GPU TECHNOLOGY CONFERENCE

Building a Life-Size Automultiscopic Display Using Consumer Hardware

Andrew Jones, Jonas Unger*, Koki Nagano, Jay Busch, Xueming Yu, Hsuan-Yueh Peng, Oleg Alexander, Paul Debevec

USC Institute for Creative Technologies
*Linköping University

Stereoscopic

Automutiscopic

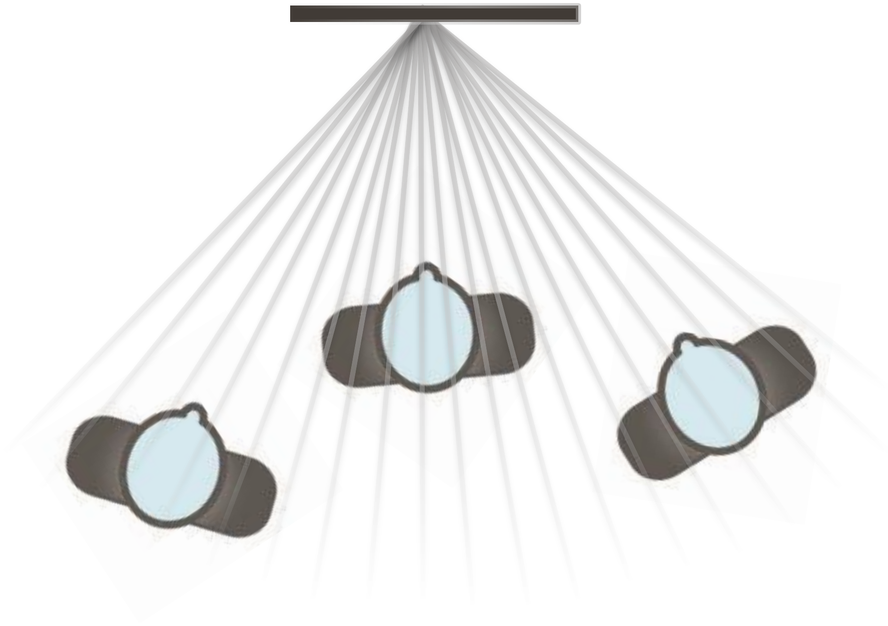

Automutiscopic

How do we capture, render, display automultiscopic content?

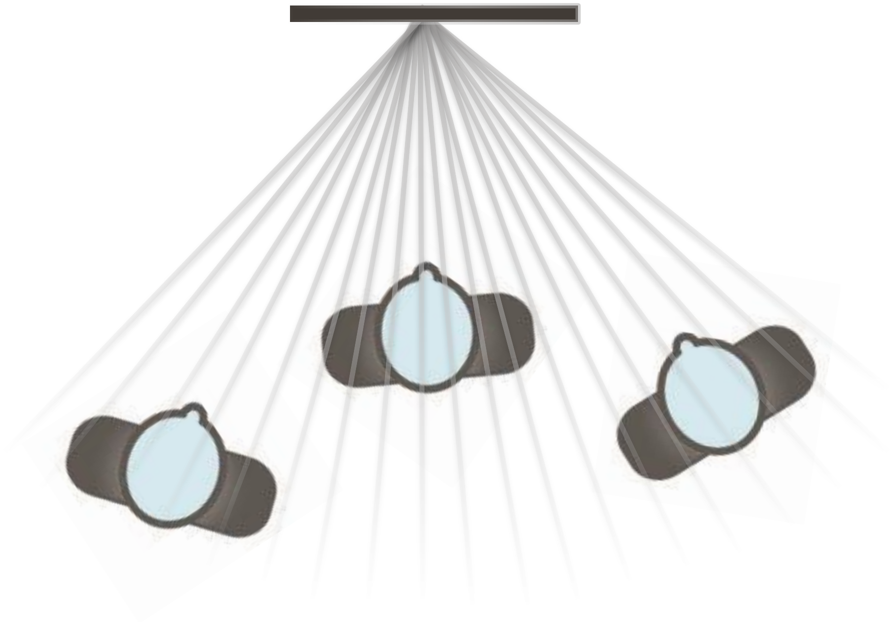

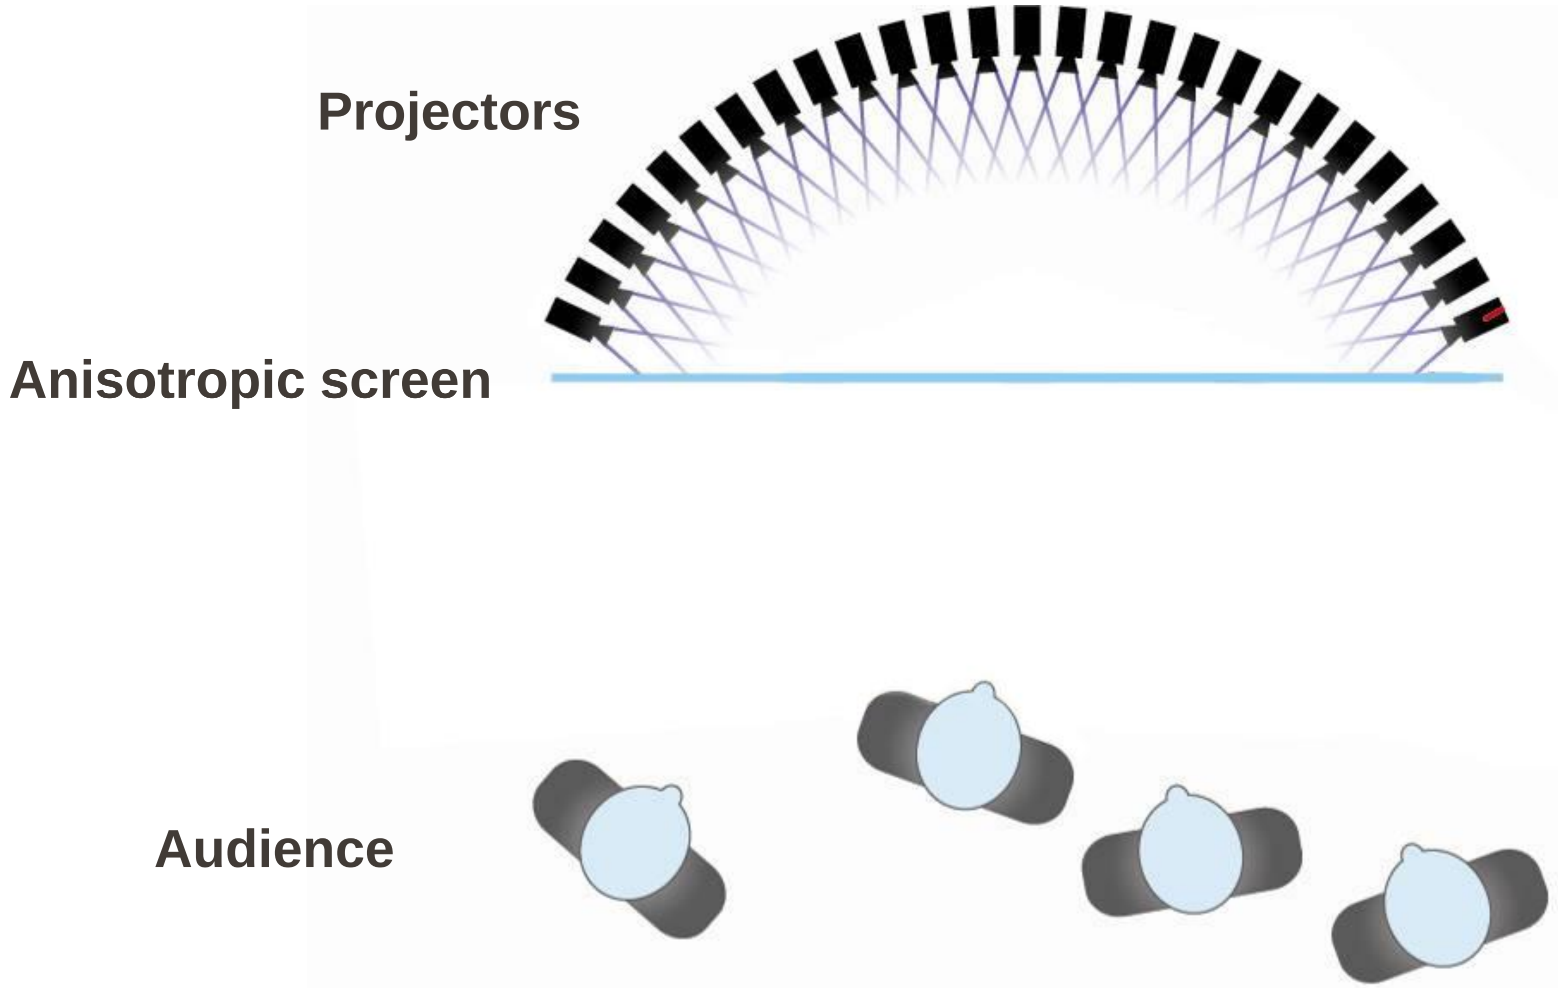
Projectors
Anisotropic screen
Audience

1st prototype
Focus on face

2nd prototype
Full-size bodies

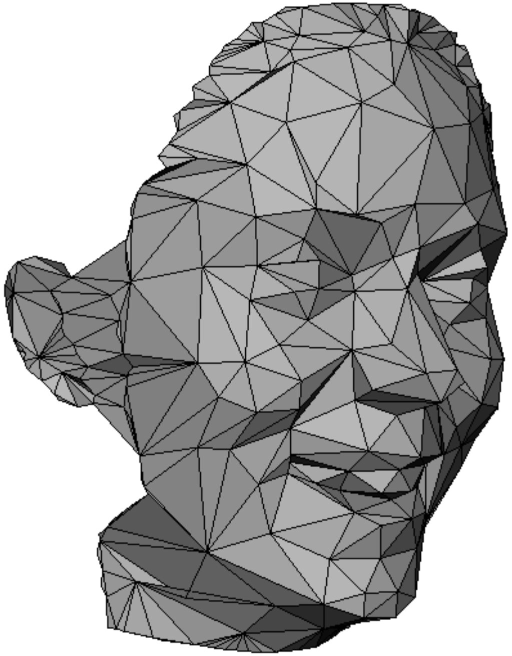

3D Geometry

custom vertex shader

Image-based Light Fields

custom pixel shader

Bandwidth

1920 x 1080 x 60 fps x 360º x 24 bit = **134GB / sec**

Large number of output streams

Data transfer to GPU

Our Approach

- Distribute rendering across multiple GPUs and computers
- Scalable, additional projectors increases field of view

