Light-Field Displays:
Technology and Representation of 3D visual information

Péter Tamás Kovács
Holografika

JPEG PLENO Workshop
Warsaw, Poland
23 June, 2015

Holografika

- Hungarian company, active in the holographic, opto-electronic and 3D display technologies
- Manufacturing next generation holographic displays and developing other 3D related technologies (3D acquisition, 3D conversion, rendering, 3D applications) with proprietary IP
- Experienced team for optical design, hw & sw, etc.
- Patent portfolio
- Red Herring Top 100 Europe in 2005, WEF Technology Pioneer company of 2006, shortlisted for World Technology Award 2006, nominee for the EU ICT Prize in 2007, silver Best Exhibit Award at 2008 ICT Lyon, One To Watch @ GPU Technology Conference 2014

What is possible?

- *„We can make miracles immediately, for the impossible you should wait a few days...”*
- Law of physics – *today...*
  - Is it impossible or just difficult?
  - Easy self-check to look like an expert: “*Could the light beams reach my eyes, as shown?”*
  - No one ever saw a light ray making a turn in free space – “*Where did those light beams come from?”*

The window

- The goal of displaying is to provide perfect representation of real/synthetic scenes
  - Life-like view
  - The ultimate display will be like a window...

- True 3D displaying - reconstructing the light-field as present in the natural view
  - Producing light beams with the same parameters the human perception is capable to process: direction, position, intensity, color (but polarization, phase)
  - „…Let the display work, not the brain…”

Light Field

- General representation of 3D information that considers a 3D scene as the collection of light rays that are emitted or reflected from 3D scene points. (M. Levoy, and P. Hanrahan, “Light Field Rendering”, 1996.)
- The visible light beams are described with respect to a reference surface (screen) using the light beams’ intersection with the surface and angle.
- The LF is defined as a function of position (2 or 3 parameters) and direction (2 parameters): L(x,y,[z],Θ,Φ)

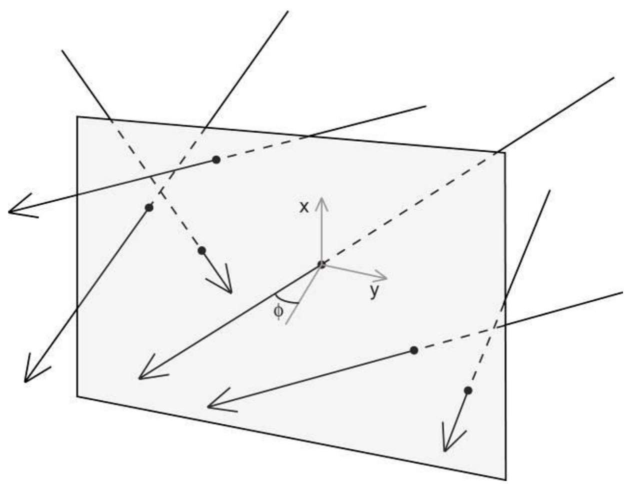

3D Displaying basic rules

- Emission range - FOV
- Number of independent beams in the range
- Discrete directions – but how many?

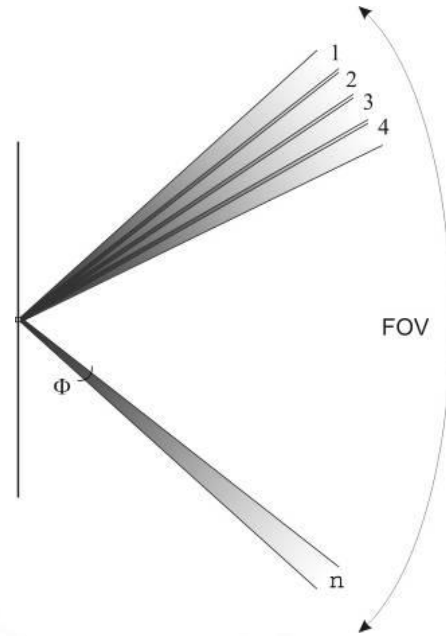

- Horizontal parallax or full parallax
- Reducing the number of light beams by omitting the vertical parallax

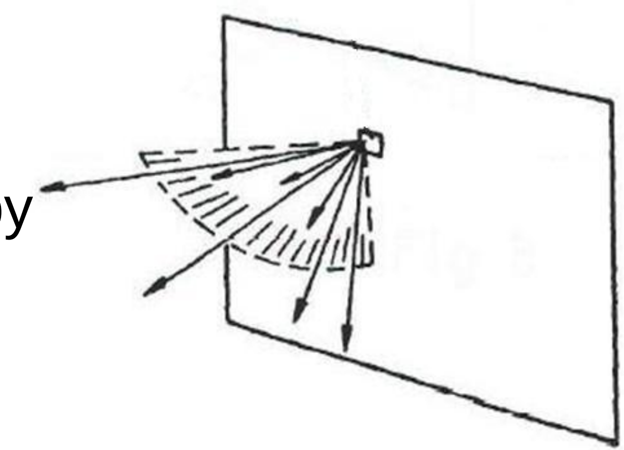

3D Displaying basic rules

- Direction selective light emission
  - Common for all systems having a screen, conforming with current displaying conventions (also true for the outer surface of volumetric systems)
  - Not true for glasses-based, as separation is achieved in a different way
- General approach: to create a light emitting surface, where we are able to emit different light beams from each point in a controlled way

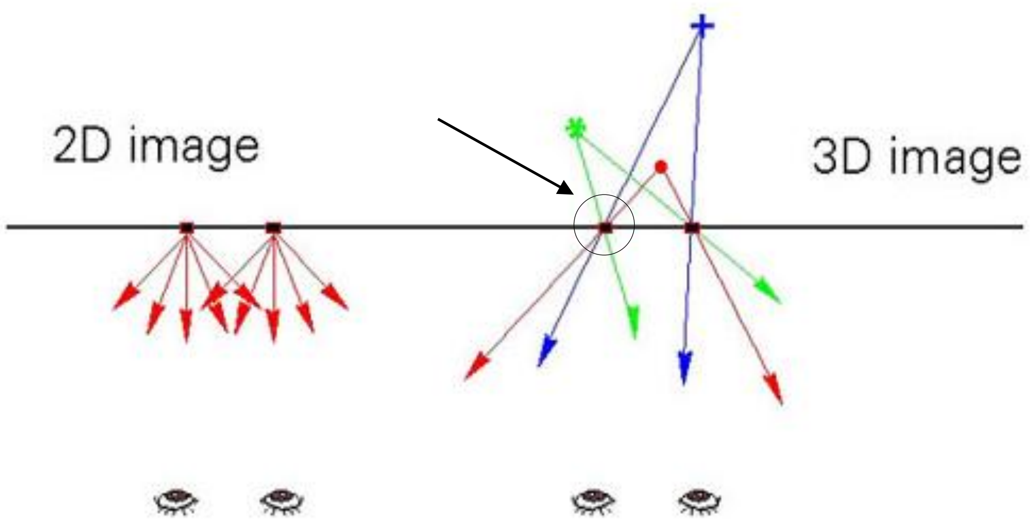

Light-field displays

- HoloVizio by Holografika
- Optical modules
  - Project light beams to hit the points of a screen with multiple beams under various angles of incidence
- Holographic screen
  - Direction selective property
- Light field reconstruction instead of views
- Emission angle geometry determined
  - The projected module image not a 2D view of the final 3D image
  - Each view of the 3D image comes from more modules
  - Continuous motion parallax - no discrete border between views

