Published in the Proceedings of the ASPRS Conference, St. Louis, Missouri, April 23rd to 27th, 2001.

DEVELOPMENT OF A SOFTWARE PROTOTYPE FOR THE GEO-REFERENCING AND VISUALIZATION OF NON-METRIC TERRESTRIAL PHOTOGRAPHY IN A GIS ENVIRONMENT

Stephen Rawlinson, Y.C. Lee, Y. Zhang
Geographical Engineering Group
Department of Geodesy and Geomatics Engineering
University of New Brunswick
P.O. Box 4400, Fredericton, New Brunswick, E3B 5A3, CANADA
stephen.rawlinson@unb.ca

ABSTRACT

Conventional approaches to adding realism in a GIS environment involve the development of complicated 3D geometric models through the use of sophisticated computer hardware and software. While these approaches provide for some benefits with regard to increased user comprehension, they are often limited due to the complexity of their creation and inability to provide realistic visual cues for the user. This is especially significant in the development of real-time navigation systems, where the uninitiated user must be able to quickly and efficiently interpret directions provided on a computer display. This paper focuses on the integration of digital terrestrial photographs in a map-based (GIS) environment acquired with a non-metric camera (conventional and/or digital) which can be readily understood and appreciated by the non-specialist. Standard stereo-photographic techniques are used to link 360-degree panoramic virtual environments to a dynamic GIS database within a software prototype. The linked panoramic and map interface allows for user query and interaction. Techniques and results are outlined for the creation of the system, including: acquisition, processing (data reduction), and visualization. Ease of use and low cost were primary considerations for the development of the prototype. The overall performance of the system is considered and future development work is explored.

INTRODUCTION

Virtual reality (VR) is gaining in popularity as a useful visualization technique. VR systems allow for the creation of virtual environments, which place users in a computer simulated environment allowing for interaction (El-Hakim *et al.*, 1998). VR systems have traditionally been developed and designed within the computer graphics community, for example CAVE (Cruz-Neira *et al.*, 1993); and the virtual workbench (Kruger *et al.*, 1995). These systems render 3D geometric models generated from secondary sources such as 3D digitizing tools, rangefinders, and stereo photogrammetric techniques. Surface texture shading or environment maps are subsequently introduced to the models to increase realism (Kang, 1998).

Geometric Modeling

The above VR approach, referred to as geometric modelling (GM), has been adopted by the GIS and cartographic communities and is a growing area of active research (Germs *et al.*, 1999; Hearnshaw and Uniwin, 1994; Huang and Lin, 1999; Rhyne, 1997; Unwin, 1997). Three-dimensional VR GIS is largely focused on the visualization of geographic scenes to mimic human perspective views (Raper *et al.*, 1999). The popularity of GM VR and GIS can be largely attributed to the decreasing cost and increasing availability of powerful rendering hardware and software, in conjunction with a general awareness in these communities that 3D visualization dramatically increases the level of understanding for the end user. Compared with standard 2D planimetric maps oriented to the north, 3D scenes present almost unlimited viewing perspectives. The availability of commercial GIS software products supporting 3D visualization, such as ESRI's ArcView 3D Analyst extension and ERDAS's VirtualGIS, and the development of a 3D "geographic" modelling language, Virtual Reality Markup Language, or VRML, typify this trend (ESRI, 2001; ERDAS, 2001; VRML, 2001).

Image-based Rendering

Recently however, a new VR approach, called image-based rendering (IBR), has emerged that renders photo-realistic views depending on the user's observation location (Chen, 1995; Szeliski and Shum, 1995; McMillan, L. and Bishop, G., 1995). Views are represented as a mosaic or collection of images and new views created by interpolating and/or reprojecting input images onto target surfaces such as cylinders or spheres (Szeliski and Kang, 1995). As Kang (1998) suggests, this contrasts with the GM approach where the typical rendering process relies on modeling transformation, view transformation, culling (deciding on and displaying what is theoretically visible), and finally hidden surface removal. This is an important difference since increased realism requires increasingly complex geometric models, and thus the cost of rendering in a GM VR can be high since rendering time is a function of the scene complexity. In fact, the GM approach is well known to require "laborious modeling and special purpose software" for effective realistic view rendering (Chen, 1995). A comparison of the GM and IVR approaches in the context of GIS is presented in Table 1.