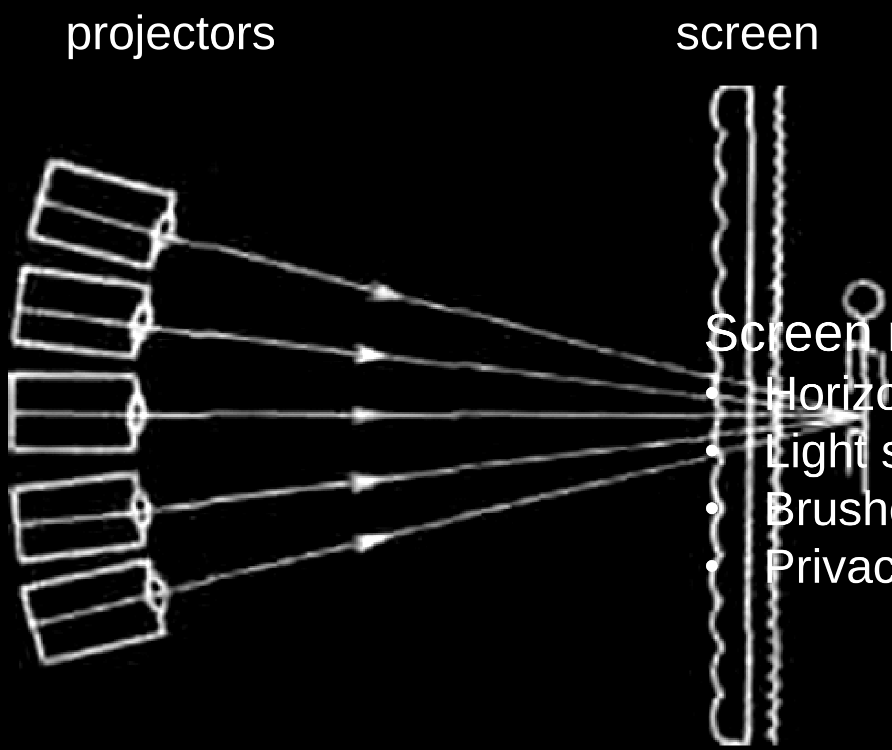

Screen materials:

- Horizontal lenticular
- Light shaping diffusers
- Brushed metal
- Privacy filters

Takanori Okoshi, *Three-Dimensional Imaging Techniques*, Academic Press 1976
Fig. 5.5(b), “projection-type three-dimensional display”, p. 131

Anisotropic Projector Arrays

[Agocs et al. 2007]

[Kawatika et al. 2012]

[Yoshida et al. 2011]

GPU TECHNOLOGY CONFERENCE

Projector Array

- 72 TI DLP Pico
  - 480 x 320 Resolution
  - Mini HDMI input
- 1.66° Angular Resolution
- 110° Field of View

Anisotropic Screen

- 40 lines per inch Lenticular screen from Microlens Inc.
- 1° horizontal x 60° vertical diffuser from Luminit Co.

Graphics Cards

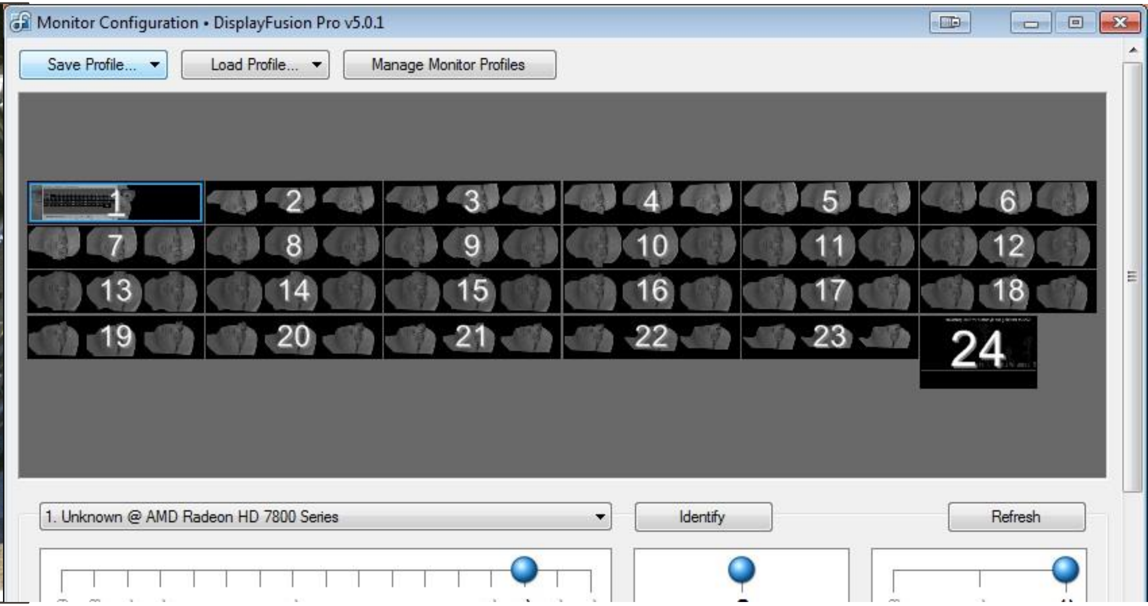

AMD Radeon 7870 graphics cards,

4 x 6 Mini DisplayPort outputs = total **24** outputs

DisplayFusion (nView, Ultramon)

Video Splitters

24 Matrox TripleHeadToGo video splitters

– 1 DisplayPort input, 3 DisplayPort outputs each

DisplayPort 1.2

- Multi-Stream Transport (MST)
- Appear as separate displays
- Each display can have different resolution/refresh rate etc
- Each graphics card still has upper bound for total number of streams

