Advanced Applications in Stereographic, Panoramic, and Fulldome visualisation.

Technical session: Content creation and integration

Paul Bourke
iVEC@UWA

Contents

- Visualisation Laboratory University Western Australia.

- Introduction to 3D to 2D image projections.
  - Standard perspective projection.
  - Cylindrical projection.
  - Spherical (equirectangular) projection.
  - Fisheye projection.

- Stereoscopic displays.
  - Multiple planar display walls.
  - Omni-directional stereoscopic cylindrical display.

- Hemispherical displays.
  - Planetarium dome.
  - iDome.

- Camera based image and video capture.
  - Stitching multiple images.
  - Fisheye lenses.
  - LadyBug-3 camera.
  - Rotating slit stereoscopic cameras.

- Computer generated.

- Project examples.

Re-Actor

CAVE 0.5

iDome

MediaDome

AVIE

Visualisation Laboratory UWA

- Science visualisation: Use of advanced computer graphics to provide insight into research data.

- Primarily employs the sense of sight, so we should leverage the human visual system.
  Three characteristics of our visual system not engaged by standard computer displays are
  - Depth perception for exposing geometrically complicated relationships.
  - Peripheral vision when exploring inside datasets.
  - High resolution displays for large datasets, reduces the context/detail conflict.

- Outcomes include:
  - Revealing new features in the data, scientific discovery.
  - Exposing problems or errors in the data.
  - Presenting research to peers through papers and conferences.
  - Conveying science to a non-expert audience, public outreach and creating educational material.

Stereopsis

Peripheral vision

Resolution/fidelity

Introduction to 3D to 2D image projections.

- Image projection: the process by which a 3D scene (real or synthetic) is mapped onto an image plane (2D).
- Traditional drafting projection categories are shown on the right.
- Computer graphics is mostly concerned with orthographic and perspective projections.
- There are an infinite number of possible projections but only a handful have practical advantages for representing scenes for digital displays.
- As will be shown there are other projections required to support the immersive displays discussed here.

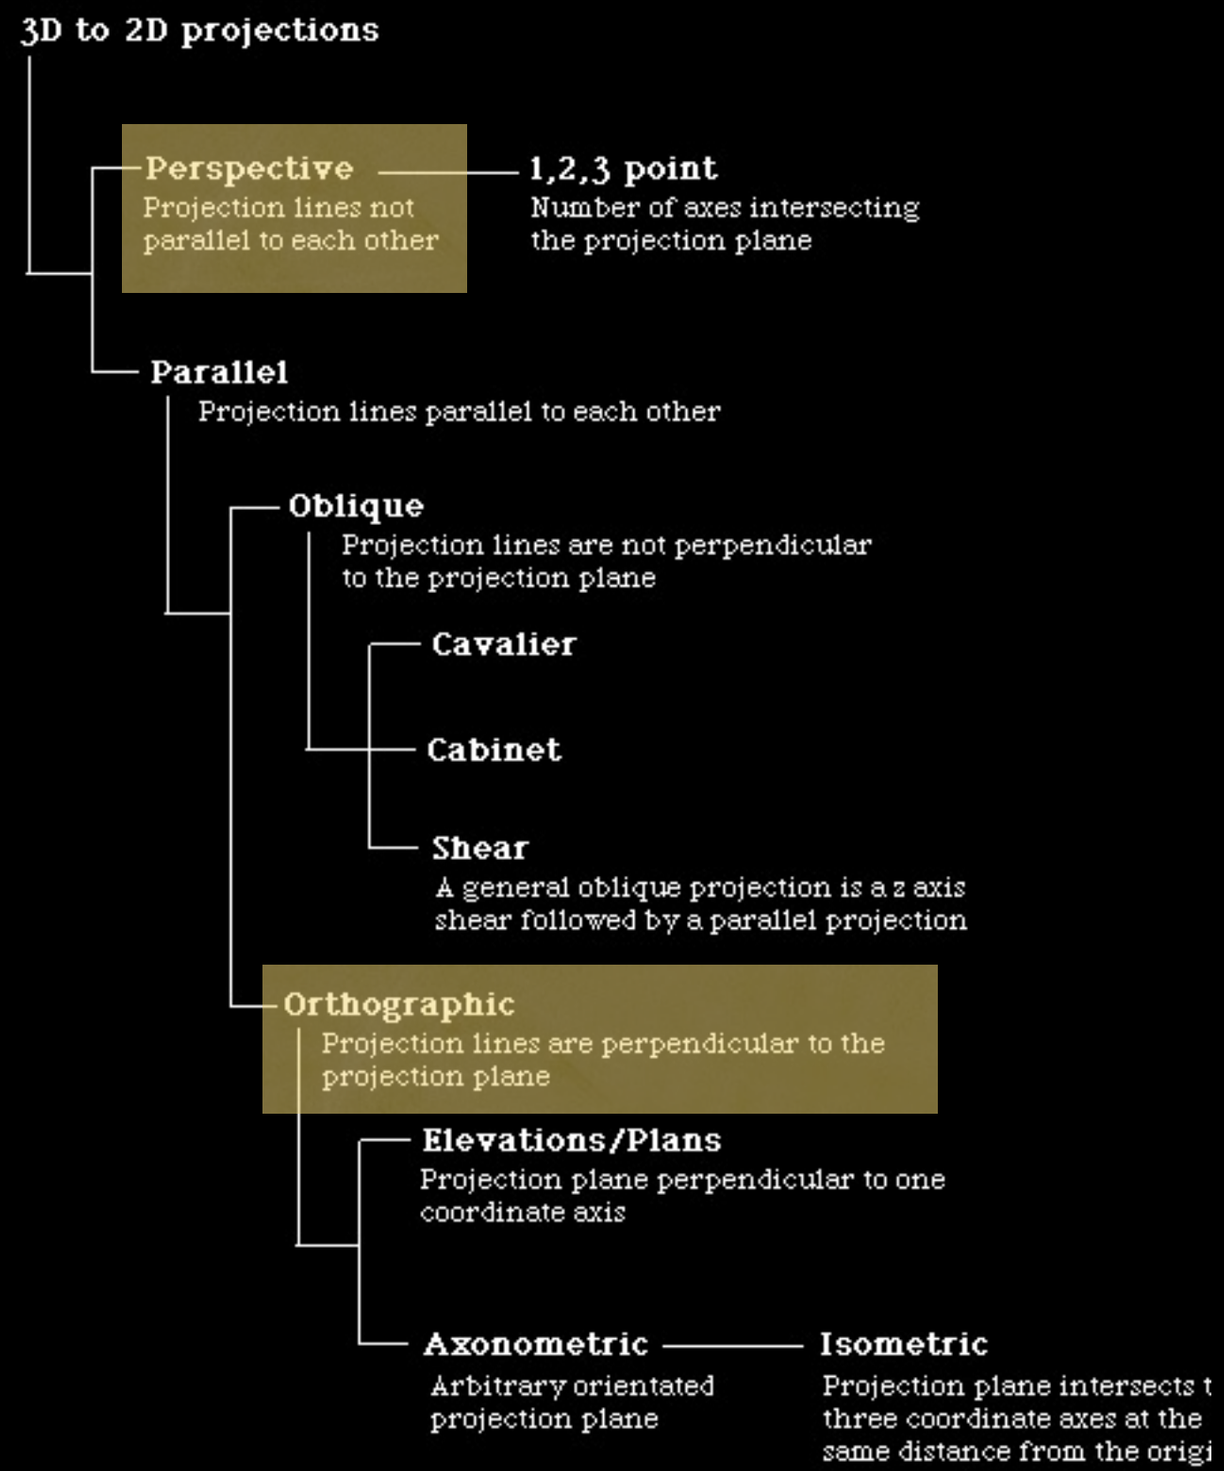

Standard perspective projection

- The correct way to think about standard idealised perspective projection is to imagine looking through a window on a real, or virtual, world.
- Draw a straight line from the camera/observer to an object in the world, where the line intersects the projection surface is where the object appears on the screen surface.

Camera = Observer
Projection plane = Display surface

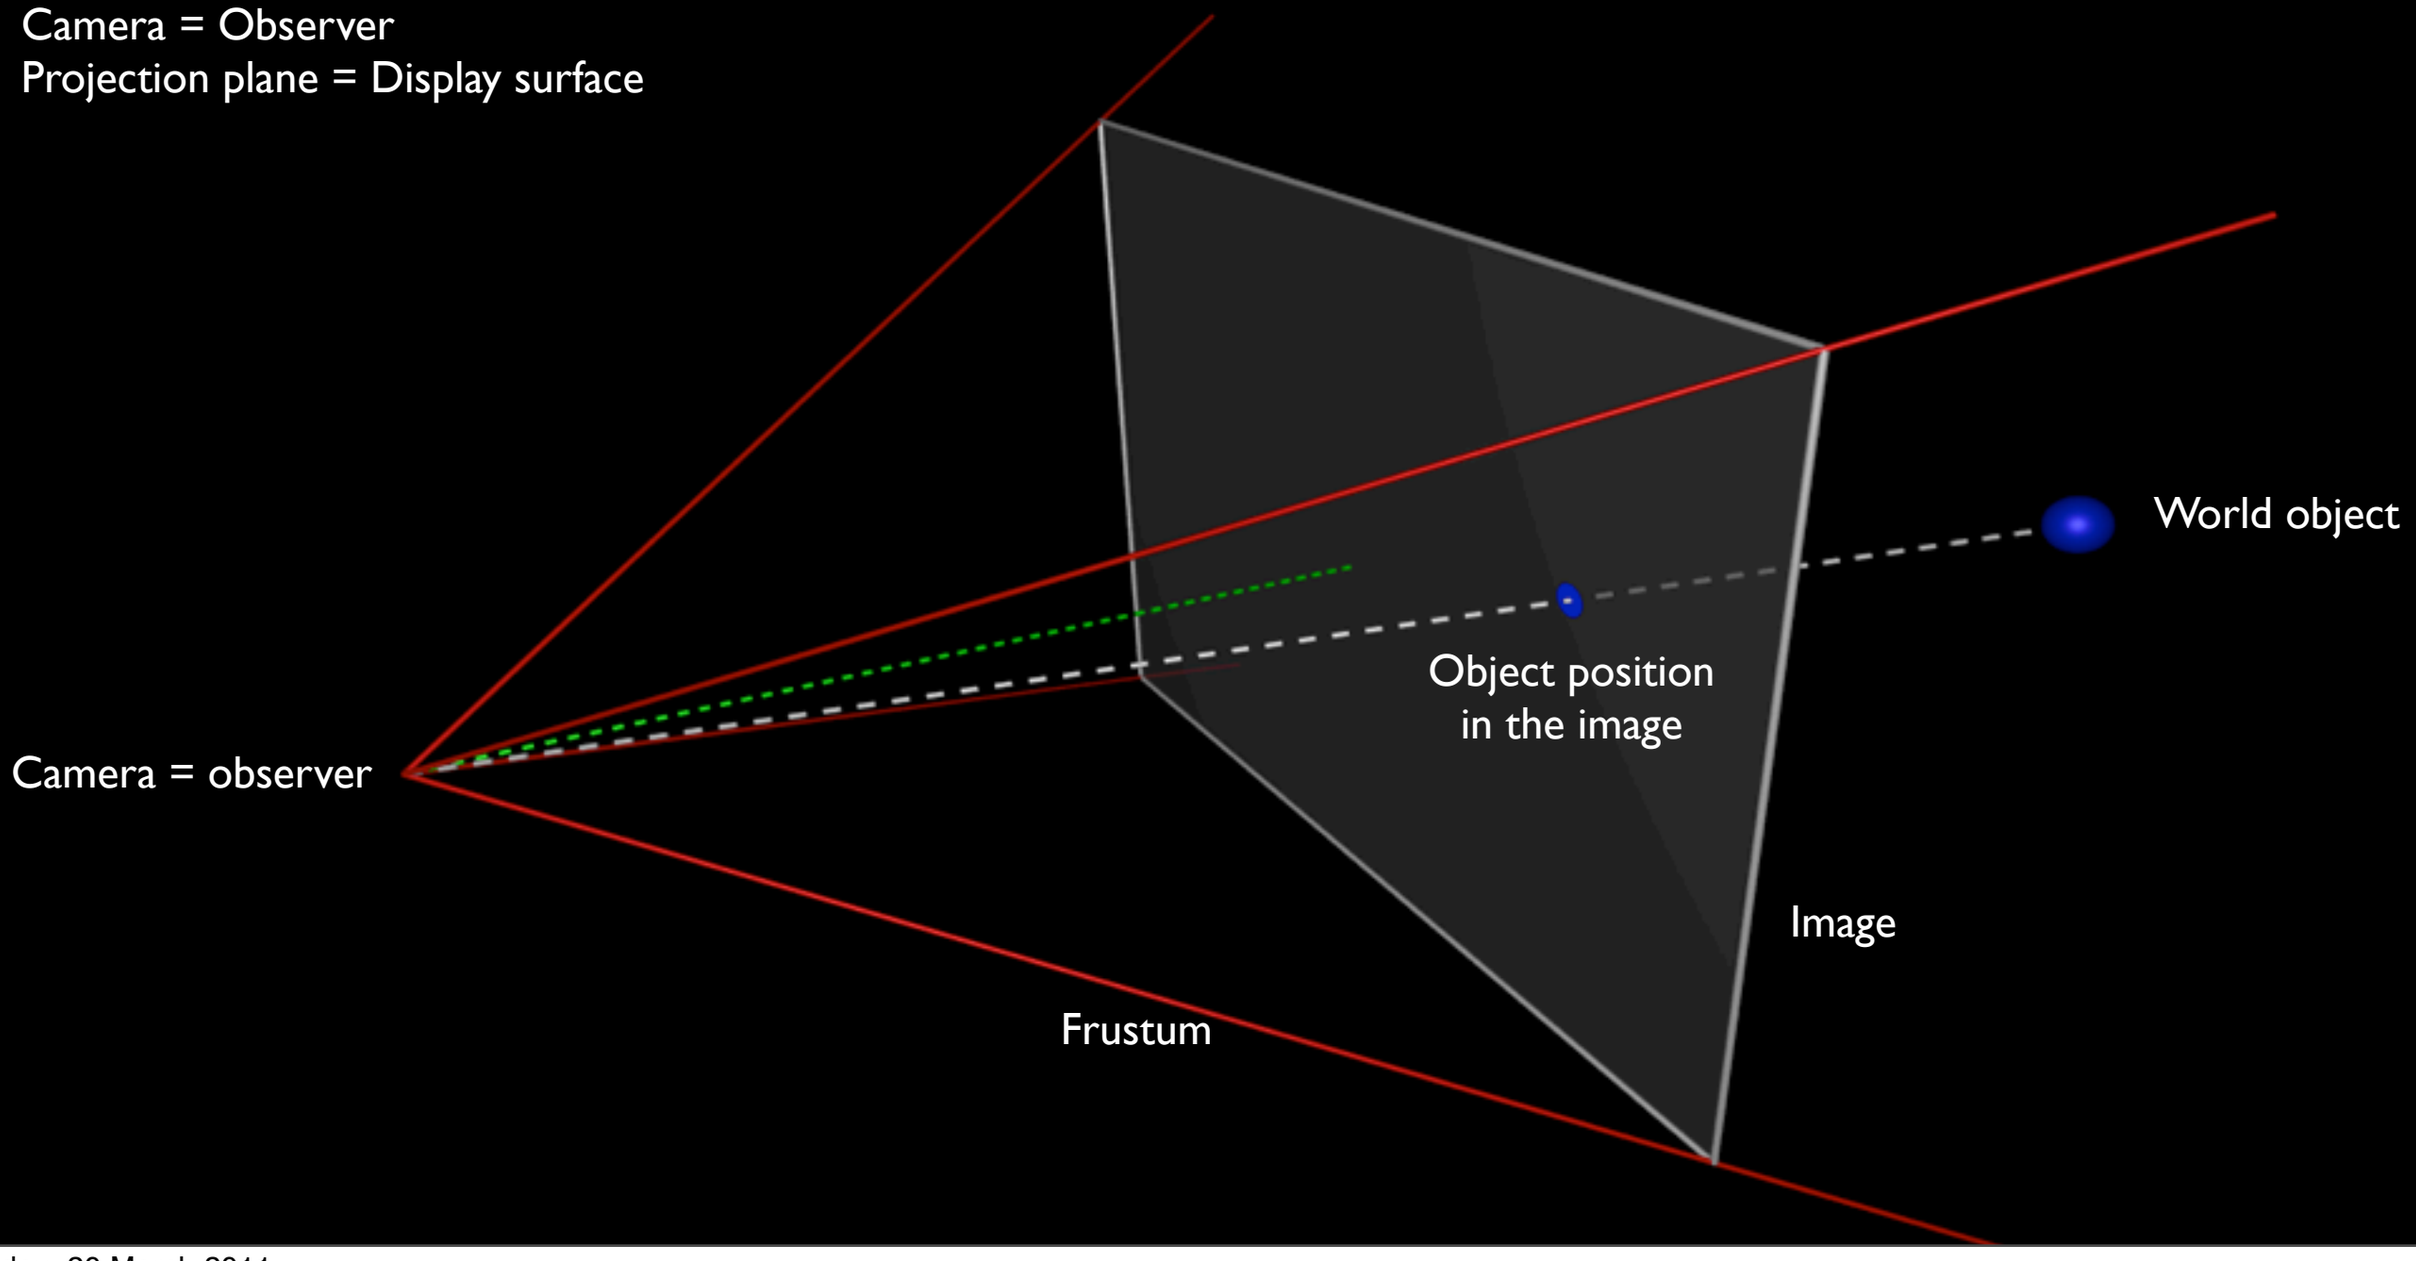

Perspective projection

- A perspective projection gets increasingly inefficient as the horizontal field of view increases.
- Many immersive displays require a wide horizontal field of view.
- These two examples have the same vertical field of view (90 degrees).
- As the horizontal FOV increases one needs more pixels to get the same effective resolution.
- Falls down completely at 180 degrees.

100 degrees

140 degrees

Cylindrical projection

- Solves the issue of efficiently capturing a wide horizontal field of view.
- The view of the world is now through a cylindrical window.
- A cylinder and plane are topologically equivalent so can represent the surface undistorted as a rectangular image, usally called a panorama.
- Doesn't need to be a full 360 cylinder.

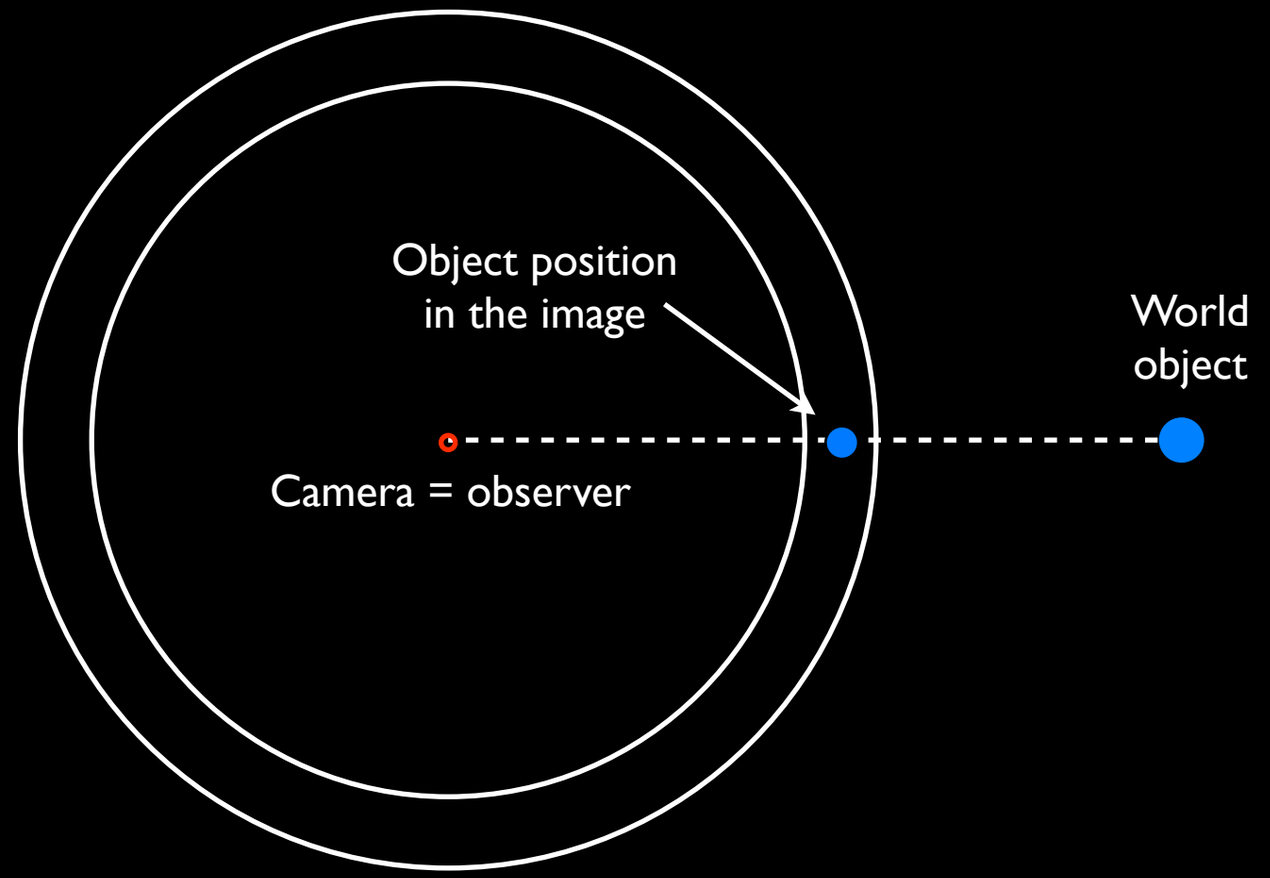

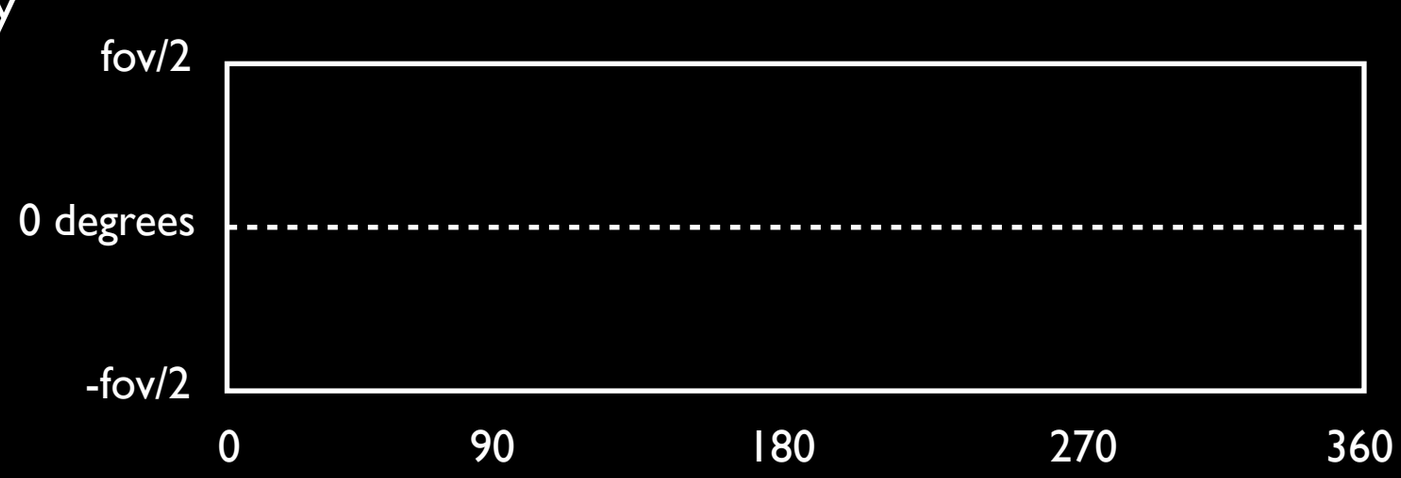

Cylindrical projection

- A cylindrical projection becomes increasingly inefficient as the vertical field of view increases.
- In the example below one needs twice the pixels to capture 30 additional degrees vertically.