Table 1. Comparison of the geometric modeling and image-based rendering approach to virtual reality (adapted from Kang, 1998).

Geometric modeling approach	Image-based rendering approach
Complex 3D geometric data structures	Set of images
Conventional rendering	Reprojection/Interpolation
Sophisticated hardware/software for added realism	Realism function of input scenes
Expensive inputs	Inexpensive inputs
Query support	Limited query support
Link to GIS well developed	Link to GIS less developed

Panoramic Virtual Reality

Perhaps the most widely known and available IBR technique is panoramic virtual reality, or PVR. This novel VR approach allows for complete 360 degree panning and viewing around a given observation point by warping a set of input images to simulate a user's perspective view. The set of input overlapping images are generally acquired around a rotation point by consecutively panning a camera until complete 360-degree coverage is obtained. These images are subsequently stitched together and warped onto a cylinder to form a continuous mosaic (refer to Figure 1). Using a standard desktop PC and appropriate software, realistic scenes can be rendered (re-projected from the cylinder onto a plane) "on-the-fly" (Apple, 2001; IPIX, 2001).

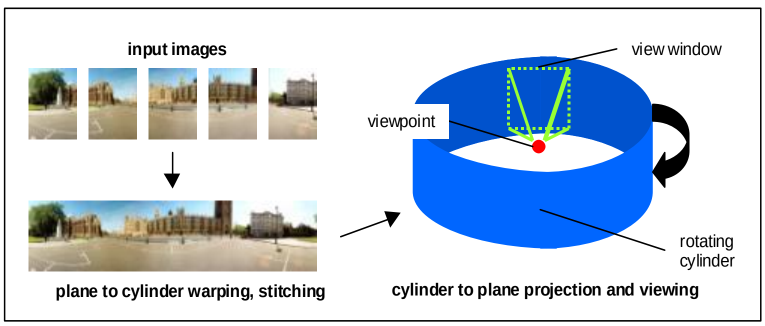

Figure 1. Overview of cylindrical panoramic imaging process

In addition, static "hot-spots" can be created that identify pixel regions on a panoramic image that support additional interaction, such as WWW navigation or activating actions (Chen, 1995). The "hot-spot" concept, while seemingly useful in providing GIS linking capability, are simply user defined pixel regions and thus have no geographically referenced meaning.

Unlike the GM approach, the integration of panoramic VR and GIS is less developed. For example, Chapman and Deacon (1998) used panoramic imagery to supplement traditional 2D and 3D CAD databases; while Dykes (2000) integrated panoramic imaging to a geographic base to provide bearing information in the context of a virtual field course. While these approaches are advantageous over GM techniques due to their simplicity in design and their added realism, they fail to effectively take advantage of the full potential of a dynamic link between a photo-realistic VR environment and a spatial database. The design of a system that generates valuable coordinate information within the PVR environment for linking with a GIS is the focus of this research.

A New Approach to VR GIS

Due to the aforementioned inadequacies of existing panoramic virtual reality and GIS integration approaches, the objective of this research is to develop and test a complete methodology for acquiring, processing, and displaying panoramic images that are linked to a GIS environment. The idea here is to construct a prototype that presents a user with two views of a scene (a standard 2D overhead view and an interactive 360 degree panoramic view) that are dynamically linked such that user interaction in one view is reflected in the corresponding view. Based on this concept, valuable spatially linked attribute information from the GIS can be displayed (Figure 2). In addition, the following are key design considerations: 1) low cost; 2) easily available inputs (no object control and little or no calibration); 3) simplicity and ease of use for the non-specialist; 4) adequate accuracy (+/- 1-2 metres) for the purposes of GIS integration and, 5) robustness and reliability. It is anticipated that a full working prototype would find use in interactive touring and navigation guides, as well as planning and view shed visualization.