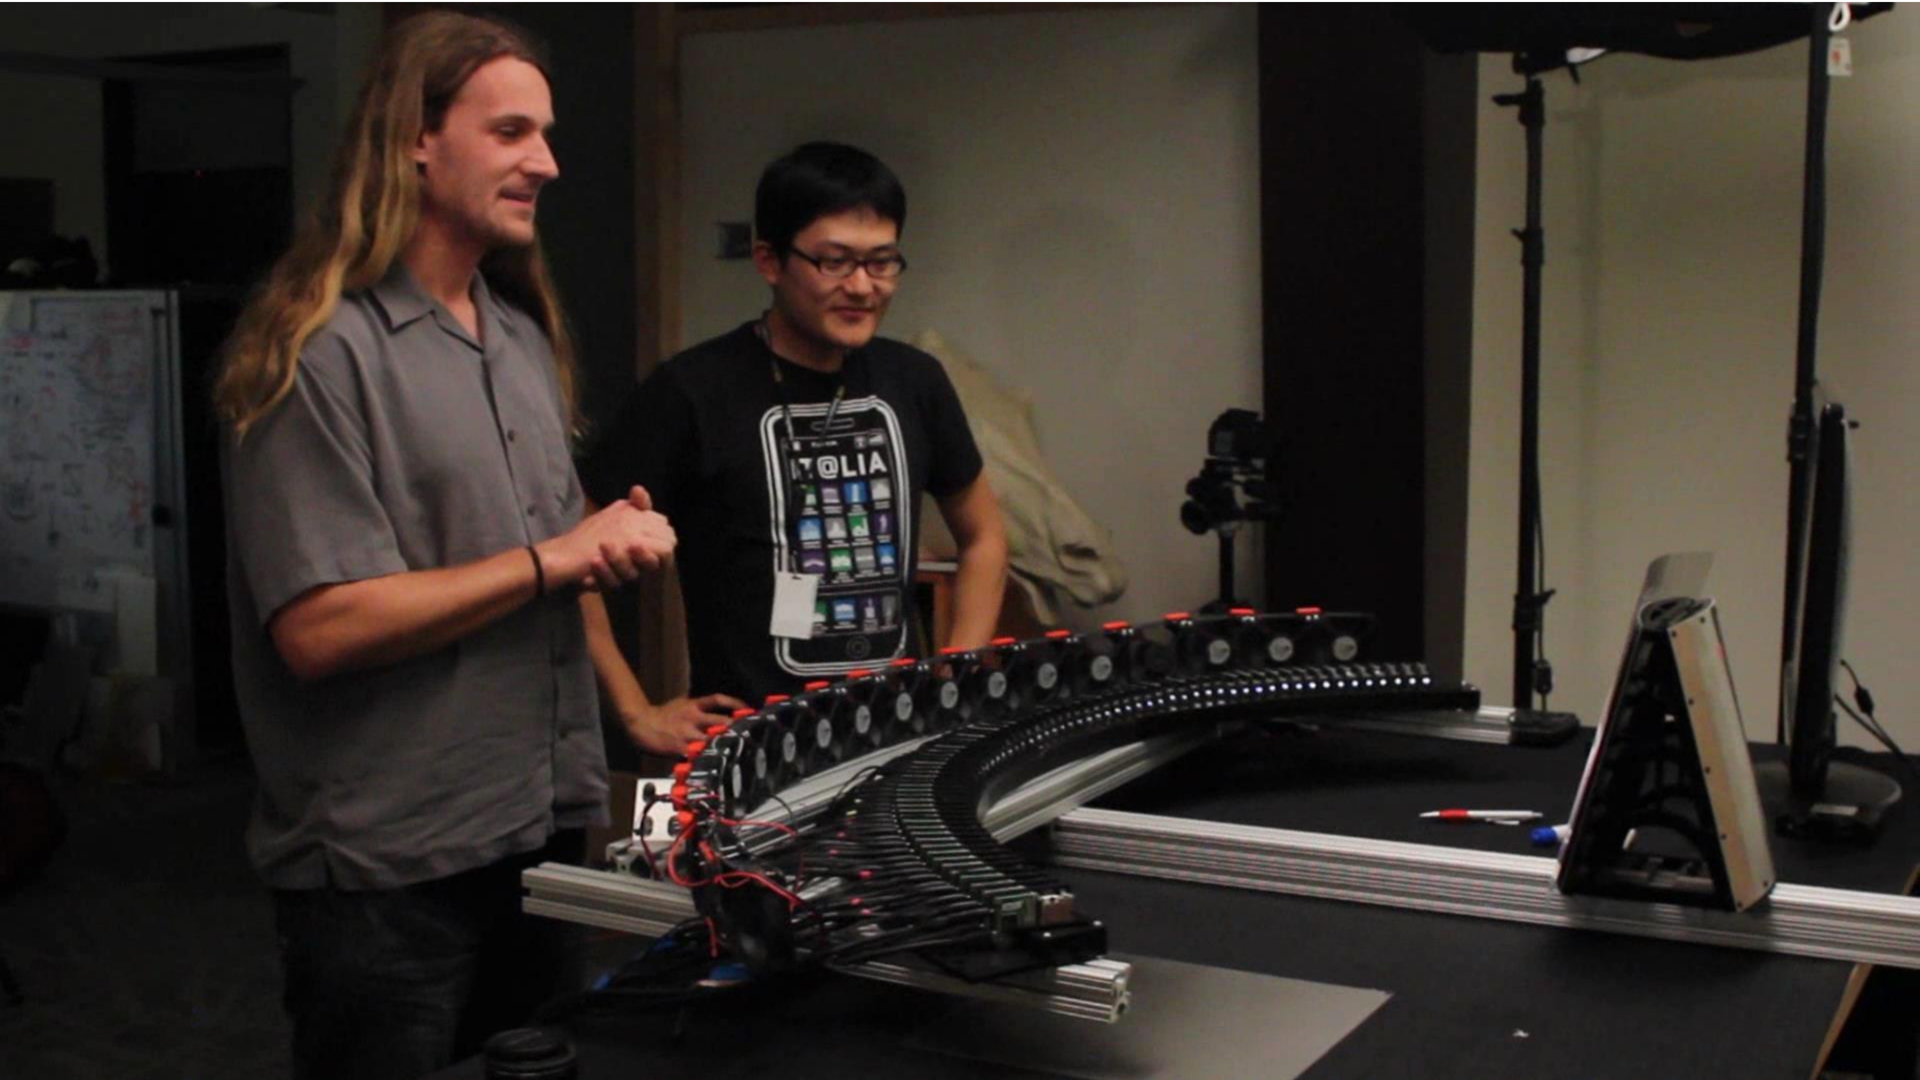

Multiple-center of projection

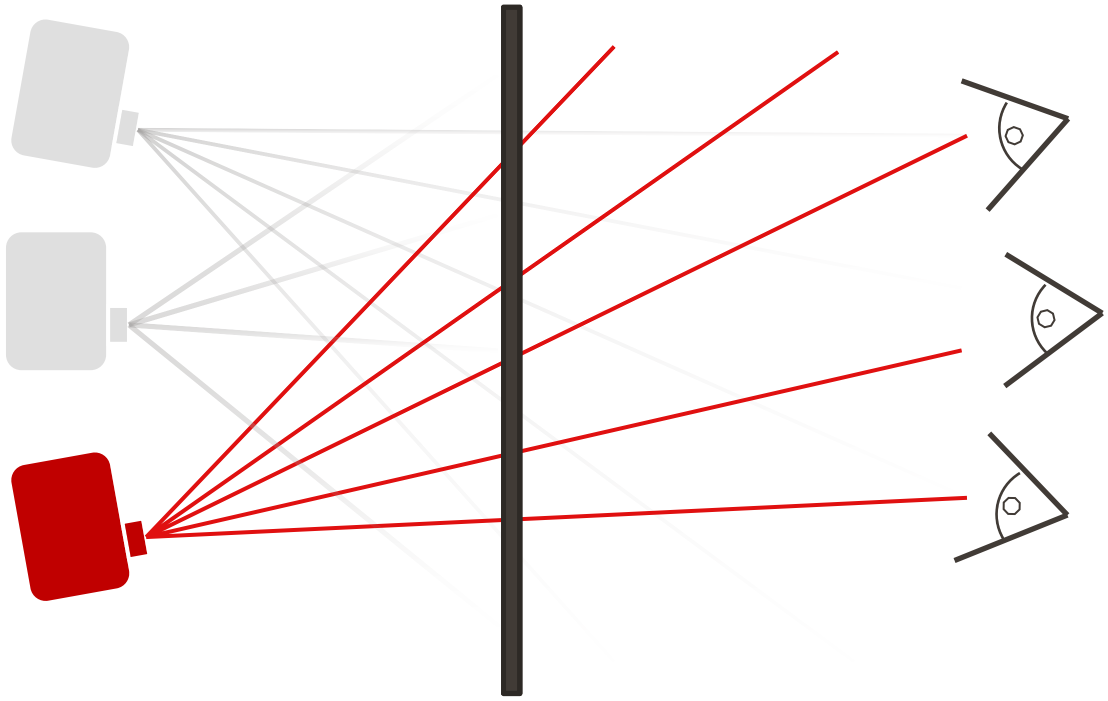

Every pixel rendered from different viewpoint

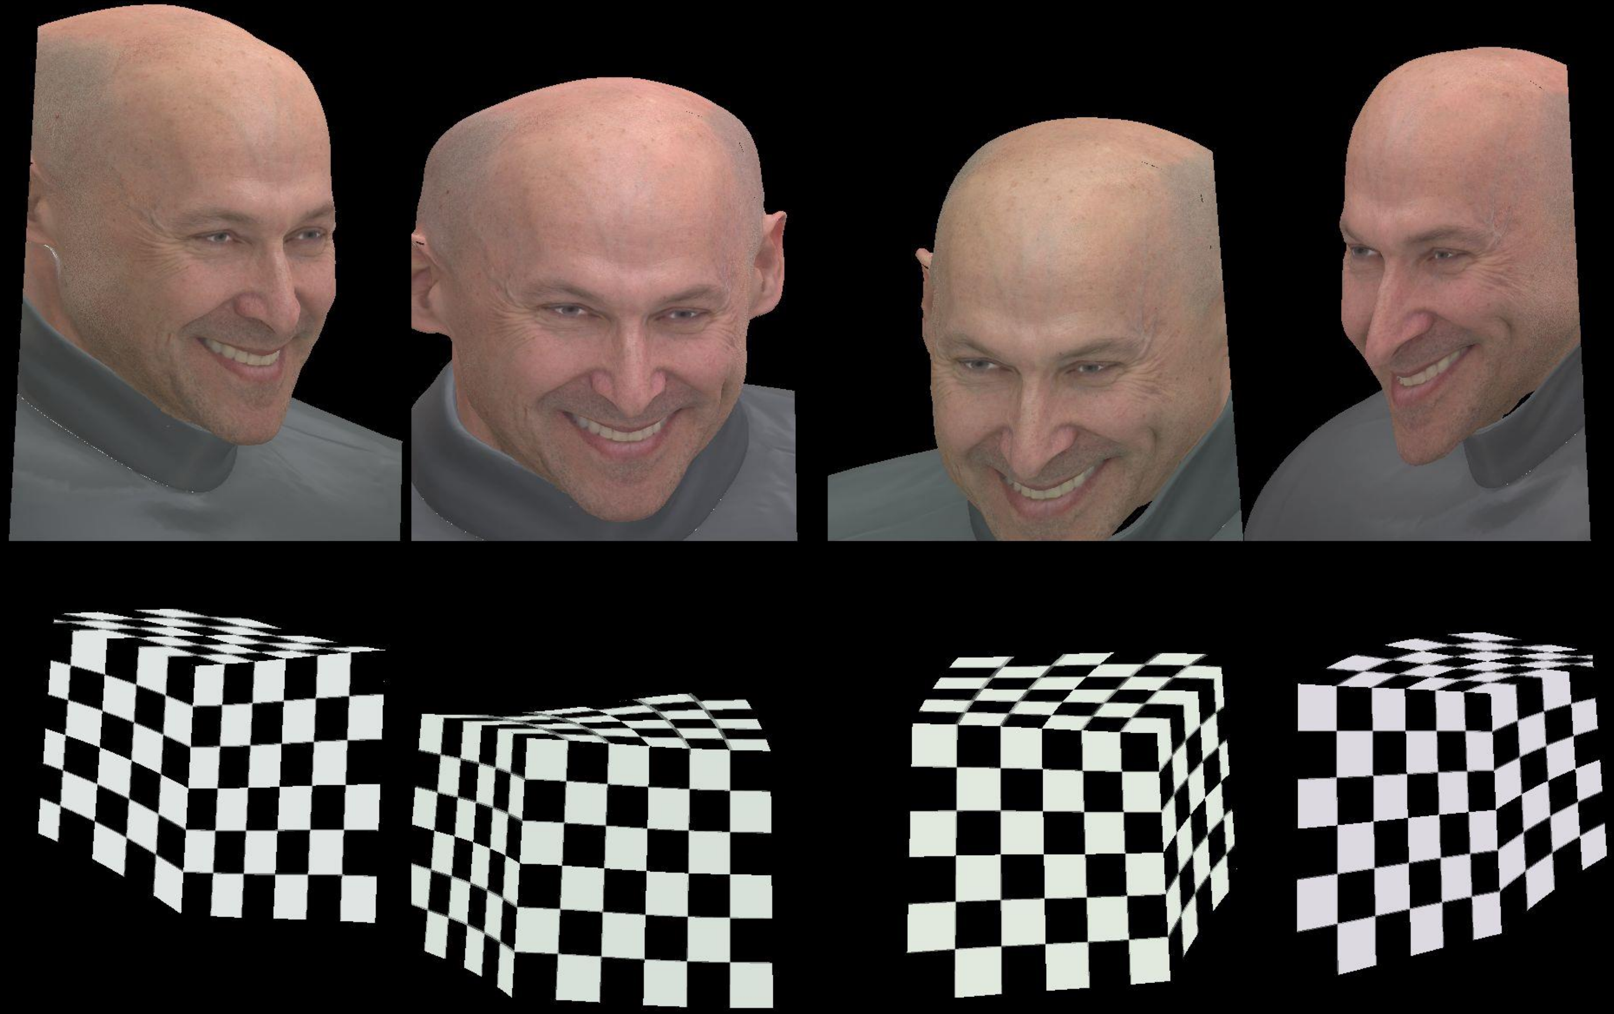

Vertex projection

- For each vertex, find corresponding viewer
- Project back onto screen from view point

