Volumetric Hyper Reality: A Computer Graphics Holy Grail for the 21st Century?

Gavin Miller
Apple Computer, Inc.

- Structure of this Talk
  - What makes a good holy grail?
  - Review of photo-realism
  - Limitations of Virtual Reality
  - Review of 3D display technology
  - Thought experiment of "Hyper Reality"
  - An interesting special case

- Properties of a Useful Holy Grail
  - The goal must be inspiring
  - The goal must be attainable
  - Incremental steps should be useful
  - The goal should be slightly vague
  - The goal should not eat you

- Photo-Realism as a Holy Grail
  - The ability to create images of 3-Dimensional imaginary objects which are indistinguishable from photographs of the corresponding real objects

- The Goal must be Inspiring
  - Both useful and magical
  - Applications to film industry, industrial design, training etc.
  - Artistic medium reminiscent of the Renaissance

- The Goal must be Attainable
  - Not physically impossible if you have enough equipment
  - Required new display devices
  - Required better computers

- Incremental steps should be Useful
  - Wire-frame graphics
  - Smooth shaded polygons
  - Textured polygons
  - Renderman style algorithms
  - Global illumination

- The Goal should be slightly Vague
  - Photo-realistic pictures of what?
  - What effects can be left out?
    - Depth of field, motion blur?
  - Complexity versus fidelity
  - Motion versus imagery

- The Goal should not Eat You
  - Photographs no longer "true"
  - Addictive games medium
  - Spectacle eclipses narrative
  - Pedantic realism vs. artistic inspiration or the appropriate illustration style

- Augmented Reality
  - The seamless integration of synthetic objects with live video imagery in an immersive environment
  - Real objects inside Virtual Reality

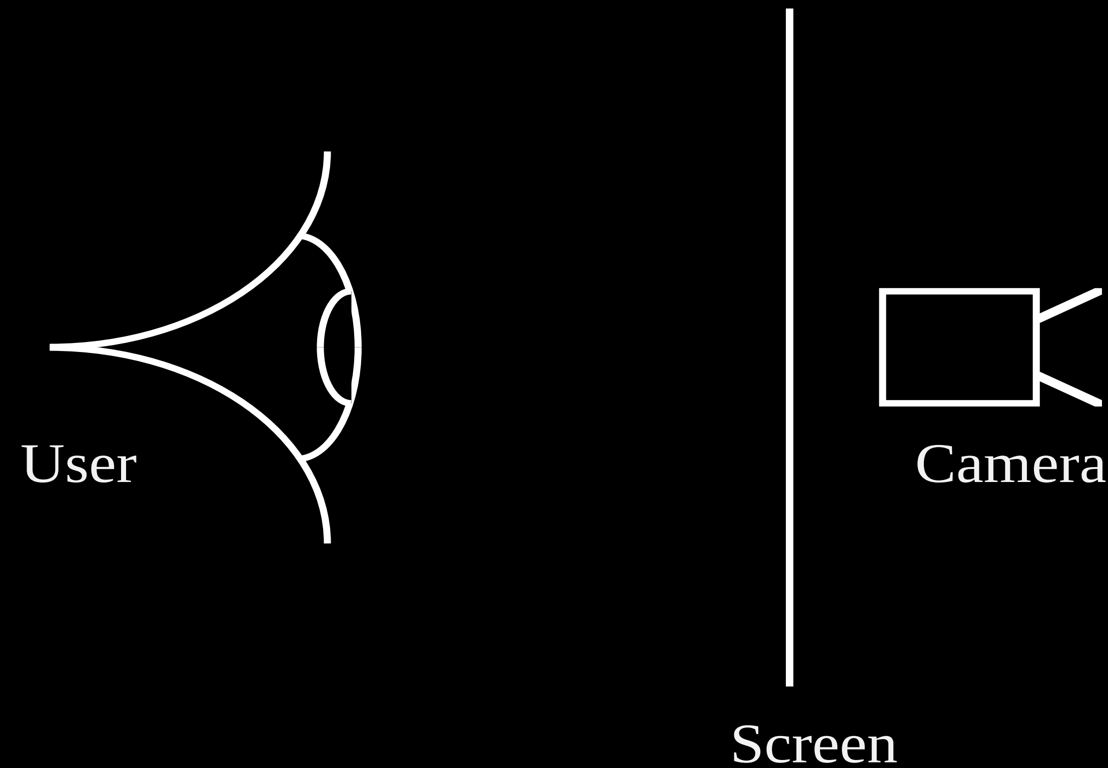
User
Camera
Screen

- Hidden Surface Elimination
  - Do compositing using known order
  - True hidden surface elimination requires depth cameras
  - For opaque matte surfaces all information is knowable from the eye position

- Shadows from Synthetic To Real Objects
  - Use video projectors as light sources
  - Strobe the lighting to test direct illumination from each light source
  - Effects of interreflection not computable

- Shadows from Real To Synthetic Objects
  - Requires a depth camera at each light source
  - Compute visibility of synthetic surface points against the measured depth maps

- Specular Reflections of Real Objects
  - Synthetic specular surfaces may reflect portions of the image not visible directly from the eye position
  - We need to gather appearance information from other locations

- Reflective spheres in the scene
  - Place the reflective spheres at the centers of objects
  - Use the eye cameras to capture reflection information
  - Reflection map results onto synthetic objects

- Problems with reflective spheres
  - Foreground objects may obscure views of the reflective spheres
  - Practical for movies where the scene is carefully planned
  - Inappropriate for real-time augmented reality

- Environment Cameras
  - Capture video images in all directions
  - Apply to small reflective curved objects
  - Apply to large reflective planar objects

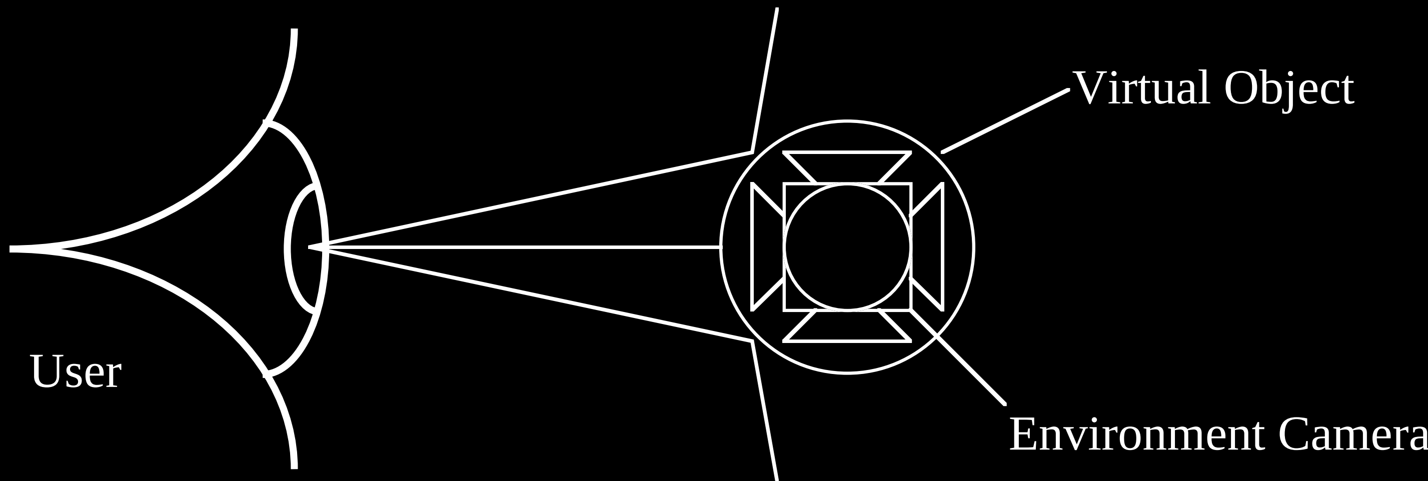
Virtual Object
User
Environment Camera

- Small Reflective Objects
  - Environment camera is centered at the location of the virtual object
  - As object moves, camera must track it
  - Independent of viewer's location

