To appear in IEEE/RSJ International Conference on Intelligent Robotic Systems (IROS'98)

A Photo-Realistic 3-D Mapping System for Extreme Nuclear Environments: Chornobyl

M. Maimone, L. Matthies

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91109
mark.maimone@jpl.nasa.gov

J. Osborn, E. Rollins, J. Teza, S. Thayer

Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213
sthayer@ri.cmu.edu

http://www.frc.ri.cmu.edu/projects/pioneer/

Abstract

We present a novel stereoscopic mapping system for use in post-nuclear accident operations. First we discuss a radiation shielded sensor array designed to tolerate extended cumulative dose using 4× shielding. Next, we outline procedures to ensure timely, accurate range estimation using trinocular stereo. Finally, we review the implementation of a system for the integration of range information into a 3-D, textured, metrically accurate surface mesh.

1 Introduction

In the early morning of April 26, 1986 the worst nuclear accident in history occurred in reactor Unit 4 of the Chornobyl power station [1]. After a heroic initial response to cope with the immediate emergency, an external containment structure, or "shelter" (often referred to as the "sarcophagus" in the West) was constructed over and around the damaged reactor building in little more than six months. An international effort to remediate Unit 4 is being mounted, but major difficulties will have to be overcome.

Specifically, the integrity of the shelter, which rests upon the damaged structure of Unit 4, is dubious. In light of this deterioration, any remediation activities must ensure that no further environmental damage occurs and that areas of major structural weakness are identified and reinforced. To aid in this activity, a robot equipped with a functional three-dimensional mapping system is being constructed to provide the shelter operators with a safe, reliable, and effective means to document the current state of the sarcophagus interior.

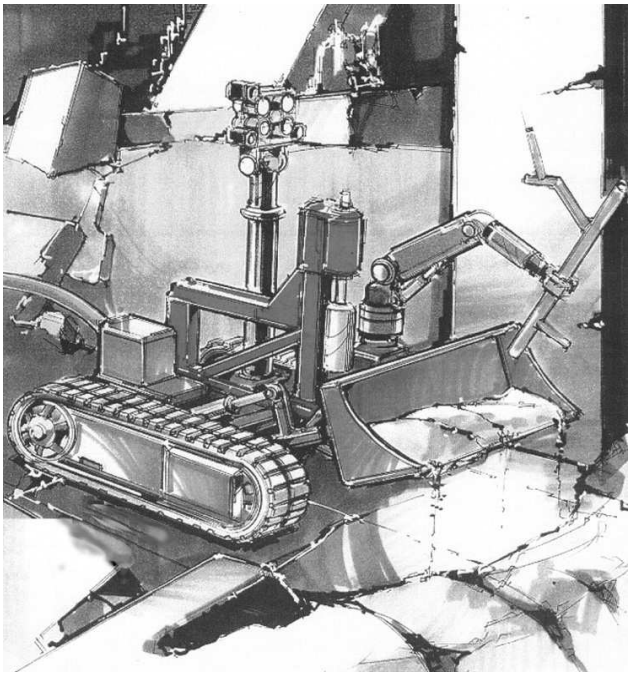

Figure 1: Visualization of the Pioneer robot

The **Pioneer** robot (designed by RedZone Robotics, Inc.) consists of a tracked vehicle approximately 1.2 × 0.75 m with a centrally located mast 1.4 m high [9]. The two tracks, operated by electric motors, can perform skid steering and climb a slope of 45 degrees. The vehicle supports one of several payloads including the mapping sensor package discuessed in this paper. Other sensors mounted on the vehicle

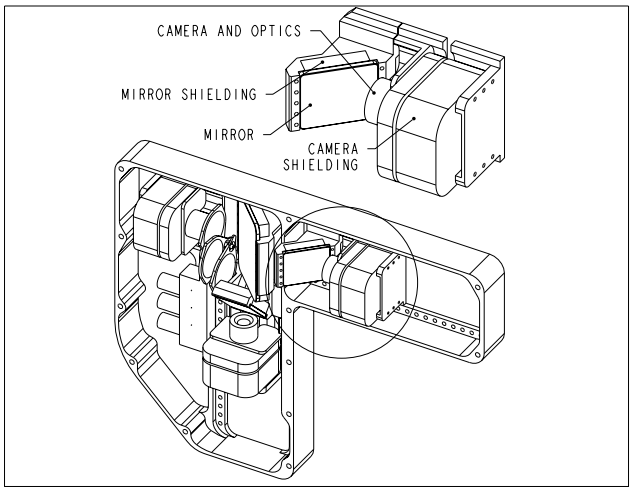

Figure 2: **Mapper Camera System Assembly:** Detail shows typical camera assembly.

