

Bremen



Virtual Reality & Physically-Based Simulation


VR Displays,
Stereo Rendering,
Display Issues



G. Zachmann

University of Bremen, Germany

cgvr.cs.uni-bremen.de





Depth Cues

- Motion parallax: apparent motion of objects relative to each other, when observer moves
- Occlusion
- Stereopsis (binocular/stereo vision)
 - Important, but not the most important depth cue
- Accommodation & convergence
- Defocus blur (a.k.a. **blur gradient**)
- Perspective (see CG1)
- Lighting & shading
- Relative size / familiar size
- Texture gradient



- **Focus** = adjustment of the eyes' lenses to adapt to different distances
 - So that the fixated object appears sharp on the retina
 - A.k.a. **accomodation**
- **Convergence** = counter-rotating 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 appears on the center of the retina (has highest resolution)
 - A.k.a. just **vergence**
- **Stereopsis** = "vision with two eyes"
 - The mechanism in human vision for sensing depth




- Stereo blindness: ~10% of general population
- Some people can actually turn their eyes to *divergence*:



- Convergence causes **disparity δ** between corresponding points on the retinas:

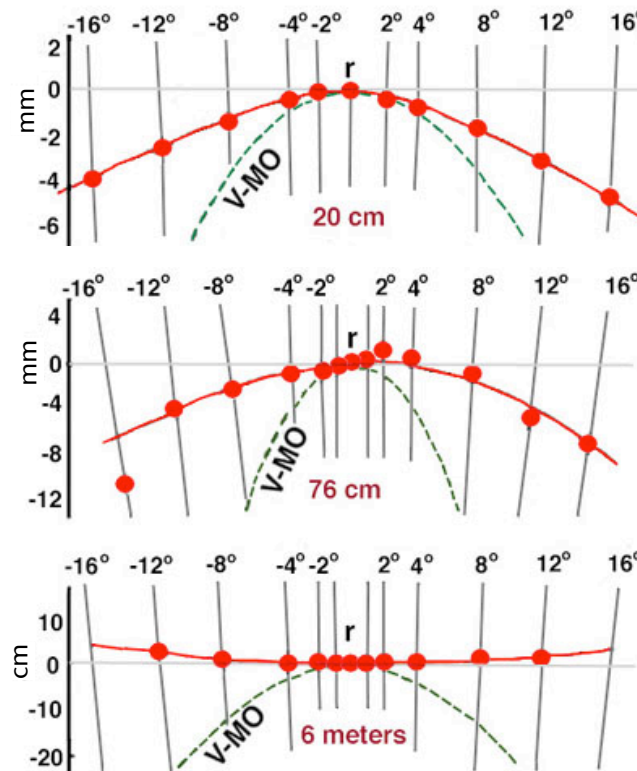

$$\delta = \delta_2 - \delta_1 = \gamma - \alpha$$




- Horopter** = locus of points in space with same apparent depth as the fixated object = points with 0-disparity

The Shape of the Horopter

- Mathematical construction → 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

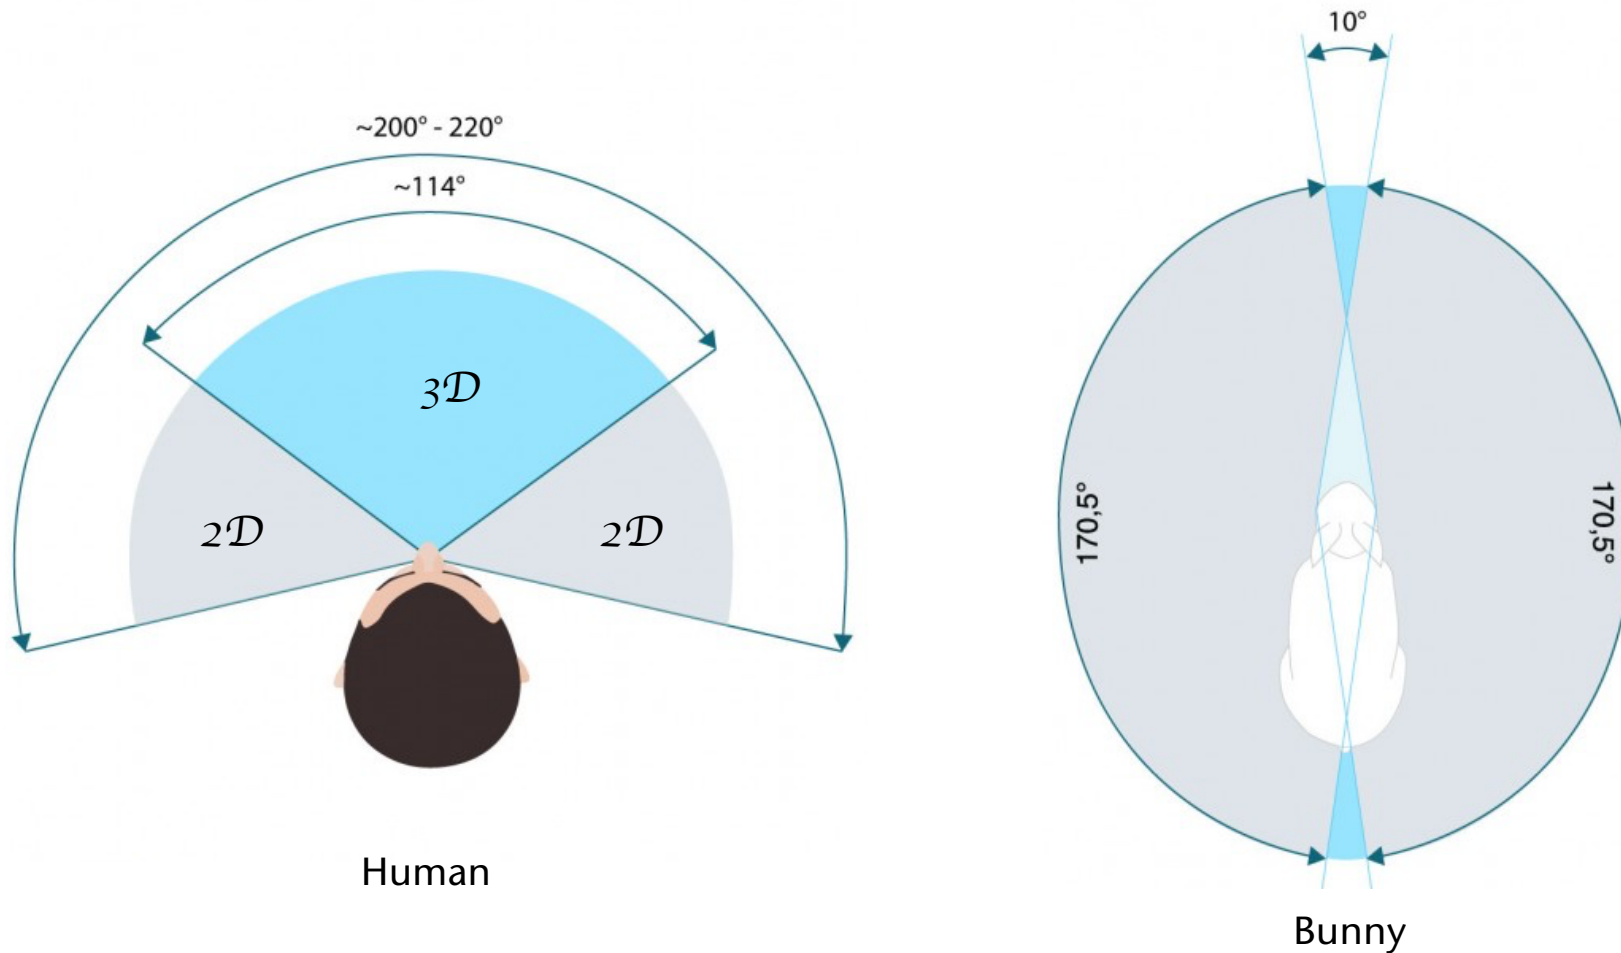


- There is a zone/range of depth around the horopter, where the brain is able to fuse the double image of an object
→ Panum's Area of Fusion




Limitation of Human Stereopsis

- Stereoscopic vision works just up to a few meters (< 6 m, ca.)
- Does not work in the left & right periphery:



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?

- One channel, two senders & receivers → 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:

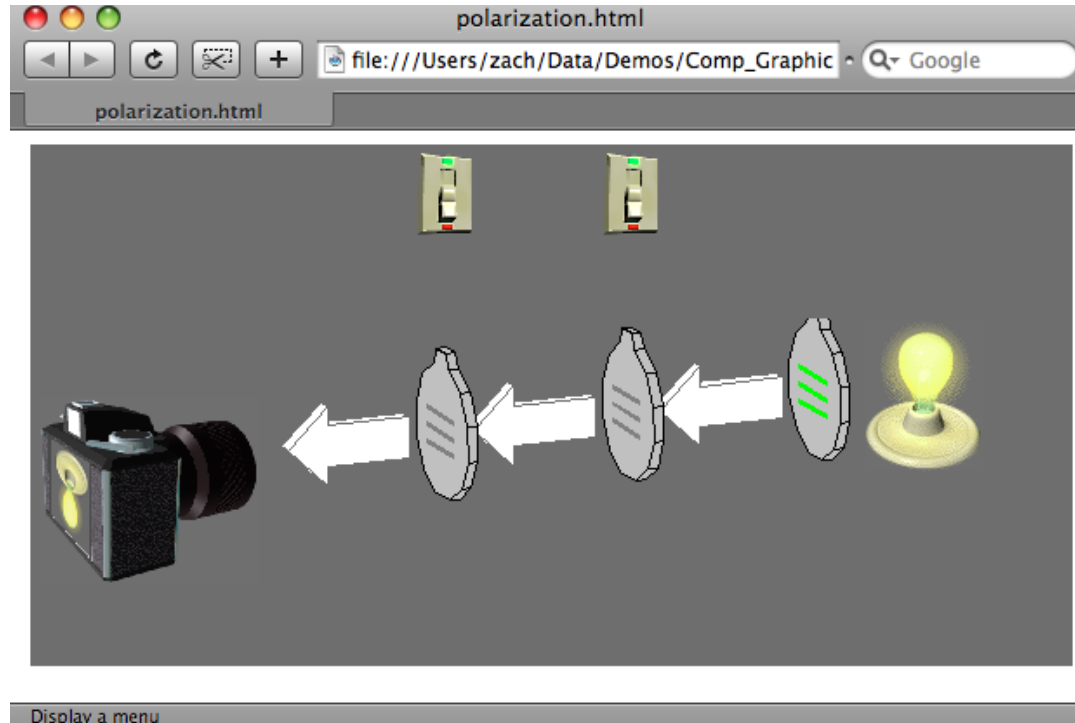
- Any direction perpendicular to direction of travel of light