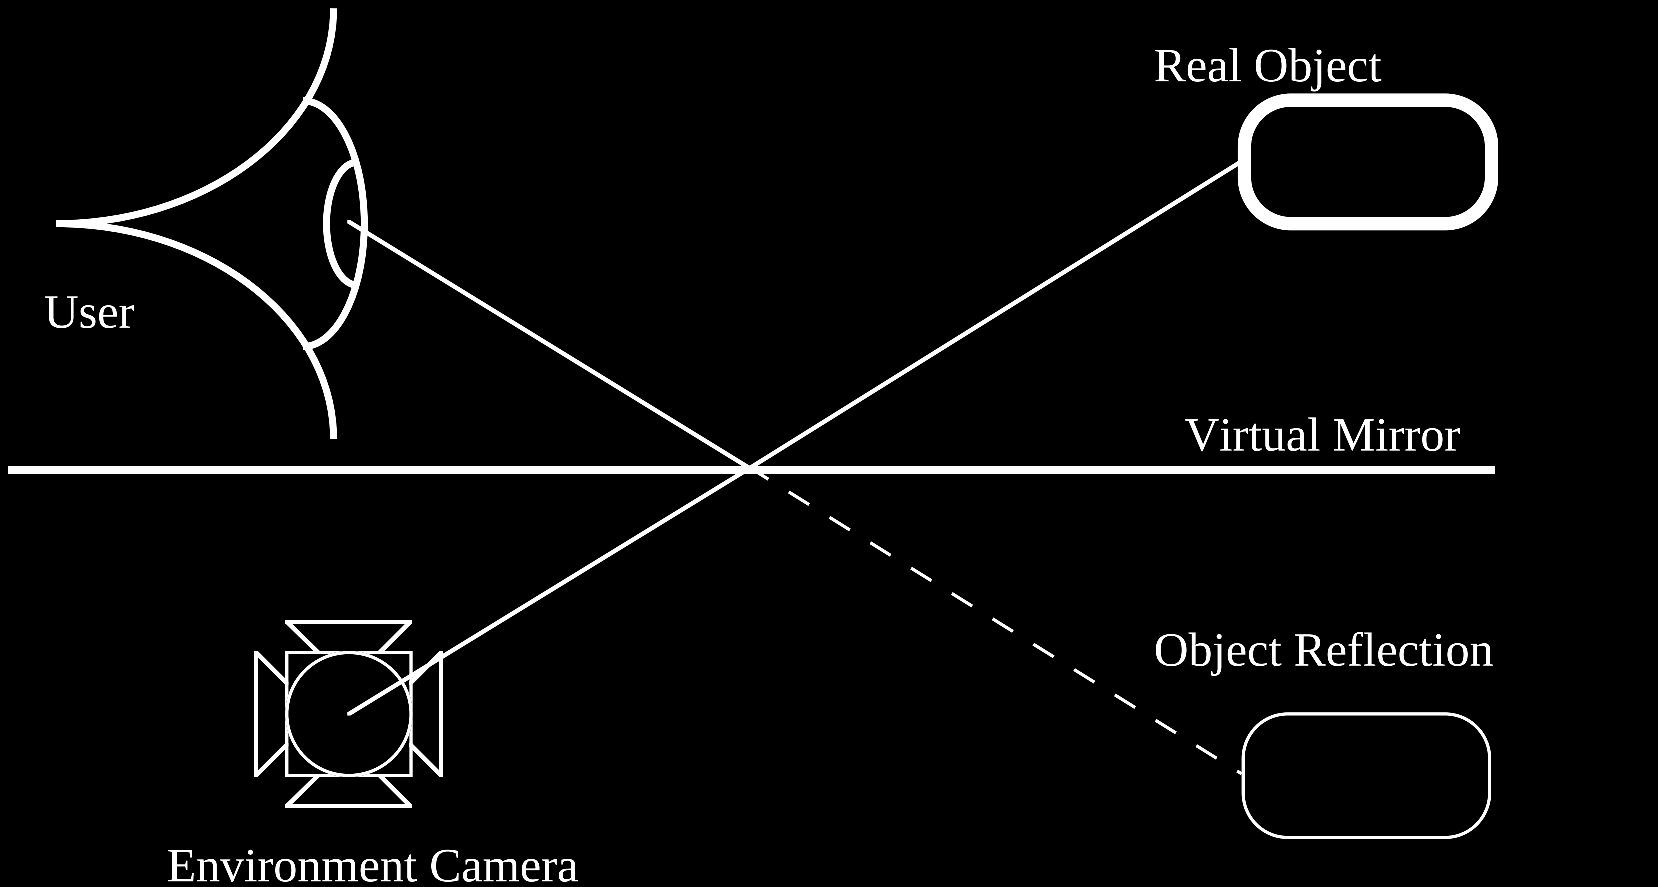
Real Object
User
Virtual Mirror
Object Reflection
Environment Camera

- Planar Reflective Objects
  - Environment camera is at the reflection of the eye position in the plane of the reflective surface
  - As viewer position moves, the camera must track it
  - If the plane moves, the camera must track it too

- 3D Display Technology
  - Virtual objects inside Physical Reality
  - Support more 3D perceptual cues than a standard screen
  - Localized in space so high resolution
  - Possibly less nausea and fatigue

- Additive Volumetric Displays
  - Create the illusion of glowing dots in 3D space
  - Have true motion parallax
  - No special glasses
  - Multi-user
  - Correctly handles changing focus of user's eyes (accomodation)

- A Spinning Light Emitting Diode Array
  - Each angular position of the array corresponds to a slice of data
  - Rotating the array at more than 30 revs per second gives visual fusion
  - Viewable from all directions

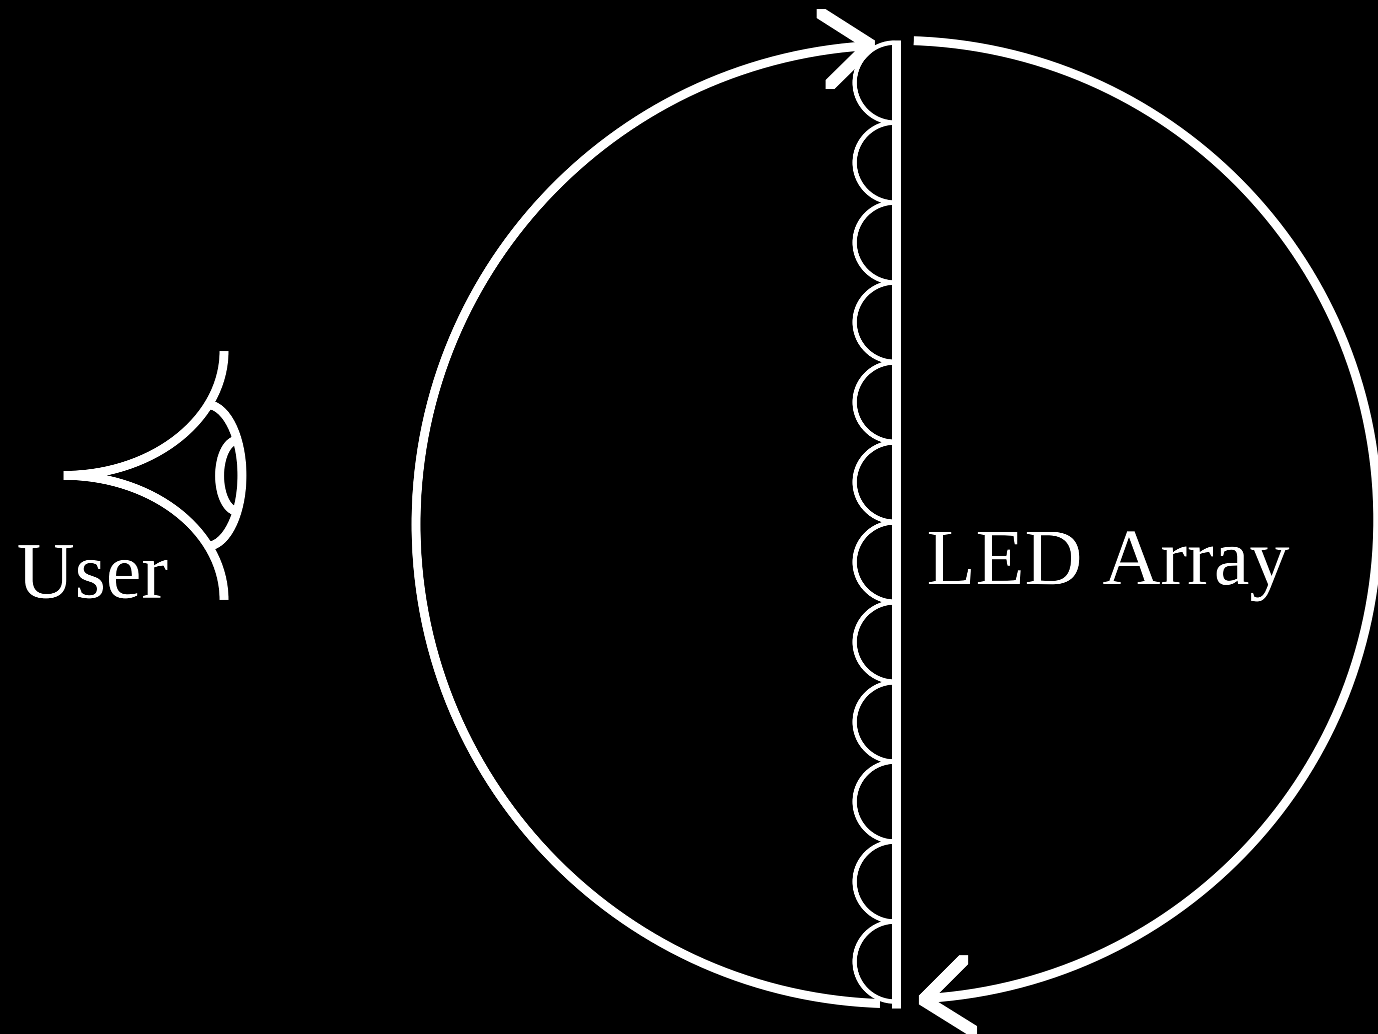
User
LED Array

- A Vibrating Light Emmitting Diode Array
  - Each linear position of the array corresponds to a slice of data
  - Vibrating the array at more than 30 cycles per second gives visual fusion
  - Direct mapping to the volume data
  - Restricted viewing angle

