up to 100 meters and wherein each measurement can be made in less than 0.005 seconds and wherein the laser has an energy per pulse of no more than 0.2 micro joules and an average power of no more than 1 milliwatt. If the object is provided with retro reflectors and where the range of operation is up to one mile.

[0033] The subject invention further includes an apparatus for acquiring three dimensional information from a remote object comprising: a scanning laser module for measuring position information of the object; a video module for capturing image information from the object; and a processor for rendering a model of the object which includes the position information and the image information.

[0034] The video image information can be collected in a spatially coincident manner with the measurement of position information. The video image information can be collected from points adjacent to the points where position information is obtained.

[0035] The subject invention further includes an apparatus for obtaining positional information about surface points of a three dimensional object comprising: a scanning module for measuring three dimensional position information about an object; a video module for capturing and displaying image information from the object; and a processor operating with the scanning and video modules and permitting the use of the image information captured by said video module to aid in targeting the scanning module. The processor can function to specify a portion of the object to be targeted by the scanning module by dragging the image of an outline over the video image of the area to be targeted.

[0036] The subject invention further includes an apparatus for obtaining positional information about surface points of a three dimensional object comprising: a scanning module for measuring three dimensional position information about an object; a video module for displaying image information obtained from the scanning module; a processor operating with the scanning and video modules and permitting the use of the image information displayed by said video module to further refine the targeting of the scanning module.

[0037] The subject invention further includes an apparatus for obtaining positional information about surface points of a three dimensional object comprising: a scanning module for measuring three dimensional position information about an object, said scanning module including a laser for emitting a beam of visible radiation; and a processor for controlling the scanning module and wherein said laser can be manually positioned so that the visible beam will target the portion of the object to be scanned in response to a control signal from the processor.

[0038] The subject invention further includes a system for calibrating the measuring electronics in a device which requires monitoring frequency changes in a light beam generated by a laser used to measure distance to an object, wherein said beam has frequency chirp imposed thereon comprising a single mode optical fiber with one end thereof positioned to receive light from the laser; a detector positioned to receive light which has traveled through and exited the fiber in combination which light from the laser which has not traveled through the fiber, said detector for monitoring the changes in the frequency of the combined beam; and processor for determining the linearity of the chirp on the beam based on uniformity of the frequency changes measured by the detector and using the result to calibrate the measuring electronics.

[0039] The fiber can have a known length and includes a partial reflector on said one end and at least a partial reflector on the other end, and wherein light reflected from said one end of the fiber which has not traveled in the fiber is measured by the detector and wherein the processor further functions to calculate a theoretical length of the fiber based on the frequency changes measured by the detector and compares that calculated length with the known length of the fiber to calibrate the measuring electronics.

BRIEF DESCRIPTION OF THE FIGURES

- [0040] FIG. 1 is a block diagram of a system in accordance with an embodiment of the invention.
- [0041] FIG. 1A shows the overall flow of how one may use an embodiment of the invention to scan an object, organize acquired points, fit geometric shapes to the organized point, manipulate the fitted geometric shapes, and display the resulting manipulated geometric shapes.
- [0042] FIG. 2 is a more detailed block diagram of the system of FIG. 1.
- [0043] FIGS. 3A and 3B show the physical arrangement of the FDV of the FIG. 1 system, and also shows how the FDV is coupled to the tripod by a fork mount.
- [0044] FIGS. 4A and 4B show an example coordinate system relative to the FDV of the FIG. 1 system.
- [0045] FIG. 5 is a block diagram of one embodiment of an FDV in accordance with the invention.
- [0046] FIG. 6 is a block diagram of the optical transceiver of the FIG. 5 FDV.
- [0047] FIG. 6A shows a dual mirror arrangement of the scanner shown in FIG. 6.
- [0048] FIG. 7 is a block diagram which shows an embodiment of the laser.
- [0049] FIG. 7A is a block diagram of an embodiment of the beam expander shown in FIG. 6.
- [0050] FIG. 8 shows an embodiment of the duplexer.
- [0051] FIG. 8A shows a partially-reflecting duplexer.
- [0052] FIG. 9 shows an embodiment of the window of the FIG. 8 duplexer.
- [0053] FIG. 10 is a flowchart that shows calculations performed by the FDV DSP.
- [0054] FIGS. 11A and 11B show a unidirectional scan pattern and a bi-directional scan pattern, respectively.
- [0055] FIG. 12 is a block diagram of an embodiment of the FDV processor.
- [0056] FIG. 13 is a block diagram of example circuitry for determining a desired position of an FDV mirror.
- [0057] FIG. 14 is a block diagram of an example signal conditioning and energy integration circuit of the timing circuit shown in FIG. 12.
- [0058] FIG. 15 is a detailed block diagram of the system of FIG. 1.
- [0059] FIGS. 16A and 16B show two windows used to operate the CGP.
- [0060] FIGS. 17A and 17B show a targeting box and a point cloud.
- [0061] FIG. 18 shows a point cloud from the surface of a horse sculpture.
- [0062] FIG. 19 shows the point cloud of FIG. 18 color mapped with the laser return intensities.
- [0063] FIG. 20 shows a cloud of points from a corner feature.
- [0064] FIG. 21 shows the cloud of points of FIG. 20 and a polygonal lasso used for manual segmentation.
- [0065] FIG. 22 shows the cloud of points of FIG. 20 segmented into four subgroups, three subgroups on the surfaces of planes and a subgroup of edge points that are not part of the plane.
- [0066] FIG. 23 shows the cloud of points of FIG. 20 rendered as a triangulated mesh.
- [0067] FIG. 24 shows the corner feature of FIG. 20 with planes fit to the groups of cloud points.
- [0068] FIG. 25 shows a point cloud from the surface of a cylinder.
- [0069] FIG. 26 shows a cylinder primitive that was fit to the points shown in FIG. 25.
- [0070] FIG. 27 shows a cloud of points from the surfaces on a piping system.
- [0071] FIG. 28 shows cylinder primitives that were fit to the points shown in FIG. 27.
- [0072] FIG. 29 shows the completed piping model, after extending pipes and adding elbows.
- [0073] FIG. 30 shows the result of a corner fit, giving three planes, three lines, and a vertex.
- [0074] FIG. 31 shows a cylinder primitive in a scene.
- [0075] FIG. 32 shows the cylinder form FIG. 31 extended to meet adjacent objects.
- [0076] FIG. 33 shows a cloud of points from the surface from a variety of objects.
- [0077] FIG. 34 shows a model containing primitives that were fit to the points shown in FIG. 33.
- [0078] FIG. 35 shows configuration of a frequency adjustable laser.
- [0079] FIG. 36 shows block diagram of conventional FM chirp lidar.
- [0080] FIG. 37 shows block diagram of self-calibrating FM chirp lidar.
- [0081] FIG. 38 illustrates the relative timing at which a large and a small pulse cross a predetermined threshold.
- [0082] FIG. 39 illustrates one circuit for measuring pulse energy.
- [0083] FIG. 40 illustrates another circuit for measuring pulse energy.

DETAILED DESCRIPTION

[0084] A. Overview

[0085] 1. Overall System

[0086] FIG. 1 is a block diagram that illustrates the invention in its broadest aspect. Referring to FIG. 1, a Field Digital Vision (FDV) module 10 includes a scanning sensor for scanning an object 20 and for sensing the position in three-dimensional space of selected points on the surface of the object 20. The FDV module 10 generates a point cloud 30 which represents the sensed positions of the selected

points. The point cloud **30** also represents other attributes of the sensed positions, such as reflectivity, surface color and texture.

[0087] A Computer Graphics Perception (CGP) module **40** interacts with the FDV to provide control and targeting functions for the FDV module **10** sensor. In addition, using the point cloud, the CGP module **40** recognizes geometric shapes represented by groups of points in the point cloud **30**, and the CGP module generates a CGP model **42** that represents these geometric shapes. From the CGP model **42**, the CGP module **40** generates a further model usable by computer-aided design (CAD) tools **50**. The CAD tools may be conventional.

[0088] FIG. 1A shows the overall flow of how one may use an embodiment of the invention to scan an object, organize acquired points, fit geometric shapes to the organized point, manipulate the fitted geometric shapes, and display the resulting manipulated geometric shapes.

[0089] 2. FDVModule Overview

[0090] Referring to FIG. 2, the FDV **10** includes a scanning laser system (lidar) **210** that scans points of the object **20** and that generates a lidar data signal that precisely represents the position in three-dimensional space of each scanned point. The lidar data signal for groups of scanned points collectively constitute the point cloud **30**. In addition, a video system **220**, preferably including both wide angle and narrow angle CCD cameras, is provided. The wide angle CCD camera of the video system **220** acquires a video image of the object **20** and provides, to the CGP **40** via a control/interface module **230** of the FDV **10**, a signal that represents the acquired video image.

[0091] In response to user input relative to the signal that represents the acquired video image, the CGP **40** provides a scanning control signal to the lidar **210**, via the control/interface module **230**, for controlling which points on the surface of the object **20** the lidar **210** scans. More particularly, the scanning control signal provided from the CGP **40** controls an accurate and repeatable beam steering mechanism to steer a laser beam of the lidar **210**.

[0092] In addition, the narrow angle CCD camera of the video system **220** captures the texture and color information, and provides this captured information to the CGP **40**.

[0093] 3. CGP Module Overview

[0094] Referring still to FIG. 2, the CGP **40** is constituted of a data processing system (e.g., a notebook computer or a graphics workstation) and special purpose software that when executed config.s the CGP **40** data processing system to perform the FDV **10** control and targeting functions, and also to perform the CGP model **42** generation functions.

[0095] i. FDV Control

[0096] The CGP **40** controls the scanning lidar **210** of the FDV **10** by providing a lidar control signal to the FDV **10** that controls which points of the object **20** the FDV **10** scans. User input is provided to the CGP **40**, which defines what portions of the object **20** to scan, and at what resolution.

[0097] ii. Model Generation

[0098] Each data point in the point cloud **30** generated by the FDV **10** represents both distance to a corresponding laser

impingement point from an FDV **10** "origin point" and the angle from the origin point to the laser impingement point. The CGP software config.s the CGP **40** computer to process the data points of the point cloud **30**, generated by the lidar **210** as a result of scanning the object **20**, to display and visualize the scanned portions of the object **20**. More specifically, the CGP software config.s the CGP **40** computer to recognize geometric shapes in the object **20** ("graphic perception") and, using these recognized geometric shapes, to perform geometry construction, 3D model construction, 3D visualization, and database functions for automated acquisition or manual input of object attributes, generation of plans, sections, and dimensions, data query, and CAD interfaces, and networking options.

[0099] B. Details

[0100] 1. FDV Module Detail

[0101] FIG. 5 is a block diagram of one embodiment of an FDV **10** in accordance with the invention. A lidar transceiver **502** includes a laser, transmit optics, receive optics and detector for generating ranging and intensity data. A scanning system **504** includes dual orthogonal scanning mirrors, galvo motors, and encoders for steering the laser beam and determining the azimuth and altitude angles of the laser beam from the positions of the mirrors. A wide angle video system **506** generates targeting video information and a narrow angle video system **507** generates color and texture information. Control/interface circuitry **230** handles the exchange of data between the FDV **10** and the CGP **40**.

[0102] If the laser beam is quasi-CW, always on with either intensity modulation (AM) or wavelength modulation (FM), the distance to the object **20** can be inferred by any of a number of techniques involving demodulation at the transceiver **502**. If the laser is pulsed, the distance to the object **20** is usually measured by the time of flight from the transceiver **502** to the object **20** and back. Other laser modulation schemes may be used.

[0103] In the time-of-flight embodiment, the laser is preferably of the type disclosed in U.S. Pat. Nos. 5,132,977; 5,386,427; and 5,381,431, assigned to Massachusetts Institute of Technology. In particular, the beam generated by such a laser has special properties such as being capable of producing pulse widths less than Insec.

[0104] A particular embodiment of the laser which has been used is particularly suitable for precision lidar since:

[0105] 1. The short pulsewidth provides high accuracy, since radar theory shows that the accuracy is proportional to the inverse of the pulse time.

[0106] 2. The laser itself is physically quite small, especially useful for portable applications.

