

The Quality of Stereo Disparity in the Polar Regions of a Stereo Panorama

Daniel Sandin¹, Haoyu Wang¹, Ahmad Atra¹, Richard A. Ainsworth², Maxine Brown¹, Thomas A. DeFanti³

¹Electronic Visualization Laboratory, University of Illinois at Chicago, ²Ainsworth & Partners, Madison, Wisconsin, ³California Institute for Telecommunications and Information Technology, University of California San Diego

Abstract

Without considering the pose of the viewer there are problems with displaying stereo correctly at the top and bottom poles of a 4 Pi steradian panorama. This has led to a commonly held position that one cannot handle stereo adequately at the poles of a panorama. In the tracked case, where the pose of one viewer's head is known, it is possible to improve the stereo experience for that viewer utilizing variations on standard VR display methods. It is also possible to display an approximately correct projection without knowing the viewer's pose. There are practical issues with the projections used and the camera placement and stitching methods that affect the quality of stereo disparity, especially at the poles. This paper visually analyzes issues in the display of stereo panoramas with standard VR display methods on both HMDs and panel or projection-based walk-in VR displays.

Dominant Full Sphere Panoramic Representations

In the equirectangular projection, X and Y values of the X and Y indices are mapped into the latitude and longitude on a sphere. See Figure 1.




Figure 1. Equirectangular projection.

In the cube map, a perspective projection from the center of the cube onto its six faces is performed. See Figure 2.




Figure 2. Cube Map projection.

In both Figure 1 and Figure 2, the green and red lines represent the stereo projections of a vertical line in the scene. In the equirectangular projection, if you imagine taking the left- and right-eye view and mapping them to the same sphere where the red and green lines represent the left and right projection of a vertical line, one can see that the disparity between the left and right images becomes zero. Even more disturbing, if you look directly at the pole, the disparity between eyes will reverse at the pole causing a pseudoscopic image in part of the scene. In the cube map projection, Figure 2, one can see that there is a stereo discontinuity between the top and side walls. This cannot be fixed by rotating the top 90°. This will just move the discontinuity to the back wall. It is apparent from these examples that one can't just paint the stereo pairs from these projections onto the surface of either a cube or sphere and have the stereo work everywhere.

The Acquisition of Panoramas

Three authors of this paper have been shooting panoramas for over 10 years. We acquire panoramas using two micro 4/3 cameras in a parallel stereo configuration, Figure 3, mounted on a two-axis rotary "GigaPan Epic Pro Robotic Pano-head" platform [1]. We typically take 60 or more stereo pairs, Figure 3, and stitch them into two equirectangular projections, Figure 4. We often acquire images at over 60 pixels per degree (20/20 vision) and process them into over 20,000 x 10,000 pixel panoramas. A more detailed description of image acquisition is available in the paper [2].



Figure 3. Camera rig and captured images.




Figure 4. Equirectangular projection of left image of the stereo pairs

The Display of Panoramas in Virtual Reality CAVE(s)

The method of display used at the Electronic Visualization Laboratory (EVL) and Qualcomm Institute (QI) at Calit2/UCSD is to make 2 textures from the left and right equirectangular projections. The left and right eye have infinity coincident. We show the left texture on a large sphere for the left eye. We show the right texture on the same large sphere for the right eye. Then we do the normal VR projections for the display of the sphere. By “large sphere” we mean much larger than the VR display device.




Figure 5. Left, features at infinity and right, features closer than infinity.

For objects at infinity, i.e., zero disparity between the left- and right-eye images, disparity is generated by the VR projection process. See Figure 5. For features closer to the viewer, the disparity in the image is added to the disparity at infinity. As in real life, the standard VR protection of faraway objects will move with the viewpoint. The moon always follows you. This aspect of the standard VR projection extends to the difference in eye positions. The position of the left eye’s sphere (in green) is one half the interocular distance to the left. And the position of the sphere for the right eye’s view (in red) is one half the interocular distance to the right. This produces the disparity between the two images equal to the interocular distance. For objects closer than infinity, the disparity of features in the image is added to the disparity at infinity.

There are a number of affordances that come with this display method. One, in the real world, if you tilt your head to the right and left, the horizontal disparity between features of the world rotates with your head. This is also true in this display method because the sphere will move with the eyes. Two, faraway objects will move with the viewer’s position. Of course, objects close to the viewer also move with the viewer’s position which is not as in the real world. But if the images don’t move with the viewer’s motion, motion parallax would clearly fix the objects to the surface of the screen. Three, the equirectangular projection which maps X and Y positions in the image to latitude and longitude on the sphere is supported in virtually all 3D display systems.

Method Used to Generate Stereo Examples

At the time of this writing, we didn’t have access to a CAVE with a floor or ceiling. So, the examples were generated in the CAVE2 by rotating the polar areas up onto the vertical screens. We photographed the left- and right-eye images by placing the correct polarizing filter for each eye. For this printed paper, I am utilizing the anaglyph method of showing stereo. Figure 6 shows the CAVE2 and the photographic process.



Figure 6. CAVE2 and photographic process.