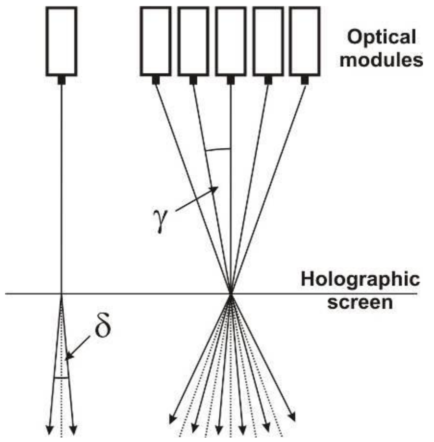

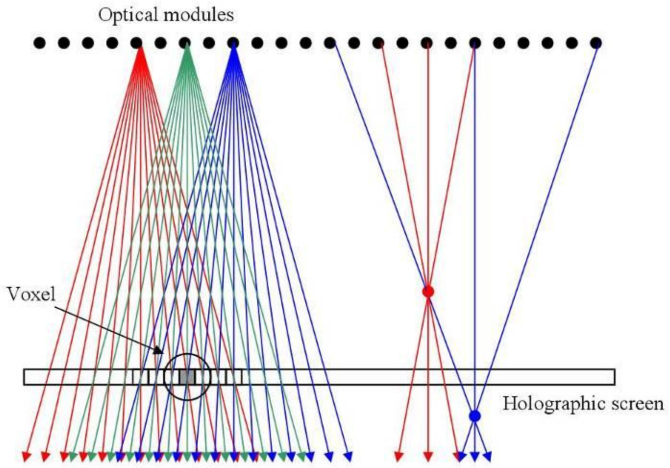

Light-field displays

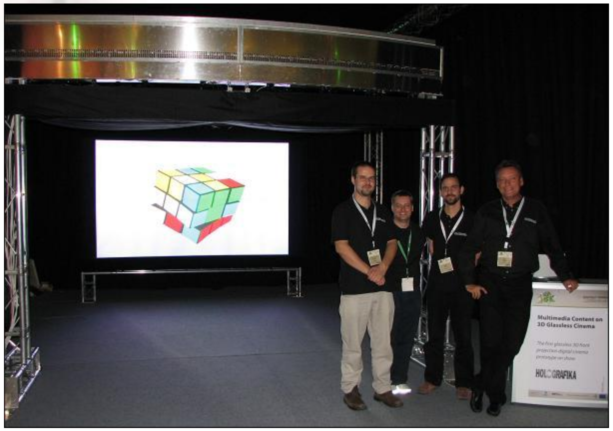

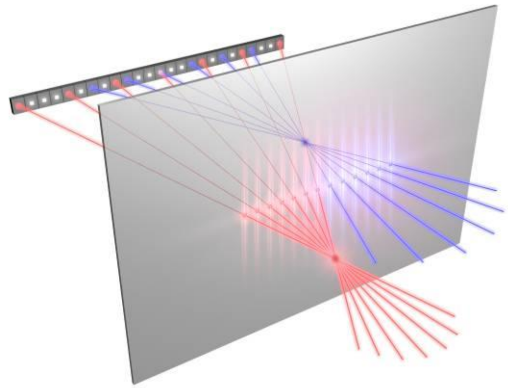

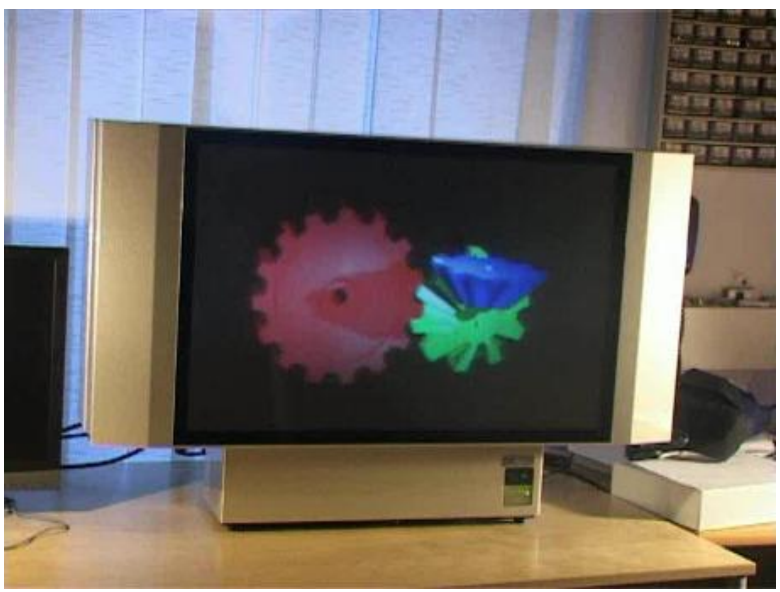

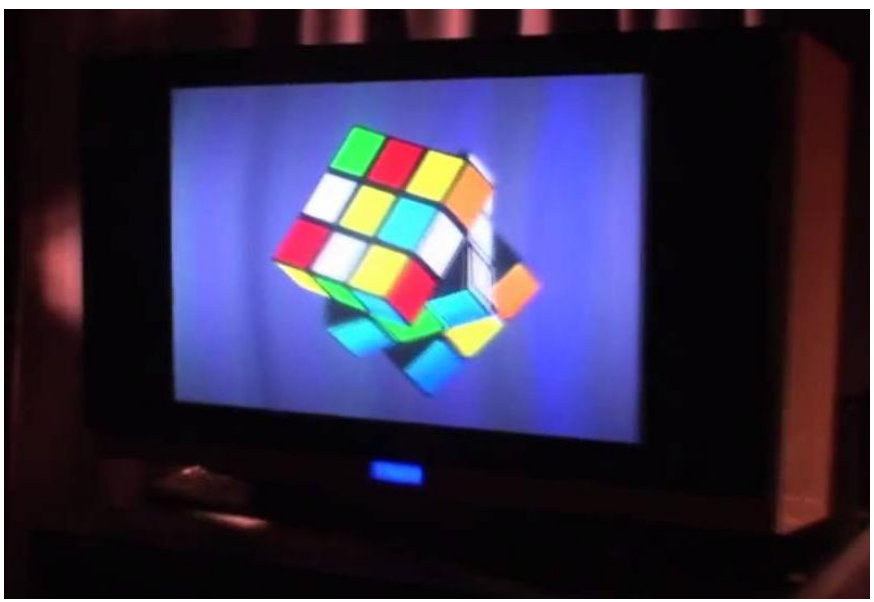

HoloVizio displays

- The large-scale HoloVizio system
  - **HoloVizio 721 RC**
    - Earlier: 640RC, 720RC

    73 Mpixel, 72”, (16:9)
  - 70 degrees FOV
  - Angular resolution 0.9 degrees
  - 2D equivalent image resolution 1280x768
  - LED colors
  - Gigabit Ethernet input
  - Control system+PC based render cluster

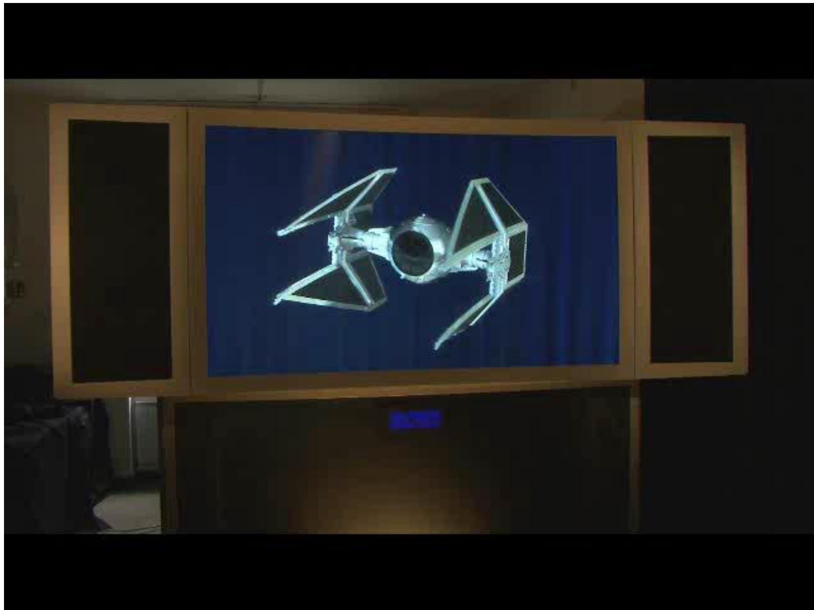

HoloVizio displays

- The digital signage HoloVizio kiosk system

  - **HoloVizio 361H**
    - Earlier: 240P, 360P

    33 Mpixel, 52”

  - 50 degrees FOV
  - 2D equivalent image resolution 1280 x 720
  - LED colors
  - Gigabit Ethernet input
  - Control system + 3 PC
  - Only 0.5 m deep!