Figure 2. Conceptual representation of GIS and panoramic virtual reality integration

The research outlined in this paper presents an alternative approach for GIS and IBR virtual reality integration that provides a true link between the image scene and the GIS database through a prototype georeferenced panoramic imaging environment. This system takes advantage of simple stereo photogrammetric principles and image processing techniques that provide proof of concept for seamless virtual reality and GIS integration. This paper presents 1) the prototype system description; 2) accuracy considerations of the system, and; 3) future directions of the research.

SYSTEM DESCRIPTION

The prototype system outlined in this paper consists of a simple tripod, a non-metric off-the-shelf camera (conventional Ricoh FF-3 35mm automatic focus camera, nominal focal length = 35mm), and a measuring device (optional), along with a series of software modules developed by the Geographical Engineering Group, UNB. A set of stereo-pairs corresponding to a complete 360-degree rotation around a desired viewpoint are first acquired using a tripod mount. The images are entered as input into the software which 1) warps and stitches the imagery into a cylindrical panoramic mosaic; 2) processes the stereo-pairs for further distance calculations (through space intersection); and, 3) displays and renders the mosaicked imagery into the integrated panorama and GIS system for subsequent user query and interaction. The linking of the GIS and the panoramic viewer is accomplished internally through the automatic calculation of geographic coordinates from the input stereo-pairs. At the time of writing, the entire process is not completely automated and requires direct user interactions. However, it is expected that future research will yield a fully automated system. Further, no object space control or camera calibration is currently necessary. Details of the process are documented the following sections.

Image Acquisition

The prototype system is based on the cylindrical panorama model due to the simplicity in image acquisition and the relative ease in projecting from the cylinder to the plane (and vice-versa). Stereo-pairs and panoramic image sequences can be taken simultaneously using a specially adapted stereo rig or by moving the tripod in carefully determined intervals to ensure consistent arc distance around a complete horizontal circle (Figure 3). The left camera position stays fixed about its nodal point while the right camera position traces the arc in successive steps at radius r

(where r = the stereo base). The number of photographs taken in a function of the camera's field of view and the percentage overlap of successive shots.

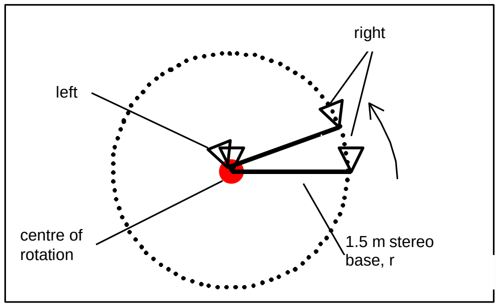

Figure 3. Overhead view of rotating panoramic and stereo-pair image acquisition

In the testing of this prototype, a single tripod (without a nodal head) and camera setup with 1.5 metre stereo base separation was used. In total, 24 photos were acquired at two different testing locations (test sites A and B). Successive left stereo pair images have a consistent overlap (50%) to ensure effective mosaicking for subsequent panoramic warping. In this design, the right image pair is used exclusively for subsequent stereo model and space intersection calculations. Ideally, the left image should rotate about its nodal point (optical centre) to eliminate, through the use of a panoramic head, the potential for parallax in the sequence of left images. The introduction of parallax *within* the left stereo pair image sequence can make it difficult to stitch the sequence together. Furthermore, the left and right optical axes should be parallel and perpendicular to the surface plane. In practice, a slight misalignment of either the nodal point or axes was unavoidable but can be tolerated. As such, a slight x-parallax was noticeable in the sequence of photos taken by the left camera for both test sites A and B. This did present some complications in panoramic warping and alignment.