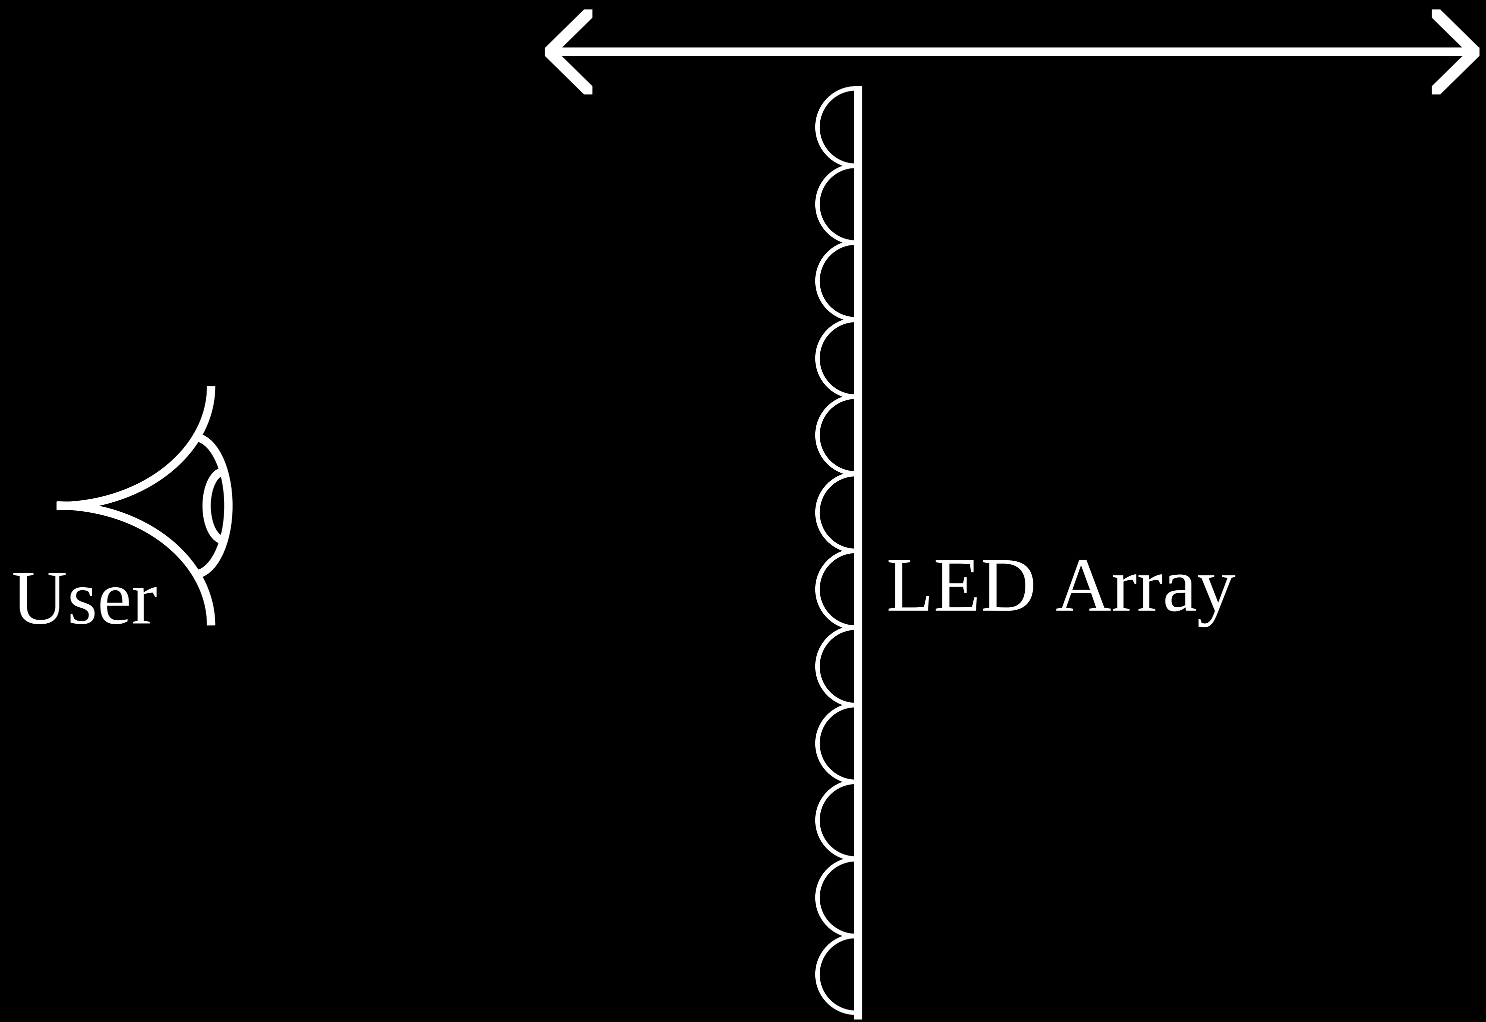
User
LED Array

- A Varifocal Mirror
  - A Cathode Ray Tube screen is seen reflected in a vibrating mirror
  - Each mirror position corresponds to a slice of the image data
  - The viewing angle is limited
  - The mirror is noisey

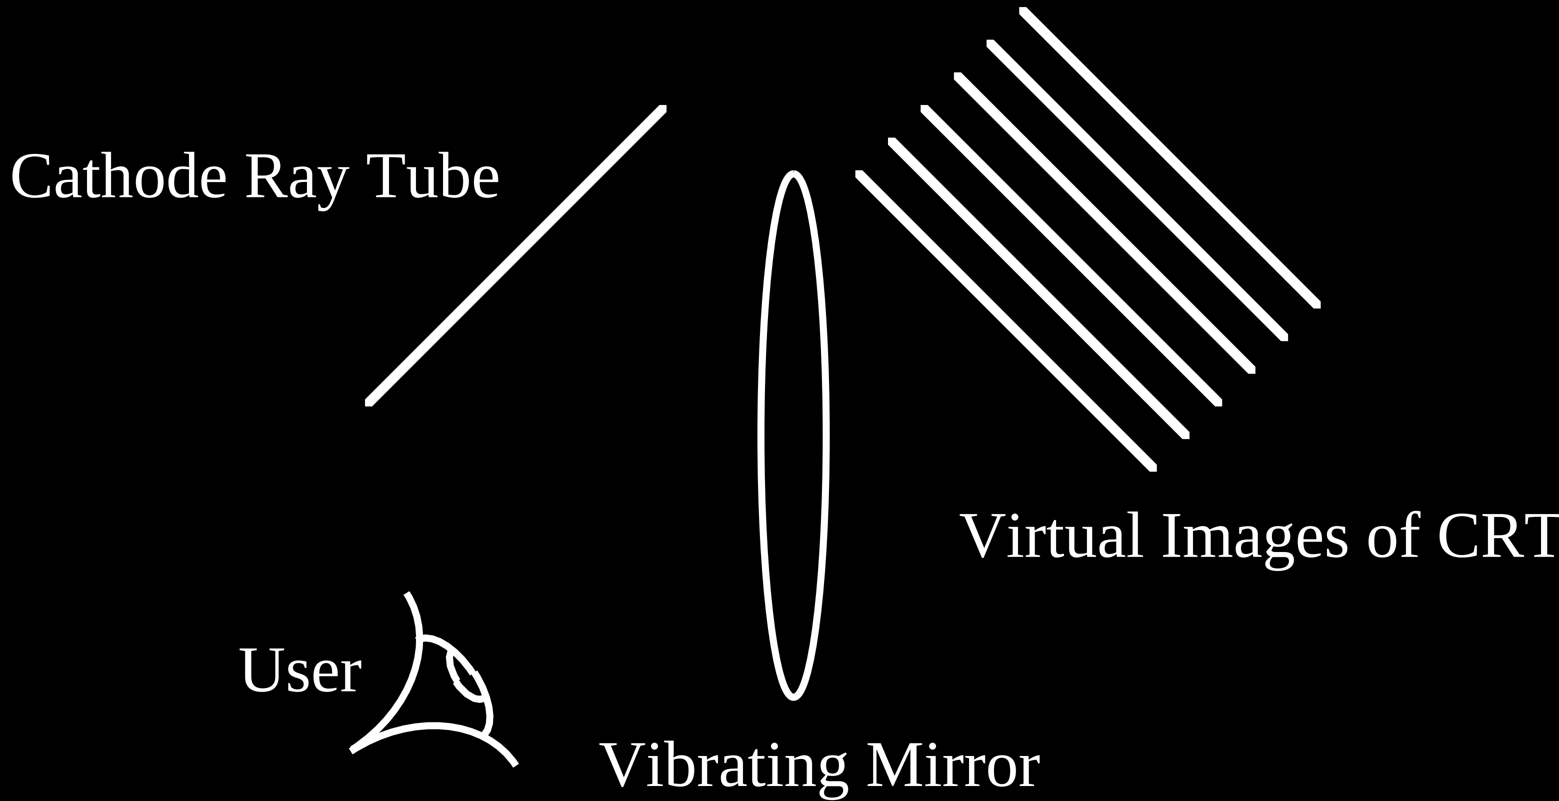
Cathode Ray Tube
Virtual Images of CRT
User
Vibrating Mirror

- A 1D Stack of Liquid Crystal Mirrors
  - A Cathode Ray Tube screen is seen reflected in one of a stack of LCD mirrors
  - Each mirror position corresponds to a slice of the image data
  - The mirrors are "switched on" in sequence scanning out the depth

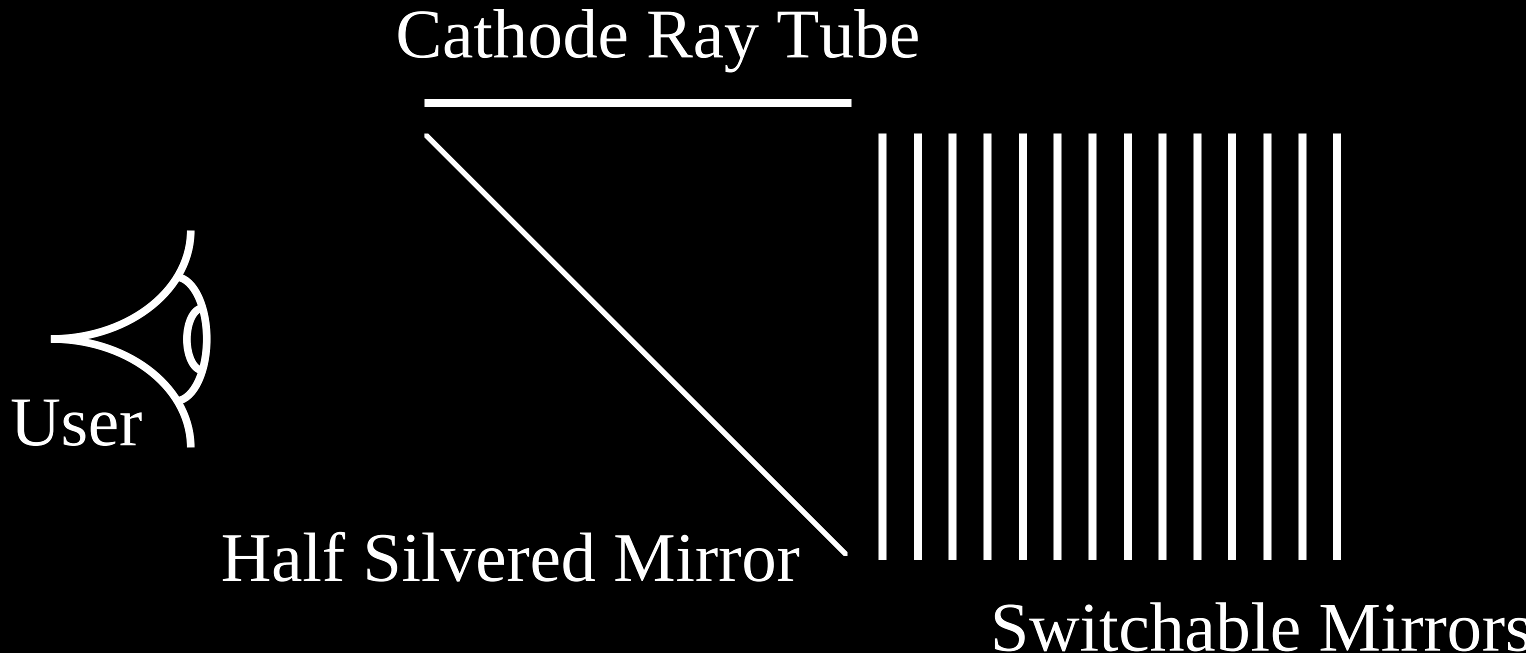
Cathode Ray Tube
User
Half Silvered Mirror
Switchable Mirrors