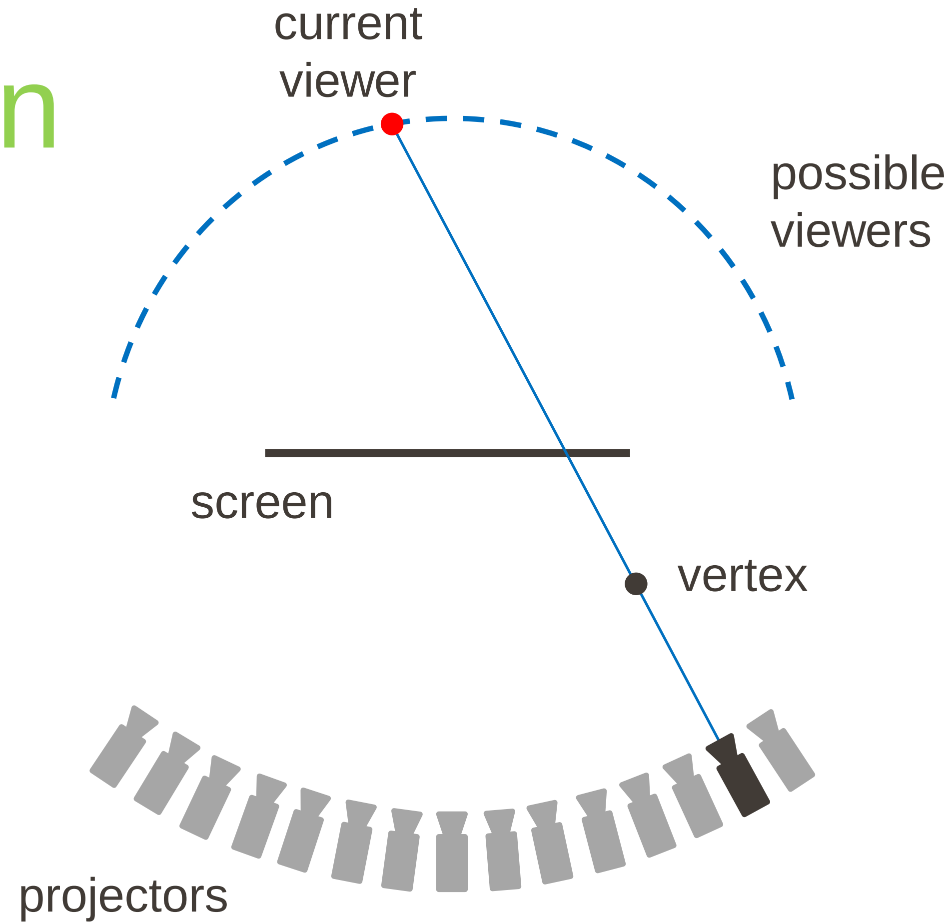

Vertex projection

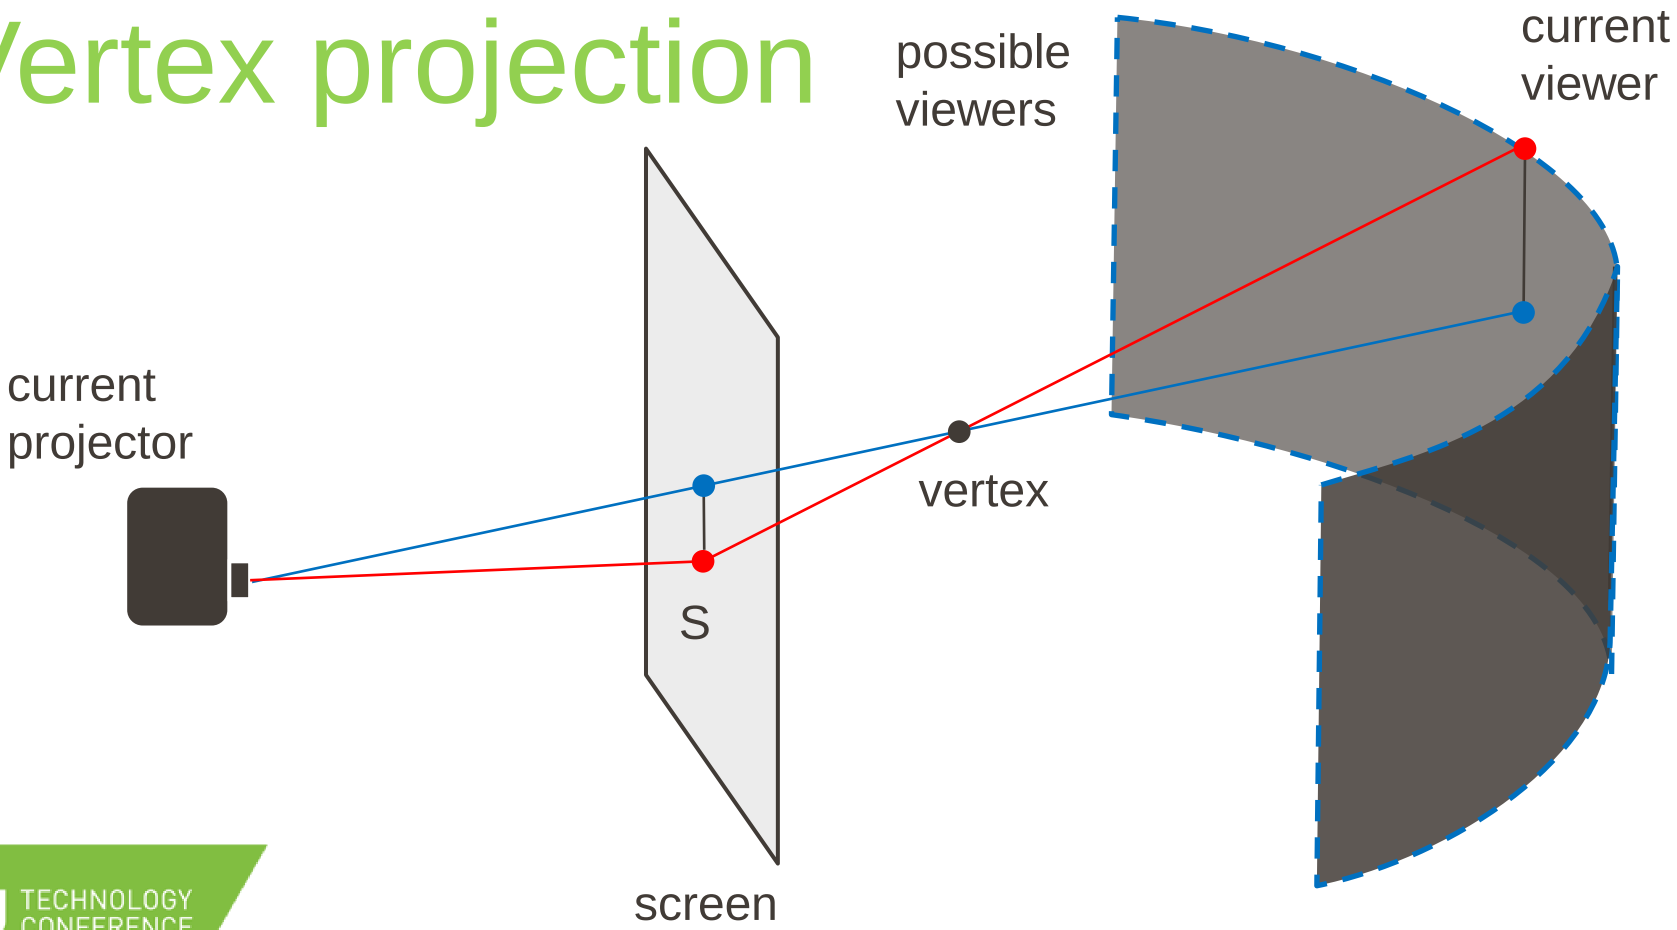


INK

Multiple viewers

- Sum of weighted Gaussians
- Can revert back to default height and distance
- Falloff distance ≈ width of shoulders

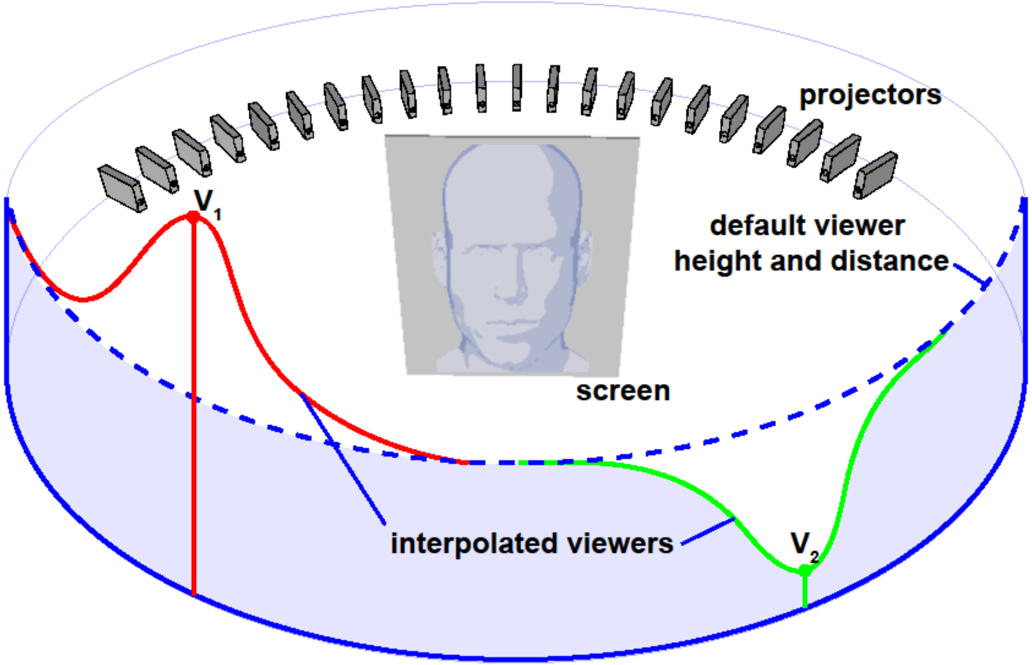

Anisotropic Projector Arrays

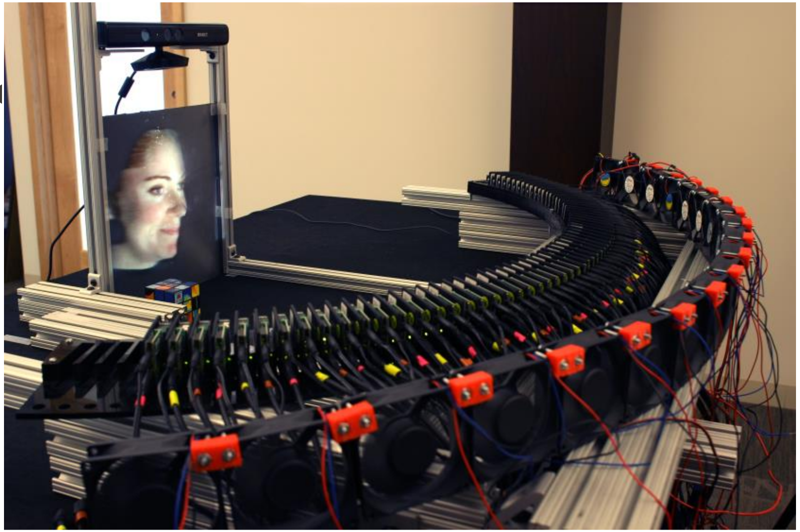

Jones et al. “Interpolating Vertical Parallax for an Autostereoscopic 3D Projector Array”. SPIE Stereoscopic Displays and Applications 2014

