Virtual Reality

VR Displays & Stereo Rendering

G. Zachmann

University of Bremen, Germany

cgvr.cs.uni-bremen.de

The Kinds of VR Displays

- Autostereo Monitor
- *Head-Mounted* Displays (HMDs)
- *Head-Coupled* Displays (HCDs)
- *Immersive projection* displays (IPDs)
  - *"Powerwall"*
  - *Workbench*
  - *Cave*
- "Exotic" displays:
  - Retinal displays
  - Holographic displays
  - ...

Stereo Monitor (mostly Autostereo)

- Sometimes called "*Fishtank VR*"
- Advantages:
  - Inexpensive
  - Resolution up to 1900 x 1600
  - Well accepted by users
  - No special requirements on the environment/setting
  - Some 3D capabilities
- Disadvantages:
  - Small Field-of-View (FoV)
  - No immersion
  - Very limited working volume
  - "*Stereo frame violation*" is very common

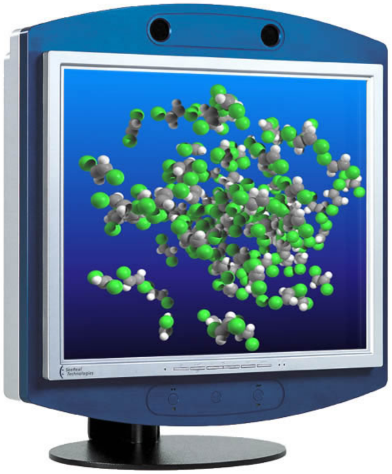

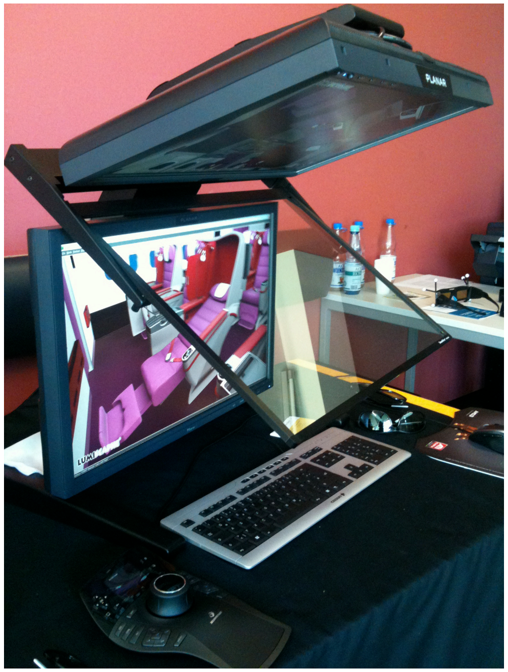

- Interesting things you can do with a simple monitor: the "Reachin-Idea"

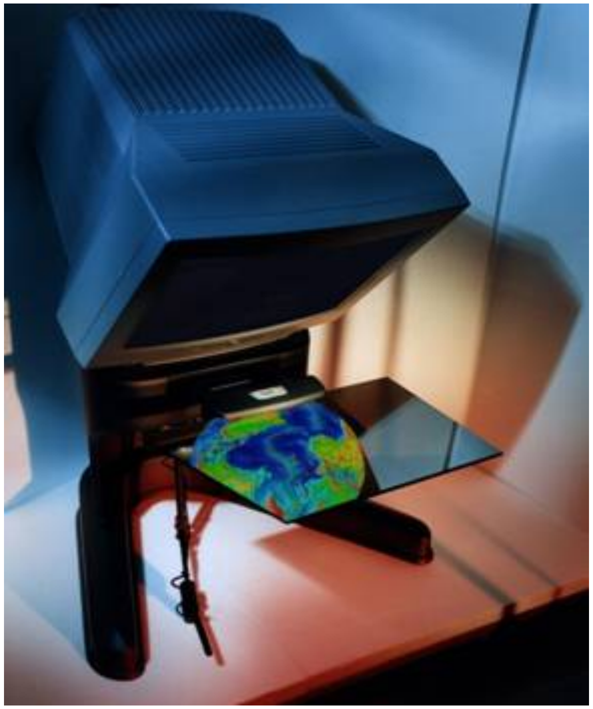

- The problem with a small FoV: there is no immersion!

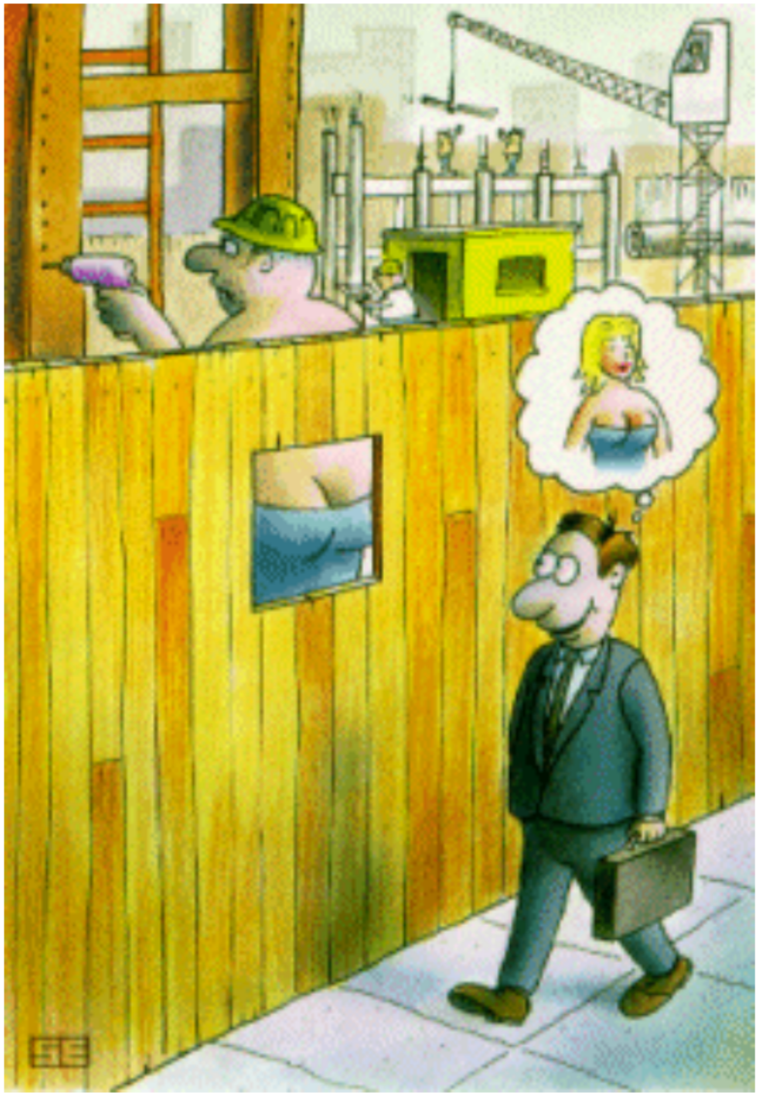

Head-Mounted Diplays (HMD)

- First "true" VR display
- Technologies / characteristics:
  - HMDs using LCDs (sometimes CRTs)
  - Weight:
    - Small FoV → lightweight; large FoV → heavy
- Advantages:
  - Kind of "surround display"
  - Very good immersion
  - No "*stereo frame violation*"
  - Large working volume
  - *Low-end models* are inexpensive
  - Almost no special requirements on the working environment

- Disadvantages:
  - Uncomfortable when used for a prolonged time ("*invasive interface*")
  - Distortions
  - Real environment is shut off (good for immersion, bad for collaboration and self-embodiment)
  - Manipulation of real controls is difficult (e.g., in mockup of cockpit)
  - Every participant needs an HMD (bad: expensive, good: everybody has correct perspective in VE)
- Actually, HMDs hve been invented a long time before "VR"

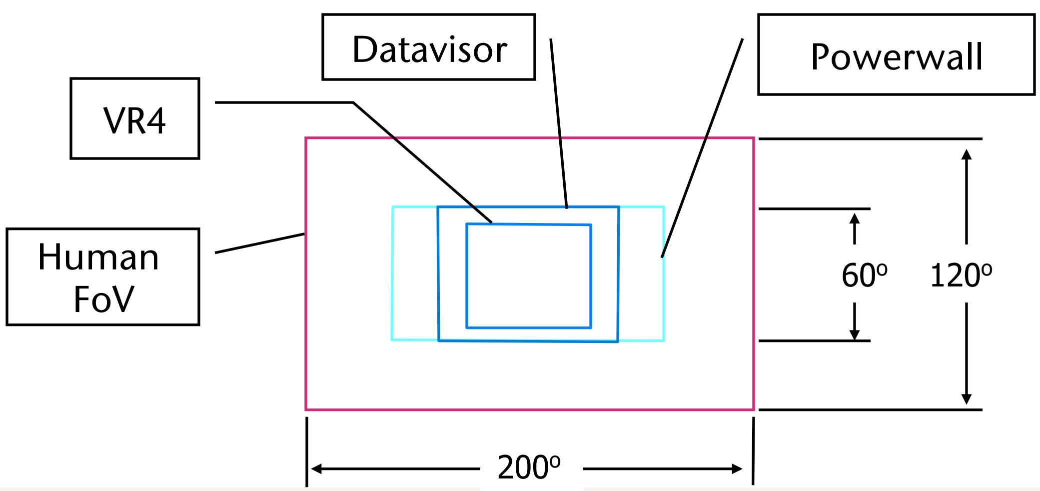

Bell Helicopter, 1967

Head Coupled Displays (HCD)

- HCD = HMD mounted on a "boom"
- Advantage of HCDs over HMDs:
  - Possible to quickly "take the display off" for a moment; or users can just take a "quick peek" into the VE
  - Low weight on the head
  - Extremely good tracking comes built-in
- Disadvantages compared to HMDs:
  - Smaller working volume
  - One hand is always occupied
  - Inertia

→ Failed to gain market share

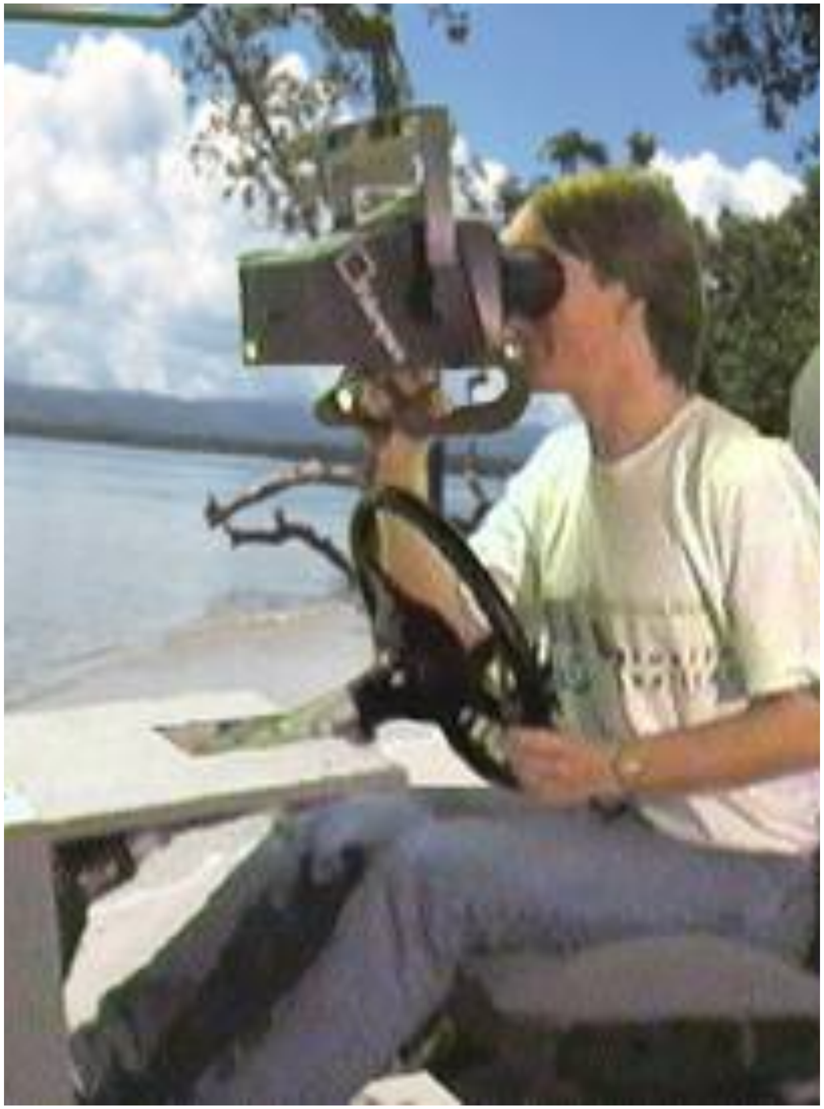

Immersive Projection Displays / Technology (IPD / IPT)

- Idea is (somewhat) similar to cinema theaters
- Setup: 1–6 walls on which VE is projected
- *Powerwall* = 1 wall (e.g., 3x6 meters)
- *Workbench* = 1 horizontal display surface (table)
- *Holobench, L-Shape* = 2 display surfaces, 1 vertical, 1 horizontal
- *Cave* = 3–6 walls

Powerwall

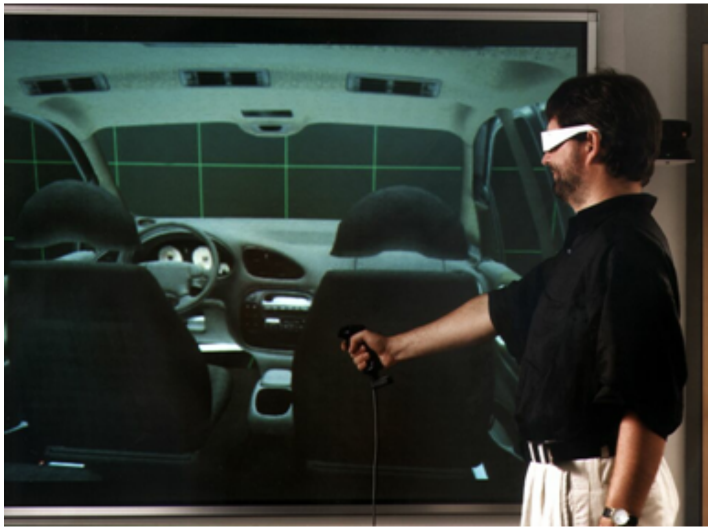

Powerwall with back projection

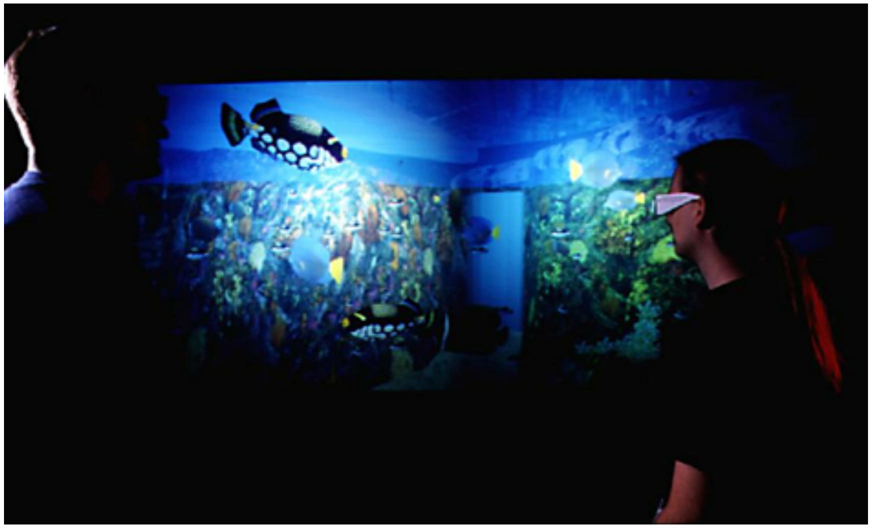

Powerwall with front projection,
(problems with that: edge blending, hot spots)

- "*HeyeWall*", Darmstadt:
  - 24 tiles, 48 PCs
  - Total resolution: 18 Mio pixels (6144 x 3072) in stereo

