

# Omnidirectional Stereoscopic Panoramas

Paul Bourke

Volker Kuchelmeister (UNSW)

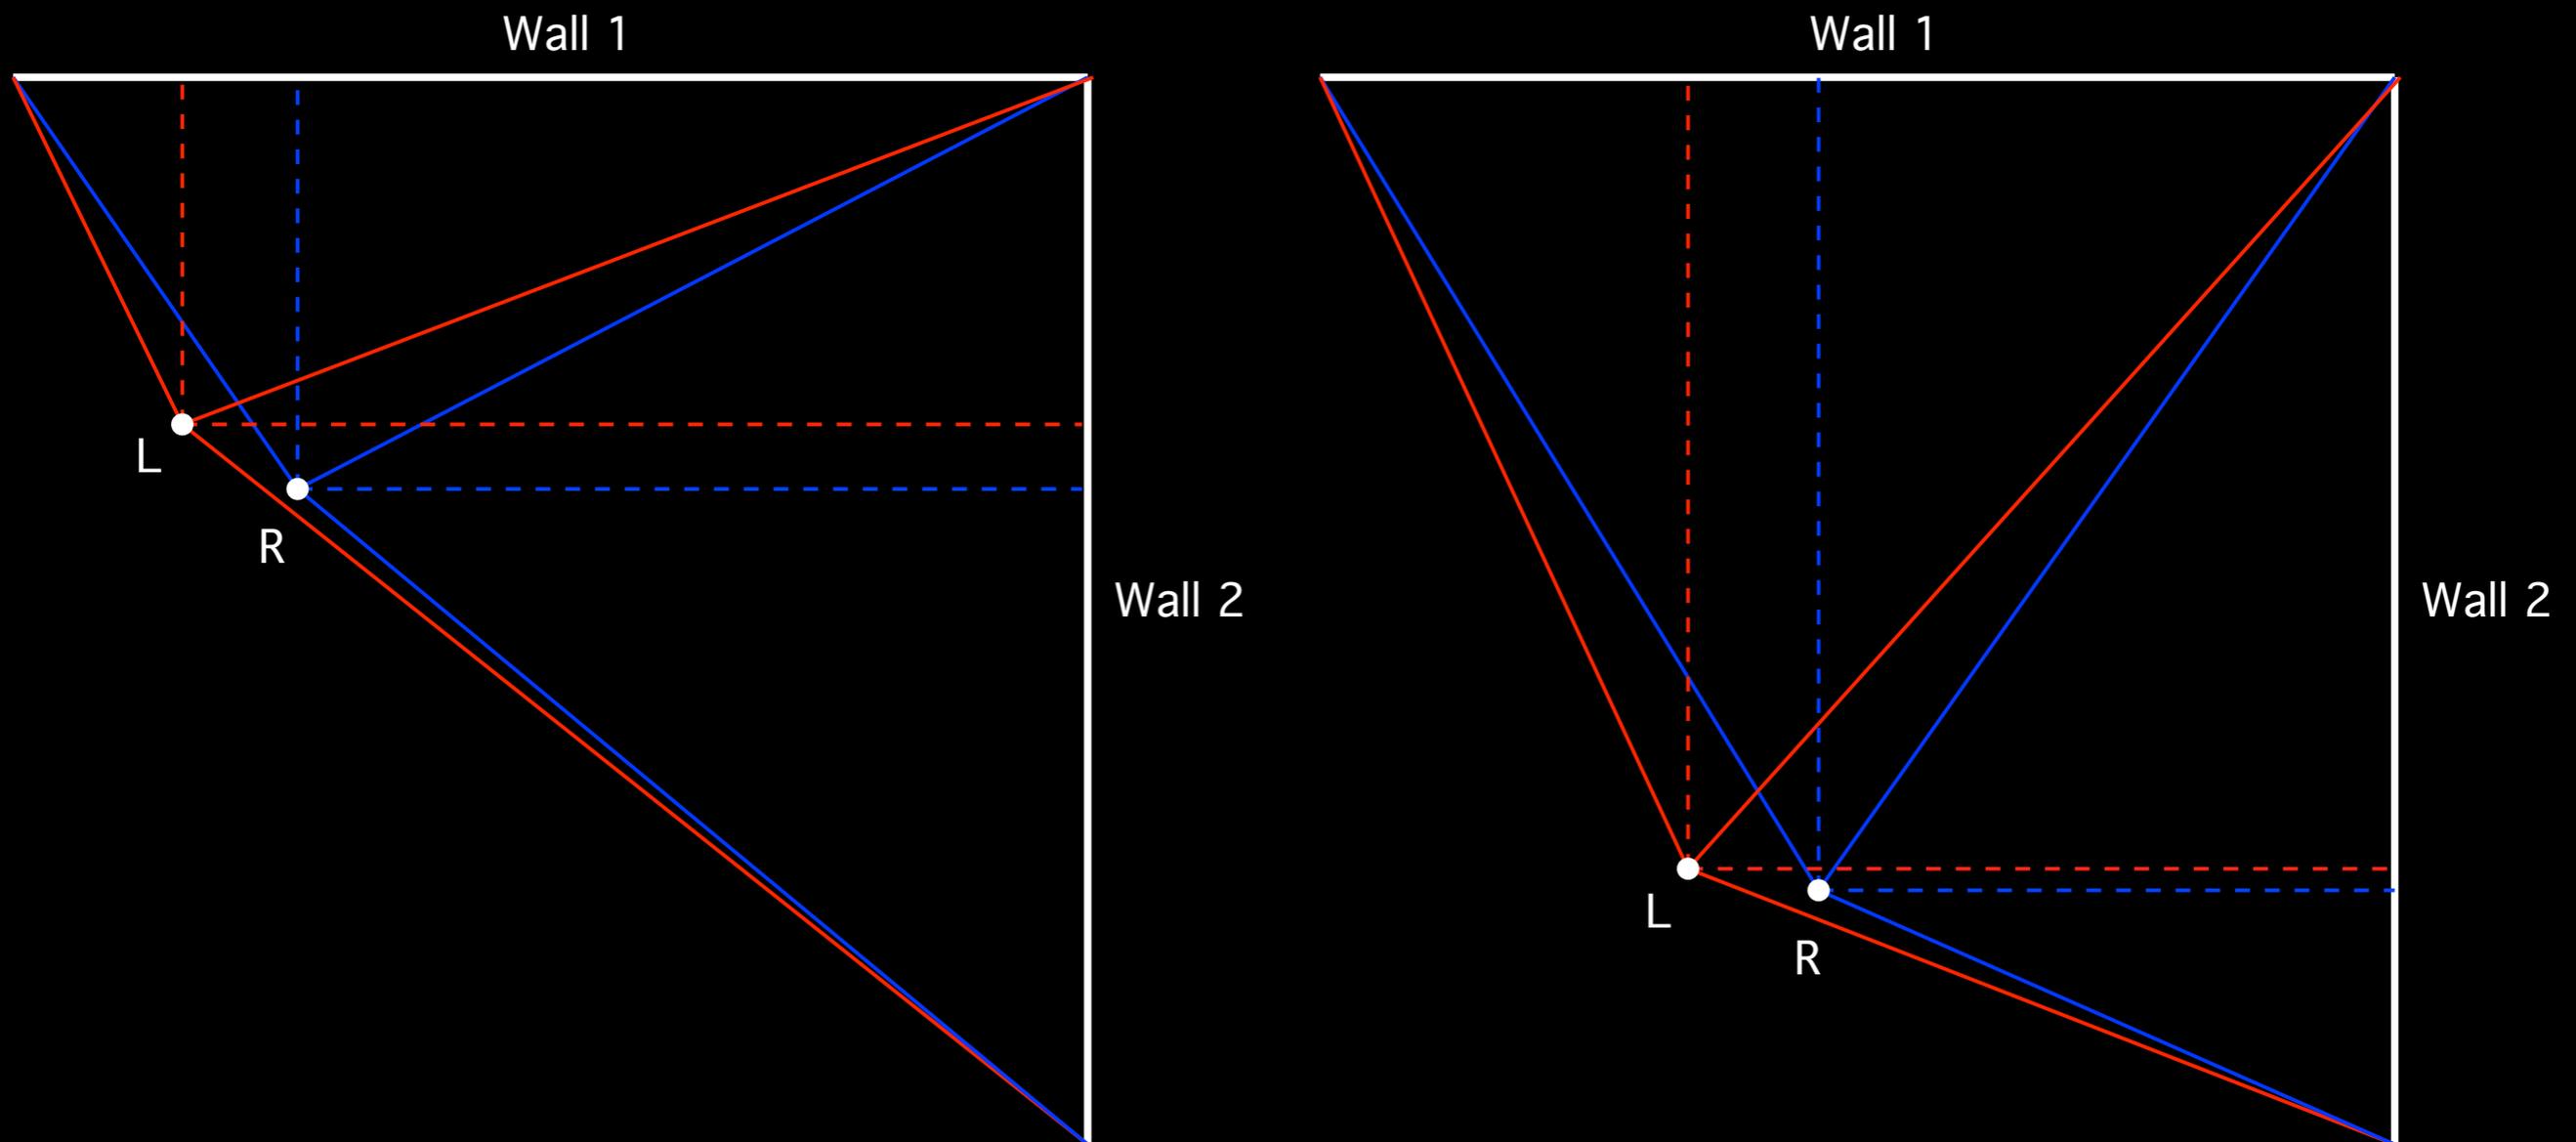
Collaboration with City University (Hong Kong)