[0107] 3. It has a diffraction limited beam, which implies that the spot size at a distance is not limited by the properties of the light, but only by the quality of the optics used to collimate or focus it

[0108] 4. Since the wavelength is quite short (532 nm), and Rayleigh range is inversely proportional to wavelength, the beam can be kept small over a large distance interval. In fact with a 1 cm exit aperture, the beam will remain less than 6 mm over 50 m.

[0109] In one preferred embodiment, the laser beam is directed by the orthogonal scanner mirrors to a laser impingement point on the surface of the object 20. The range can be determined by any of a number of conventional lidar techniques. For example, the “time of flight” of a laser pulse to travel from the laser to the surface of the object 20, and then back to the detector, is determined. The range is determined based on the constant speed of light, with appropriate adjustments being made for atmospheric factors.

[0110] A system in accordance with the present invention can provide high ranging accuracy at high acquisition rates. For example, at 100 m ranges, a 1 mm accuracy can be achieved on a single shot basis, with anywhere from 1000 to 5000 data points being acquired per second.

[0111] In other embodiments, a chirp lidar may be employed. The essential component of a chirp lidar can be modulated with a linear change of wavelength over time. Thus, the wavelength of the light emitting from the laser will be given by $\lambda(t)=k(t-t_0)+\lambda_0$. In practice, such a laser is commonly manufactured by making a composite of two materials, a typical laser gain media such as NdYAG (3510), and substance which can have its index of refraction altered electrically such as Lithium Niobate (3520). (See FIG. 35) This effectively changes the length of the laser cavity, and hence alters the emitted wavelength. Commercially available lasers can be modulated up to about 100 GHz with a voltage modulation of about 1 kV, so the frequency of the light will be approximately $f(t)=k(t-t_0)+f_0$.

[0112] Referring to FIG. 36, in a typical FM chirp system, a portion of the light emitted by the laser 3610 is sampled and recombined at the beam splitter 3630 with the light returning from the target 3620. Since the light is delayed by the amount of time to contact the target and return, the light returning from the target will have a lower frequency than the light which is sampled from the laser. This difference will be apparent in the output of detector 3610 which measures the intensity of the combined beams. If the frequency ramp of light is exactly linear and the laser has a coherence length much greater than the distance to the target, combining the beams will produce a constant measured frequency from the detector which is proportional to the range: $f=k_1*d+d_0$. Chirped Yag lasers as described above have a coherence length of about 20 km, but the chirp is not linear, and this has severely limited the accuracy of existing FM chirp lidars.

[0113] Referring to FIG. 37, a huge improvement in accuracy can be realized by adding a system to calibrate every range measurement. A fiber is prepared which has a partial reflector 3771 on one end, and a nearly total reflector 3772 on the other. Now a portion of the light emitted by the laser 3710 is sampled and recombined at beam splitter 3740 with the light returning from the target 3720, and the intensity is measured by a detector 3760. An additional sample of the beam emitted from the laser is sampled by beam splitter 3730 and introduced into the fiber at the partially reflected end 3771. The beam propagates a fixed distance down the fiber and reflects off the end face and is recombined with the beam which is reflecting off the partially reflecting face 3771, and is measured with a second detector 3750. The linearity of the chirp is then measured by determining the deviation from a constant frequency of the output of detector 3750, and this information is used to

correct for the effects of the nonlinear chirp in the output of detector 3760, which corresponds to the target range measurement.

[0114] Referring to FIGS. 3A and 3B, in the embodiment, the FDV 10 is physically housed in a box 330 made of metal or other suitable housing material. The box 330 is suspended from its side panels by a fork mount mechanism 310. The fork mount system is supported on a turntable 340, and the turntable 340 is mountable on a tripod 320. Using the fork mechanism 310, the FDV 10 can be rotated horizontally (“azimuth rotation”) and vertically (“elevation tilt”, or “altitude”). Generally, a position of the tripod 320 is referred to as a “setting” or “positioning”; the rotation and tilt of the FDV 10 in the fork mount 310 is referred to as “pointing” or “orientation”. A “view” is generally associated with a given setting and orientation.

[0115] The fork mount system 310 preferably includes high accuracy azimuth and altitude rotation measuring devices (e.g., conventional theodolite-type optical or electronic encoders) to provide precise rotation and tilt data of the FDV 10. This feature can allow the automatic integration of scans taken from the same tripod 320 setting, but with a different orientation of the FDV 10. In the event these devices are not used, and for scans taken from different settings and orientations, these scans can be integrated using techniques described later in this disclosure.

[0116] It should be noted at this point that, while conventional surveying instruments should be leveled prior to operation in order to operate properly, this is not a requirement of the FDV 10. This is due to the novel methods of this invention as embodied since its own internal coordinate system and the procedures utilized in its software that take advantage of its method of acquiring position data. The system, however, does have the ability to be leveled and to be used in a manner similar to a traditional theodolite. The Cartesian coordinates of the FDV 10 in the example embodiment are shown in FIGS. 4A and 4B.

[0117] Referring still to FIGS. 3A and 3B, in one embodiment, two orthogonal mirrors of the FDV 10 provide a field of view of approximately 40° by 40° (“Field of View”, or “View” is defined as the maximum size of the area projected by the laser maximum deflections of the beam in degrees). The field of view can be increased or decreased by resizing the mirrors and certain parts of the optical train. The fork mount described above is utilized to allow pointing of the FDVs 40°×40° field of view anywhere over a projected sphere thus affording a wide range of flexibility in imaging large objects or groups of objects from the same setting. Other mounting methods may be used to accomplish the same purpose.

[0118] High accuracy and repeatability electronic encoders read the rotational angles of the orthogonal mirrors, and the readings of the mirror rotation angles are precisely timed to coincide with the range reading. Preferably, the system is Class II FDA eye safe. A first embodiment has 6 mm spatial accuracy over a range of ±50 m. In another embodiment, autofocus capability and 5-6 picosecond electronics are included, which extends the system’s range accuracy to ±1 mm and ±1 mm spatial accuracy over ±50 m. The range (and accuracy) of the system can be significantly influenced by the choice of eye safety classification selected, but these limitations are not inherent limitations of the invention itself.

[0119] The following is a description of the key components of preferred embodiment of the FDV 10. A block diagram of the optical transceiver 502 of the FDV 10 is shown in FIG. 6. The optical transceiver 502 transmits an optical pulse to a spot on object 20, and receives a reflected optical pulse from the object 20. Given the constant speed of light, the optical transceiver calibrates the distance to the spot on the target.

[0120] Referring to FIG. 6, the laser 602 fires an optical pulse which lasts less than 250 psec, in response to an external command provided from a laser controller 604. The laser 602 produces a pulse, preferably at a wavelength of 532 nm, within 100-300 μ sec after an external signal emanating from a digital signal processor which provides central control of real time events. The time delay is a complicated function of recent laser history and environmental conditions. This function is not completely known at present. However, a software algorithm, which is described elsewhere, is employed to estimate the time delay with adequate accuracy for the required measurements.

[0121] The laser beam output of the laser 602 is transmitted through a beam expander 606 that is focused to adjust the size of a light spot that will eventually impinge upon a point on the object 20. The focussed optical pulse is then transmitted through a duplexer 608, which is an optical system for aligning the outgoing optical path with the incoming optical path. The duplexer 608 directs a significant first portion of the light energy of the outgoing optical pulse to a spot on the object 20 via a scanner 614, but a second, much smaller portion, of the light energy of the outgoing optical pulse is directed to a receiver telescope 610. The portion of the outgoing optical pulse that propagates to the object 20 impinges on the spot on the object 20, and some of the energy of the optical pulse is reflected off the object 20 in a direction back to the duplexer 608. The returning optical pulse is directed by the duplexer 608 to a receiver telescope 610, which focuses the received energy onto a detector 612. The detector 612 converts the received optical pulse energy into electrical energy, and the output of the detector 612 is a series of electrical pulses, the first (which is generated by the detector 612 in response to the second, small portion, of the transmitted pulse not directed toward the object 20) occurs at a short fixed time (i.e., fixed by the length of the optical path through the beam expander 606, the duplexer 608 and the receiver telescope 610) and the second of which occurs as light energy returns from the object 20. Both the second, small portion of the transmitted pulse not directed toward the object 20, and the return optical pulse reflected from the spot on the object 20, are provided to the timing circuit 616 which calculates the time of flight to the spot on the object 20. The range to the spot on the object 20 can then be readily calculated from the calculated time of flight.

[0122] FIG. 7 is a block diagram which shows an embodiment of the laser 602. The heart of the laser system 702 is a conventional laser chip 702 that include two bonded crystals coated with antireflective dielectric coatings on their faces. The laser chip 602 is pumped with a solid state diode 704 operating at $808.5 \text{ nm} \pm 3 \text{ nm}$. The output frequency of the diode pump 704 is adjusted by changing its temperature with a thermoelectric cooler 706. The temperature of the diode pump 704 is measured with a thermistor 708, and the measured temperature is fed back into the diode power

supply 710. The required temperature varies with each individual diode, but it typically ranges from 20° to 30° C.

[0123] The output power of the diode pump 704 is typically 1 watt, launched into a $100 \mu\text{m}$ core glass fiber. When continuously pumped, the output of the crystal laser 602 is approximately 35 mW average power at $1.064 \mu\text{m}$, which corresponds to $2.4 \mu\text{J}$ pulses lasting about 280 psec at a repetition rate of 15 kHz. The multimode fiber is preferably terminated by an SMA905 solid brass connector, with the crystal of the laser chip 702 glued to one end of the connector with an optical resin. Thus ensures adequate thermal dissipation from the crystal of the laser chip 702, keeping the crystal 702 within the temperature range required for most efficient operation.

[0124] A piece of KTP frequency doubling crystal 712 is held within a few millimeters of the face of the laser chip crystal 702. This provides an ultimate output from the laser 602 having a 12 mW average power at 532 nm, which corresponds to $0.8 \mu\text{J}$ pulses lasting approximately 218 psec. This ultimate output from the laser 602 is nearly diffraction limited (i.e., one which has theoretically minimum divergence, given a specific wavelength and waist diameter), with an apparent waist diameter of $561 \mu\text{m}$.

[0125] Embodiments of the invention which meet FDA Class II eye safe system design specifications are potentially more commercially viable. In order to meet this specification, the maximum energy per pulse that can be transmitted at 532 nm is $0.2 \mu\text{J}$. With this restriction, the average power transmitted is largely dependent on the pulse repetition rate, and is given by the following table

CLASS	ENERGY PER PULSE	AVERAGE POWER	REPETITION RATE
I	.2 μJ	.39 μW	1.95 Hz
IIA	.2 μJ	3.9 μW	19.5 Hz
II	.2 μJ	1.0 mW	5 kHz
IIIA	.2 μJ	5.0 mW	25 kHz

[0126] In one embodiment of the invention, the beam expander 606 is entirely conventional (e.g., Melles Griot model number 09LBM013, 10x beam expander). The transceiver 502 has cross axis accuracy which is proportional to the size of the laser beam impinging on the intended target. The base design of 6 mm accuracy has a simple beam expander. The laser 602 can be collimated with a fixed 10x beam expander 606 which has an aperture of $<1 \text{ cm}$ to produce a beam whose $1/e^{2*power}$ beam width is less than 6 mm over a range of 50 m.

[0127] FIG. 7A shows a further embodiment 750 of the beam expander 606 that includes features which allow the system of the invention to measure ranges at an accuracy of approximately 1 mm at 50 m. This is because the impingement spot on the object 20 of the laser beam expanded by a conventional beam expander is collimated, and produces a spot of no more than 6 mm over a 50 m range. However, a beam can be focused through a 50 mm aperture to a spot of size no more than 1 mm over a 50 m range—but the spot will be much larger at other ranges. Thus the beam expander 750 of a system having 1 mm accuracy at 50 m includes a movable optical element 752 which can change the size of

the focused spot. Additionally, the beam expander **750** includes an adjustable aperture **755**, and means for controlling the adjustment, so that the distance, from the laser, over which the beam stays at 1 mm in diameter remains approximately constant. The minimum diameter spot produced with a diffraction limited lens of focal length f and diameter D is $d_0=2f\lambda/D$. The Rayleigh range of the focused spot, which is the depth of focus of the beam, is given by $b=2\pi\omega_0^2/\lambda=2\pi f^2\lambda/D^2$. Thus, if f/D is held constant, the depth of focus will not be a function of the range of the focused spot, f .