Panoramic Warping

The 48 images scenes were processed, developed, and scanned commercially using the Kodak PhotoCD system to a digital image resolution of 1536 x 1024 (Kodak, 2001). The left set of stereo pairs were then projected on a cylinder and stitched to form a complete mosaic using software developed by the authors.

Projection of a plane to a cylinder. The mapping of a plane to a cylinder is a well-known and understood geometric concept. In fact, there are numerous commercially available products capable of warping a sequence of overlapping images into a cylindrical panoramic image. However, for the purposes of this research, it was felt that these products were not easily modified or adapted. This presented problems in subsequent processing steps, such as stereo matching. The algorithm developed for plane-to-cylinder warping is as follows: given a pixel in the projected image, the corresponding pixel location (and thus set of RGB brightness values) in the planar image is computed. The following sets of equations are used to convert from planar to cylindrical (x,y) coordinates:

$$
\begin{aligned}
x &= pd\,\frac{x'}{y'} + x_{centre} \\
y &= pd\,\frac{y'}{z'} + y_{centre} \\
&\text{where :} \\
x' &= \sin\left(\frac{x_{cylinder} - x_{centre}}{pd}\right) \\
y' &= \frac{y_{cylinder} - y_{centre}}{pd} \\
z' &= \cos\left(\frac{x_{cylinder} - x_{centre}}{pd}\right) \\
pd &= \text{estimated principal distance} \\
x_{centre}, y_{centre} &= \text{image centre (plane) in pixels} \\
x_{cylinder}, y_{cylinder} &= \text{cylinder coordinates in pixels}
\end{aligned}
\qquad (1)
$$

Warping and Alignment. Each of the images (12 for each test location corresponding to the left camera position of the stereo pair) were subsequently warped and then semi-automatically aligned using a similar methodology as that outlined by Lucas and Kanade (1981) in a pair-wise fashion. Aligning the images is necessary for a seamless panoramic output image and in practice only a simple *x*, *y* translation was required. The determination of the translation relationship is essential since it provides an automatic way of stitching the input images into a complete mosaic. The approach used here incorporates image intensity information to determine a set of matching locations in the overlapping areas of each image. Given a pair of images to be matched the approach equates to a normalized cross correlation function in the form:

$$\rho = \frac{\sigma_{xy}}{\sqrt{\sigma_{xx}\sigma_{yy}}}, -1 \le \rho \le 1$$
$$\bar{x} = \frac{1}{n^2}\sum_{i=1}^{n}\sum_{j=1}^{j} x_{ij}, \qquad \bar{y} = \frac{1}{n^2}\sum_{i=1}^{n}\sum_{j=1}^{j} y_{ij} \tag{2}$$
$$\sigma_{xx} = \frac{1}{n^2}\sum_{i=1}^{n}\sum_{j=1}^{j} x^2{}_{ij} - \bar{x}^2, \quad \sigma_{xy} = \frac{1}{n^2}\sum_{i=1}^{n}\sum_{j=1}^{j} x_{ij}y_{ij} - \bar{x}\bar{y}, \quad \sigma_{yy} = \frac{1}{n^2}\sum_{i=1}^{n}\sum_{j=1}^{j} y^2{}_{ij} - \bar{y}^2$$

In order to compute this function, a template window is shifted pixel by pixel across a larger search window, and in each position the cross-correlation coefficient between the template window and the corresponding region of the search window is computed. The maximum of the cross-correlation defines the position of best match between the template and the search window. A 7 by 7 pixel template and a correlation threshold of 0.95 provided effective translation estimates.

Matching points must be extracted from the input imagery with manual intervention if the images are considerably misaligned. Of the 12 image pairs processed at each test location (=24 in total), 14 pairs required manual intervention. A good initial selection of a high contrast and well-defined pixel region was essential for effective alignment. Further designs are expected to be completely automatic through the incorporation of an initial selection algorithm (using intensity gradient information).