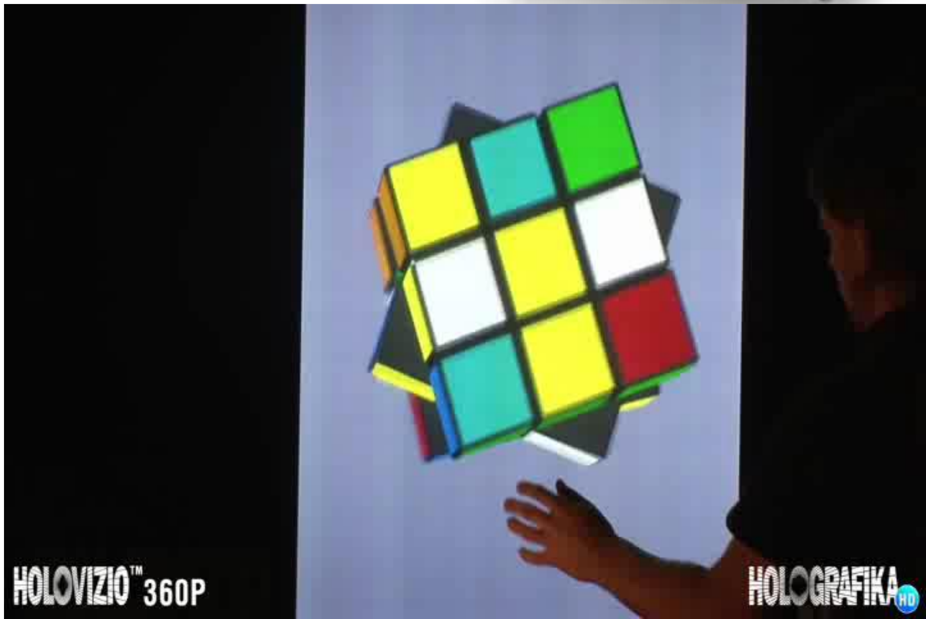

HoloVizio displays

- The world first glasses-free 3D cinema system
  - **HoloVizio C80**
  - 3,5 m reflective holoscreen (140”)
  - 63 Mpixels
  - LED based 3D projection unit
  - Exceptional 1500 Cd/m^2 brightness
  - 40 degrees FOV
  - PC based render cluster
  - No perspective distortion
  - 2D compatible
  - Fitting cinema rooms, 3D simulators
  - Also with portable mechanics for event rental

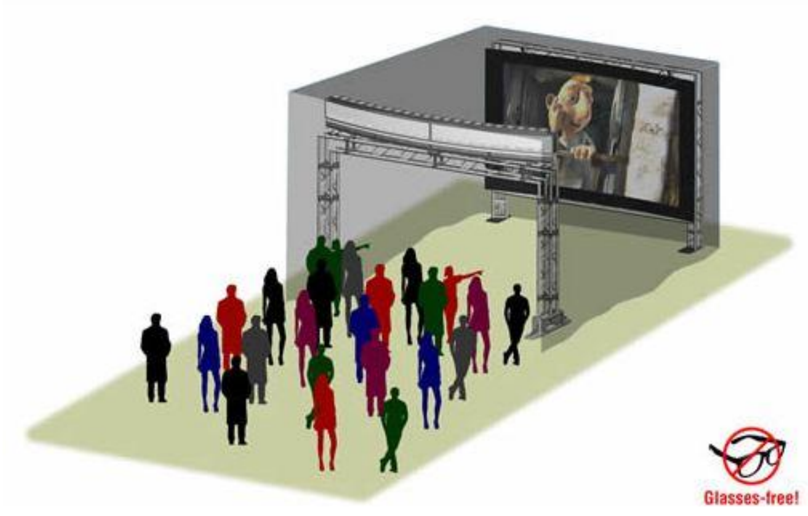

HoloVizio displays

- The full-angle HoloVizio monitor
  - **HoloVizio 80WLT**
    - Earlier: 96ND, 128WD, 128WLD
  - 78 Mpixel, 30” (16:10)
  - ~180 degrees FOV
  - Total freedom in 3D experinece
  - No invalid zones
  - No repeated views

  - 2D equivalent image resolution 1280 x 768 (WXGA)
  - LED colors
  - Multiple DVI inputs

Viewing angles of 3D display technologies

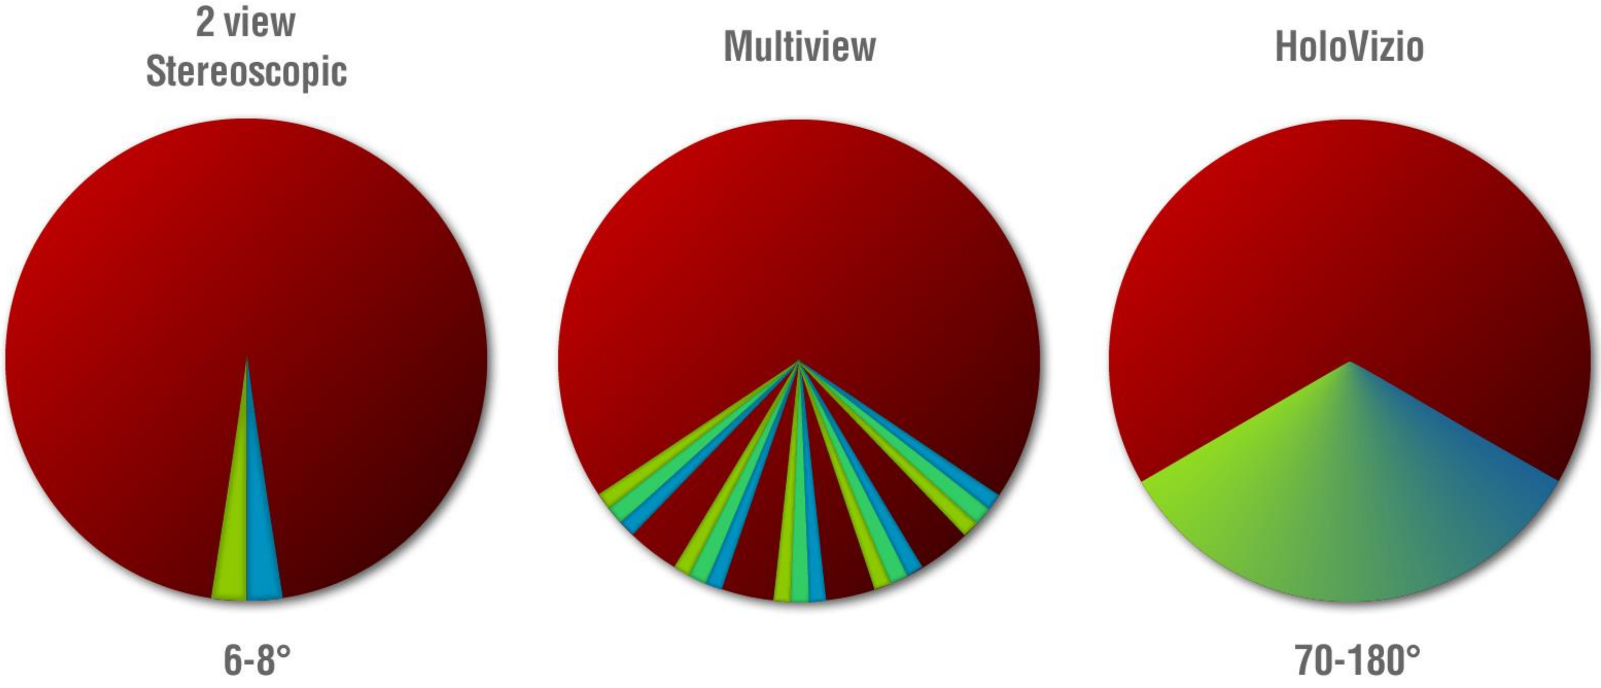

HOLOGRAFIKA

How to feed?