2. Circular polarization:

- Left-handed / right-handed polarization






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


"Color Multiplexing"

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









- Monochrome images:
 - Render left & right images
 - Convert to grayscale \rightarrow L, R
 - Merge into red & cyan anaglyph image $I(r,g,b)$ by assigning

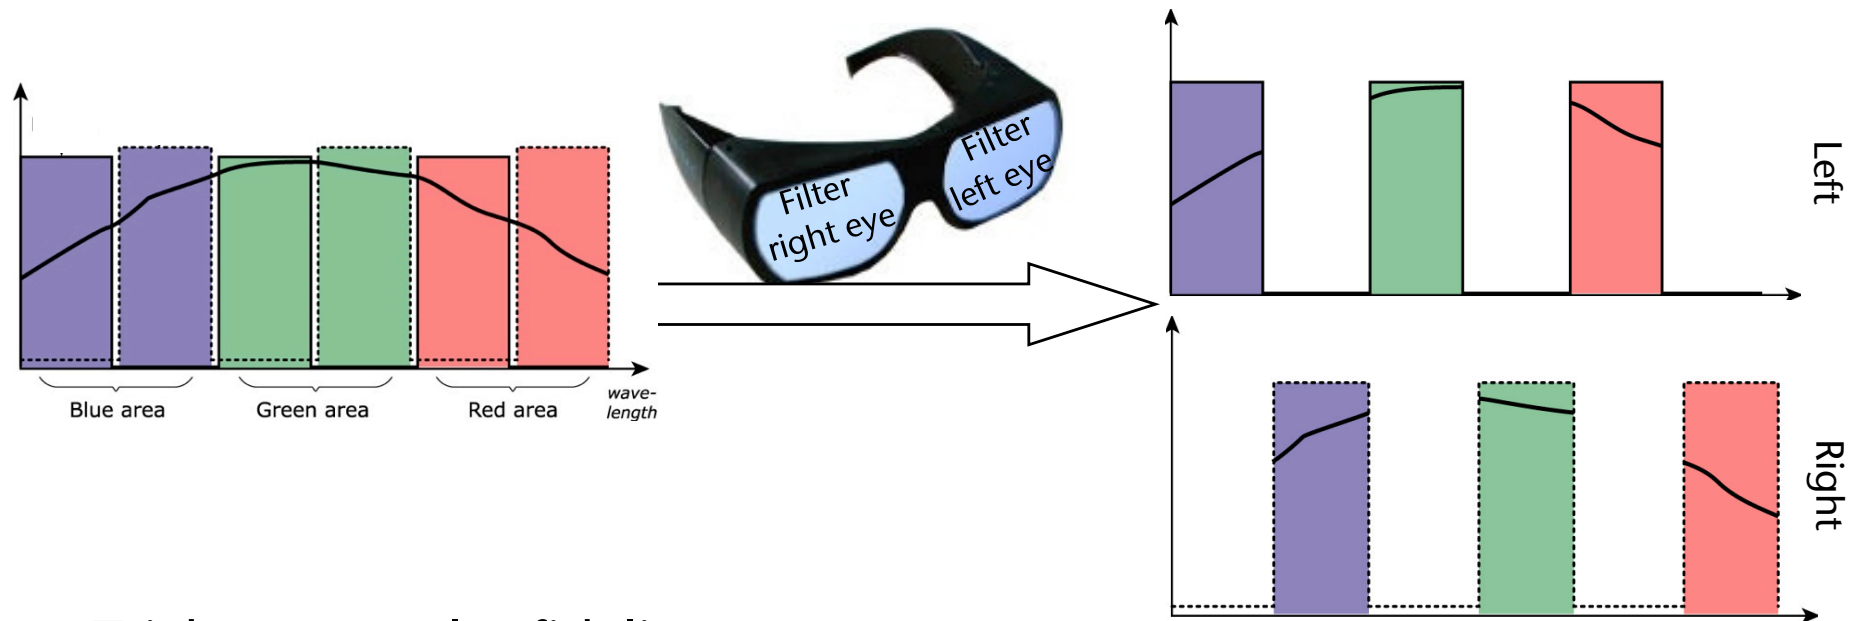
$$I(r) = L, \quad I(g, b) = R$$

- Full color anaglyph images:
 - Render left & right images, but do not convert to grayscale \rightarrow L, R
 - Merge into red & cyan anaglyph image:

$$I(r) = L(r), \quad I(g, b) = R(g, b)$$

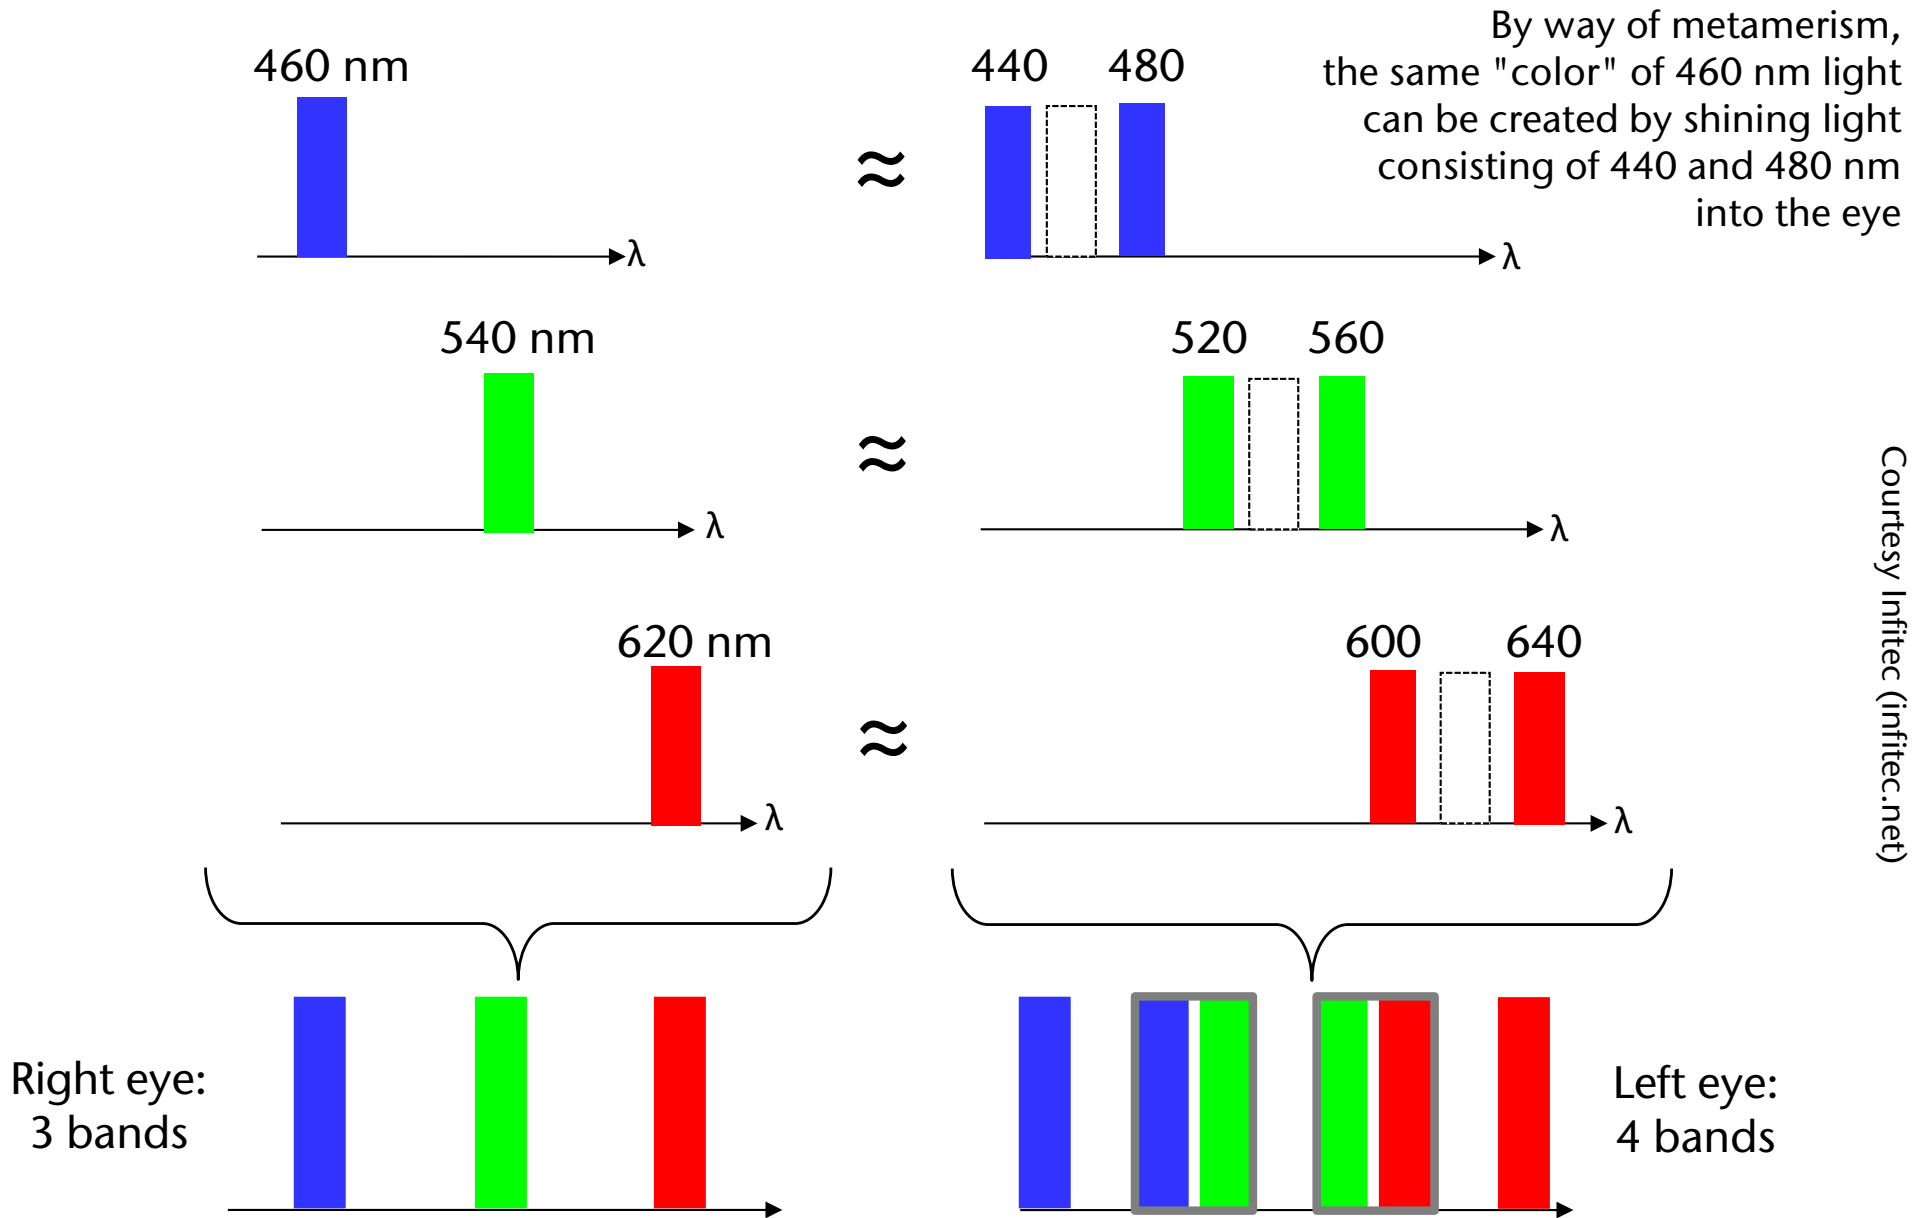
Multiplexing by Wavelength (Infitec)