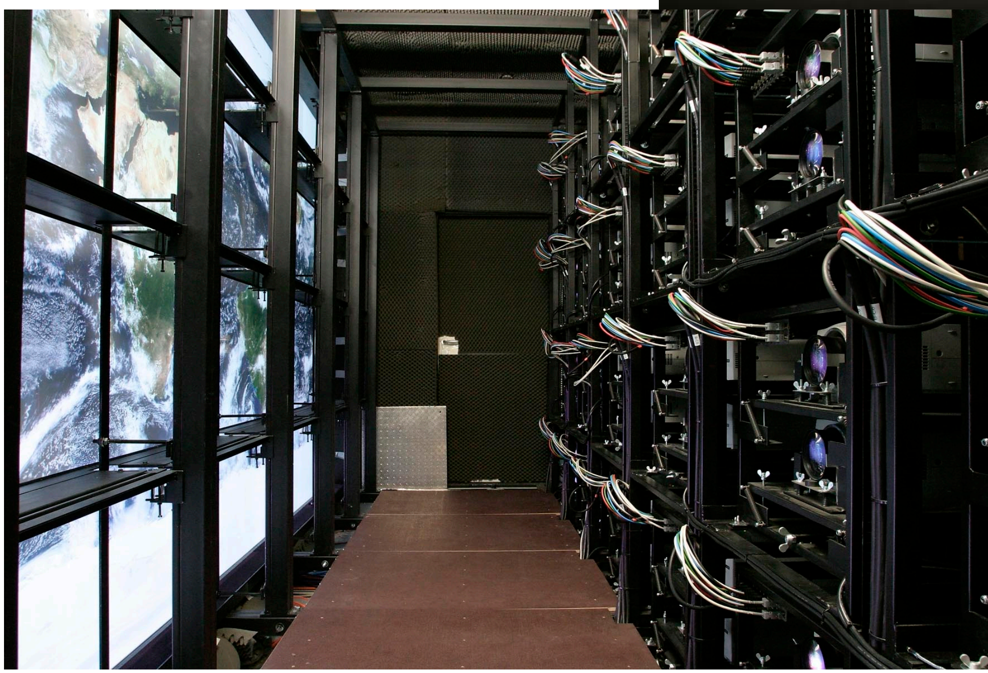

Example Application: Virtual Conference Room

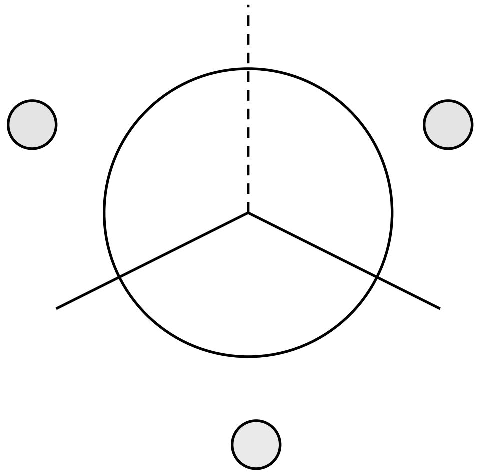

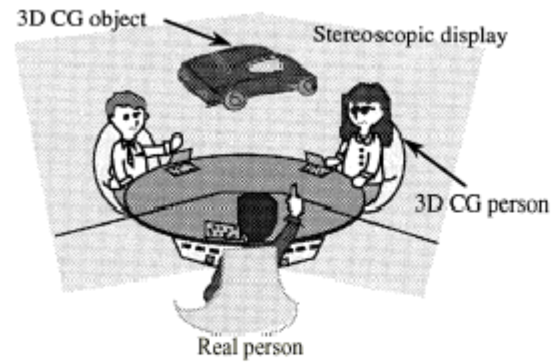

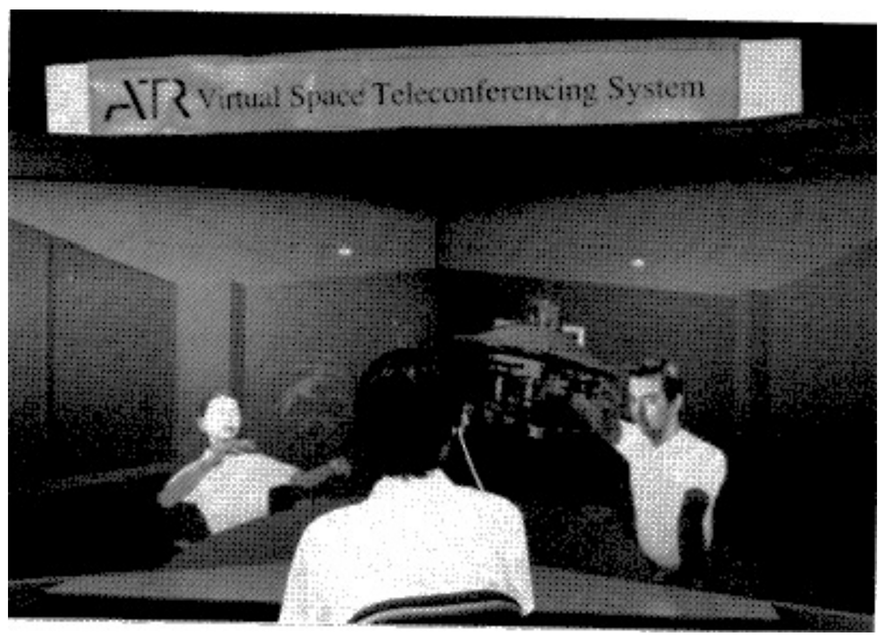

Result: *1 shared workspace*,
by way of coherently adjoining
"desktop IPD"

Workbench, L-Shape, etc.

Workbench

Principle of the workbench

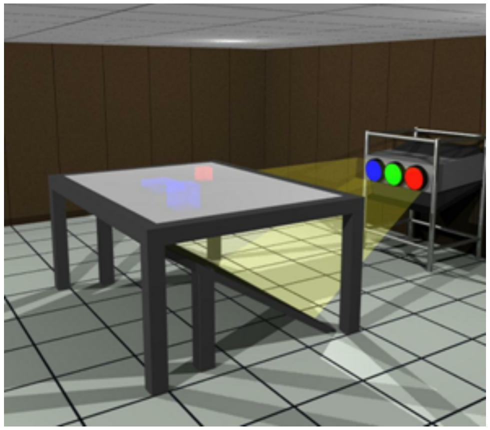

Tilting "workbench"

Holobench

Cave

3-wall cave

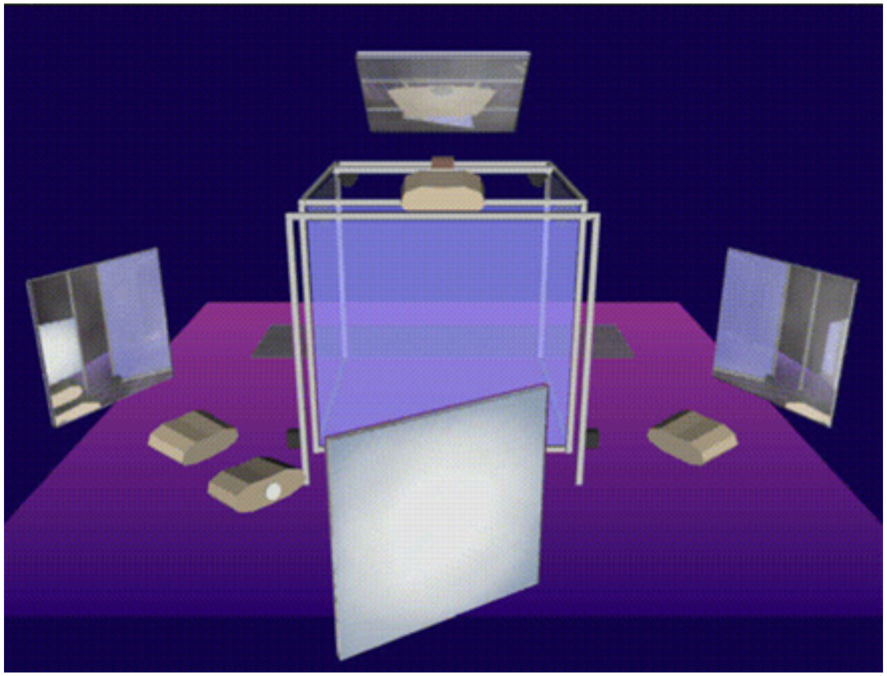

Schematic of the arrangement of the mirrors

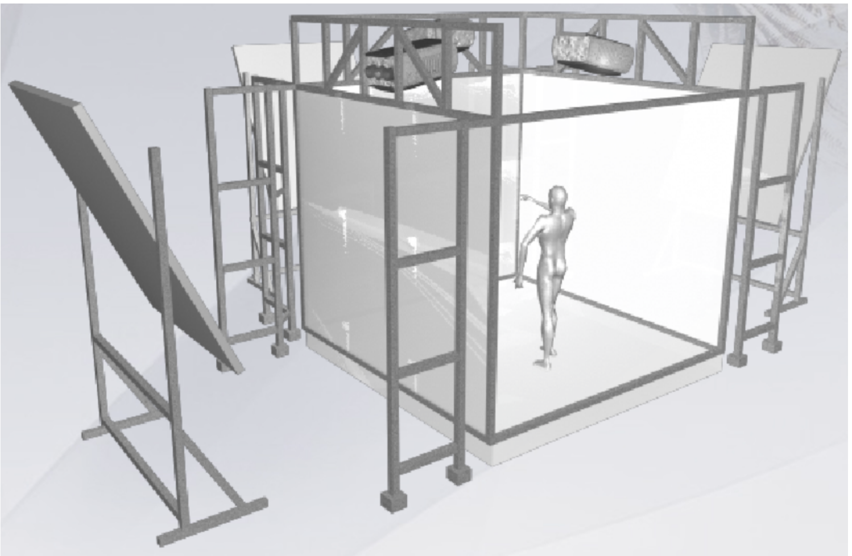

5-wall cave, FhG-IGD, Darmstadt

6-wall cave, Alborg, DK

RealityDeck - Immersive Giga-Pixel Display

- Developed at Stony Brook U, New York
- 308x 30" LCD displays
- 2560 x 1600 resolution per display
- 1.2+ billion pixels of resolution in total
- 40'x30'x11' physical dimensions
- 85 dual quad-core, dual-GPU cluster nodes

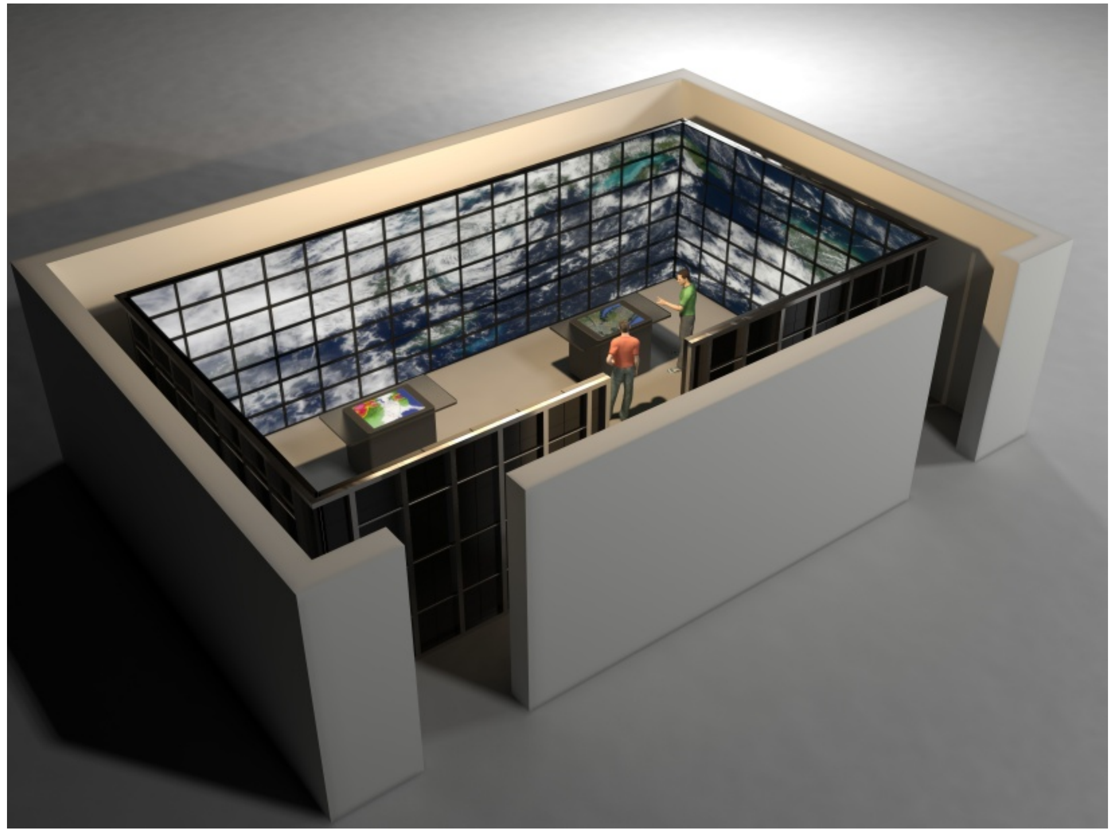

http://www.cs.stonybrook.edu/~realitydeck/

Curved Screens

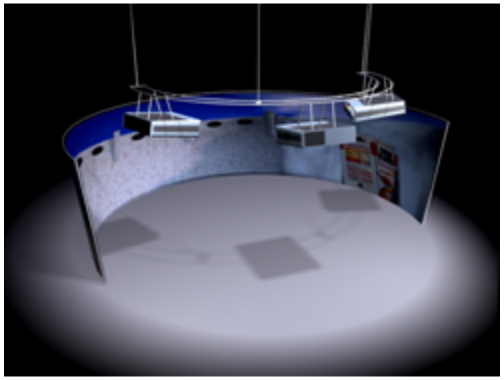

Curved Screen made out of 3D-TVs

- Idea: construct the walls of a Cave out of a (small) number of 3D TVs
- Advantage: reconfigurable relatively easily (just put the walls on wheels)

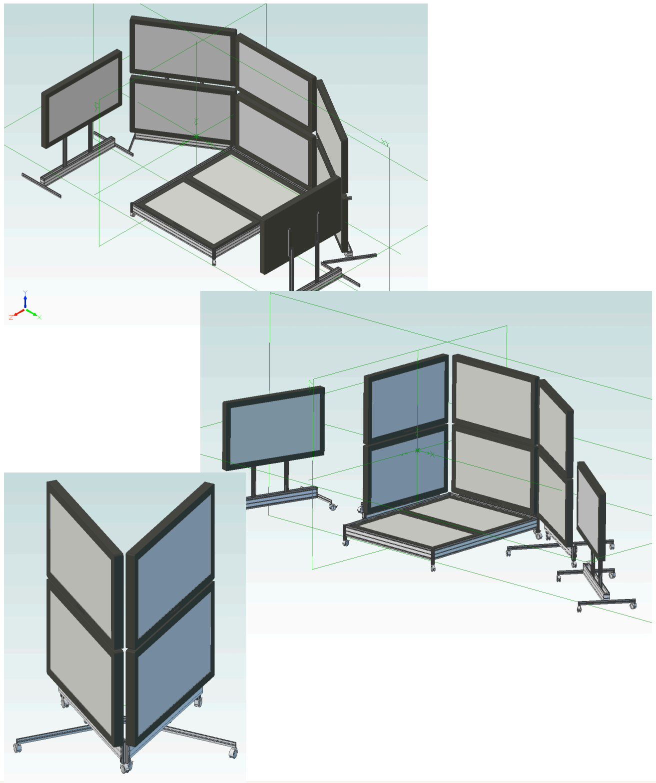

Personal Domes

- Example: Wii + Dome + MacBook Pro

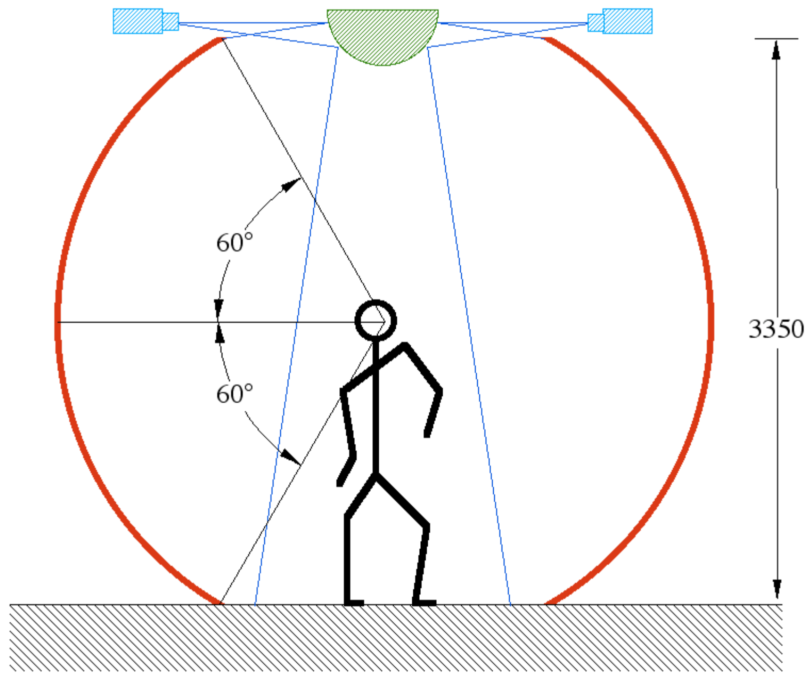

Source: Paul Bourke, University of Western Australia, http://local.wasp.uwa.edu.au/~pbourke/