- A 1D Stack of Holograms
  - Each hologram reflects just one orientation of incident light
  - Each slice of the stack has a corresponding projector
  - The slices add without time multiplexing (can be slides)

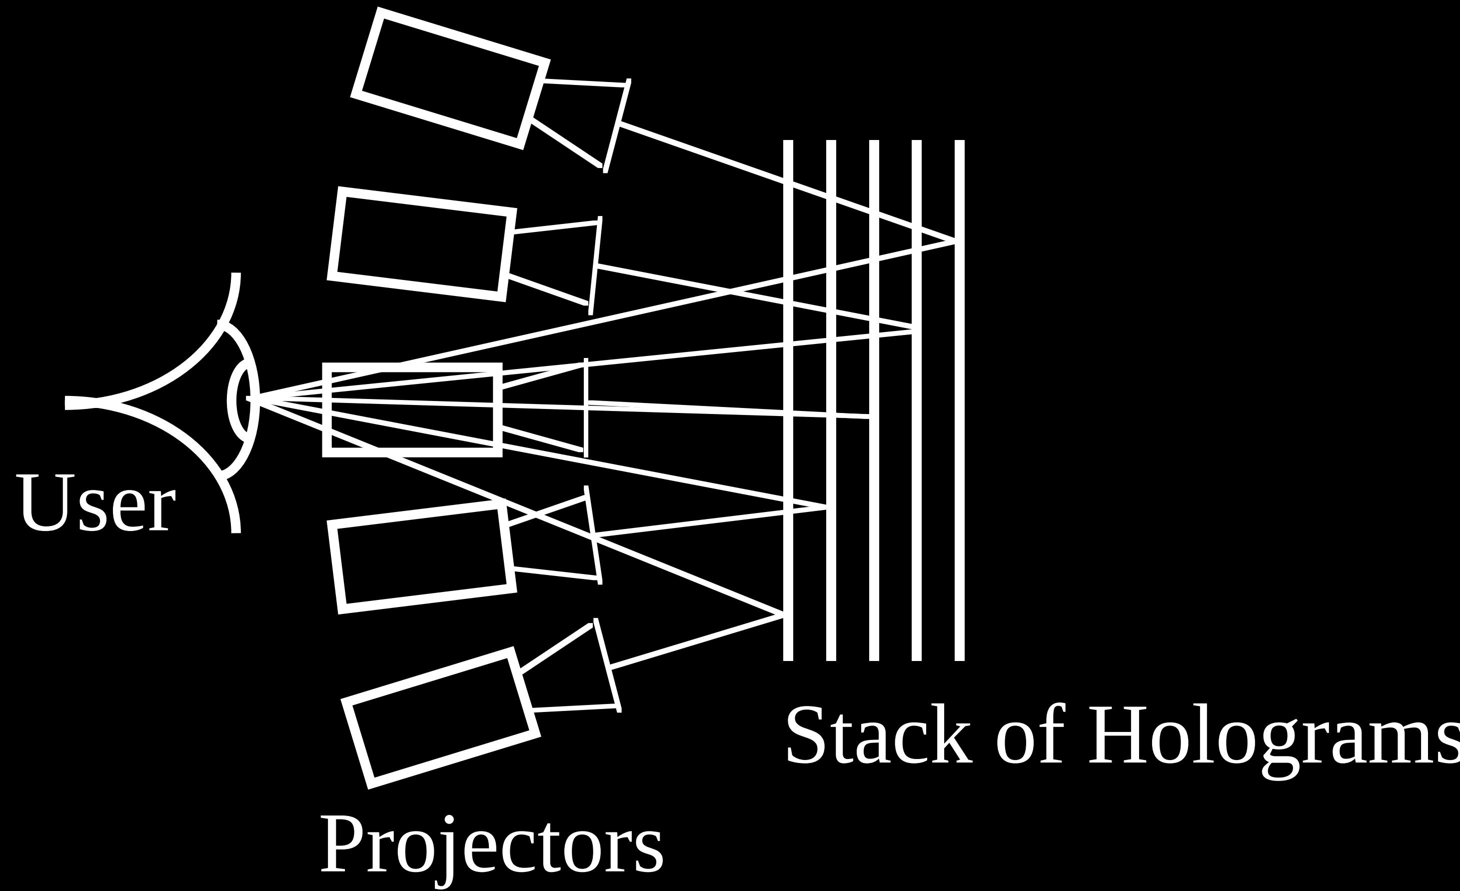
User
Projectors
Stack of Holograms

- The Stacking Problem
  - Many slices of nearly transparent material become opaque
  - The opacity increases exponentially
  - Differences in refractive index cause more problems
  - Demos do not scale

- Laser Projection on Spinning Helix
  - Helix sweeps out the volume
  - Laser projects down onto the surface
  - Laser is deflected using an acousto-optical crystal
  - Large volume with wide view angle

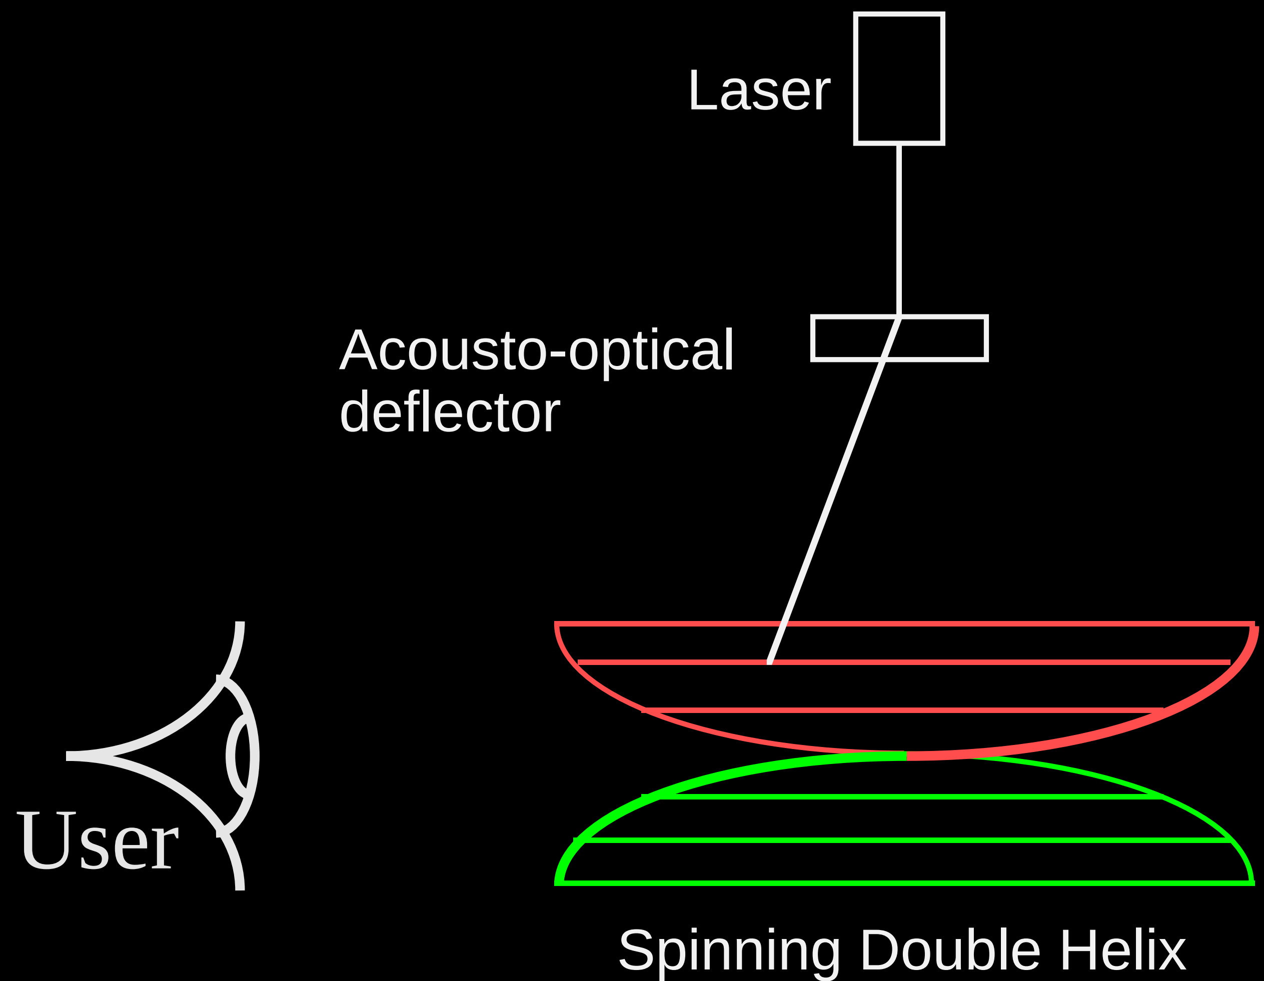
Laser
Acousto-optical
deflector
User
Spinning Double Helix

- Problems with volumetric additive displays
  - No hidden surface removal
  - Makes shading and texturing difficult
  - Displays become more confusing as their capabilities increase
  - Useful for diagrammatic display
  - Not the path to flawless realism

- Stereoscopic displays
  - Shows a different image to each eye
  - Can have shading and texturing, specular effects etc.
  - Motion parallax is through head tracking and instant update

- Red-Green Stereo
  - Color filters on each eye
  - Only gives monochrome images
  - Strange color switching effects

- Active Shutter Glasses
  - Light valve on each eye
  - Opens one then the other
  - Needs a display with twice the frame rate
  - Rest of environment is also attenuated


User
Screen
Shutter Glasses

- Active Shutter Screens
  - Polarizing shutter on the screen
  - Passive polaroid glasses
  - Glasses are very cheap
  - Screen display is expensive
  - Rest of environment is brighter without flicker