- Generalization of anaglyph stereo:
 - Partition whole spectrum into 6 (narrow) bands
 - Left & right eye get filters with interleaving band passes
 - Other names: Dolby3D, spectral comb filter



- Tricky part: color fidelity

Improvement: Utilize Color Metamerism



Autostereogram (Single Image Stereogram)

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



Underlying "depth image"

- The Pulfrich effect:
 - Discovered by Carl Pulfrich, German physicist, 1922
 - Dark stimulus in the eye arrives later in the brain than a bright stimulus
- Viewing instructions: put sunglasses or similar darkening filter over *one* eye, the other eye remains naked



http://www.youtube.com/watch?v=1mnWl_u_zBg

- Head-Mounted Displays (HMDs)
 - Head-Coupled Displays (HCDs)
 - Immersive projection displays (IPDs)
 - Autostereo Monitor
 - Desktop setups
 - E.g. Autostereo monitors, zSpace, or "reach-in"
 - "Powerwall"
 - Workbench
 - Cave
 - "Exotic" displays:
 - Retinal displays
 - Holographic displays
 - ...
- A.k.a. **World-Fixed Displays**

Stereo Monitor

- 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)
 - Very little immersion
 - Very limited working volume
 - "*Stereo frame violation*" is very common





52" Autostereo Display



zSpace

- Interesting things you can do with a simple monitor: the "Reach-in idea"



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

Head-Mounted Displays (HMD)

- First "true" VR display
- Technologies / characteristics:
 - HMDs using LCDs or OLEDs
 - Weight:
 - Small FoV → lightweight; large FoV → heavy
- Advantages:
 - Kind of a "surround display"
 - In theory, very good immersion
 - No *stereo frame violation*
 - Large working volume
 - Almost no special requirements on the working environment



Around 1992

Virtual Research



Around 1984

Other Models (as of 2017)



Oculus VR / Facebook



HTC Vive



Sony's PlayStation VR



"Sword of Damocles" (1965)




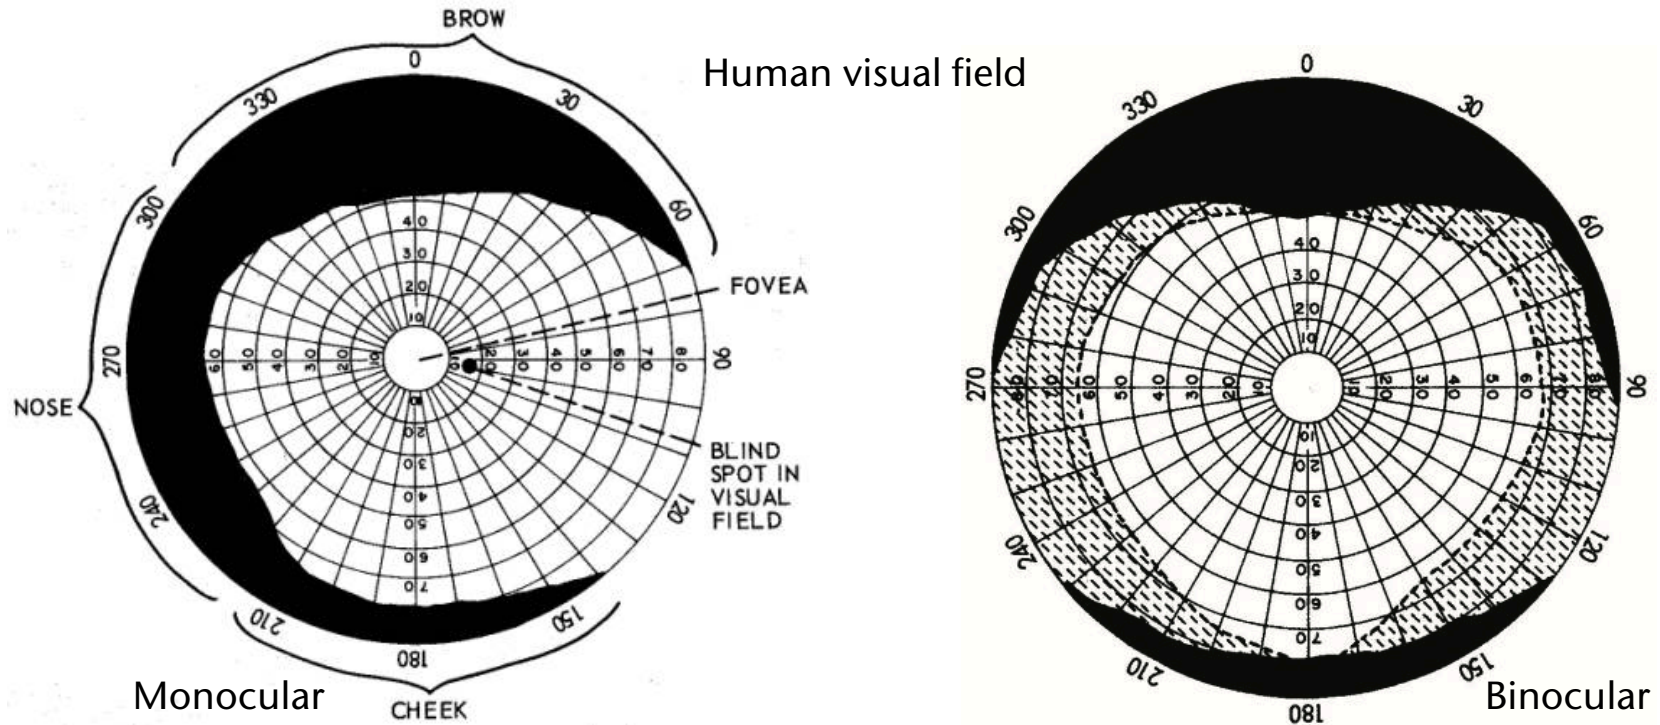
NEO VR by Immersion



Moon by Royole


- Uncomfortable when used for a prolonged time ("*invasive interface*")
- Distortions (can be corrected somewhat by pre-distortion)
- 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)

The Field-of-View Problem of HMDs



The Resolution Problem of HMDs

- Human visual acuity:
 - 1 photo receptor (cone) = 1 arc min = 1/60 degree
- Display needed for a "retina" HMD:
 - $150^\circ \times 135^\circ$ with $1/60^\circ$ resolution = 9000 x 8100 pixels per eye
- Challenges:
 - Bandwidth: moving the data at 60 Hz from GPU to display
 - Miniaturize display panels with 73 Mio pixels



- Potentials:
 - "Foveated rendering"
 - Requires end-to-end latency of < 10 ms
 - Control game using eye gaze direction
- Dynamically move the zero-parallax plane?
- Control focus depth for depth-of-field rendering?
- Make eye contact with virtual avatars (NPC)?
 - So they "notice" and look back at you
- Shoot enemies in games just by looking at them?




My Vision

- Wireless HMD with very wide field-of-view and SLAM-based tracking like HoloLens
- Can someone build that for me please?



Head Coupled Displays (HCD) – Out-Dated

- 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




- 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

Large-Screen Projection Walls (Powerwalls)

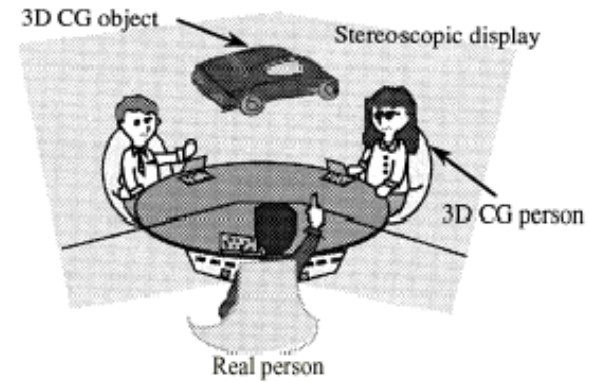



© Immersion


- "*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 IPDs"



Workbench, L-Shape, Holobench, etc.