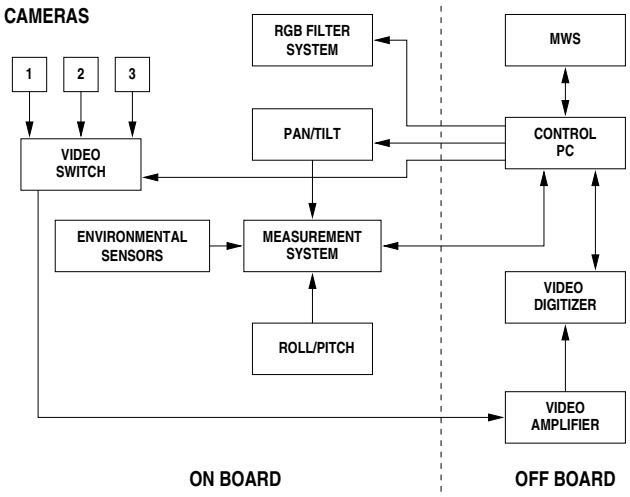

Figure 3: **Schematic of the Pioneer Stereo Mapping System:** Components in the left half of the figure are mounted on board the robot.

measure gamma and neutron radiation, ambient temperature and relative humidity. Electrical connections between the vehicle and the control electronics are via a tether of 100 meters in length. In addition to the mapping cameras on the vehicle, a color vidicon camera with zoom lens is used for teleoperation. Finally, a "core-boring" assembly designed at Carnegie Mellon and an associated material analysis system from the Jet Propulsion Laboratory provides structural information.

The primary goals of the **Pioneer** mapping system are to:

- Perform a complete, accurate three-dimensional mapping of the designated target areas.
- Co-register auxiliary sensory information with the 3-D surface maps.
- Integrate color texture information with the 3-D surface maps to provide a photo-realistic rendering of the target facilities.

2 A Rugged Stereo Rig

The hardware of the mapping sensor must withstand an environment which includes radiation, decontamination procedures, and the vibration and shock resulting from vehicle motion. The gamma radiation field will result in a 10^6 R total accumulated dose over the planned robot mission with dose rates up to 3.5 kR/hr. Average dose rates in Unit 4's room 305 range from 500 - 1000 R/Hr giving the robot a mean lifetime of 1000 hrs (maximum). Neutron dose rates experienced by the components are from 1 to 1000 cm^{-2} sec^{-1}. Electronic components on the vehicle must be either radiation hardened, shielded, or considered sacrificial. The induced effects of radiation must be negligible for the expected dose rates, or transient so that a given mapping session is not interrupted by equipment failure or poor performance of a component. Components deemed sacrificial must be inexpensive enough to be replaced frequently and must perform within acceptable limits over their expected lifetime. The system must be designed to allow any sacrificial components to be easily replaced under field conditions. The mapping camera system is removable as a unit so that needless exposure of the cameras is avoided when only non-mapping operations are performed by the vehicle.

The raw data for the **Pioneer** mapping system is provided by three CCD cameras using stereo vision. This has several advantages over competing technologies such as laser scanners, radar sensors, and structured light rangers. Stereo imagery operates passively, requires no moving parts (except for the pan/tilt unit on which it is mounted), requires little power, allows processing to be performed off-board, requires a relatively small number of images, and enables direct registration of video images onto range maps.

2.1 Camera Selection

The requirements of the mapping system and the limitations on the number of conductors contained in the 100 meter vehicle tether combined to preclude the use of commercially available radiation hardened cameras. Fixed image geometry with respect to the optical coordinates is necessary to preserve stereo calibration. Vidicon or other tube based imaging technologies, although radiation hard, do not have sufficient image stability to be used for precision stereo applications because their deflection ciruitry can drift when affected by fluctuating external magnetic fields. Alternatively, solid state detectors are available with appropriate radiation hardness and stable image geometry, but require complex control interconnection between the hardened sensor and non-hardened support circuitry. Therefore, this design reduced the amount of on board electronics by using readily available CCD board level cameras inside a shielding envelope.