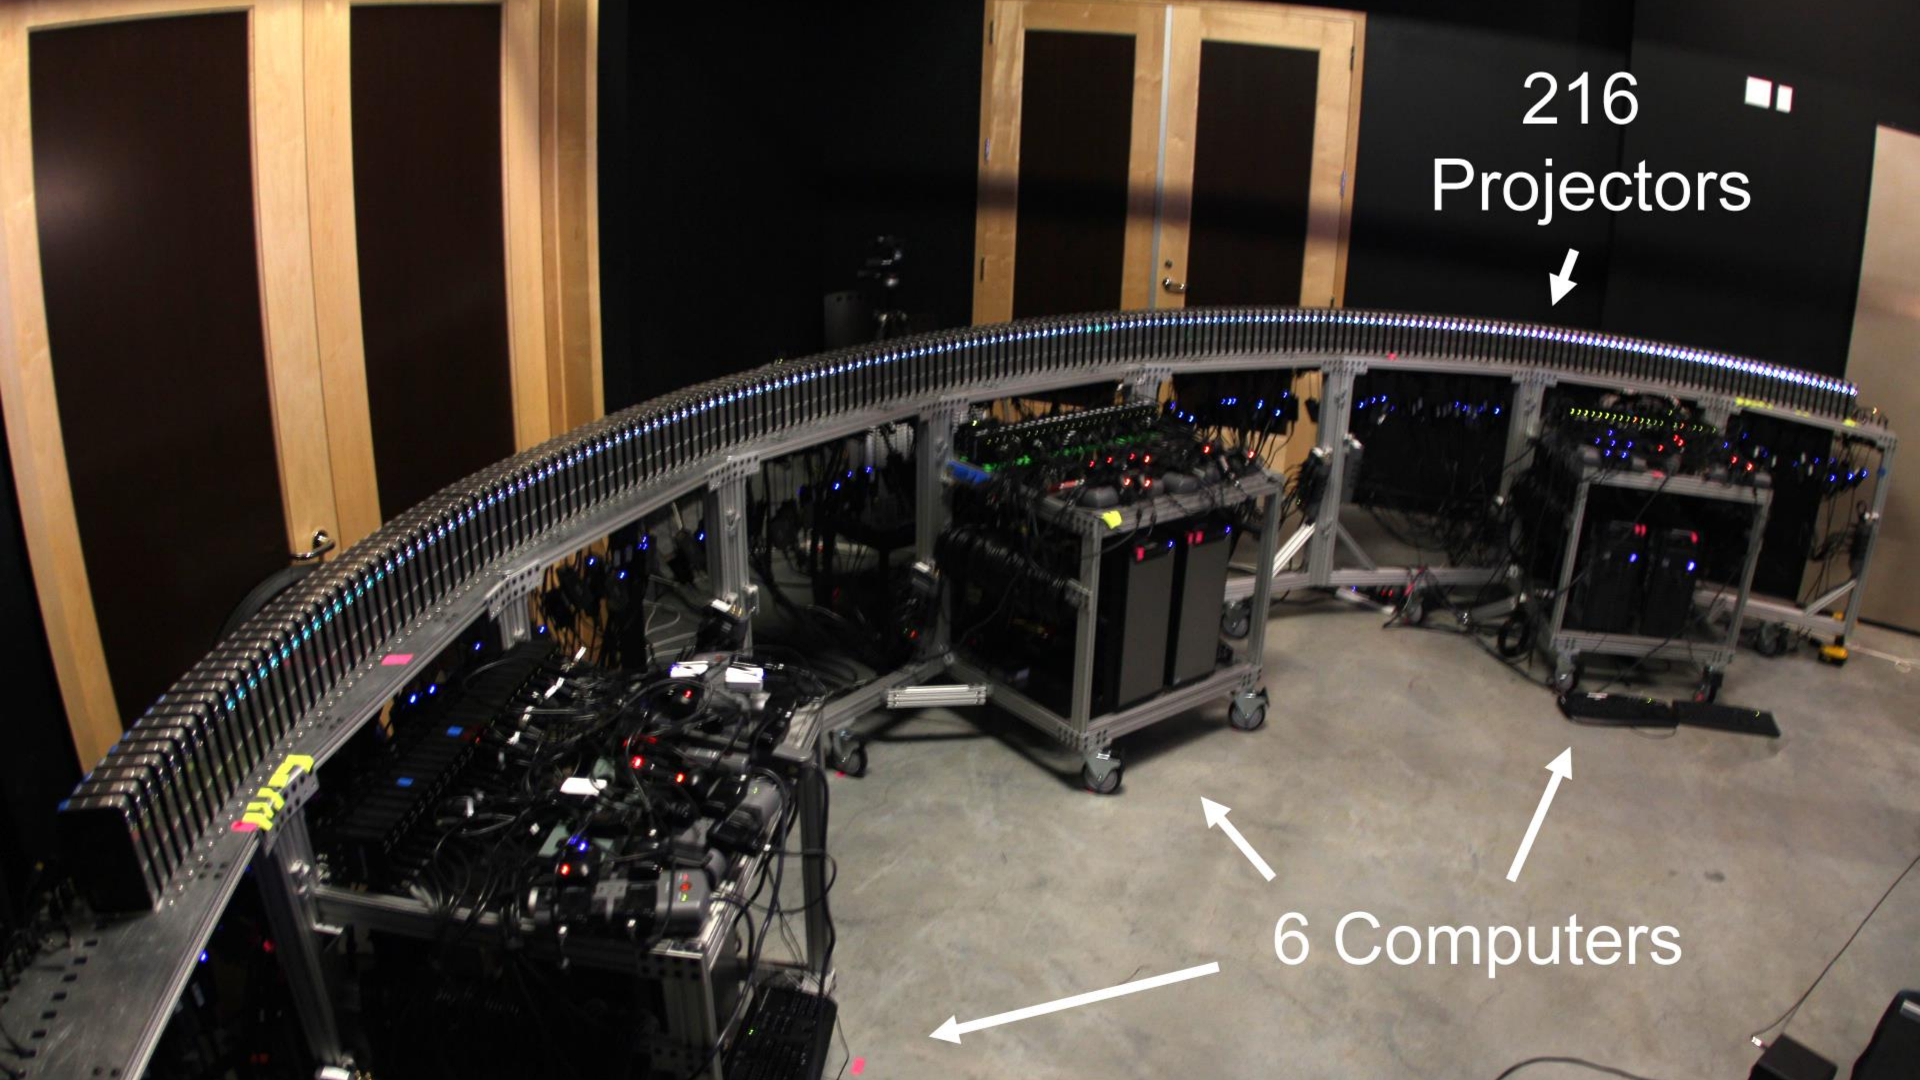
216
Projectors
6 Computers

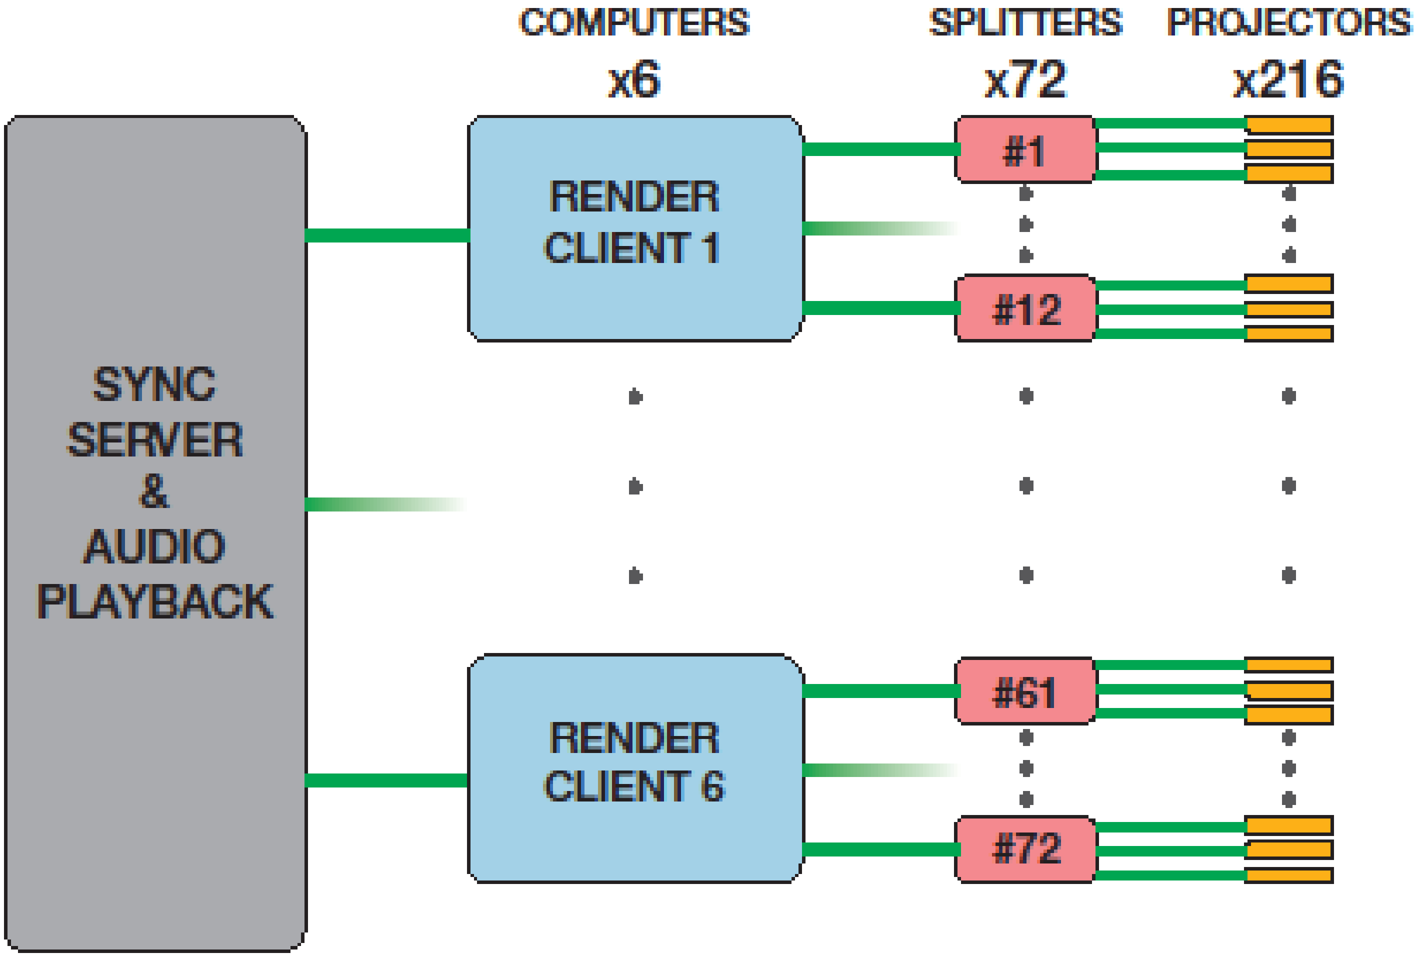
COMPUTERS
x6
SPLITTERS
x72
PROJECTORS
x216
SYNC
SERVER
&
AUDIO
PLAYBACK
RENDER
CLIENT 1
RENDER
CLIENT 6
#1
#12
#61
#72

Vivitek Qumi projectors

- 1280 x 800 pixels
- LED light source
- 300 Lumens
- Low power, small size
- ~$300 each