# Outline

- Motivation:  
Capture of so called “Omni Directional Stereoscopic Panorama image pairs” (ODSP).
- Past work:  
Collaboration with Sarah Kenderdine & Jeffrey Shaw
  - Existing camera, RoundShot
  - Place Hampi and YER-Türkiye (Place Turkey)
  - Software: “panoalign”
  - Papers
    - Capturing Omni-Directional Stereoscopic Spherical Projections With A Single Camera. Published in the Proceedings 2010 16th International Conference on Virtual Systems and Multimedia. ISBN: 978-1-4244-9025-7, pp 179-183. Presented at VSMM 2010. Seoul, South Korea. Oct 2010.*
    - Omni-Directional Stereoscopic Fisheye Images For Immersive Hemispherical Dome Environments. Proceedings of the Computer Games & Allied Technology 09 (CGAT09), Research Publishing Services, ISBN: 978-981-08-3165-3, pp 136-143. Presented at Computer Games & Allied Technology 09, Singapore, May 12, 2009.*
- Current:  
Collaboration with Volker Kuchelmeister (UNSW) and City University (Hong Kong).  
Applications to heritage and archaeology site recording.  
Develop hardware and software for a digital solution.

# Head tracking for stereo in immersive environments

- We are used to watching stereo3D on a single flat screen while seated.
- If one moves around the 3D content gets squashed and warped in strange ways.
- And, you can't look around objects.
- Reason: the view frustums need to change depending on position and viewing direction.



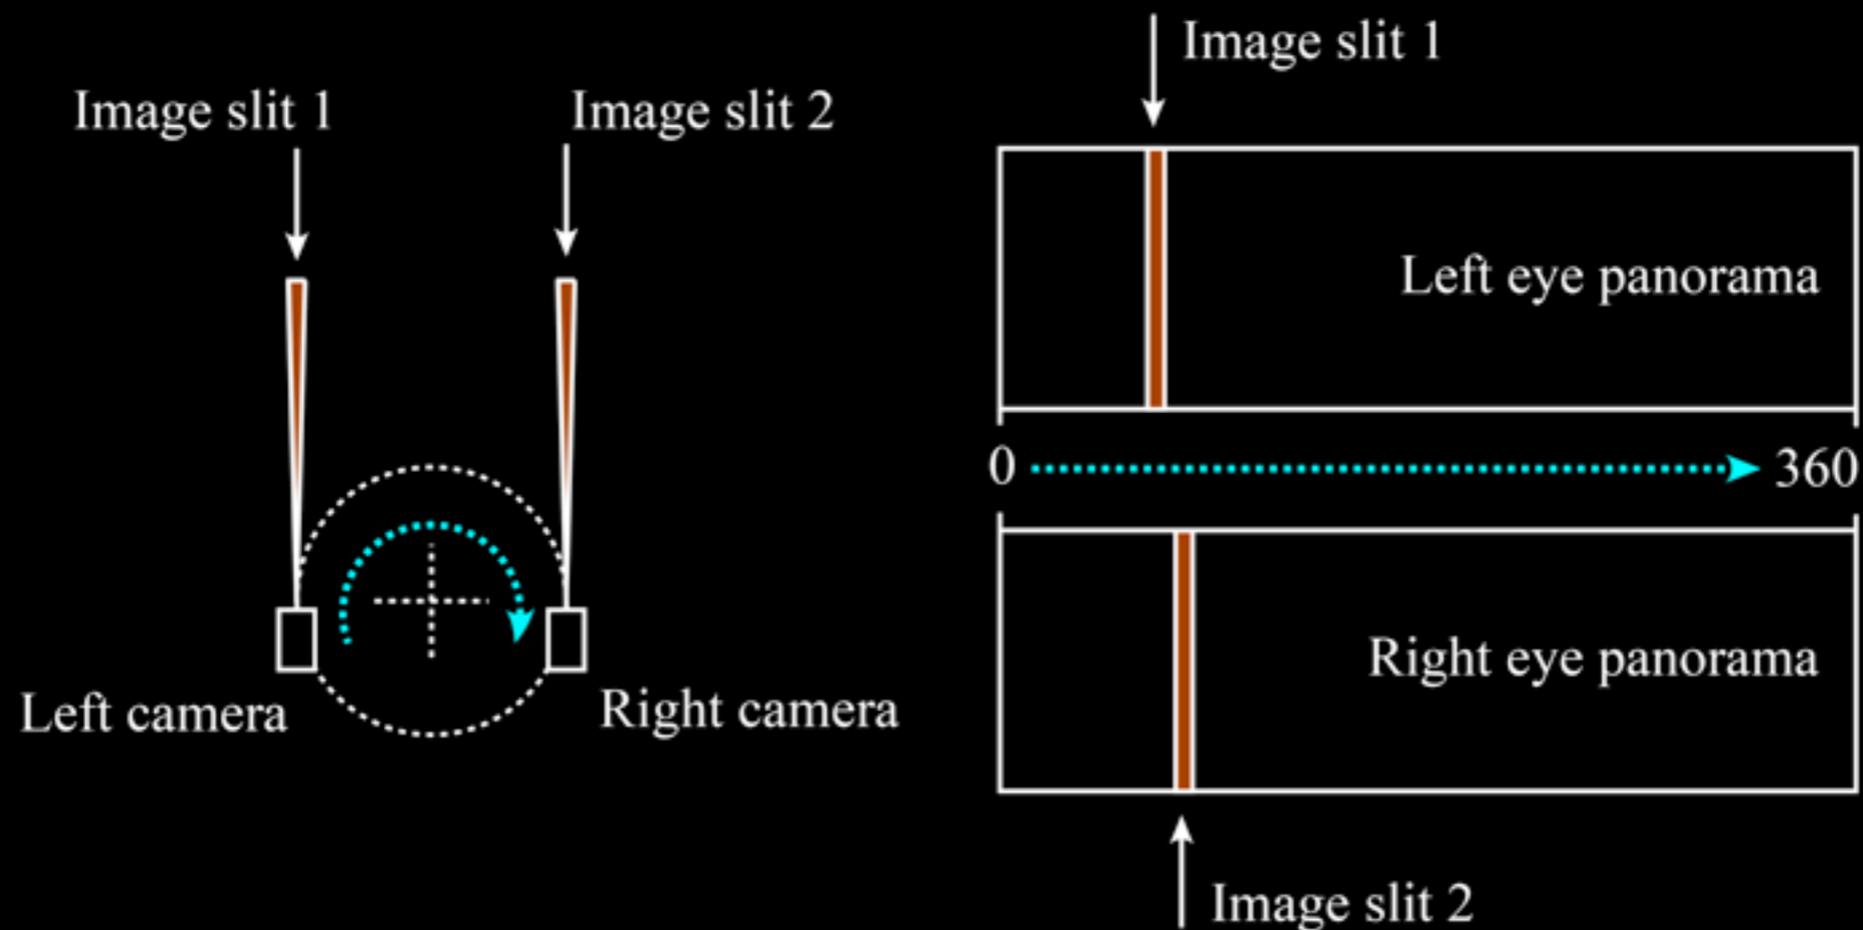
# Cylindrical displays

- AVIE (Advanced Visualisation Interactive Environment): 10m diameter. UNSW - UTS - CityU
- 1/2 cylinder at the Curtin Data Visualisation Facility (CVDF).
- CAVE2 at Monash University
- Obviously suited to multiple participants all likely looking in different directions.
- Can the need for head tracking be relaxed and still provide an acceptable stereoscopic experience?



# Omni-Directional Stereoscopic Panoramas (ODSP)

- Concession: Only require the stereoscopy is strictly correct for the portion directly in front of the viewer.
- The error increases towards the edge of the field of view, but
  - glasses limit the FOV anyway
  - and we don't see stereo in our far field
- The camera model has two (vertical slit) cameras rotating about the midpoint.
- Computer generated ODSP implement the camera model directly.