90 degrees

120 degrees

Spherical projection

- Solves the issue of efficiently representing wide vertical field of view in cylindrical projections.
- Imagine a sphere around the camera, an object in the scene intersects the sphere at a particular longitude and latitude.
- A sphere and plane (image) are topologically different so there is no undistorted way of unwrapping the sphere to a plane.
- Usual approach is a so called equirectangular projection (spherical projection to region of a plane). This is also the most common texturing mode for a sphere in computer graphics.

Hemispherical projection

- A spherical projection maps the entire scene (360 degrees longitude, 180 degrees latitude) onto an image plane. What if one only needs half the scene?
- A fisheye projection is one way of flattening half of a spherical projection onto an image plane.
- Fisheye projection is not limited to 180 degrees.
  Can be less than 180.
  Can be up to 360 degrees but becomes increasingly inefficient.

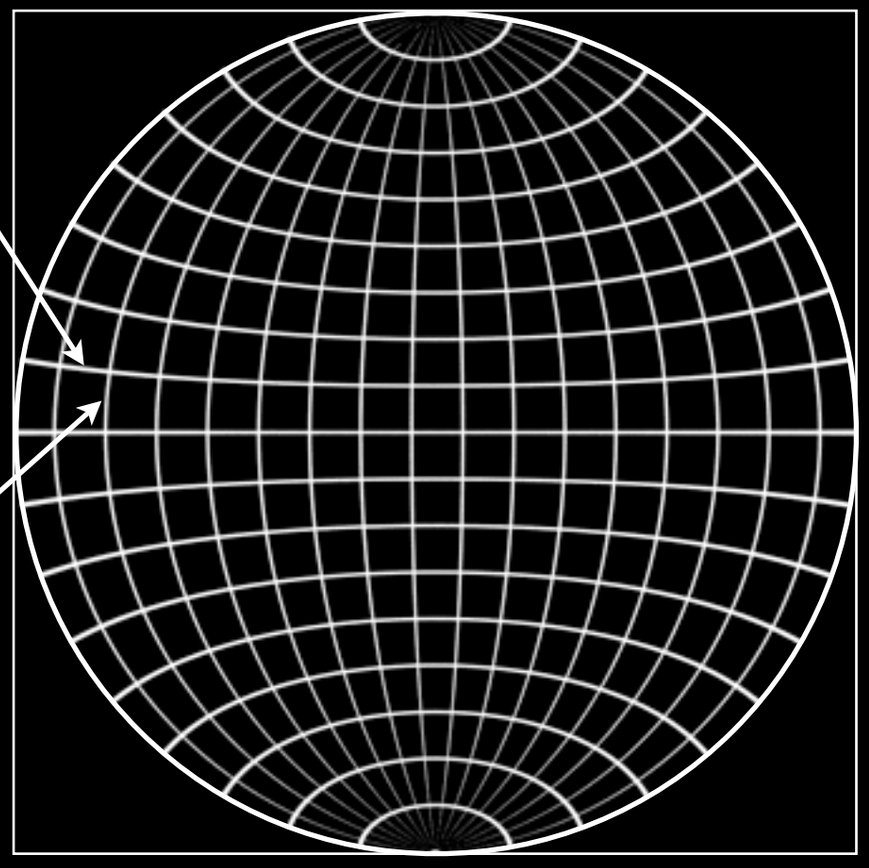

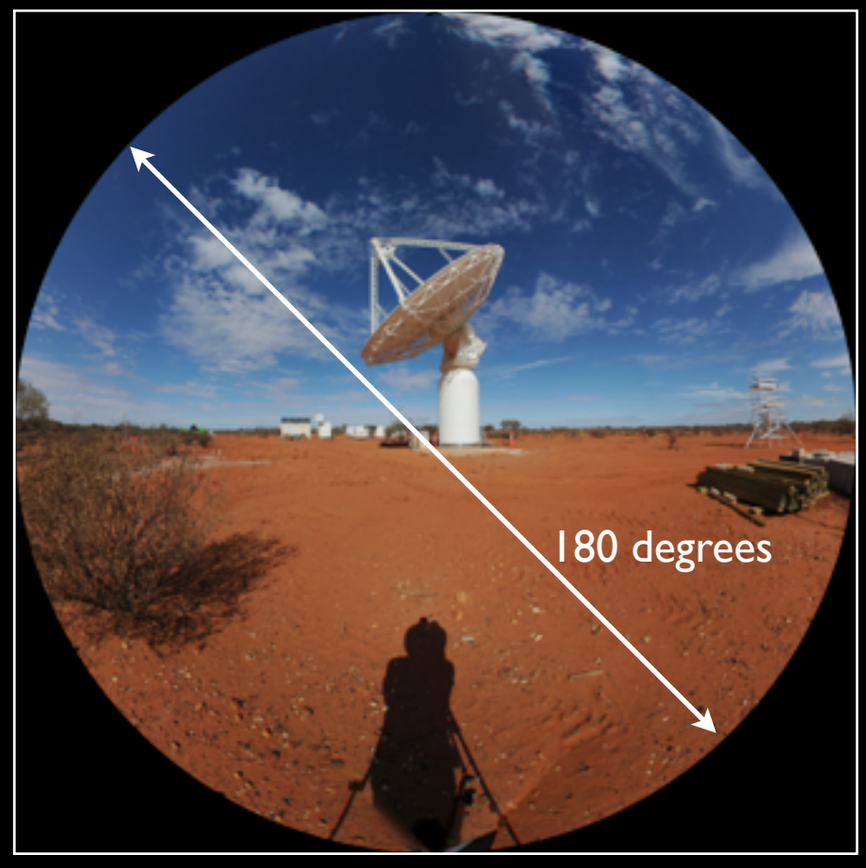

Stereoscopic perspective projection

- Extension of perspective projection to stereoscopic projection is straightforward, one perspective frustum per eye.
- The cameras horizontally offset and as such have asymmetric frustums.

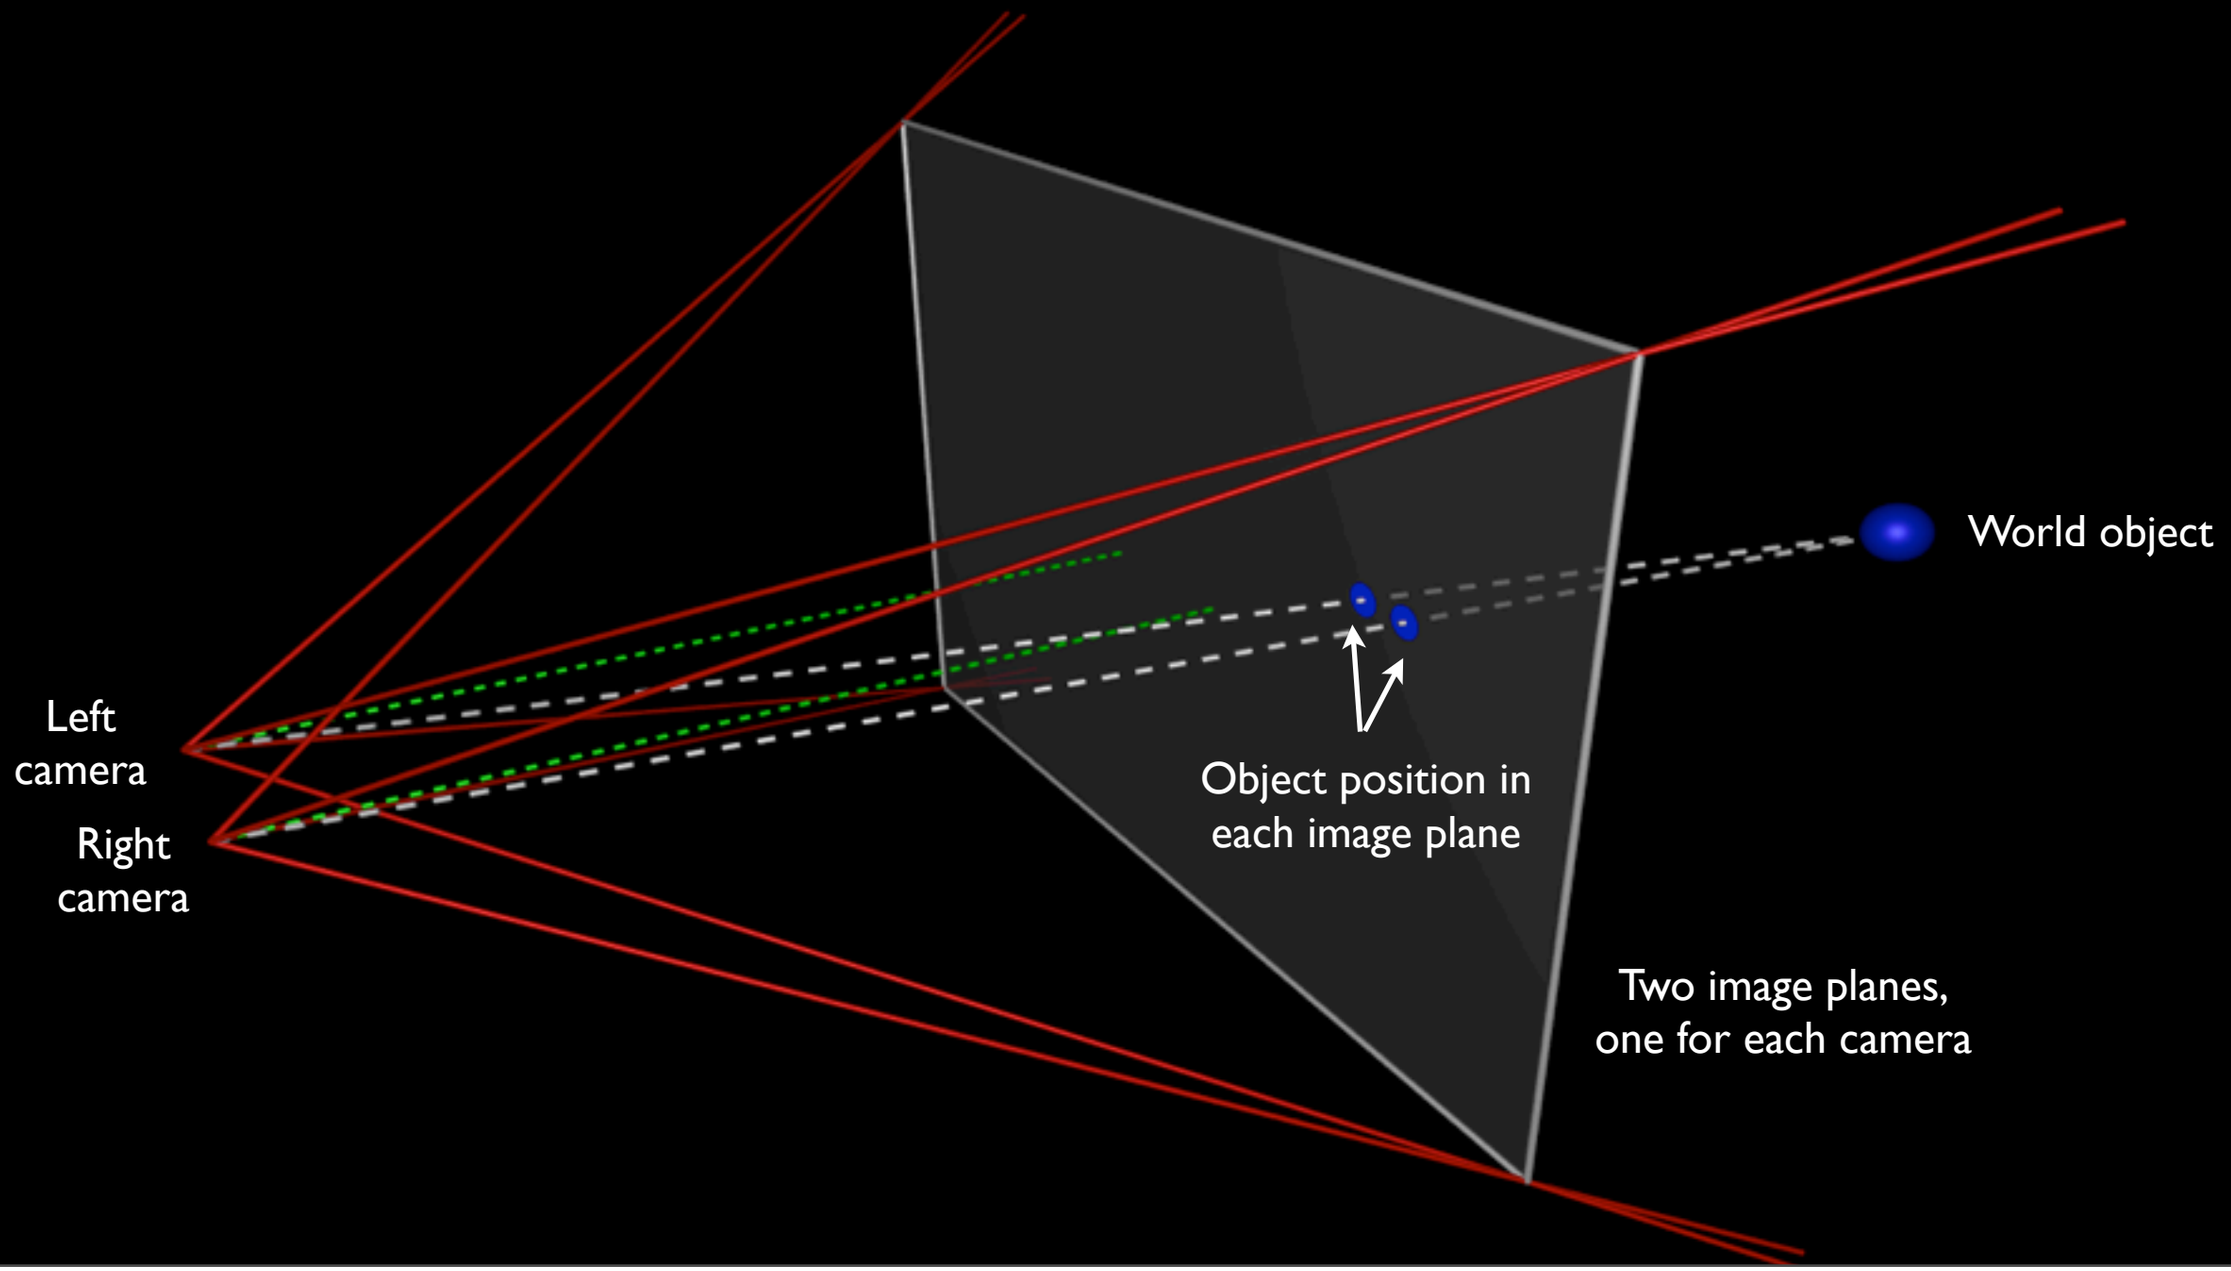

Stereoscopic perspective projection

- The correct mental picture is to imagine viewing the world through a rectangular window.
- The consequence for stereoscopic projection is the concept of an offaxis (asymmetric) view frustum.
- Two parallel cameras are not the same as rotated cameras, rotated projection planes results in a keystone type effect.
- Rotated cameras have often been used mainly due to limitation with the underlying technology, introduces vertical parallax towards the corners of the image.

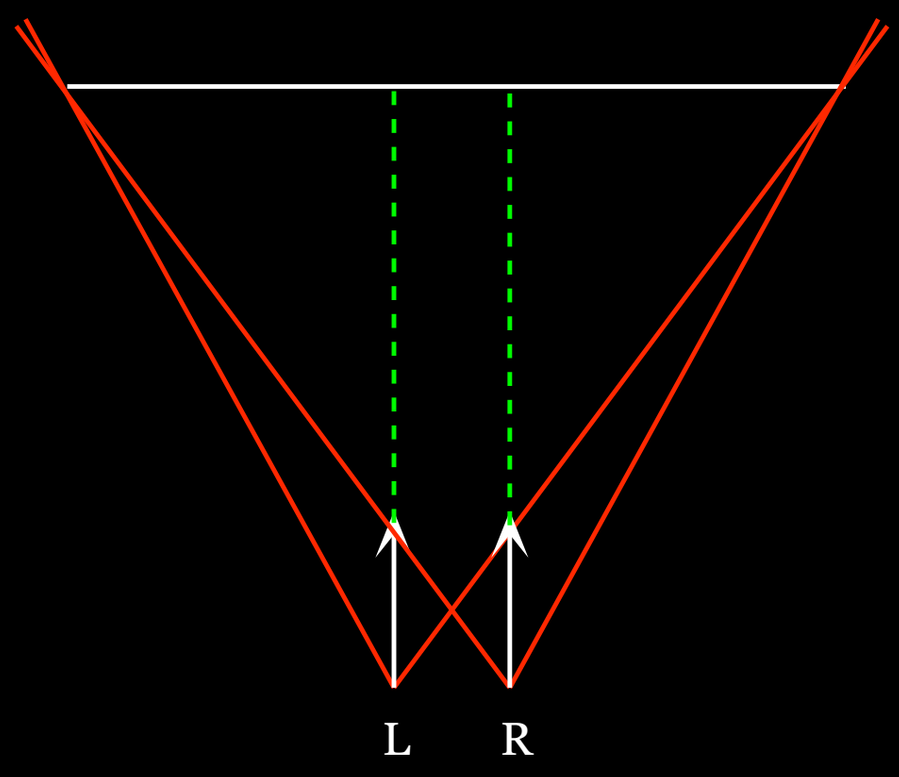

Diagram of offaxis frustum

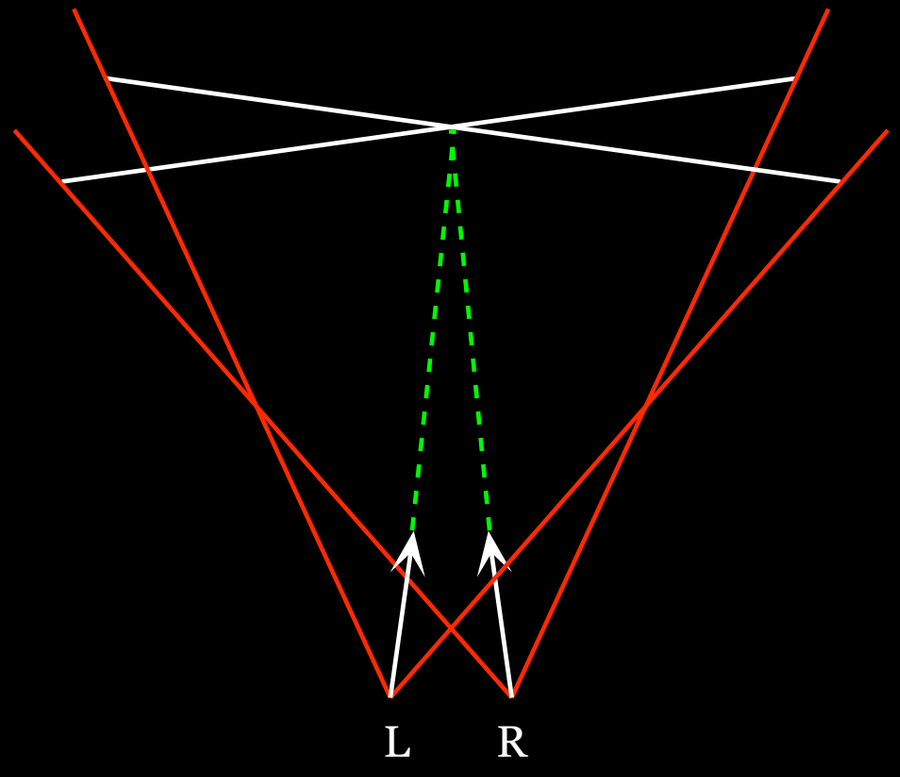

Diagram of rotated cameras

Stereoscopic perspective projection

- A stereoscopic image pair is only strictly correct (depth perception) from a single viewer position with respect to the display surface.
- Viewing from all other positions results in a distorted view.
- Correct depth perception also depends on the viewing direction, obviously an issue for an audience.
- In virtual reality realtime environments this can be corrected for by head tracking and adjusting the frustums in realtime to preserve the correct view through the window.
- Head tracking is clearly for single person experiences and is not an option for photographic or filmed content.