- Integrate existing applications / display 3D models in real-time
  - Interactive applications
- Display pre-rendered 3D models
  - High quality images
- Display real-life 3D scenes
  - Light field streaming

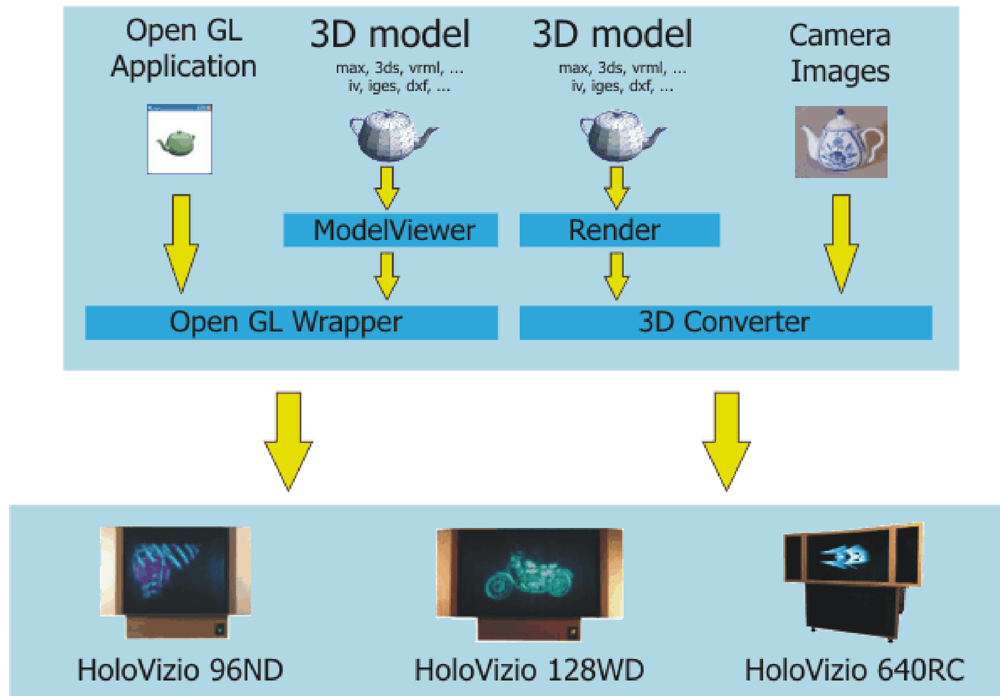

Integrate Existing Apps

- OpenGL wrapper for HoloVizio display systems
- Invisible OpenGL layer, with the same interface
- No application support is required!