Workbench



Holobench




Tilting workbench




Principle of the workbench



3-wall cave




Schematic of the arrangement of the mirrors



5-wall cave, FhG-IGD, Darmstadt



6-wall cave, Alborg, DK






5-sided CAVE at University RWTH Aachen



Disney Imagineering's DISH


RealityDeck - Immersive Giga-Pixel Display

- 308 x 30" LCD displays
- 2560x1600 resolution per display
- 1.5 Giga 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




- Usually, with wall-sized screens (curved or not), some kind of **edge blending** and color correction between projectors is necessary





Curved Screen made out of 3D-TVs

- Idea: construct the walls of a Cave / curved powerwall 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/>



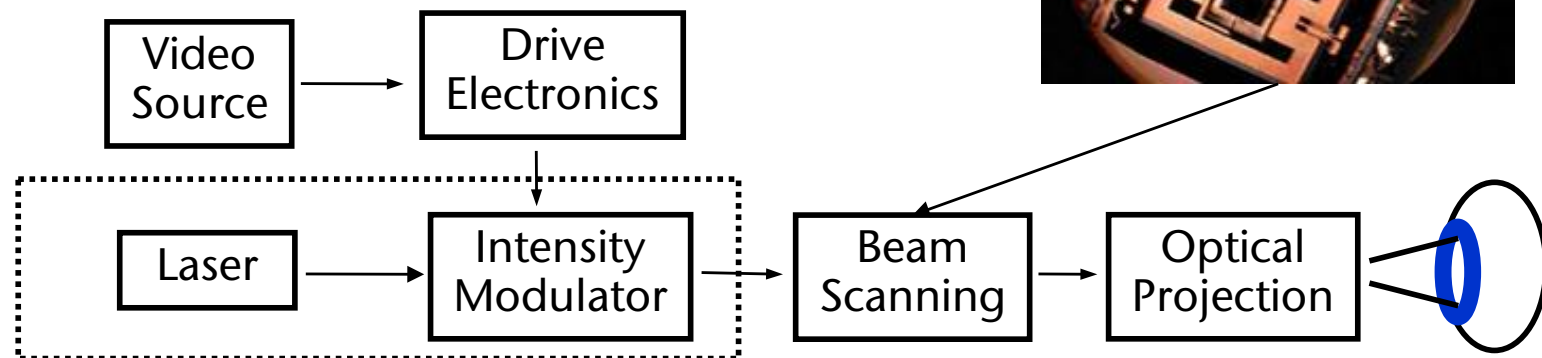

- A modern "Sensorama":



Immersa-Dome from Aardvark Applications

- Advantages:
 - Large resolution
 - Large *field-of-view*
 - "*Non-invasive*"
 - No isolation of the real world
 - (Can accommodate 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)

- 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, bad for VR
 - Small power consumption





Retinal display



Design study

Holographic / Volumetric / POV Displays

- Hologram = can reconstruct real 3-dimensional image
- Advantages:
 - Provide correct perspective/view from *every* angle!
 - Coherence between accommodation and convergence
 - Depth of field (Tiefen(un-)schärfe)
- Holographic displays: algorithmic computation of holograms
- Problems:
 - Staggering amount of computational work
 - Colors
- Volumetric displays: voxels are projected onto a rapidly rotating surface covering a volume
- A.k.a.: **Persistence of Vision** Displays
- Problems:
 - Size of data (e.g. 100 mega-voxels = 1000x1000x100 display resolution)
 - Occlusions?

- Example volumetric display:
 - $198 \times 768 \times 768 \approx 100$ million voxels
 - Frame rate: 20 Hz



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




FogScreen




DisplAir




- The "Janus" display of KAIST, Korea:
 - Utilizes [persistence of vision](#)
 - See-through display with touch interaction for collaboration
 - Each person on either side gets their own, possibly different image



- "Everywhere displays":




Recap: Perspective Projection in OpenGL



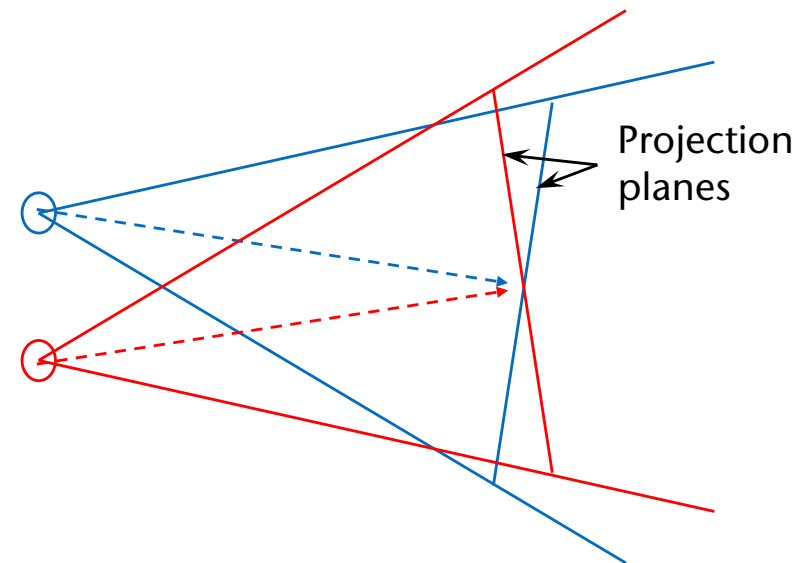

```
glFrustum( left, right,
           bottom, top,
           near, far );
```

Stereoscopic Projection

- Parallax on the screen
→ disparity in the eyes



- Wrong way: converging view vectors
 - Problem: vertical parallax!




Parallax Not Well Done



Correct Stereoscopic Projection


- Parallel viewing vectors
- Off-center perspective projection (a.k.a. "off-axis projection")
- Important stereo parameters:
 - IPD and ZPP



- Thought experiment: imagine a single line emanating from 1 m in front of you, away from you to infinity
 - What stereo image do you get?
 - What happens, if the IPD increases?
 - What happens, if you move the ZPP closer or further away?

Hypo- and Hyper-Stereo

- In monoscopic filming/display, cameras just have these parameters:
 - Field-of-View, focal length (film), ...
- In stereoscopic filming/rendering, cameras *in addition* have:
 - **Interaxial separation** (a.k.a. IPD)
 - Zero-parallax plane
- **Hypo-Stereo**: $\text{Interaxial} < \text{IPD} \rightarrow$ dwarfism effect
- **Hyper-Stereo**: $\text{Interaxial} > \text{IPD} \rightarrow$ gigantism effect
- Can make sense for macro/micro scenes



Interaxial Separation between lenses, a.k.a. Stereo Base, a.k.a. Interocular separation, (a.k.a. IPD for human eye)

Standard Stereo



Hypo Stereo and Dwarfism Effect




Hyper Stereo and Gigantism Effect



Computation of the Frustum

- Given: i = interpupillary distance $\div 2$,
 w/h = aspect ratio, α = horizontal FoV,
 n = near plane, z_0 = zero-parallax depth




- Task: determine **left/right/top/bottom** for `glFrustum()`
- Assumption (for now): no head tracking \rightarrow cyclop's eye is in front of the center of the viewport
- Example:
 compute **left** for left eye


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


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

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



Video: Stereo Projection in the Analog World






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

Convergence-Focus Conflict

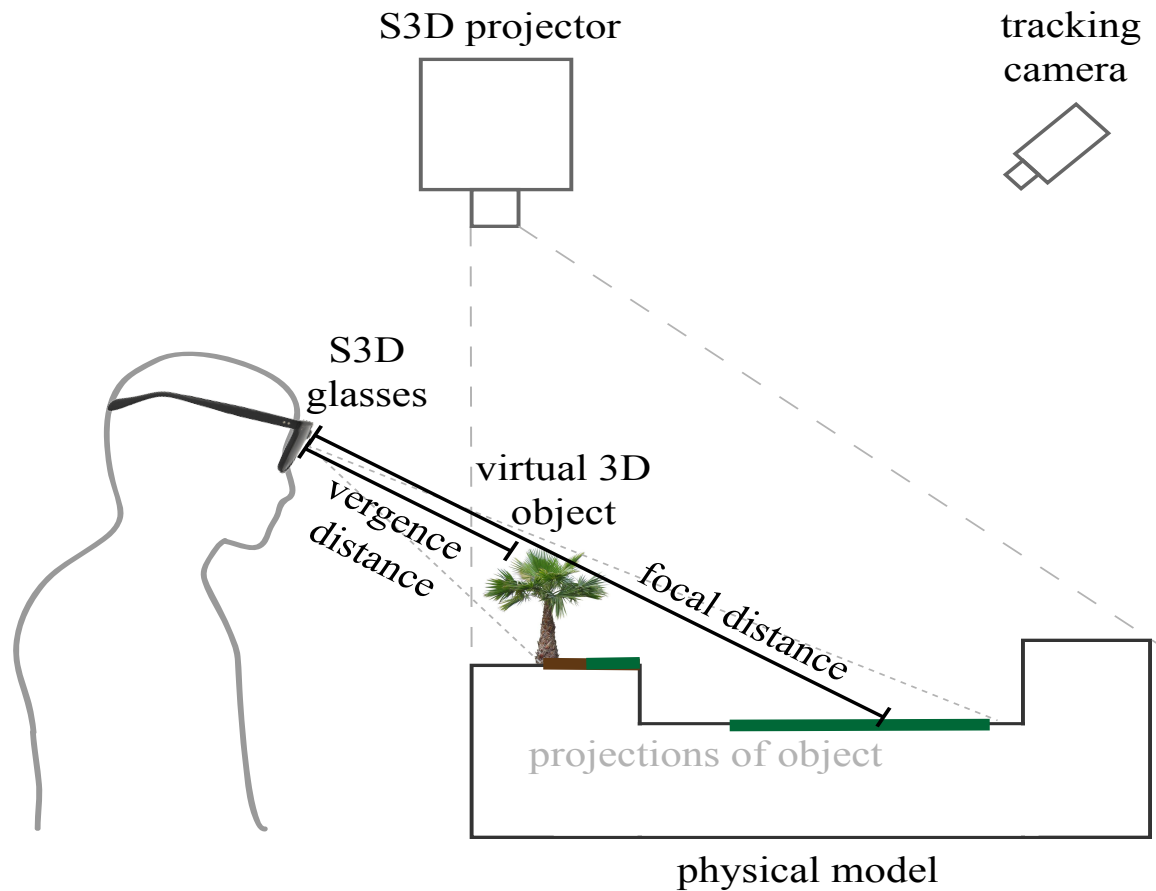
- Experimental evidence shows: the brain computes a weighted average of multiple depth cues, including focal depth
- With stereoscopic displays, our eyes receive inconsistent depth cues:



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

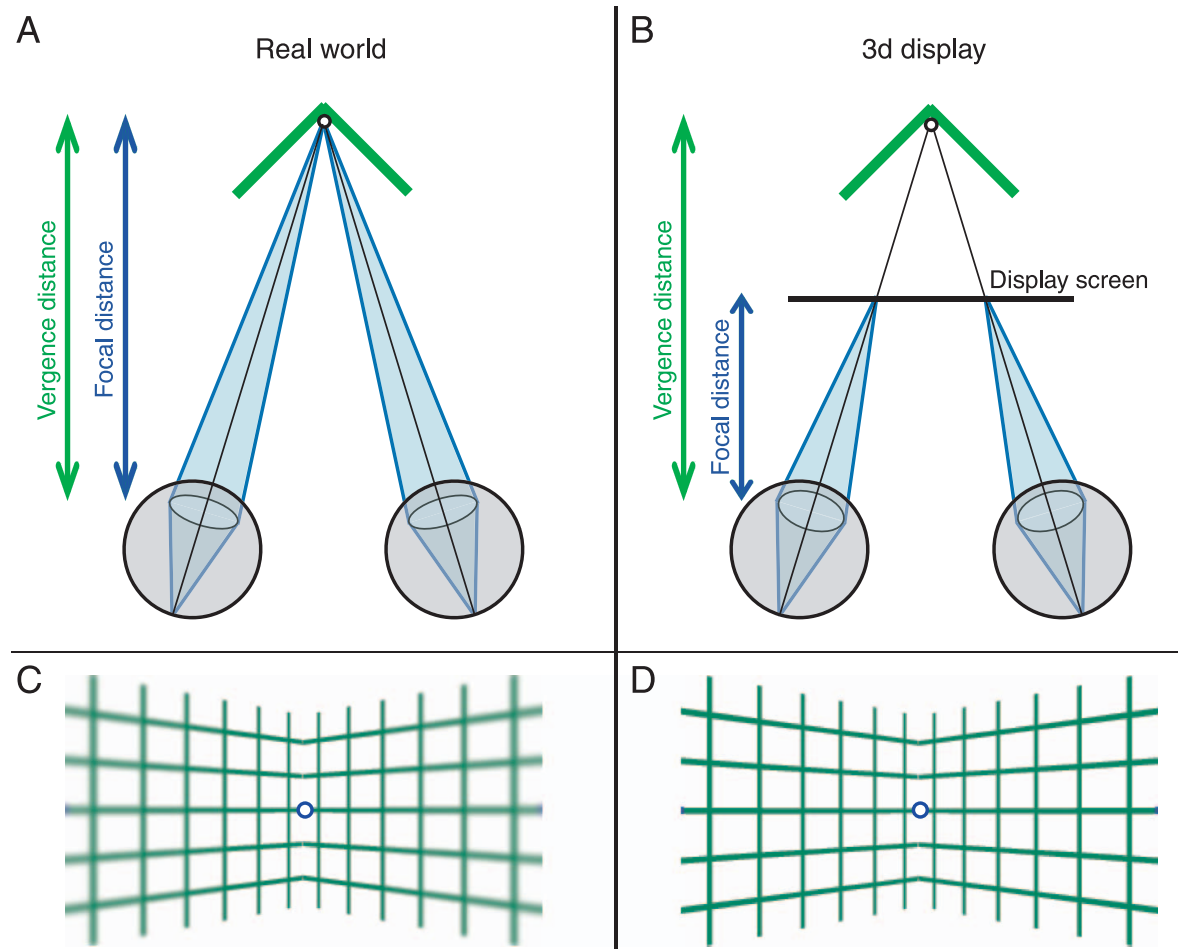
Watt, Akeley, Ernst, Banks: "Focus cues affect perceived depth", J. of Vision, 2005]

- This problem is potentially aggravated in projection-based AR, even if geometric correction by eye tracking is done




- Another depth cue: blur
 - The eye (brain) can estimate (relative) depth from the amount of blur

- If no depth-of-field is being rendered, then our eyes perceive different depth cues:



Example Stereogram

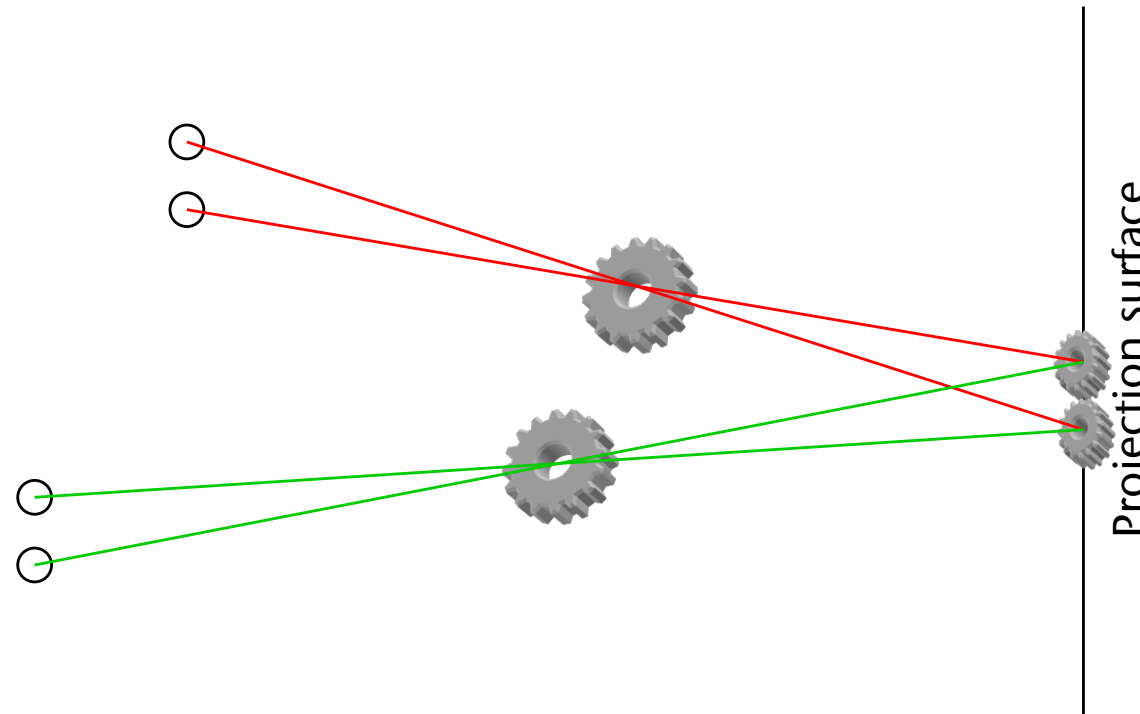
- The following image appears to be 3-dimensional, if you can decouple focus (=accomodation) and convergence (you have to scale the slides so that the statues are about 5-7 cm apart)



Postcard from 1868

Stereo is (Usually) a "One Man Show"

- Why are stereoscopic images correct only for 1 viewpoint?




- One problem: images (e.g., on a powerwall) shift and move for the un-tracked user when the tracked user moves


Further Problematic Case: the Cave

- Pertains to segmented curved screens, in general:

User's eye is different from virtual camera




User's eye matches virtual camera perfectly





Also, users next to the tracked one will see this

- Two effects that can occur together:
 - *Clipping*
 - Depth from stereoscopic image
- Object is clipped, although apparently *in front of* the projection surface!
- Consequence: conflicting *depth cues*
 → stereo violation (a.k.a. window violation)
- Example: lower left corner of the anaglyph mars image




Too Much Parallax

- Assume you created a stereo image for a small desktop display
- Then, you run the app on a big screen:



... then you decide to run your application on a BIG display!



Same can happen if you sit too close in a 3D movie



Guidelines for Stereo Rendering

1. Make parallax not too big! (common error of novices)
 $\pm 1.6^\circ \sim \text{parallax} \leq 0.03 \cdot (\text{distance to projection wall})$
2. Single object \rightarrow put zero-parallax plane at its center
3. Complete VE \rightarrow 1/3 negative parallax, 2/3 positive parallax
4. Keep objects with negative parallax away from the border of the projection surface

The Model of a User's Head for Precise Viewpoint Tracking

M_e = viewpoint transformation


M_s = current sensor reading, relative to its zero calibration

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

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

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


$$M_e = T_{l|r} M_{er} M_{rs} M_s$$



Coherent Virtual Workspaces

- Assume the situation: one stereo display wall, several users in front of it
- Problem with a single-tracked projection (stereo or mono) and multiple users: only the viewpoint of the *tracked* user is correct, only she will see a correct image!
- One of the problems: communication via pointing fails

Image's perspective is correct for the user




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

Benefit of Correct Projection for All Users

- With *perspectively correct* projections for *all co-located* users, the shared 3D space will become **coherent** for all users

- Consequence: direct communication (including *pointing!*) in **co-located collaborative virtual environment** is possible
 - Note: 80% of all human communication is non-verbal





Solution: Correct (Stereo) Projection for Multiple Users

- Probably only possible for a small number of users
- *Temporally multiplexed* (shutter glasses):
 - Framerate for multi-user stereo = framerate for mono $\times 2 \times \text{\#User}$
 - Light intensity reaching each eye gets extremely low
- Infitec for several users:
 - Each user gets glasses with slightly shifted comb filters
 - With n users we need $2n$ different comb filters \rightarrow extremely narrow bands, $2n$ projectors needed
- *Spatially multiplexed*
- Combination of the above

Spatial Multiplexing

- Projection 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

- Combination of active and passive stereo, plus ingenious utilization of field-sequential projectors