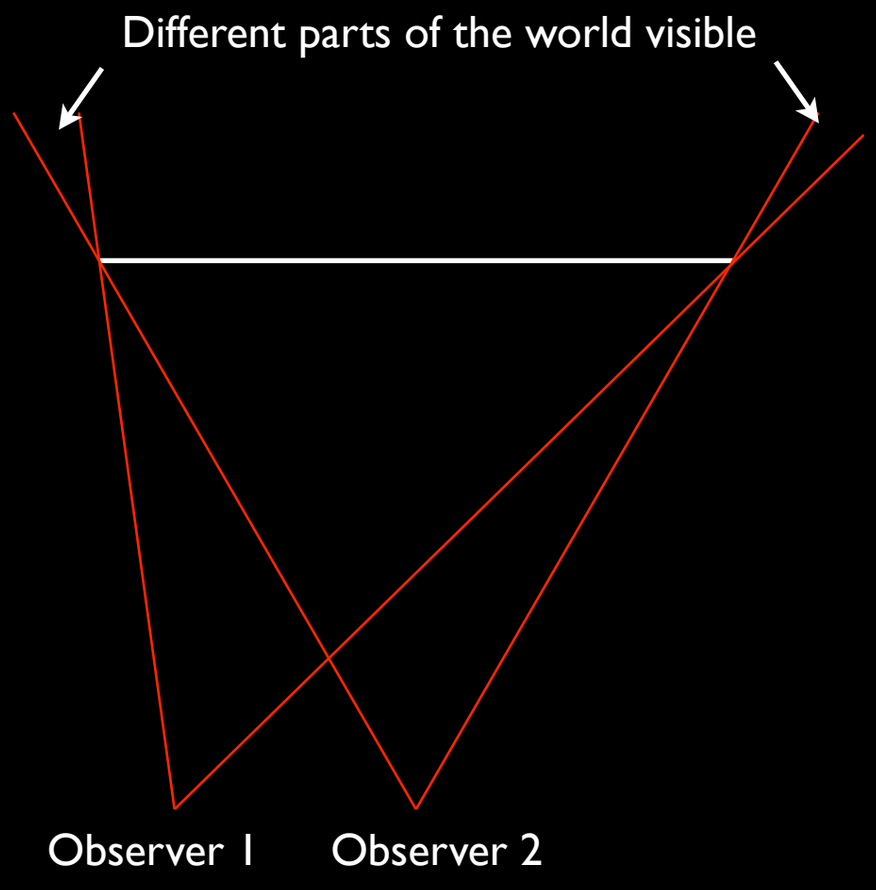

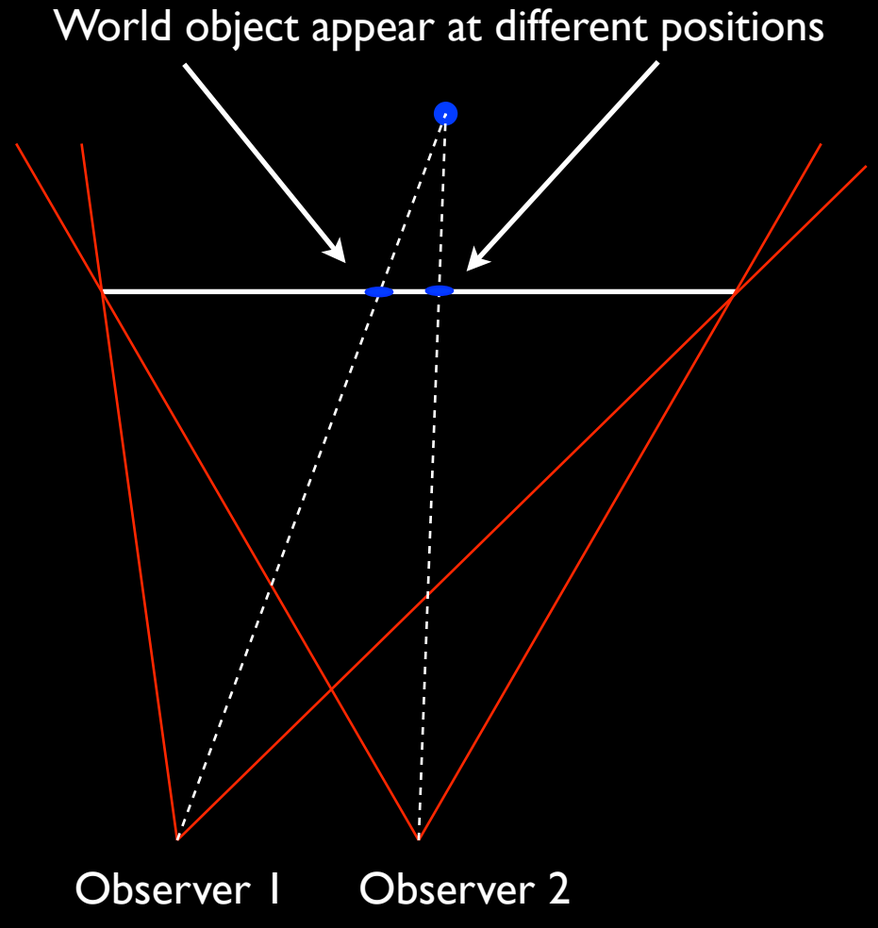

Stereoscopic perspective projection: parallax

- The difference in position of the projection of a world object onto the left camera and right camera image plane.
- For stereoscopy there is ideally only horizontal parallax, our eyes are offset horizontally not vertically.
- Positive parallax features appear behind the screen.
  Negative parallax features appear in front of the screen.
- Maximum positive parallax is equal to the camera/eye separation and occurs for objects at infinity.
- Maximum negative parallax is infinite, a key consideration is acceptable negative parallax. Dependent (among other things) on the degree of ghosting in the presentation system, the age of the viewer, etc.

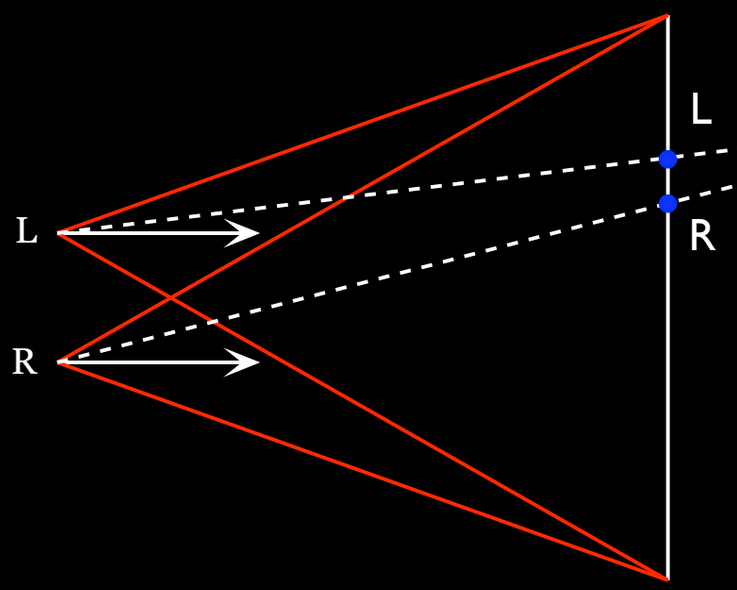

Positive parallax
Object behind projection plane

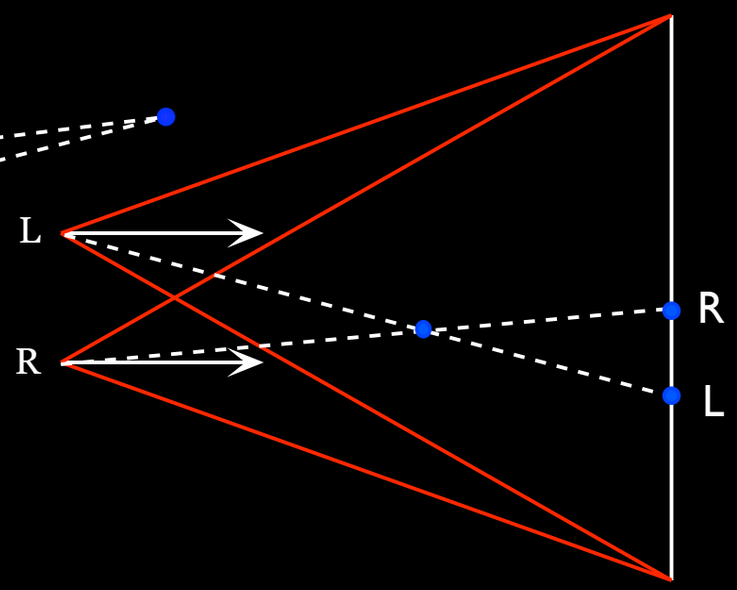

Negative parallax
Object in front of projection plane

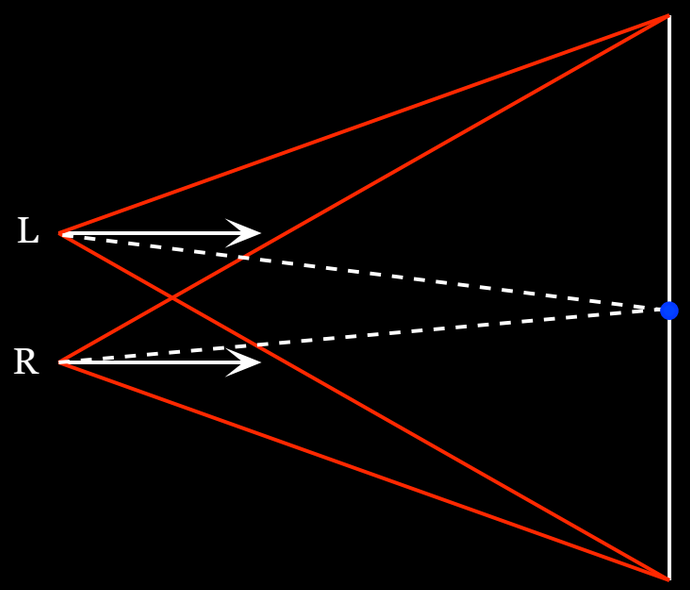

Zero parallax
Object at projection plane

Multiple walls

- Same principles apply. For 2 stereoscopic walls there are 4 distinct view frustums.
- Note the dependence of the frustums on the viewers position and their viewing direction.
- For an undistorted view these are single person experiences which require head tracking.
- Without realtime head tracking the content is always created with a single assumed viewer position and view direction.
- Acceptable viewing by a group can be acceptable if the viewers are located in the position and are generally looking in the same direction as the content assumes.

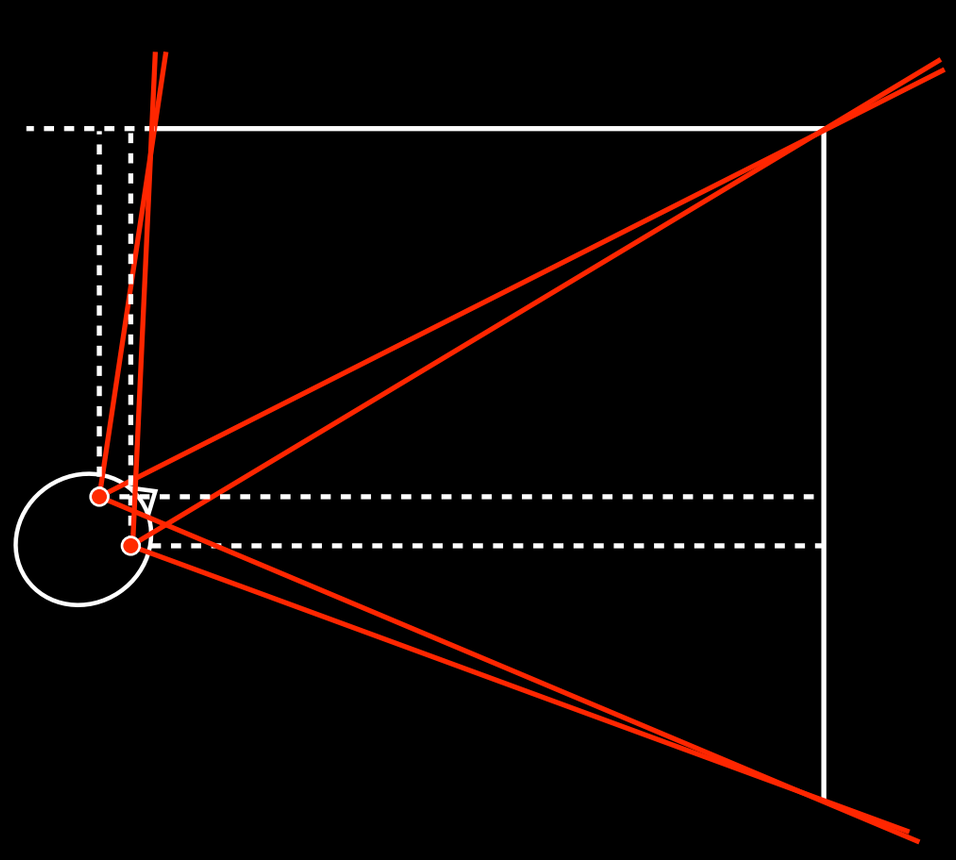

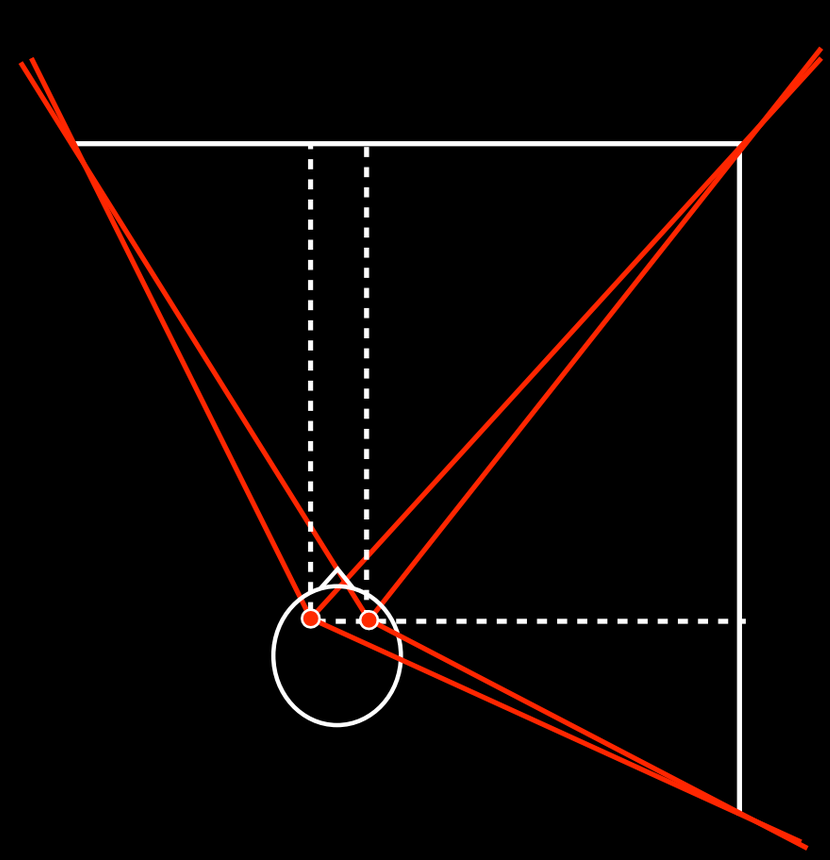

Omni-directional cylindrical stereoscopy

- By making a small concession it has been known for some time that one can create stereoscopic cylindrical image pairs that can be viewed without head tracking or viewed simultaneously by multiple people all looking in different directions.

- Still only strictly correct for a single viewing position (generally the center of the cylinder) but have relaxed the view direction constraint.

- The stereoscopic imagery is only strictly correct directly along a vertical strip in front of the viewer and is increasingly incorrect away from this strip. Acceptable because the frames of the glasses limit the field of view to this region.

- An immersive effect is still achieved because one can still perceive imagery on the extreme left and right, past the glass frames. Not in stereo but neither is it in real life.

- The camera model has two (vertical slit) cameras rotating about the midpoint of their respective nodal points. Can capture (still) images photographically with a single camera (dual slit).

Left eye

Right eye

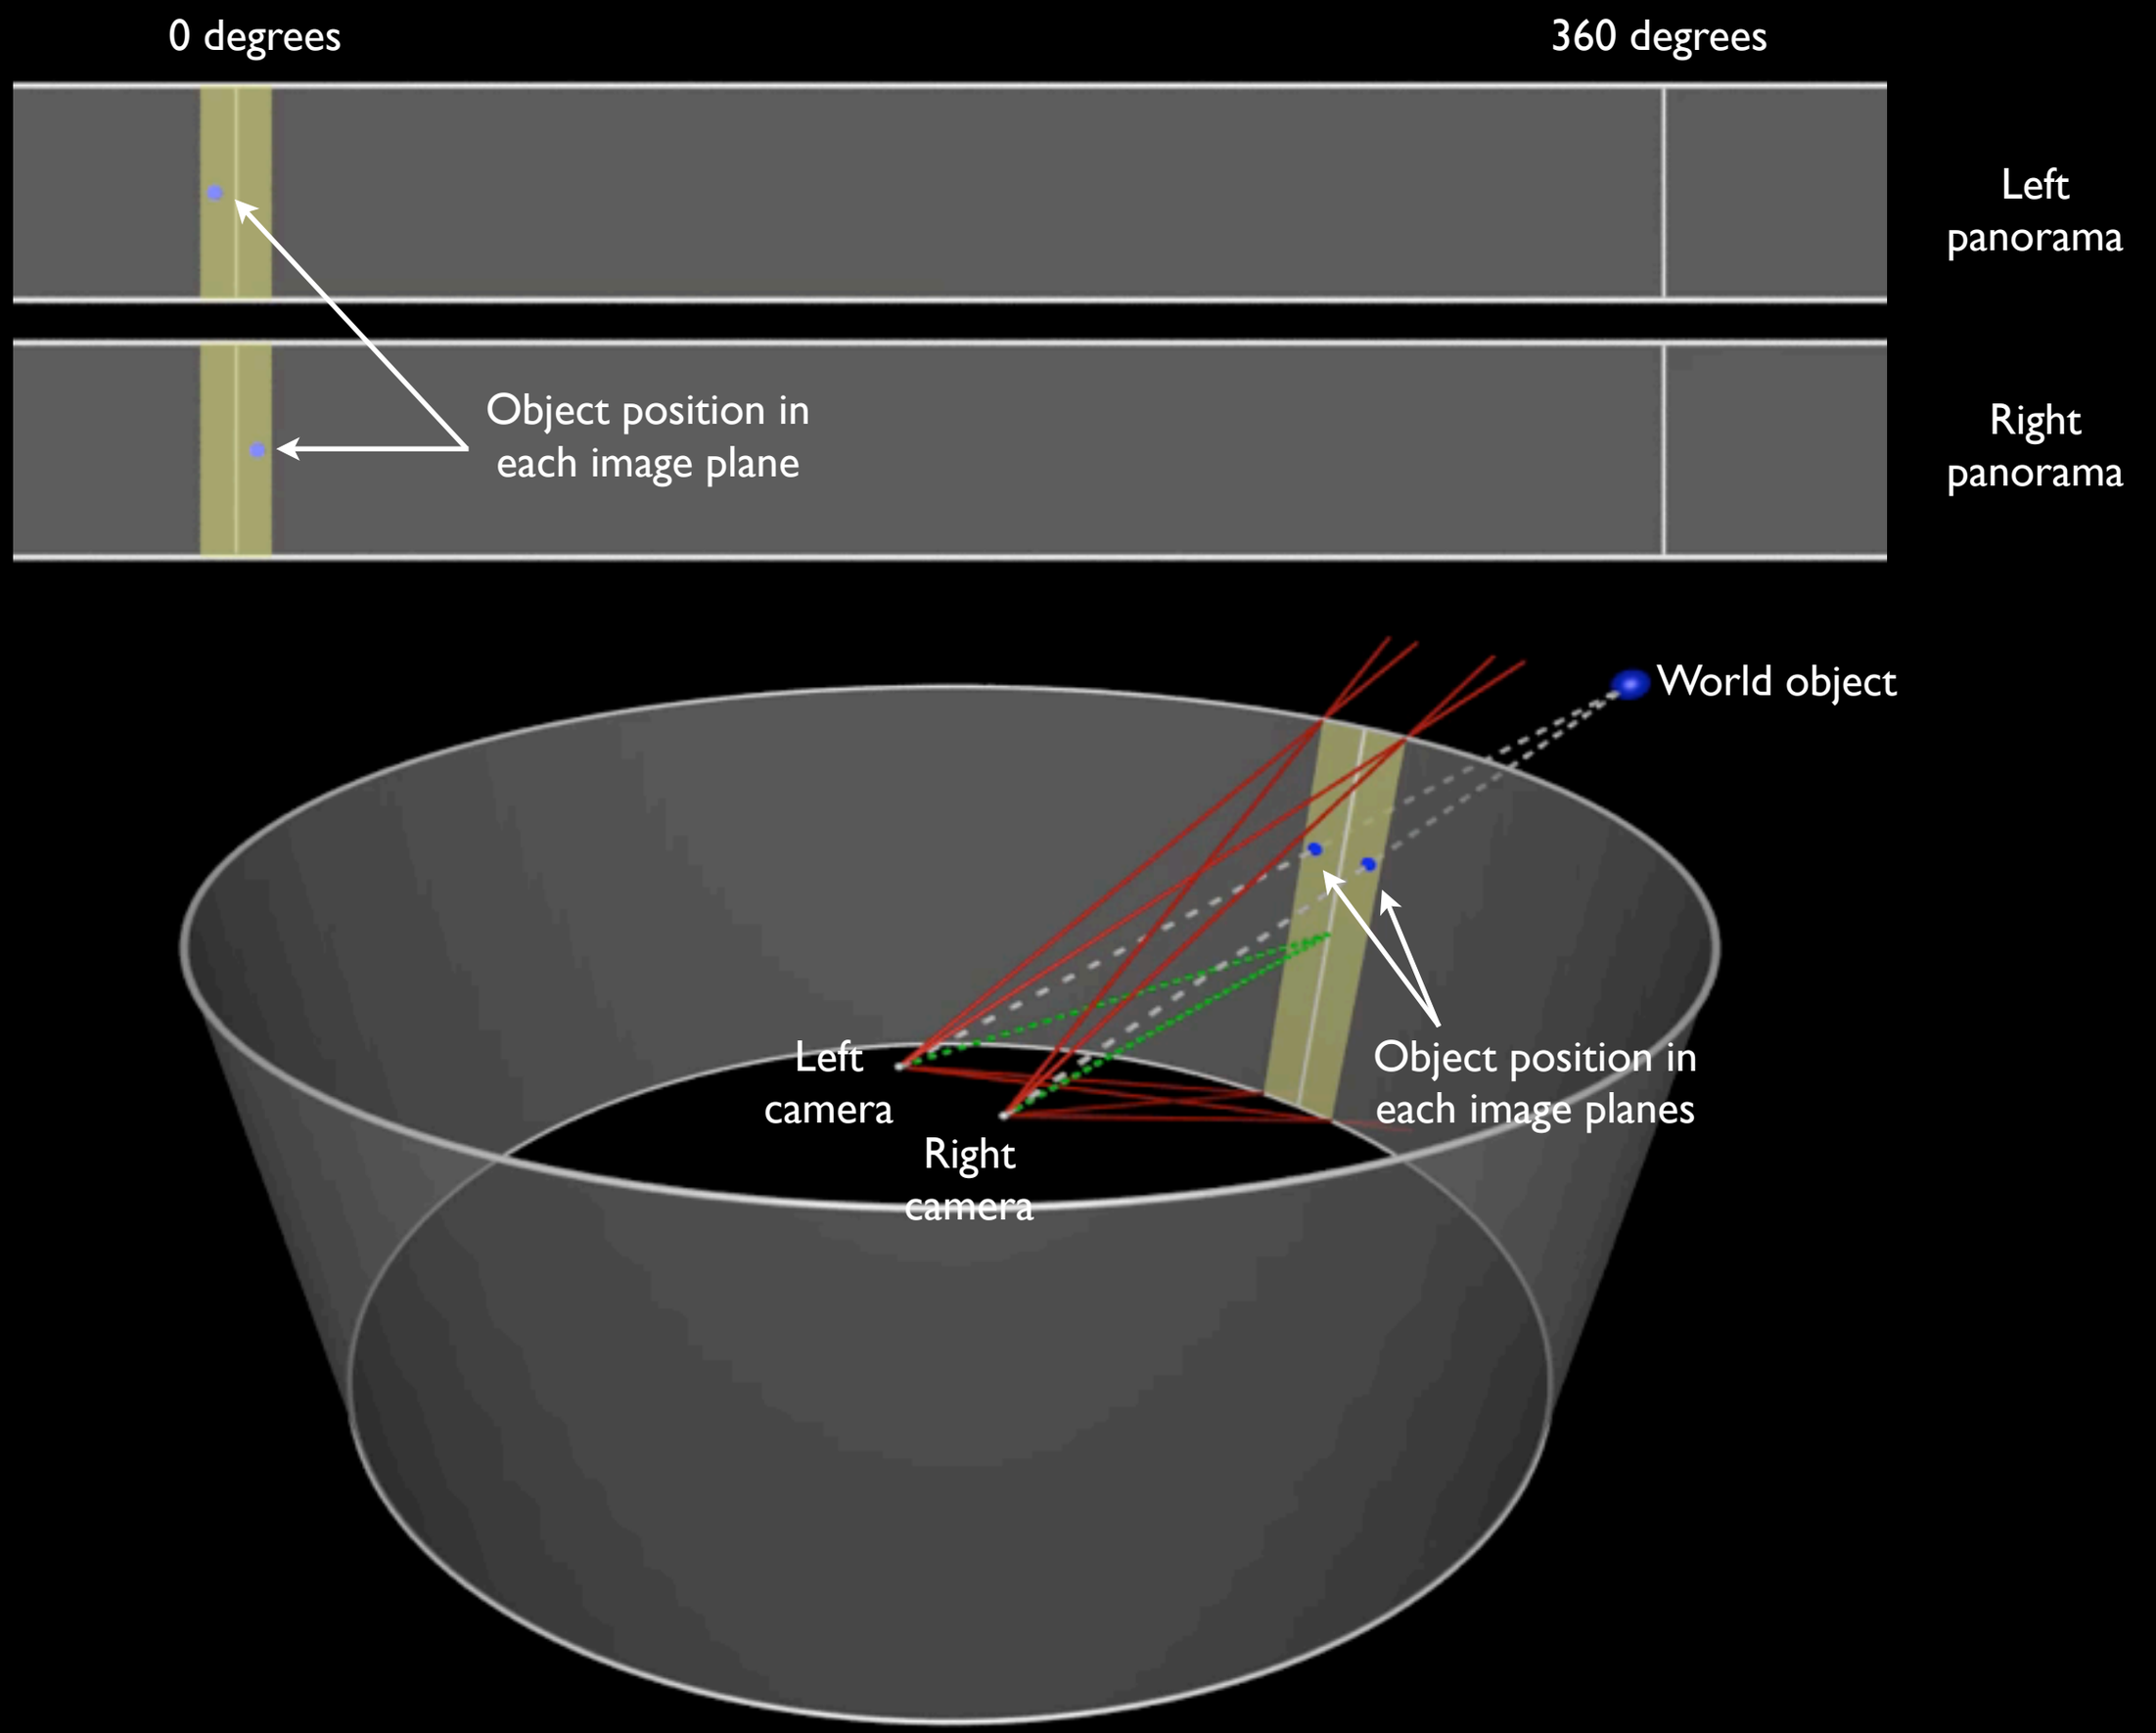
0 degrees
360 degrees
Left panorama
Right panorama
Object position in each image plane
World object
Left camera
Right camera
Object position in each image planes