GIZMODO

- A modern "Sensorama":

Immersa-Dome from Aardvark Applications

VirtuSphere

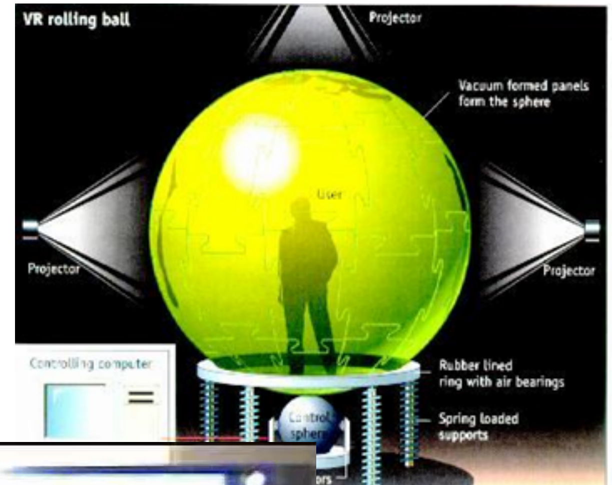

Studie

www.virtusphere.com


SPIEGEL ONLINE

Advantages and Disadvantages of IPDs

- Advantages:
  - Large resolution (currently up to ca. 1600 x 1280 per tile)
  - Large *field-of-view*
  - "*Non-invasive*"
  - No isolation of the real world
  - (Can accomodate Several Users)
  - *Cave*: turning the head results in small changes of the images

    → problem of latency is reduced / not so prominent
- Disadvantages:
  - Size
  - Price (lots of projectors, lots of graphics cards)
  - Precision, calibration
  - Potentially "*stereoscopic violation*"
  - Correct view only for one *viewer* (unless a massive amount of hardware is used)

Retina Displays (retinal displays)

- Idea:
  - Use the human retina as the display surface (all images from the outer world end up there anyway)
  - Use a laser to write the image by scanlines into the eye
- Advantages:
  - Can be miniaturized (potentially)
  - High contrasts, high brightness
  - Good for *see-through* displays
  - Small power consumption

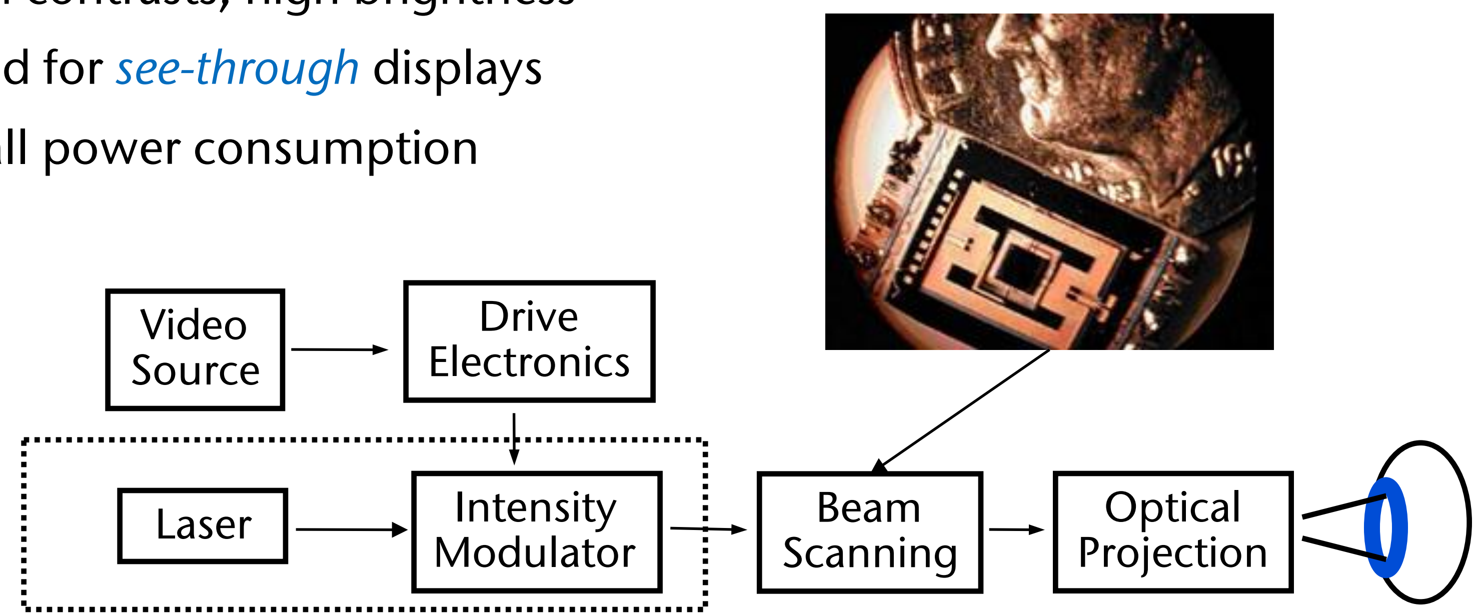

Retinal display

Design study

Holographic / Volumetric Displays

- Real 3-dimensional displays
- Advantages:
  - Provide correct perspective/view from *every* angle!
  - Coherence between accomodation and convergence
  - Depth of field (Tiefen(un-)schärfe)
- Holographic displays: algorithmic computation of holograms
- Problems:
  - Staggering amount of computational work
  - Colors
- Volumetric displays: voxel are projected into a volume (as opposed to a surface)
- Problems:
  - Size of data (e.g. 100 mega-voxels = 1000x1000x100 display resolution)
  - Occlusions?

- Example volumetric display:
  - 198 x 768 x 768 ≈ 100 million voxels
  - Frame rate: 20 Hz


Unusual Display Surfaces

- Fog ("fog screen"):
  - Laminar, non-turbulent air flow
  - Water droplets are "sandwiched" within the air flow

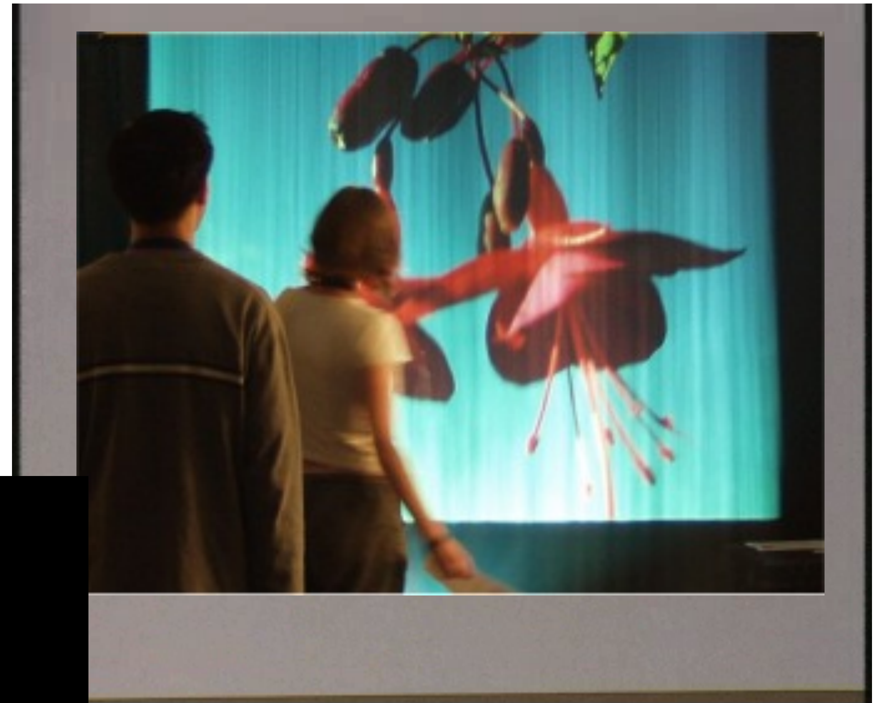

FogScreen

Olwal et al.: Consigalo: Multi-user, Face-to-face Interaction on an Immaterial Display

- "Everywhere displays":

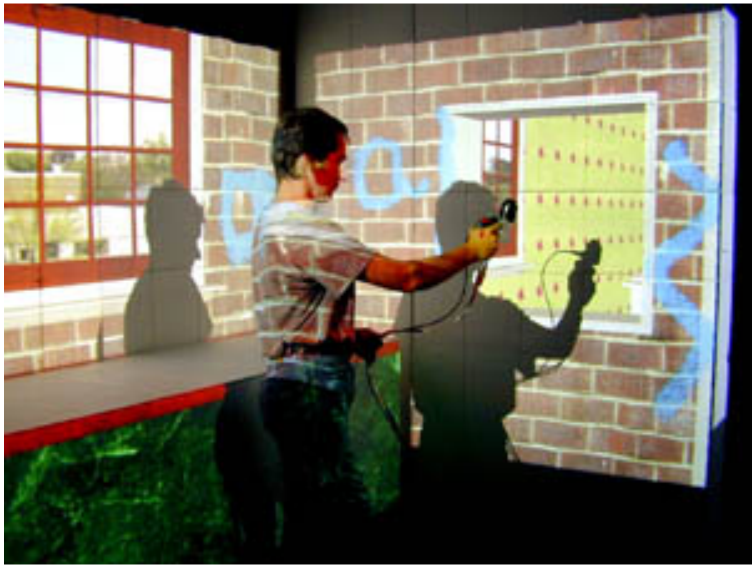

The History of Stereo Images

- Euklid (4th century BC)
- Sir Charles Wheatstone (1838)

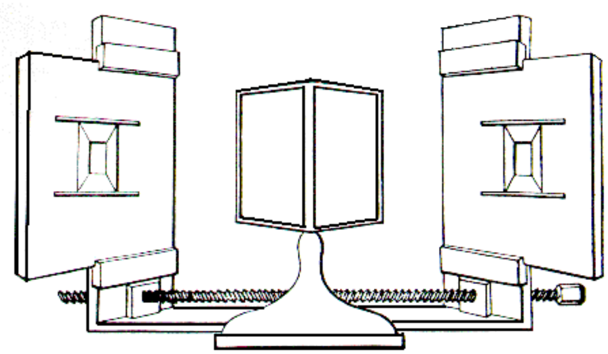

- 1860: 1 million Stereoscopes sold

- 1950-ies:

- Today (demo):

How to Project Stereo with only *one* Display Surface?

- Need some kind of *Multiplexing*

1. Temporal Multiplexing ("active stereo"):
   - Typically 1 projector (e.g. monitor)
   - Project/render alternatingly left/right image
   - Synchronously, switch left/right glass of *shutter glasses* to pass-through
   - Shutter glasses run with 120 Hz → 60 Hz framerate

2. Multiplexing by polarization ("passive stereo"):
   - Usually 2 projectors displaying on same surface
   - Project left/right simultaneously but with different Polarization of the light
   - Polarization glasses let only left/right images pass, resp.

- Kinds of polarization:
  1. Linear polarization:
     - Any direction perpendicular to direction of travel of light

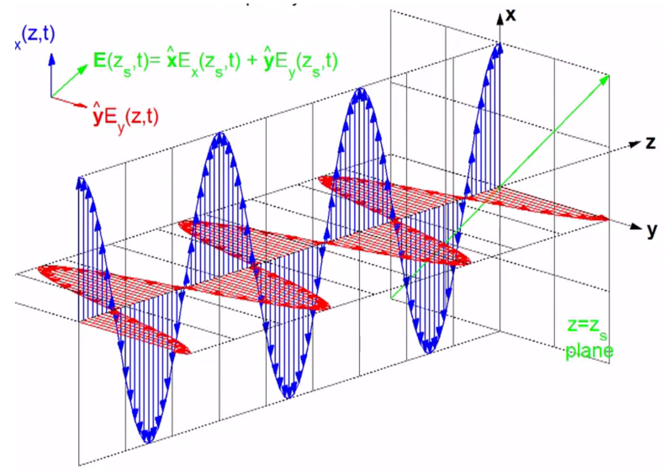

  2. Circular polarization:
     - Left-handed / right-handed polarization

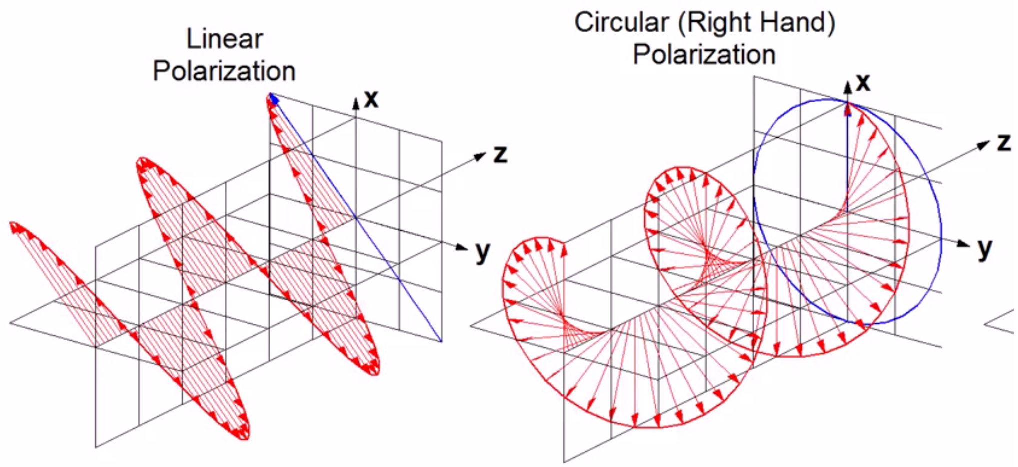

Demo

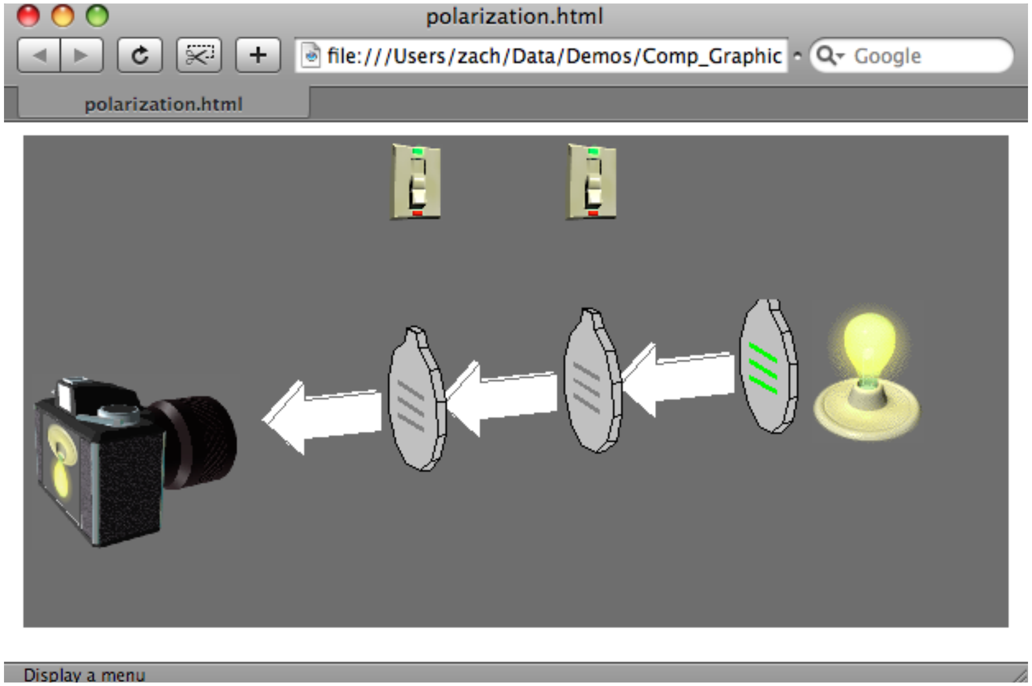

http://www.colorado.edu/physics/2000/applets/polarization.html

"Color Multiplexing"