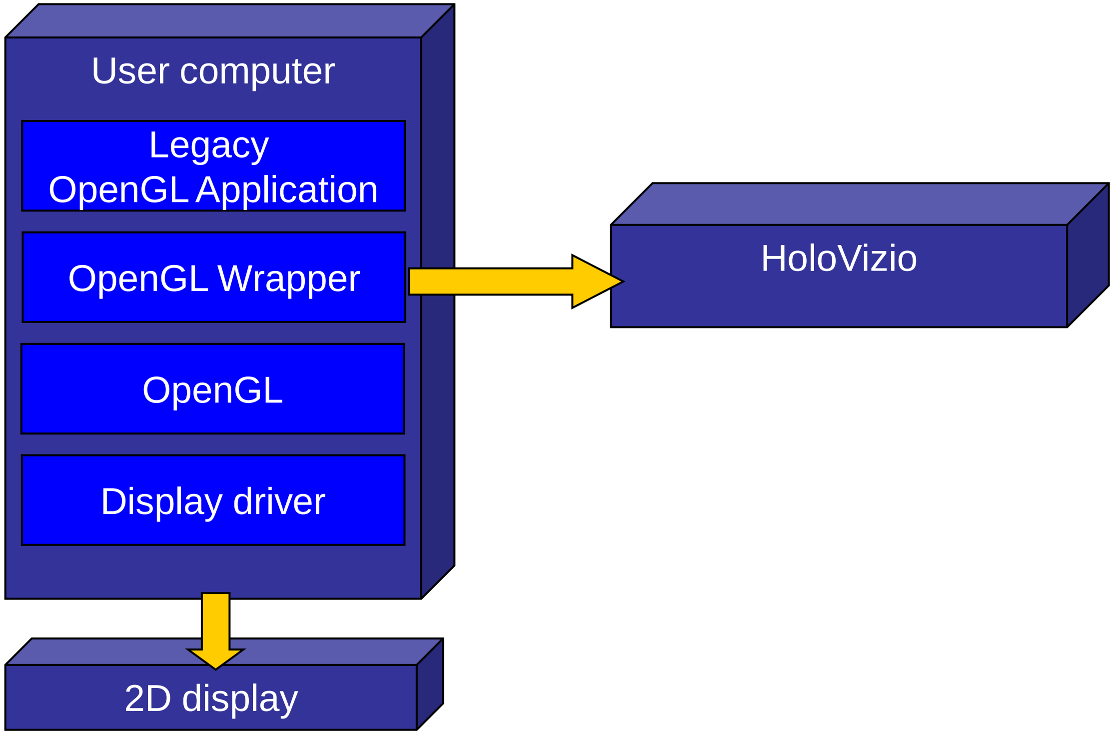

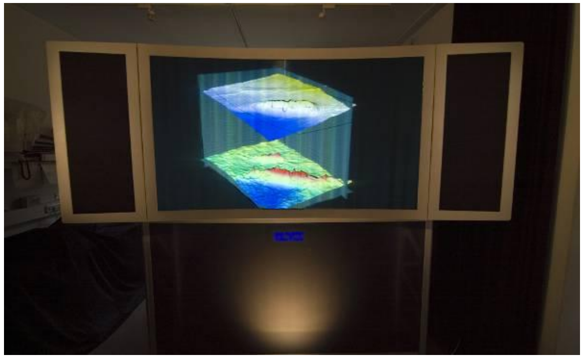

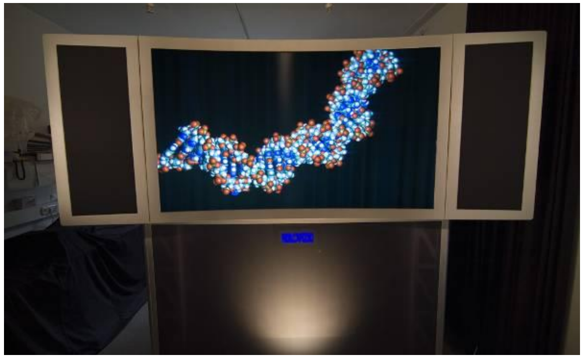

Display Pre-Rendered 3D

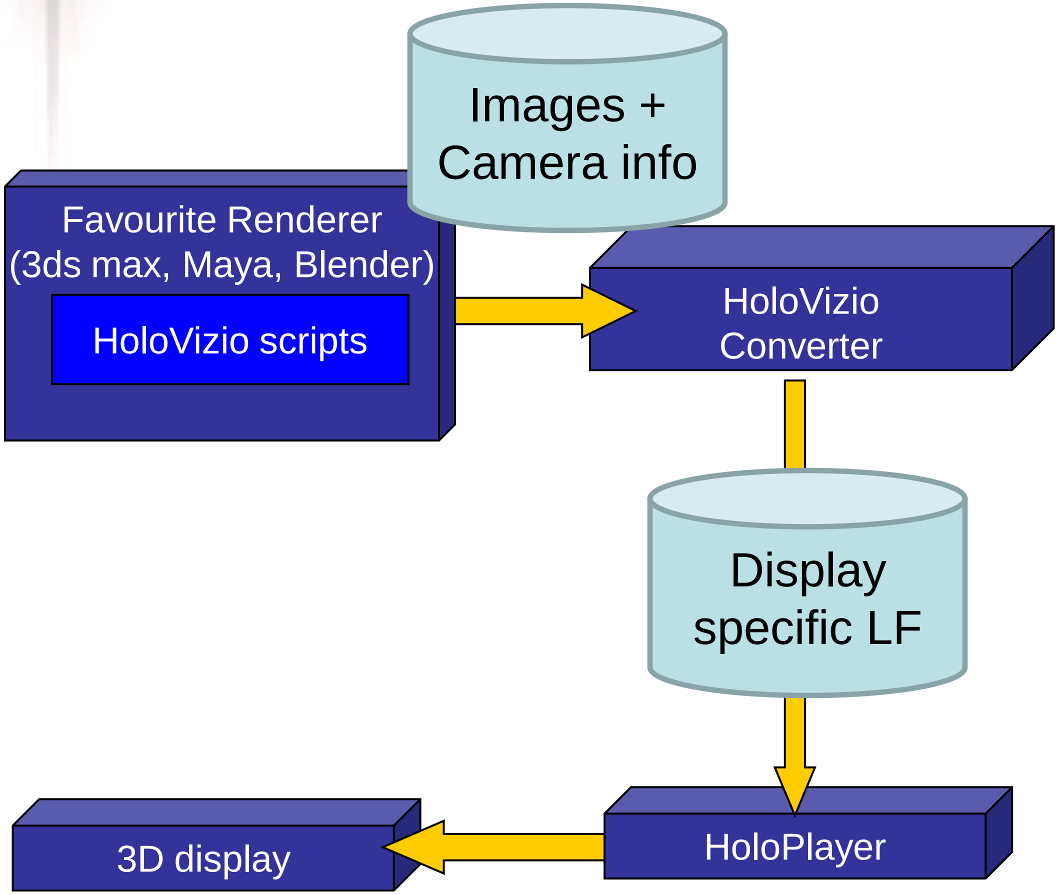

Live 3D transmission

- Linear camera array prototype
  - 27 cameras (more is better)
  - 640x480 at 15 FPS
  - 960x720 at 10 FPS
- Camera calibration
  - Mapping of camera pixels to light rays
- Transmission
  - MJPEG streams from cameras
- Rendering
  - Real-time LF conversion
  - Captured rays to displayed rays
- Display
  - Any HoloVizio display
- Depth-assisted all-in focus rendering results in a dramatic quality improvement

Acquiring image-based 3D data

- Multiple cameras
- Still camera, rotating object
- Still object, rotating camera
- Light-field cameras