[0128] As the beam focus is changed, the elements should stay sufficiently aligned so as to prevent the beam from changing direction by more than a fraction of 1 mm at 50 m, or this will appear as an error in the placement of the point in space. In order to minimize this beam wander, a linear servo motor **754** (see FIG. 7A) is employed for controlling the position of the focusing mechanism, and a transducer provides position feedback. The lens **752** is mounted in an annular ring **753**, which prevents it from rotating or misaligning while it is being translated.

[0129] Duplexer

[0130] An embodiment of the duplexer **608** is shown in FIG. 8. The optical system of the duplexer **608** is configured such that the outgoing beam from the beam expander **606** to the scanner **504** is coaxial with the return beam reflected from the object **20**. Thus, only one scanner **504** need be provided. In the embodiment of the duplexer shown in FIG. 8, a window **802** is provided, with a 50% beamsplitter **804** attached over the window **802**. When an optical pulse is transmitted from the laser **602** and through the beam expander **606**, the pulse impinges upon the beam splitter **804**. Most of the light energy of the pulse is reflected off the beam splitter **804** and is passed to the scanner **504**, but some of the optical pulse proceeds through the beam splitter **804** and impinges upon a low reflectance beam block **806**. Due to the low (but non-zero) reflectance of the beam block **806**, a small fraction of the optical pulse hitting the beam block **806** returns to the beam splitter **804** and is reflected into the receiver **610**.

[0131] Moreover, as an optical pulse returns from the object **20**, since only the central portion of the return pulse is obscured by the prism **804**, most of the light impinging on the window **802** makes its way to the receiver **610**.

[0132] Partially Reflecting Duplexer

[0133] Referring now to FIG. 8A, for the 1 mm accuracy embodiment, a partially-reflecting duplexer **850** is employed. With this duplexer, a fraction of the light pulse provided from the beam expander into a beam stop **652** and reflects off the duplexer window **850** to the receiver telescope **610**. The remainder of the light pulse proceeds to the object **20**. Most of the return light pulse from the object **20** continues on through the window **850** and is collected by the receiver telescope **610**. The window **850** is AR coated on the receiver side, and partially mirrored on the laser side. The entire window **850** is used to steer the outgoing beam, since a 50 mm aperture is required to focus the spot to 1 mm at 50 m. The partial reflectance is chosen in view of the laser transmission power and the applicable eye-safe classification level. For example, if the laser transmission power is four times the allowable level of the applicable eye-safe level, then the partial mirroring is chosen to reflect 25% and absorb 75%.

[0134] Referring now to FIG. 9, in the 6 mm embodiment, improved efficiency can be achieved in collecting the return optical pulse if only the center of the window **802** is coated **904** to reflect the outgoing pulse, and the remainder of the window **802** is anti-reflective coated **906**. In this way, the return optical pulse is not reflected out of the receiver by the part of the window **802** that is antireflection coated **906**.

[0135] Preferably, the laser **602** emits a strongly polarized beam so that the reflective coating **904** can be optimized to have slightly different reflection coefficients for the two planar polarizations (20%-S and 30%-P). In such an embodiment, the power of the beam impinged onto the object **20** can be fine tuned merely by physically rotating the laser body.

[0136] Receiver Telescope

[0137] Referring again to FIG. 6, after the returning pulse has passed through the duplexer **608**, it is collected by the receiver telescope **610**, which optimizes the amount of signal provided to the detector **612**. The receiver telescope **610** may be a simple 50 mm aperture lens. The lens is preferably selected so that the variation in pulse energy entering the detector **612** does not change as a function of the distance to the object **20** over the range of distances for which the instrument is designed. A multiple element lens can be designed to minimize the variation in received pulse energy as a function of range somewhat more effectively than a single element lens. That is, at the greatest expected distance, the focal length of the lens is such that all the incoming light, which is effectively collimated since it is generated by a point source in the far field, is focused to completely fill the detector **612**. As the object **20** becomes closer to the telescope **610**, the spot of return light becomes larger than the detector **612**. The power incident on the detector **612** increases as the square of the distance from the telescope **610** to the object **20**, up to the maximum expected distance. Moreover, the power returning from the object **20** decreases as the square of the distance from the telescope **610** to the object **20**. Thus, in practice, these two effects approximately cancel each other. This minimizes the variation in optical power incident on the detector **612** over the range of anticipated use. In the 1 mm option, the receiver optics can be improved in some cases by using a two element, adjustable focus, Newtonian telescope (e.g., similar to the 1 mm beam expander).

[0138] Detector

[0139] The detector **612** converts optical pulses to electrical pulses which can be processed by the elapsed time measurement electronics (timing circuit **616**). In one embodiment, the detector **612** is an avalanche photodiode (APD) with greater than 1 GHz electrical bandwidth. In addition to the time between the start and any stop pulses, the intensities of all the pulses are recorded. The intensity information is used to make a correction to the range measurement derived from the timing information.

[0140] Scanner

[0141] The scanner **504** may be conventional. The scanner **504** directs the outgoing pulses from the duplexer **608** to a desired position on the object **20** and directs the incoming return pulse into the receiver telescope **610**. The scanner **504** directs light to the narrow field ccd camera **507** to collect color and texture in the immediate vicinity of the scanned

laser points, which provides for precise registration of color and texture obtained with the lidar acquired point geometry. In one embodiment, the scanner **504** includes a dual mirror arrangement (see **FIG. 6A**) for beam steering, although any conventional high accuracy and repeatability beam steering mechanism may be employed. The dual mirror arrangement includes two mirrors which are rotated on orthogonal axes by moving coil motors. These motors have an integral position decoder which has angular repeatability of less than 1 microradian. The mount for the scanners is integrally formed with the supports for the laser and other optics. This system provides 40 degrees of optical motion in both altitude (elevation) and azimuth at several Hertz.

[0142] Electronics

[0143] A. Timing Circuit

[0144] Another embodiment of the scanner **504** mechanism consists of a single mirror rotating about a central axis, mounted on a rotating turret. In this configuration, the physical coordinate system would be spherical, with the faster (due to the less inertia) mirror providing the elevational angle and the more slowly rotating turret providing azimuthal motion. A system such as this could provide a field of view of more than 90 degrees in a vertical plane and a full 360 degrees in a horizontal plane (both planes being relative to some chosen scanner coordinate system).

[0145] Electronics

[0146] Ranging Electronics

[0147] The function of the ranging electronics is to compute the range from the FDV **10** to the object **20** based upon the electrical output of the detector **612**. Several possible methods may be used, including a demodulator in the case of a quasi-CW modulated laser system. For the preferred time of flight embodiment, an interval timer (timing circuit) measures the relative time interval between an initial (start) pulse reflected directly into the receiver **610** by the duplexor **608**, and the pulse reflected off of the object **20** back into the receiver **610**.

[0148] Reflectivity Electronics

[0149] In many cases, it is useful to know not only the position in space of a point on the object **20**, but also know the reflectivity (at some particular wavelength) of that point. The reflectivity electronics measure the amount of light reflected from the object **20** into the receiver **610** and detector **612**. This data can be used to provide corrections to the range information as well as the information on the material and/or finish of the surface of the object **20**.

[0150] Digital Signal Processor

[0151] A digital signal processor integrated circuit controls all the time critical is functions of the FDV—scanner control, laser firing. It also provides fast floating point computation capability for making geometry corrections, calibration corrections, and video lens corrections, and video compression. The digital signal processor is interrupted at regular time intervals, typically about 10 usec. At each of these time intervals, a check is made to see what real time calculations are outstanding.

[0152] Scanner Control

[0153] The electronics for the scanner are a simple precision PD controller which are driven by a digital signal from the DSP. When driving this system quickly, there is noticeable lag in the ability of the scanner to follow the driving signal. However, the controller circuit does not have an error signal output. An external precision analog differential amplifier provides an error signal (the difference between the command signal and the actual displacement), which is sampled by the DSP at low resolution. The DSP then computes the exact scan position by computing the sum of the command signal and the error signal. The advantage of this method is that it requires only a low resolution A/D converter and a precision D/A converter, rather than a far more expensive precision A/D.

[0154] The digital signal processor generates the trajectories for the analog scanner controller, and makes measurements of the difference between the desired trajectory and the actual position. It predicts the time at which the laser pump is turned on so that the laser will fire at the desired angle. These predictions are made at regular time intervals. **FIG. 10** is a flow chart that shows the calculations performed at each time interval.

[0155] Trajectory Computation

[0156] The user defines areas within the view of the scanner that are to be scanned, and indicates the density of points to sample within the scanned region. There are several scan patterns which can be used, and these require a specific pattern of mirror movement, known as the trajectory. The objective of picking a good trajectory are the conflicting needs of doing the move quickly and accurately. Accurate movement requires minimum torque, which would otherwise deform the apparatus. This limits the speed with which motions can be made. At equal time increments, a calculation is performed to determine the current position of each mirror. The particular calculation used depends upon the type of scanning employed.

[0157] Raster Scanning

[0158] When the desired scan field is a polygon, one of two raster scanning patterns is used. In the first, scanning is uni-directional (i.e., always proceeds from left to right, or right to left, on parallel lines). **FIG. 11A** shows such a unidirectional scan pattern. In between scan lines, the scan mirror retraces to the beginning of the next line without making any range measurements. The retrace can proceed quite quickly since no measurements are being made during the retrace.

[0159] A slightly more efficient means of raster scanning is bi-directional, in which scanning is also performed on the retrace. **FIG. 11B** shows such a bi-directional scan pattern. This is not as efficient as it might seem because the retrace time is used for other calculations, and because the resulting scan pattern is not as regular.

[0160] Both raster scanning methods require traversing a straight line in the minimum time, starting at zero velocity and ending at zero velocity. The torque applied to the mirror is proportional to the angular acceleration, which must zero at the beginning and end of the scan since the mirror is at result. It can be shown that the trajectory that makes such a minimum energy move between two points is given by the sum of a straight line and a full cycle of a sin. However, this is closely approximated with much less computation by the

minimum degree polynomial, with boundary conditions $p(t_0)=p_0$, $p'(t_0)=0$, $p''(t_0)=0$, $p(t_1)=p_1$, $p'(t_1)=0$, and $p''(t_1)=0$ which is the fifth order polynomial: $p(t)=(p_1-p_0)t^3(6t^2-15t+10)+p_0$, where $t'=(t-t_0)/(t_1-t_0)$.

[0161] Spiral Scanning

[0162] A disadvantage of raster scanning is that since the speed of the trajectory is varying, the scanning efficiency is not optimal. A spiral pattern can achieve a constant speed trajectory which permits a uniform point distribution.

[0163] Seeking

[0164] In addition to scanning a range image, the system is capable of performing a number of functions which are common in surveying. The scanner can be made to search for important features, or locations of high reflectivities. This allows the system to perform normal surveying functions by finding a target whose location is approximated identified, and reporting its exact angles and position.

[0165] Angle Calibration

[0166] The capacitive encoders in the moving coil motors have tremendous repeatability, but relatively poor accuracy. A number of calibration activities need to be continuously performed to ensure system accuracy.

[0167] Before use, each scanner is calibrated over its complete range of angles. At a number of discrete temperatures, a map is created and stored of the measurements of apparent angles for thousands of accurately measured points using an external resolver that is traceable to NBS standards. The DSP linearly interpolates between these measured points on every angle measurement. Preferably, the accuracy of angle measurement is improved by determining scale or offset errors in the encoder during operation. Commercially available scanners can drift significantly with environment changes. This results in a shift in the effective zero and full scale range of the angle measurement, while maintaining the overall shape of the calibration curve obtained by making careful laboratory measurements before the scanner is installed in the system. The environmental effect is reduced by providing a means for determining when the scanner is at a known and repeatable angle. In one preferred embodiment of such a system, two optical references which are fixed with regard to the case of the instrument are aimed at the back of each scanning mirror. There are a variety of mechanisms for providing the optical reference, but in one preferred embodiment, a pair of autocollimators are aimed at a reflective surface on the back of the scanning mirrors and will provide a highly repeatable measurement of when the mirror is normal to the axis of each autocollimator. Each autocollimator gives a reference angle to within approximately 10 μ rad. Periodically, the scanner is moved under computer control to the position at which the mirror is closest to being normal to the autocollimator axis, and the apparent angle is measured. The measurements are compared with the measurements taken when the scanners were calibrated, and a linear correction is calculated and applied to every subsequent measurement.