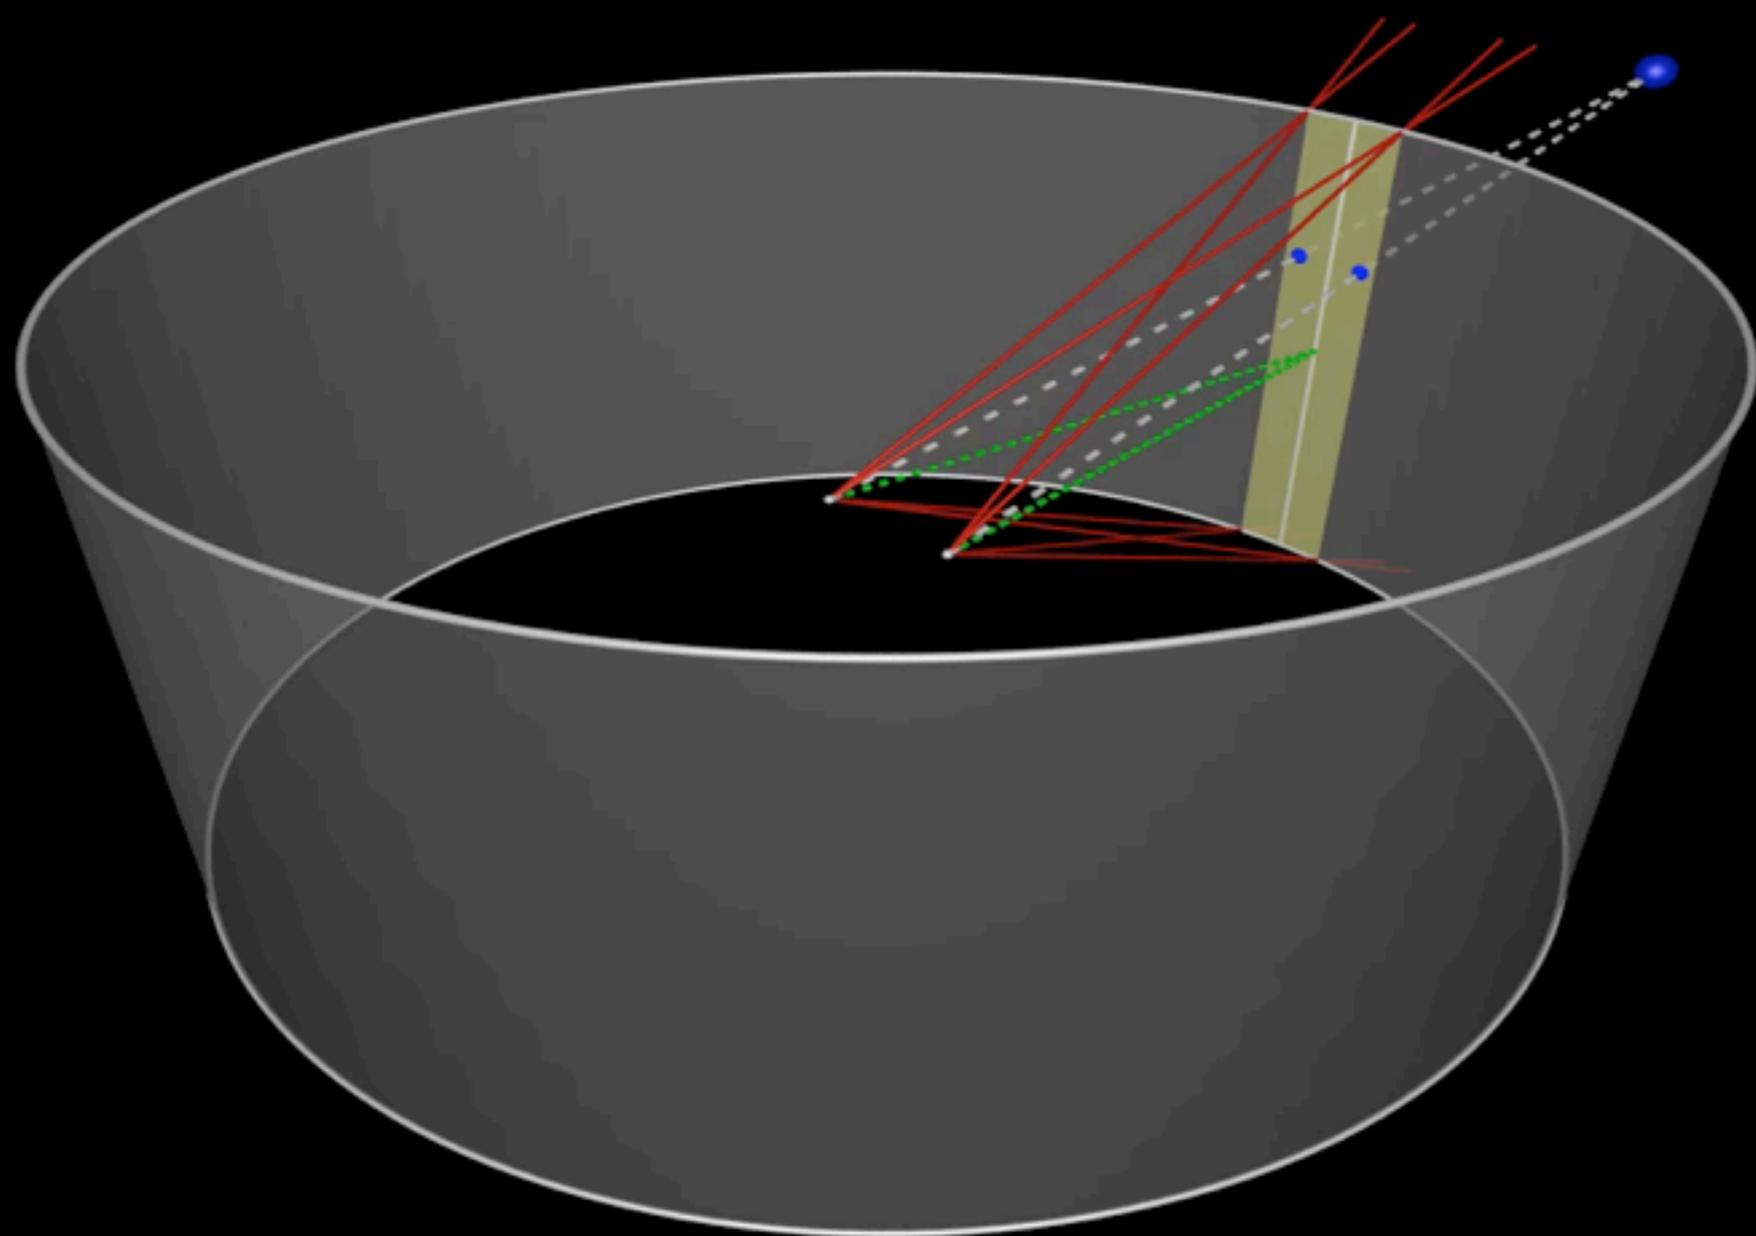




Left eye panorama



Right eye panorama



Pair of rotating slit cameras

# Place Hampi



Left eye



Right eye

# YER-Türkiye



Left eye



Right eye

# Roundshot camera

- Uses two rolls of 70mm film.
- Requires drum scanning and then alignment of the stereo pairs and colour calibration.
- Developed tools to do this in 2004.
- Capable of high resolution.  
~ 100,000 pixels horizontally
- Film stock getting harder to get
- Not sustainable.