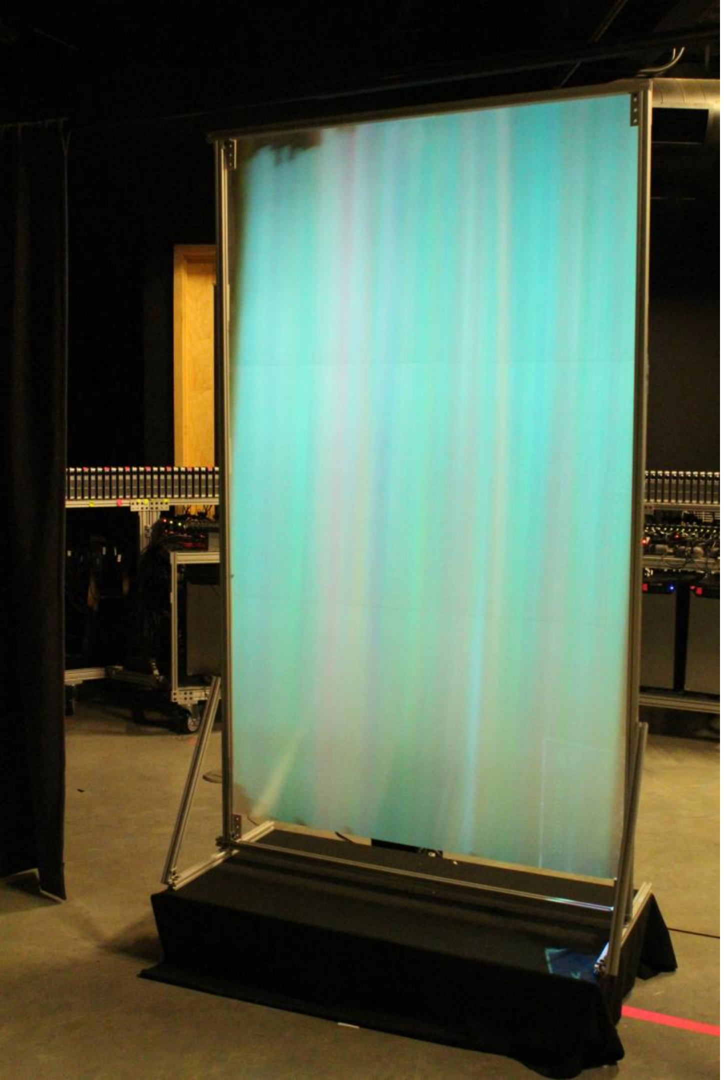

The Anisotropic Screen

1° horizontal x 60° vertical diffuser from Luminit Co

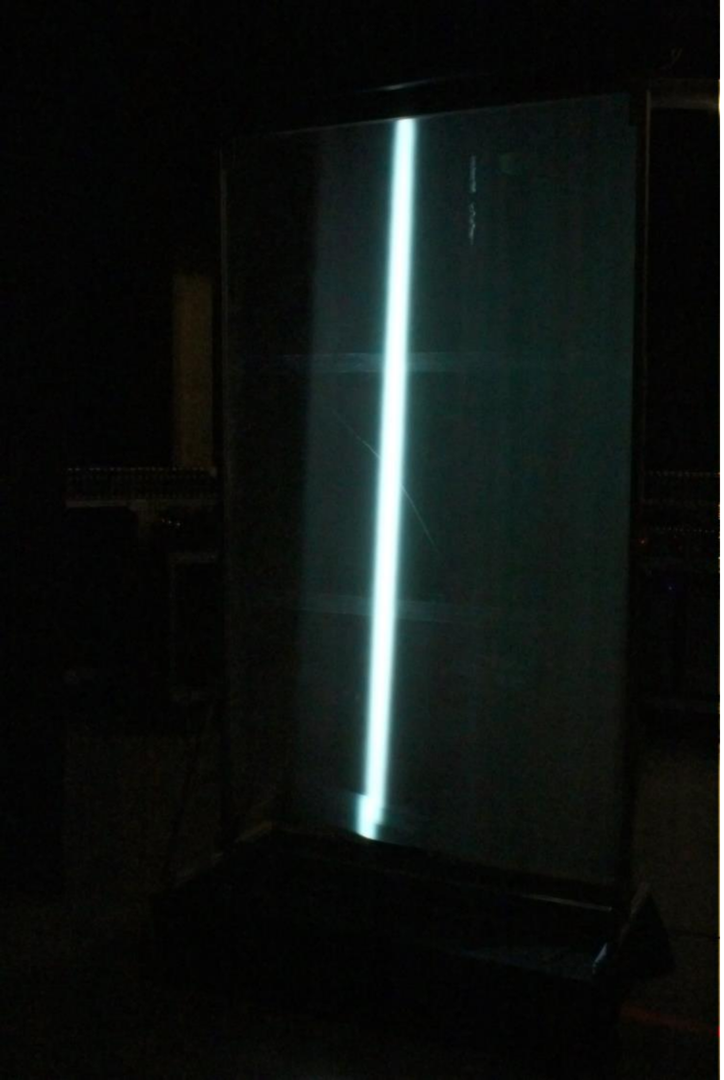

The Anisotropic Screen

Light from each projector is scattered as a vertical stripe

The Anisotropic Screen

Light from each projector is scattered as a vertical stripe

The Anisotropic Screen

Each view is composed of multiple projector stripes

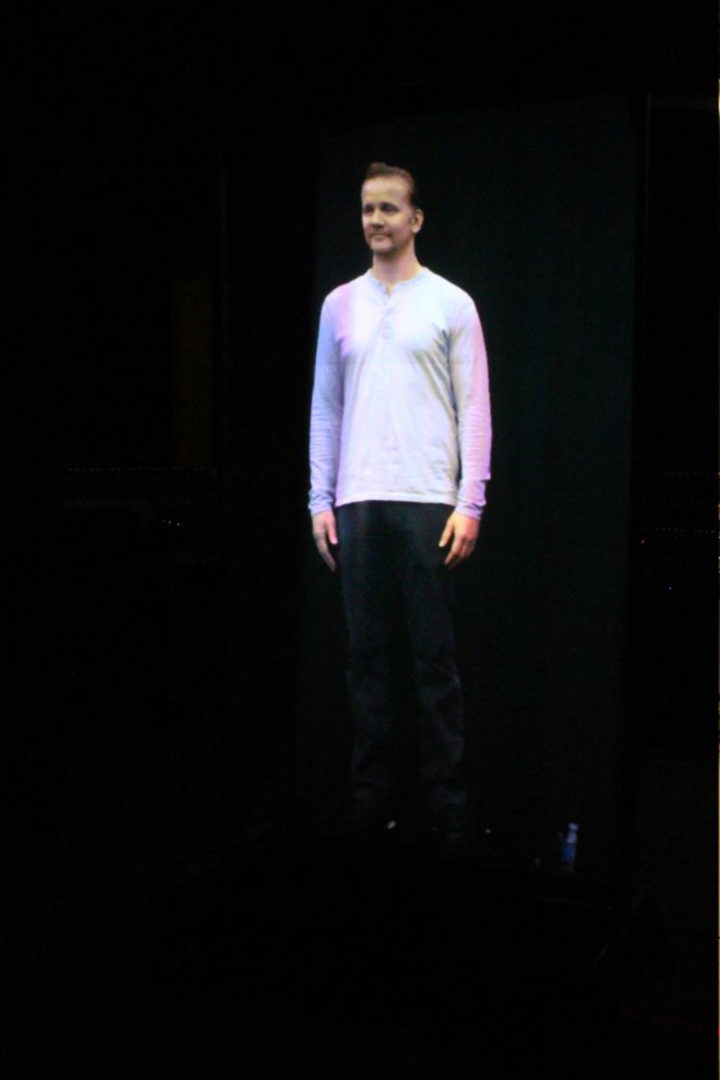

The Anisotropic Screen

Each view is composed of multiple projector stripes

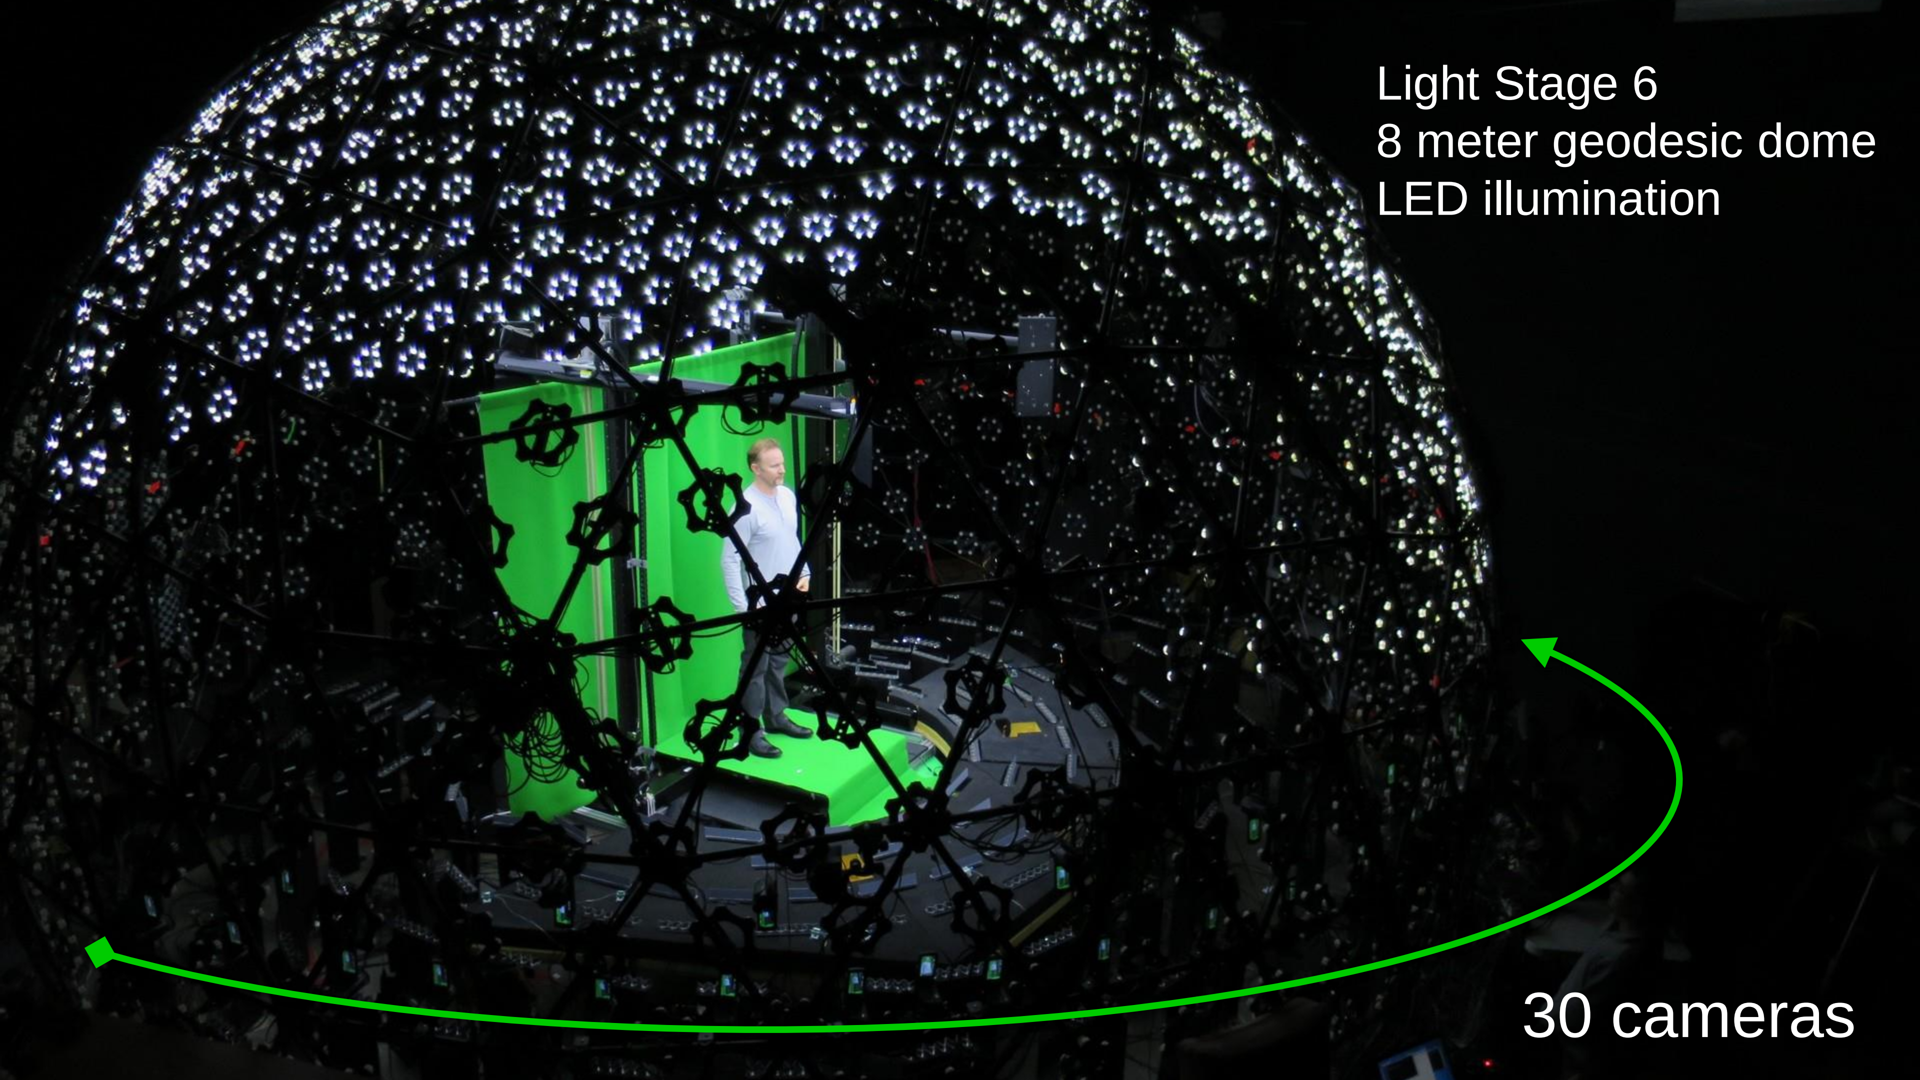
Light Stage 6
8 meter geodesic dome
LED illumination
30 cameras