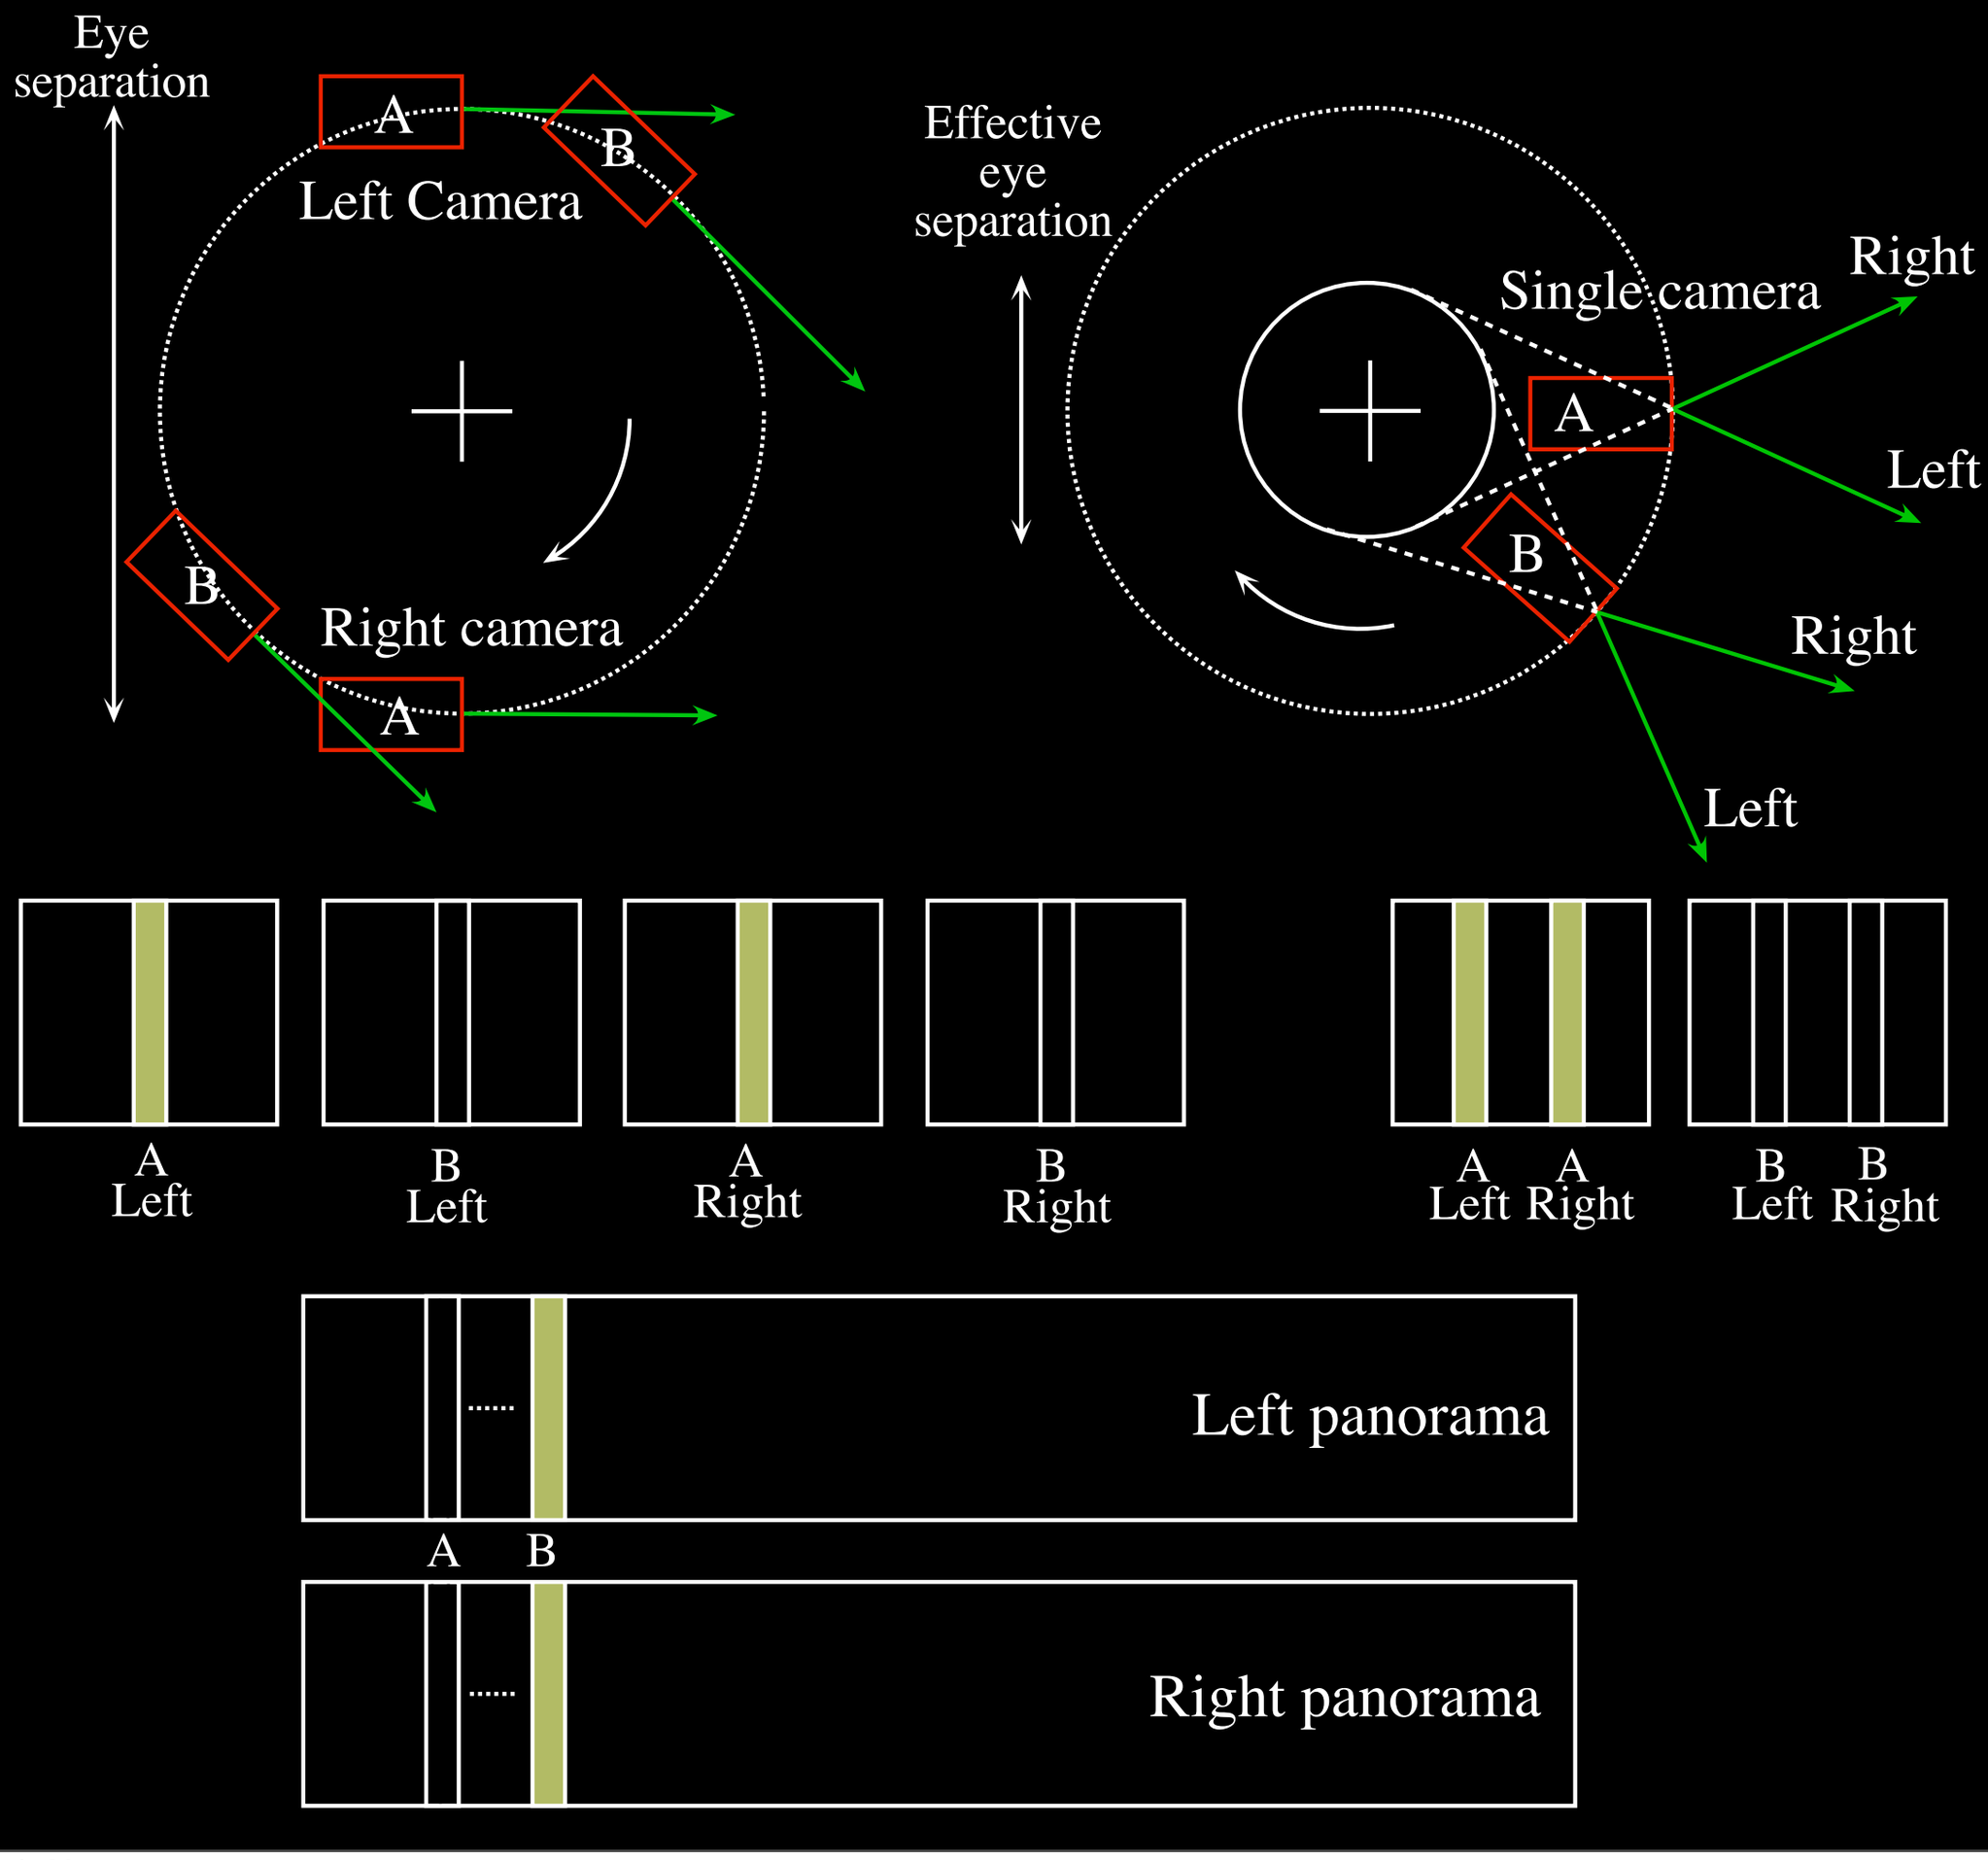
Eye
separation
A
B
Left Camera
B
Right camera
A
Effective
eye
separation
Single camera
Right
A
Left
B
Right
Left
A
Left
B
Left
A
Right
B
Right
A
Left
A
Right
B
Left
B
Right
.....
Left panorama
A
B
.....
Right panorama

Hemispherical displays

- Traditionally hemispherical displays were planetariums.
- Digital planetariums today are immersive theatres capable of presenting high resolution digital video.
- A fisheye projection is the natural way to represent the imagery for a hemispherical display.

5m inflatable

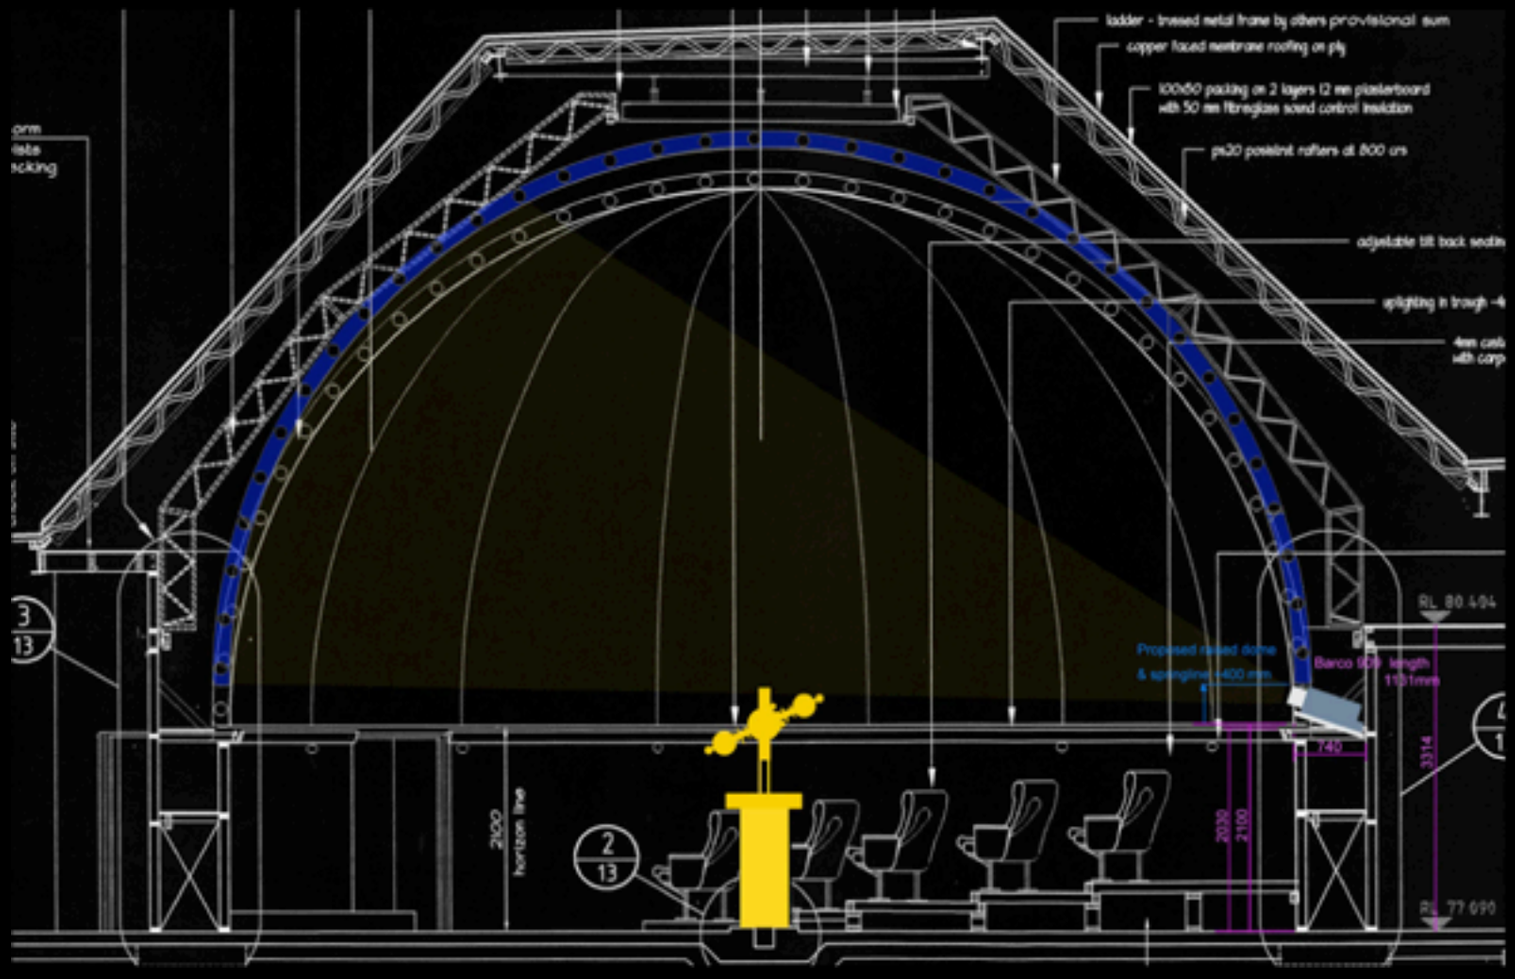

Fixed, generally steel mesh

Fibreglass (Courtesy Mirage3D)

iDome: personal immersive display

- A small, personal, upright hemispherical surface. 3m diameter.
- 180 degrees horizontally by 145 degrees vertically.
- A projector and fisheye lens is less convenient for the iDome because it would be located where the viewer should be.
- Spherical mirror (innovation by the author) conveniently locates the projector behind the base of the dome.

iVEC@UWA Visualisation laboratory

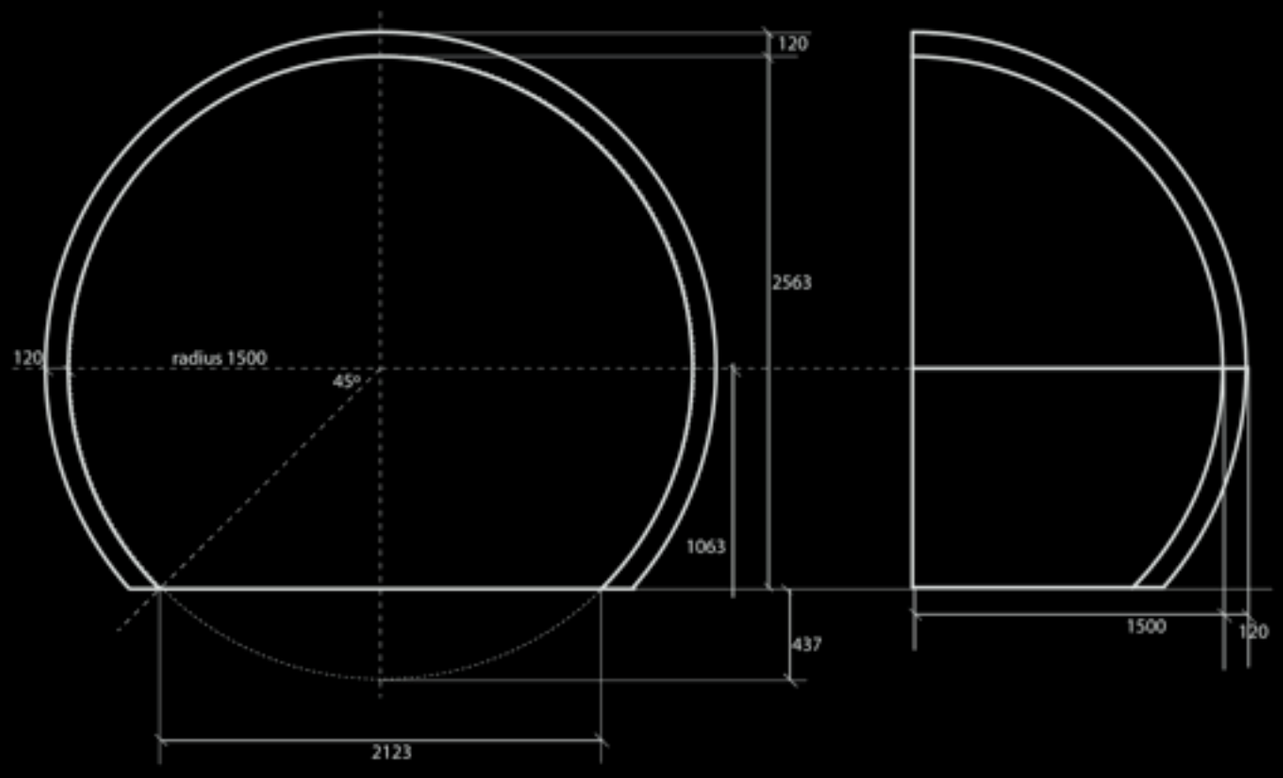

Courtesy Giorgio Marchetto

Considerations

- Angle of the dome and the audience position is critical to the content process.
  - Planetarium: 0 degrees.
  - Tilted (eg: iMAX): commonly 30 degrees.
  - iDome: 90 degrees.

- For example, content for a planetarium is generally unsuited to an iDome, and visa-versa.