# Developing a digital alternative

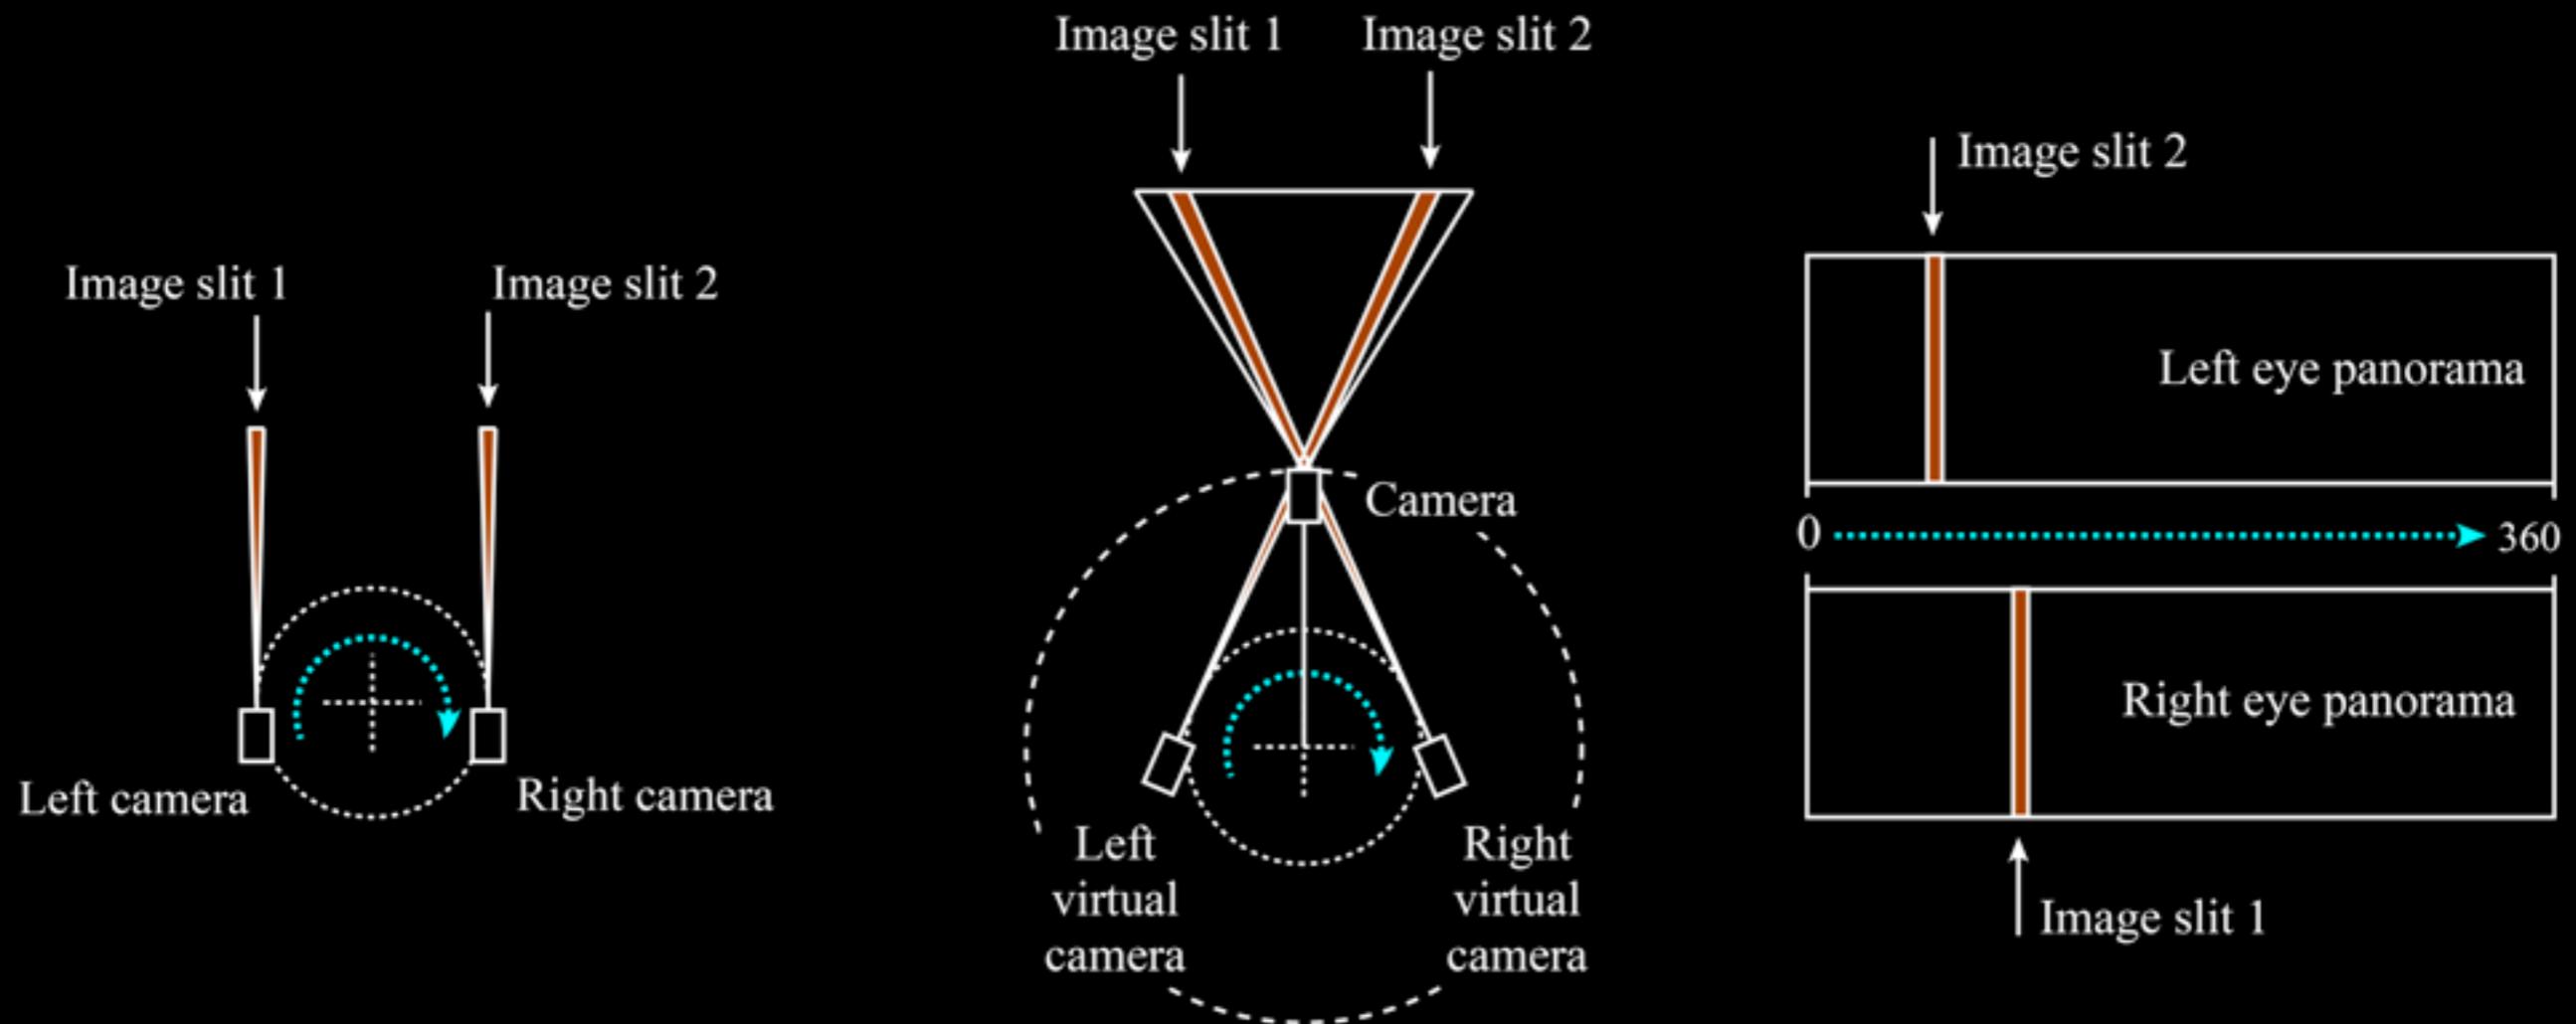
- Challenge/tradeoffs: high enough resolution and fast enough capture.
- The former is required for increasingly high resolution cylinders. eg: CAVE2
- The latter determines how much movement is allowed in the scene. Perhaps surprisingly some movement is allowed since the capture is localised.



- Two recent developments make it possible
  - 4K video cameras, eg: Red
  - fast continuous capture SLR still cameras

# Single camera

- Perhaps surprisingly, can capture the stereo pair with only a single camera.
- Choosing two columns from each image gives the same results as two cameras.
- Very powerful consequence in that the interocular separation can be changed in post production by choosing the separation between the columns selected.