NOV 24 2013

Light Field Sampling

0.625 degrees between projectors

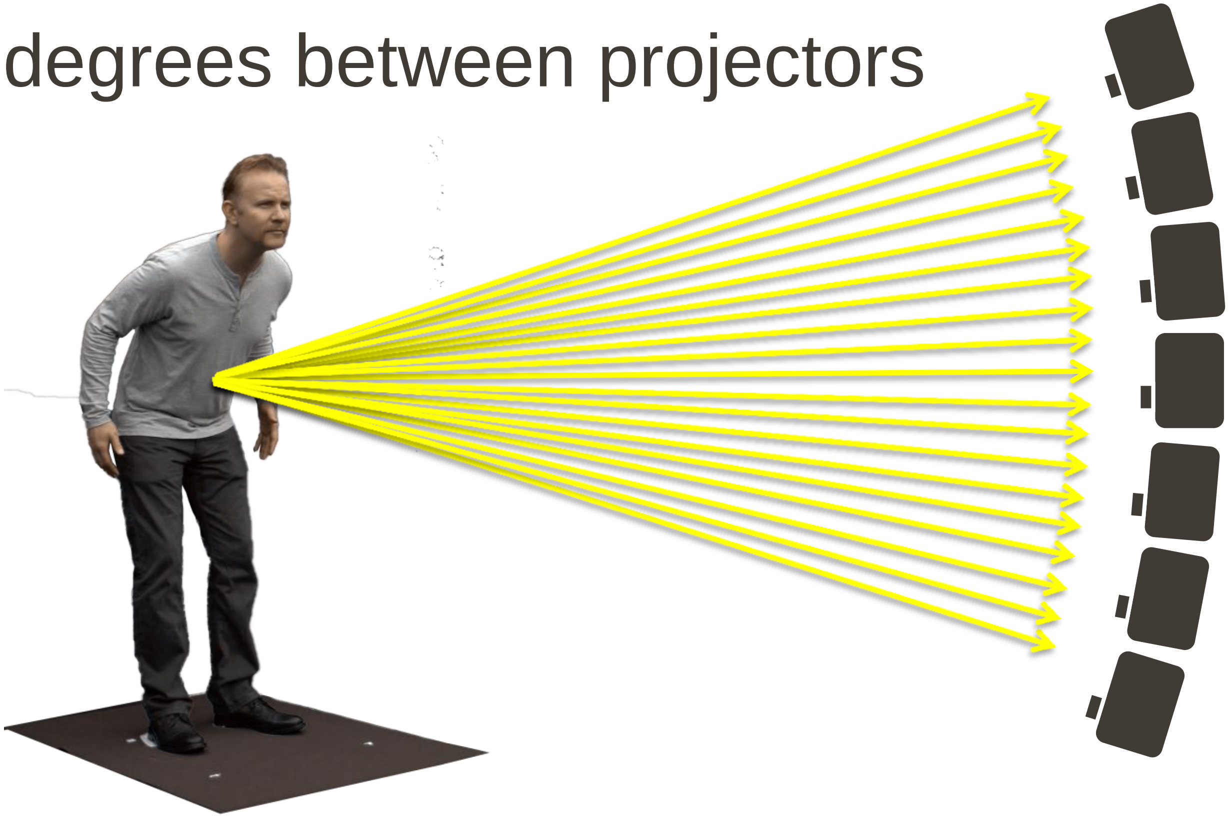

Light Field Sampling

1.75 degrees between eyes at 2 meters

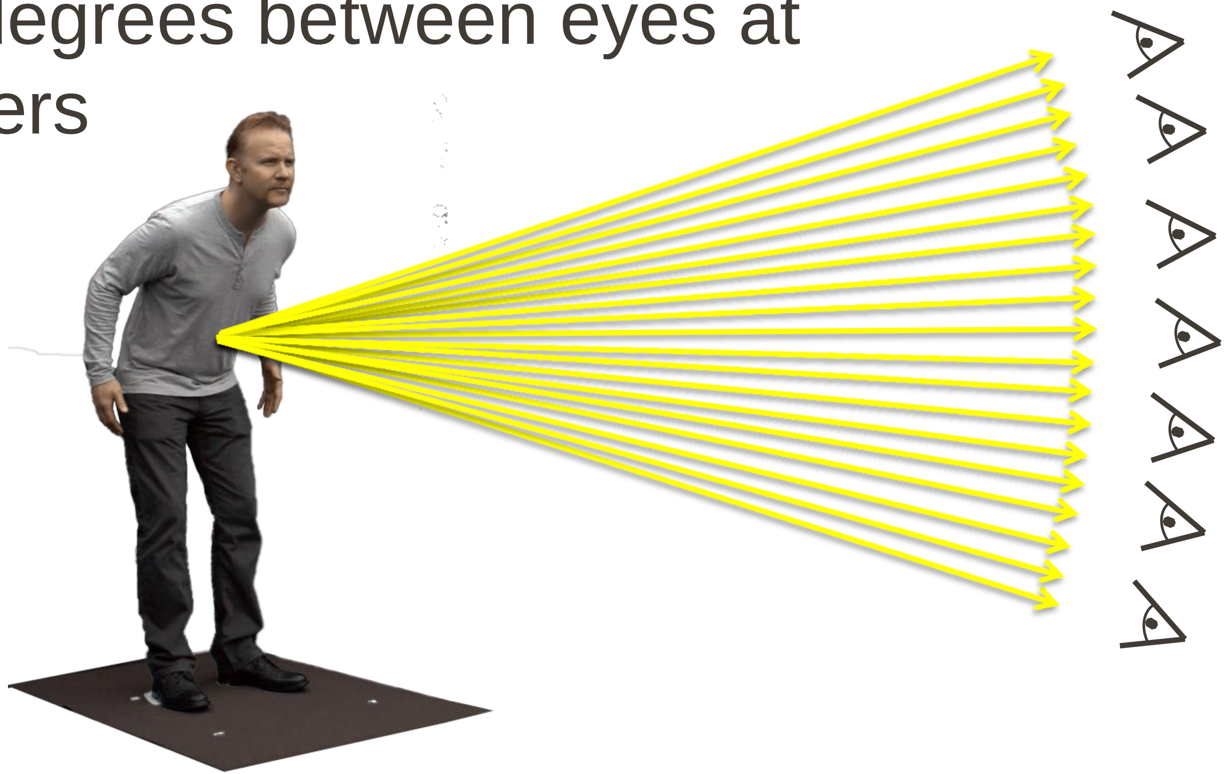

Light Field Sampling

6 degrees between cameras

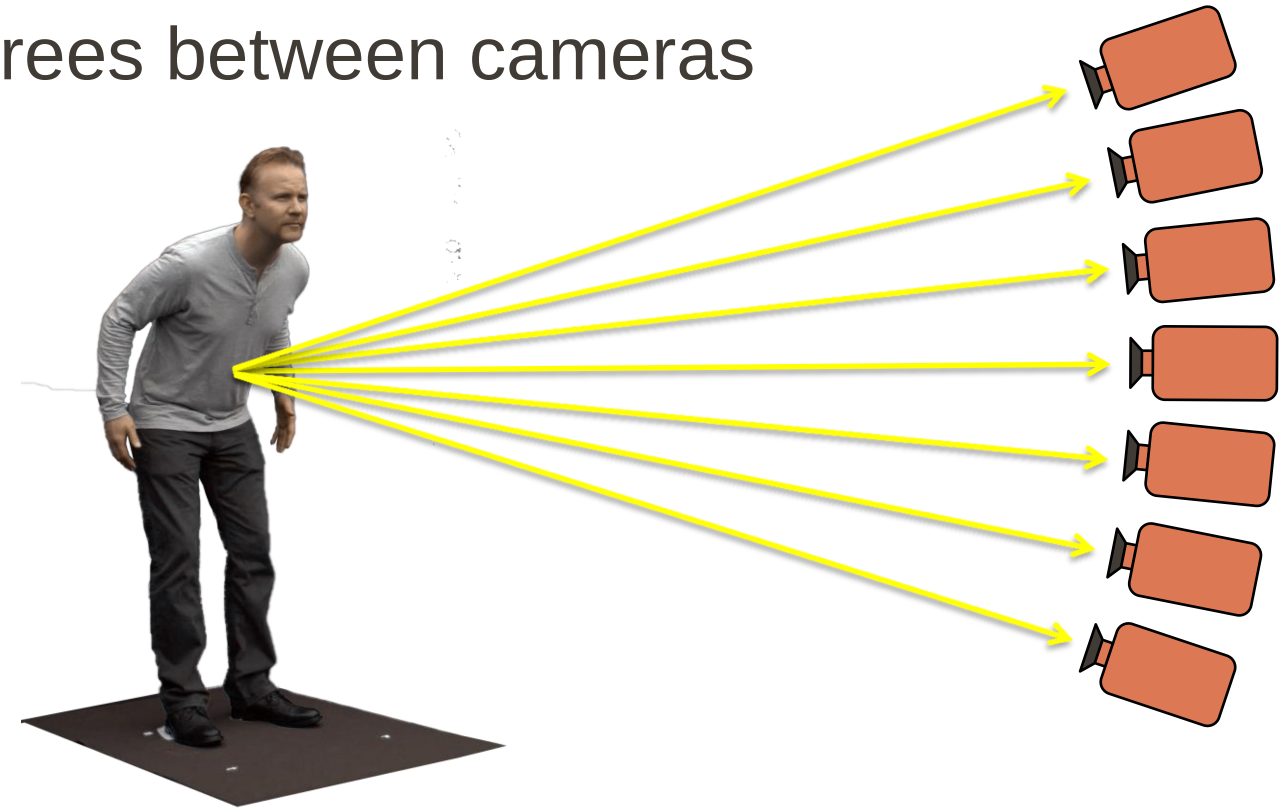

View Interpolation

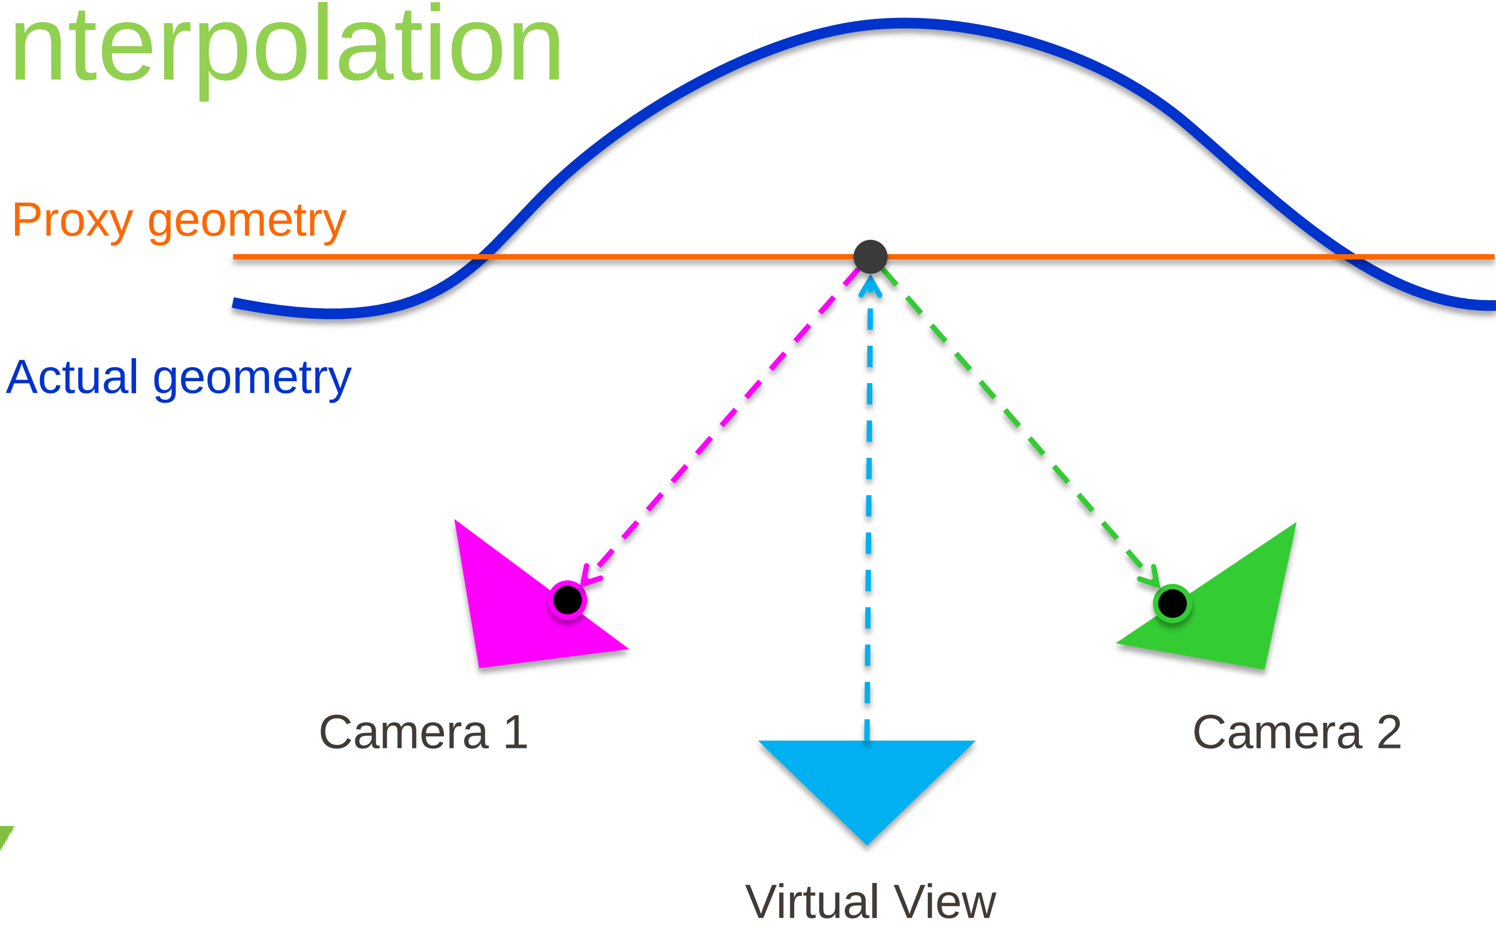