The amount of shielding is constrained by the capabilities of the pan/tilt unit carrying the cameras, but provides sufficient attenuation to result in satisfactory performance over the limited lifetime of the cameras. Tests were performed in which board cameras were exposed to a Cs 137 gamma source with dose rates up to 3 kR/hr with noticeable induced noise. The useful lifetime of these cameras is at between 50-100 hours at the expected dose rates. In addition, the component cost of the cameras is sufficiently low to be neglected over the course of the mission.

The three cameras composing the stereo sensor array are rigidly mounted within a sealed, aluminum housing in the shape of an inverted "L" as shown in Figure 2 , which is itself mounted on the pan/tilt head. This housing is machined internally to provide prealigned mounting locations for the cameras allowing the baseline of the stereo images to be varied continuously from 15 to 30 cm during the development phase of the project. The centrally located camera is fitted with a filter wheel that can place any of three filters in the optical path to generate a composite color image of the scene. Not shown in the figure is the front section of the camera housing containing three ports of fused silica. Also not shown are a set of four waterproof quartz halogen lamps mounted on the pan/tilt unit. Two lamps are mounted above and two below the camera array.

Figure 3 is a block diagram showing the portions of the system mounted on board the robot and off board at the control station. Angular position of the pan/tilt unit, robot orientation and environmental data are acquired using the on board measurement system. The control PC is able to select which signal from any of the three cameras will be acquired via the video switch, control the filter mechanism and drive the pan/tilt unit. The control station consists of two computers: a PC supporting a video acquisition board and the mapping work station (MWS in the figure).

2.2 Shielding

The shield of each camera consists of 1.25 cm lead, covering the CCD board camera, lens, and a mirror. This thickness of lead results in a 4× reduction in incident gamma radiation. An aperture is provided in the shield for the lens. The mirror is placed at 45 degrees to the optical path in front of the shield aperture and folds the optical path by 90 degrees. In this design the lead behind the mirror shields the camera and optics from radiation incident along the longitudinal axis of the lens. Radiation will reflect off the mirror into the optics, but also at a 4× reduction.

2.3 Baseline and Field-of-View

The field of view was selected by balancing two competing goals. Horizontal precision in the depth maps increases as the field of view narrows, suggesting that a smaller field of view would be better. But the number of images required to completely map an area increases dramatically as the field of view shrinks, so a wider field of view is desired for more timely mapping operations. 35 degrees was found to be an appropriate compromise.

The baseline required between the stereo cameras is determined from the desired depth resolution of the system and its field of view, according to this approximation:

$$b \approx \frac{z^2 r \Delta\phi}{\Delta z} \quad (1)$$

where z is the distance of a world point from the camera head (5 m maximum), r is the sub pixel resolution expected from the stereo system (0.2), $\Delta\phi$ is the approximate field of view of each pixel (1 milliradian), and Δz is the desired depth resolution (3 cm). Thus for this application, a baseline of approximately 16 cm is most appropriate.

3 Trinocular Stereo Processing Suite

The stereo vision sensor provides the raw data required by the **Pioneer** mapping system to perform

accurate world modeling. The software component of this sensor has been demonstrated to provide automatic range data with enough accuracy to enable autonomous vehicle navigation from binocular stereo imagery [8]. However, the indoor setting of this project offers unique challenges, so certain parts of the automatic system have been made accessible to the user, providing the ability to adjust the environment and tailor the stereo parameters to individual sets of stereo images. Also, a third camera was added to help reduce disparity artifacts that might occur at horizontal object boundaries; these are much more common in man-made environments than in natural terrain.

3.1 Automatic Stereo Processing

Effective stereo processing is enabled by accurate camera calibration. Our camera model is Gennery's extension of the Yakimovsky/Cunningham linear model [10], and includes radial lens distortion terms [7]. Camera calibration parameters are aquired by imaging a specially constructed target (three 1 meter^2 faces of a cube, each with a 10x10 grid of features) and running the calibration software. This target will be delivered as part of the **Pioneer** mapping system, to allow for field verification of calibration and recovery from any on-site changes to the camera head.

The stereo algorithm used by the **Pioneer** mapping system is a coarse-to-fine correlation procedure using 2-D windows. It will not be described in detail here, but can be found in [8].

3.2 Range Expert

To assist the user in the acquisition of appropriate map information, a range image evaluator will provide feedback concerning the quality of the raw images and their resulting range data. Some examples of perceived problems and suggested solutions include:

Pixels are saturated: Turn the light level down
Pixel intensity range too small: Turn the light level up
Sparse range data: Adjust lights to improve image texture, or reposition vehicle closer to target, or Verify calibration.