# Red on the Roundshot motor



# Canon 5D MkIII on the Roundshot motor



# Syrp motorised rig

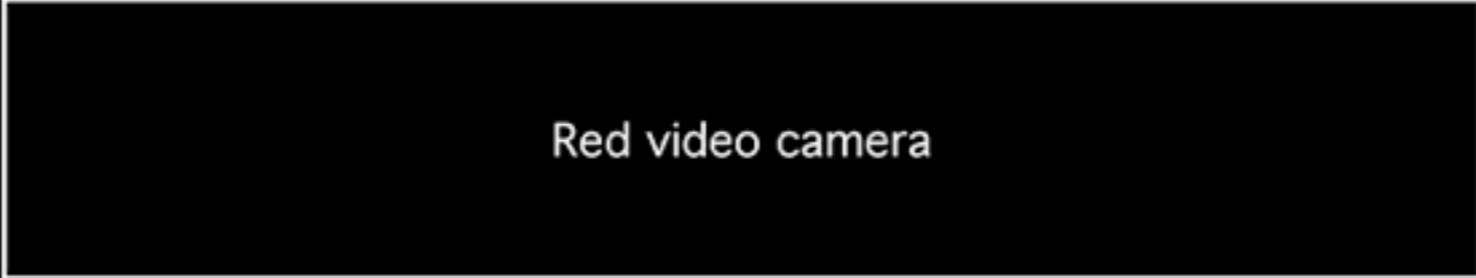


Volker Kuchelmeister (UNSW)

# Resolution / speed tradeoff

1900 pixels  ~20 second scan

The diagram shows a small white rectangular box labeled "HD video camera". To its left, a vertical double-headed arrow indicates a height of "1900 pixels". To its right, the text "~20 second scan" is written.

4000 pixels  ~40 second scan

The diagram shows a medium-sized white rectangular box labeled "Red video camera". To its left, a vertical double-headed arrow indicates a height of "4000 pixels". To its right, the text "~40 second scan" is written.

5700 pixels  ~60 second scan

The diagram shows a large white rectangular box labeled "SLR". To its left, a vertical double-headed arrow indicates a height of "5700 pixels". To its right, the text "~60 second scan" is written.

Note: Image height is determined by the lens being used.

# Pipeline

Images from camera

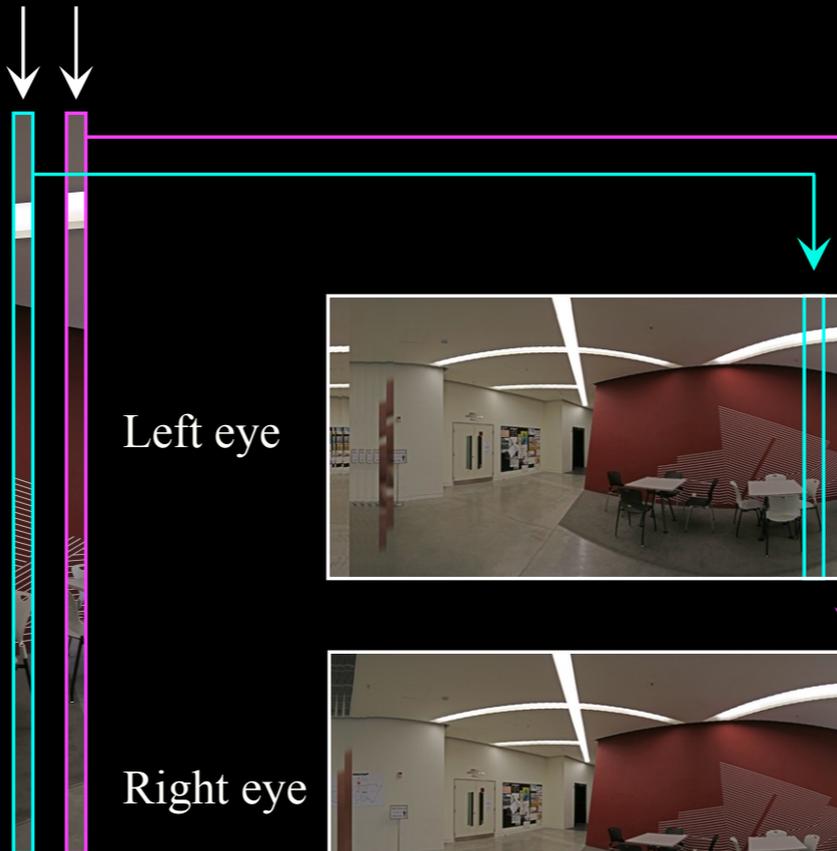


Images rectified and rotated



Canon 5D: 5760 high  
Red Scarlet: 4096 high

2 vertical slices extracted from each image.



Vertical slices each combined into a left and right panoramic image pair

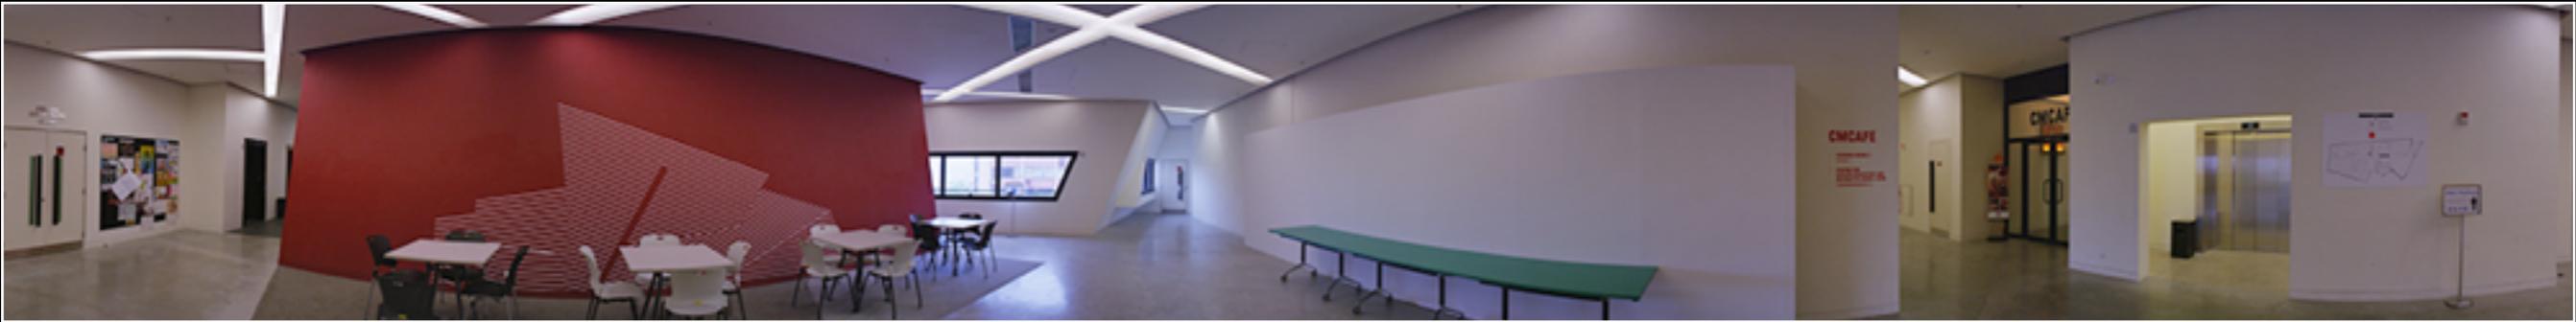
Slices defined by their width and offset from the center of the image. In turn those parameters are calculated from the geometric and optical properties of the camera rig along with the desired interocular distance in order to result in a seamless panorama pair.