- Recap from CG1:


field-sequential RGB with DLPs



1. Modification: remove color wheel

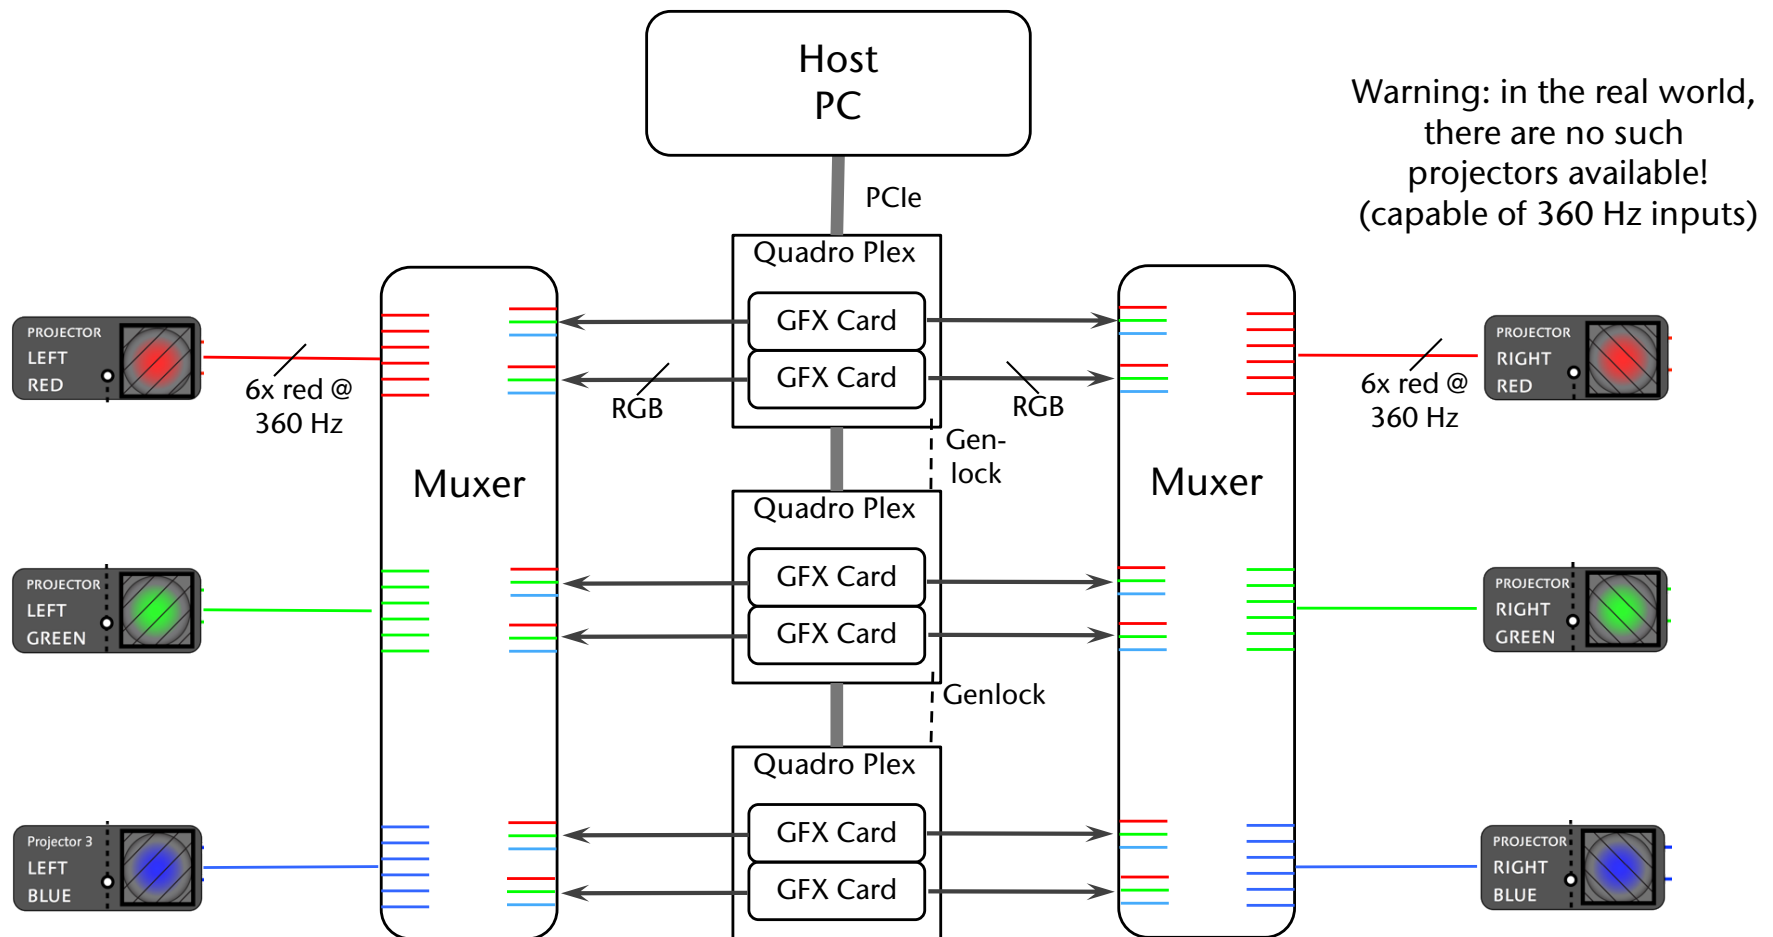
2. Modification: each user gets shutter glasses that additionally has left/right polarization filters

- Must be fast enough to prevent cross-talk!

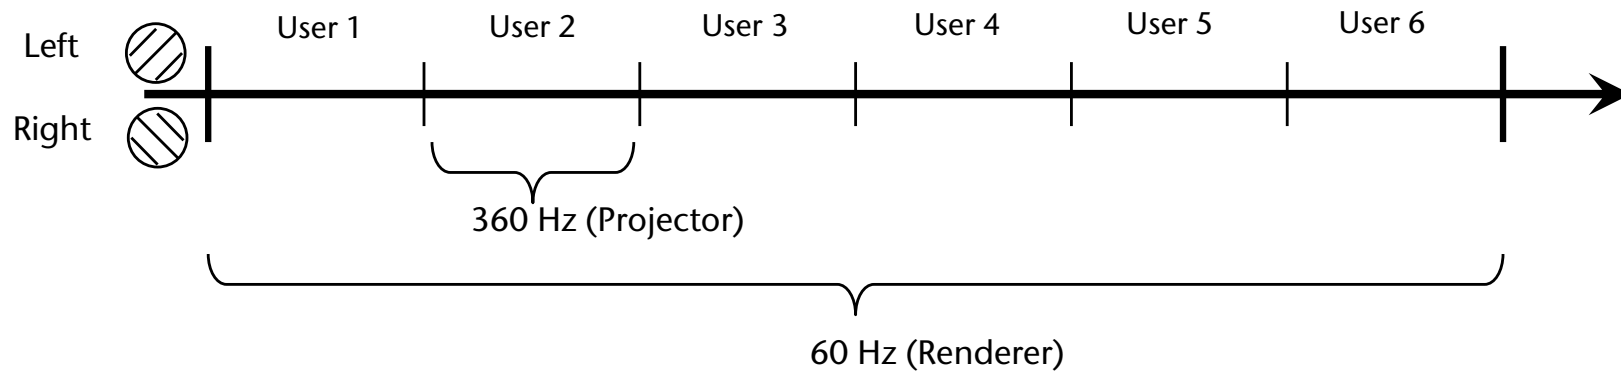


The Hardware in Principle

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



- Timing:



- Demo application:






Agrawala et al. 1997



Arthur et al. 1998





Kitamura et al. 2001



Agócs et al. 2006


- **Workspace Awareness** = "up-to-the-moment understanding of the other person's interaction with the shared workspace" [Gutwin & Greenberg, 2002]
- **Factors / questions:**
 - Who is participating / interacting? (People)
 - What are they doing ? What will they be doing next? (Actions / Intentions)
 - What can they see? Where can they have effects? (Perception / Influence)

- 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





Automultiscopic Display

- Like a lightfield / holographic display, but views/images differ only along horizontal viewpoint changes



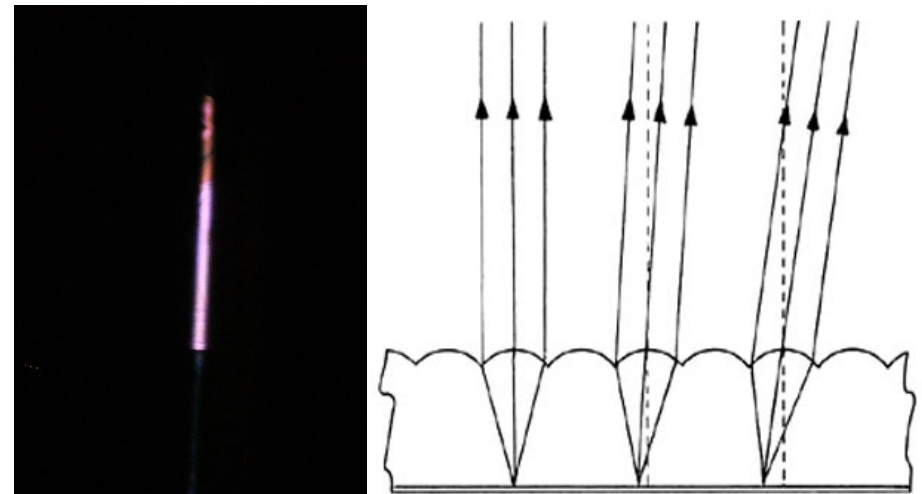
- Special screen sends images from projectors *only* in **one direction** with a very small scattering angle (1°)




Screen = anisotropic light shaping diffuser; scatters light vertically, maintains narrow horizontal blur

1° degree horizontal blur

- If number of cameras < number of projectors → video streams for "in-between" projectors must be interpolated from neighboring streams
- Bandwidth: $1920 \times 1080 \times 24 \text{ bits} \times 60 \text{ FPS} \times 216 \text{ cams} = 80 \text{ GB/sec}$
- Synchronization between all GPUs (swapbuffers) and all projectors (VSYNC)
- Lenticular screen with small horizontal diffusion angle:
 - From a specific viewing direction, the light from a single projector appears as a single stripe of light