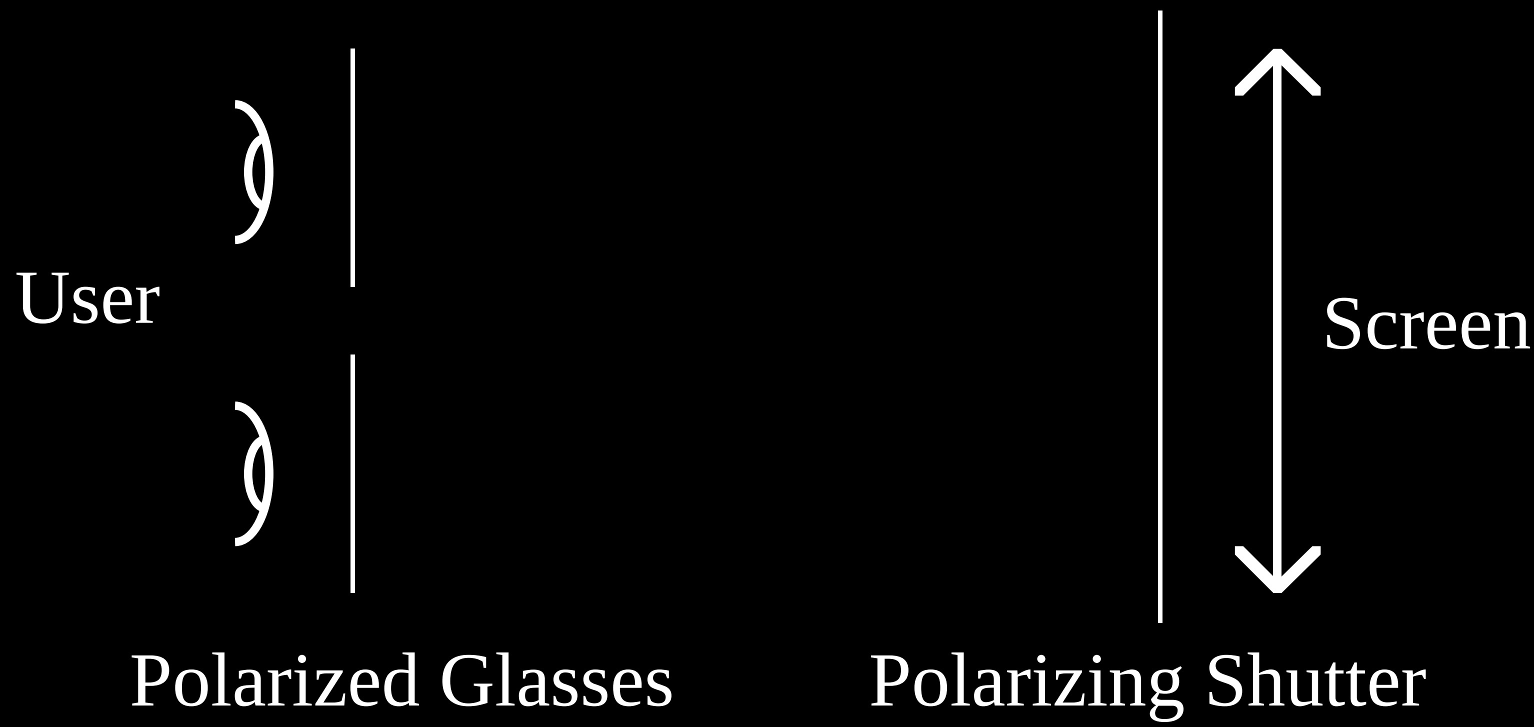
User
Screen
Polarized Glasses
Polarizing Shutter

- Lenticular Stereo Screens
  - Cylindrical lenses in front of pairs of pixels
  - Need head in the correct position
  - No special glasses
  - Sacrifices horizontal resolution for stereo effect
  - Can move lenses to do head tracking

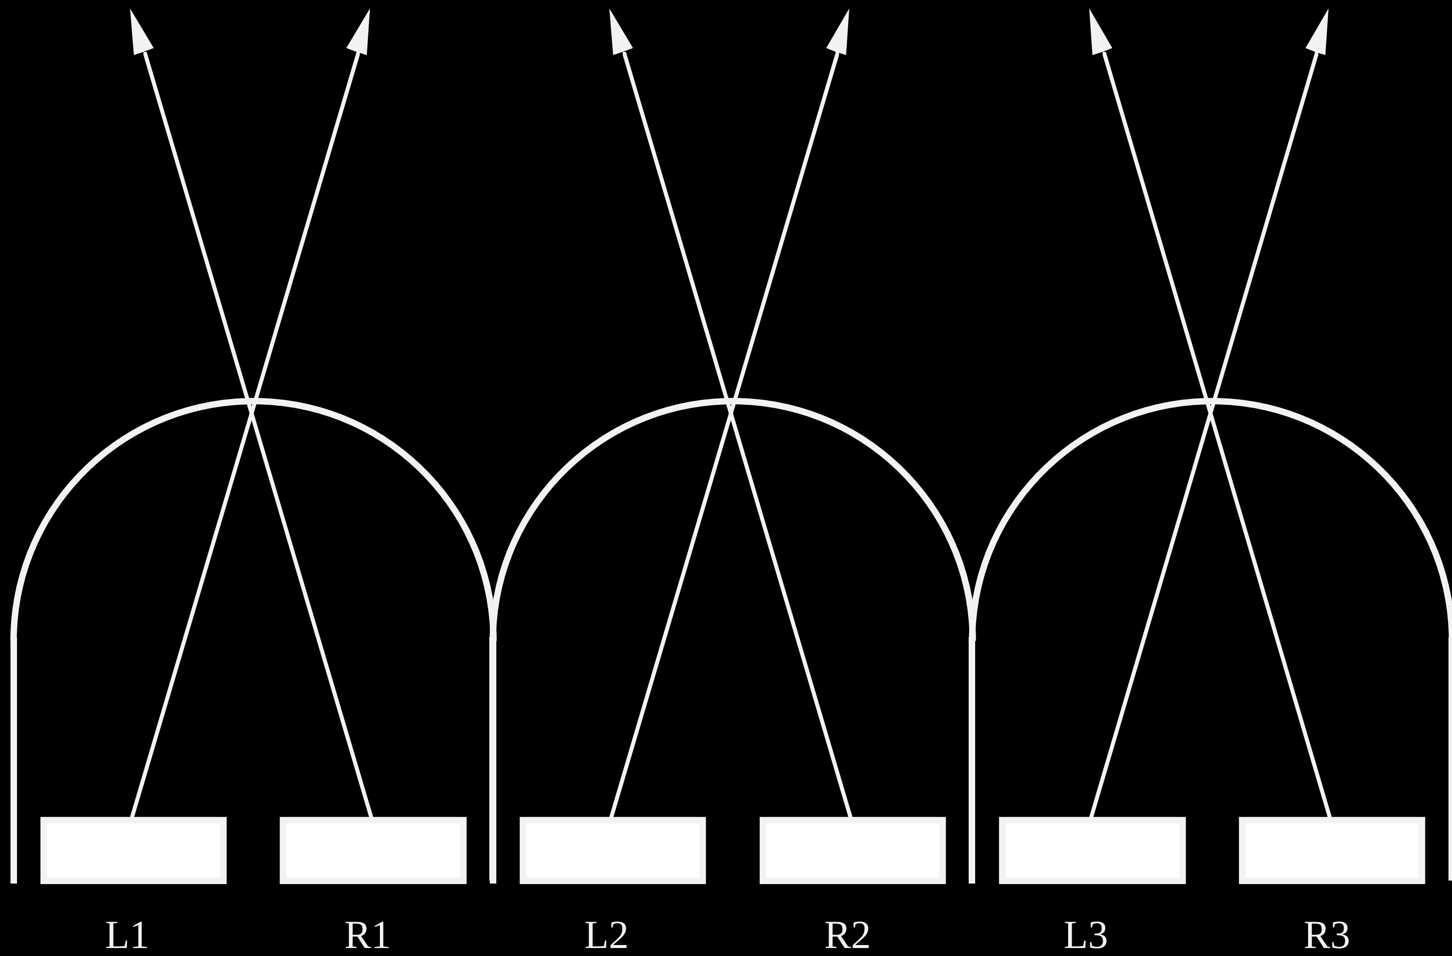
L1
R1
L2
R2
L3
R3

- L.C.D.s with Directional Backlighting
  - Backligting is in vertical strips corresponding to a zone of projection
  - Switches between left and right zones
  - Increased frame rate for stereo
  - Can default to 2D display
  - Can have more zones for more angular resolution (integral displays)