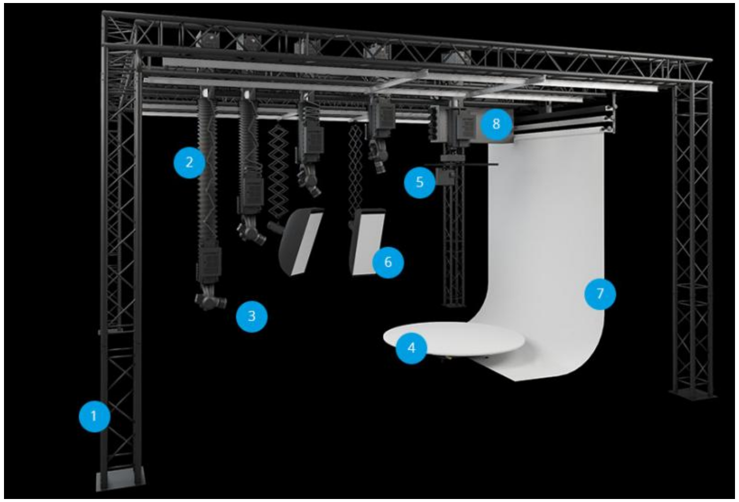

MDS Witikon


Camera array


Packshot creator

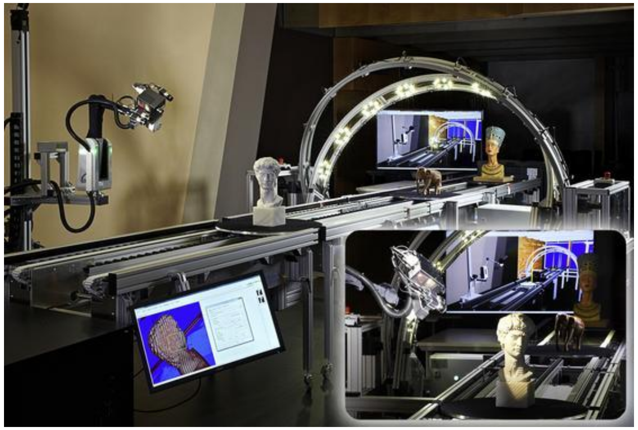
Cultlab3D

A note on baselines and angles

- Light-field cameras
  - Narrow baseline
    IF distance is large

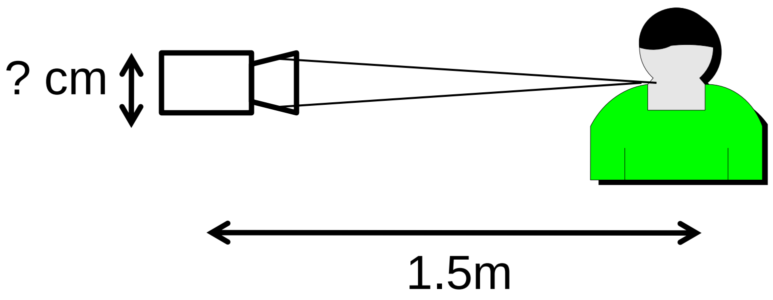

- Camera arrays
  - Wide baseline
  - Wide capture FOV

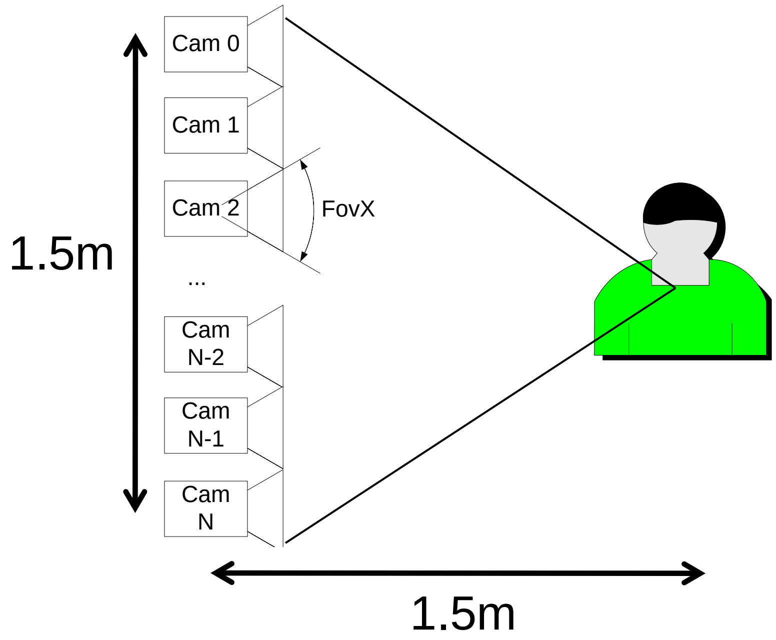

A note on baselines and angles

- Light-field cameras
  - Na
- Cam
  - Wi
  - Wi

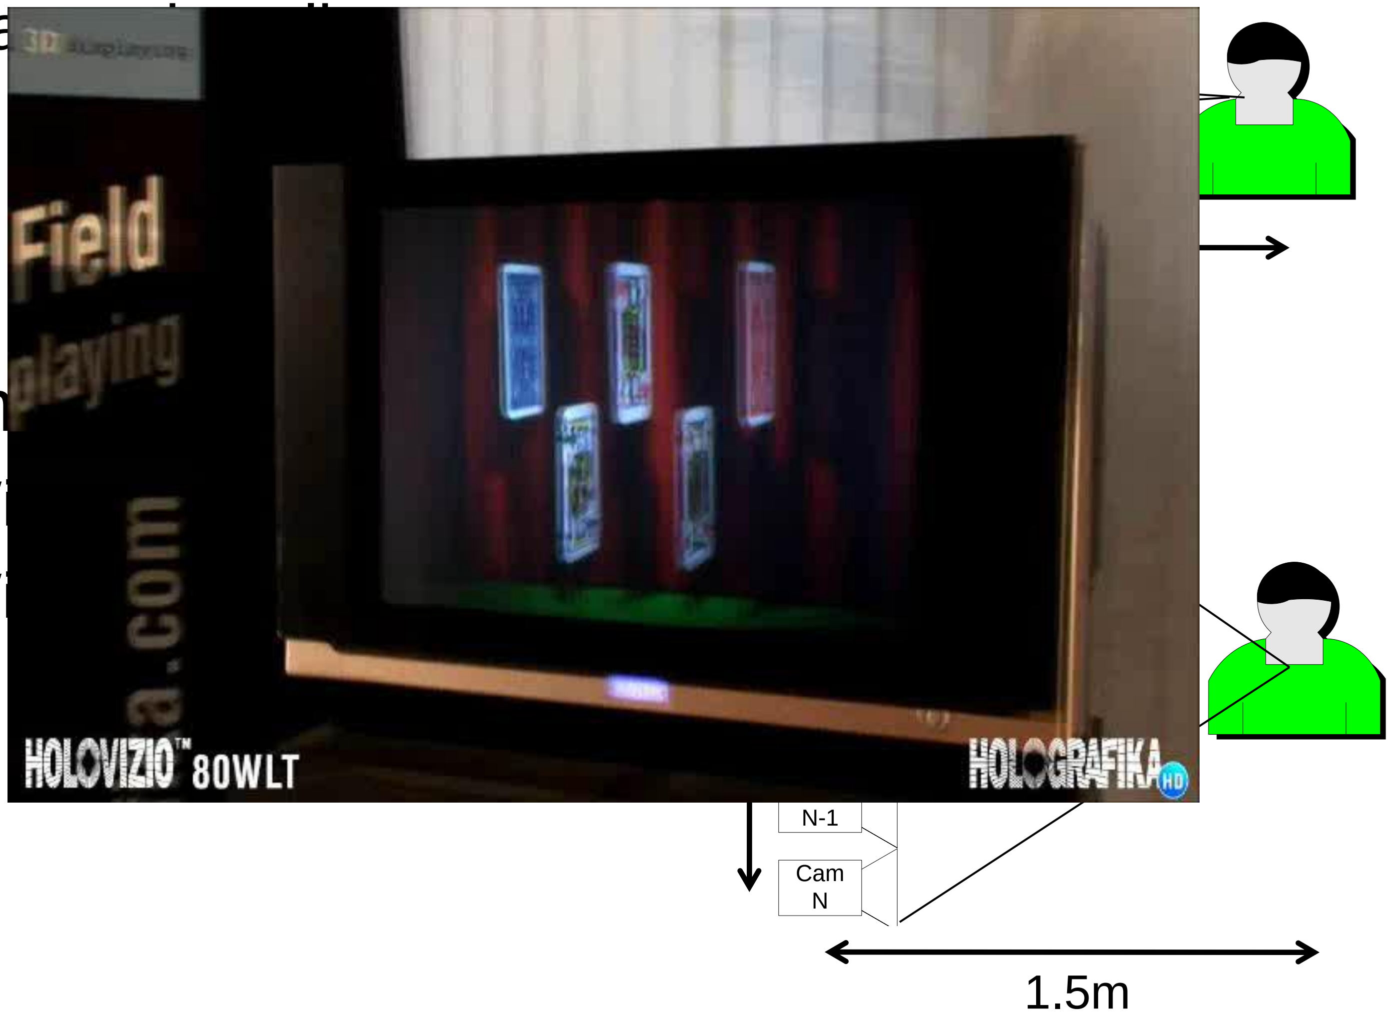

A note on baselines and angles

- Light-field cameras
  - Wide angle if object is close
  - Such applications exist, but it’s not for recording people
- Camera arrays
  - Wide baseline
  - Wide capture FOV

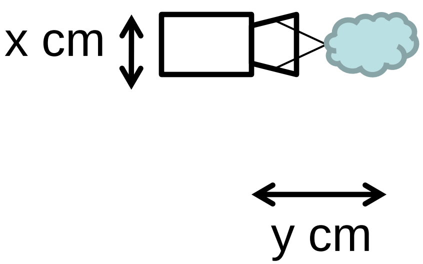

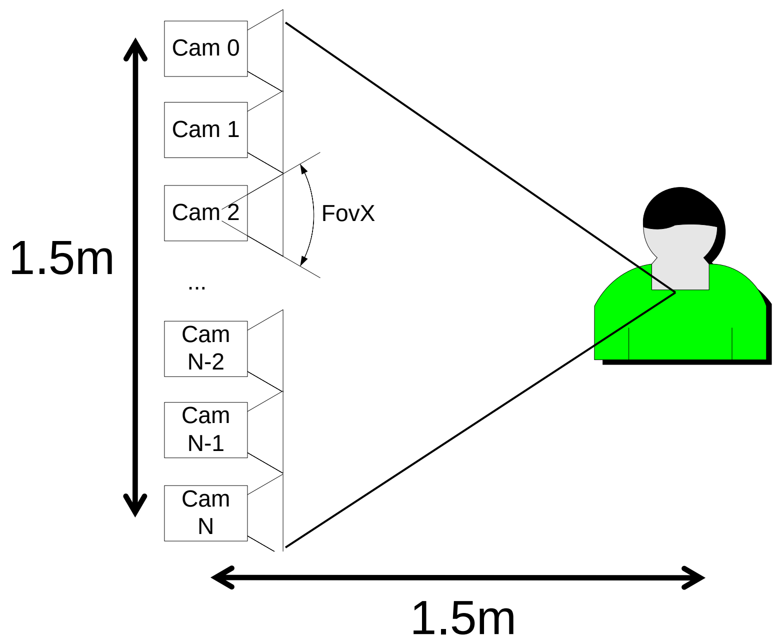

Light-field slices

- The projected image not a 2D *view* of the 3D scene
- Light rays emitted by a single projection module are spread over screen positions and viewing directions, thus cannot be seen from a single viewing position
- A projection module shows a *LF slice*

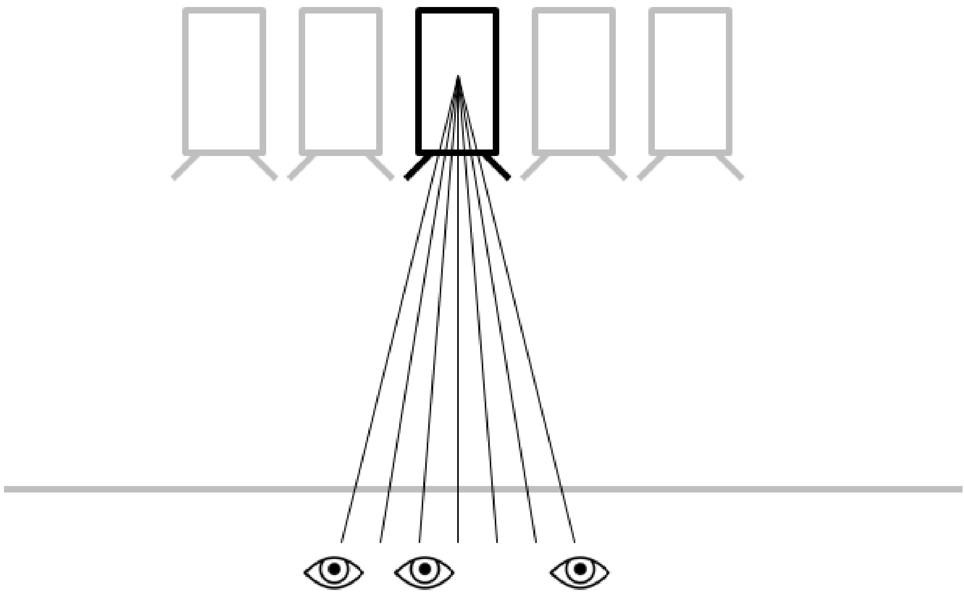

Distributed rendering

- Distributed LF generation
  - High number of pixels / light rays
  - Many inputs to feed directly
- Distributed optical system – distributed rendering
- GPU enabled rendering clusters provide scalability
  - GPU outputs correspond to projection modules
  - One output → one LF slice

    OR
  - One output → adjacent LF slices
  - One rendering node → adjacent LF slices