Following the determination of the translation between image pairs, resampling is performed using the nearest neighbour technique since it provides for optimal efficiency while still preserving original input values (useful for subsequent stereo matching). Blending and cropping follow to generate a smooth and seamless output image. Results (Figure 4) indicate that some experimentation with the blending and resampling may be required in overlapping areas that are slightly shifted due to moderate parallax.

Figure 4. Sample results of automatic warping, stitching, and blending for cylindrical panorama creation for test site B

Space Positioning

It is a well-known photogrammetric principle that if images of an object point appear in two or more images, the position of the point can be determined. This, of course, only holds if the camera positions used to acquire the images and the optical axes directions are known. Distances from the left camera position to any image object (X_p, Y_p, Z_p) selected by the user are automatically calculated using the following sets of equations:

$$X_p = \frac{Y_p}{pd} x_1$$

$$Y_p = \frac{Bpd}{x_1 - x_2}$$

$$Z_p = \frac{Y_p}{pd} y \tag{3}$$

where :

x_1 = x position of the object in the left image (panorama)

x_2 = x position of the object in the right image

pd = estimated principal distance

B = stereo base separation

These equations constitute the "normal case" and assume that the two camera positions have optical axes parallel to each other and perpendicular to the stereo base separation B. To calculate the real-world distance of an object in the image, a pixel location (representing x_1) is first selected by the user in the panoramic image. This x, y pixel position of the object in the panoramic image is converted to original input image coordinates through a reverse process similar to that outlined in (1). This provides the x,y location of the selected pixel in the original left camera image. In regions of overlap, two potential coordinates may be found. In this case, both locations are used in the calculations. Next, the matching point in the right image (x_2) is found using an adaptive image matching approach.

Stereo Matching

Stereo matching, often referred to in the literature as the correspondence problem, is an essential method for automatically obtaining depth information from a stereo pair. Matching essentially amounts to searching one image for the matching location in another image. This is a similar problem to that discussed above for panoramic stitching. However, there are keys distinctions that require alternative approaches. Unfortunately, the technique used previously for panoramic stitching did not provide reliable matches in initial testing. This was likely a result of the fact that this approach relies on a high contrast and well-defined initial pixel region to increase the likelihood of a strong match in the corresponding overlapping image. However in this scenario, the user, and not the computer, selects the initial object in the image with which to match. To try and alleviate this problem, a cross-correlation adaptive template matching technique was designed. Using a coarse window size at first, the best match is found that assists in refining the search at progressively finer levels. While further testing and modifications are necessary, results thus far suggest that this approach provides more reliable matching.

In practice, each stereo pair is no more than 3-4 scan lines offset in the y direction, and thus the matching can be reduced to a 1-dimensional search problem. This is advantageous since matching is performed on-line and computational costs can be high if not adequately constrained. Currently, matching can be achieved in under one second on a Pentium III processor with 198Mbytes of RAM.

Visualization

Numerous panoramic viewers are currently available that can display an image warped to a cylinder. However, as was the case for panoramic warping, only minor modifications are possible with a commercial viewer. This presents an obstacle for effective integration with another software product, such as a GIS. The visualization interface is currently a work in progress. A viewer is being developed by the authors that is similar to viewers currently available yet offers the advantage of being completely customizable. A fully integrated graphical user interface is also currently being developed and is a significant focus.

ACCURACY CONSIDERATIONS

The accuracy of the prototype developed in this research rests on many factors: the stereo rig setup, the camera and lens assembly, the automated stereo matching estimates, as well as the commercial developing and scanning process. These are discussed in the following sections.