Stereo Examples of the Display Method

Figure 7 shows that the stereo disparity of faraway objects is equal to the interocular distance. This image is not in stereo but has the left and right eyes’ images superimposed on the screen. Figure 8 shows the same image without me in anaglyph stereo.



Figure 7. The left and right eyes’ images superimposed showing the disparity of faraway objects is equal to the interocular distance.




Figure 8. The left and right images in anaglyph of Figure 7.

Figure 9 shows the disparity is correct for features behind and in front of the display screen. In the CAVE2, the screen is about 11 feet from the camera. This is, of course, stereo near the equator of the projection where one would not expect stereo problems.




Figure 9. An anaglyph showing the disparity of objects far and near.



Figure 11. An anaglyph showing stereo is correct for the zenith if rotated.

Stereo performance at the poles

Figures 10 and 11 show the image of the ceiling at the zenith and the stereo is at least approximately correct. The ceiling is estimated to be about 30 feet high. This means that the disparity encoded in the image is small and the dominant stereo effect comes from the standard VR projection. In general, faraway objects will be rendered with excellent stereo for normal VR viewing conditions whether or not they are on the poles.




Figure 10. An anaglyph showing stereo is correct for the zenith.

There is, however, a problem with the stereo for objects that are close to the camera near the poles, as seen in Figure 12.




Figure 12. An anaglyph showing the disparity of a ceiling about 5 feet away from the camera. The stereo clearly is not correct.

In Figure 12, the pole is close to the bottom of the frame. The horizontal disparity at the bottom of the frame is near zero. By the middle of the frame, there is horizontal disparity but also a great deal of vertical disparity. The vertical disparity of a stereo image should always be zero. By the top of the frame, the stereo is approximately correct. A more pathological case is demonstrated in Figure 13.




Figure 13. An anaglyph showing the disparity of a ceiling about 2 feet away from the camera. The stereo clearly is not correct.

In Figure 13, the pole is just below the center of the frame and the disparity there is zero. For an object that close, there should be a large disparity. In addition, one can see that the disparity is organized as a rotation about the pole. It should primarily be a right to left displacement. This pattern is very reminiscent of the behavior of disparity at the bottom of an equirectangular projection as demonstrated in Figure 1. This behavior might be caused by the equirectangular projection, but it also could be caused by the arrangement of cameras combined with the stitching methods. There is similar behavior on the floor, as shown in Figure 14.




Figure 14. An anaglyph showing the behavior of disparity on the floor. The stereo clearly is not correct.

In Figure 14, the rotary behavior of disparities is clearly shown around the pole at the bottom of the frame. By the top of the frame the stereo is approximately correct with very little vertical disparity.

Display-Omni-Stereo Display

We use the term “display-omni-stereo” to distinguish it from the “omni-stereo” term commonly used in the description of a method of construction of 3D panoramas [3]. In the display method normally used in the CAVE, both for stereo panoramas and normal 3D models, the stereo experience is only correct for the tracked person. A non-tracked person looking in the opposite direction of

the tracked person will see pseudo-stereo. A person looking at right angles from the tracked person will see disparity equal to zero.

For this reason, when we have groups of people in the CAVE, we use another display technique, display-omni-stereo. In this case we, optionally, continue tracking the viewer’s position. But we calculate the head pose as if the viewer was looking directly at the center of each panel making up the CAVE. Since modern CAVEs are made up of many relatively small panels, the display still looks good. And, more importantly, the stereo is more correct for all viewers. There is some discontinuity between panels for objects that are close to the tracked viewer. These become noticeable when objects are within an arm’s reach.

In the classic CAVE, the display screens are large, encompassing a 90° angle of view. So, the display-omni-stereo technique has to operate differently. We have developed a more correct raytracing technique where we assume the viewer is looking directly at each pixel. This avoids discontinuities at the display frame’s edges, but it doesn’t operate in real time. A technique that does operate in real time for the classic CAVE is to divide the display into a series of vertical rectangles of full height but sufficiently small width and do the standard VR projection. These can operate in real time. The paper “OmniStereo for Panoramic Virtual Environment Display Systems” describes this approach in detail [4]. When we apply the display-omni-stereo technique to stereo panoramas, the stereo is correct in the equatorial region for all viewers near the center of a CAVE looking in any direction and continues to be approximately correct in the polar regions for distant objects. For close objects on the poles, the distortions are similar to the pose tracking case, except that the pole is not always directly below the viewer.

Conclusions

The motivation for this work arises from having shown these panoramas to thousands of people in CAVEs with floors and the experience is well received. Even stereo experts don’t complain about the stereo on the floor. I’ve been aware that the stereo on the floor is not correct, so I wondered why does it seem to work?

In our display method, for the position tracked case, the biggest problem is always covered by your feet! In this display method, the whole sphere moves with your tracked motion. The moon always follows you. In addition, because of the tracking, the participant is always looking away from the pole in the direction where the stereo is at least approximately correct. The tracked person, without being in a contorted position, can’t look in the direction where the stereo is reversed. In addition, the stereo becomes correct moving away from the pole and connects without discontinuity to the walls of the CAVE. Although for close objects like a floor, the stereo is compromised, but it still contributes strongly to the feeling of immersion and the feeling of connection to the displayed panorama.