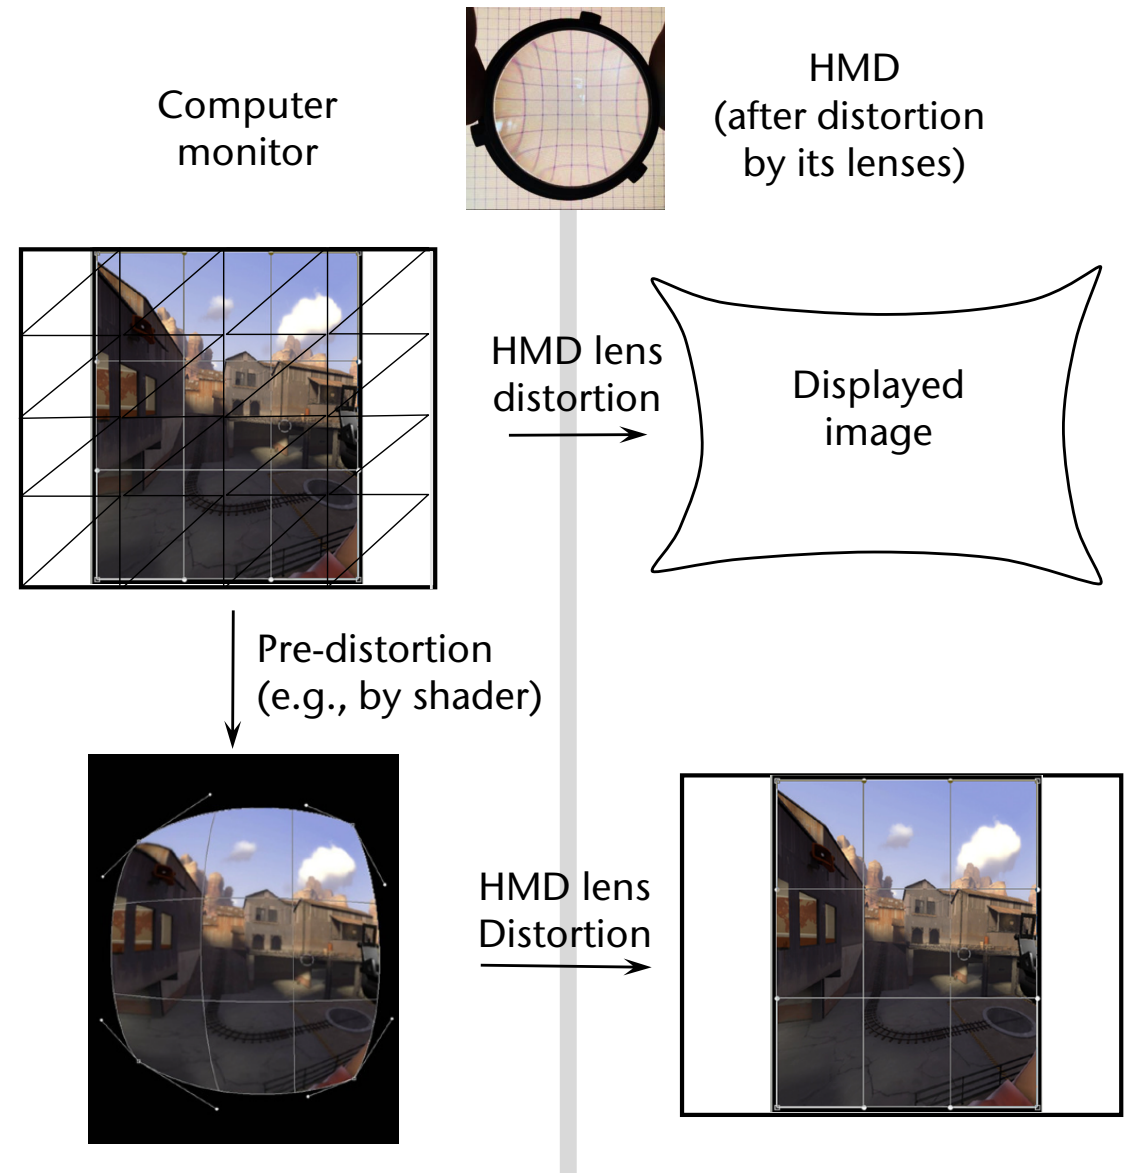




- Advantage: unlimited number of viewers
- Disadvantages:
 - Expensive (lots of projectors), and needs lots of space
 - Does not work with tilted heads (eyes must be aligned with the lenticular lenses)

Rendering on HMDs

- Optics in HMDs usually cause some amount of distortion
 - Especially the Oculus Rift
- Idea: **pre-distortion** (using multi-pass and texturing or shaders)



Optional


One of the Hard Requirements for VR / AR

- Images **must appear fixed** in space, no matter how users move



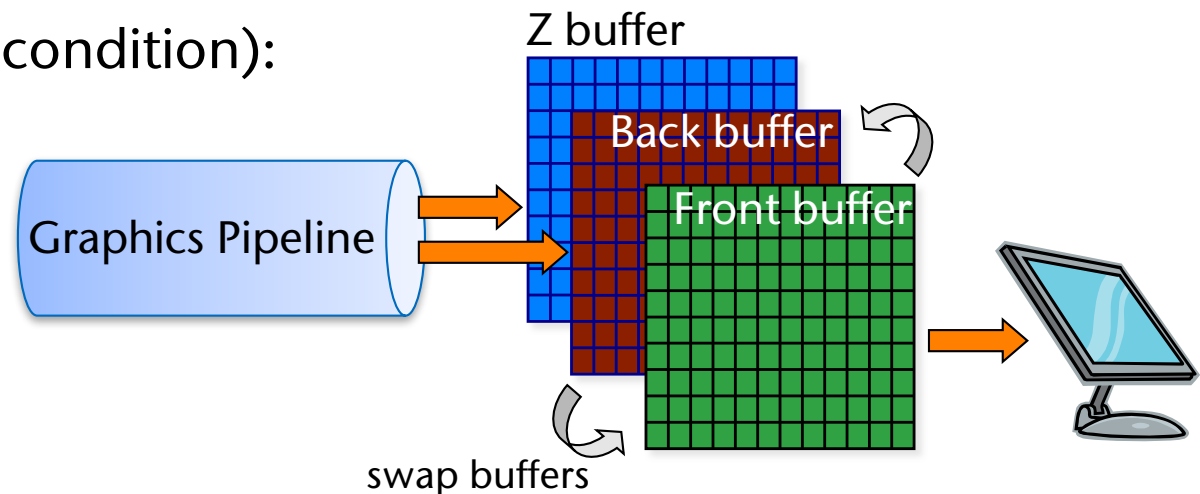
- This is hard even for a still environment!
 - Reason 1: latency (later)
 - Reason 2: display persistence (in the following)

- Current displays are always raster displays:



- Double buffering to prevent flickering (i.e., a race condition):



- Best point to swap the buffers?
- VSYNC




Low Persistence vs. Full Persistence



- Definition: **persistence** (in displays) = length of time that a pixel on a display remains emitting light after it has been switched on / energized
- Persistence in
 - CRT's: phosphor gets energized by electron beam, illuminates, then decays → **low persistence** / short persistence
 - LCD's & LED's: pixel can be illuminated virtually infinitely → **full persistence** / long persistence


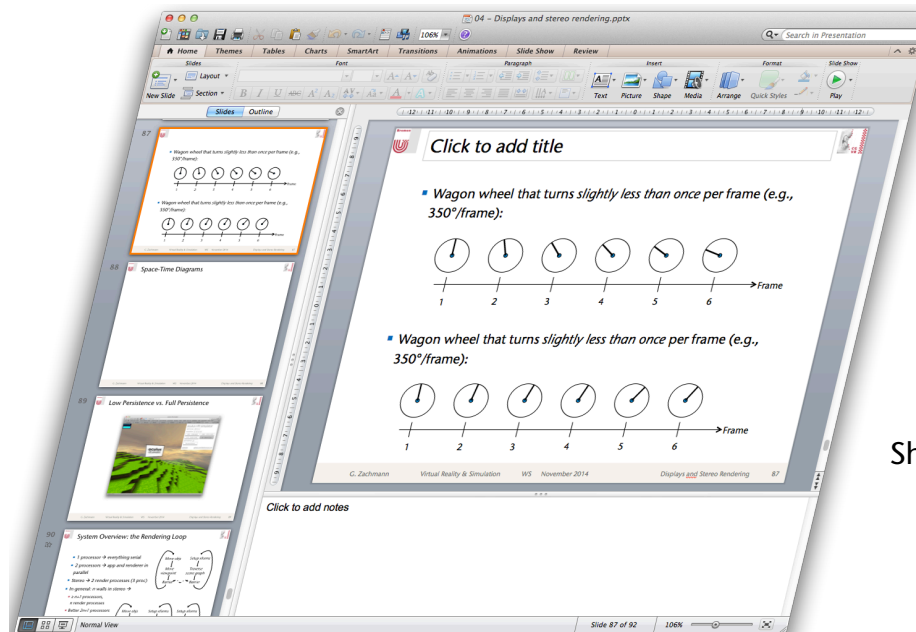




High-speed video in slow-motion, comparing an LCD and a CRT display

Optional A Simple and Quick Experiment

- Grab a window with high-contrast borders with the mouse and drag it left and right with medium speed; with your eyes, follow one of the vertical borders (a.k.a. **eye tracking**)
- What (shape) do you see?




Direction of motion of the window

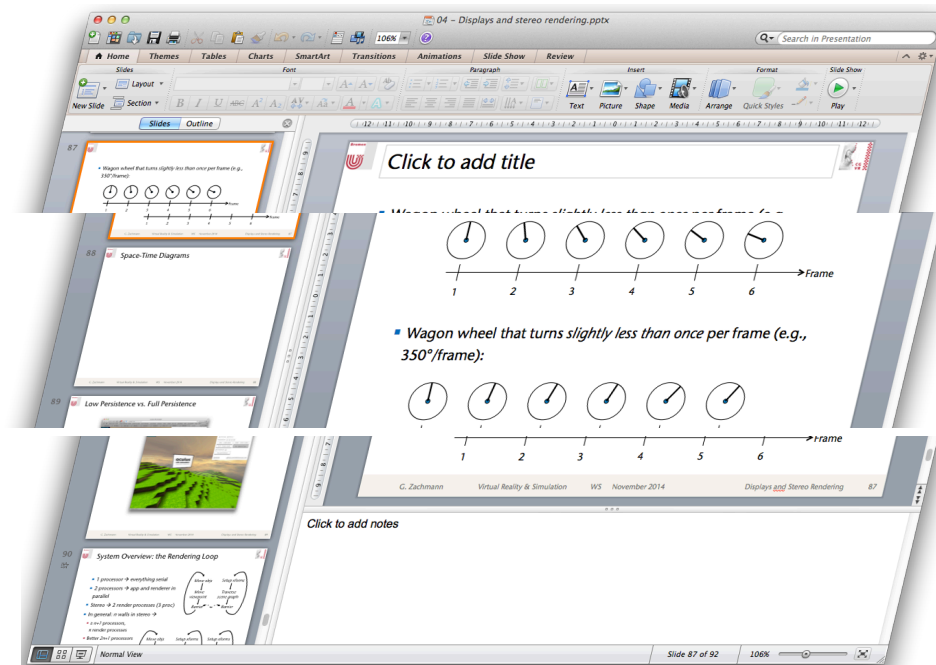
Shape exaggerated

- How can you explain this?