Camera calibration

The necessity of camera calibration in non-metric high precision close-range photogrammetry is well established. Camera calibration typically involves the determination of the interior and exterior orientations of the camera and stereo-rig setup respectively. Interior orientation establishes the geometrical relationship between the perspective centre and image plane (principal point, principal distance), while the exterior orientation defines the position and orientation of the image in object space (X, Y, Z object space coordinates, and x-tilt, y-tilt, and swing) (Derenyi, 1996). The process of relative orientation is one aspect of exterior orientation that establishes the orientation of one camera to the other, and thus is based in a local coordinate system (in this case with origin at the left stereo camera position). However, for the purposes of testing this prototype system, a calibrated camera setup (interior and relative orientations) was compared with a non-calibrated setup to determine whether the non-calibrated setup could provide adequate accuracy. A non-calibrated rig is advantageous since it can find use with the non-specialist.

Testing approach. The approach used in this camera calibration testing is as follows: 1) a new set of photographs were acquired at a testing range designed explicitly for camera calibration so that the interior orientation of the camera could be determined; 2) these interior orientation parameters where used to determine the relative orientation parameters for one of the previously collected stereo pairs, and; 3) various distances were calculated (space intersection) using the calibrated and non-calibrated setup and then compared.

Five photographs were acquired of a geodetically surveyed test range using the same camera and commercial development process used for acquisition stage above. The test range consists of 61 3D points on 2 wall planes that have been previously surveyed to acceptable levels of accuracy (95% confidence level) for the purposes of camera calibration (Liu, 1991). Using an iterative Direct Linear Transformation (DLT) self-calibration methodology (Heikkila and Silven, 1998), the principal distance, x and y principal point offsets (x_o, y_o), as well as radial and tangential distortion coefficients ($K1$, $K2$, $T1$, $T2$) were calculated for each photograph (Table 2). This modified DLT technique has the added benefit of directly modeling lens distortion parameters in the overall interior orientation solution.

Table 2. Interior orientation parameters

Photo #	Principal distance (mm)	x_o (mm)	y_o (mm)	*K1*	*K2*	*T1*	*T2*
1	35.4076	0.2920	0.1632	1.09e-4	-3.24e-8	2.11e-5	8.27e-5
2	35.4845	0.4072	0.1012	1.79e-4	-3.92e-8	-5.70e-5	6.98e-5
3	35.7997	0.3126	0.2979	1.64e-4	-2.75e-8	-7.30e-5	2.96e-5
4	35.7063	0.4912	0.2516	1.49e-4	-2.32e-8	-8.47e-5	1.39e-5
5	34.8050	0.3423	0.1014	1.32e-4	-1.42e-8	-6.67e-6	3.95e-5
mean	35.4406	0.3691	0.1831	1.47e-4	-3.53e-8	-4.01e-5	4.71e-5
RMSE	0.3894	0.0809	0.0889	2.74e-5	9.44e-9	4.53e-5	2.85e-5

The mean values of the principal distance, principal point offsets, and distortion parameters of the 5 calibration images were used to determine the relative orientation of each of the first stereo pair from test location B through the direct analytical calibration approach (independent pair relative orientation).

Distance Estimate Comparison

The stereo rig setup is problematic since it is unlikely that the assumptions of "normal case" photogrammetry hold. However, it is worthwhile experiment to test what level of accuracy can be achieved.

Testing Approach. Twenty-two (22) separate distance calculations corresponding to 22 distinct measuring points were obtained on stereo pair 1 from test site location B. These points were selected in the panoramic viewer and distances to the left camera position were calculated automatically based on (3). Further, the same 22 distances from the left camera position were calculated using space intersection using a calibrated and non-calibrated setup

(Figure 4). In this testing, corresponding points were selected manually. A comparison of these distances to manually selected matching points (calibrated and non-calibrated cases) is shown in Figure 4.

Manual selection of matching points. Using the manual selection approach, these results suggest that the non-calibrated stereo setup used in this research can provide adequate accuracy (within the +/- 1-2 metre threshold) for distances shorter than 60 metres from the left camera position (RMSE = 1.28). However, beyond 60 metres error values increase to well over acceptable thresholds (RMSE > 10.0). This is not surprising since the setup outlined in this paper was not rigourous and errors far greater than the measuring accuracy exist in the space intersection solution. In general, closer objects were more accurately determined since more pixels in the image define the object. Currently, a simple prototype stereo rig is being developed that may provide increased distance range through a more permanent rotating "boom" setup. As a result, periodic calibration may be required for supporting applications needing longer range distance calculations. This approach has the added benefit of being easier and more efficient to use.