[0168] In an alternative embodiment, a pair of mechanical stops is provided just past the normal range of motion of the scanning mirror. Periodically, the mirror is driven until it touches a mechanical stop. Then, the scanning mirror is driven with a known current, which corresponds to a known

force. The mirror arrives at equilibrium at a very repeatable position, and this is used to calculate a linear correction to the mirror calibration curves.

[0169] Range Calibration Fibers

[0170] The timing circuits have a certain amount of offset and scale drift with time and temperature, and a provision has been included to compensate for these variations. When an optical pulse is emitted from the laser 602 a small amount of the energy is sampled by a beam splitter 810 and introduced into a single mode optical fiber 830 by focusing the beam using a lens 833 on the fiber face 831. The other face of the fiber 832 is arranged so that the beam which comes out of it is collimated into a beam which enters the lidar receiver 610. The fiber can either be produced so that its length does not vary with temperature, or its variation in length with temperature can be accurately characterized. When a single mode fiber is used, the variation in propagation delay will be less than a few picoseconds and the pulse shape emitted by the fiber will be nearly identical to the pulse shape going into the fiber. Periodically, the timing circuits are used to measure the propagation delay through the fiber, and corresponding adjustments are made to the range measurements taken from external surfaces.

[0171] The fibers can be manufactured so that the end at which the pulse is launched 833 and from which the pulse is emitted 834 are partially reflecting. When this is done, the pulse enters 833 and is propagated to the opposite end 834, at which point only some of the energy is released and the rest returns to the first end 833. Again, a fraction of the light is emitted, and the rest reflected, which eventually is emitted into the receiver. This process repeats until the remaining energy in the fiber falls to negligible levels. The result is a sequence of pulses, commonly 3-10, being applied to the receiver, which have delays all repeatable to within a few picoseconds. Periodically, the timing circuits are used to measure these pulse trains from the fiber, and corresponding adjustments are made to the range measurements taken from external surfaces.

[0172] Range Walk Calibration

[0173] The lidar system measures the range of surfaces by timing the delay between the laser pulse being emitted and returning from the surface. This delay is measured electronically by imposing a sample of the outgoing pulse, and the return pulse on an optically sensitive electronic detector 612 embedded in the receiver 610. In one embodiment, the electronic timing circuit measures the time between when the outgoing pulse exceeds a set threshold voltage, and when the return pulse exceeds the same voltage. The outgoing pulse will be the same intensity within a few percent. However, many surfaces vary greatly in the amount of light that will be reflected. The result is that the apparent relative time for two pulses which occur at the same range but have different intensities may appear to be at different ranges. The measured time for a small pulse 3810 to first exceed the threshold level will be later than the measured time for a large pulse 3830 to exceed the same threshold, even when the pulses return from objects at the same range. Thus highly reflective objects or objects at distances of maximum transceiver sensitivity will appear slightly closer. This creates an apparent "range walk" as a function of intensity. The range walk can be corrected if the shape of the optical return is always the same and the energy of the return is known. The

extremely repeatable shape of the pulses generated by the passively Q-switched microchip laser makes this possible.

[0174] Part of the timing circuit estimates the energy in each detected pulse. A table of corrections is maintained to improve the range estimates. Two different circuits have been employed to make a measurement of the pulse energy for this purpose. The first is a gated integrator, the gate being open at the beginning of the pulse, and closed at the end. The signal is applied to a comparator 3920 which closes the switch 3930 when the signal exceeds a selected level, and closes it when the signal falls below the same level. The signal is also applied to a delay 3910, and the output of the delay goes through the switch 3930 when it is closed, and is applied to the integrator 3940 over the period of time the switch is closed. The delay is chosen to compensate for the time lag in the comparator and switch. When the pulse is complete, the value of the integrator is sampled by an analog to digital converter 3950. The second consists of an integrator with a time constant scaled to the width of the pulse 4010, followed by a peak detector 4020 which has a time constant much longer than the pulse width. The output of the peak detector is sampled shortly after the pulse is detected.

[0175] Periodically, the timing circuit is used to measure a sequence of pulses which have been delayed by the single mode fibers 830 used to calibrate the offset and scale factors associated with the time circuits. Additionally, the intensity of these pulses are varied over a broad range by a variable attenuator 820. By altering the amount of light coupled into the fiber, the energy of the detected pulses can be varied over the dynamic range of the receiver, at one particular time delay. The intensity and the measured time delay values produce a map of the range walk correction required for each intensity, and this correction is applied to subsequent measurements. This correction can provide accuracy of 1 mm over the dynamic range of the instrument, particularly as a result of the great repeatability of the laser pulse waveform. This function is then used to correct the measured range of external surfaces as a function of light intensity returned from those surfaces.

[0176] Geometry Calculation

[0177] The output of the FDV after a range scan consists of points in spherical coordinates with respect to a coordinate system in the scanner. However, the raw data consists of mirror angles and time intervals. The DSP computes the spherical coordinates of the scanned points by talking into account scanner geometry (mirror thickness, optic axes, mirror offsets, etc.) and all the appropriate calibration adjustments.

[0178] Laser Control

[0179] Delay Prediction

[0180] The digital signal processor is responsible for controlling the firing of the pulsed laser, but it can only do so indirectly. The processor has control of the timing for starting the pump diode, which causes the passive q-switch to fire after saturation has occurred. However there is a variable delay between turning on the pump and having the laser fire. The delay is a function of junction temperature, which in turn is a complex function of ambient temperature and recent history of laser firing. The delay generally ranges between 100-300 usecs.

[0181] Fortunately, it is primarily necessary to know the scanning mirror angle at the precise moment the laser fires. After the laser has been fired just a few times, the pump delay does not change quickly if the firing rate does not change quickly. As a result, accuracy of a few microseconds can be achieved by estimating the next pump delay to be the same as that in the previous firing cycle. The digital signal processor measures the pump delay by reading an internal counter when the pump is started and when the laser actually fires, causing an interrupt. Since the interrupt latency is less than a microsecond, this becomes the timing accuracy to which the pump delay can be measured.

[0182] A more sophisticated dynamic model of the thermal properties of the laser could lead to slightly enhanced scanning pattern regularity, but is probably equally limited by the time resolution of the processor interrupts.

[0183] Firing Control

[0184] Given a time vs. angle trajectory for a scanning axis, $w(t)$, a desired angle to fire the laser, and an interrupt interval Δt , the decision to fire the laser amounts to computing the time at which point the pump diode is started.

[0185] Computer Control

[0186] The FDV is designed to perform under the control of a remote host computer which contains graphical controls for a user to specify areas to be scanned. The remote machine controls the FDV through a bi directional serial byte stream, which is effected in any of a number of media: Ethernet, EPP parallel port serial port A processor in the FDV is assigned the task of decoding messages, and scheduling the required activity. FIG. 12 is a block diagram of the FDV processor.

[0187] Host Communications Interface

[0188] The host machine acts as a master, sending a well defined message protocol to command the FDV. When actions are completed, the FDV responds with data and status information. Among the actions which can be requested are:

[0189] Point the scanner

[0190] measure a distance

[0191] range scan a box

[0192] fire the laser n times

[0193] take a video image

[0194] Scanner Control

[0195] Referring to FIG. 13, in normal operation, each scanner in the dual mirror system requires a 16 to 18 bit digital word to set the desired position, which is applied to a precision digital to analog converter to create a voltage proportional to the desired position. However, there will be some error between the position commanded by the output of this converter and the actual position of the scanner, which is reflected by the output of the position encoder. A precision difference signal is generated, and the difference is measured to 12 bit accuracy. This provides an economic method of making 18 bit position measurements while only using an inexpensive 12 bit converter.

[0196] Commercially available galvo scanners have microradian repeatability, but have relatively poor scale and

offset performance, particularly over temperature. A calibration mode has been incorporated into the system to permit making measurements at two precise angles, and using the two measured data points the offset and scale drift of the scanner can be calculated.

[0197] Two methods have been developed for this purpose: an optical and a mechanical means. In the mechanical method, the scanner shaft is gently placed against one of two mechanical stops, and the current in the scanner controller is adjusted to a specific value, which provides a known force. The position signal is adjusted until there is no position error, and this gives the calibrated position measurement. In the optical method, two autocollimators are aimed at the back of the scanner mirrors, which have also been polished and mirror coated. When the scanner mirrors are exactly aligned with one of the collimators, the output from the split photodetector in the autocollimator is balanced. By placing the scanner in each of these precise angles in turn, an offset and scale correction for the scanner encoder can be calculated.

[0198] Timing Circuit

[0199] The purpose of the timing circuit is to provide the relative time between the start pulse and the stop pulse, in picoseconds. There are two subsystems in the timing circuit: a signal conditioning and energy integration circuit (an embodiment of which is shown in FIG. 14), and a time interval analyzer. Both communicate directly with the DSP. Initially, systems have been produced with a commercial timing instrument, the Stanford Research Systems SR620 time interval analyzer. The interface to this instrument is through an IEEE488 interface. In a preferred embodiment, the communications interface to the Stanford Research Systems SR620 time interval analyzer is IEEE488.

[0200] A custom time interval measurement circuit has been developed which utilizes a separately patented interpolation technology. The circuit employs a clock, typically operating at >100 mhz, which is used to make a coarse count of 10 nsec intervals between stop and start pulses. Additionally, there is an interpolator which divides each 10 nsec coarse count into 1000 smaller increments, providing 10 psec resolution. This system has approximately 5 psec jitter. Differential time measurements can be made with less than 20 psec RMS error, which corresponds to about 3 mm. This circuit communicates with the DSP using a dedicated serial bus, and employs a packet protocol: the DSP arms the circuit by sending a single byte. When the timing circuit completes its task, it sends a sequence of bytes which represent both the time delay between start and stop pulses, and the intensity of each pulse.

[0201] Laser Firing

[0202] The DSP has three lines for laser control: one starts the laser pump, the second indicates that the laser has fired, and the third indicates that the return pulse from a target has been detected. When the laser fires, the DSP samples the analog pulse amplitude signal. This happens typically within 1 μ sec.

[0203] Video

[0204] For targeting, the user is provided on the host a video representation of the scene from which he can choose

a portion to be range scanned. In most cases this will correspond to the scene rendered in ambient illumination.

[0205] Capture

[0206] One way the video is captured is by using the scanner to aim single sensitive detector across the scene with the laser turned off. This permits acquiring an image which has very accurate spatial alignment with subsequent range scans. However, image capture can be quite slow in comparison to commercially available cameras.

[0207] A second approach is to utilize standard commercial CCD video cameras to acquire an image. One CCD camera with a wide angle lens is aligned with the range scanner with as small an offset as possible. A second camera with a 5 degree field of view is placed so that its optic axis is coaxial with the transceiver. Thus, a much smaller field of view is accessible through the scanner, and can be scanned with the same resolution as the transceiver. This allows targeting small or distant objects.

[0208] Alignment

[0209] The wide angle lens introduces a fish-bowl effect in the image captured by the CCD sitting behind the lens. Straight lines in the world are not straight in the image. This distortion increases with the distance from the center of the lens. This distortion is removed by comparing the image the camera produces when aimed at a carefully designed and printed calibration target image. The difference in the anticipated image and the recorded image provides the information needed to warp subsequently acquired images to eliminate the distortion.

[0210] Compression

[0211] Each video image is compressed prior to transfer. Currently we are using JPEG standard image compression. It is relatively fast, and creates reasonably small compressed images for communication. Another desirable feature is that the algorithm operates on blocks, which permits us to do interleave image capture, alignment, compression, and transmission in parallel—significantly enhancing throughput.

[0212] Point Video

[0213] A second camera, with a narrow field of view (e.g., approximately 5°) is placed such that it is coaxial with the scanning laser beam. The field of view is adjusted so that the pixel resolution is approximately the same as the voxel resolution of the lidar system. The camera can be operated while the laser is activated. When this is done, a small group of pixels will be illuminated by the laser, and the centroid of these pixels will correspond to the point which would be measured by the lidar. When a video image is capture it can be mapped onto a surface which is estimated by a lidar scan.

[0214] Computer Graphics Perception (CGP) Software

[0215] Referring to FIG. 15, the CGP 40 is a software system that runs on a CGP Computer 1500 and communicates with the FDV 10. The CGP 40 runs on many different types of computers, including laptops and workstations. The CGP 40 operates on a computer 1500 with a suitable display device 1510, such as a color graphic display terminal, a suitable character input device 1520, such as a keyboard, and a suitable pointing device 1530, such as a mouse. The software can use any number of standard 3-D graphics

rendering libraries to interactively present the acquired 3-D data in a window on the display device. The portion of the CGP 40 user interface that involves 3-D view manipulation and data projection into a window is handled by the 3-D library.