Left eye

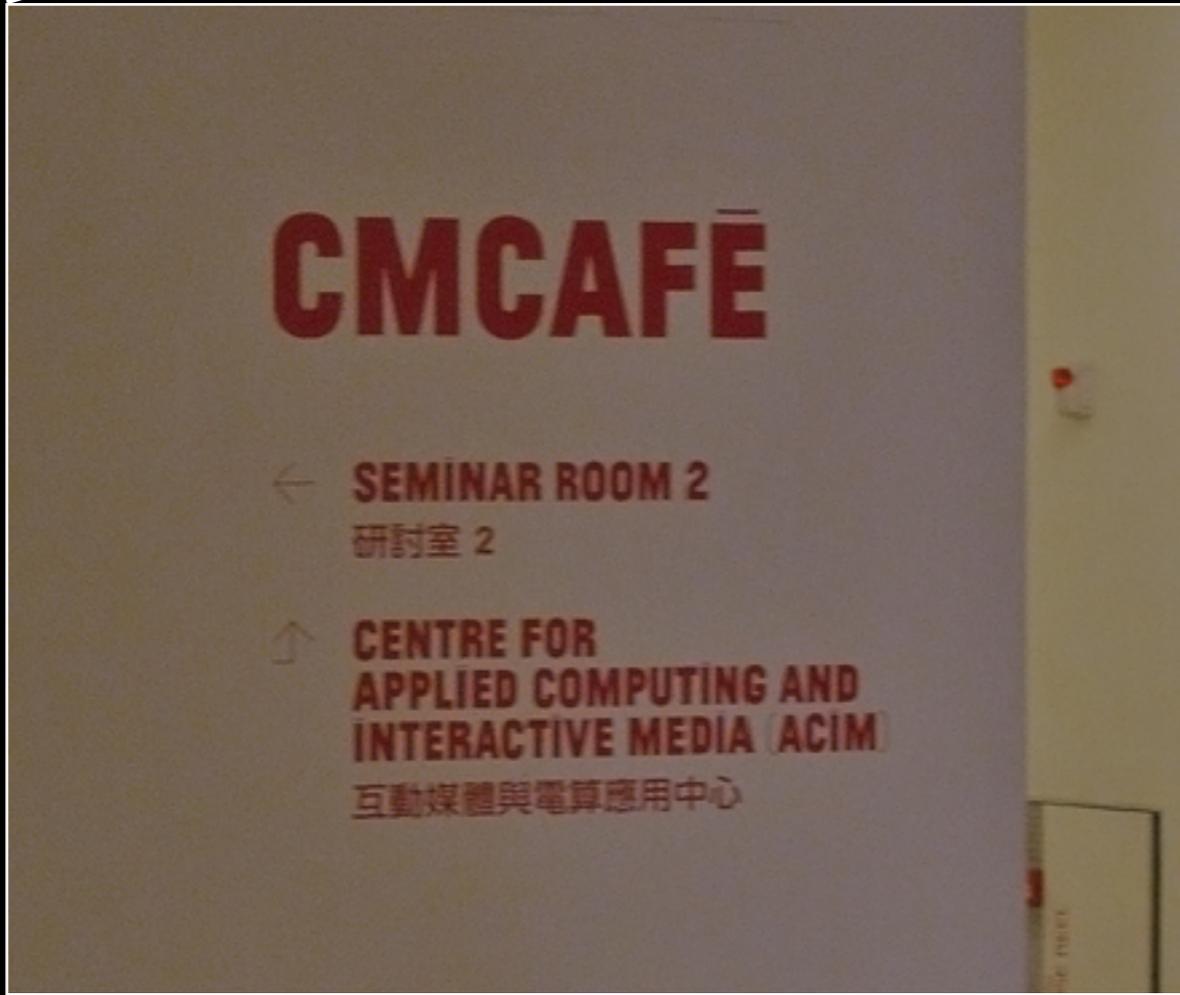


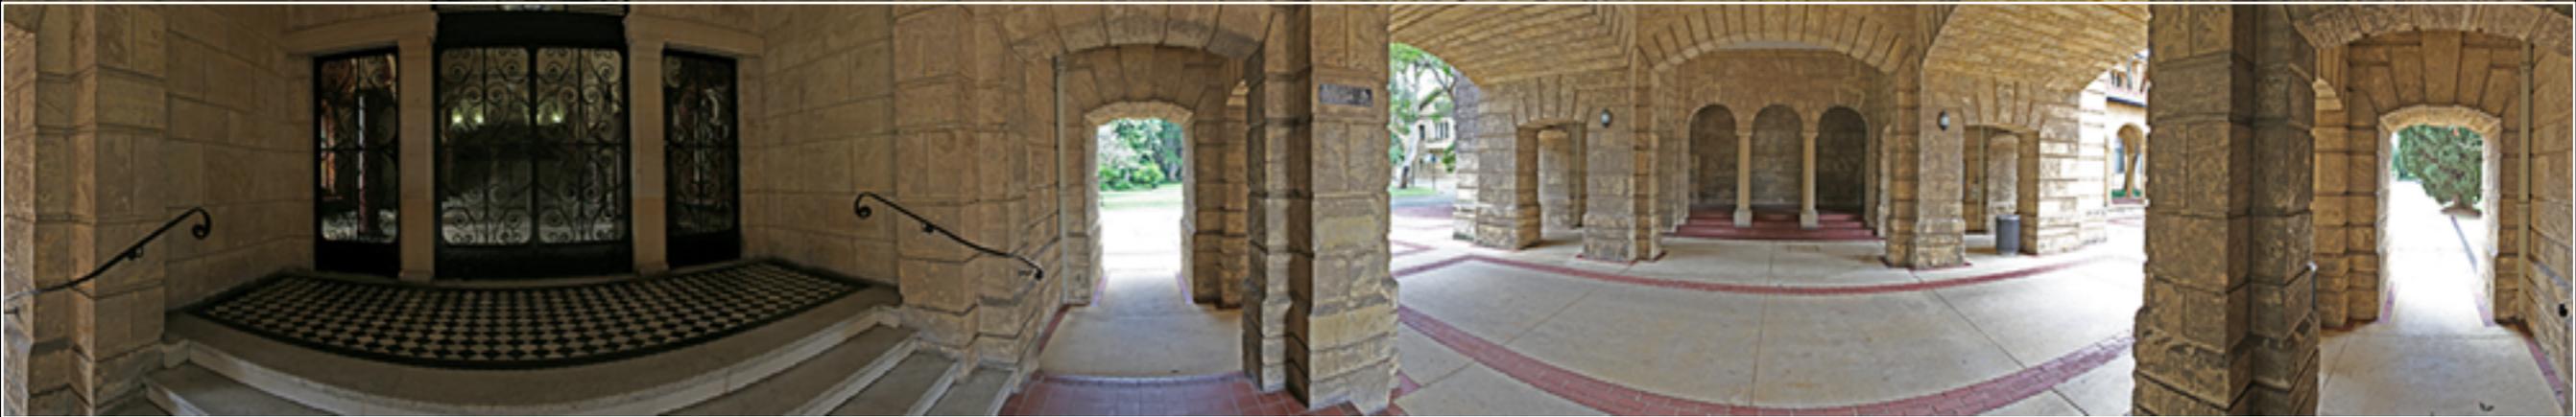
Right eye





40,000 pixels wide





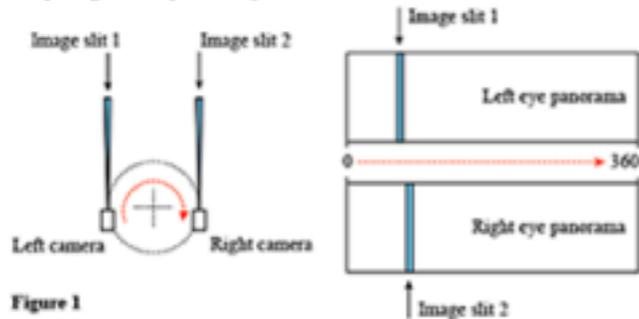


# Poster to Siggraph Asia

Highly immersive displays are those that are both stereoscopic 3D enabled as well as filling the viewers field of view. Traditionally these are built from multiple planar walls that surround the viewer, the most well known being the CAVE [1]. Due largely to the discrete nature of the walls, head tracking is required in order to create the correct stereo pairs on each wall, and as such, these displays are inherently single person experiences.

Another category of immersive displays are based upon a cylindrical surface, such as the AVTE [2] (Advanced Visualisation and Interaction Environment) and the CAVE2 [3]. These generally encompass a larger space and are capable of hosting a number of participants all potentially looking in different directions. This raises the question of whether the need for head tracking can therefore be relaxed. The solution is generally referred to as omni-directional stereo, that is, an acceptable stereoscopic 3D experience can be enjoyed by multiple participants without the need for head tracking.

The creation of synthetic omni-directional stereoscopic 3D panoramas for cylindrical displays is well understood [4] but the capture of digital panorama photographs is challenging. Successful film based cameras [5] have been built which are based upon the optics illustrated in figure 1 with the basic principle being a pair of slit cameras rotating about a central location exposing two strips of film, one for each eye. This is the same model typically employed for computer generated panorama pairs.



This poster presents exploration into a viable digital alternative for capturing a stereoscopic panorama pair using off-the-shelf components. There is additionally the requirement to capture these panoramas at sufficiently high resolution to match the growing resolution of modern cylindrical displays.

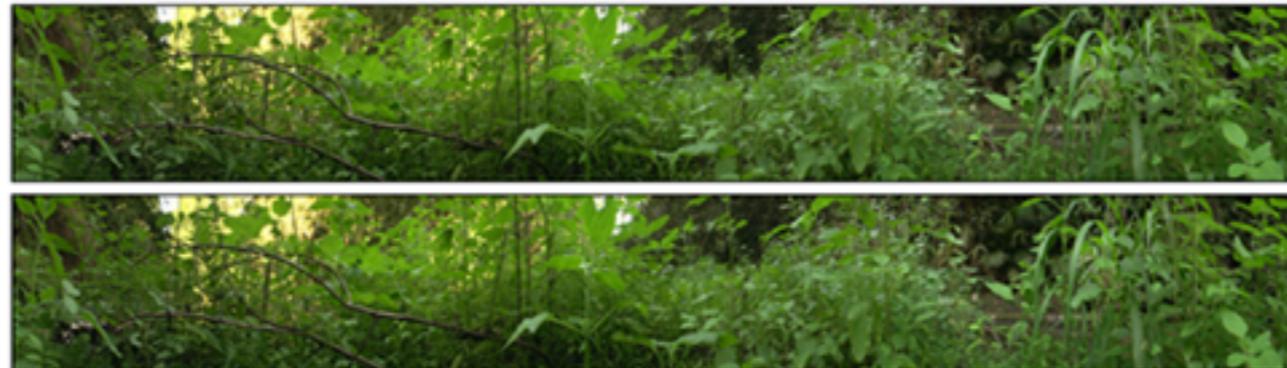
It has been previously proposed that, perhaps surprisingly, only a single camera is required [6,7,8,9,10]. This is achieved by taking a large number of photographs with only a single camera rotating at an offset position from a central axis, and pointing perpendicular to the axis of rotation. The final panorama image pairs are assembled in post-production by mosaicing precisely selected columns from the source images. Figure 2 illustrates that if two image slits are extracted from the frames captured by such a perpendicularly arranged camera, then the strips form the same images as would be obtained from a pair of slit cameras. This not only reduces the cost but removes many of the complications of stereoscopic video capture employing two cameras, such as camera calibration and colour matching.