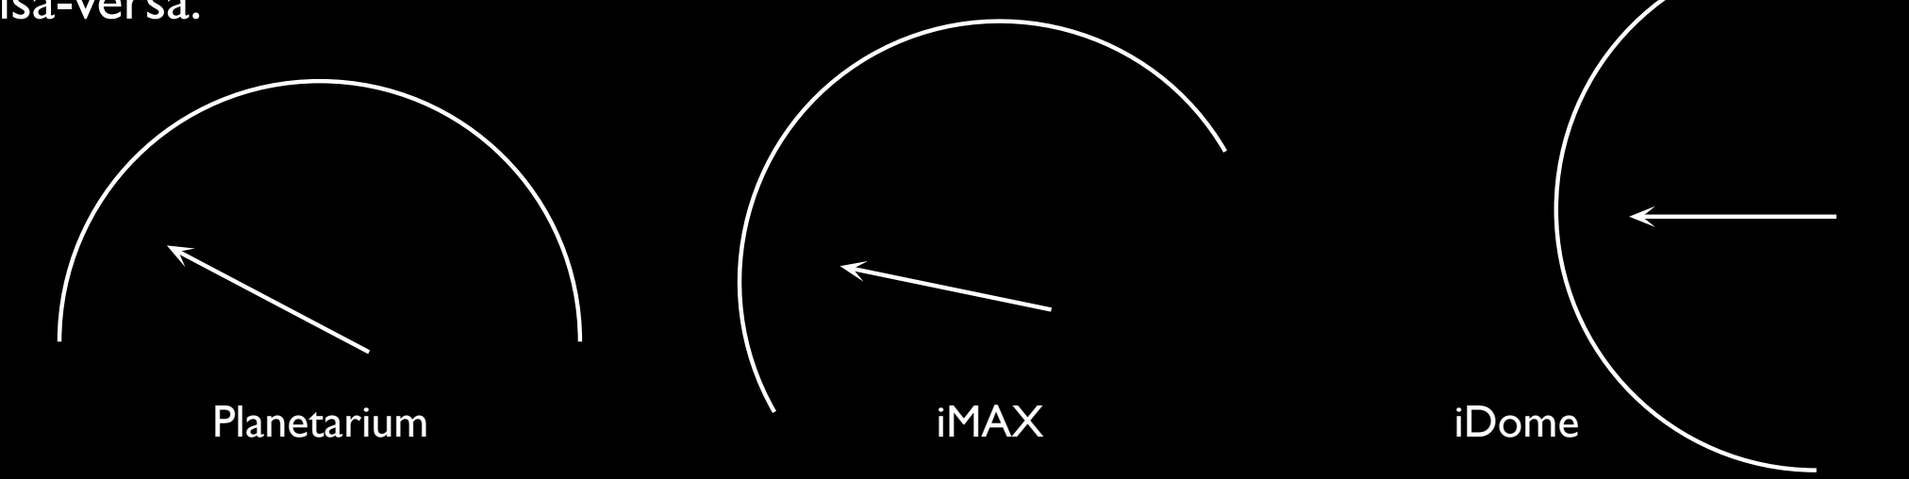

Planetarium iMAX iDome

- Content can be created for any viewer position, all other positions will see distorted views the degree of which depends on their distance from the intended position. eg: straight lines will not appear straight.

Content creation: Camera based image and video capture.

- Stereoscopic filming: rich history and increasingly commodity.
- High resolution panorama capture of still scenes using image stitching.
- Panorama capture of near still scenes using rotating camera rigs.
- Fisheye photography of still images and timelapse for dome displays.
- Spherical video capture using the LadyBug range of cameras.
  - generating reasonable cylindrical projection movies.
  - generating reasonable resolution fisheye movies.
- Capturing sufficiently high resolution fisheye video for todays digital domes is still a challenge.

Remember

- The model for how to form the 3D to 2D projection is based upon the display surface being a window on the world to the viewer.
- All content is only strictly correct for a single viewer position and direction.
  All other viewers will receive a distorted view.

Stereoscopic filming

- A very long and rich history in both stereoscopic photography and filming.
- Many home made rigs have been built over the years.
- The recent commoditisation of stereoscopic capable television and monitors is seeing a similar commoditisation in video cameras.
- Still exist challenges for variable eye separation: stereoscopic filming of distant landscape vs small scale objects. For visual effect the camera separation is related to the scale of the material.

IMAX underwater camera

Panasonic semi-pro video camera (2010)

Home made rigs

Stitching multiple images

- Higher resolution spherical panoramic images can be acquired using a large number of shots.
- One takes multiple overlapping images at different latitudes and longitudes. These are stitched together and blended to form a single seamless spherical projection.
- This can be automated with motorised units such as the Gigapan, and others.
- In general a regular scan in longitude and latitude is performed even though this is inefficient towards the north and south pole.

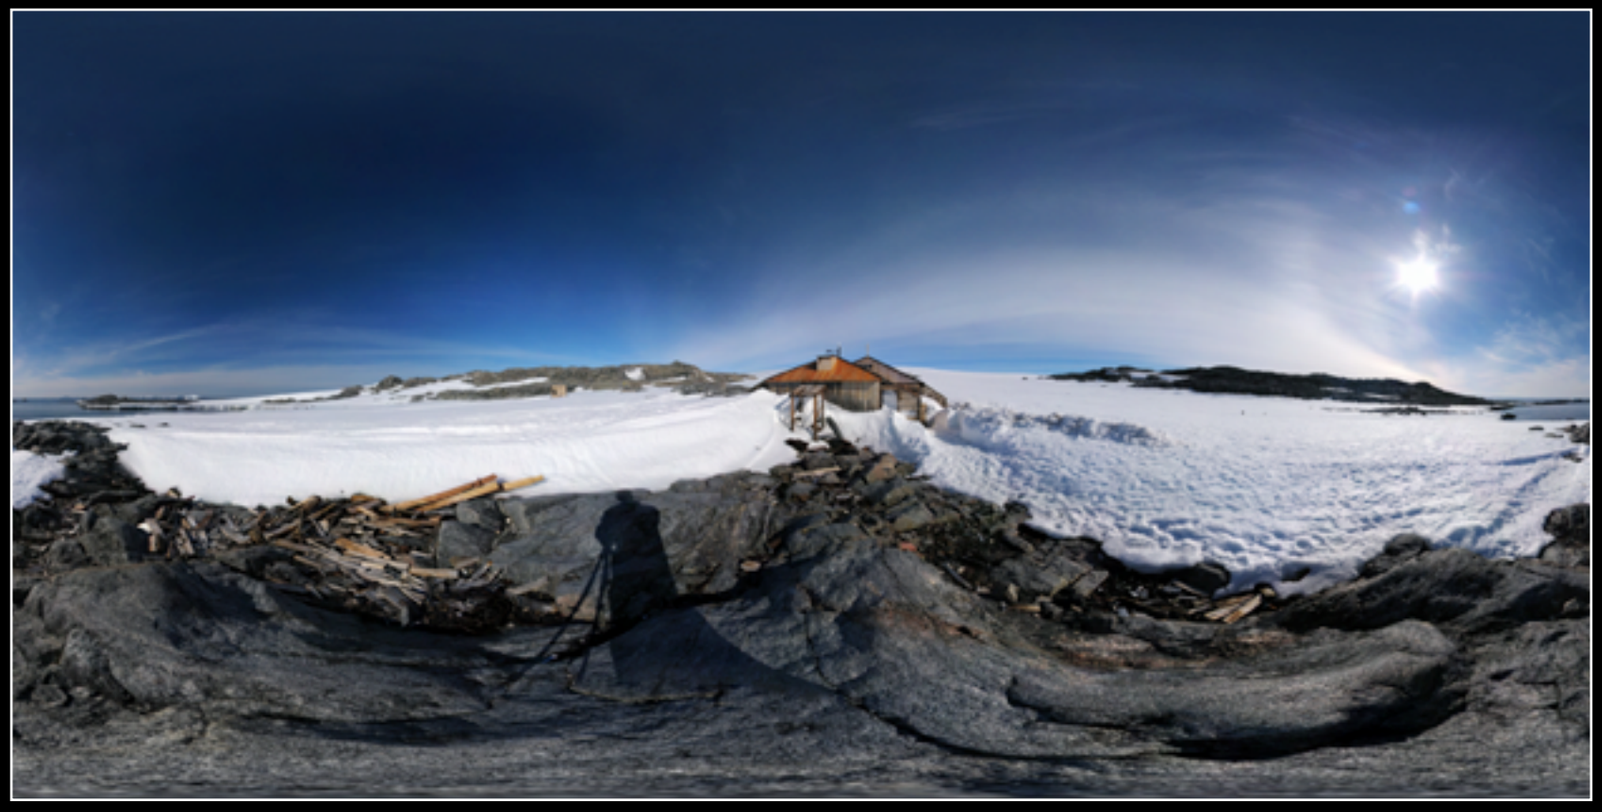

Mawsons hut (Antarctica), courtesy Peter Morse.

Gigapan mount

Stitched photography example

- The only way to get resolution above the highest resolution camera sensors available.
- Don't expect stitching errors with some of the more recent machine vision algorithms if the camera is panned about the nodal point.

Hurleys darkroom, Mawsons hut (Antarctica), courtesy Peter Morse. 40,000 by 20,000 pixels

Rotating slit camera

- Direct implementation of the omnidirectional cylindrical geometry showed earlier.
- Left and right pair captured together so some modest movement in the scene is acceptable.

Roundshot camera

Left eye

Right eye

Ephesus. Courtesy Sarah Kenderdine, Jeffrey Shaw

Fisheye lens

- Note the distinction between ultra wide angle (often called fisheye in the photography industry) and circular fisheye.
- 180 degree fisheye is not always required, for example planetarium audience below the the rim of the dome, or iDome viewer position back from the rim.

Canon EOS 5D MkII and Sigma 8mm fisheye lens

170 degree wide angle fisheye

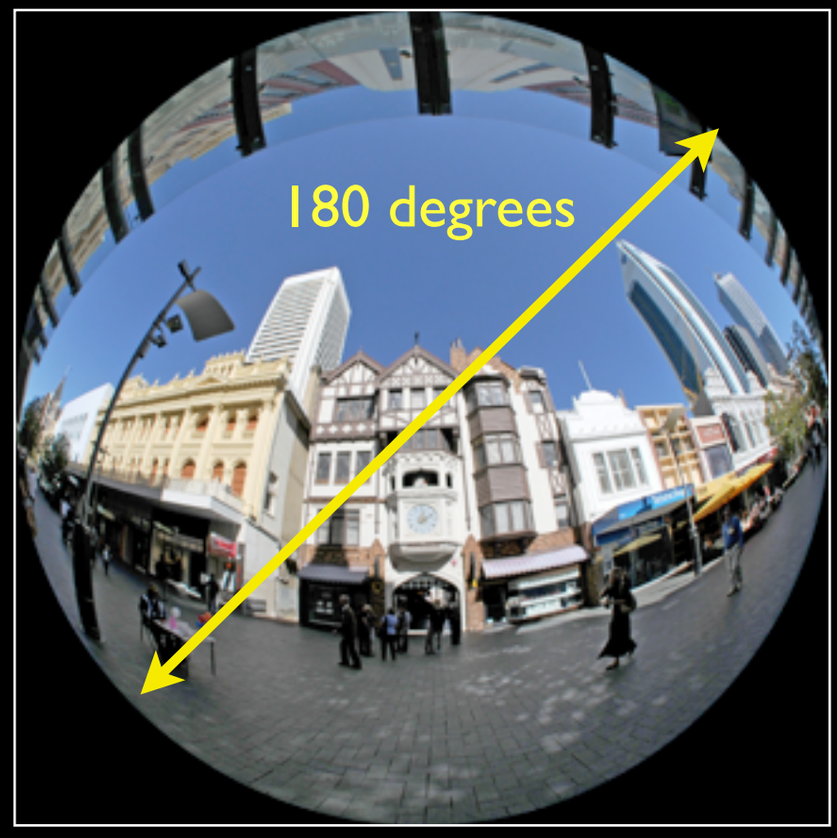

Circular fisheye

Sensor size

- Need to consider relationship between lens image and sensor size.
- Same applies to video cameras and fisheye lens.
- Truncated fisheye not uncommon for planetariums (cut off at the back) and the iDome (cut off at base). Achieves higher resolution.

Full sensor utilisation from
Canon EOS 5D MkII and Sigma 8mm fisheye lens

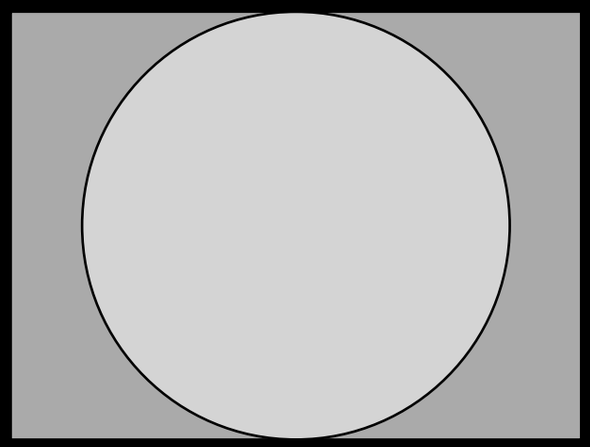

Optimal

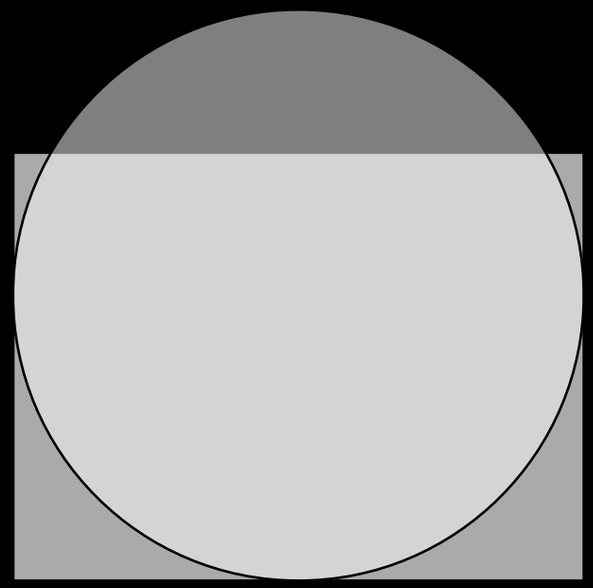

Optimal for truncated projection

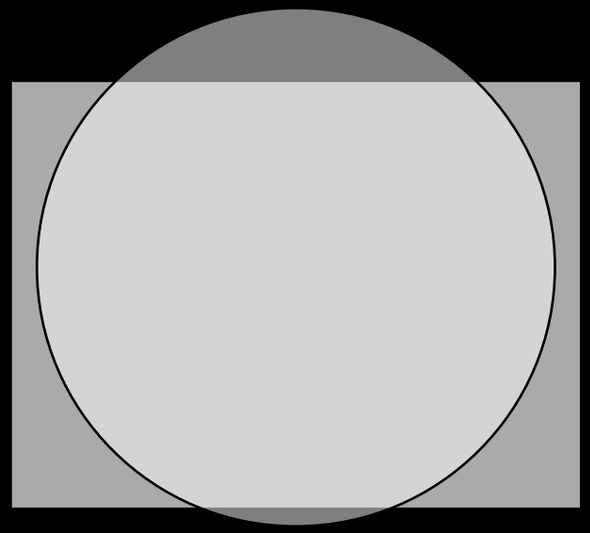

Too large

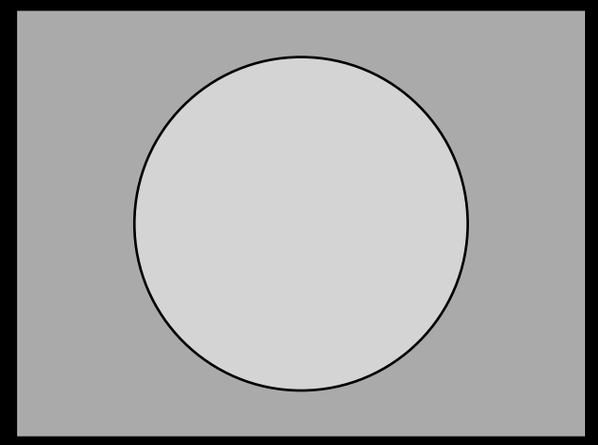

Too small (Inefficient)

LadyBug-3 camera

- 6 cameras, 5 on the horizontal and one pointing upwards.
- Each camera is 1600x1200 pixels.
- Native resolution: 5400x2700 pixels.
- Frame rate: 16fps at full resolution.
  (LadyBug-2 was lower resolution but 30fps).
- Captures a full 360 degree field of view horizontally and about 150 degree field of view vertically.

LadyBug-3 example

LadyBug-3 stitching errors

- The cameras are close together but their nodal points are not coincident.
- This results in a stitching error.
- Possible to stitch at one radius perfectly but stitching error increases as objects are more distant to this chosen radius.
- A common problem for multiple camera rigs.

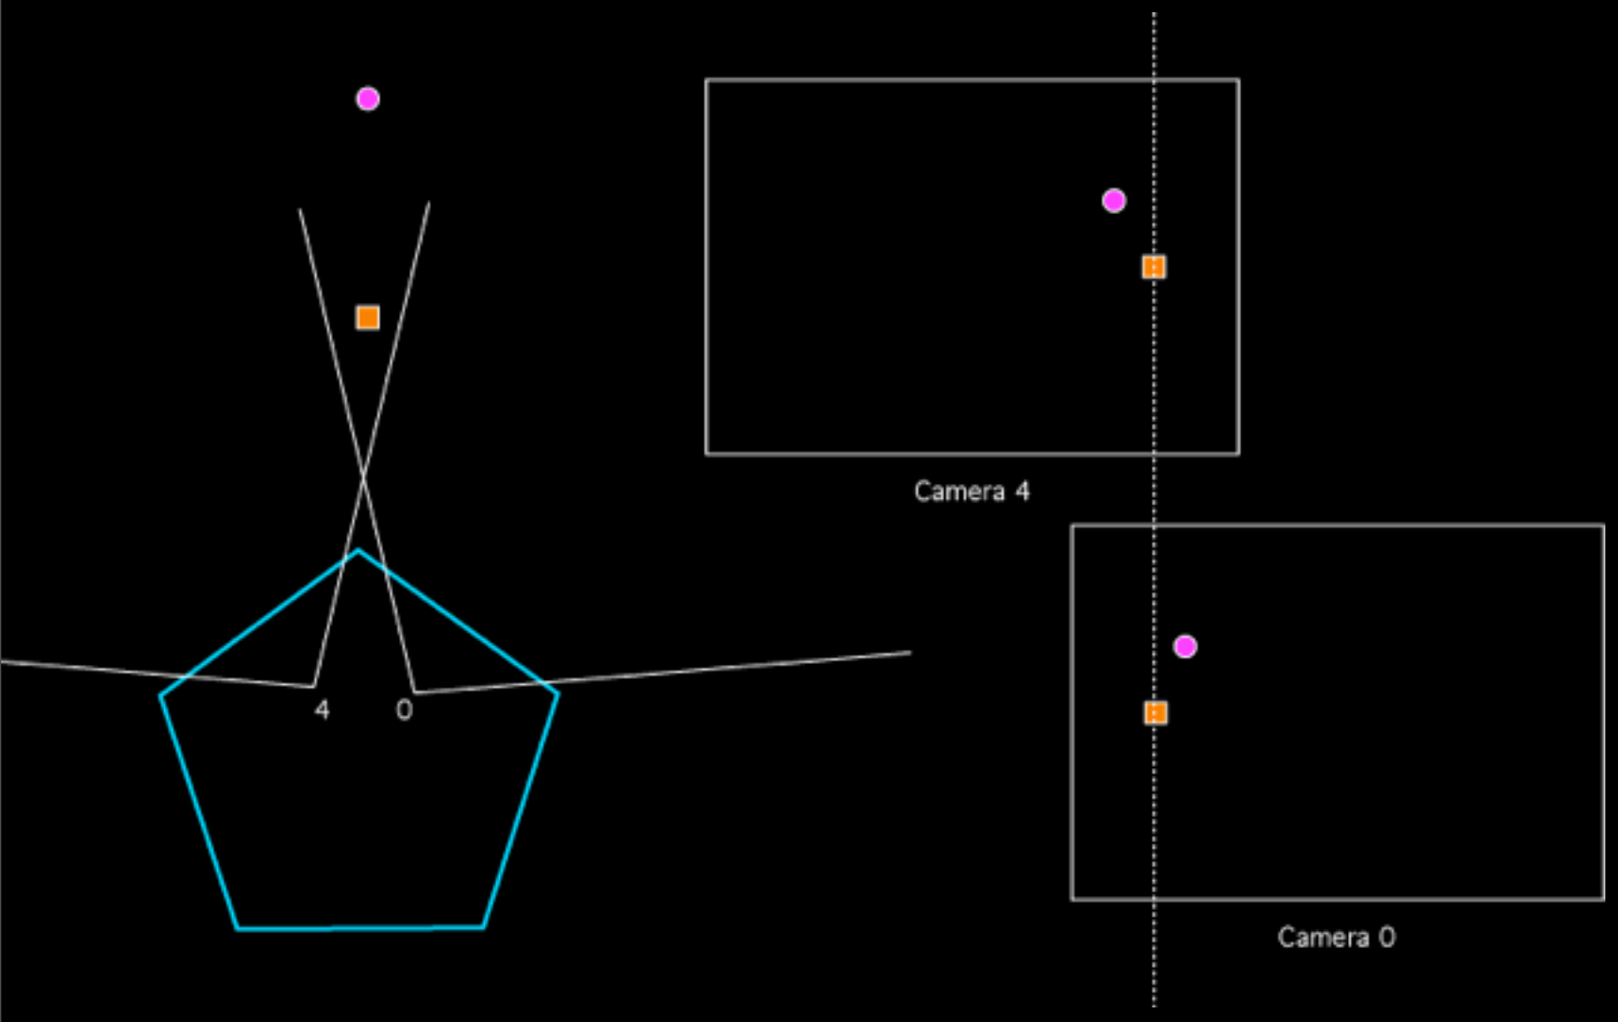

Micoy

iCinema

LadyBug-3: Creating cylindrical and fisheye video

- Since the entire scene is captured as a spherical projection any other projection can be derived.
- In particular cylindrical and fisheye projections, but also perspective.
- An infinite number of fisheye projections are possible, similarly any number of vertical FOV cylindrical panoramas.

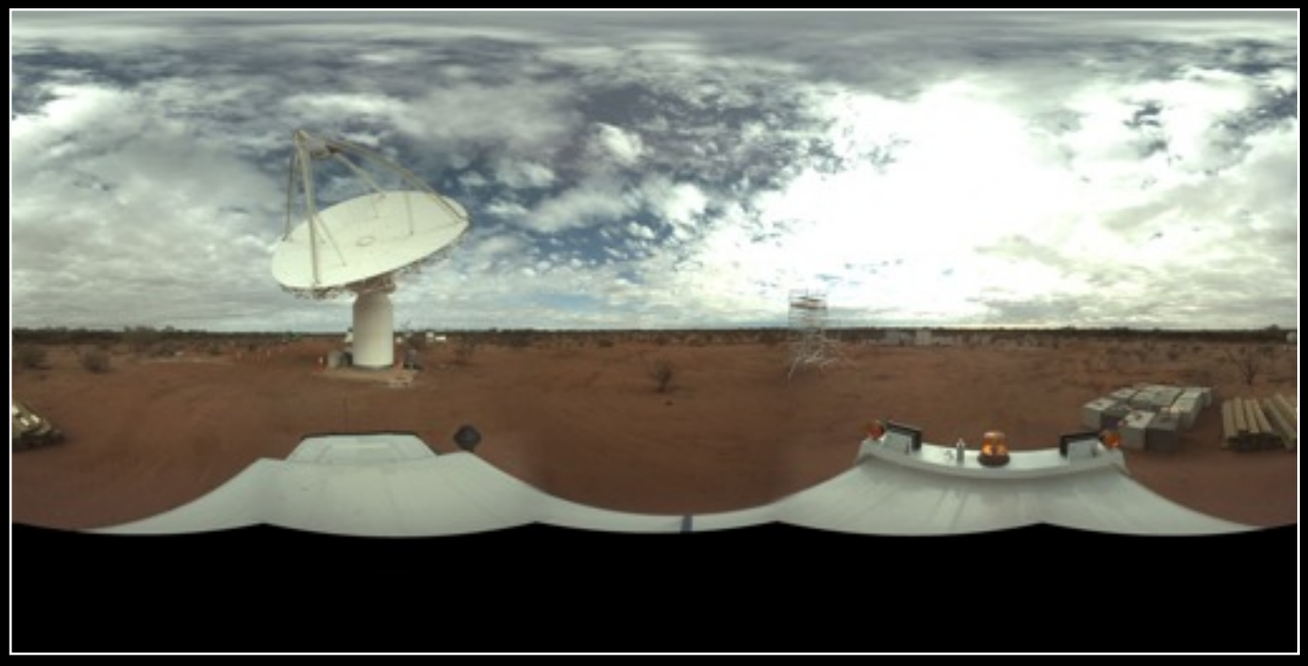

Spherical

Fisheye

Cylindrical

Computer generated

Realtime

- Realtime APIs such as OpenGL and DirectX only support orthographic and standard perspective projections.