[0216] The CGP 40 performs real time 3-D data acquisition and modeling in the field. The CGP's 40 functionality includes high level FDV 10 control, targeting and data acquisition; display and visualization of scanned points; surface segmentation and fitting; manual 3-D model construction; 3-D visualization; interaction with part and model databases; and the ability to export data in standard data exchange formats to other CAD systems for further processing. The integration of hardware and software, as described here, enables major improvements in productivity and quality in the overall process of three dimensional modeling of reality.

[0217] With reference to FIG. 1A, the data acquisition and modeling process divides into the following steps: FDV 10 control, point acquisition, segmentation, geometry fitting, modeling by manipulating the geometry, scene registration with or without warping, model annotation, and geometry display and query.

[0218] With reference to FIGS. 16A and 16B, the foregoing operations may be performed in at least two graphic display windows. One window 1610 (FIG. 16A) displays a video image of the target scene used to define regions to be scanned by the FDV 10 while the other window 1620 (FIG. 16B) displays an interactive 2D projection of the 3-D model consisting of the scanned points and constructed surface geometry as well as other information about the scene. Additional windows may be used to provide multiple views of the data. In addition, the CGP 40 provides additional windows for controlling the FDV 10 hardware and for setting and displaying the status parameters of the system.

[0219] Scan Control

[0220] Referring to FIG. 15, prior to using the integrated hardware/software system the FDV 10 is positioned to point in the direction of the objects 20 of interest.

[0221] Scan control is the process of indicating which portions of the scene that are visible to the scanner are to be scanned. Different parts of the visible scene can be scanned at different densities, since simple geometric objects, such as planes, cylinders and spheres can be accurately modeled with a fairly low number of scan points. Therefore, the region in front of the scanner is often captured in multiple scans, rather than in one high resolution scan. Only regions with high levels of detail need high resolution scans.

[0222] Referring to FIG. 17A, one of the means of scan control is the use of a video image 1710 of the scene acquired from the FDV 10. Using a pointing device, such as a mouse, one can indicate the region to be scanned by any number of methods, such as dragging out a rectangle 1720 on the video image. The CGP 40 instructs the FDV 10 to measure the range of whatever object exists at the center of the user specified targeting region to assist in specifying the scan density, since the angle between points is determined by both the desired density on the surface and the distance from the scanner. A means for specifying the desired scan parameters, such as a dialog box, is provided and allows the user to specify the scan parameters in a variety of ways,

including point density, point spacing, or total number of points in each of the vertical and horizontal directions.

[0223] The CGP 40 then translates the region and scan resolution information into a set of commands for the FDV 10. These commands are communicated to the FDV 10 using a means of communications, such as a TCP/IP network connection, and the acquired data is also returned to the CGP Computer 1500 using the same means.

[0224] Additional scans at different densities can be initiated in the same way, or one can use previously scanned data points rather than the video image to specify new scan regions. If the view of the scanned data is oriented so that it is exactly aligned with the scanner direction, then a scan region can be indicated by methods such as dragging out a rectangular box. When the data is aligned to the scanner in this way most of the 3-D information is difficult to see, therefore, the software can display the points with the intensity of the returned laser light at each point color mapped as described in the next section. The intensity information is often sufficient to identify objects in the data window, so that new scan regions can be defined. Alternatively, the user can model and/or color some of the objects in the scene to help locate regions of interest in the window. Using the data window to define new scan regions avoids any parallax errors, since the view is aligned with the scanner.

[0225] Scan control can also be achieved by using the pointing device to move the laser beam and highlight points in the actual scene. Any number of methods could be used to describe the desired scan region by moving the laser beam and identifying points of interest by a user action, such as clicking a mouse button. Methods could include operations such as: indicating a bounding box by moving the laser to diagonally opposite corners of the desired scan regions; indicating the top, bottom, left and right bounds of the scene; indicating a sequence of points that represent the bounding polygon of the scan region; indicating the center of the scan region and using other means, such as dialog boxes to describe the extent of the desired scan region.

[0226] Point Acquisition

[0227] With reference to FIG. 17B, the data returned by the FDV 10 consist of the coordinates of the points and their intensity values. In one preferred embodiment, the scanning is performed in such a way that the data returned lies in an ordered grid of three dimensional points 1730. Viewed from the scanner, these points appear as a regular rectangular grid, much like a bitmap. However, each point consists of its coordinates in three-space as well as the intensity of the reflected laser pulse at that location.

[0228] Each point returned is displayed in the data window 1620 as it is transmitted by the FDV 10. The CGP 40 lets the user interactively change the 3-D view of the data while the data is arriving to get a better idea of the spatial layout of the data. Also, to help visualize different features in the data, of the CGP 40 can allow each point to be color mapped from the intensity of the reflected laser pulse at that location. A scan cloud 1810 is shown in FIG. 18 representing the surface of a horse sculpture. Instead of having all the points in a single color, as shown in FIG. 18, one can map different laser return intensity values to different colors, and produce a multicolored scan field 1910 as shown in FIG. 19.

The intensity color mapping provides considerable extra surface feedback to the user and this is useful in both targeting and modeling, as described later.

[0229] The ordered grid of points generated is referred to as a scan field. Multiple, possibly overlapping scan fields may be gathered and simultaneously displayed in the manner described above. The data structures within the CGP 40 maintain the list of scan fields, so each data point is always associated with a scan field. The scan fields usually contain data points from the surfaces of many different objects, so they need to be partitioned into smaller groups of points, as described in the next section.

[0230] Segmentation

[0231] Segmentation is the process of grouping together points that were scanned from the surface of the same object. The points from a single object may be a small portion of a scan field, or may occur across multiple scan fields. The segmentation process may be manual, as described below, or automated, as described later in the auto-segmentation section.

[0232] Referring to FIG. 20, the first step of the manual segmentation process is to select one or more scan fields 2010 that contain scan points on the object of interest. Selecting one or more scan fields can be performed by any of the conventional means, such as using a pointing device, possibly together with keyboard keys. Selecting a scan field selects all of the points in the scan field. The group of points resulting from this step form a pool of candidate points that can now be trimmed to remove points on other objects. Each point in the pool is initially marked as selected, and the operations described below can be used to toggle the point states between selected and deselected.

[0233] Referring to FIG. 21, scan points 2010 from a desired object surface can be cut out using one or more lasso operations from possibly different views. The user can manipulate the view direction as needed to give a clear view of the desired subset of the point pool. The user then uses the pointing device to draw a polygonal lasso region 2110, which divides the screen into two regions: the interior and exterior of the lasso polygon. The following operations are supported: mark all points in a region as selected and all of the other points as deselected, mark all of the points in a region as selection without affecting the other points, and mark all of the points in a region as deselected without affecting the other points. The lasso operation may be repeated as many times as necessary to refine the selection, possibly changing the view of the scene between lasso operations. The user can cut the currently selected set of points out to form a new point set. The new point set acts like a scan field in that it is, and can take part in the fitting operations described in the next section. In FIG. 22, three new groups of points 2210, 2220 and 2230 have been created using the manual segmentation method described here, and some points near the intersection of the planes are left from the original cloud of points.

[0234] Geometry Fitting

[0235] In one preferred embodiment, the CGP 40 can contain many geometric primitives that can be used to simulate the actual surfaces of the objects scanned. The geometric primitives include any number of standard graphics primitives, such as triangulated meshes, planes, cylin-

ders, spheres, torii, lines, and points. The simplest form of geometry fitting involves using a triangulated mesh to connect the scan points to show the surface features of the objects scanned. The scan cloud 1810 in FIG. 18 can be meshed 2310 and rendered as shown in FIG. 23. Since the scan data are acquired in a regular grid, it is simple to create a triangular mesh by connecting neighboring points. The user can also set discontinuity tolerances in depth and angle to avoid meshing adjacent points separated by more than the specified threshold. Breaking the mesh in this way provides a more realistic looking surface, referred to as a shrinkwrap surface, because artificial mesh surfaces at occlusion edges do not occur. A wide variety of known mesh operations can be applied to the resulting mesh, such as smoothing (noise reduction) and mesh simplification (to reduce the mesh density in smooth areas that do not require a fine mesh grid). Mesh vertices may also be colored with information such as intensity.

[0236] As stated above, the CGP 40 includes many standard geometric primitives. Before fitting the points to such objects, the point clouds must be segmented as described above. Once segmented, each group of points represents a single surface that can be fit by a geometric object. The fitting can be guided by the user, who may know the type of shape to be fit. For instance, after scanning the corner of a room it is clear to the user that the points on a wall can be fit by a plane and the points on a pipe can be fit by a cylinder, so the user can request the fit of a specific object. It is also possible to semi-automate this process to identify which shape best fits a particular point group.

[0237] Fitting a plane to a set of points is a simple problem that has many well-known solutions. The extent of the patch used in the CGP 40 to represent the plane can be determined by the convex hull of the points in the plane. For instance, the three point groups 2210, 2220 and 2230 shown in FIG. 22 can each be separately fit to planes 2410, 2420 and 2430 as shown in FIG. 24 using any available fitting algorithm.

[0238] Many standard approaches are available for fitting more complex shapes. In one preferred embodiment, two phases are involved: a parameter estimation phase to get a starting point, and an optimization phase, where the parameters are varied to minimize the overall error. The total error is the sum of the squares of the distance between each scan point and the nearest point on the surface of the object being fit. The optimization phase uses conventional optimization methods to reduce the error between the object, as defined by its parameters, and the data given by the scan points.

[0239] A cylinder fitter can convert a cloud of points 2510 as shown in FIG. 25 into a cylinder object 2610 as shown in FIG. 26. All fitted objects, including the cylinder, reference the original points that were used to fit the object. The user may choose to view the resulting cylinder 2610 or the original points 2510, or both, at any time. Using manual or automatic segmentation methods, it is possible to convert the scan clouds, 2710 in FIG. 27, representing many cylinders into the best fit cylinders 2810 shown in FIG. 28. Once each cylinder's diameter and axis are established, it is possible to manually or automatically add elbows 2910 in FIG. 29 to complete the modeling process.

[0240] A cylinder is described by five parameters: a normalized vector describing the cylinder axis (two independent parameters), the radius, and two additional parameters

that are used to locate the line of action of the cylinder axis in space. The length of the resulting cylinder can be determined by projecting the scan points onto the cylinder axis and noting the extreme values of this projection.

[0241] Two novel methods for estimating cylinder parameters are implemented in one preferred embodiment. The first way to find initial parameter estimates for a cylinder is to find approximate surface normals, as described in the auto-segmentation section. If all of the normals are set to unit length, then they can all be considered to be vectors from the origin to a point on the surface of the unit sphere. If one uses each normal vector and its to accumulate a group of points on the unit sphere, then one can fit a plane through the resulting group of points. The resulting plane normal is roughly parallel to the cylinder axis. Given the cylinder axis and the plane from the previous step, one can project the scan points onto the plane. The projected points will be well described by a circle in this plane, since the plane is normal to the cylinder axis. A best fit circle can be calculated using the projected points on the plane to give an estimate of the cylinder radius. The center of the circle on the plane can be converted to a 3-D point to give a point on the cylinder axis.

[0242] The second way to estimate the cylinder parameters is to fit the set of points to a quadric surface, which is described by the implicit equation:

$$F(p)=0=c_1p_1^2+c_2p_2^2+c_3p_3^2+c_4p_1p_2+c_5p_1p_3+c_6p_2p_3+c_7p_1+c_8p_2+c_9p_3+c_{10} \quad (1)$$

[0243] where $p=\{p_1, p_2, p_3\}$ is a point on the quadric surface.