- Explanation:
 - Assumptions, for sake of simplicity
 - Monitor is a CRT (for LCD, the argument works, too, but a bit more complicated, at least with full-persistence)
 - Graphics hardware waits for VSYNC before scanning out framebuffer
 - The eye's fixation line moves with constant speed across display
 - Because scanlines are displayed one after another, pixels with same screen x coordinate are projected onto positions on the retina with different retinal x position!




- If your graphics hardware does *not* wait for VSYNC, you might see something like this: *tearing & shearing*




The Wagon-Wheel Effect (Temporal Aliasing)




- Wagon wheel with a rotation that is *slow relative* to the FPS:



- Wagon wheel that turns *once per frame*:




The Wagon-Wheel Effect (Temporal Aliasing)




S.a. "Wagon-Wheel Effect" in Optische Täuschungen

- Wagon wheel that turns *slightly less than once* per frame (e.g., $350^\circ/\text{frame}$):



- Wagon wheel that turns *slightly faster than once* per frame (e.g., $370^\circ/\text{frame}$):



Digression: Wagon Wheel Effect for Slow-Mo with Regular Camera



- Experiment setup:
 - Strobe light at $60 \pm$ Hz
 - Water droplets coming out of faucet at 60 Hz
 - Regular camera at 60 Hz (with very short shutter open period)



Space-Time Diagrams





- Space-time diagram = graph showing positions of objects as a function of time
 - In general, they are 4-dimensional
- Example: 3D space-time diagram of a journey on a 2D map
- Simplification in the following: consider only the x-position of objects → 2D space-time diagrams
- Example: a point staying still on the x-axis





- Example: a point moving with constant speed along x

- A point in VE being moved steadily by a simulation along x with constant speed, being rendered on a monitor with full persistence
 - Remember: "sample-and-hold" display

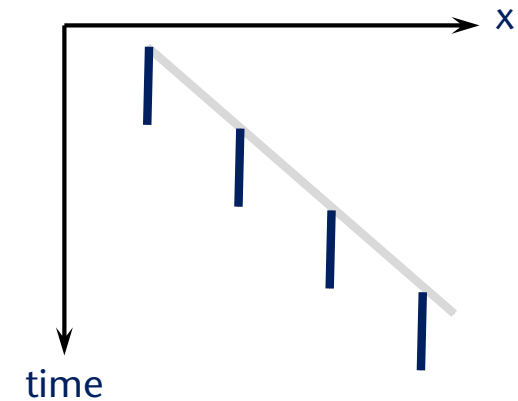


- Same as before, except *half persistence* display
 - A point being moved steadily by a simulation along x with constant speed, being rendered on a display


- Same again on *low persistence* display
 - E.g. CRT's, laser displays
 - LCD's and OLED's can be turned into low persistence displays (reduces brightness significantly)





- Recap: same experiment (moving point)



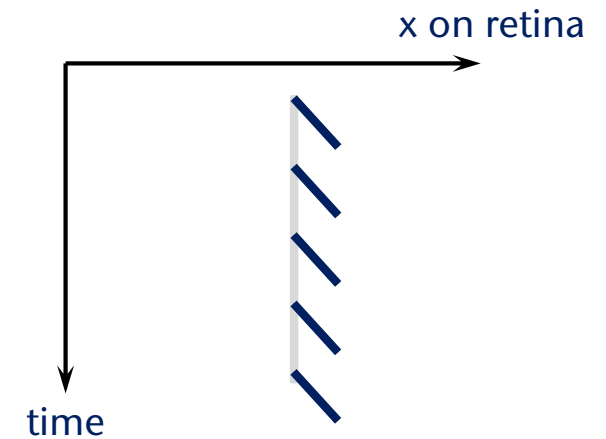
- Consider a slight change:
 - Point is moving in a VE along x at constant speed
 - Full-persistence monitor renders it at n FPS
 - Eye is *tracking* the virtual point (i.e. following its position)
 - What is the space-time diagram of the image of the virtual point *on the retina*?



- Called **judder** and/or **smear**
- Effect of doubling the frame rate:
 - Still **smear**, but less "smeared out"
- Consider this case:
 - User is wearing an HMD
 - Point moves constantly in the VE
 - Eye tracks point such that eye fixates always the same pixel, i.e., HMD and eye do not move relative to each other; instead, user turns head
 - Assume no latency in HMD tracking & renderer
 - Space-time diagram of image of point on retina?



- Reverse case:
 - Virtual point is fixed in virtual space
 - User turns head & HMD at constant speed
 - Eye tracks the point, i.e., fixates it




Example from an Oculus simulator




With Oculus DK2, you can press F4 in the demo "Titans of Space" to toggle between low- to full-persistence

Animation of the Cause for Judder




How Blurry is Your Display?

- First, fixate the upper UFO with your eyes: you should see stationary black & white vertical stripes, with some grey squares moving by
- Then, track the lower UFO with your eyes – what do you see now?



- Can you explain this effect?



Use external Quicktime Player
in case of too much judder by Powerpoint

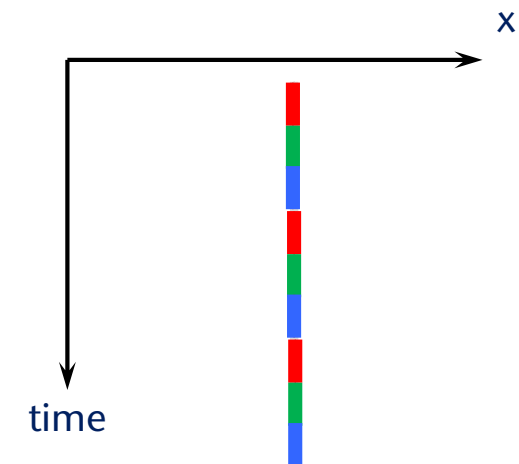
It Can Get Worse

- **Field-sequential color** (FSC) displays: first, only the red channel of all pixels of the frame buffer is transferred (and displayed), then the green channel, then blue channel



- Reduces cost, size, wires, ...
 - E.g., Google Glass used field-sequential color [2014]

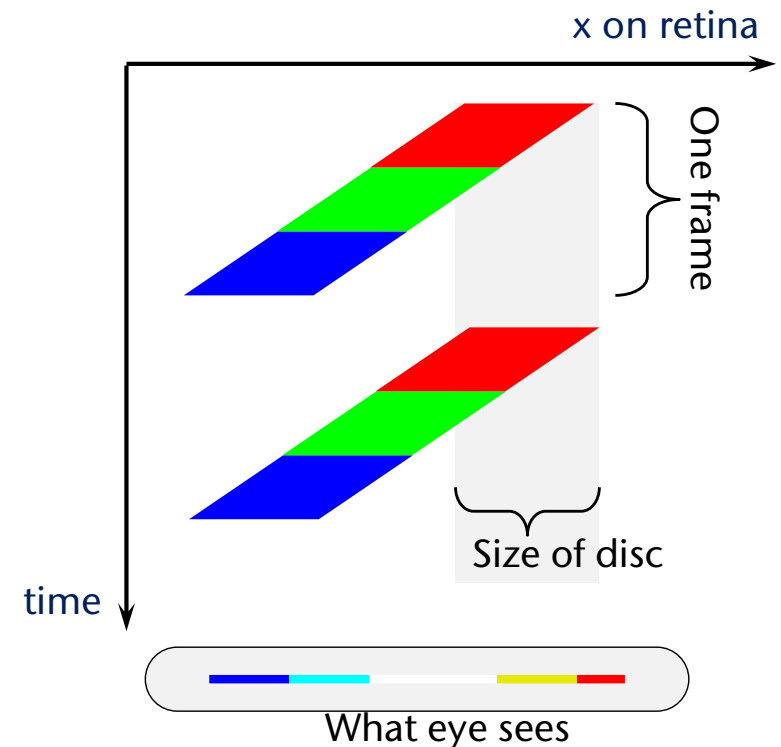
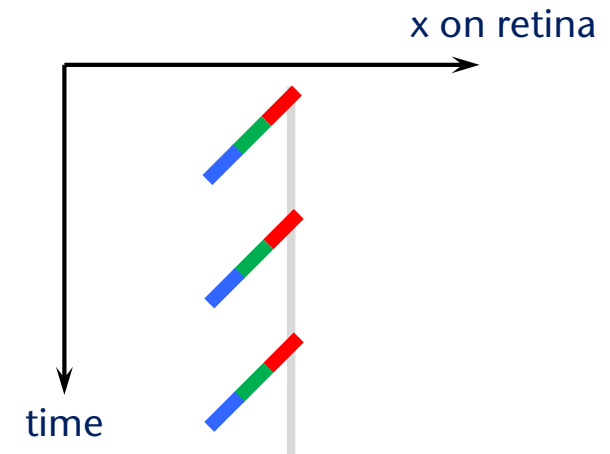
- Space-time diagram of a stationary point on an FSC monitor



- Space-time diagram of a moving virtual point's position *on the retina*, with eye tracking its image on an FSC monitor

- Space-time diagram of a moving disc on the retina rendered on an FSC monitor, tracked by the eye
 - Result: smear and *color fringes*!

- Similar stuff happens in HMD!



Optional

Possible Side-Effects of Low Persistence



- Low-persistence might introduce other problems
- **Strobing**: perception of multiple copies of the same object
 - Smear can hide strobing artifacts
- The short light bursts of a low-persistence display could interact/disturb saccadic masking
 - **Saccadic masking** = eye is effectively blind (to some degree) during a saccade
 - Consequence of the interaction: brain might lose frame of reference
→ visual instability
- Lots of perceptual research needed, and good engineering idea!
 - The 1000 Hz display & rendering pipeline?