The equirectangular and cube map projections are equivalent except for sampling efficiency. They are both point projections from a center through the sphere or a cube. Yet, they predict stereo problems in different positions. It is plausible that the stereo problems at the poles come from the camera arrangement or the stitching technique.

The partial success of the omni-stereo display approach for faraway objects suggests that it may be possible to successfully display full 360° x 180° panoramas without tracking the pose of the viewer.

With the exceptions of the display-omni-stereo discussions, the observations and conclusions in this paper also apply to head mounted displays.

Future work

The arguments about the quality of stereo in this paper have been primarily visual arguments. In future work, we would like to investigate algebraic arguments. We plan to work with modified stitching methods and perhaps different camera arrangements to attempt to solve the problems of stereo at the poles. In addition, we plan to investigate extended display-omni-stereo methods to obviate the need for tracking.

Acknowledgements

This publication is based on work supported by the National Science Foundation (NSF) award #CNS-1456638 for SENSEI. We also would like to express our appreciation for the significant contributions of Ahmad Atra, Serena Schultz, Hayley Christianson and Bartosz Kupiec, as part of the NSF Research Experiences for Undergraduates (REU) program. *CAVE* and *CAVE2* are registered trademarks of the University of Illinois Board of Trustees.

References

- [1] <http://gigapan.com> – Website of EPIC robotic panoheads, used to create gigapixel panoramas
- [2] R.A. Ainsworth, D.J. Sandin, J.P. Schulze, A. Prudhomme, T.A. DeFanti, M. Srinivasan, "Acquisition of stereo panoramas for display in VR environments," Proc. of SPIE-IS&T Electronic Imaging, Vol. 7864, January 2011.
- [3] V. Couture, M.S. Langer, S. Roy, "Analysis of disparity distortions in omnistereoscopic displays," ACM Transactions on Applied Perception, vol. 7, issue 4, Article 25, 13 pages, July 2010.
- [4] A. Simon, R. Smith, R. Pawlicki, "OmniStereo for Panoramic Virtual Environment Display Systems," IEEE Virtual Reality Conference 2004, Chicago, IL, 2001

Authors Biography

Daniel Sandin is Director Emeritus of the Electronic Visualization Laboratory (EVL) and a Professor Emeritus in the School of Art and Design at the University of Illinois at Chicago (UIC). Currently, Sandin is a visiting Senior Research Scientist at EVL. Sandin is particularly interested in the creation of VR artworks that involve the natural world and as built environment, camera image materials, rich human interaction and mathematical systems.

Haoyu Wang is a PhD candidate in the Department of Electrical and Computer Engineering at University of Illinois at Chicago. He works as Research Assistant for SENSEI (SENSor Environment Imaging) project in Electronic Visualization Laboratory since 2015. In this project, he mainly focuses on the construction of efficient and automatic stereoscopic panorama stitching system. His current work includes panorama depth correction, stitching quality evaluation and generation of the stereoscopic 360 video.

Ahmad Atra is a Solution Developer at Avanade. Atra studied Computer Science with a Software Engineering concentration and an Economics minor at University of Illinois at Chicago. Ahmad contributed to the SENSEI project as funded Undergraduate Research Assistant through NSF's Research Experience for Undergraduates program. His current work focuses on implementing custom solutions using tools and technology within the Microsoft ecosystem.

Richard A. Ainsworth is President of Ainsworth & Partners Inc. and has extensive interests in photography and computer software design. He has published five books in the computer field. His software products have been marketed by Microsoft, IBM, Tandy Corp, Sinclair Research Ltd., and

Simon & Schuster. In collaboration with Daniel J. Sandin and Thomas A. DeFanti, Ainsworth created and produced the CAVEcam system for photographing spherical images in full stereo.

Maxine Brown is the Director of the Electronic Visualization Laboratory (EVL) and the Software Technologies Research Center at University of Illinois at Chicago. Brown has been active in the ACM SIGGRAPH organization and in SIGGRAPH and ACM/IEEE Supercomputing conferences. In recognition of her services to UIC and the community at large, Brown is a recipient of the 1990 UIC Chancellor's Academic Professional Excellence award; 2001 UIC Merit Award; and 1998 ACM SIGGRAPH Outstanding Service Award.

Thomas A. DeFanti, PhD, is a research scientist at the Qualcomm Institute, University of California, San Diego, and a distinguished professor emeritus of Computer Science at the University of Illinois at Chicago. He is recipient of the 1988 ACM Outstanding Contribution Award and was appointed an ACM Fellow in 1994. He shares recognition along with EVL director Daniel J. Sandin for conceiving the CAVE virtual reality theater in 1991.

JOIN US AT THE NEXT EI!

IS&T International Symposium on

Electronic Imaging

SCIENCE AND TECHNOLOGY

Imaging across applications . . . Where industry and academia meet!



- **SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •**
- **INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •**

www.electronicimaging.org