- Simple version: *Anaglyph stereo* (red-green stereo)

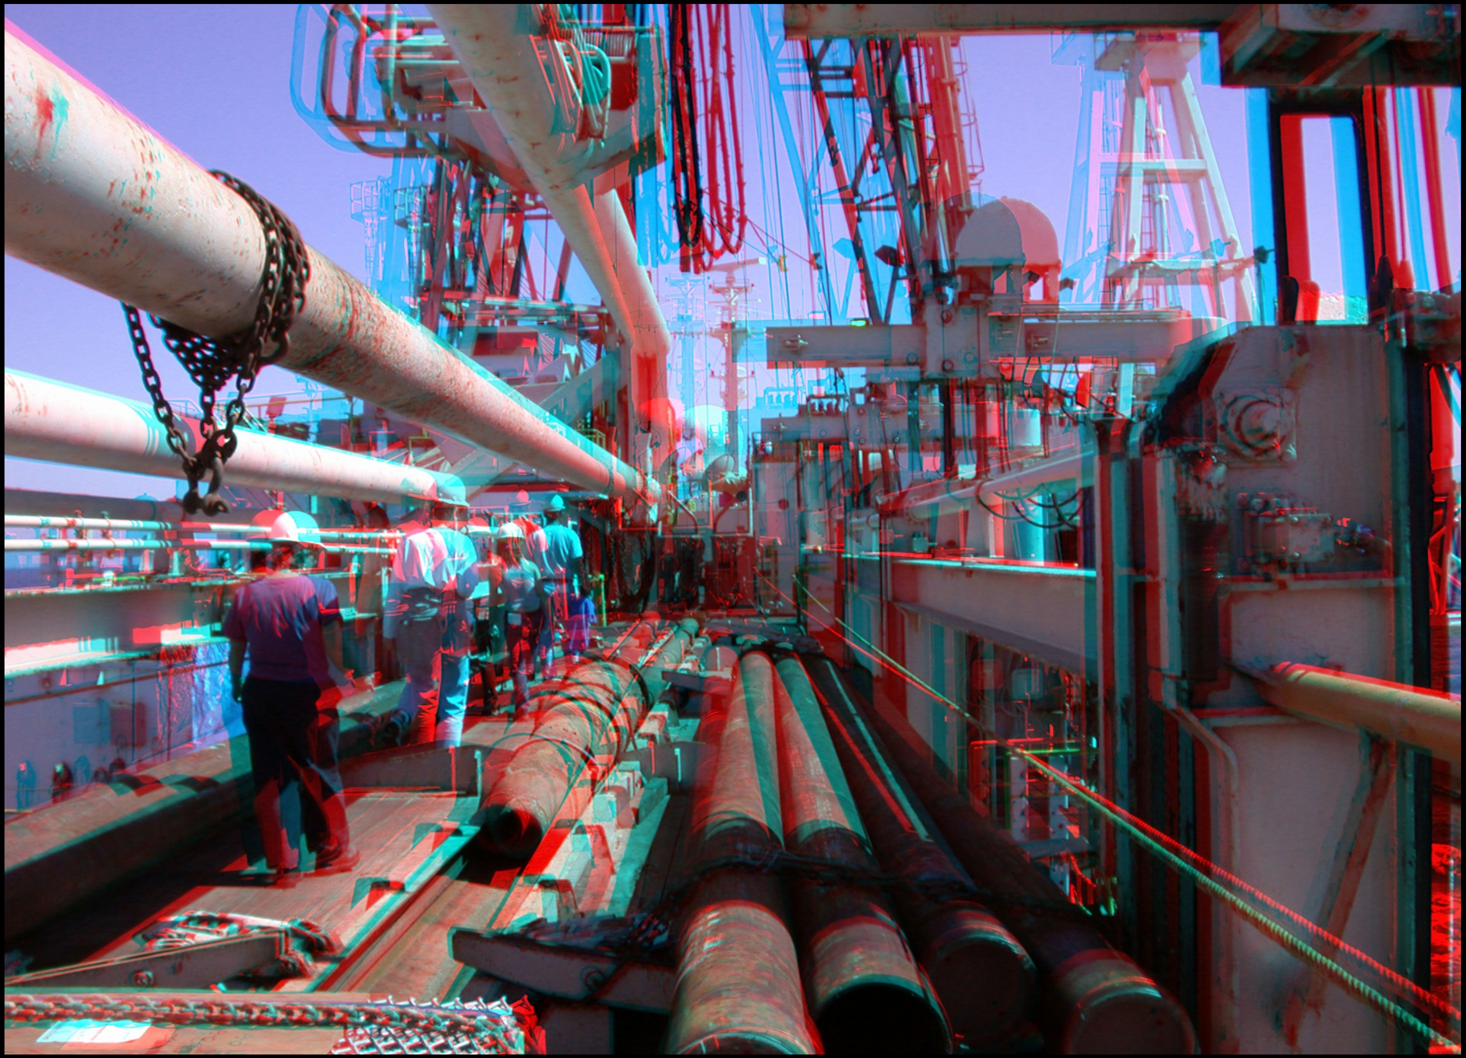


RESOLUTION-

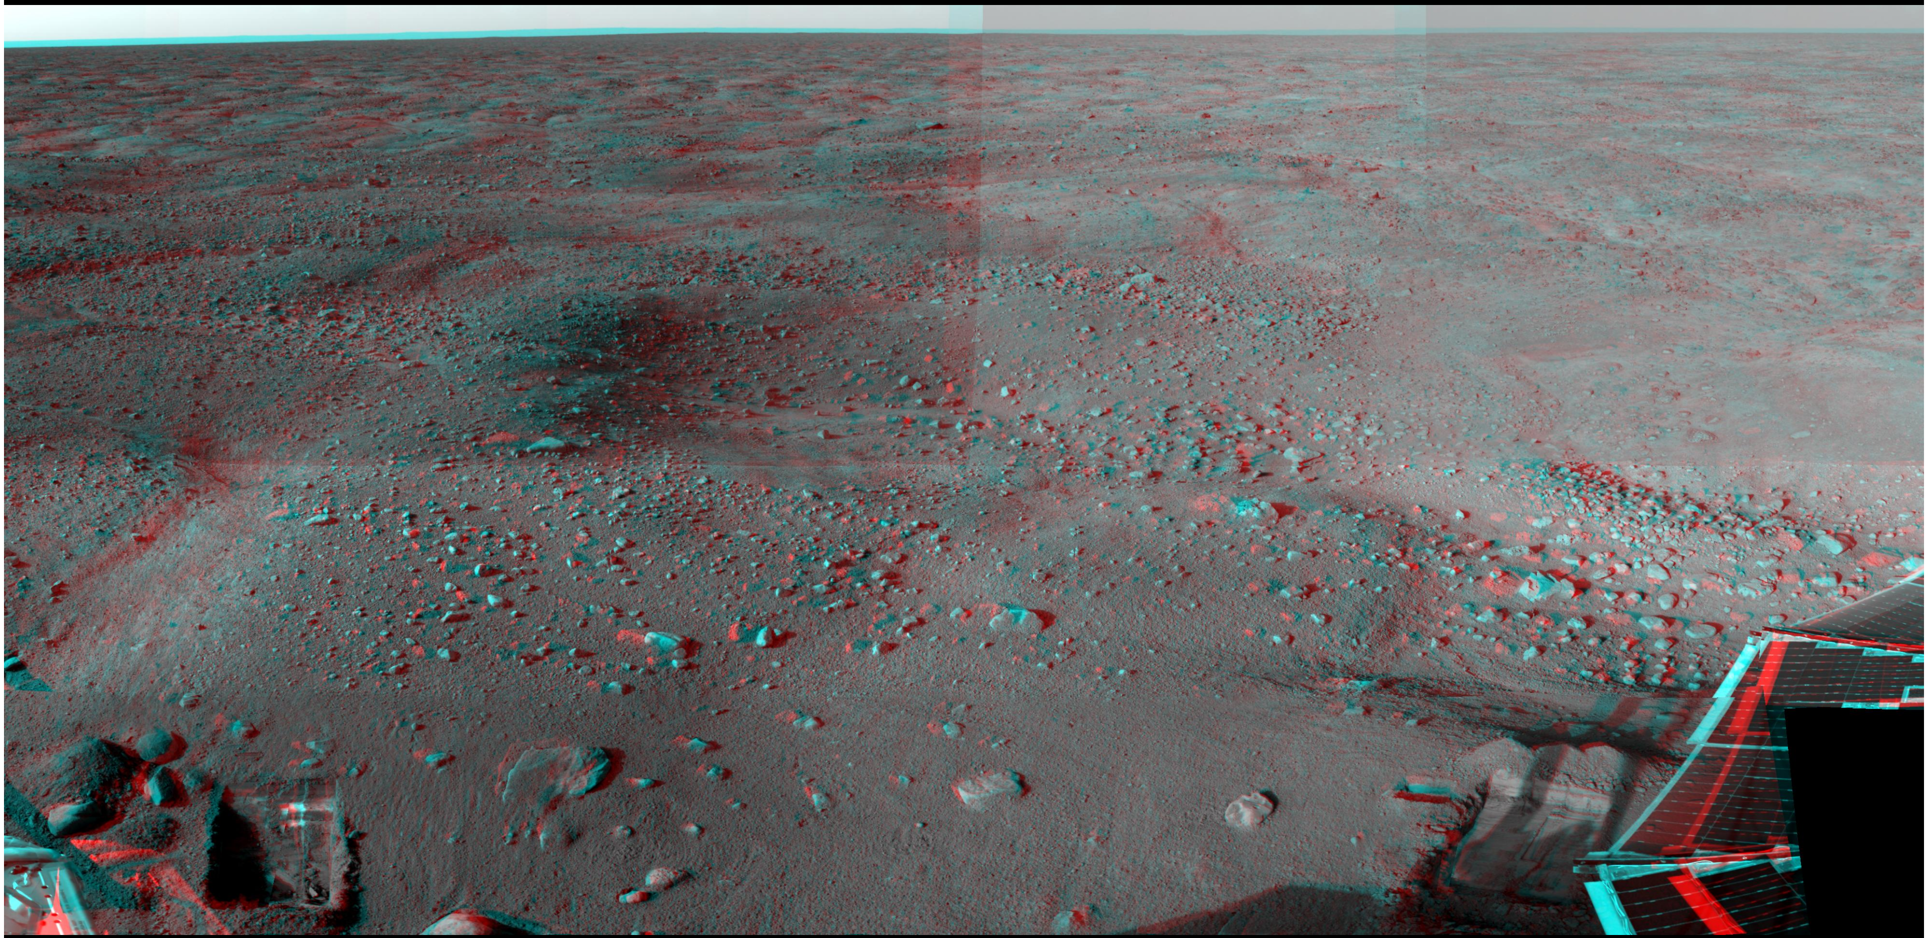

- Generalization ("*Infitec*", Dolby3D, spectral comb filter, wavelength multiplex):
  - Each of the primary colors must pass through a narrow band pass filter
  - Left & right eye get filters with interleaving band passes

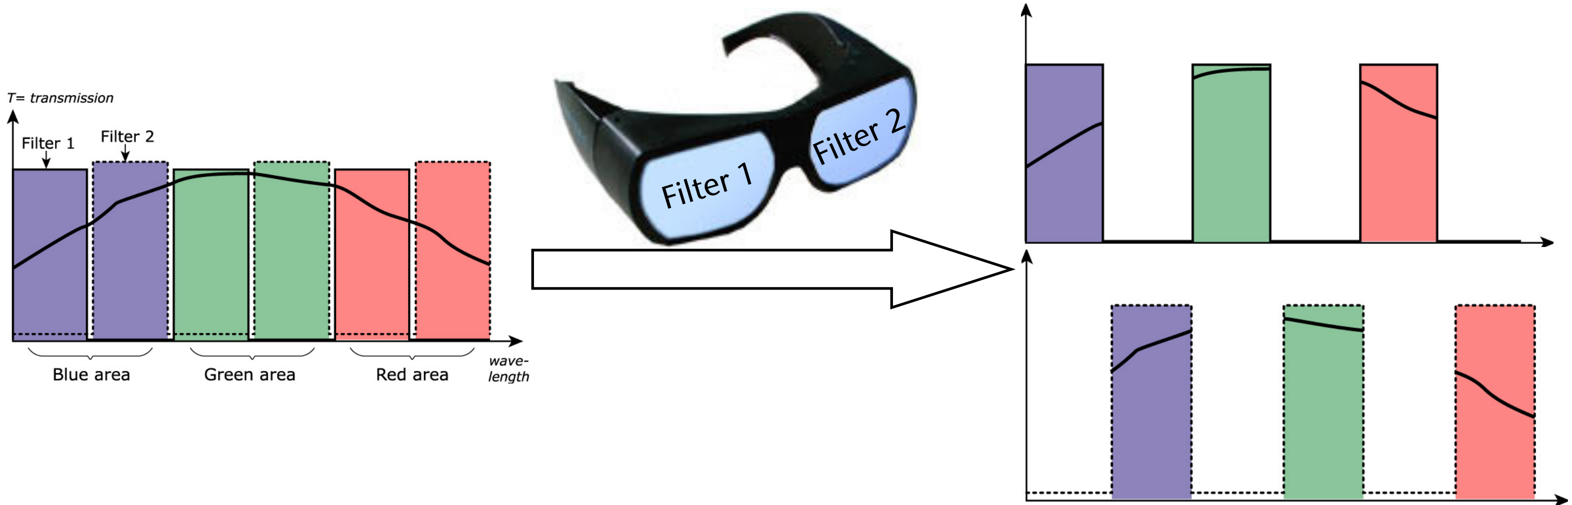

- Problem:
  - Color fidelity

The Problem of Multiple User and a Single IPD

- Problem with a single-tracked projection (stereo or mono): only the viewpoint of the tracked users is correct, only she will see a correct image!
- Example:

Image's perspective is correct for the user

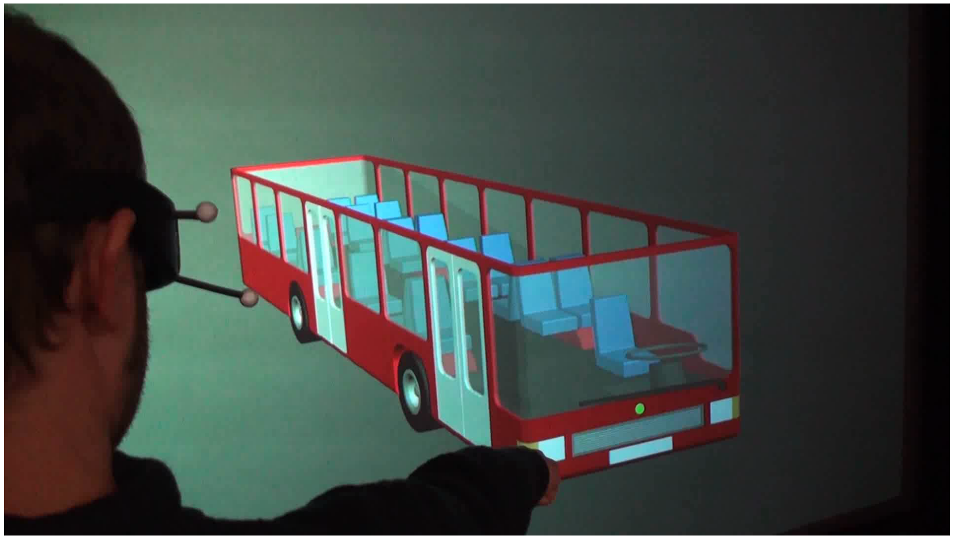

Image's perspective is correct for the (real) camera

Correct (Stereo-)Projection for Multiple Users

- Probably only possible for a small number of users
- *Temporally multiplexed*:
  - Framerate for multi-user stereo = Framerate for mono * $2 \cdot$ #User
- Infitec for several users:
  - Each user gets glasses with slightly shifted comb filters
  - With n users we need $2n$ different comb filters → extremely narrow bands
- *Spatially multiplexed*
- Combination of the above