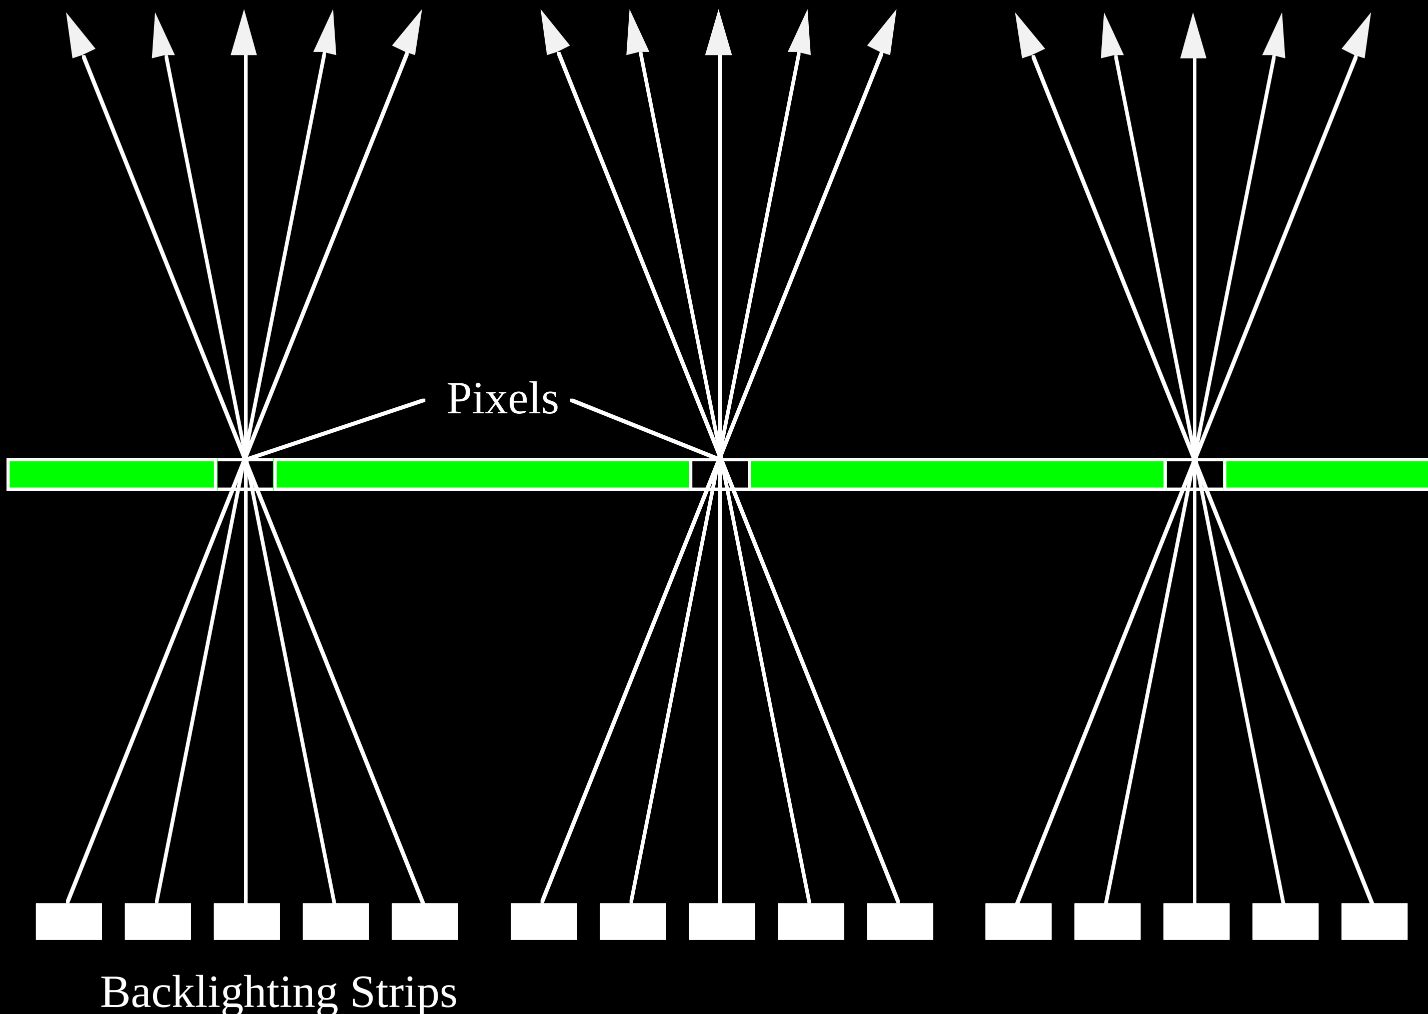
Pixels
Backlighting Strips

- Integral Displays
  - Display many images multiplexed in orientation (rather than depth)
  - Create the illusion of motion parallax
  - Allow hidden surface removal, shading and other optical effects
  - A generalization of stereo displays
  - The more images the better

- Types of Integral Display
  - Unidirectional displays only multiplex the images over horizontal orientations
    - Needs tens of big images
  - Omnidirectional displays multiplex horizontally and vertically
    - Requires one image per pixel
    - Needs thousands of small images

- CRT with LCD moving slot
  - LCD shutter has a moving slot
  - Each slot position has an image behind it
  - Much higher frame rate
  - Can have several slots
  - Can display 16 images simultaneously

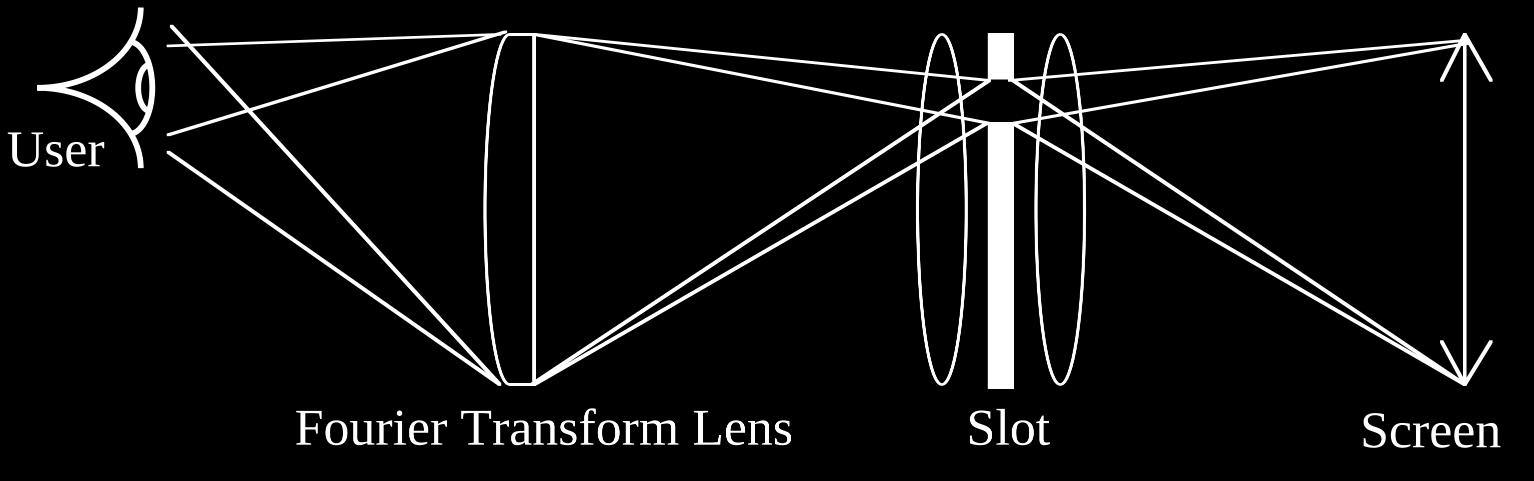
User
Fourier Transform Lens
Slot
Screen

- Virtual Vertical Slit using Hologram
  - A spinning hologram creates the effect of a spinning cylinder with a slit in it
  - Use a static high frame-rate laser projector to generate the images

- 2D to 1D interleaving using a hologram
  - Spread a square block of pixels over a horizontal range of orientations
  - Sacrifices horizontal and vertical resolution equally to get horizontal motion parallax
  - Uses standard refresh rates
  - Good for high resolution LCDs

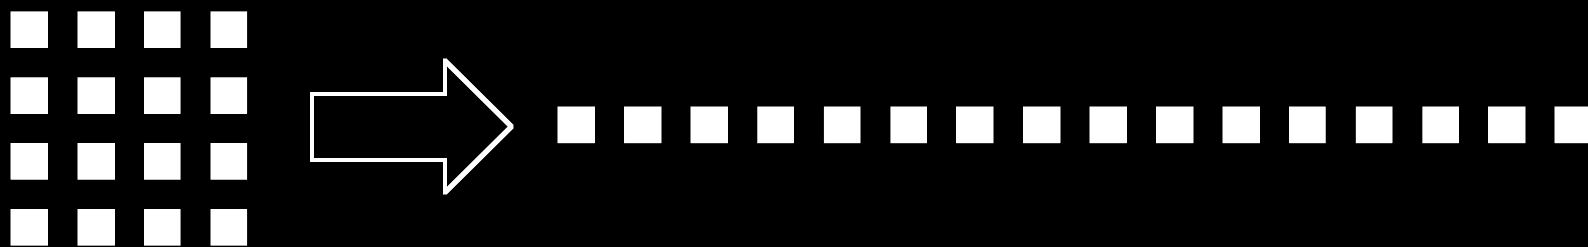

- Omnidirectional display using retro-reflectors and multiple projectors
  - Retro-reflective screen only reflects back in direction from which it is illuminated
  - Place the projectors appropriately so that the viewer sees an image corresponding to that view direction