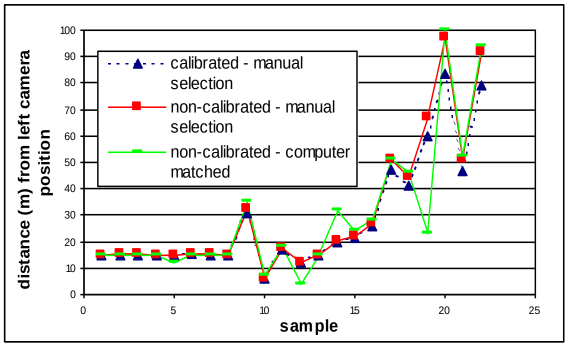

Figure 4. Comparison of manually-derived (calibrated, non- calibrated) and computer-derived distances (stereo pair 1, test location B)

Computer-matched points. The RMS error term for the computer matched versus the calibrated manual selection is high (> 10.0) and exceeds both the measuring accuracy as well as the aforementioned objectives of the prototype. However, as can be seen from these results, blunders in the distance calculations greatly increase the overall inaccuracy. Four computer-derived distances were greater than 12 metres from their calibrated manually derived estimates. There appears to be no systematic over or under evaluation of the computer matched distances. On closer inspection of the input stereo pair, it was revealed that poorly derived computer-matched distances were either in regions of homogenous pixel intensity, or in areas experiencing temporal de-correlation. Temporal de-correlation results from the non-simultaneous acquisition of the stereo pair. Thus, although the scenes overlap, environmental conditions changed and are not consistent between images (examples include: changing atmospheric conditions, or people moving in and out of the scene). Further, it was revealed that the maximum cross-correlation value computed in these regions did not exceed 0.65. Removing these obvious blunders results in an RMS error term of 3.01, which is closer to the desired level of accuracy.

SUMMARY AND CONCLUSION

The prototype described in this paper is a work in progress. However, the results of this initial testing show that +/- 3 metre accuracy (compared with calibrated, manually selected matches) can be achieved for distance estimates under 60 metres using a completely un-calibrated stereo and camera setup, and a computer assisted cross-correlation template matcher. Further, it is suggested here that currently available panoramic software processing tools are not adequate for the purposes of GIS integration. This unfortunately increases the development time and effort necessary for building a fully integrated GIS and PVR system. However, it should be noted that this research has shown a valuable proof of concept for the introduction of panoramic virtual reality into a GIS.

FUTURE RESEARCH

Research and development will continue towards the goal of a fully working prototype. Specifically, the results from this paper suggest that a periodically calibrated stereo rig setup may be useful for applications requiring greater than 60 metre distance calculations. Further, an error propagation model is currently being designed and incorporated as error ellipses directly on the user interface. It is anticipated that this would assist the user in evaluating the quality of the calculated distance estimate. Further, a more reliable matching algorithm will be developed, especially for use in homogenous and low contrast pixel regions.

ACKNOWLEDGEMENT

This research is partially funded by a Natural Sciences and Engineering Research Council of Canada (NSERC) Individual Research Grant.

REFERENCES

Apple. (2001). Apple Computer, Inc. http://www.apple.com

Chen, S.E. (1995). QuickTime VR – an image-based approach to virtual environment navigation. In*: Proceedings of SIGGRAPH '95 Computer Graphics Conference, ACM SIGGRAPH, August*, pp. 29-38.

Chapman, D., and Deacon, A. (1998). Panoramic imaging and virtual reality – filling the gaps between the lines. *ISPRS Journal of Photogrammetry & Remote Sensing*, 53: pp. 311-319.