[0244] One can then take $c_{10}=-1$ since the equation is implicit and perform a least squares fit with all of the data points to determine the other nine parameters. After determining the best fit quadric surface for a given set of points the next step is to find a point actually on the new surface (p_s) that is in the vicinity of the other points. This is achieved by finding the centroid of the scan points (p_c) and then finding the closest point on the surface of the quadric to give p_s . The normal to the surface at point p_s can be determined using:

$$N_p=D_1p_s+D_2 \quad (2)$$

[0245] where

$$D_1 = \begin{bmatrix} 2c_1 & c_4 & c_5 \\ c_4 & 2c_2 & c_6 \\ c_5 & c_6 & 2c_3 \end{bmatrix}, D_2 = \begin{bmatrix} c_7 \\ c_8 \\ c_9 \end{bmatrix}$$

[0246] Two unit vectors, u_1 and u_2 , are then found such that they are normal to both each other and N_p . These vectors form a basis for the surface at the point under consideration, and additional vectors on the surface can be found as:

$$v_\alpha = u_1 \cos\alpha + u_2 \sin\alpha, 0 \leq \alpha \leq 2\pi \quad (3)$$

[0247] The unit principle vectors v_α are then found by determining the rotation α that satisfies:

$$v_\alpha \cdot (N_p \times D_1 v_\alpha) = 0 \quad (4)$$

[0248] There are two solutions to this equation that yield the orthogonal unit principal vectors v_1 and v_2 . The surface curvatures κ in these two principle directions are then given by:

$$\kappa(v_i) = -v_i \cdot \left(\frac{D_1 v_i}{\|N_p\|} \right) \quad (5)$$

[0249] For cylindrical surfaces one of the principal curvatures will be near zero, and the radius of the cylinder is the reciprocal of the absolute value of the nonzero curvature. A method to determine the radius (r) and axis of the cylinder has been described, and only the location of the axis needs to be determined. A unit surface normal can be calculated as $\hat{n} = N_p / \|N_p\|$. The sense of the normal can be adjusted so that it points towards the interior of the cylinder by ensuring that $\hat{n} \cdot (p_c - p_s) > 0$ since the centroid of the points lies on the interior of the cylinder. A point on the axis is then given by $p_s + r \hat{n}$. These starting parameters can then be used in a minimization process to find the best fit parameters for the cylinder.

[0250] The novel method described above for curvature estimation using the quadric surface formulation is further used in a novel way for automated object type determination. If the points being fit are well represented by a plane then both principal curvatures will be near zero. If the points being fit are from a cylindrical surface then one curvature will be near zero and the other will be nonzero. If the points are from a sphere then both curvatures will be nonzero and their magnitudes will be approximately equal. Combining the automatic detection of object type and the auto-segmentation algorithm, described later, allows the CGP 40 to have a novel method for automatic fitting of many objects that occur in typical scanned scenes.

[0251] A further use of the curvature estimation is sphere fitting, which is achieved by using the quadric surface approach to approximate the radius and location of the center point, and then using a four parameter (center point and radius) minimization to reduce the error between the sphere model and the measured points. The novel method described above for finding a point on the axis of a cylinder is also used in the preferred embodiment to find the center of a sphere.

[0252] The segmentation techniques disclosed above can be used to create a variety of useful fitting tool based on combinations of the previously described shapes. For instance, a corner, consisting of an intersection of three planes which may or may not be orthogonal, is a very common feature to scan. Knowing that the specified point group contains three intersecting planes, such as the points 2010 in FIG. 20, the points are automatically segmented, using the technique described later, into three subgroups of points that each lie on separate planes. Then, any available plane-fitting algorithm can be used to fit the planes to the scan points in each group. Unlike the more general auto-segmentation algorithm described later, if one knows that the corner object consists of three planes, the fitting algorithm need not try to fit cylinders, spheres or other objects and check which produces the best fit, only planar fits are required. Referring to FIG. 30, the corner fitting tool not only fits the planes 3010, 3020, 3030 but can also intersect them to complete the corner. As shown in FIG. 30 additional useful information can be given to the user by introducing lines 3040, 3050, 3060 at the intersection of pairs of planes, and a vertex point 3070 that represents the location of the

corner. This vertex point is much more accurate than a single scan point, because each plane is fit using many data points, and the vertex is generated by intersecting the planes. The above novel method for automatically creating a corner with its intersection lines and vertex can be used as a tool of the CGP 40.

[0253] Each object stores information about the quality of the fit, so that the user can query the object and examine the mean, standard deviation, and worst 10 errors. Knowing the accuracy of the FDV 10, the CGP or the user can then decide if an error has been made during the fit. Errors can arise when the wrong type of primitive is fit, or when extraneous points that were not actually scanned from the desired surface remain in the data set. In addition, the objects store their geometric parameters so the users can query for radius, length or other values of interest.

[0254] In addition to the class of general object fitters, which are given close to no initial information other than the points to fit, there is a class of fitters that can take advantage of foreknowledge about objects in the scene. An area in which such foreknowledge exists is that of the construction industry, where parts used follow standards of dimensions and design. For example, the external diameter of pipes from a particular manufacturer may come in five different sizes: 4", 5", 6.5", 8", and 10". This information typically resides in tables that describe relevant attributes of these parts. The cylinder fitter can take advantage of the information in these tables to significantly reduce the solution space to be searched: the fitter need only search for solutions involving cylinders of one of those diameters. Another way to use such table lookups is to have the fitter come up with a general solution, then match against entries in the object tables to find the entry with the closest parameters. For example, a pipe fit by a cylinder with 7.8" diameter would be matched against the 8" entry in the table from the example above; the user (or fitter) then has the option of refitting an 8" cylinder to the pipe, or accepting the 7.8" cylinder. Yet another use is for the user to manually select a specific entry (or collection of entries) from the table and tell the fitter to use its parameters in the fit, which also reduces the fitter's solution space (which can decrease the time taken).

[0255] Modeling

[0256] The fitting of geometric primitives, as described in the previous section does not usually complete the modeling process. It is often the case that only a portion of the objects surface is scanned, such as one end of a cylinder or a portion of a wall, and further operations are required to complete the 3-D model. Modeling is the process of completing the construction of the 3-D model given some fitted geometric primitives.

[0257] Many common CAD operations such as extension, intersection (mutual extension) and trimming are available in the CGP 40. For example, the cylinder 3110 in FIG. 31 does not initially extend to the plane of the floor 3120. In FIG. 32 the cylinder 3220 has been extended to the floor plane 3120. These simple operations enable rapid completion of portions of the model from the objects created by geometry fitting. For instance, given three planar patches that were fit from scan data near a corner, it is easy to mutually extend the three planes to complete the corner feature.

[0258] Object extensions can be accomplished in several ways. One way is to select the geometric object to be

extended and tell it to extend to a subsequently selected object. The nature of the extension is determined by both the type of object to be extended and the second object selected. For example, a cylinder extends the end closer to the second object along its centerline until its end intersects with the infinite plane defined by the second object's geometry (in the case of a planar patch, the infinite plane is that of the patch's plane, and for a cylinder, the infinite plane is that containing the centerline and as orthogonal to the extending cylinder as possible).

[0259] Another way is to make use of object handles, which are nodes that the user can grab. These handles are tied to an object's definition (position, orientation, and size) where appropriate, and by moving a handle, the object's definition changes accordingly. Again, taking the cylinder as an example, the same extension described above can be accomplished by grabbing the handle on the end to be extended, and then moving the handle (and extending the cylinder) to the desired position. A handle's motion depends on the part of the object with which it is tied; a handle on the centerline of the cylinder is constrained to move only along that centerline, while a handle on the boundary of a planar patch is constrained to move only inside the plane of the patch. For some objects, handles may be inserted and removed, changing the definition of the shape of the object (for example, handles on a planar patch have a one-to-one correspondence to vertices on the planar patch's boundary). Other handles can provide rotational control over an object. The control of handles is interactive and dynamically updates, so the user can see the intermediate results of the redefinition.

[0260] A new operation, called merging, has been developed to allow different portions of a single object surface to be joined to form a single object in the CGP 40. It is often the case that one's view of an object is obscured by other objects in front of it. For instance, the view of a back wall in a room may be divided into two pieces because of a column in the foreground. A scan of the region will result in different groups of points on the particular object. If auto-segmentation is used, as described later, rather than manual methods where a user knows that the points belong to the same object, then separate point groups would be formed. Each point group would then be fit to a separate object, resulting in multiple pieces of the same surface. The two objects in the CGP 40 that are known to be on the surface of the same feature, such as the two planar patches of wall obscured by a column, can be merged to form a single object. Each object stores a reference to the data points that define it, so when a merge request is received a new geometric fit is performed on all the underlying data points that were part of the constituent geometries to achieve the best overall fit. This novel method of increasing the accuracy of object fitting is used in the merging operations of one preferred embodiment of the invention. The characteristics of the two primitive objects that were merged do not affect the outcome of the merge; only the underlying point positions are considered.

[0261] Using the manual or automatic methods, the user can take a cloud of points 3320 from the surfaces of many objects, such as the points on the pyramid 3310 shown in FIG. 33 and convert them to a set of geometric objects 3420, such as the planar face of the pyramid 3410 in FIG. 34. The modeled scene shown in FIG. 34 accurately represents the

features of the original objects that were scanned, and allows measurements to be made between any particular locations in the modeled scene.

[0262] Scene Registration

[0263] The initial position of each scan point is described in a local coordinate system whose origin is that of the FDV 10, and whose axes are fixed relative to the FDV 10. Therefore, scan fields taken without moving the FDV 10 are inherently registered in that all the scan points use the same coordinate system. However, if the FDV 10 is moved between scanning operations, additional effort is required to transform the data to use a shared coordinate system. This is a six parameter rigid body transformation of a data set, involving three translation and three rotation parameters, and is easy to apply to the data once the transformation is known.

[0264] A novel process is used to register the scan fields from different FDV 10 positions. The novel registration process requires the user to identify pairs of points, lines or planes in two different scenes that represent the same feature. It is also possible to use different features, such as the back and front face of a wall, if the user knows the offset between them. The planes and lines are converted to point pairs, as described below, and the process operates entirely on point pairs. The points used for registration may be actual scan points or may be constructed points, such as a corner at the intersection of three planes.

[0265] Given a set of point pairs, the registration process searches the set of candidate points for three pairs that are not colinear. Using the three point pairs, one can construct the transformation required to convert the coordinate system in one view to that used in the other view, which in turn can be used to transform the scan points to all share a single coordinate system. For convenience, the process will be described in terms of the first data set remaining fixed while the second data set is transformed to use the coordinate system of the first. The process works equally well if the user fixes the second data set and transforms the first.

[0266] The first step is to apply a rigid body translation to the second data set to make the first pair of points coincide in terms of their x, y and z components. The second step is to rotate the second data set about its first point until the lines formed by points one and two in both data sets are colinear. The third step is to rotate the second data set about the line established in the previous step until the planes defined by points one, two and three in both data sets are coplanar.

[0267] Once an initial estimate is made one can use all the point pairs and an error minimization method to reduce the sum of the squares of the distances between each point pair.

[0268] In order to use the point registration method described above, the CGP uses a novel method to convert lines, planes, and planes with offsets to sets of point pairs. Whenever a nonzero plane offset is present the new points introduced are shifted in the second scene to a position where they will match exactly with the corresponding points in the first scene. The replacement of planes and lines with points makes it simple to write the error function for minimization, since only point errors are involved, rather than angular and distance errors simultaneously.

[0269] In replacing planes and lines, one can only introduce points that are at location relative to the user specified

objects, since the origins of the two data sets are different. For instance, introducing a new point pair in a plane at the location closest to the origin would not result in points that actually match in space, since the origin is arbitrary. However, introducing a point pair at a plane-line intersection will give matching points in the two data sets. Some pairs of objects, like parallel lines, should not be used to introduce new points so an angular tolerance, called ATOL below, is used to ignore poor object pairs. ATOL is initially set to ten degrees but other values can be used to generate fewer or more artificial point pairs as needed. The point pairs are introduced in the following order:

[0270] For all plane-line pairs where the angle between the line and plane is greater than ATOL, introduce two new point pairs. The first new point is inserted at the intersection point of the line and plane, and the second point pair is inserted along the line of action at a fixed distance away from the first point, here taken to be the minimum of the line lengths in the two views.

[0271] For all pairs of planes and points, introduce a new point pair on the plane such that the plane normal passes through the new point and specified point.

[0272] For all plane pairs whose normals are at least ATOL apart, generate a new line segment along the intersection of the planes and make the line segments length equal to the minimum extent that any plane has along the line. The new line segment has no direction, but has both length and position information. After this step the planes are no longer needed.

[0273] For all pairs of lines and points, introduce a new point on the line at the location where it is closest to the specified point.