A powerful consequence of this mode of capture is that the interocular separation can be adjusted in post production simply by choosing the distance between the two image slits extracted. Choosing the "best" depth for a stereoscopic image is as much a creative decision as a technical challenge for the photographer.

## Capture of Omni-Directional Stereoscopic Panoramic Images

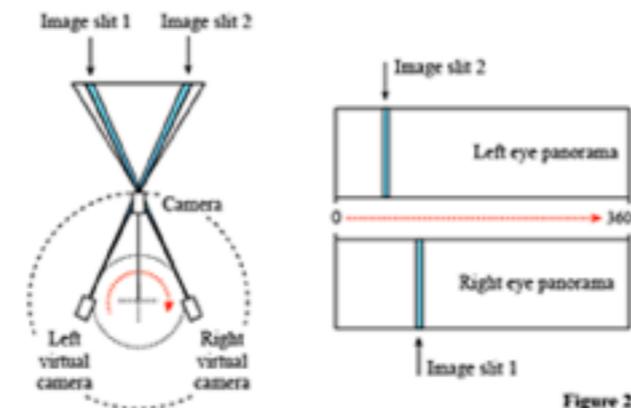
Paul Bourke<sup>1</sup>, Völker Kuchelmeister<sup>2</sup>

<sup>1</sup>IVEC@UWA, The University of Western Australia. (paul.bourke@uwa.edu.au)  
<sup>2</sup>iCinema Centre for Interactive Cinema Research, University of New South Wales

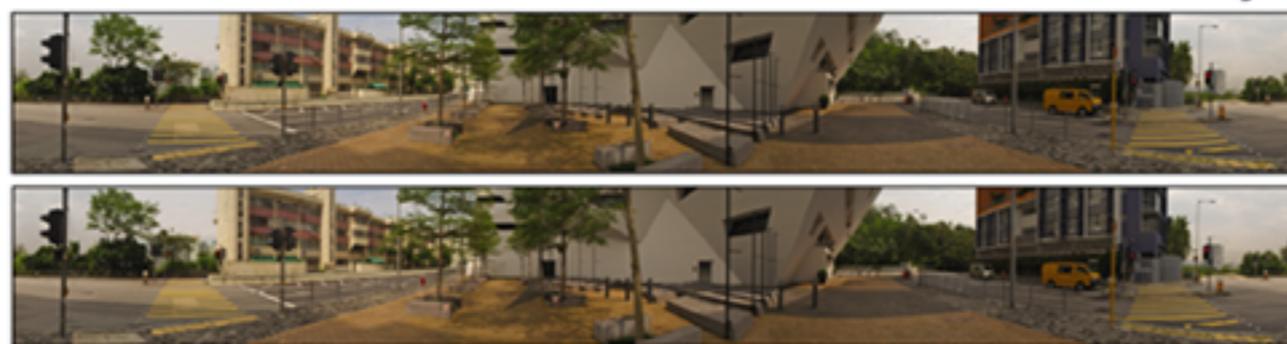


Such control in post production over the disparity and perceived depth is normally not possible and has the important consequence of being able to adjust content for different screen dimensions, optimise viewing comfort, as well as providing significant creative control.

Implementing the configuration shown in figure 2 requires a motorised unit to automatically rotate the offset camera. If long capture times are possible then the head may rotate discretely but typically it is rotated continuously.



There are two options for the camera, it may be either a video camera or a SLR still camera. The relative merits of the former can result in faster acquisitions times at the expense of image resolution. Two examples of prototype cameras built for this work are shown in figure 5. Cameras being tested include, but are not limited to, the Red Scarlet, Canon 5D MkIII and Sony FS700.



Two technological developments that allow sufficient resolution capture are HD or 4K video cameras, providing 2K or 4K pixel high panoramas, and SLR cameras with fast continuous capture modes (eg: 6 fps of the Canon 5D Mk III) resulting in 5.5K pixel high panoramas. Note that in both cases the camera is typically mounted in portrait mode and the horizontal resolution is dictated by the field of view of the lens.

The current and future work will explore and document the relative trade-offs between the parameters involved, how they can effect the quality of the panorama pair, and how they can be used to achieve particular artistic goals. The key parameters are the radial offset of the camera, the separation of the extracted slits, the capture time and the intended resolution/fidelity of the final panorama. This future work will be in the form of guiding principles presented both quantitatively and qualitatively to practically assist in the effective capture of these image pairs for high resolution stereoscopic 3D capable cylindrical display environments.

### Acknowledgements

The work was supported by IVEC through the use of advanced computing resources located at the University of Western Australia. It was conducted in collaboration with the School of Creative Media, City University, Hong Kong.

### References

1. CRUZ-NEIRA, C., SANDIN, D., DEFANTI, T., KEYON, R., HART, J. (1992). *The CAVE: Audio visual experience automatic virtual environment*. Communications of the ACM, Volume 35 Issue 6, June 1992. pp 64-72. ACM New York, NY, USA
2. MCGINITY, M., SHAW, J., KUCHELMEISTER, V., HARDJONO, A., FAVERO, D.D. (2007). *AVTE: A Versatile Multi-User Stereo 360 degree Interactive VR Theatre*. Proceedings of the 2007 workshop on Emerging displays technologies: images and beyond: the future of displays and interaction. San Diego, California. Article No 2. 2007. ISBN 978-1-59593-669-1. 2007.
3. CAVE2 <http://www.mechdyne.com/cave2.aspx>
4. BOURKE, P. D. (2006). *Synthetic Stereoscopic Panoramic Images*. Lecture Notes in Computer Science (LNCS), Springer, ISBN 978-3-540- 46304-7, Vol 4270, pp 147-155.
5. RoundShot, Roundshot panoramic film cameras by Seitz. <http://cameras-wiki.org/wiki/Seitz>
6. PELEG, S., PRITCH, Y., BEN-EZRA, M. (2000). *Cameras for stereo panoramic imaging*. in Proc. IEEE Computer Vision and Pattern Recognition, Vol 1, pp 208-214.
7. PELEG, S., Ben-Ezra, M. (1999). *Stereo panorama with a single camera*. IEEE Conference on Computer Vision and Pattern Recognition, pp 395-401, Ft. Collins, Colorado.
8. BEN-EZRA, M., PRITCH, Y., PELEG, S. (2001). *OmniStereo: Panoramic Stereo Imaging*. IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol 23, No 3, pp 279-290.
9. GURRIERI, L. E., DUBOIS, E. (2013). *Acquisition of omnidirectional stereoscopic images and videos of dynamic scenes*. SPIE Journal of Electronic Imaging, Volume 22 (3), 2013.
10. ISHIGURO, H., YAMAMOTO, M., TSUJI, S. (1992). *Omi-Directional Stereo*. IEEE Transactions on Pattern Analysis Machine Intelligence 14 (2), pp 257-26.

# Poster to Siggraph Asia (Courtesy Karina)

## Capture of Omni-Directional Stereoscopic Panoramic Images

Paul Bourke, IVEC@UWA, The University of Western Australia. (paul.bourke@uwa.edu.au)

Volker Kuchelmeister, iCinema Centre for Interactive Cinema Research, University of New South Wales

Highly immersive displays are those that are both stereoscopic 3D enabled as well as filling the viewers field of view. Traditionally these are built from multiple planar walls that surround the viewer, the most well known being the CAVE [1]. Due largely to the discrete nature of the walls, head tracking is required in order to create the correct stereo pairs on each wall, and as such, these displays are inherently single person experiences.