LF conversion

- Input: many 2D views
- Output: LF slices in display-specific arrangement
- One LF slice is composed of pixels from many views

  and backwards
- One 2D view contributes to many LF slices

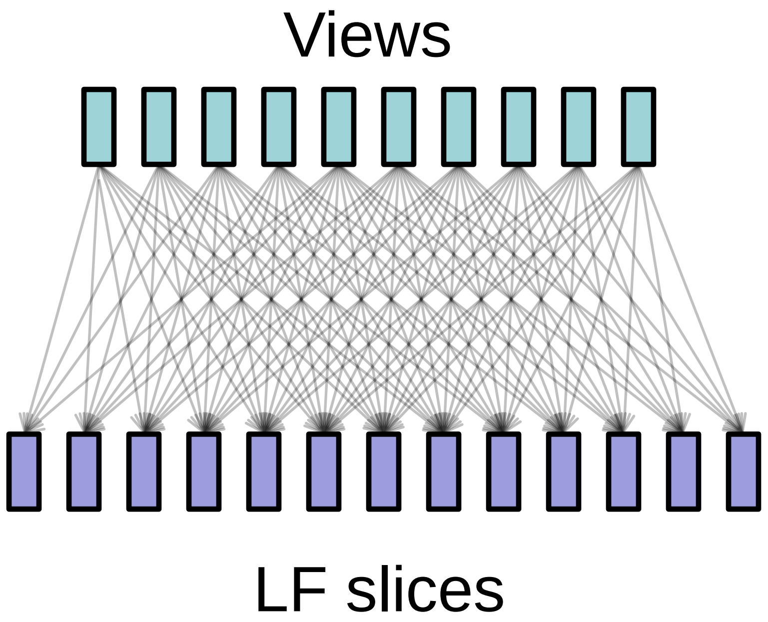

LF conversion

- Input: many 2D views
- Output: LF slices in display-specific arrangement
- One LF slice is composed of pixels from many views

  and backwards
- One 2D view contributes to many LF slices

- Display independent LF

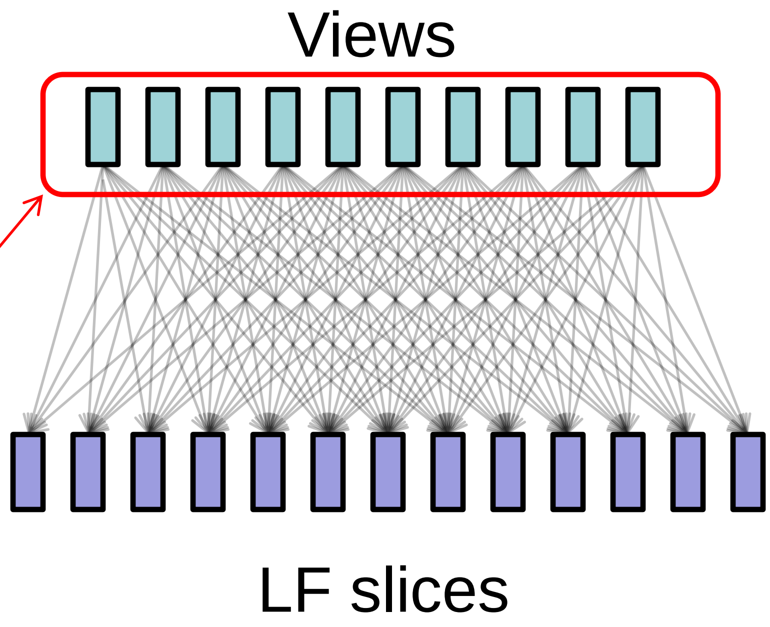

LF conversion

- Input: many 2D views
- Output: LF slices in display-specific arrangement
- One LF slice is composed of pixels from many views

  and backwards
- One 2D view contributes to many LF slices

- Display independent LF
- Display specific LF

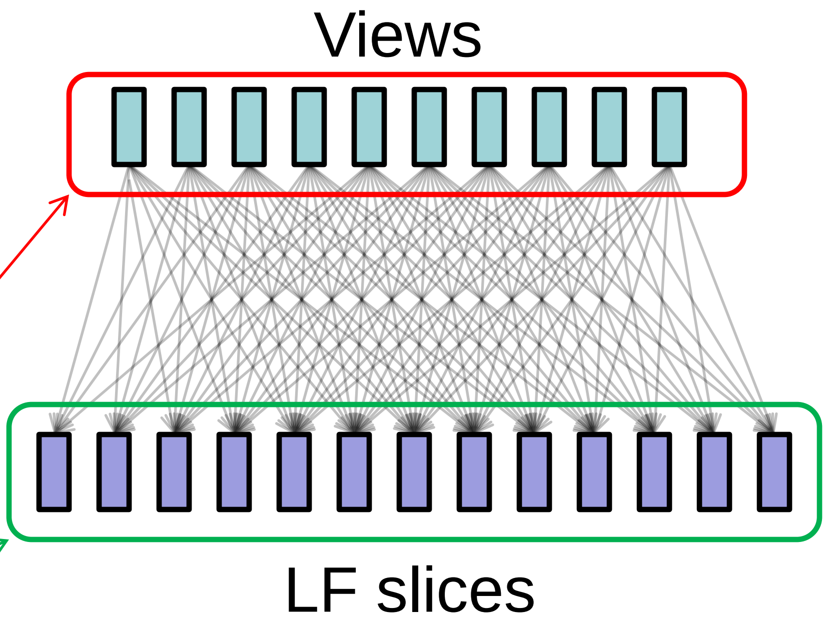

Use cases

- Playback of pre-processed LF content
  - Stored in display-specific LF
  - LF slices correspond to the layout of the distributed rendering and projection system
  - Example: 80 x WXGA LF slices
  - Visualized directly on the targeted display
- Real-time LF video transmission
  - Target different displays
  - Different Field Of View
  - Stored in display independent LF (views)
  - Needs to be converted to display-specific LF at the end user in real-time

Compressing display-specific LF

- Simplest solution is to simulcode LF slices (e.g. 80x1)
  - Decode them in parallel on rendering nodes independently
  - Decoding e.g. 8 times H.264, JPG or JPG2000 is possible on a multi-core CPU
  - This approach works today
  - Does not exploit similarities between LF slices
- Compress all LF slices into a single multiview stream (e.g. 1x80)
  - Single processing node should be able to decompress the stream
  - Decoding complexity increases with number of views
  - Decompression in real-time is too resource intensive
- Compress N adjacent LF slices (e.g. 10x8)
  - Reasonable tradeoff that could map well to the distributed rendering cluster

Processing display-independent LF

- Processing node needs to
  - Decompress views
    - Decompressing all views is too much
  - Perform LF conversion
- Not all views are required by a single node
  - But many
- Number of views used depends on display's FOV
  - How many views we need to decompress for the whole display?
  - Simulated LF displays with a range of FOVs

FOV (degrees)	27	38	48	59	69	79	89
No. views used	42	44	46	48	50	52	54

  - Decompressing this many views is still prohibitive
  - How to reduce decoding workload?
  - Not all pixels are used from the views


View


Nodes 4,5,6


Nodes 0,5,9

Processing display-independent LF

- Not all pixels are used from the views
- Why decode pixels that are left unused?
  - Why not just transmit parts that are actually used?
- Used area increases with FOV
  - Still, large portions could be skipped
  - To support different FOVs, in-view scalability is needed
- Slicing / tiling the image vertically could do the trick
  - With a specific granularity
  - By losing some coding gain
  - If the encoder / decoder supports this kind of slicing
  - If the decoder supports skipping slices

27º FOV

59º FOV

89º FOV

Nonlinear camera setups

- A linear setup cannot capture the scene from all around
  - Think of a 180º FOV 3D display
- Inter-view similarity is less / different with arc setups
  - MVC, MV-HEVC, 3D-HEVC are designed for parallel linear cameras
- With wide-angle LF displays, compressing views from arc rigs is important