Cruz-Neira, C., Sandin, D.J., and DeFanti, T.A. (1993). Virtual reality: the design and implementation of the CAVE. In: *Proceedings of SIGGRAPH '93 Computer Graphics Conference, ACM SIGGRAPH, August,* pp. 135-142.

Derenyi, E. E. (1996). *Photogrammetry: the concepts*. Department of Geodesy and Geomatics Engineering, University of New Brunswick, Canada.

Dykes, J. (2000). An approach to virtual environments for visualization using linked geo-referenced panoramic imagery. *Computers, Environment, and Urban Systems*, 24(1): pp. 127-152.

El-Hakim, S.F., Brenner, C., and Roth, G. (1998). A multi-sensor approach to creating virtual environments. *ISPRS Journal of Photogrammetry & Remote Sensing*, 53: pp.379-391.

ERDAS. (2001). *ERDAS Virtual GIS*, ERDAS, Inc. Atlanta, GA.

ESRI. (2001). *ArcView 3D Analyst*. ESRI, Inc, Redlands, CA.

Germs, R., Van Maren, G., Verbree, E., and Jansen, F.W. (1999). A multi-view VR interface for 3D GIS. *Computers & Graphics*, 23: pp. 497-506.

Hearnshaw, H.M., and Unwin, D. (1994). *Visualization in Geographcial Information Systems*. Wiley, Chichester, pp. 243.

Heikkilä, J. & Silvén, O. (1997) A Four-step Camera Calibration Procedure with Implicit Image Correction. In: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'97), San Juan, Puerto Rico*. pp. 1106-1112.

Huang, B., and Lin, H. (1999). GeoVR: a web-based tool for virtual reality presentation for 2D GIS data. *Computers & Geosciences*, 25: pp. 1167-1175.

IPIX. (2001). Internet Pictures Corporation. http://www.ipix.com

Kang, S.B. (1998). Geometrically valid pixel reprojection methods for novel view synthesis. *ISPRS Journal of Photogrammetry & Remote Sensing*, 53: pp. 342-353.

Kodak. (2001). http://www.kodak.com

Kruger, W., Bohn, C-A., Frohlich, B., Schuth, H., Strauss, W., and Wesche, G. (1995). The responsive workbench: a virtual work environment. In: *IEEE Computer Society IEEE Computer July*, 28(7): pp.42-48.

Lucas, B.D., and Kanade, T. (1981). An iterative image registration technique with an application to stereo vision. In: *Proceedings of the 7th Imaging Understanding Workshop*. pp. 121-130.

Liu, X. (1991). *Calibration of non-metric stereo cameras*. Thesis: Department of Geodesy and Geomatics Engineering, University of New Brunswick, Canada.

McMillan, L., and Bishop, G. (1995). Plenoptic modeling: an image-based rendering system. In: *Proceedings of SIGGRAPH '95 Computer Graphics Conference, ACM SIGGRAPH, August*, pp. 39-46.

Raper, J., McCarthy, T., and Williams, N. (1998). Georeferenced four-dimensional virtual environments: principles and applications. *Computers, Environment, and Urban Systems*, 22(6): pp. 529-539.

Rhyne, T.M. (1997). Going virtual with geographic information and scientific visualization. *Computers & Geosciences*, 23(4): pp. 489-491.

Szeliski, R., and Kang, S.B. (1995). Direct methods for visual scene reconstruction. In*: IEEE Workshop on Representation of Visual Scenes, Cambridge, MA, June*, pp. 26-33.

Szeliski, R., and Shum, H.Y. (1997). Creating full view panoramic image mosaics and environment map. In: *Proceedings of SIGGRAPH '97 Computer Graphics Conference*, *ACM SIGGRAPH*, *August*, pp. 251-258.

Unwin, D. (1997). The virtual field course. In: *Second International Conference on GIS in Higher Education, Columbia, MD, September 14*, pp. 4.

VRML. (1996). VRML 2.0 Specifications. http://www.vrml.org