[0274] For all pairs of lines separated by an angle greater than ATOL, introduce four new pairs of points. The new points are the ends of line segments along the original line of action, but centered on the location of closest approach of the two lines. The distance between the new line points is equal to the minimum length of the line segment lengths along that line of action from the two data sets. After this step the lines are no longer needed.

[0275] The result of the plane and line replacements as described above is a set of point pairs that retains the direction information associated with the original planes and lines. The augmented set of point pairs can then be used for the registration process that is described above.

[0276] After registration of the two scenes, primitives from the two individual views which represent the same physical object can be combined using the merging technique described previously. In particular, matching planar patches representing the same surface can be combined into one extended planar patch. Similarly, pieces of matching cylindrical surfaces can be merged to form a single cylinder.

[0277] Warping Data Sets

[0278] The registration process described above is a rigid body transformation that does not modify the relative locations of objects within either data set. After registration, most of the point, line or plane pairs that were identified will

still have small errors, since the minimization process reduces the total mean square error. A novel method is presented that allows the user to force the identified pairs to exactly match by deforming the scene volumes.

[0279] As with any measured data, there is some level of error associated with each scan point location. The magnitude of the error associated with a point location will vary with the measuring technology used, but some error will always be present. Since the data under consideration here describes surface features of objects, the data errors will manifest themselves as surface irregularities. For example, a set of points acquired from an actual planar surface may not all lie on a plane, but will have some small scatter away from the real plane location. Calculating a best fit plane through the set of measured points may not give the real plane location or orientation due to the errors in the point data set.

[0280] The errors in recovered features, such as planes, cause errors in the relationships between the recovered objects as well. For instance, if data is collected from two planes that have an exactly ninety degree angle between them, the best-fit planes generated from the data points may not be exactly ninety degrees apart. Similarly, cylinders that were parallel in the real scene may result in best-fit cylinders that are not parallel after fitting from the scanned points. These inconsistencies in the recovered features, that occur due to measurement errors, will appear whether the data points are collected from a single scan position or are a union of scans from a variety of different positions.

[0281] The lack of fit problem may actually grow if several different sets of scan data are registered using a relative system. If a series of sequential scans are collected, and each scan is registered with respect to some recognizable sequence of data points in a previous scan, then the absolute errors in each scan may grow. If at the end of the sequence of scans the locations of features are exactly known, then one must adjust the scanned data points so that they fit the known locations. In surveying both the 2-D closure problem and the 3-D benchmark matching problems are similar in nature to the problems described above. In the surveying closure application, when one surveys a sequence of locations and arrives back at the starting location one typically finds that through cumulative measurement errors the starting and finishing locations are not at exactly the same location. The closure error, which is the distance between the starting in finishing locations, is distributed using well known surveying techniques throughout the other data points collected such that the first and last end points meet after the correction is made. Similarly, when surveying benchmarks of known location are introduced into a surveying data set the data set must be adjusted to accommodate the known benchmark locations. Both the closure problem and the benchmark matching problem can be solved by the method described here since they can be described in terms of displacement constraints.

[0282] The novel method described here to correct location errors in measured 3-D data sets and distributes the errors throughout the point sets by applying solid mechanics principles to a volume surrounding the data points. The method provides a technique for satisfying a wide variety of displacement constraints on 3-D data sets and also distributes the measurement errors throughout the data sets. The process of deforming the data sets to achieve these goals is

called warping. The displacement constraints can be specified in terms of both control points, whose absolute coordinates are known in space and do not move, and tie points, which represent the same location in two or more data sets, but whose absolute location is unknown. One can describe constraints involving more complex objects, for instance, line segments by specifying two points, and reality, the constraints themselves can be used to anchor multiple volumes together. Mechanics principles allow one to determine the minimum energy deformation of the volume that satisfies the stated constraints, which mimics what would actually happen to a real deformable object subjected to the same constraints.

[0283] In a particular embodiment, the warping method uses principles of solid mechanics to deform a volume containing points of interest in order to satisfy a set of constraints applied to the data points. Not only are the constraints satisfied, but the effects of the initial location errors are spread throughout the volumes operated on.

[0284] The finite element method is used to apply the principles of solid mechanics to the volumes enclosing the points. The volume is discretized into a set of points or vertices and a set of elements that connect to the vertices. Four node tetrahedral elements are used to discretize the volume.

[0285] The first step of the process is to collect the set of constraints that apply to one or more data sets. During this phase, one must identify constraints that are to be satisfied by the warping process. These constraints include the identification of points that represent the same physical location in different data sets (tie points), like the corner of a cube, which are to appear at the same location when the warping process is completed. Some of the tie points may not be points that were scanned in the original data set, but may be constructed from groups of other points. For instance, if one had a series of points that represented three planes intersecting at a corner, then one could fit three planes to the points, and use the resulting corner point as a tie point. The constraints are specified in terms of pairs of objects, such as points, lines and planes, as well as the desired offset and angle between them. The two objects involved in the constraint can be contained in a single data set or can occur in different data sets. Within a single data set, one could specify that lines or planes remain parallel, or that the distance between two points be a specified amount. Between multiple data sets, one could write similar constraints, or indicate that the features seen in two data sets represent the same object. One might also know the actual location of some points very accurately (benchmarks) and constrain points in the data set to lie at the known locations. Using these benchmark points to anchor different points in the data sets enables the closure problem to be solved, since the data sets will be warped so that the measured data points move exactly to the desired control point locations and the errors in the data set will be smoothed over all volumes.

[0286] The second step in the warping process is to register all the data sets involved, as described in the previous section.

[0287] The third step in the warping process is to select a volume that surrounds the region of interest, and describe the volume in terms of a set of new points. The region that can be scanned by the FDV 10 is called the view volume and

is shaped like a pyramid, with the tip of the pyramid located at the origin of the scanning device. A pyramidal shape can be used to bound the view region for the purpose of warping, and the pyramid is easily described by five points, using the same coordinate system as the data points. These new points do not become part of the data set, but are used in the warping process. The convex hull of these points represents the new volume surface, and should enclose all the data points on the interior. This operation is performed separately for each data set.

[0288] The fourth step is to mesh each of the data volumes. Meshing involves filling the volume with finite elements that leave no gaps and that do not overlap. The finite elements span between the points or vertices that have been defined on the volume boundary and those that are involved in constraints on the interior. The points in the data set do not all need to be included in the warping process, only those that are used in constraint specifications and those that define the volume boundary need to be used. The elements in the initial mesh may be of poor quality due to their shape. Long sliver elements, for instance, are known to give poor results in finite element analysis. Therefore, the meshing process is actually iterative. New points are inserted into the mesh, and then old elements are removed and new elements are introduced so that the mesh quality improves. This iterative process continues until one is satisfied with the overall quality of the mesh. In one preferred embodiment, four node tetrahedral elements are used. The initial mesh is constructed by applying a 3-D Delaunay triangulation on the starting set of points. The iterative process identifies poorly shaped elements using an element quality measure, and introduces new points and remeshes the region. The process terminates when all elements meet a minimum quality criteria. The preferred implementation uses longest edge bisection to introduce new points that improve the mesh, but other methods can be used.

[0289] The fifth step processes the constraints described in step one into a system of linear constraints. In the preferred embodiment, the final system of constraints is linear in terms of the nodal displacements at the vertices of the tetrahedral elements. The desired form of the constraints is:

$$Cu=q \tag{6}$$

[0290] The matrix C contains constant coefficients. The number of rows of C is equal to the number of constraints in the system. The vector u represents the 3-D displacements of the vertices of the tetrahedral elements. The vector q contains constant coefficients. If the constraints are homogenous then each element of q will be 0. The form of constraint specification given in Equation (6) allows arbitrary linear multipoint (involving more than one vertex) constraints.

[0291] The conversion of the constraints specified in step one into the form shown above depends on the type of constraints involved. For two points to be tied together the constraint would be:

$$p_1+u_1=p_2+u_2 \tag{7}$$

$$\text{or } u_1-u_2=p_2-p_1 \tag{8}$$

[0292] In these equations, p_1 and p_2 are vectors from the origin to the vertices of interest, while u_1 and u_2 are the displacements of the same vertices during warping. Equation (7) demands that the final location of each vertex, equal

to the starting point plus the displacement during warping, be equal to the final location of the other point Equation (8) is in the form of Equation (6), with $q=p_2-p_1$ and results in three linear constraints, in terms of the x, y, and z components of the nodal displacements. Expanding Equation (8) into three equations in the form of Equation (6) then gives:

$$\begin{bmatrix} 1 & 0 & 0 & -1 & 0 & 0 \\ 0 & 1 & 0 & 0 & -1 & 0 \\ 0 & 0 & 1 & 0 & 0 & -1 \end{bmatrix} \begin{Bmatrix} u_{1x} \\ u_{1y} \\ u_{1z} \\ u_{2x} \\ u_{2y} \\ u_{2z} \end{Bmatrix} = \begin{Bmatrix} 0 \\ 0 \\ 0 \end{Bmatrix} \tag{9}$$

[0293] Other constraints, like the distance between two points, are non-linear in nature. The nonlinear constraints can use the existing geometry of the system as well as small deformation assumptions to produce linear multipoint constraints. For example, to specify the desired distance between two points to be some specified value x, one could determine the vector v_{21} between the final points locations:

$$v_{21}=(p_2+u_2)-(p_1+u_1) \tag{10}$$

[0294] and then specify the desired length of the vector:

$$\|v_{21}\|=x \tag{11}$$

[0295] or, using the vector dot product:

$$v_{21} \cdot v_{21}=x^2 \tag{12}$$

[0296] Both Equations (11) and (12) are nonlinear in terms of the displacements of the nodes, u_1 and u_2 . To linearize the constraint we can specify the desired length along the original line of action to be equal to the desired offset:

$$\text{let } n_{21} = \frac{(p_2 - p_1)}{\|p_2 - p_1\|} \tag{13}$$

$$v_{21} \cdot n_{21}=x \tag{14}$$

$$\text{or } [(p_2+u_2)-(p_1+u_1)] \cdot n_{21}=x \tag{15}$$

$$u_2 \cdot n_{21}-u_1 \cdot n_{21}=x-\|p_2-p_1\| \tag{16}$$

[0297] The term on the right hand side of equation (16) is the desired distance between the points minus the current distance between the points. The x, y and z components of n_{21} are constraints. Equation (16) can be written as a single constraint in the proper form as:

$$\langle -n_{21x} \quad -n_{21y} \quad -n_{21z} \quad n_{21x} \quad n_{21y} \quad n_{21z} \rangle \begin{Bmatrix} u_{1x} \\ u_{1y} \\ u_{1z} \\ u_{2x} \\ u_{2y} \\ u_{2z} \end{Bmatrix} = x - \|p_2 - p_1\| \tag{17}$$

[0298] In Step 6 the final system of linear equations is assembled. There are two parts to this step: first, assembling the element stiffnesses for each of the tetrahedral elements, and second, selecting and applying a constraint handling

technique. The calculation and assembly of element stiffnesses follow standard finite element procedures. Using constraints in the form of Equation (6) involves a constraint processing method. The Lagrange Multipliers technique can be used to introduce the effect of the linear constraints, but any other method, such as penalty or transformation techniques, could be used equally effectively.

[0299] Using Lagrange Multipliers, one introduces a new variable into the final system of equations for each constraint in the system. One then modifies the static equilibrium equations for the unconstrained system, which are given by:

$$Ku=r \quad (18)$$

[0300] In Equation (18), K is the system stiffness matrix, assembled from the individual element stiffness contributions, u is the displacement vector that is the solution of the problem, and r is a vector of externally applied loads. In this embodiment of the invention, there are no externally applied loads, so the r vector contains only zeroes. Equation (18) does not include the effect of any constraints, but these can be included using the Lagrange Multipliers technique to give the system of equations:

$$\begin{bmatrix} K & C^T \\ C & 0 \end{bmatrix} \begin{Bmatrix} u \\ u_L \end{Bmatrix} = \begin{Bmatrix} r \\ q \end{Bmatrix} \quad (19)$$

[0301] In Equation (19) K, C, u, r, and q are as previously defined, and UL is a vector containing the additional Lagrange Multiplier variables that are introduced using this method. The matrix C^T is the transpose of C, and 0 is a matrix of zeroes. The solution of Equation (19) gives the displacements u that satisfy the linear constraints described by C and q. Note that these constraints may be linearizations of nonlinear constraints, and the nonlinear constraints might not be satisfied by the solution at this point.