Spatial Multiplexing

- Proj. surface is partitioned among users
- Consequence: interdependence between
  - Size of the *view frustum*
  - Working volume of users
  - D & radius of hole
- Example:

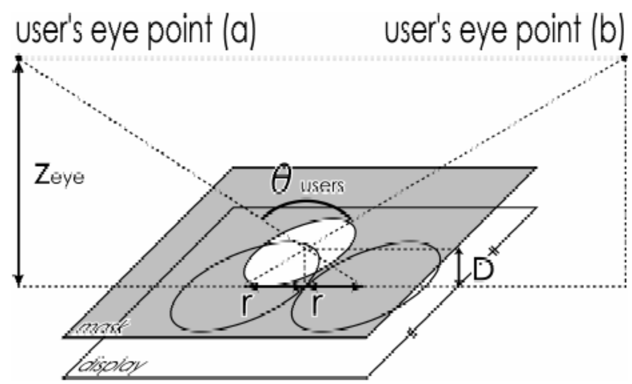

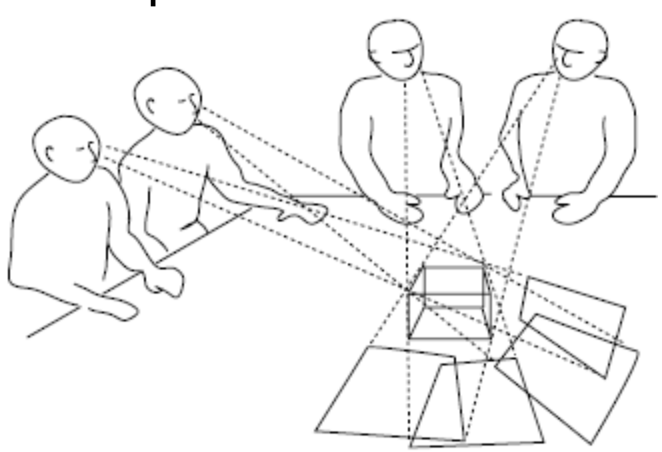

IllusionHole @ Siggraph 2001

Stereo for 6 Users

[2011]

- Combination of active and passive stereo,
  plus ingenious utilization of time-sequential projectors
- Recap from CG1:
  time-sequential RGB with DLPs

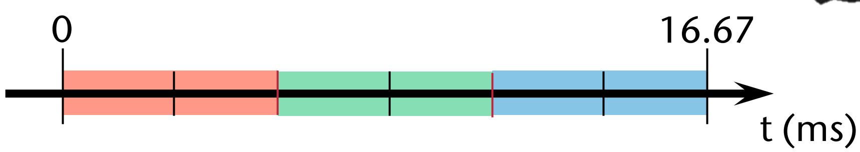

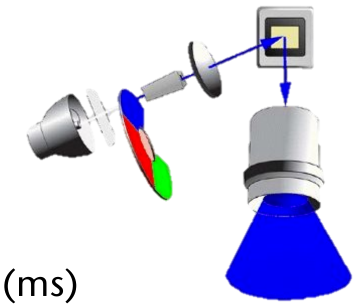

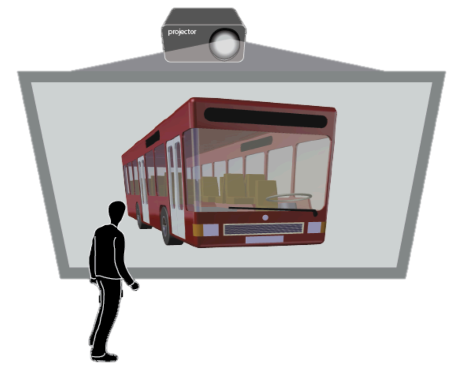

1. Modification: remove color wheel
2. Modification: each user gets shutter glasses that additionally has left/right polarization filters
   - Must be fast enough to prevent cross-talk!

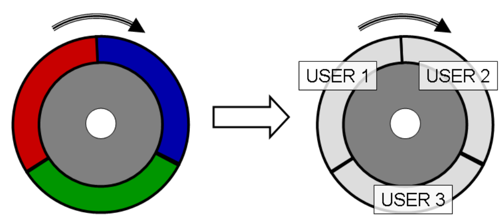

The Hardware in *Principle*

- 6 stereo video streams are generated by 6 graphics cards in 1 PC
- Distribution of the video streams to the projectors via multiplexers

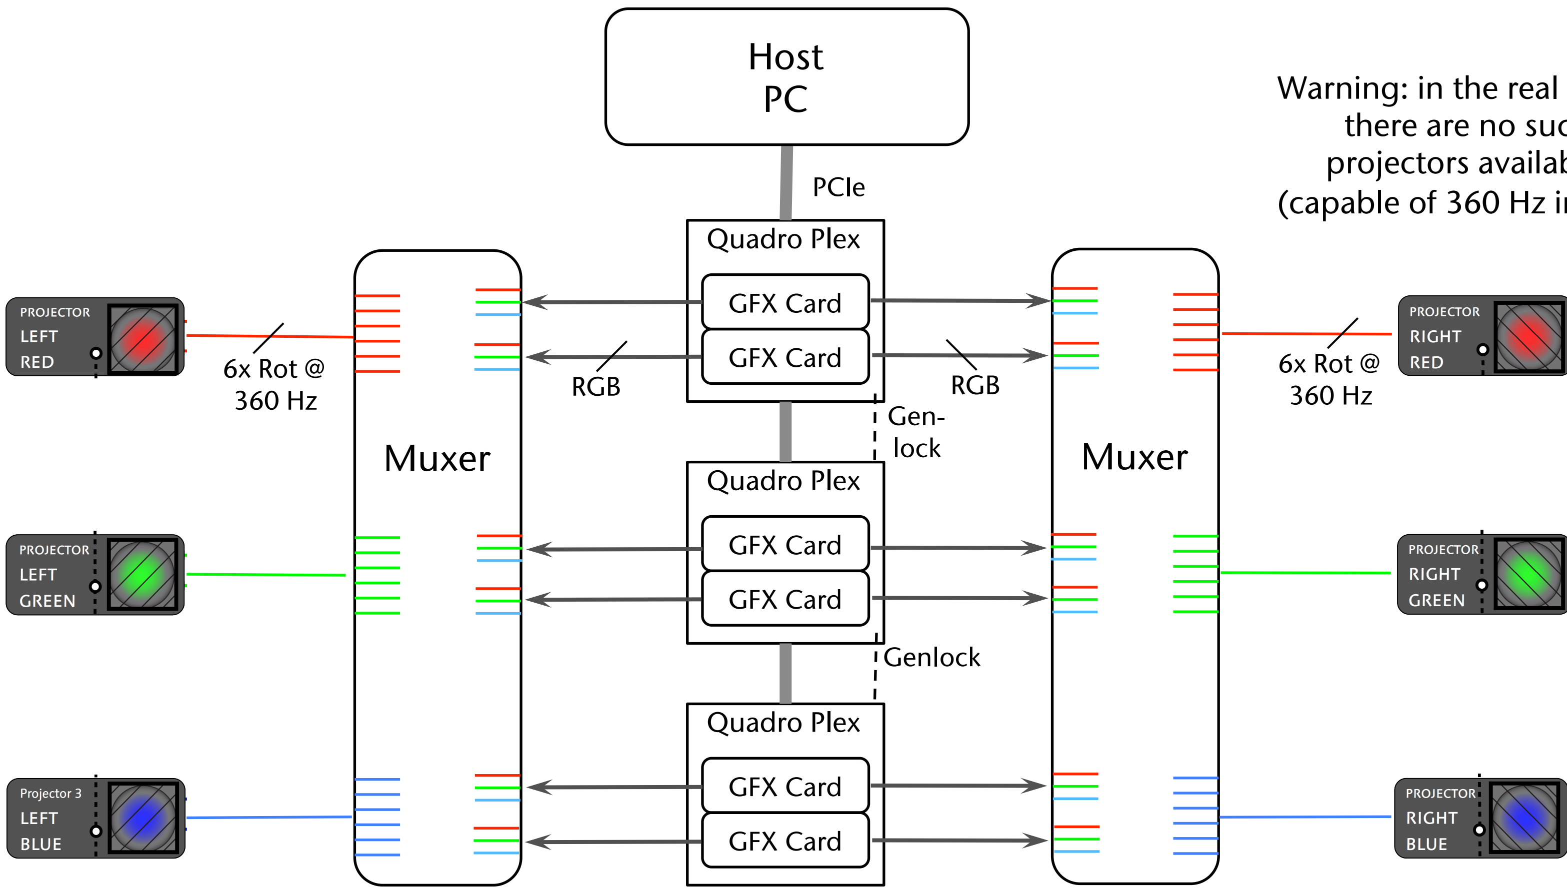

Warning: in the real world, there are no such projectors available! (capable of 360 Hz inputs)

- Timing:

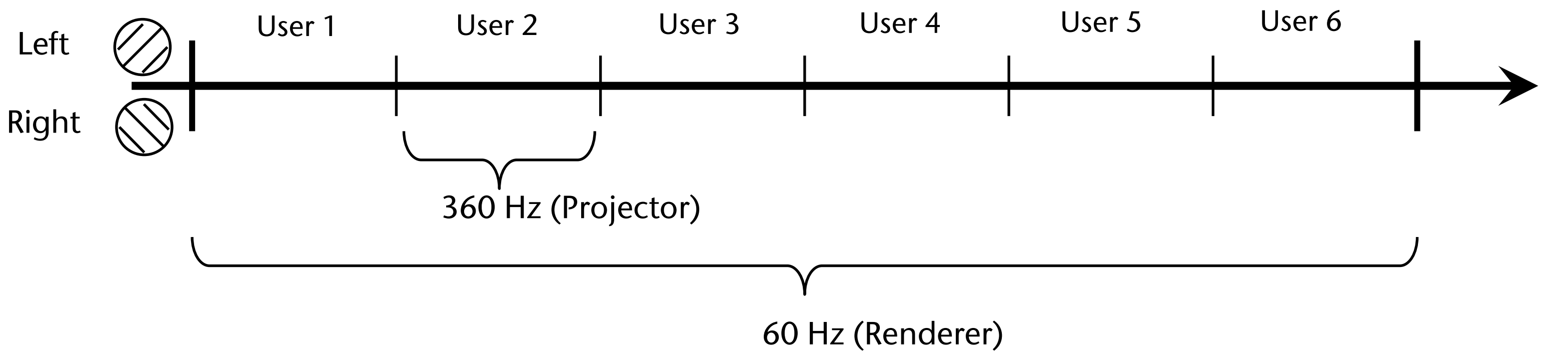

- Video:

Outlook

- With *perspectively correct* projections for *each* user, the shared 3D space will become coherent for all users
- Consequence: direct communcation (including pointing) in co-located CSCW is possible

Agrawala et al. 1997


Arthur et al. 1998

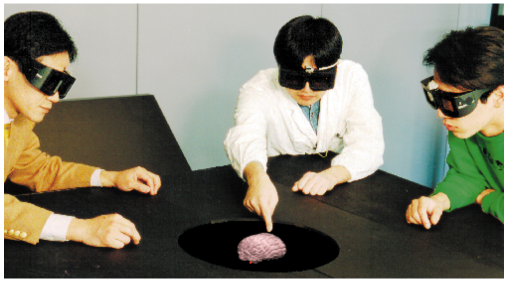
Kitamura et al. 2001

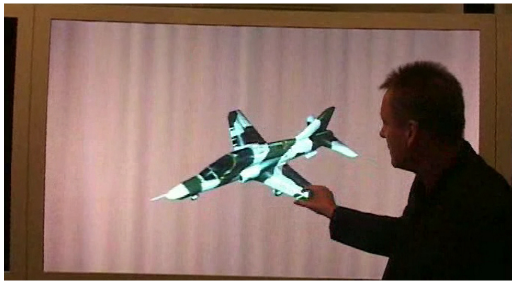
Agócs et al. 2006

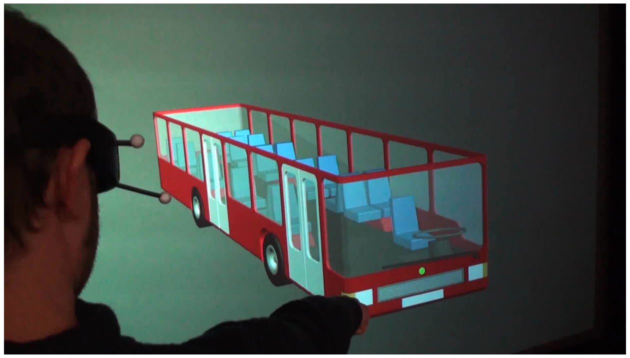

Interaction Issue with Multi-User-Stereo

- Navigation: the "navigator" controls the path for all users (and he sees only his own viewpoint!)
- Problem: the other users' viewpoint goes through walls
- Solution:
  - Adjust the paths of the other users automatically to bring them closer to the navigator's viewpoint
  - Fade away obstacles in the path of each user

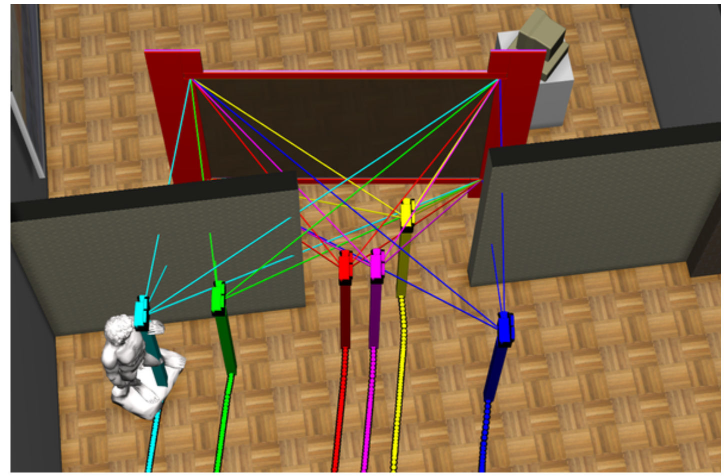

Stereoscopic Effect Based on the Pulfrich-Effect

- See slide "Pulfrich Effekt" in *Optische Täuschungen*

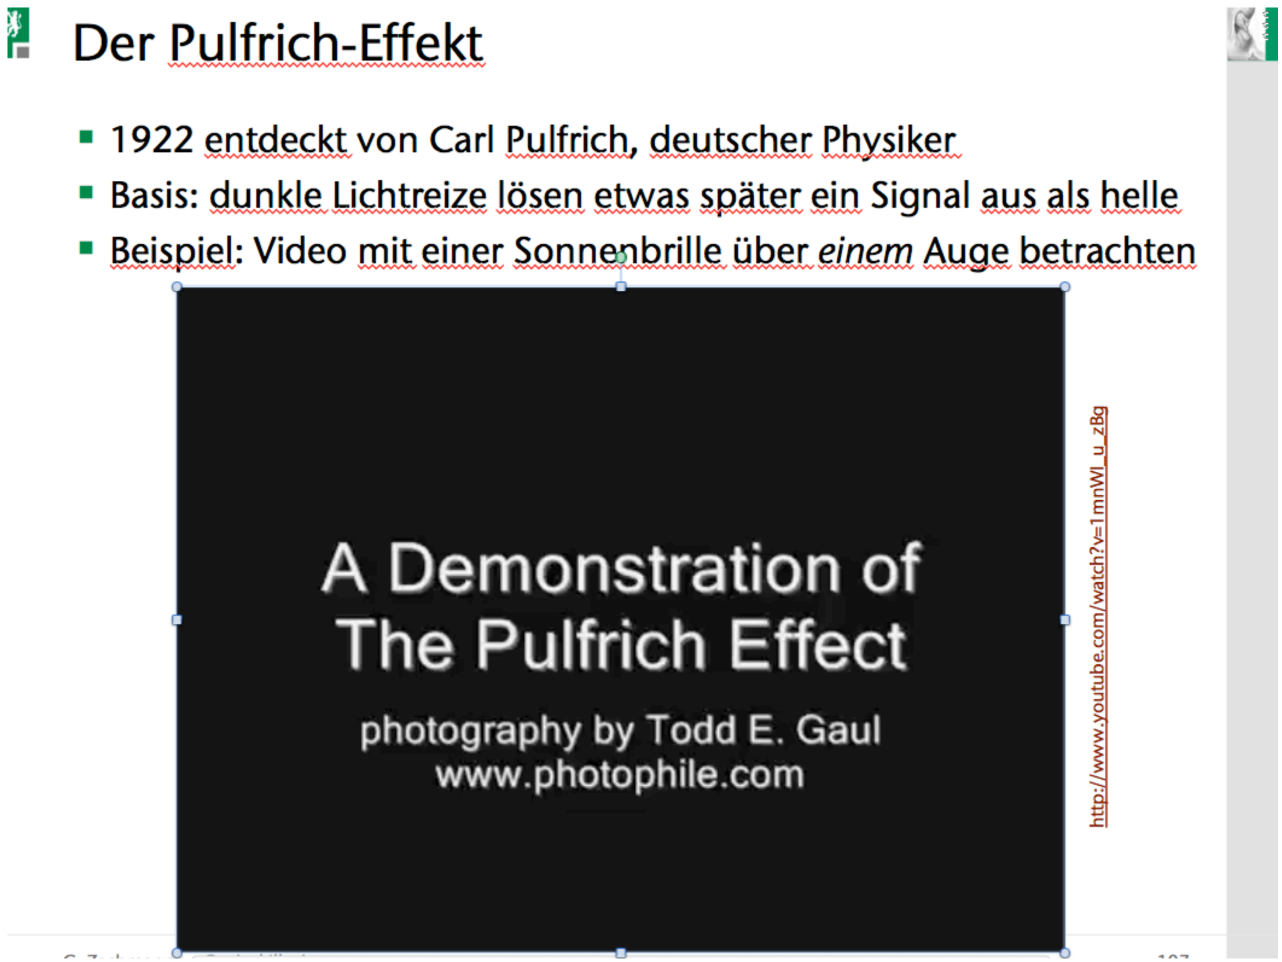

Binocular / Stereo Vision

- Important "*depth cue*" (not the most important one)
- Works just up to a few meters (< 6 m, ca.)
- Disparity between the eyes $= \delta_2 - \delta_1 = \gamma - \alpha$
- Horopter = locus of points in space with same apparent depth as the fixated object = point with 0-disparity
- Parallax on the screen:

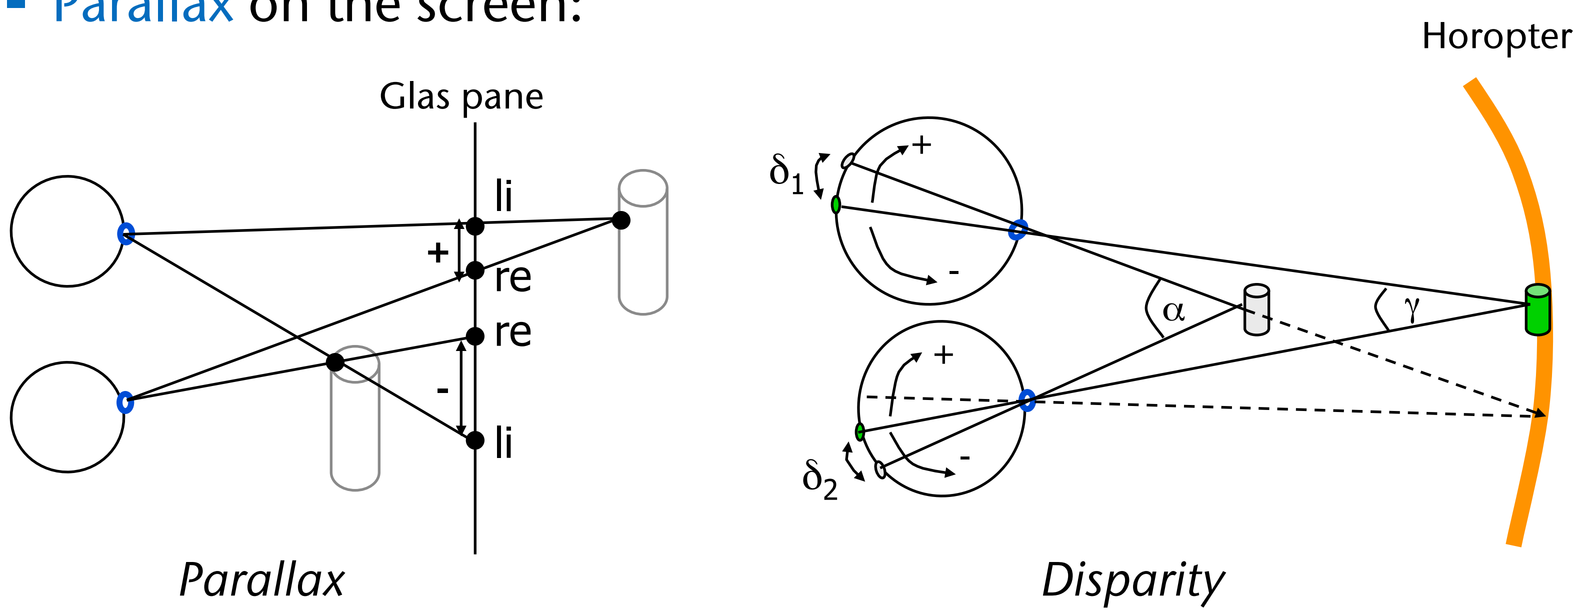

Measuring the Horopter

- First of all: not measuring it, but *constructing* it → Vieth-Müller Circle = theoretical locus of points in space that stimulate corresponding retinal points

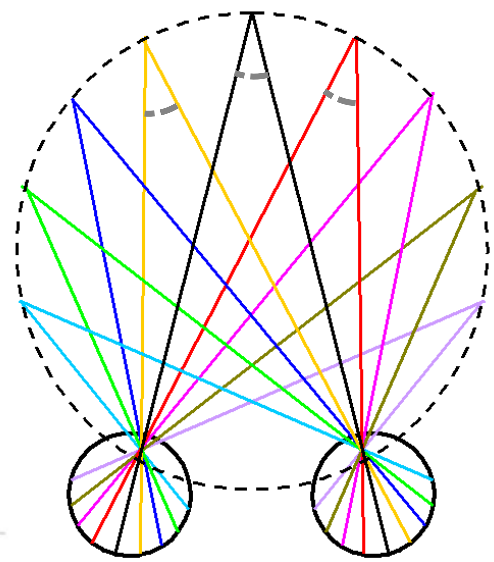

- Measuring the Horopter with the "Apparent Fronto-Parallel Plane" method:
  - Subject is asked to arrange a series of objects so that there appears to be no depth difference between them

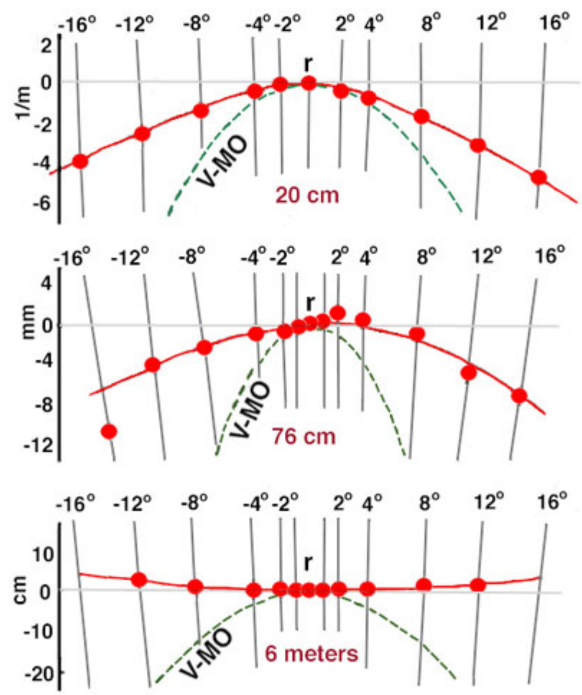

Panum's Fusional Area

- There is a zone/range of depth around the horopter, where the brain is able to fuse the double image of an object
  → Panum's Area of Fusion

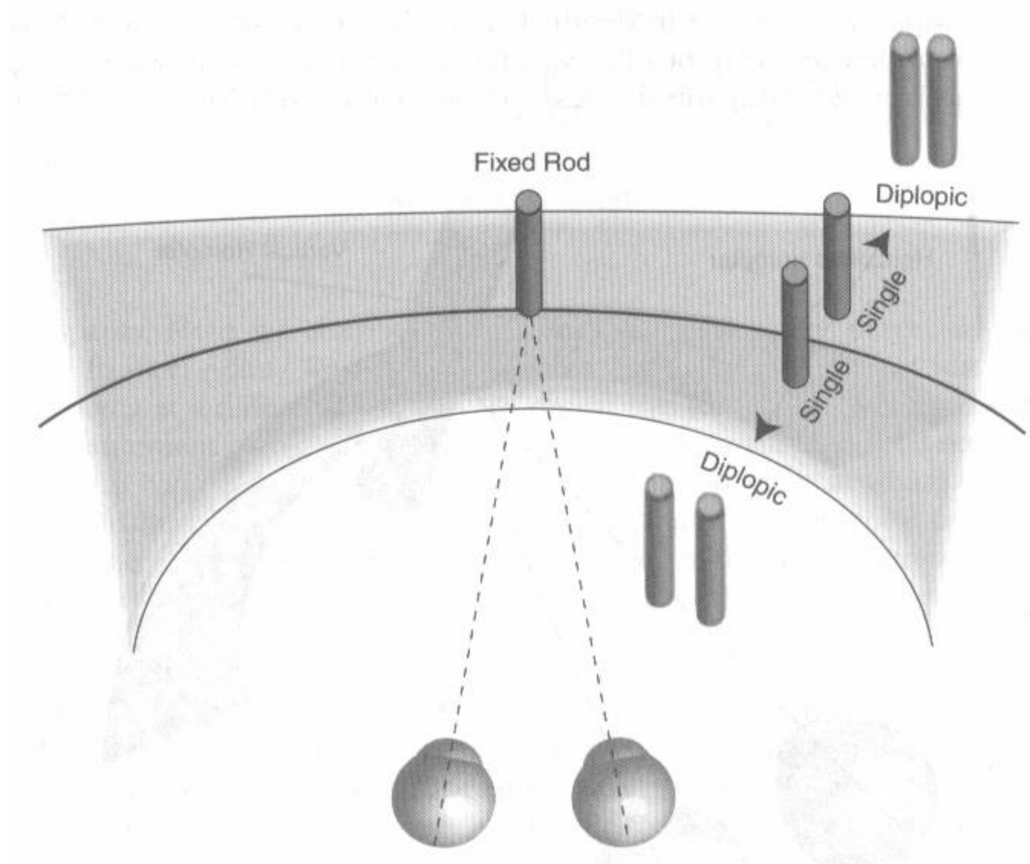

Stereo Projection

- Wrong: konverging view vectors

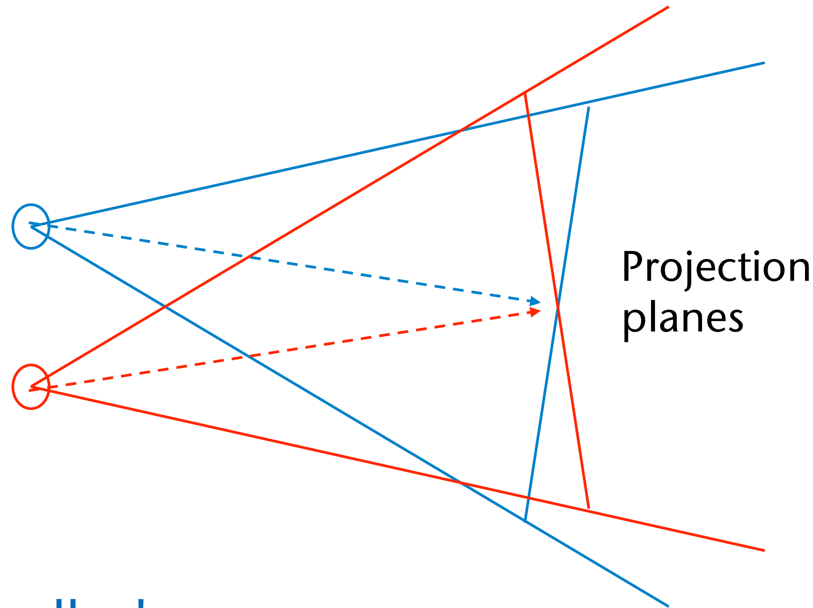

- Problem: verticale Parallax!

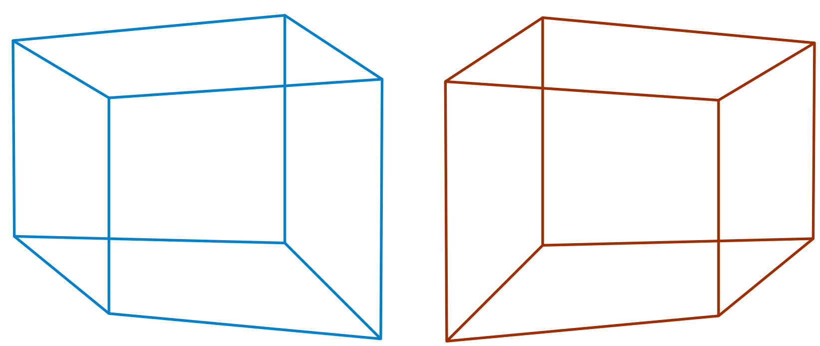

Correct Stereoscopic Projection

- Right: parallel viewing vectors
  → *off-axis perspective projection*

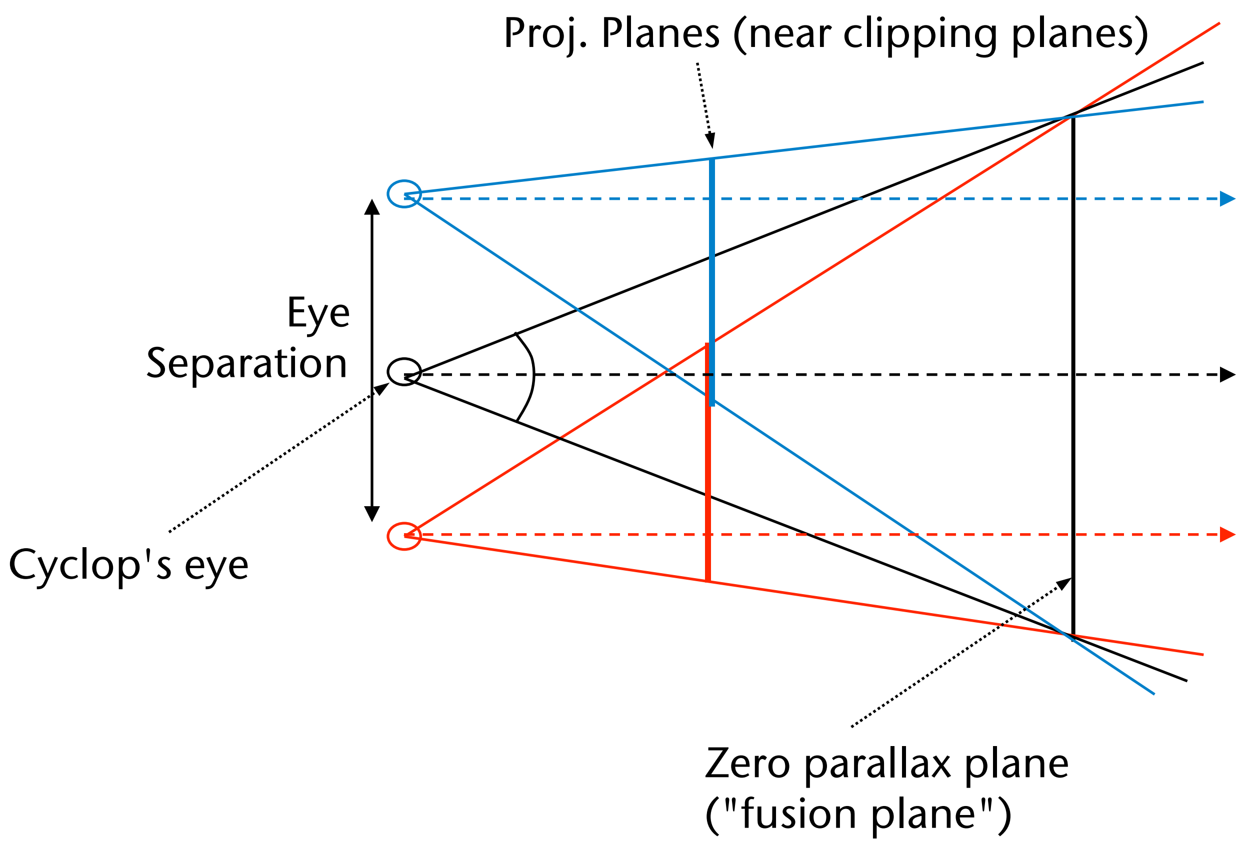

- Given i, aspect ratio w/h, horizontal FoV α, near plane n , zero-parallax depth z_0
- Determine *left/right/top/bottom* for **`glFrustum`**()
- Assumption: no head tracking, i.e., cyklop's eye is in front of the center of the zero parallax plane
- *top* and *bottom* are as usual:

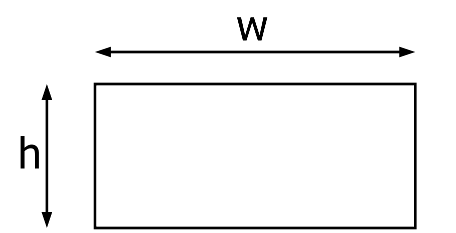

$$t = \frac{h}{w} l$$

- Example: *left* for left eye:

$$l_c = n \tan \frac{\alpha}{2}$$

$$l' - l_c = i \frac{z_0 - n}{z_0}$$

$$l = l_c + (l' - l_c) - i = l_c - i \frac{n}{z_0}$$

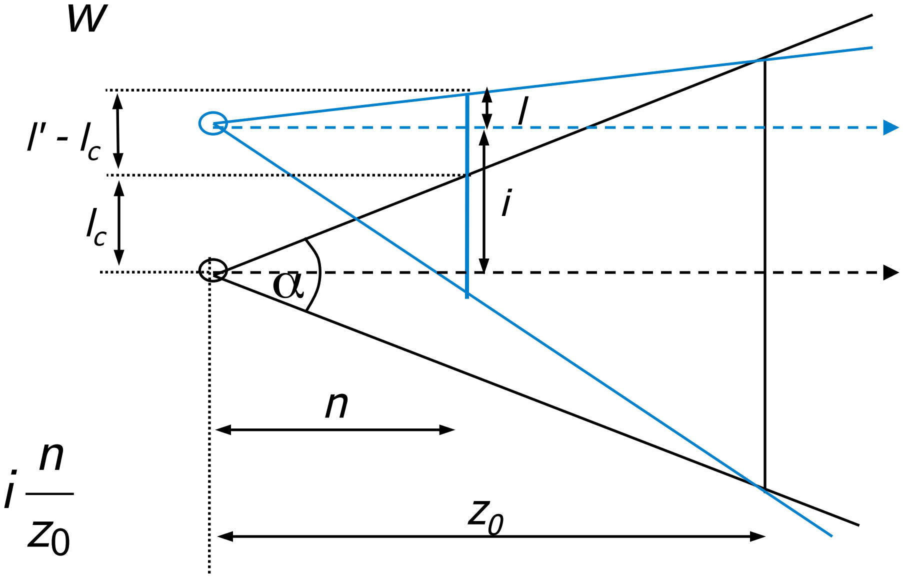

Problems with Stereo Rendering: Depth aliasing

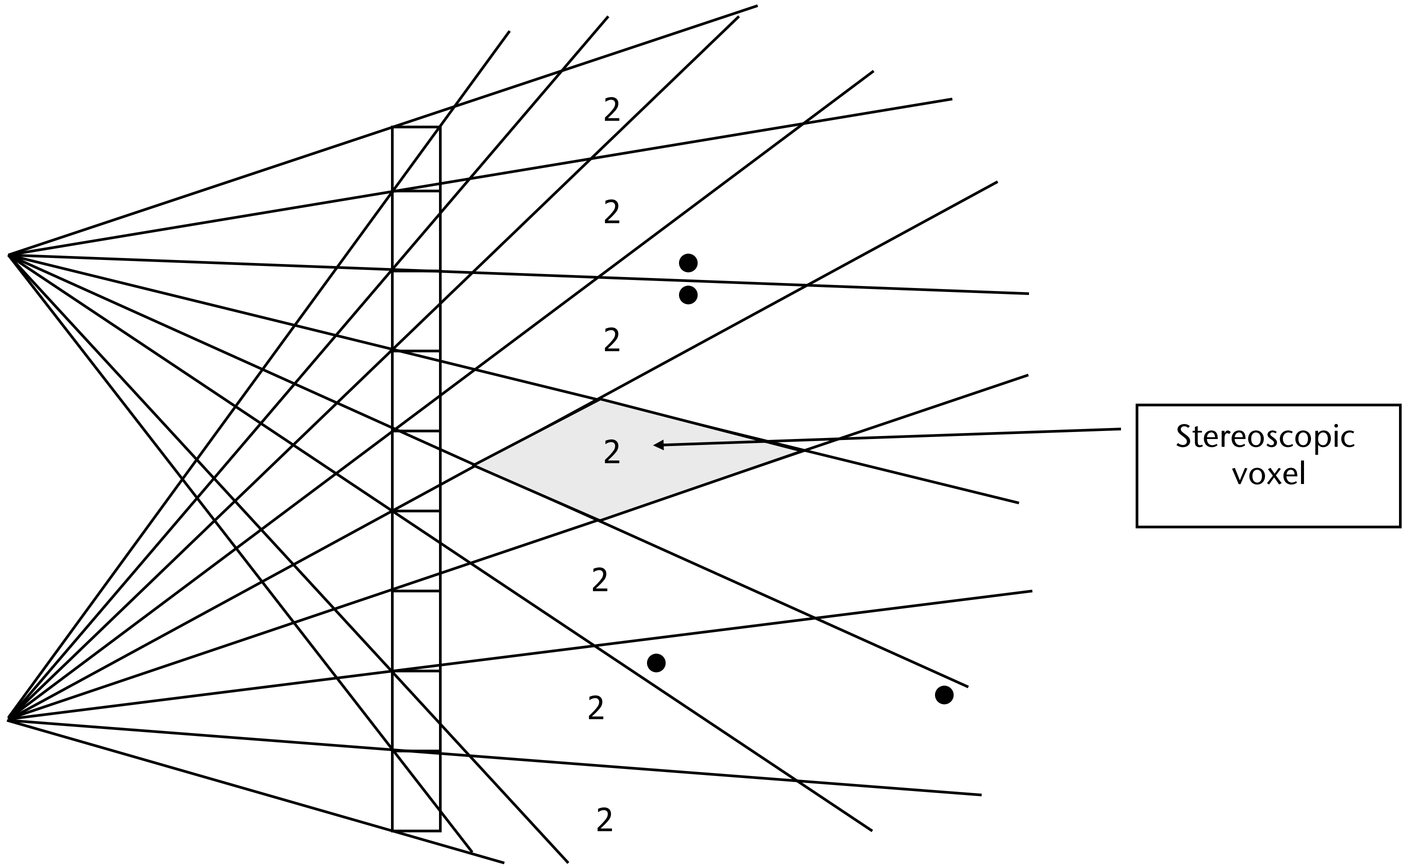

- This effect would occur, even if the Z-buffer was continuous!

Digression: Accomodation and Convergence

- Two important terms that get confused very easily
- Convergence = counterrotating eye movement (around the vertical axis), so that the optical axes of the eyes intersect at some point (fixation point)
  - So that the fixated object appear on the center of the retina (has highest resolution)
- Accomodation = adjustment of the eyes' lenses to adapt for different distances
  - So that the fixated object appears sharp on the retina
    - (Personally, I am often times confusing accomodation with focussing ;-))

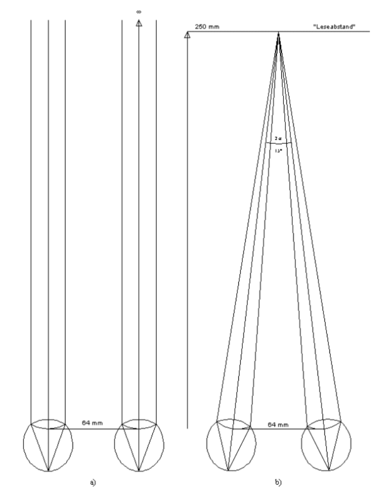

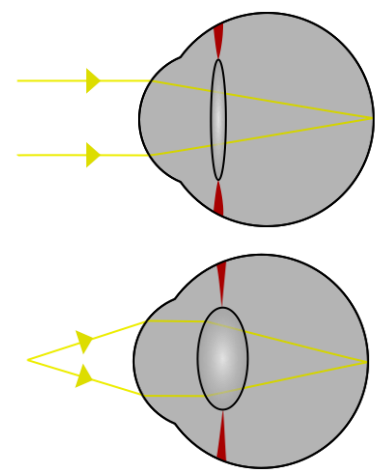

Experiment

- The following image appears to be 3-dimensional, if you can decouple accomodation and convergence:

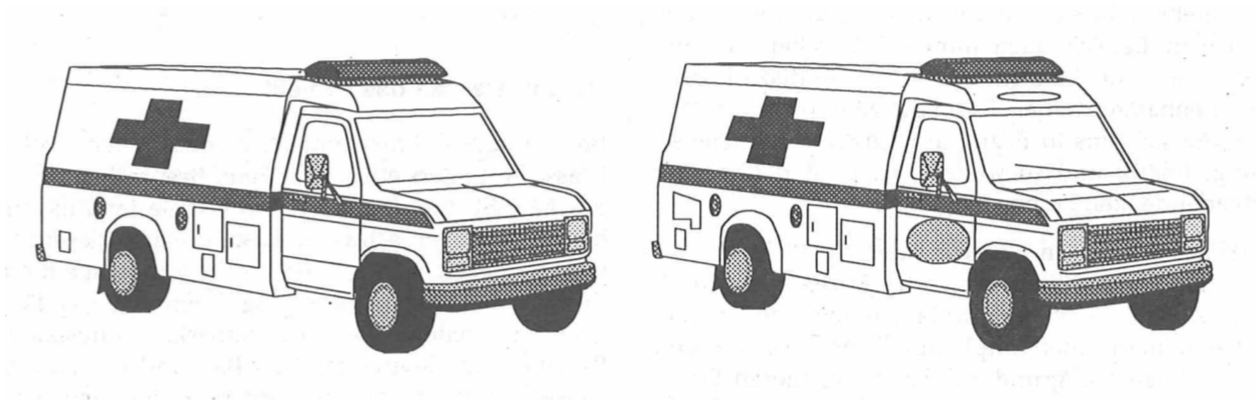

- "Magic Eye" images are constructed patterns such that corresponding points convey depth :

"depth image" for it

A Problem with Distance Estimations in Stereo Images

- Effect: in a Cave or Powerwall, objects appear more distant than they are
- My Hypothesis:

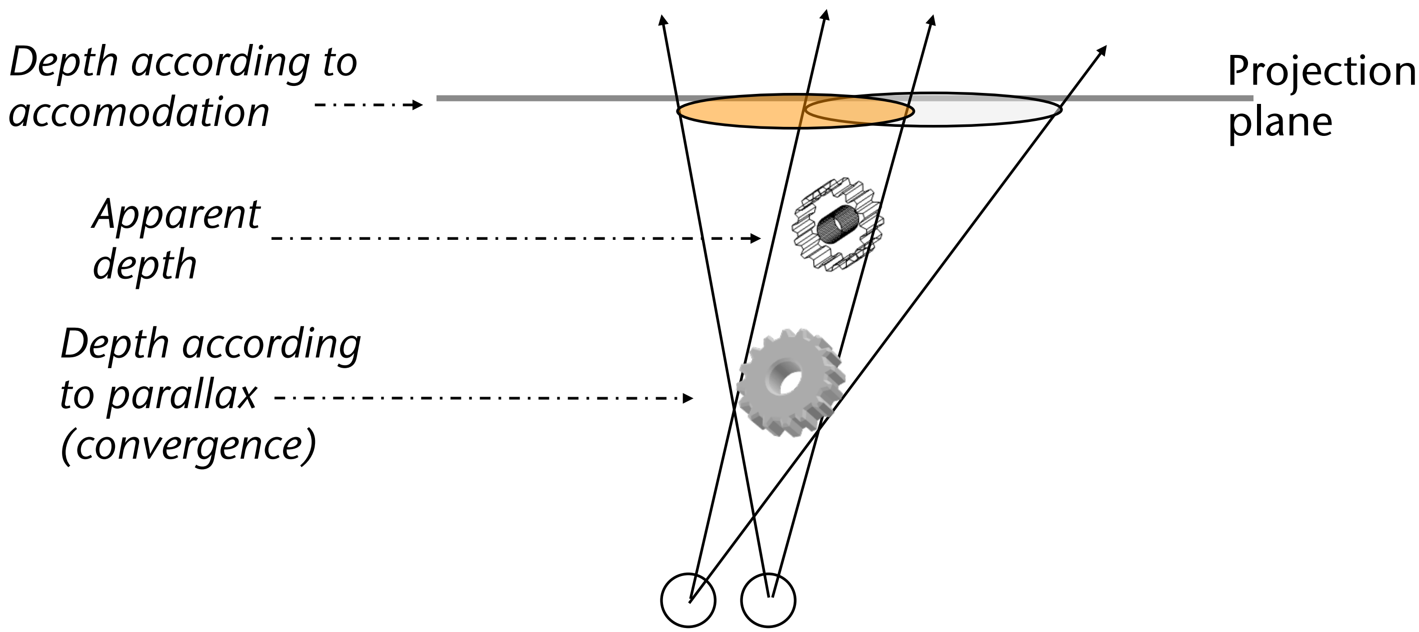

- Solution: holographic or volumetric displays
  - Which have other problems …

Stereo is a "*one man show*"

- Why are stereoscopic images correct only for 1 viewpoint?

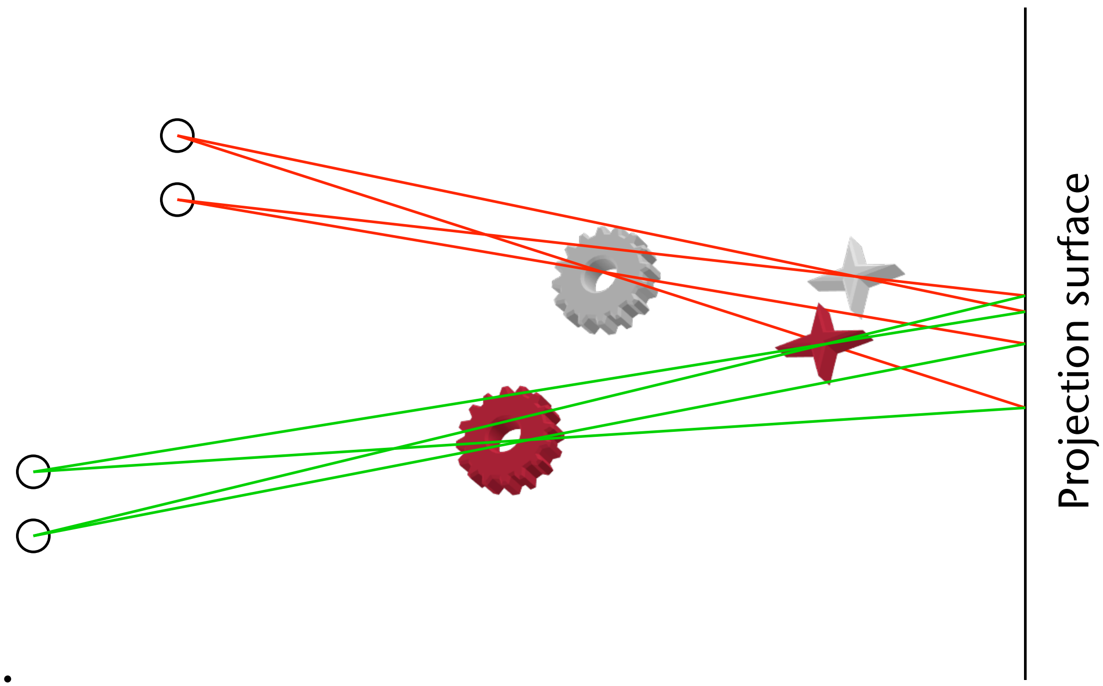

- Solution:
  - For 2 viewers: render 4 images; or
  - Holographic/volumetric displays

Stereo Violation

- 2 effects that can occur together:
  - *Clipping*
  - Depth from stereoscopic image