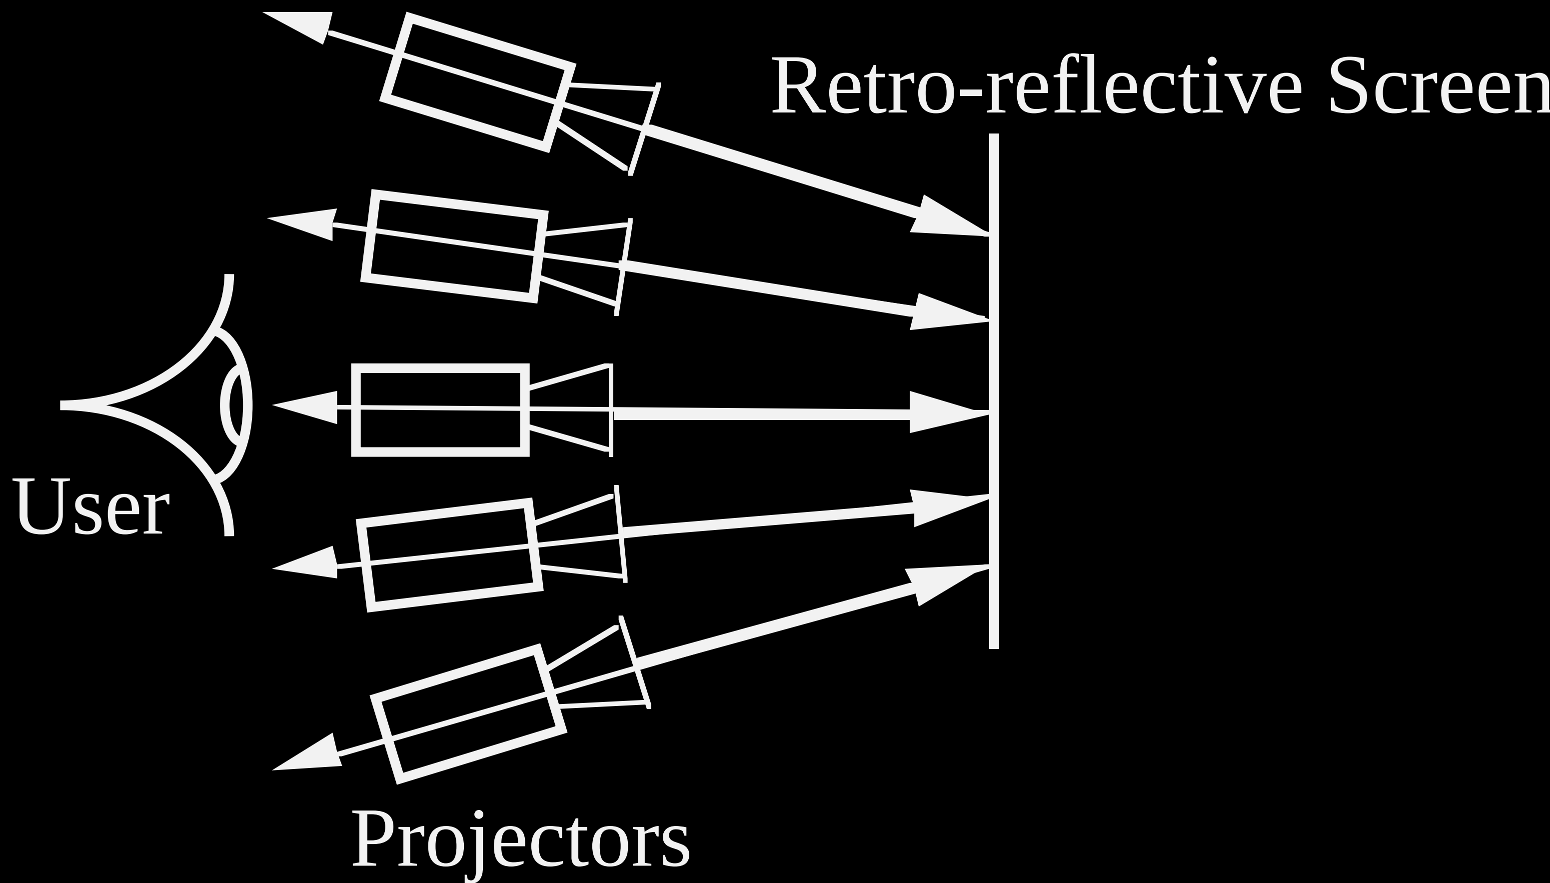
Retro-reflective Screen
User
Projectors

- Integral Photography
  - Use a fly's eye array of lenses
  - Behind each lens is a small photograph
  - Creates horizontal and vertical motion parallax
  - Direct photographic approach leads to pseudoscopic image which is inverted

- Synthetic Aperture Holography
  - Uses a travelling acoustic wave in a crystal to generate an interference pattern
  - Limited by properties of the crystal
  - Currently small field of view

- Depth of Field of Integral Displays
  - Objects in the image plane are in focus
  - As an object moves away from the plane, the focus decreases
  - Rate of decrease depends on angular resolution of the display
  - Either leads to blurring or discontinuous jumps between images

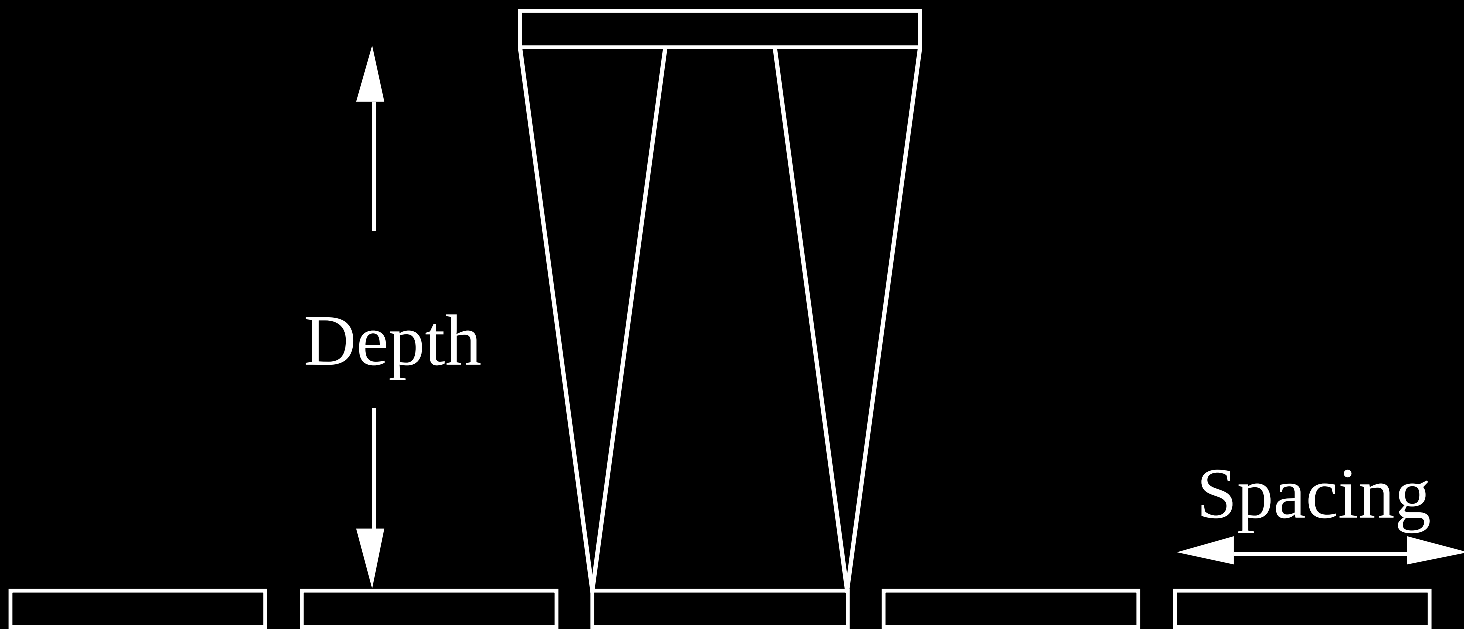

Pixel size = Pixel Spacing + Depth * Pixel Angle

- Limitations of (Integral) Photography
  - A photograph records the optical environment of the subject not the viewer
  - Metal objects reflect the camera not the viewer of the photograph
  - One's hand behind the photograph is not refracted by an image of a glass

- Reality Break
  - The rest of this talk is speculation
  - Inspired by thought experiments in physics
  - Use at your own risk
  - What about that weird talk title anyway?

- Volumetric Hyper Reality
  - Create the flawless illusion within a volume of the prescence of objects made of arbitrary materials
  - Displayed objects are indistinguishable from real objects
  - Displayed objects react to the optical environment of the display

- Volumetric Hyper Reality
  - Light shone on the display would illuminate objects within it
  - When instructed to depict nothing the display would disappear

- "The Invisible Man" by H.G.Wells
  - Eliminate all light absorbing substances (the man was albino)
  - Make the refractive index of the body go to unity using a magic potion
  - Inspired by 19th Century idea of gross physical transformations by chemistry
  - Not physically possible

- 21st Century Invisibility
  - Use surface microstructure to fake the appearance of bulk properties
  - Information flow problem "around" the interior structure
  - Requires active computational model
  - Not physically possible (yet)

- Hyper Display
  - Each hyper pixel is a hemispheric projector (like an integral photograph)
  - Each hyper pixel is a hemispheric camera (like a camera array)
  - Information flows along a bus structure to allow the light to be simulated for space internal to the structure

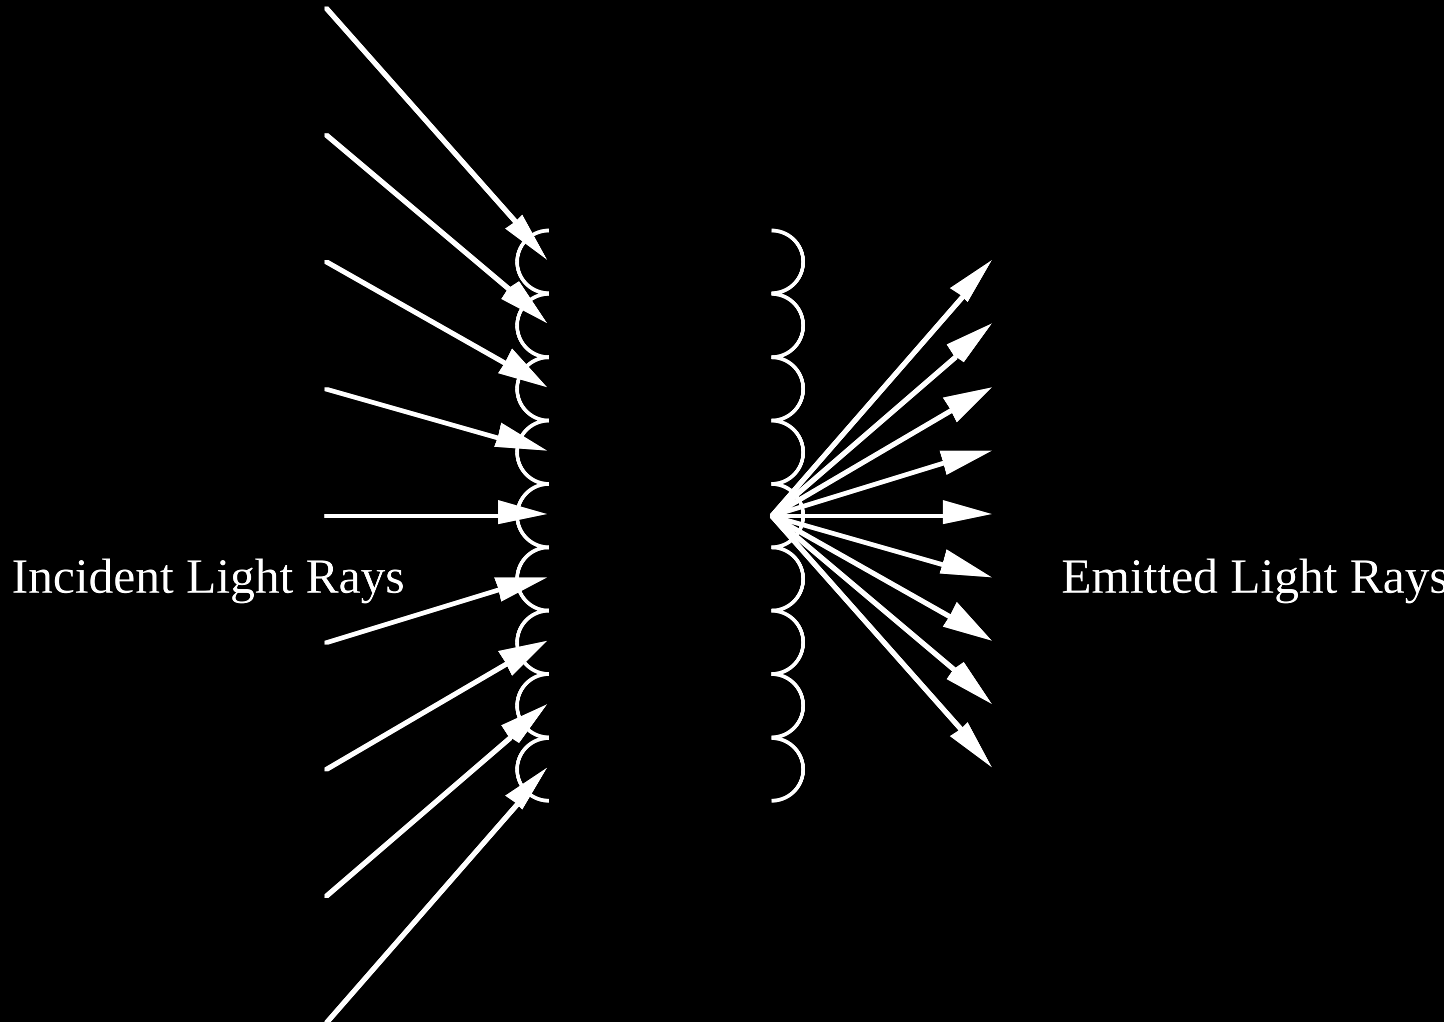
Incident Light Rays
Emitted Light Rays

- Surface Reflections
  - The hyper pixel lenses will reflect incident light
  - This light intensity can be computed for the known lens
  - Can subtract that intensity from transmitted light value
  - Will help to cancel out surface reflections