- Generally good support for offaxis frustums, this is a reflection of the use of these APIs (especially OpenGL) historically for multiple wall stereoscopic virtual environments.

- Other (cylindrical, spherical, and fisheye) projection types can be achieved with either
  - multiple rendering passes to capture the required field of view,
    these views are then combined to form the desired projection.
  - use of vertex/geometry shaders.

Precomputed

- Support for cylindrical, spherical, and fisheye is becoming increasingly common as camera plugins for commercial rendering packages. Particularly simple to implement for raytracers.

- Offaxis frustums for multiple wall stereo are still relatively uncommon.

Cubic maps

- Create spherical projections from 6 standard square 90 degree FOV perspective projections.
- Create cylindrical projections from 4 standard square 90 degree FOV perspective projections.
- Create fisheye projections from 4 standard square 90 degree FOV perspective projections.

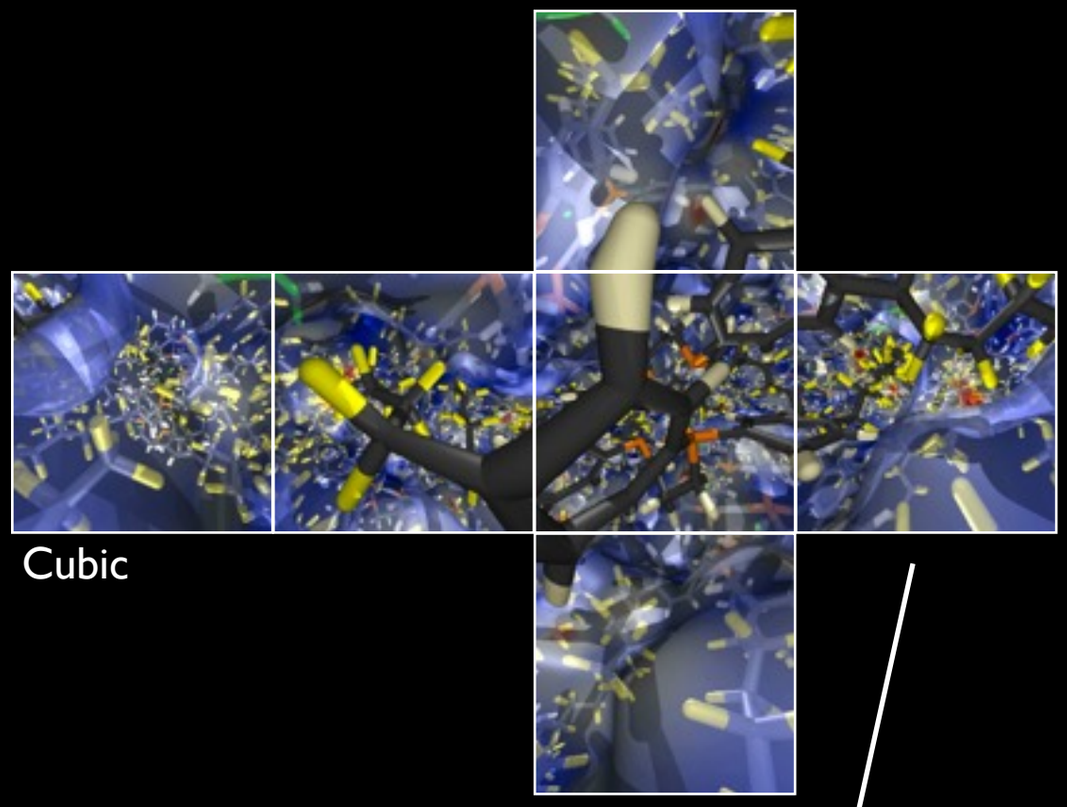

Cubic

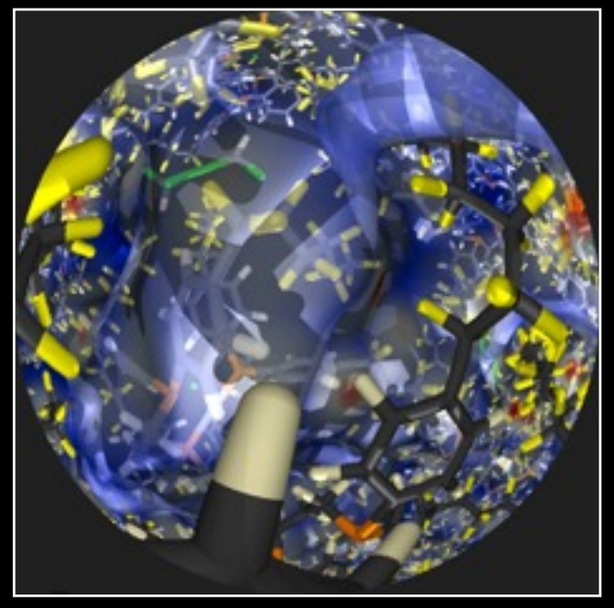

Fisheye

Cylindrical

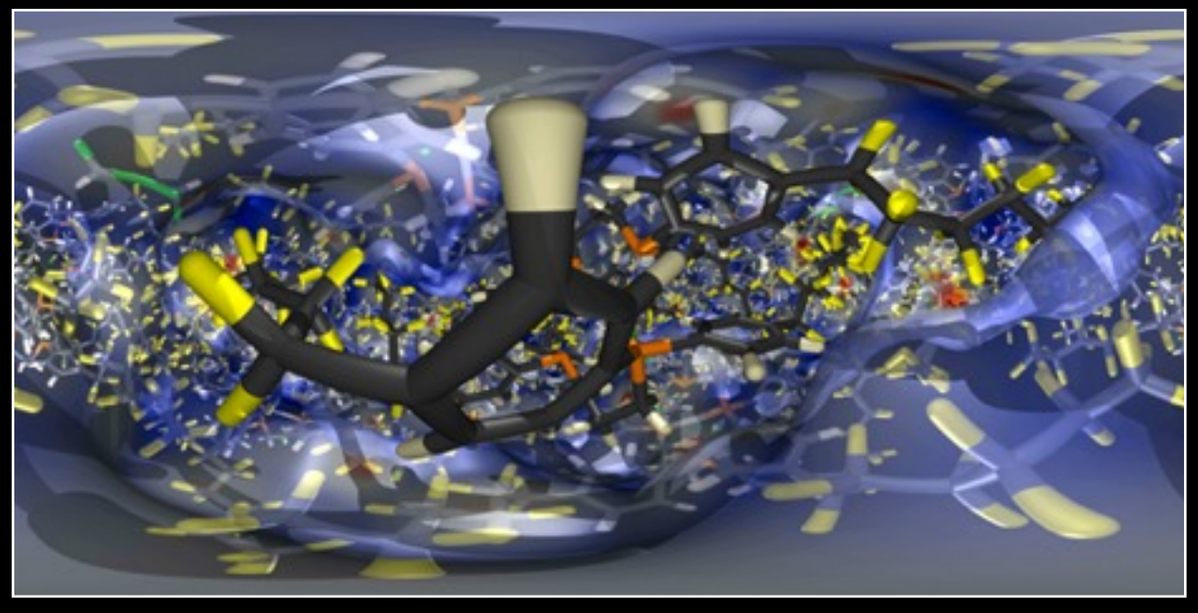

Spherical

Realtime: Multipass textures for fisheye projections

- Same idea as the previously discussed cubic map rendering but performed in realtime.
- Render to offscreen texture and apply texture to a mesh whose vertex positions and/or texture coordinates form the fisheye.
- Require 4 faces of a cube to create a full fisheye.
- Only use half the image from the top and bottom faces.

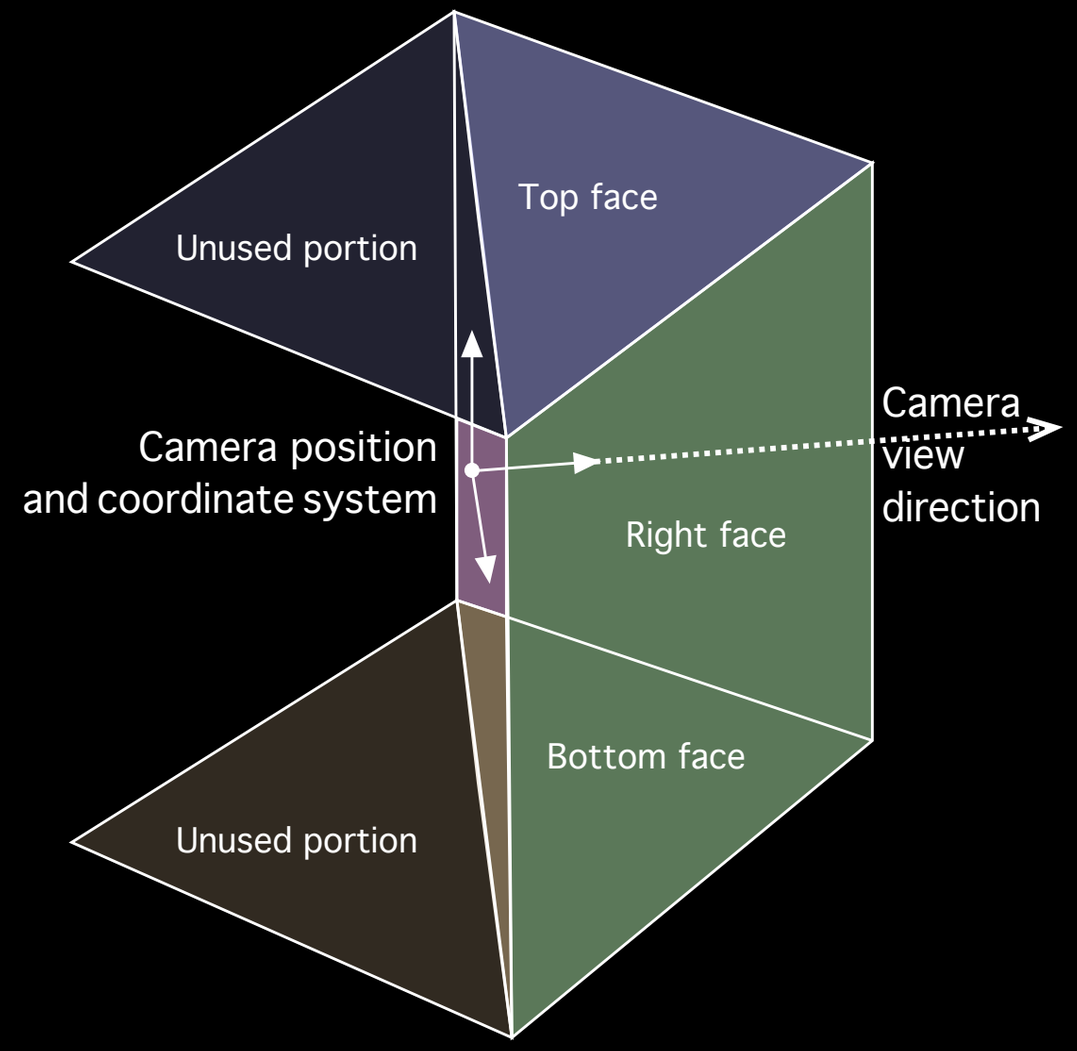

Pipeline for multiple pass texture approach

- Usually easy to integrate into existing rendering pipeline.
- Important to create offscreen textures at an appropriate resolution to avoid aliasing effects.

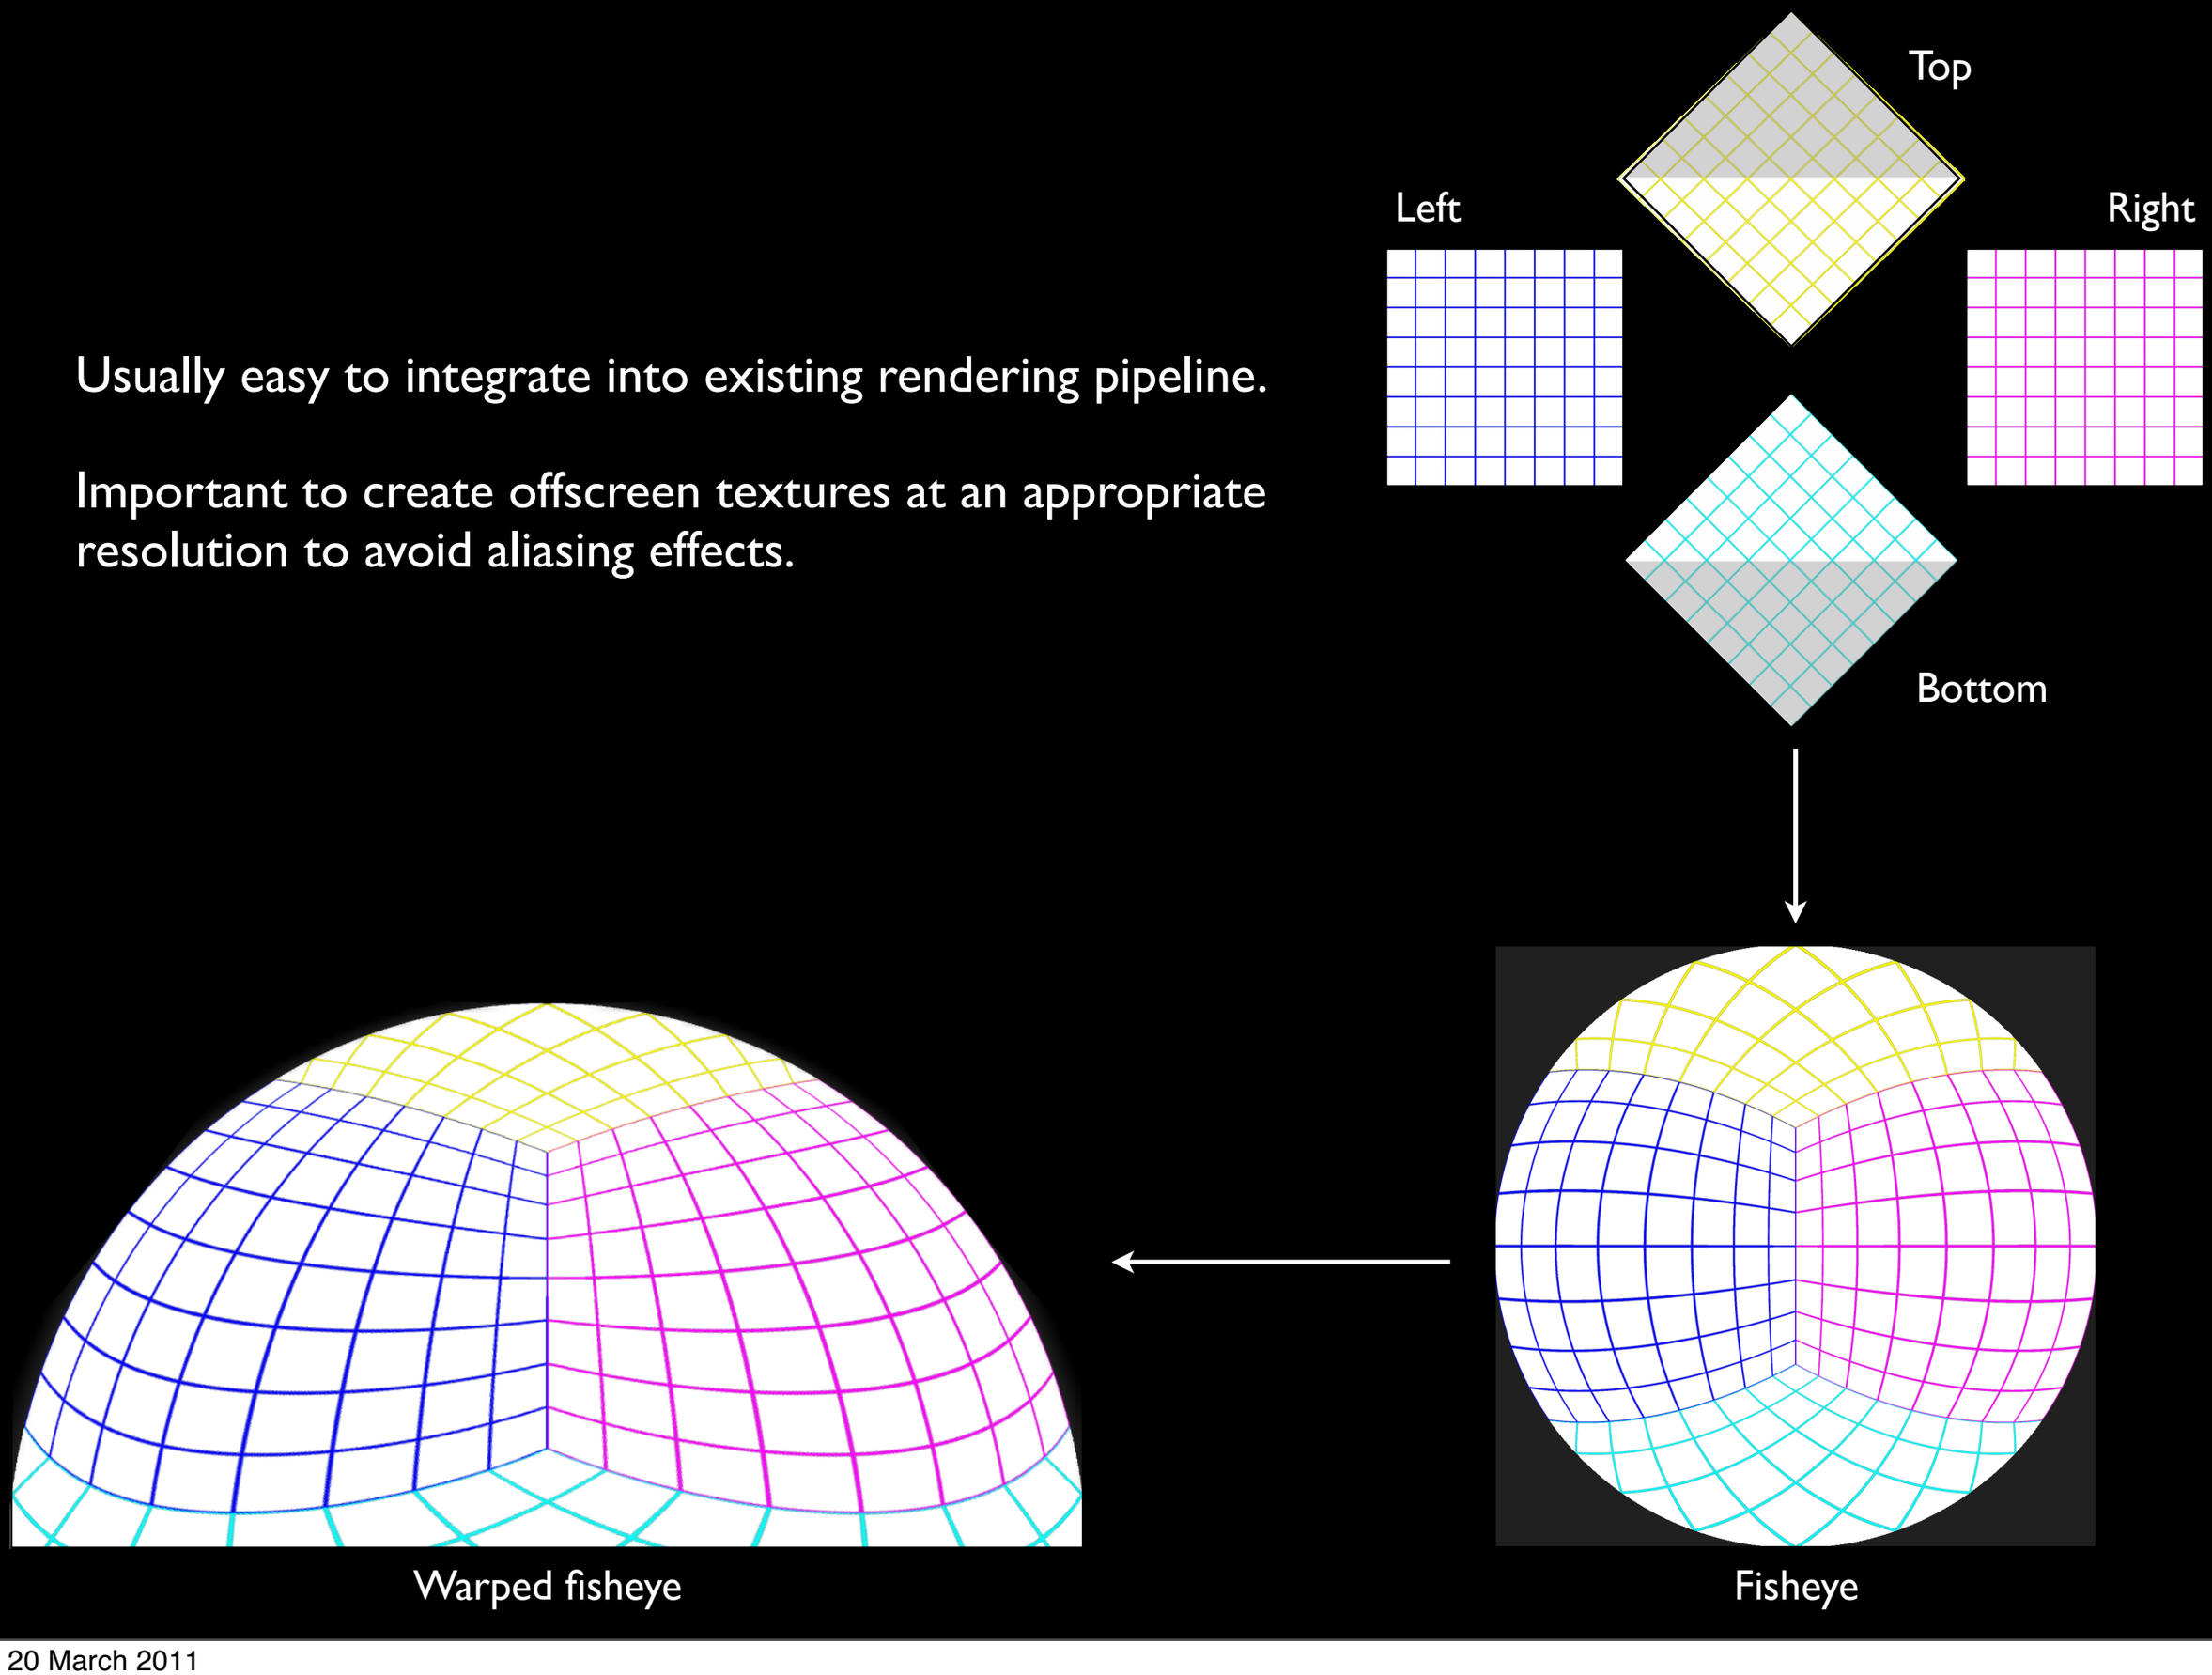

Example: Unity3D Game Engine

- Four render passes involved. Performance depends on the application, generally a factor of 2.5 penalty.
- Final texture phase to fisheye or warped fisheye. Less than 1fps performance hit on this pass.
- Limitation: billboard textures on the ground have normals aligned with camera view direction.
- Limitation: 2D effects such as a halos/flares on sun, they are applied only to the face containing the sun.