Another category of immersive displays are based upon a cylindrical surface, such as the AVIE [2] [Advanced Visualisation and Interaction Environment] and the CAVE2 [3]. These generally encompass a larger space and are capable of hosting a number of participants all potentially looking in different directions. This raises the question of whether the need for head tracking can therefore be relaxed. The solution is generally referred to as omni-directional stereo, that is, an acceptable stereoscopic 3D experience can be enjoyed by multiple participants without the need for head tracking.



Figure 3

Two technological developments that allow sufficient resolution capture are HD or 4K video cameras, providing 2K or 4K pixel high panoramas, and SLR cameras with fast continuous capture modes (eg: 6 fps of the Canon 5D Mk III) resulting in 5.5K pixel high panoramas. Note that in both cases the camera is typically mounted in portrait mode and the horizontal resolution is dictated by the field of view of the lens.

The current and future work will explore and document the relative trade-offs between the parameters involved, how they can effect the quality of the panorama pair, and how they can be used to achieve particular artistic goals. The key parameters are the radial offset of the camera, the separation of the extracted slits, the capture time and the intended resolution/fidelity of the final panorama. This future work will be in the form of guiding principles presented both quantitatively and qualitatively to practically assist in the effective capture of these image pairs for high resolution stereoscopic 3D capable cylindrical display environments.

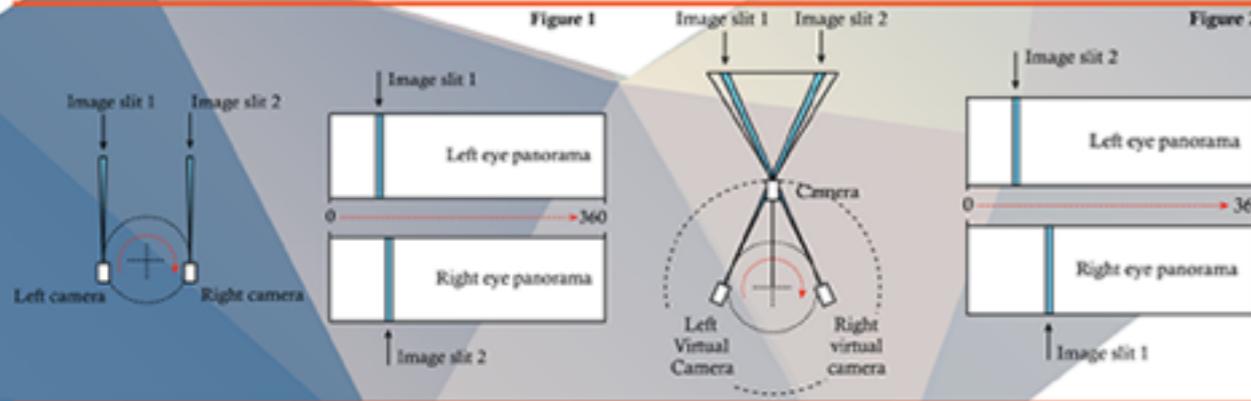


Figure 1

Figure 2



Figure 4

The creation of synthetic omni-directional stereoscopic 3D panoramas for cylindrical displays is well understood [4] but the capture of digital panorama photographs is challenging. Successful film based cameras [5] have been built which are based upon the optics illustrated in figure 1 with the basic principle being a pair of slit cameras rotating about a central location exposing two strips of film, one for each eye. This is the same model typically employed for computer generated panorama pairs.

This poster presents exploration into a viable digital alternative for capturing a stereoscopic panorama pair using off-the-shelf components. There is additionally the requirement to capture these panoramas at sufficiently high resolution to match the growing resolution of modern cylindrical displays.

It has been previously proposed that, perhaps surprisingly, only a single camera is required [4,7,8,9,10]. This is achieved by taking a large number of photographs with only a single camera rotating at an offset position from a central axis, and pointing perpendicular to the axis of rotation. The final panorama image pairs are assembled in post-production by mosaicing precisely selected columns from the source images. Figure 2 illustrates that if two image slits are extracted from the frames captured by such a perpendicularly arranged camera, then the strips form the same images as would be obtained from a pair of slit cameras. This not only reduces the cost but removes many of the complications of stereoscopic video capture employing two cameras, such as camera calibration and colour matching.

A powerful consequence of this mode of capture is that the interocular separation can be adjusted in post production simply by choosing the distance between the two image slits extracted.

Choosing the "best" depth for a stereoscopic image is as much a creative decision as a technical challenge for the photographer.

Such control in post production over the disparity and perceived depth is normally not possible and has the important consequence of being able to adjust content for different screen dimensions, optimise viewing comfort, as well as providing significant creative control.

Implementing the configuration shown in figure 2 requires a motorised unit to automatically rotate the offset camera. If long capture times are possible then the head may rotate discretely but typically it is rotated continuously.

There are two options for the camera, it may be either a video camera or a SLR still camera. The relative merits of the former can result in faster acquisitions times at the expense of image resolution. Two examples of prototype cameras built for this work are shown in figure 5. Cameras being tested include, but are not limited to, the Red Scarlet, Canon 5D MkIII and Sony FS700.

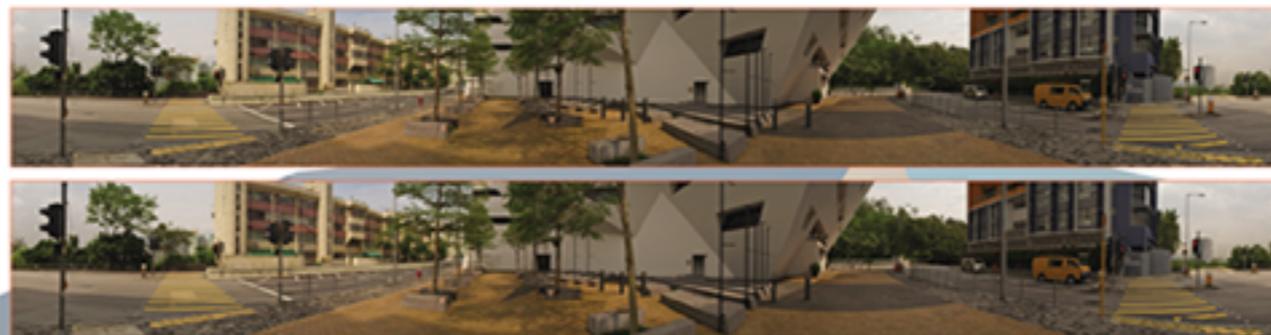


Figure 5

### Acknowledgements

The work was supported by IVEC through the use of advanced computing resources located at the University of Western Australia. It was conducted in collaboration with the School of Creative Media, City University, Hong Kong.

### References

1. CRUZ-NEIRA, C., SANDIN, D., DEFANT, T., KEYON, R., HART, J. (1992). The CAVE: Audio visual experience automatic virtual environment. *Communications of the ACM*, Volume 35 Issue 6, June 1992, pp 64-72. ACM New York, NY, USA.
2. MCINITY, M., SHAW, J., KUCHELMEISTER, V., HARDJONO, A., FAVERO, D.D. (2001). AVIE: A Versatile Multi-User Stereo 360 degree interactive VR Theatre. *Proceedings of the 2001 workshop on Emerging displays technologies: images and beyond: the future of displays and interaction*, San Diego, California, Article No 2, 2001. ISBN 978-1-59593-667-1, 2007.
3. CAVE2 <http://www.mechdne.com/cave2.aspx>
4. BOURKE, P. D. (2004). *Synthetic Stereoscopic Panoramic Images*. Lecture Notes in Computer Science (LNCS), Springer, ISBN 978-3-540-44304-7, Vol 4270, pp 147-155.
5. RoundShot, Roundshot panoramic film cameras by Seitz. <http://camera-wiki.org/wiki/Seitz>
6. PELED, S., PRITCH, Y., BEN-EZRA, M. (2008). Cameras for stereo panoramic imaging. in *Proc. IEEE Computer Vision and Pattern Recognition*, Vol. 1, pp 2084-214.
7. PELED, S., BEN-EZRA, M. (1999). Stereo panorama with a single camera. *IEEE Conference on Computer Vision and Pattern Recognition*, pp 395-401, Ft. Collins, Colorado.
8. BEN-EZRA, M., PRITCH, Y., PELED, S. (2001). Omnidirectional panoramic stereo imaging. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol 23, No 3, pp 279-290.
9. GURBERI, L. E., DUBOIS, E. (2013). Acquisition of omnidirectional stereoscopic images and videos of dynamic scenes. *SPIE Journal of Electronic Imaging*, Volume 22(3), 2013.
10. ISHIGURO, H., YAMAMOTO, M., TSUBOI, S. (1992). Omnidirectional Stereo. *IEEE Transactions on Pattern Analysis Machine Intelligence* 14(12), pp 257-26.