[0302] If penalty or transformation methods were used instead of Lagrange Multipliers, a system of linear equations different from those shown in Equation (19) would be produced, but the solution of the linear system of equations will give similar values for the displacement vector u.

[0303] In Step 7 Equation (19) is solved to give u and UL. There are many methods available to solve large systems of linear equations, and the preferred embodiment uses a symmetric solver with a profile storage scheme. The different types of solvers that could be used give essentially the same results, but optimize speed and memory usage differently.

[0304] The preferred embodiment uses a direct solver, but iterative sparse solvers could be used as well. The system of equations shown in Equation (19) is sparse, so significant speed enhancements can be achieved by selecting the proper solver. However, the results of the warping process overall are unaffected by this choice.

[0305] In Step 8, one must check if the current displacement satisfies the constraints to a desired level of accuracy. If the current deformed shape violates the offset or angle in any of the constraints collected in Step 1 by more than a user specified tolerance, then steps 5 through 7 must be repeated, starting with the new deformed shape. The linearizations of

the shape may change on each iteration since the geometry of the volume changes with the cumulative deformations. When all constraints are satisfied within the given tolerance, then one can proceed to step 9.

[0306] Step 9 uses the nodal deformations u calculated in Step 7 to determine the deformation of any point of interest within the volumes. For each point of interest, one must find a finite element that includes the point on its surface or interior. If the point is internal to an element then only one such element exists. If the point is on the surface of an element or along the edge of an element, then several elements could be considered to contain the point. Any of these elements can be selected to determine where the point of interest moves. If the point is shared between elements, then the use of any of the elements to find the point displacement will give the same results. Once an element is identified, the vertex displacements of that element are extracted from u and are used to determine the displacement of any point on the interior using an interpolation process. This procedure uses the finite element shape functions which are linear in the preferred embodiment, and is a common operation in finite element analysis.

[0307] Auto-Segmentation

[0308] The novel auto-segmentation process, as presented below, involves a similar sequence of operations to the manual modeling process described previously. A point cloud is segmented, geometric primitive objects are fit to the point groups, and then modeling operations, such as extension and intersection are used to complete the model. In this novel process, automation is applied to each of these steps, as well as the entire process, as described below.

[0309] It is possible to automatically partition the scan points into groups representing primitive geometrical shapes by using variations of common machine vision techniques. A gridded scan field is stored in a two dimensional array of points, much like a regular bitmap. The scan field differs from a bitmap in that more information is stored at each location than just a color. Each point stores its location in space, from which the distance to the scanner can be calculated, as well as the intensity of the return laser pulse. The depth information calculated from the three dimensional position stored at the points is crucial to the automated segmentation algorithm described here, even though many operations, such as filtering, rating, thresholding and thinning are commonly used image manipulation operations.

[0310] The first stage of the auto-segmentation process is to estimate the surface normal at each point in the grid. This can be achieved using many different techniques, the current embodiment of the software fits a plane to the nearest neighbors of the point in the 3 by 3 grid surrounding it. The normal of the resulting plane is taken as the normal at the center point. Each point in the grid has a normal calculated in the same way, except that edge and corner points ignore the missing neighbors in the normal calculation. The normal stored at each point is a three dimensional vector and is normalized to have unit length.

[0311] In the second phase two rating images are created by convolving standard edge detection filters over the grid. The first rating image is created by convolving the depth of the grid point with an edge detection filter to identify depth discontinuities, such as those that would occur at an

occluded edge. A variety of edge detection filters can be used, but rather than operate on color or intensity the filter operates on the depth information stored at each grid point.

[0312] The second rating image is created by convolving the normal with an edge detection filter. The normal rating image is actually composed of 3 subimages created from a convolution with the normal's x, y, and z components. The resulting three values are combined by taking the square root of the sum of the squares to give a per-point scalar value. The second rating image is used to identify normal discontinuities, such as those that would occur at the edge between a wall and a floor. Again, a wide variety of edge detection filters can be used, but the values used are normal coefficients rather than color or intensity.

[0313] Once the two rating images have been created they must separately be converted to binary images. Conventional machine vision algorithms, such as recursive thresholding can be used to achieve this task. Each point in the depth and normal rating images contains an estimate of the gradient of the depth and normal respectively. Recursive thresholding can be used to isolate the regions of highest gradient. In the resulting binary images the points in the regions of highest gradient are marked as edge points while the rest of the points are marked as non-edge.

[0314] A final binary image is created by marking a point as an edge point if it is marked as an edge point in either or both of the two binary images created by recursive thresholding above. All other points are marked as non-edge. This image contains all edge points that delineate the boundaries between groups of points on different surfaces.

[0315] The final step of the point partitioning process is to use a connected components algorithm to collect the points into groups separated by edges. Points are considered to be connected only if they are vertically or horizontally adjacent in the grid, diagonal adjacency is not used. Very simple algorithms can be used to identify the unique groups of non-edge points in the image. Each group of connected points is then cut from the initial point set to form a new group of points. The result of this algorithm is the partitioning of the point set into multiple point groups that each represents a single surface. Each of the new point groups can be fit by a geometric primitive as described in the next section.

[0316] Once the scan cloud has been partitioned into groups of scan points that lie on different surfaces, the next step is to fit objects to the desired surfaces. A variety of methods can be used to achieve this task. The current embodiment of the software can perform the object fitting process in two different ways. The first method fits a series of objects to each group of points, and selects the objects that produces the smallest distance errors between the measured points and the fitted object surfaces. The second method uses the quadric surface fit described previously, and resulting principle curvatures, to determine if a plane, cylinder or sphere should be fit to a particular point group. Other variations of these approaches could also be used, such as progressive commitment, where objects are fitted in order from simplest to most complicated, and the process stops whenever the errors associated with the particular fit drop to acceptable levels.

[0317] The last stage of auto-segmentation process extends primitive objects, where possible, to create complete

object intersections, rather than stopping at scan point boundaries. Using the gridded nature of the original data and the edge information from the point partitioning algorithm described above, it is possible to extend and intersect objects. For all edges that result from surface intersections, which are the surface normal discontinuity edges described above, one can extend the objects on either side of the edge to form an intersection.

[0318] Model Annotation

[0319] In order to compose a semantically rich 3-D model, individual parts in the above geometrical model can be annotated with additional, possibly non-geometric, information, such as material references or part numbers. This information can be entered manually through a special window for displaying object attributes.

[0320] The user may click on an individual part in the geometrical model and recover such additional information through other windows. Similarly, the user may request that all parts which meet some selection criteria are to be highlighted.

[0321] A novel method is also used for automatic model annotation. This method uses the FDV 10 to scan bar codes containing pertinent information relating to any given object. Standard bar code reading and decoding techniques are used to convert optical information to useful digital information that is associated with a given object scanned at the same time as the bar code. The captured information can be displayed as described above for the manual method.

[0322] Geometry Display and Query

[0323] The model is accessible in a variety of ways, including access through the data window 1610 where the model is rendered. Many standard graphic interface techniques can be used to manipulate the view of the model; in one preferred embodiment, a crystal ball interface is used. Any object in view can be selected using the pointing device; its geometrical attributes can then be displayed, in addition to other annotations added to the object. Traversing the data set can be simplified by placing objects in different layers, and then displaying only the layers of interest. Reducing the number of rendered-objects in this way increases the interactive performance of the program. In addition to querying the geometric properties of individual objects several standard tools for measuring distances and angles between objects can be employed. Additional standard techniques can be used for operations such as contour generation, 2-D section cutting, and automated dimensioning.

[0324] The resulting model can be exported to any of a number of CAD programs for further editing or designing. In the preferred embodiment, the CGP 40 can create a CAD file in a format compatible with several commercial CAD programs, and can then start the CAD program, having it load the generated data file. With both the CGP and the CAD program running, the user can then view and work with the model using either program.

[0325] The following documents form an integral part of this specification:

[0326] Modular Decomposition

[0327] Summary of Proposed CGP Specification

[0328] Cyrax Software Specification

[0329] Product Overview

[0330] Collection of View Slides

[0331] Introduction and Overview

1. A method of manually separating from a plurality of clouds of points, representing three-dimensional features in a scene, a subset of the points that represents a desired feature in the scene, the method comprising:

selecting all the point clouds that include at least some data points representing the desired feature; and

changing a view of the clouds and drawing a polygonal lasso to refine a selected subset of points to be included in a point sub-cloud and repeating the refining as many times as required to obtain the desired sub-cloud.

2. A method for automatically segmenting a scan field of a scene into subsets of points that represent different surfaces in the scene, comprising the steps of:

separating the scan field into a depth grid that includes depth information for scanned points of surfaces in the scene and a normal grid that includes an estimate of a normal to scanned points of the surfaces;

convolving the depth information of the depth grid to generate a depth rating image whose values represent a gradient of depth change from one scanned point to another scanned point in the scene;

convolving the components of the normal grid to generate a scalar value for each component for each point of the normal grid,

for each point of the normal grid, determining from the scalar values for the components of that particular point a gradient of the normal at that point, wherein the gradients determined for the points of the normal grid collectively constitute a normal rating image;

converting the depth rating image to a binary depth image using a recursive thresholding technique;

converting the normal rating image to a binary normal image using a recursive thresholding technique;

combining the binary depth image and the binary normal image to determine a single edge image; and

grouping subsets of non-edge points as belonging to corresponding surfaces of the scene.

3. A method as recited in claim 2 further including the steps of:

determining the type of geometric primitive that would best first each group of points; and

fitting the geometric primitive to the data points.

4. A method as recited in claim 3 further including the step of intersecting adjacent planar regions in the scene.

5. A method for fitting a point cloud representing a corner, comprising:

determining a fit of three planes to the points of the point cloud and creating the planes for a model;

determining the three lines at the intersection of pairs of planes and creating the lines for the model; and

determining the vertex point at the intersection of the three planes and creating a vertex point for the model.

6. A method for modeling a three-dimensional scene, comprising:

generating a plurality of points that each represent a point on a surface of the scene;

determining a best fit of a cylinder for a group of the points using surface normal estimates and global error minimization.

7. A method for modeling a three-dimensional scene, comprising:

generating a plurality of points that each represent a point on a surface of the scene;

determining a best fit of a cylinder for a group of the points using a quadric surface fit and global error minimization.

8. A method for modeling a three-dimensional scene, comprising:

generating a plurality of points that each represent a point on a surface of the scene;

determining a best fit of a sphere for a group of the points using a quadric surface fit and global error minimization.

9. A method for modeling a three-dimensional scene, comprising:

generating a plurality of points that each represent a point on a surface of the scene;

determining a best fit quadric surface for a group of points; and

determining which geometric primitive of a plurality of the family described by the quadric surface best fits the group of points.

10. A method for merging two geometric primitives of the same type to form a single geometric primitive of the type, comprising:

creating a new group of points by combining the points used to originally fit each of the two primitives; and

fitting the new geometric primitive using any appropriate fitting technique and the newly generated point group with points from each of the original primitives.

11. A method of registering a first model, consisting of a plurality of points and geometric primitives and having a first coordinate system, with a second model, consisting of a plurality of points and geometric primitives and having a second coordinate system, comprising:

identifying by a user common features of the first and second scenes;

identifying a transformation between coordinate systems that is responsive to the identification; and

transforming the objects of the second model so that they use the first coordinate system.

12. A method of warping, comprising:

selecting one or more models represented by a plurality of point clouds and geometric primitives;

specifying constraints on the locations of any number of points or geometric primitives;

creating an artificial volume that surrounds the points and geometric primitives in each view and assigning mechanical material characteristics to the surrounding volume;

computing a minimum energy configuration for the material in the surrounding volume in which the points or geometric primitives are embedded such that the configuration satisfies all applied constraints; and

displacing the points and primitives in accordance with the computed minimum energy configuration of the surrounding volume of material.

13. The method of claim 12, wherein the constraints are specified to eliminate closure errors.

14. An apparatus for acquiring three dimensional information from a remote object comprising:

a scanning laser module for measuring position information of the object;

a video module for capturing image information from the object; and

a processor for rendering a model of the object which includes the position information and the image information.

15. An apparatus as recited in claim 14 wherein the video image information is collected in a spatially coincident manner with the measurement of position information.

16. An approach as recited in claim 15 wherein the video image information is collected from points adjacent to the points where position information is obtained.

* * * * *