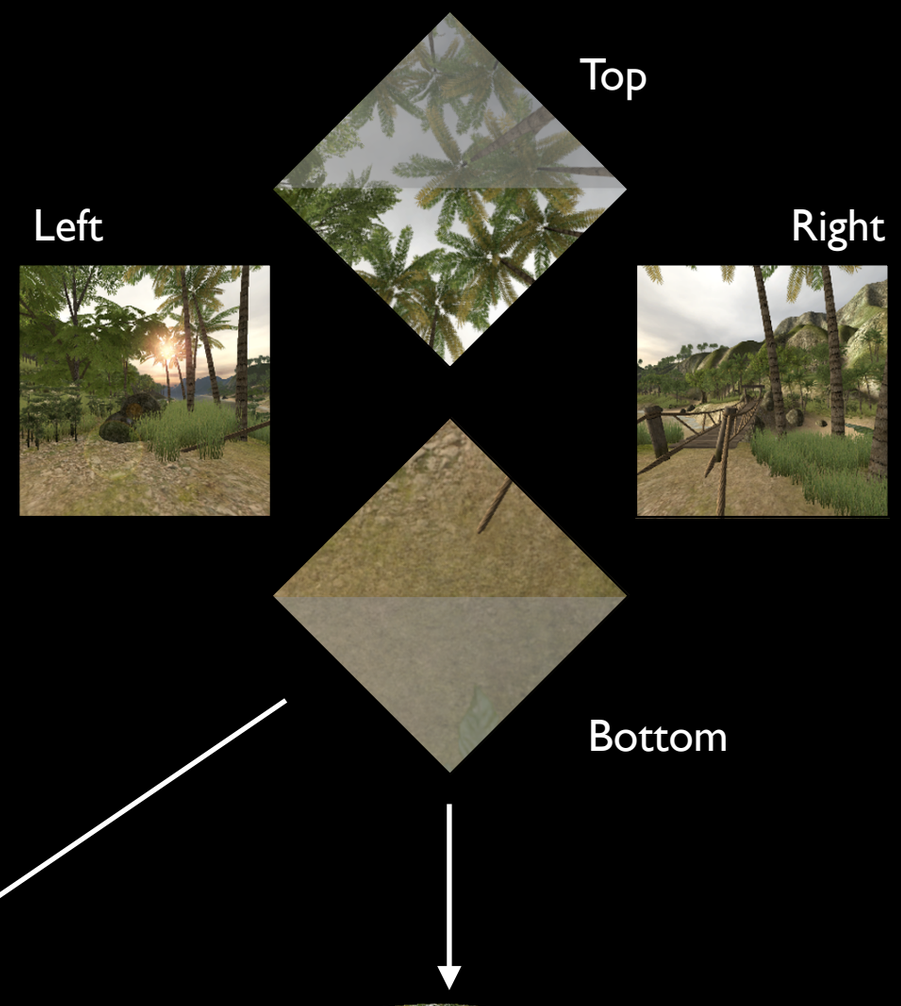

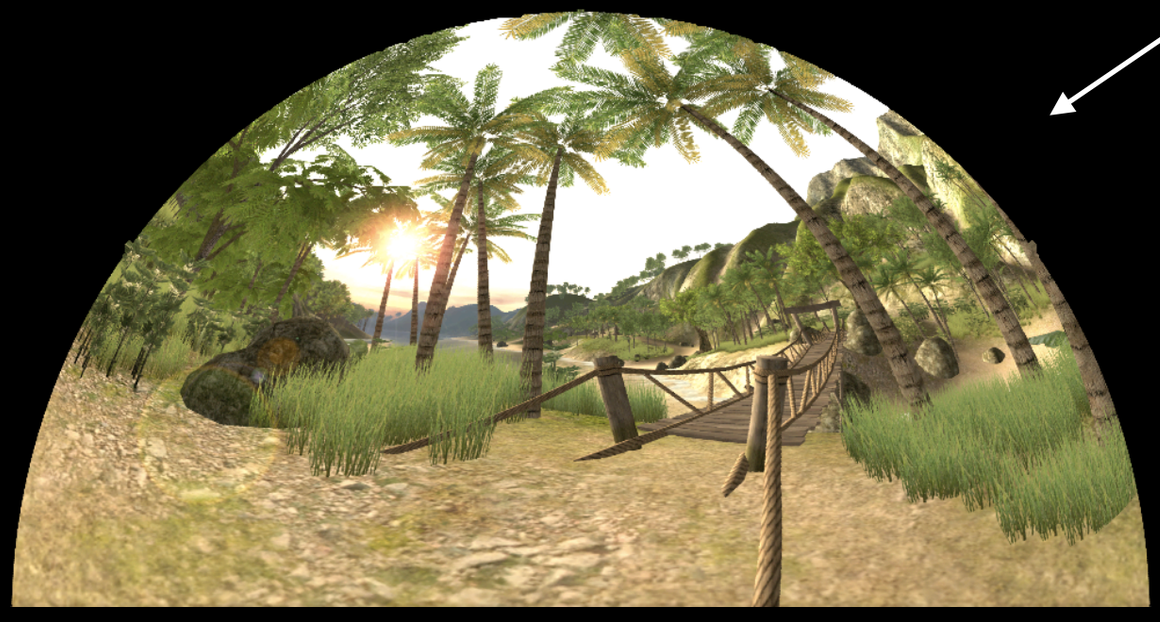

Warped fisheye

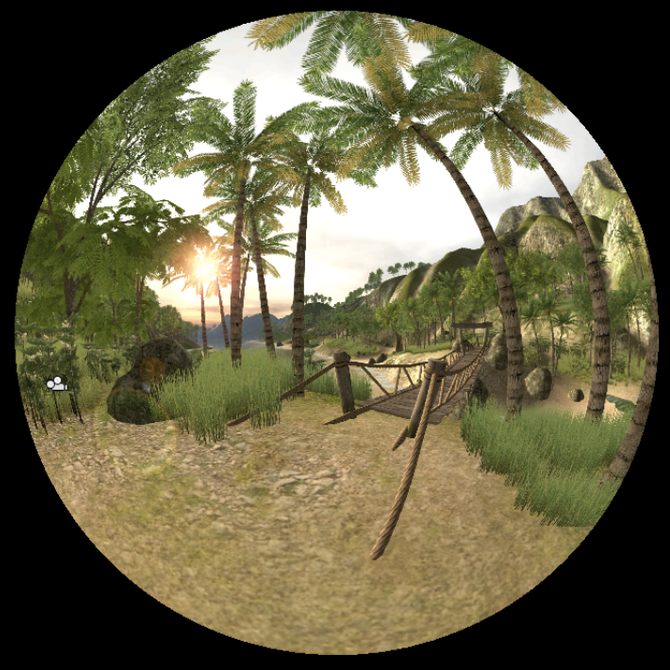

Fisheye

Example: Blender Game Engine

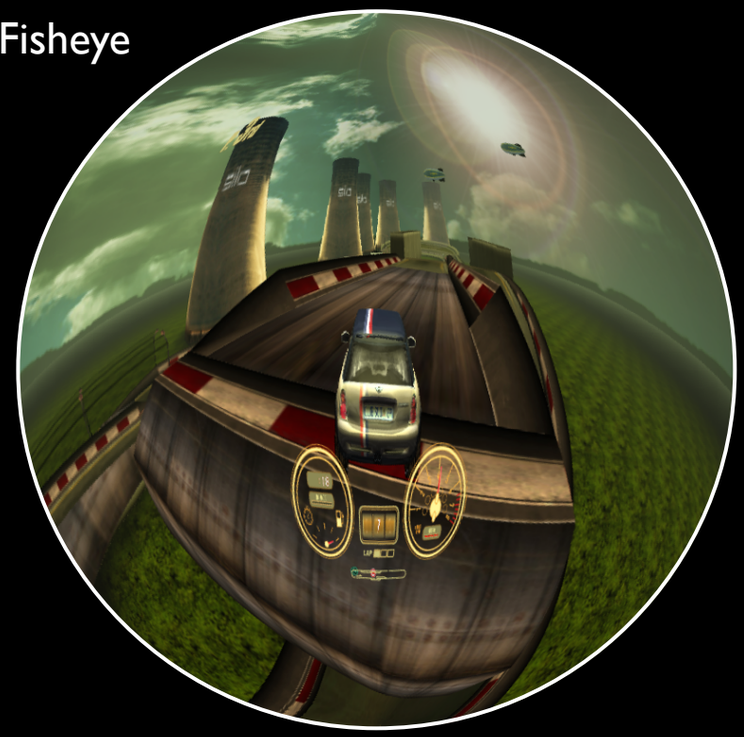

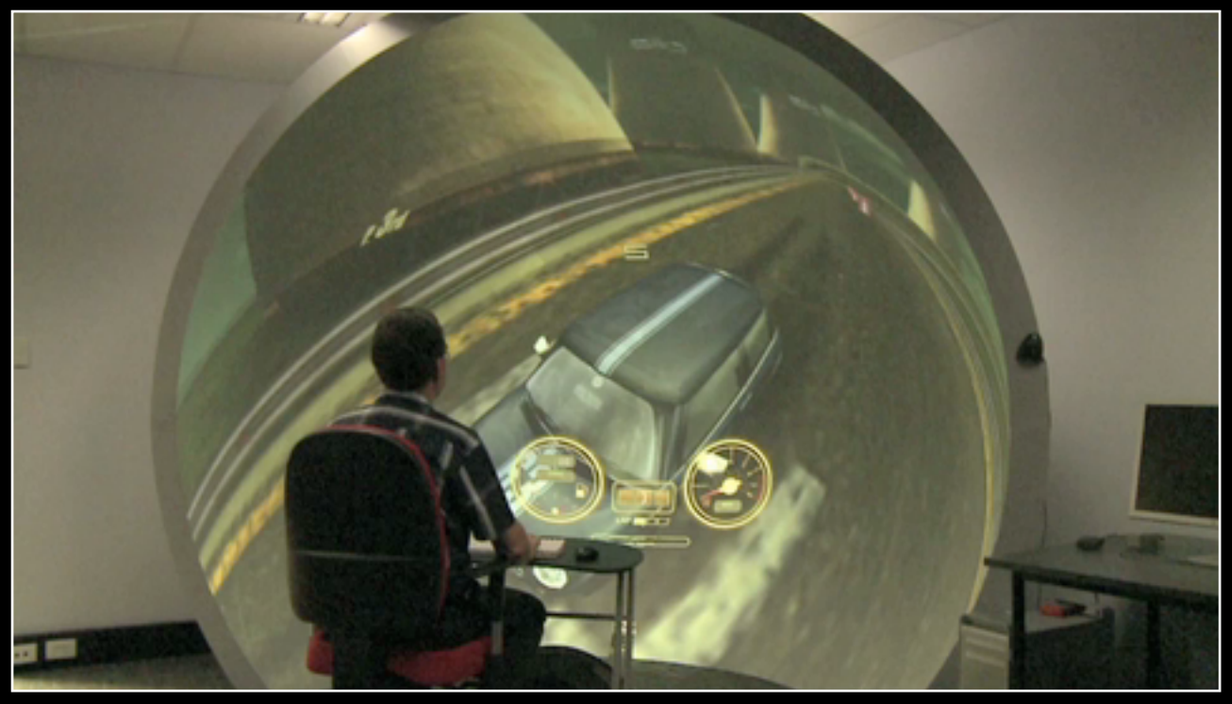

iDome

Alternative: Vertex shader

- An alternative is a single pass render where each vertex is transformed in just the right way so that the resulting orthographic projection is a fisheye.
- Problem: a “straight line” in a standard perspective projection is not a straight line in a fisheye, cylindrical, or spherical projection.
- Solution is to tessellate all geometry so the vertices of the tessellated lines with be transformed correctly. Results in a higher geometric load and the optimal solution for when to tessellate and by how much is not obvious.

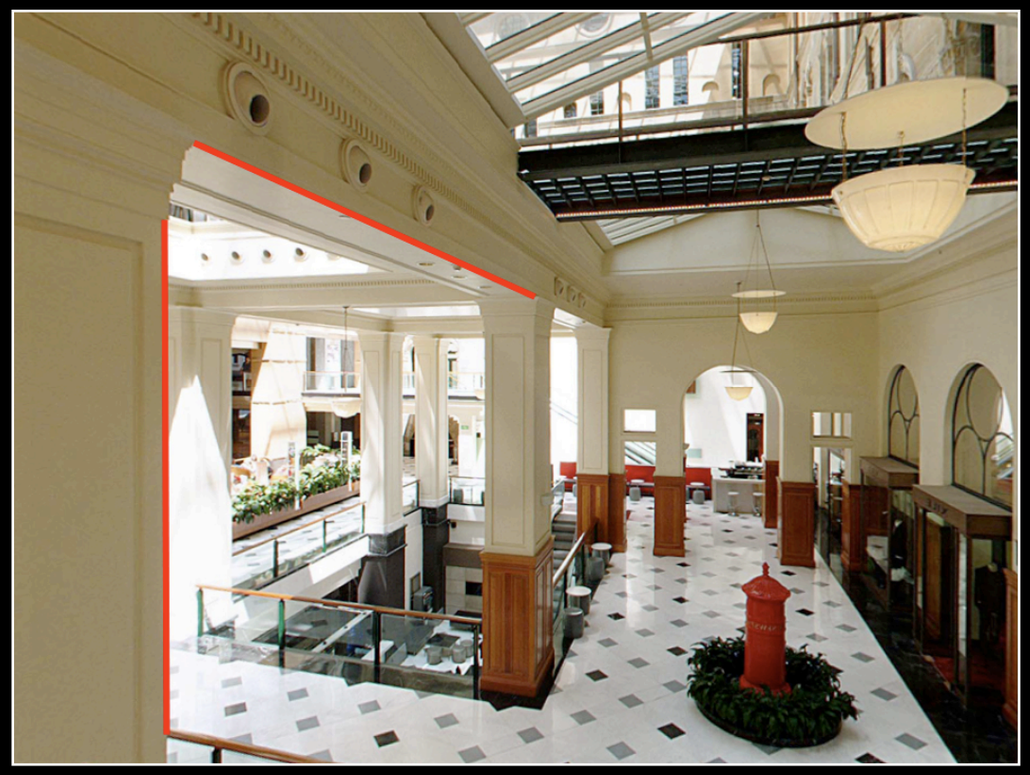

Perspective

Fisheye

Projects

- Remote operations in the mining industry.

  Streaming LadyBug-3 live to a iDome console.

- Science education.

  Interactive exploration of spherical projections, LadyBug-3 interviews and CG.

- High resolution digital fisheye.

  Spherical projections from Mawsons hut, Antarctica.

iDome: Remote operations in the mining industry

- Realtime use of the LadyBug-3 footage.
- Aim is to remove an iron ore ship loader operator from a noisy, dusty, hot, vibrating rig.
- Key requirement is they need to be able to survey the whole scene.

Operator

LadyBug-3 example footage

- Camera is mounted upsidedown, need to see into the ship hold and don't care about the sky.
- Footage streamed 2km to remote control room, operator pans the footage in realtime.
- Can actually see more than they could from the original platform of the shiploader.
- Main reported problem was the absence of depth perception.

iDome: the operators interface

Laboratory testing

iDome: Science education and visualisation

- Aim: Present the research at the ARC Centre of Excellence for Electromaterials Science.
- Presented within the Wollongong Science Centre at Wollongong University, adjacent to the Centre for Electromaterials Science.
- Themes selected from research in the Centre, see the posters below.

Software

- Participant can choose from one of eight stories from the research laboratory.
  Conducting polymers - Fibre Spinning - Looking at the Nanoworld - Nanostructured Electromaterials - Carbon Fibres - Graphene - Welcome and Invitation.

- Movie selection and navigation using Quartz Composer.

- Visualisation components rendered as spherical projections and composited into LadyBug-3 footage.

- Interestingly there are no perspective matching required when compositing spherical projections.