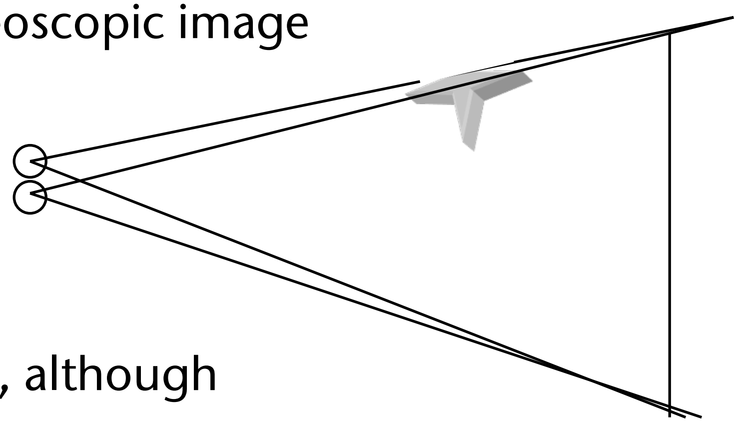

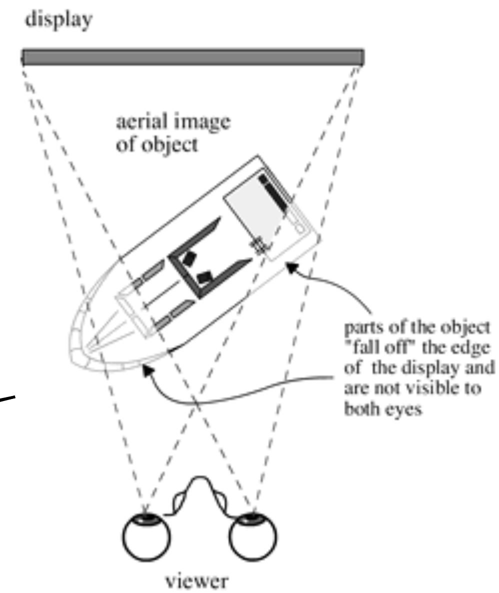

  - Object is clipped, although *in front of* the projection surface
  - Consequence: conflicting *depth cues* → *stereo violation*
- Example: lower left corner of the anaglyph mars image

Rendering onto Several Projection Walls (e.g. Caves)

- Represent the real projection surfaces by a polygonal object (e.g., cube) in the VE
- Rotate each polygon (and viewpoint) to make it perpendicular to z-axis (standard OpenGL viewing transf.)
- Determine *left/right/top/bottom*, just like previously (for single proj. wall)
- Careful with the complete viewing/ projection transformation! If set up exactly like single wall projections, then ...
  - Specular lighting will have discontinuities
  - Generated texture coords (by OpenGL) will be discontinuous

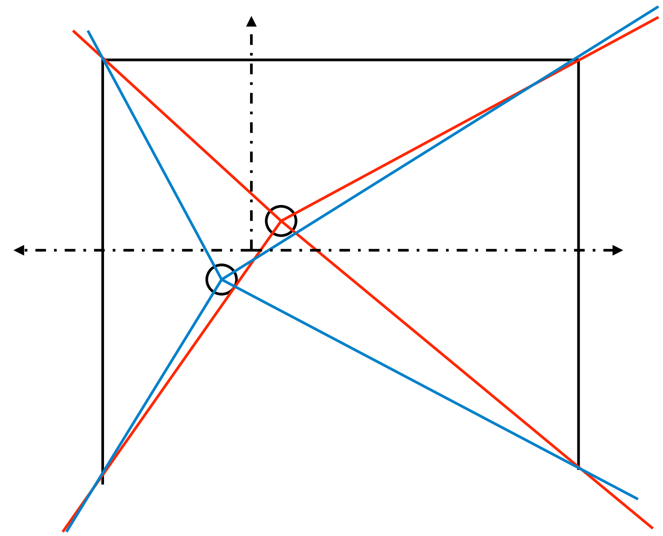

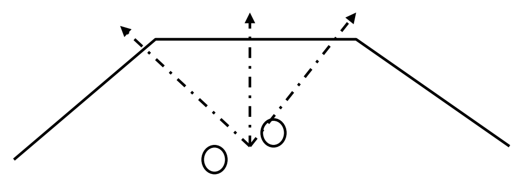

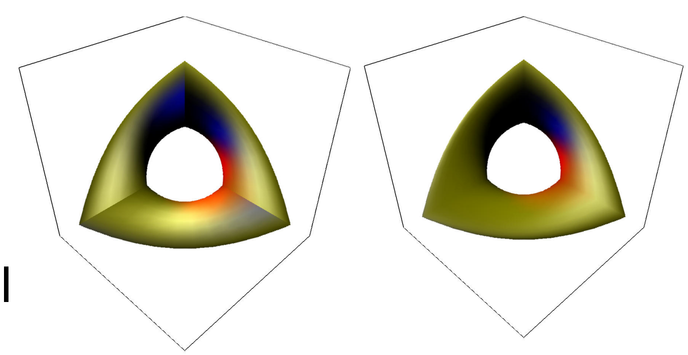

- Problem with specular lighting: an "*infinite (OpenGL) viewer*"
  - Reminder:

    $$specular : (s \cdot n)^{shiny}, s = \overline{VL} + (0,0,1)$$
  - Consequence: s "jumps" if object crosses projection walls!

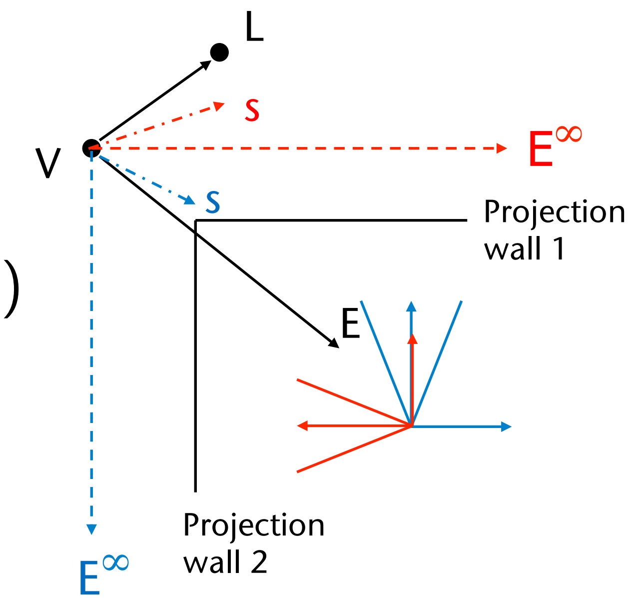

- Solutions:
  1. Use "*local viewer*" in OpenGL; then

     $$s = \overline{VL} + \overline{VE}$$

     remains consistent across walls (i.e., across projection matrices
  2. Or, multiply rotation matrix for each wall to the PROJECTION stack, instead of the MODELVIEW (which is the conventional way):

     $$\mathbf{v}_{screen} = \underbrace{M_{projection} R_{wall}}_{\text{GL\_PROJECTION}} \underbrace{M_{viewpoint} M_{world}}_{\text{GL\_MODELVIEW}} \mathbf{v}_{object}$$

The Model of a User's Head

M_e = viewpoint transformation

M_s = current sensor reading, relative to ist zero calibration

M_{rs} = transform. from head'srotational center to sensor

M_{er} = transform. from "cyclop's eye" to head'srotational center

$T^l | T^r$ = translation to left/right eye

$$\mathbf{M}_e = \mathbf{T}_{l|r}\mathbf{M}_{er}\mathbf{M}_{rs}\mathbf{M}_s$$

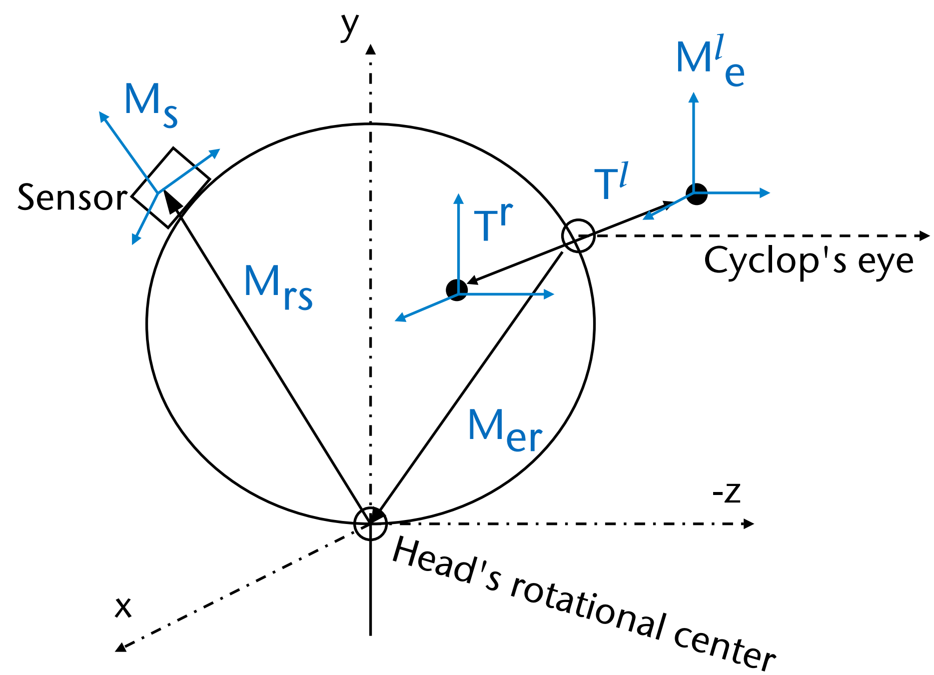

Relevant OpenGL Commands

- Initialization:

```
glutInitDisplayMode(GLUT_DOUBLE|GLUT_RGB|GLUT_DEPTH |
GLUT_STEREO);
```

- Rendering:

```
glDrawBuffer(GL_BACK_LEFT);
glClear(..)
glFrustum(...)
traverse scene graph ..
glDrawBuffer(GL_BACK_RIGHT);
glClear(..)
glFrustum(...)
traverse scene graph ..
```

- Or: render with 2 different threads into 2 different graphics cards
- Or: side-by-side stereo (2 openGL viewports in one big window)

Guidelines for Stereo Rendering

1. Make parallax not too big! (common error of novices)
   ±1.6° ~ parallax ≤ 0.03 · (distance to projection wall)
2. Single object → put zero-parallax plane at its center
3. Complete VE → 1/3 negative parallax, 2/3 positive parallax
4. Keep objects with negative parallax away from the border of the projection surface

Rendering on HMDs needs to be un-distorted

- Optics in HMDs usually cause some amount of distortion
- Idea: pre-distortion (using a texture)

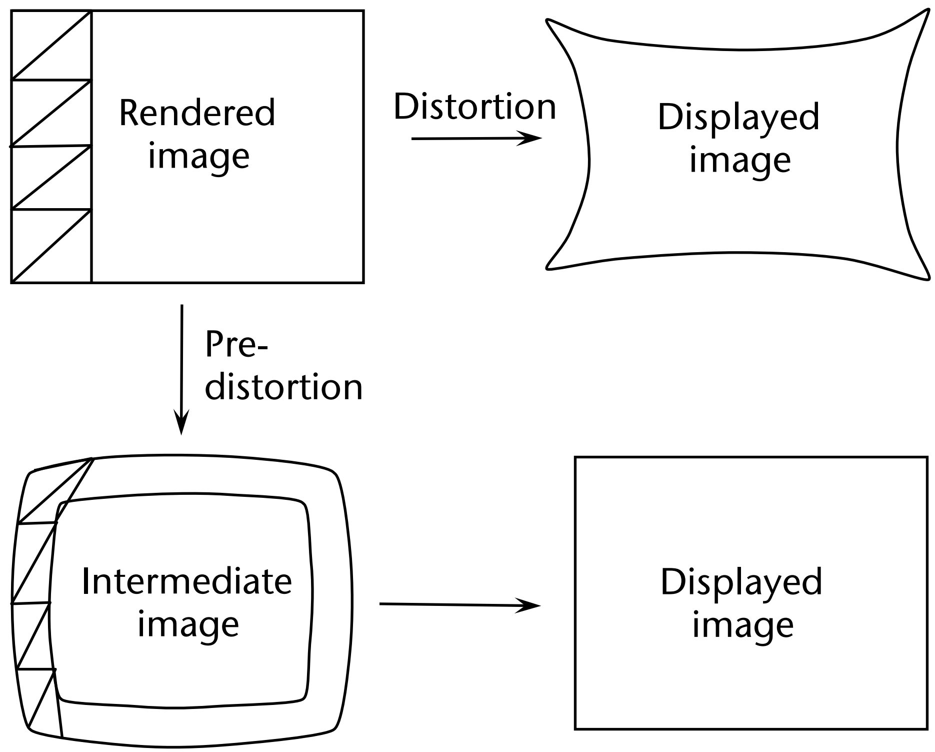

System Overview: the Rendering Loop

- 1 processor → everything serial
- 2 processors → app and renderer in parallel
- Stereo → 2 render processes (3 proc)
- In general: *n* walls in stereo →
  - ≥ *n+1* processors, *n* render processes
  - Better *2n+1* processors

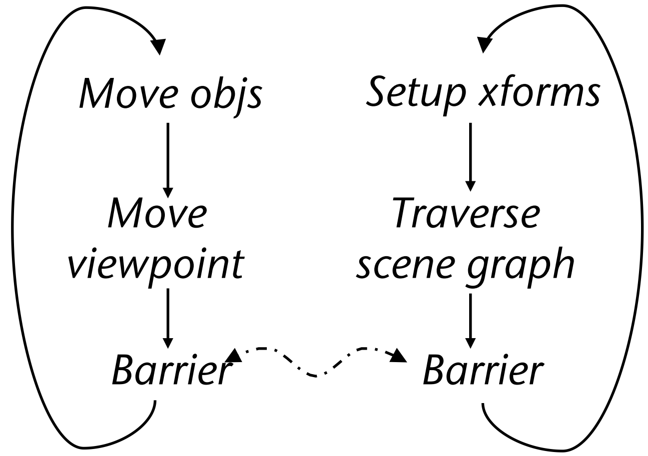

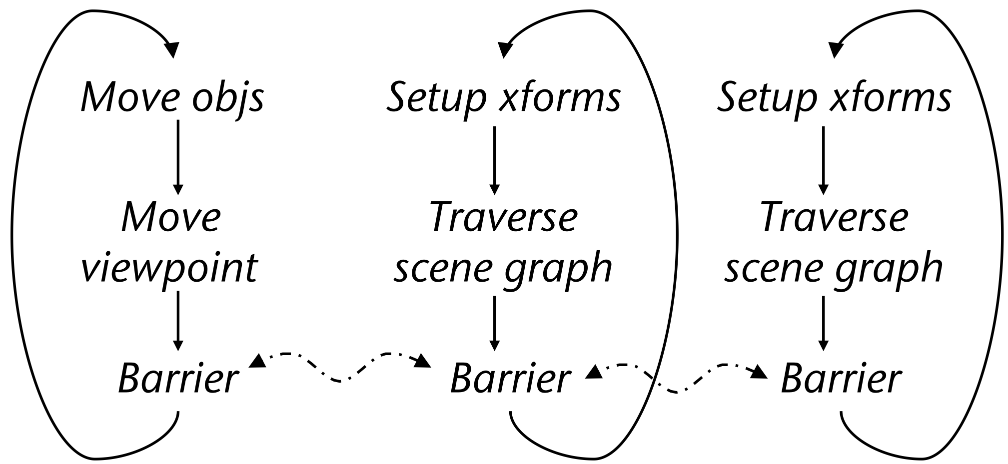

Crosstalk (Ghosting)

- Crosstalk = if one eye (also) sees the image meant for the other eye
  - Amount of crosstalk depends on technology

Simulator Sickness

- Simulator sickness = more or less of the following symptons (can sometimes occur with prolonged stay in flight simulators / virtual environments):
  - Nausea (including vomiting), eye strain, dizziness, drowsiness, blurred vision, headache, fatigue
- Cause is not entirely clear
- Common hypothesis: inconsistent sensory input to brain (e.g., mismatch between vision and vestibular organ (organ of equilibrium))
  - E.g., when staying below deck for a prolonged time
  - With latency between motion of platform and rendering in flight simulators
- Frequency: 20-40% with jet pilots
  - Occurs more frequently with experienced pilots than novices [sic]
- Other observations with mismatching sensory inputs:
  - In a rotating field when walking forward, people tilt their heads and feel like they are rotating in the opposite direction
  - If a person is walking on a treadmill holding onto a stationary bar and you change the rate the visuals are passing by, it will feel to the person like the bar is pushing or pulling on their hands