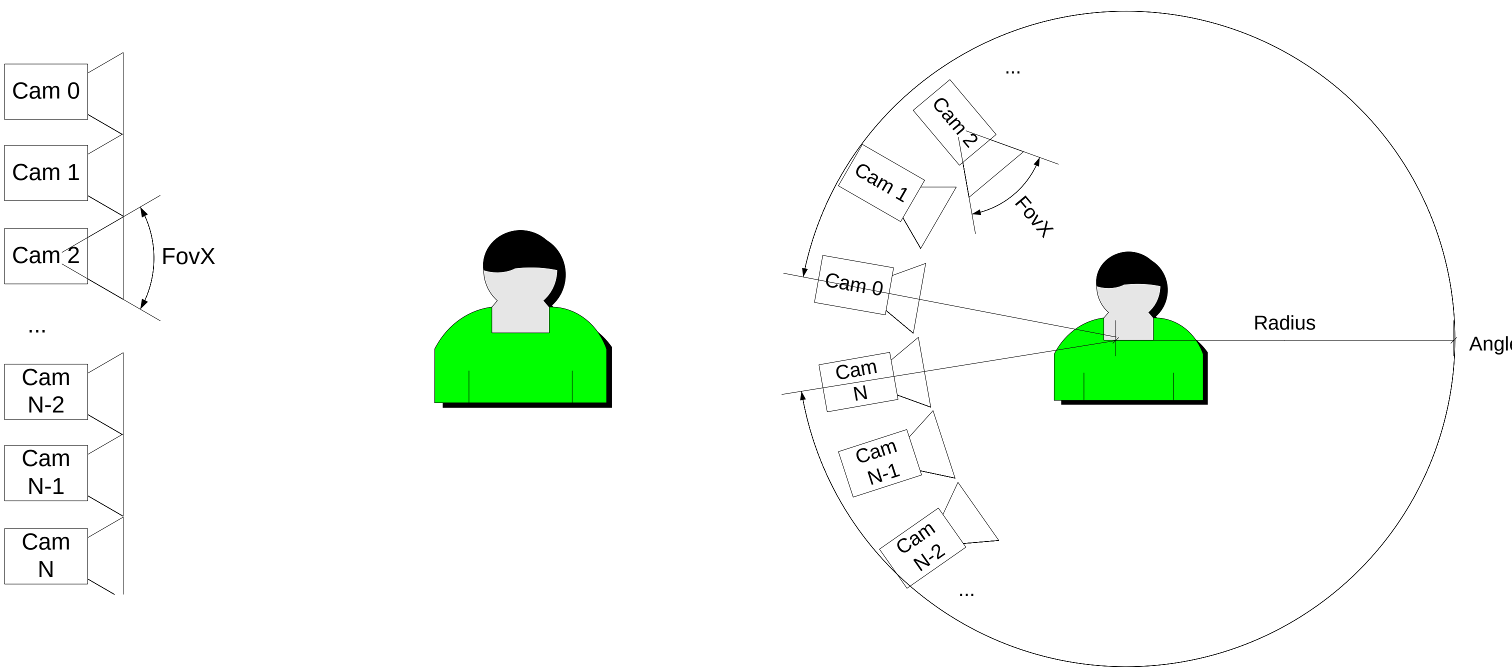

3D Light Field Content

- Computer generated animation – *solved*
  - Just rendering time, avoid 2D after effects, etc.
  - Demonstrated 1.5 minutes of Big Buck Bunny
- Live scene acquisition
  - Wide baseline camera arrangements
    - Free-viewpoint TV
  - TV programs - talk shows, permanent studio arrangements - can be solved by multiple cameras or camera arrays
  - 3D cinema - film shooting equipment to be developed
    - professional cameras getting more compact
- Philosophical and practical question

  Full acquisition (then extract information) & view reconstruction vs. sampling (loss of information) & view synthesis for production quality

3D Light Field Content

- Scalability
  - 2D
  - 2 View (stereo)
  - Narrow angle (multi-view)
  - Wide angle (Light Field)

- Backward compatibility
  - 2D – base layer
  - 3D – extraction of various number of views
- Future 3D Light Field format and standard
  - Similar to variety of resolutions and aspects today, generalized approach on the angular range
  - Narrow angle, wide angle content, up to 180/360 degrees, continuous parallax, open to full parallax

Requirements and wishes

- Desired functionality
  - Decode view individually with low overhead
  - Decode views partially with low overhead, likely starting from the scene center (slices, tiles)
  - Nonlinear (arc, circular) multi-view coding tools
- Real-time implementation!
- Develop in-between representation with better localization of rays
  - Something between views and LF slices?

Way forward

- New display models
  - Brighter, higher resolution, wider viewing angle, image quality improvements
- Towards Gigapixel displays, custom displays
  - 100Mpixel today, high pixel count systems with parallel distributed approach
- Low-level API
- Live 3D acquisition
  - 3D camera system
- 3D video formats
  - Scalable and generic representation format
  - 3D Light Field representation
- 3D Telepresence
  - Natural communication, eye contact
  - Performance and latency

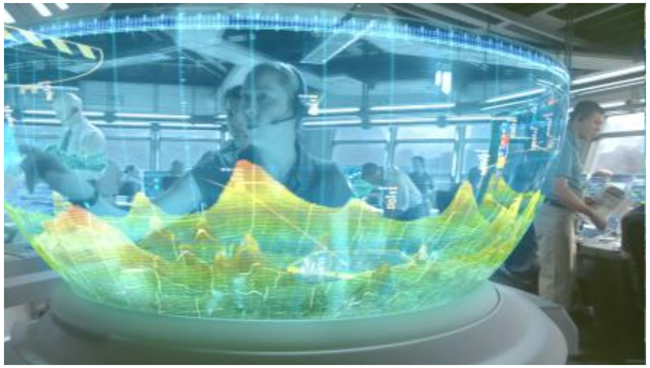

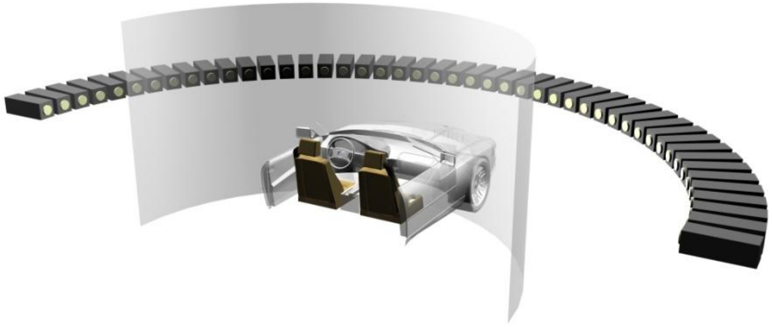

We are hiring

- Two Marie Curie positions open
- Experienced Developer / Researcher in Hardware Design / Electrical Engineering
  - 16 months
- Early Stage Developer / Researcher in 3D video coding for light-field displays
  - 36 months
- In Budapest, Hungary
- Check www.holografika.com or contact me for details

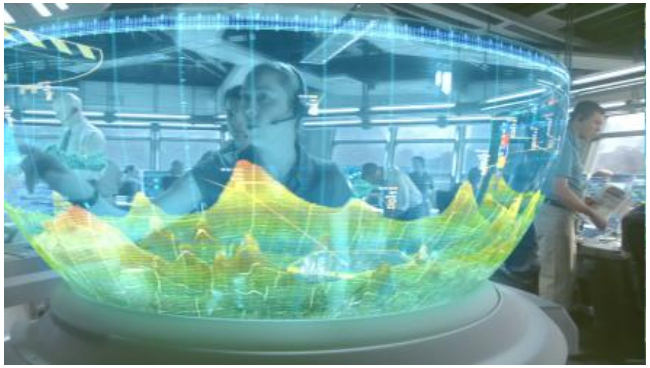

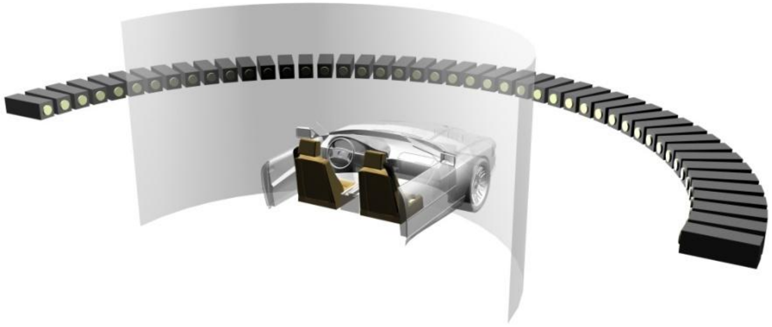

Questions?

www.holografika.com
p.kovacs@holografika.com