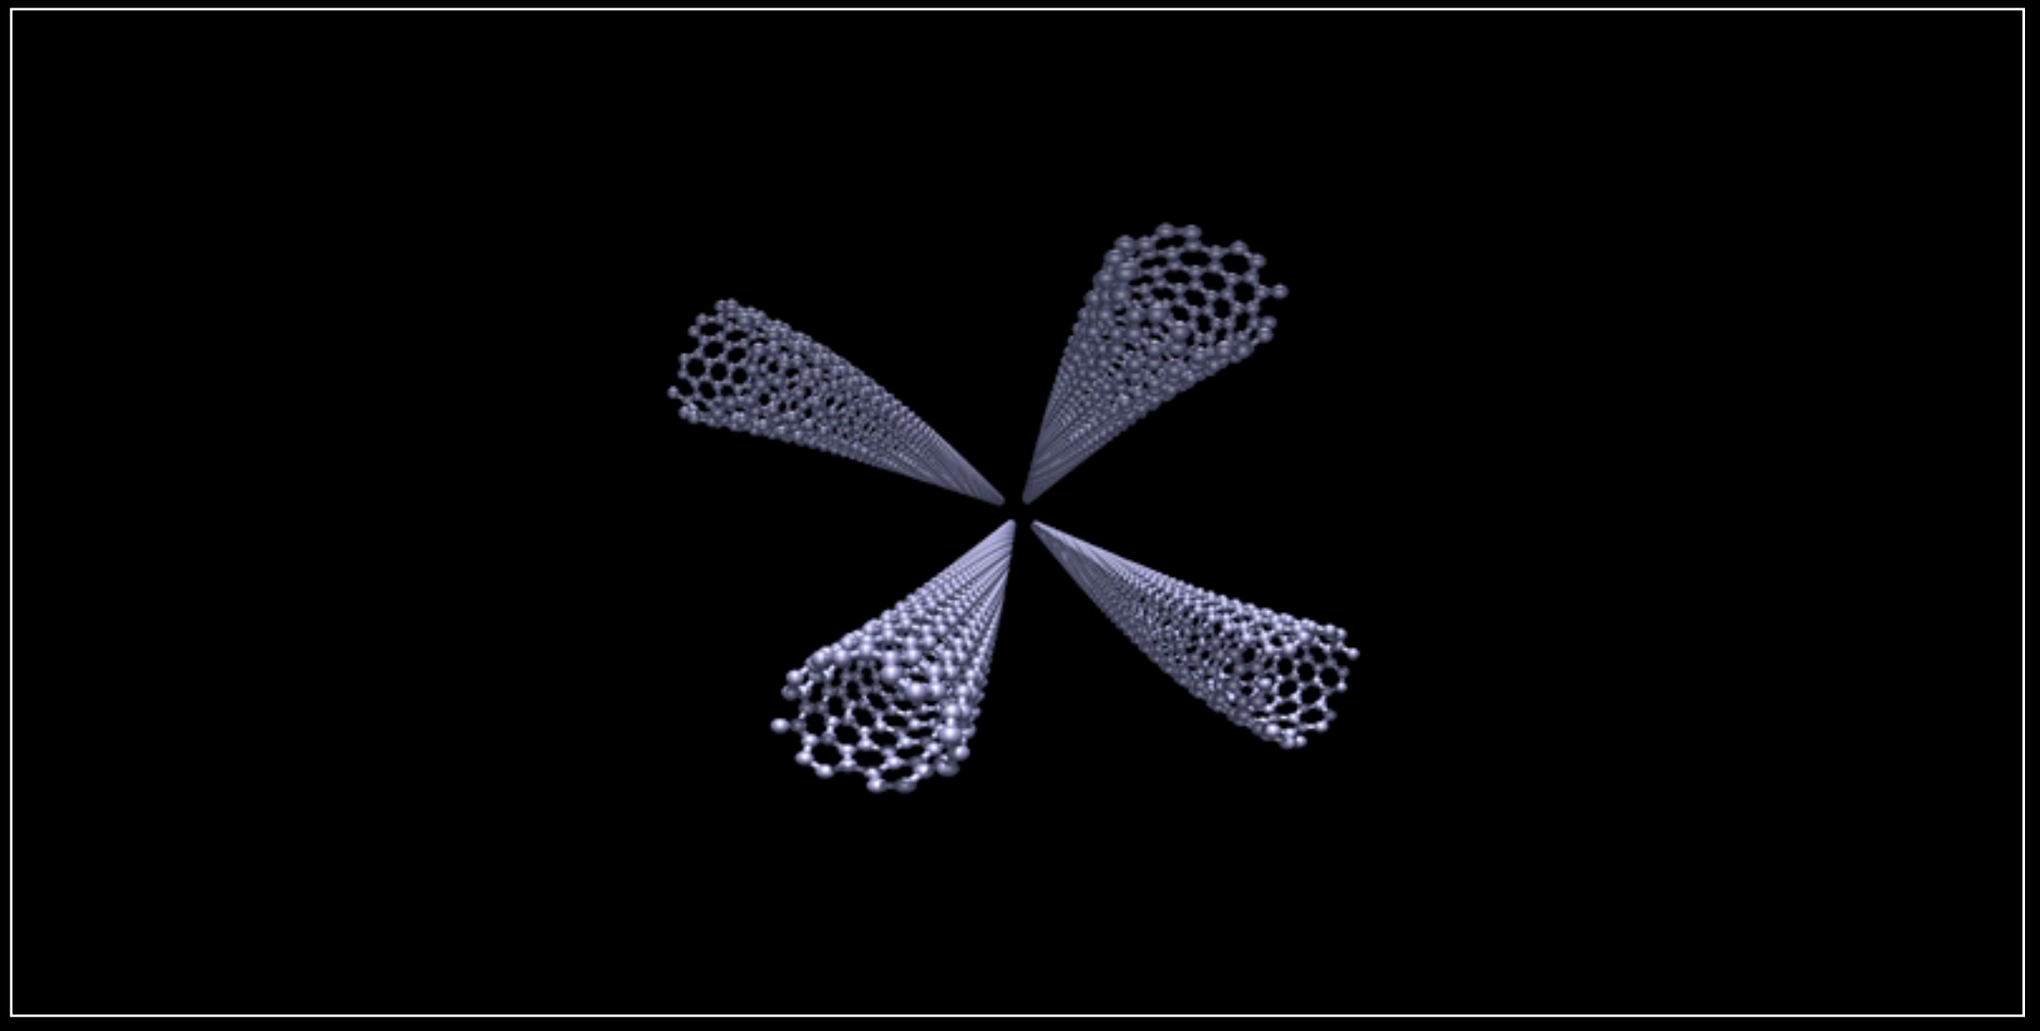

Example footage from the LadyBug-3

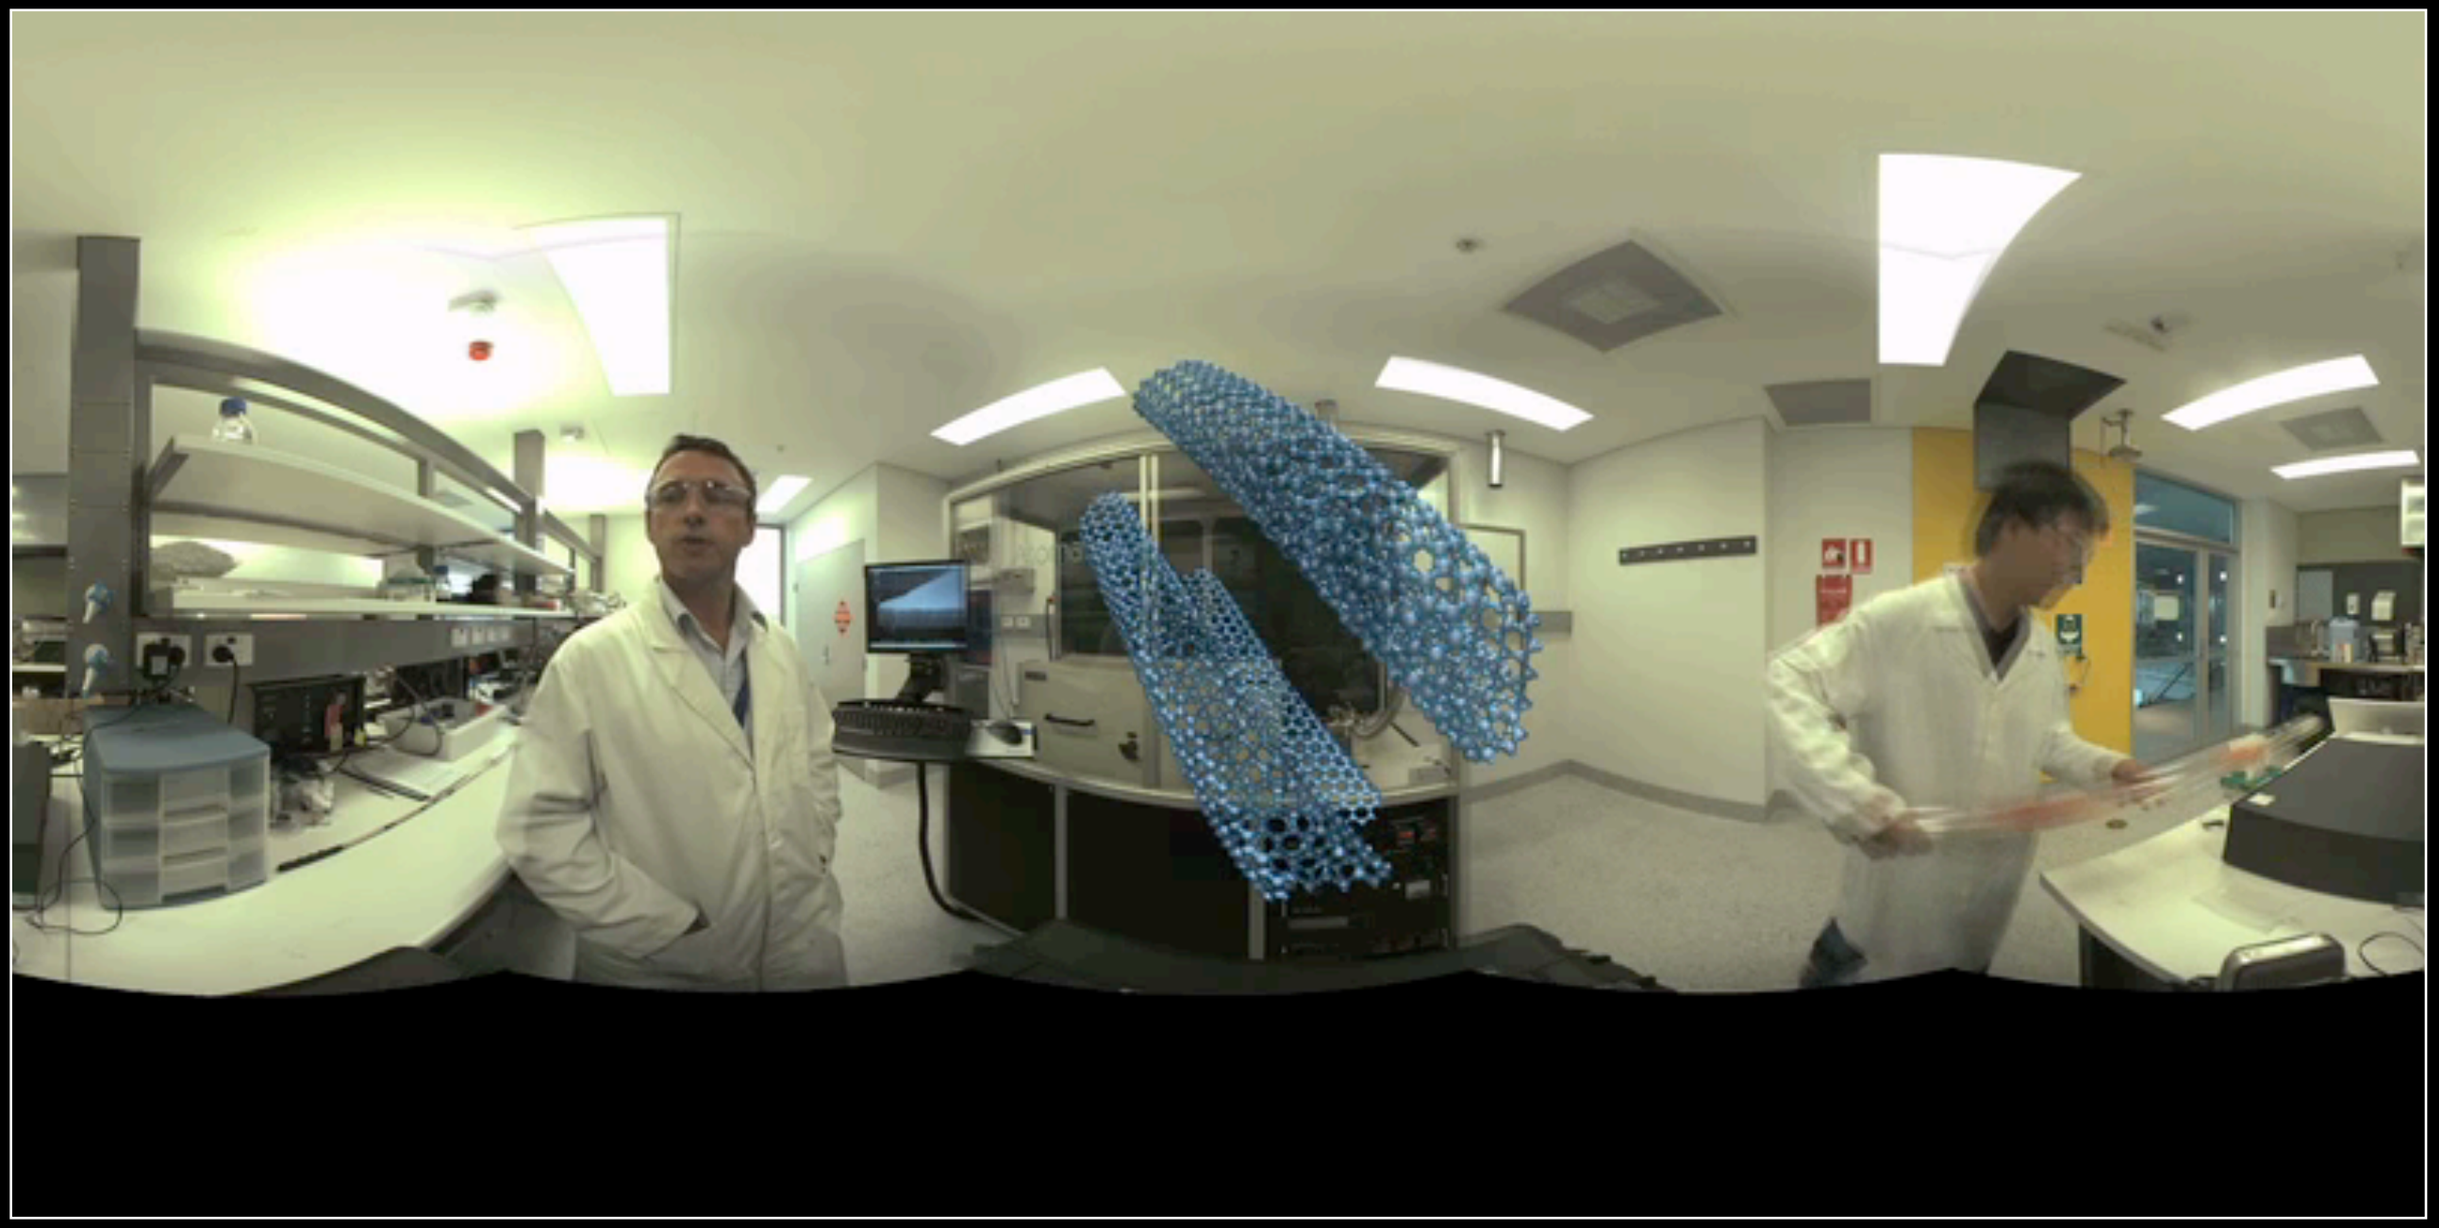

iDome as the science centre visitor interface

Planetarium: Digital capture for Antarctic heritage

- Project by Peter Morse.
- Heritage project based upon Mawsons huts, Antarctica.
- Access is difficult so required research useful images, "hundreds" of images captured at strategic nodes within the hut.

Outside the hut

Fulldome production: 4K resolution

- Initial production was created at 4K square fisheye sequence, 30fps.
- Panning around within the high resolution spherical panoramas.

Spherical 40,000 x 20,000 pixels

Fisheye 4096 x 4096 pixels

8K production

- A few planetariums in the world are now able to faithfully represent 8K x 8K fisheye movies.
- Frozen in time re-rendered at 8K, the underlying spherical panoramas contained the visual fidelity and resolution for this.
- Becomes an exercise in file management and high speed access.
- 8K x 8K is 32 times the number of pixels as 1920x1080 HD.

8192 x 8192 pixels @ 30fps

iDome interface (Prototype)

- High resolution panoramas spatially located within a model of Mawsons hut.
- Unity3D engine to provide a realtime user experience.

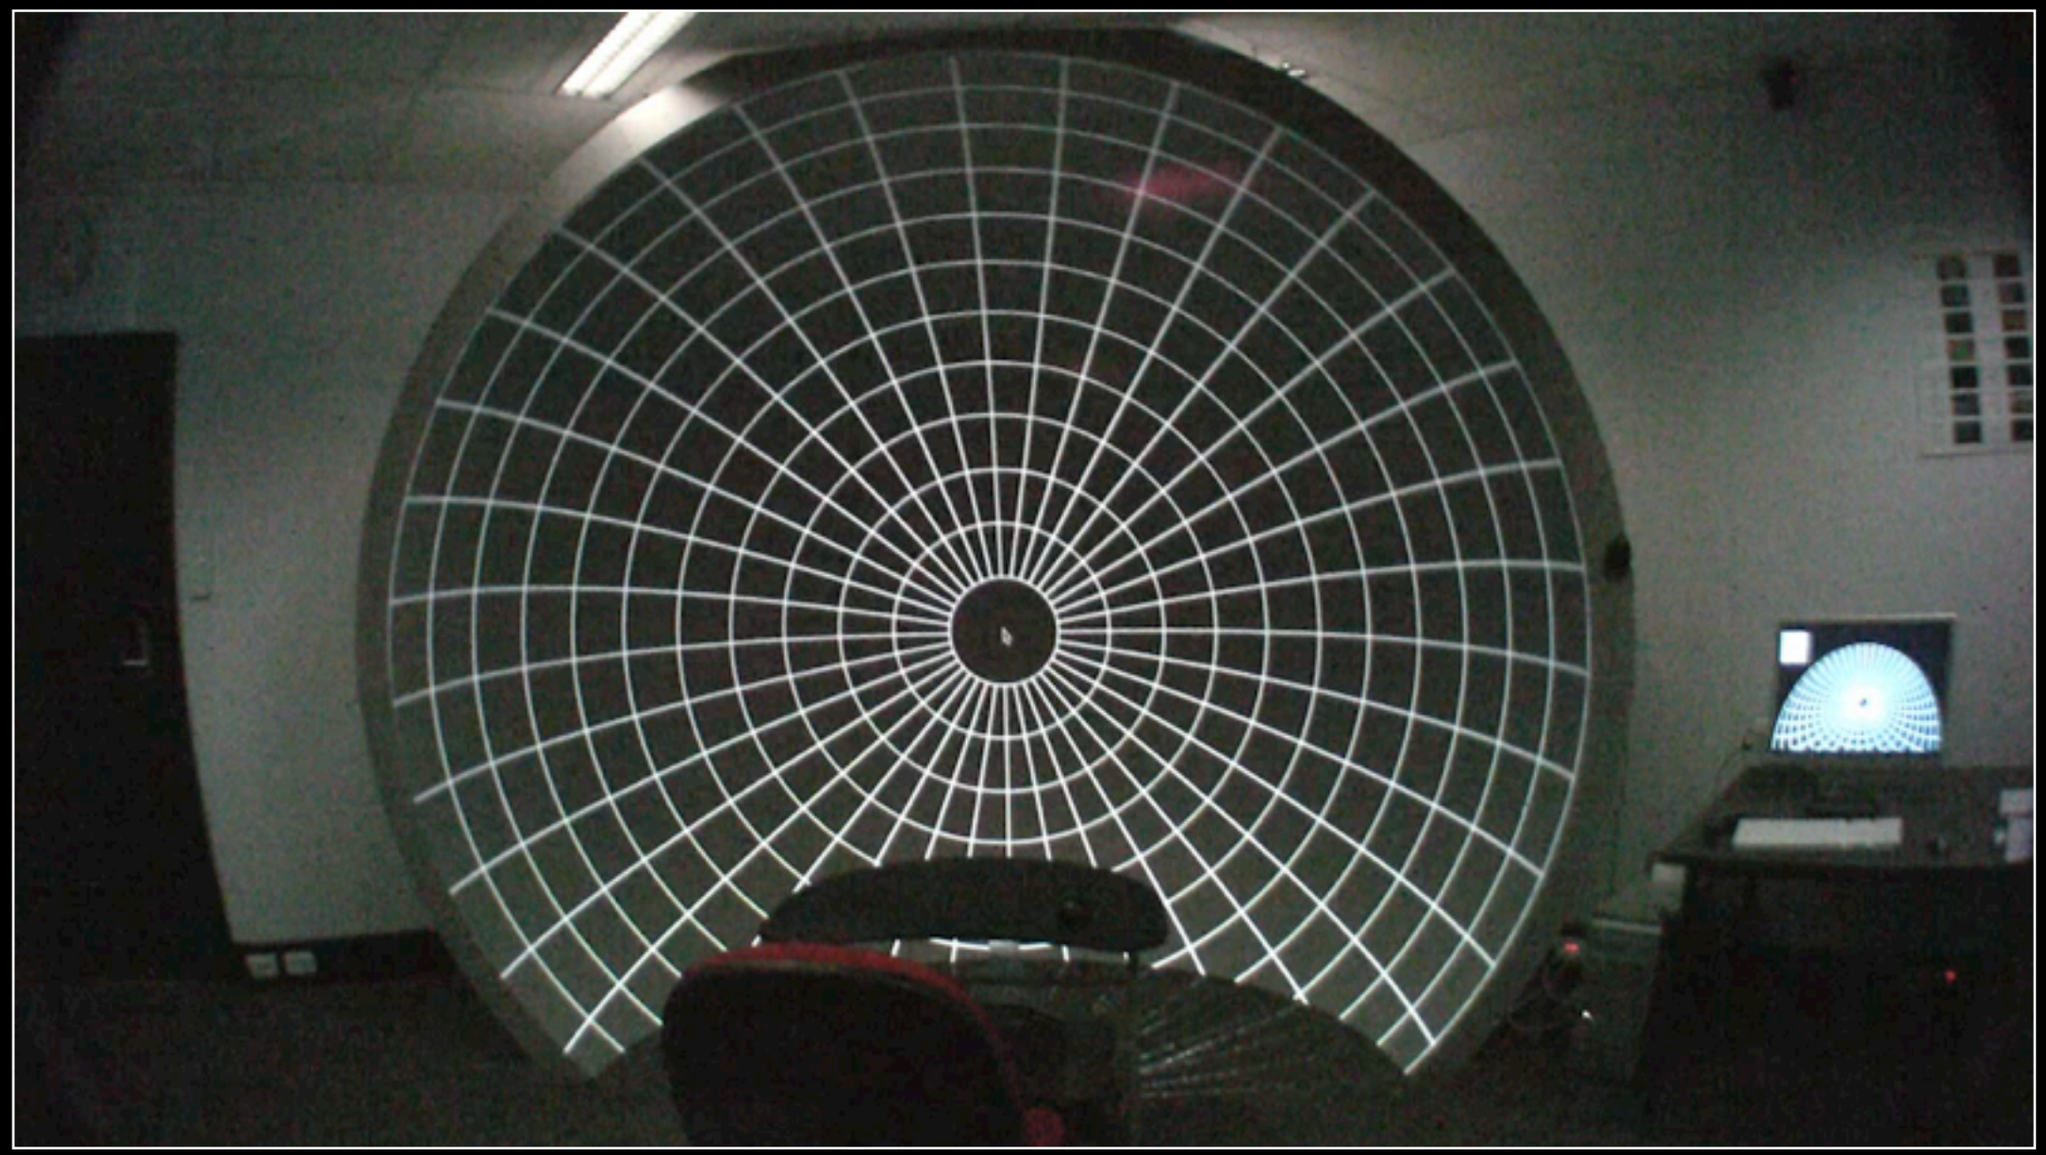

Ice Museum

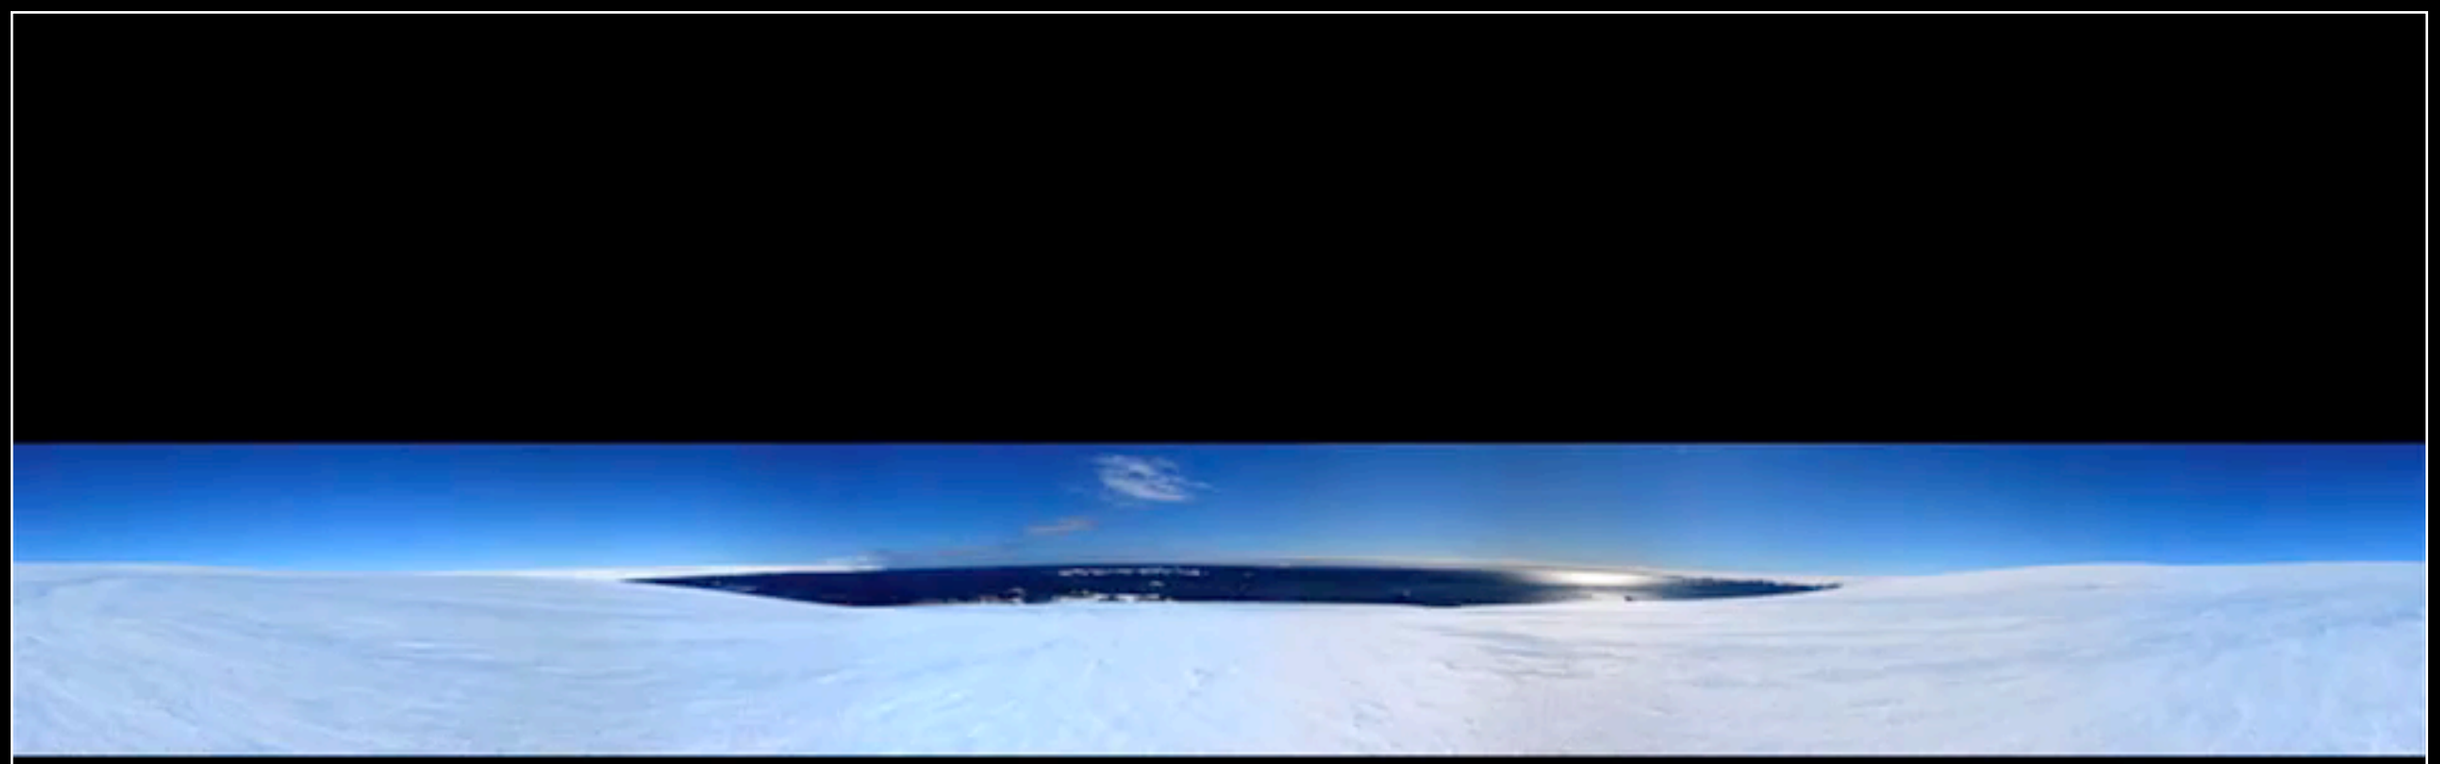

Questions and Discussion