3.3 User Controls

Several of the stereo processing parameters and environmental controls have been made available to the user to allow for post-processing refinement of the range maps.

Lighting: Much of the intended mapping operation will take place in an unlighted portion of the reactor facility. Since stereo processing depends on the quality of image texture visible in a luminance image, the user is given dimmer control of the four lights mounted on the stereo camera pan/tilt head. The range expert module will offer suggestions as to how the lighting might be adjusted to improve image quality.

Range Bounds: The mapping system was designed to provide range information to a given precision over a fixed range in front of the camera head. The user may increase these bounds if range data at greater distances is desired (at some loss of resolution), or narrow them to reduce the potential for correlation mismatch if all objects are known to lie within more restrictive bounds.

Correlation Window Size: Although not recommended, the user may change the size of the correlation window used in the stereo computation.

Pyramid level: The pyramid level indicates the size of the resulting range map. Pyramid level 0 will generate a range map at full resolution, 1 means the range map will be half as large in each dimension (and thus only one quarter the size), and so on. Higher pyramid levels reduce dependence on the calibration and speed up processing, but provide a smaller number of polygons to the mesh generator.

4 Flexible 3-D Cartography

A mapping scenario consists of maneuvering the **Pioneer** robot to positions within a particular target location. At each position an overlapping set of range and color imagery will be acquired, fully exploiting the available pan and tilt angles. The primary goal of the cartography system is to integrate an entire set of range and intensity imagery from multiple mapping positions and merge them into a textured, three-dimensional rendering of the target space. Data sets acquired from the same robot positions are trivially referenced to the local robot coordinate frame using the pan/tilt angles. However, no direct or accurate transformation between separate robot positions is made available to the mapping system from the **Pioneer** robot. In such cases where no global estimate of position is available, the mapping system has the ability to register two related data sets, and to derive the "best" possible estimate of the transformation between them. The **Pioneer** mapping system offers these capabilities through a suite of mesh processing and analysis tools.

Figure 4: **Fully Processed Textured Mesh:** Input points have been triangulated, cleaned, smoothed and resampled.

Figure 5: **Fully Processed Textured Mesh:** Input points have been triangulated, cleaned, smoothed and resampled.

4.1 Mesh Creation and Processing

Surface meshes are a common way of representing 3-D data, and when the points from a 3-D cloud have an easily recovered, common coplanar projection, such as in stereoscopic imagery, creating a simple mesh is a relatively easy process. However, creation of a well-proportioned mesh more suitable for rendering from multiple viewpoints, is a more daunting task. Pioneer's engine for mesh generation exploits Shewchuk's *Triangle* algorithm for the production of constrained Delaunay triangles [6]. The overal mesh processing in the mapping software consists of four separate stages: *Texturing*, *Cleaning*, *Smoothing*, and *Resampling*.

Texturing: Application of texture gives reconstructed scenes a more realistic feel, and allows the user to more easily recognize salient features. It is done by relating three dimensional mesh points to the co-registered color image obtained by sweeping the RGB filters over the central camera of the stereo rig.

Cleaning: A mesh may have several unwanted features upon creation, such as small, insignificant noise patches and jagged boundaries. Cleaning is the process of removing long edges, small unconnected surface patches, and restructuring the boundaries.

Smoothing: Range data can be noisy. Mesh smoothing attenuates the high frequency noise structure and (ideally) preserves important mesh structures and boundaries.

Resampling: Most 3-D vision applications perform either more efficiently or more accurately given a uniformly sampled surface at a lower resolution than the original mesh data. Resampling takes a clean, smooth mesh and outputs a lower resolution, more uniformly sampled mesh [3].

The mesh processing sequence depends only upon the acquisition of a single range image referenced in sensor coordinates. Subsequent processing stages group and assemble individual meshes into a coherent 3-D map. Figures 4 and 5 show renderings of input meshes after application of the mesh processing suite.

4.2 Position Integration Module

At a single robot position, the relative transformations between individual sensor positions can be recovered from the pan and tilt angles. Given this, a method for taking individual sets of mesh data and processing them to produce a *seamless* and *uniformly-sampled* "super mesh" is required. The fundamental technology employed by **Pioneer** uses a sensor and noise model for applying the vertices of the surface mesh into a 3-D occupancy grid [2]. Each surface mesh from the current robot position is transformed to a local sensor coordinate reference and sequentially distributed into the voxel volume. Surface normal information is used to extract a replicate, integrated surface from the grid. Overlapping texture components are resolved using the angle between the viewpoint and the surface normal of the mesh face. Figure 6 displays the integration of Figures 4 and 5 together with some additional data.