- Spherical Hyper Displays
  - Can have hyper pixels on a spherical surface
  - Gives Volumetric Hyper Reality capabilities
  - The bus structure must connect every hyper pixel to many others
  - Could use a simple linear bus structure

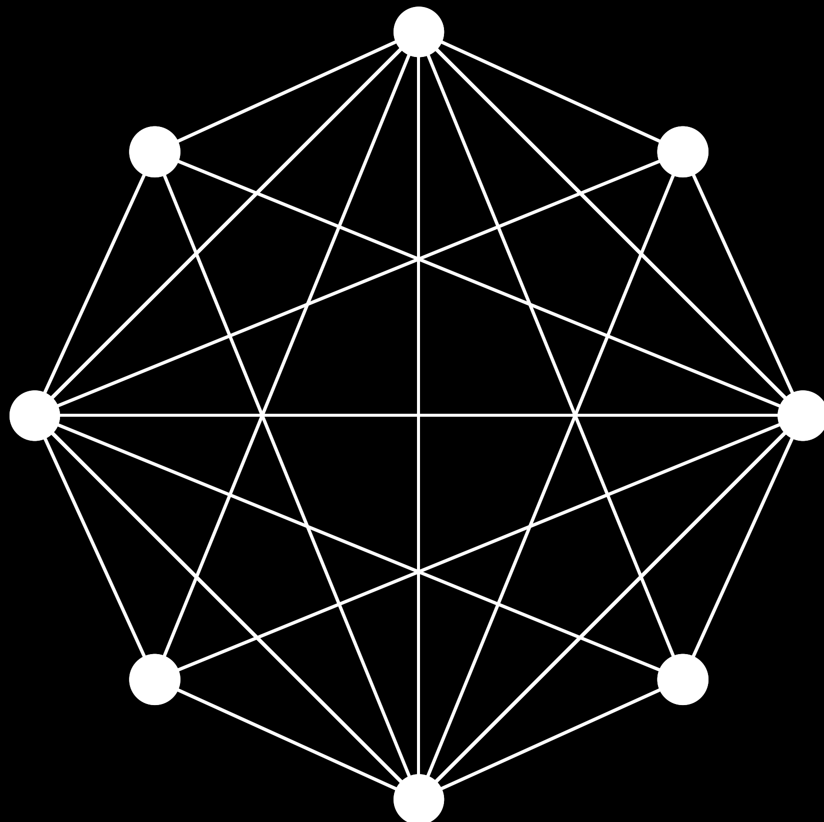

- Hyper Display for Remote Presence
  - Have a concave spherical hyper display facing inwards at real objects
  - Have a convex spherical hyper display representing that space somewhere else
  - Objects inside the concave display would reflect objects outside the convext display, true for lighting etc.

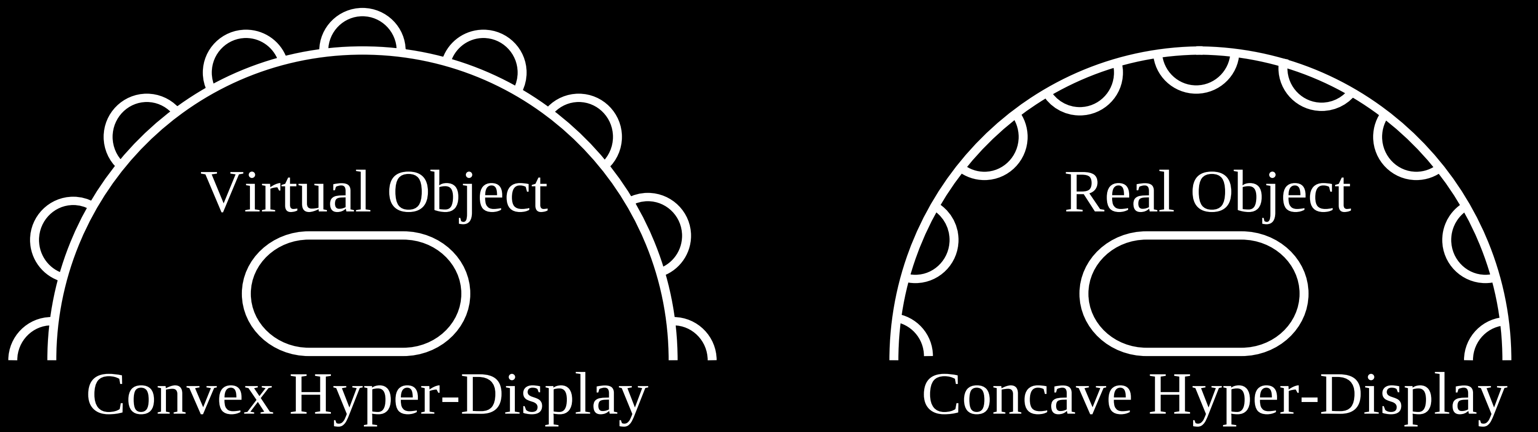

Hyper-Realistic Remote Presence

- A General Graphics Model
  - Global illumination for all hyper pixels
  - Transluminance matrix relates incoming and outgoing light
  - Matrix is fixed for a fixed model
  - Illumination updated at 60Hz for lighting and specular effects

- General Conclusions
  - We need faster computers
  - We need better displays
  - We would need a very special reason to go to this much trouble
  - Any early fruit?

- Looking Glass Display
  - Uses a single flat screen 2D display
  - Uses one or two cameras to capture the optical environment
  - Does real-time ray tracing effects
  - It's a Mac!

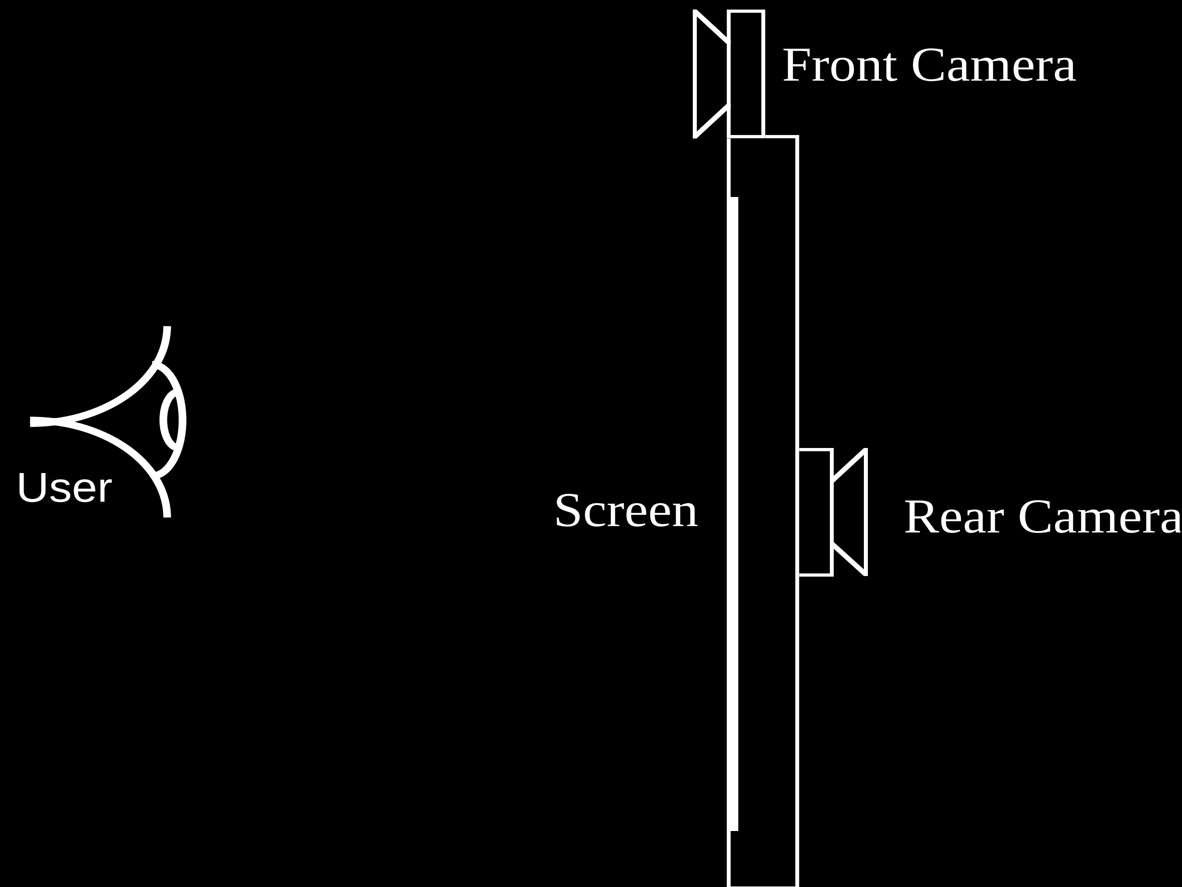

Front Camera
User
Screen
Rear Camera

- Demonstration
  - Recorded in real time
  - Processing live video feed
  - No smoke (or special hardware)
  - No mirrors (uses hyper reflections)