4.3 Registration Module

No direct means of measuring the transformation from two distinct robot positions is available from

the **Pioneer** robot; however, technology to compute the relative transformation from two featured surface meshes has been developed at Carnegie Mellon. This method derives an approximate transformation between two surface meshes through a correlation of "spin-images" obtained at oriented points on the surface of the mesh [4]. This rough transformation is then refined using a modified version of the Iterative Closest Point [5] algorithm that registers 3-D and color (texture) jointly. Figure 7 illustrates a sample panorama of the catacombs using commercially available software. It is shown for comparison with the 3-D panorama displayed in Figure 6.

These integrated meshes provide the raw data for the data visualization and analysis components of the **Pioneer** Mapper System developed by the Intelligent Mechanisms Group at the NASA Ames Research Center and the Grok Lab at the University of Iowa.

5 Summary and Conclusions

Three components of an integrated system for 3-D mapping in extreme environments were presented. A rad-hardened stereo rig and its associated electronics designed to tolerate 10^4+ R form the core of the acquisition system. An interactive trinocular stereo system outputs timely and accurate range information to a linear precision of 3 cm. Finally, a 3-D cartographic system for conversion of raw range data into integrated 3-D surface structure maps was presented and preliminary results were detailed. The **Pioneer** robot will be deployed in the Chornobyl facility during the winter of 1998/1999.

Acknowledgments

The authors would like to thank Frank Ruddy for his help in characterizing the radiation sensitivity of our cameras and the use of the Westinghouse hot cell. We would also like to thank Andrew Johnson for his mesh tool box, Todd Litwin for his help in porting the stereo code, Eric Zbinden for modifying the existing mesh processing suite to accept textured information, and the rest of the **Pioneer** mapping team at NASA Ames Intelligent Mechanisms Group, the Iowa Grok Lab, and RedZone Robotics for their input, advice, and criticisms. The work described in this paper was partially carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract to the National Aeronautics and Space Administration. The **Pioneer** project is jointly sponsored by the US Department of Energy and the National Aeronautics and Space Administration.

References

[1] A. R. Sich *The Chornobyl Accident Revisited: Source Term Analysis and Reconstruction of Events During the Active Phase,* Dept of Nuclear Engr., Massachusetts Institute of Technology Ph.D. Dissertation, 1994.

[2] A. E. Johnson and S. B. Kang "Registration and Integration of Textured 3-D Data", International Conference on Recent Advances in 3-D Digital Imaging and Modeling, Ottawa, Ontario, May 12-15, 1997.

[3] A. Johnson and M. Hebert,*Control of Polygonal Mesh Resolution for 3-D Computer Vision.* Graphical Models and Image Processing, To appear.

[4] A. Johnson and M. Hebert. *Recognizing objects by matching oriented points,* June 1997 International Conference on Computer Vision and Pattern Recognition, San Juan. pp. 684–689.

[5] P. J. Besl and N. D. McKay. *A Method of Registration of 3-D Shapes.* IEEE Trans. on PAMI, Vol. 14, No. 2, Feb 1992.

[6] J. R. Shewchuk. *Engineering a 2D Quality Mesh Generator and Delaunay Triangulator.* First Workshop on Applied Computational Geometry, Philadelphia, PA. pp. 124-133, ACM, May 1996.

[7] D. B. Gennery. *Least-Squares Camera Calibration Including Lens Distortion and Automatic Editing of Calibration Points.* Springer-Verlag, To be published.

[8] L. Matthies, A. Kelly, T. Litwin, and G. Tharp. *Obstacle Detection for Unmanned Ground Vehicles: A Progress Report.* Springer-Verlag, 1996.

[9] RedZone Robotics, Inc. *Functions and Requirements for the Pioneer System.* Document number 96055-REPT-004.2, December 8, 1997.

[10] Y. Yakimovsky and R. T. Cunningham. A system for extracting three-dimensional measurements from a stereo pair of TV cameras. *Computer Graphics and Image Processing*, 7:195-210, 1978.

Figure 6: **Integrated Mesh:** A sample integration of three meshes, including Figures 4 and 5. Note that although this is only a 2D rendering, 3D information is known about all visible pixels.

Figure 7: **Integrated Panorama:** A sample integrated panorama generated using commercial software for comparison with the 3-D mesh integration used in the previous figure.