View Interpolation

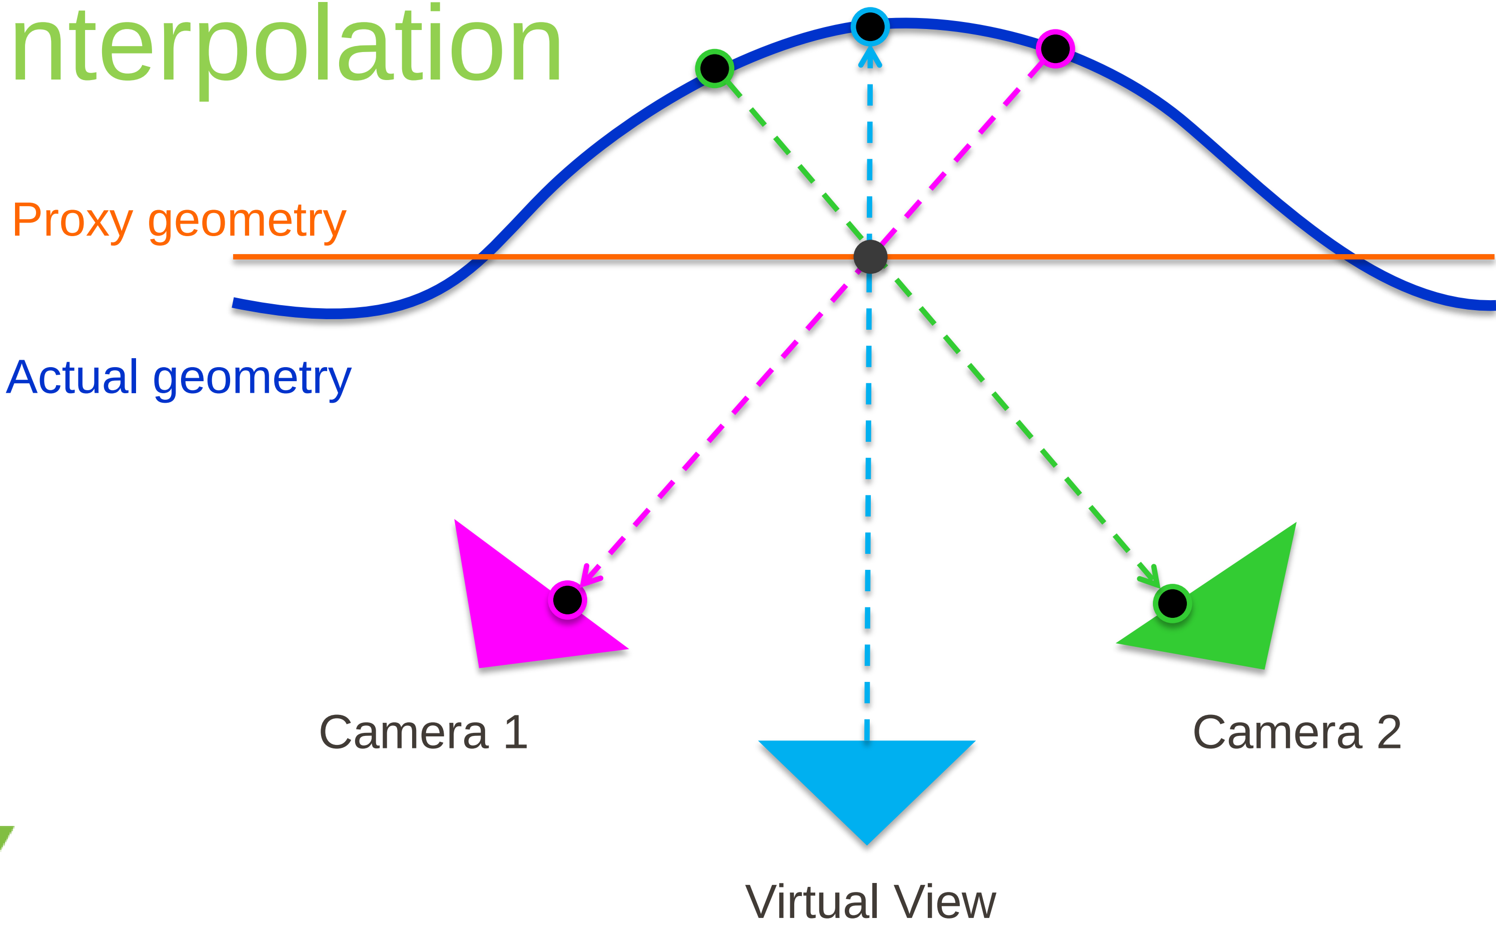


LINEAR BLENDING

Geometry Reconstruction

- Visual hulls, stereo reconstruction
- Relatively slow
- AGIsoft - 40 minutes per frame with 30 cameras

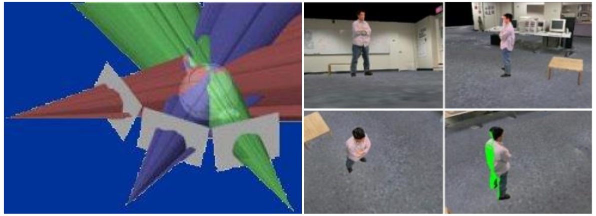

Image-Based Visual Hulls
Matusik *et al*., SIGGRAPH ‘00

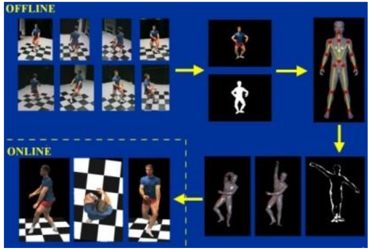

Free-viewpoint Video of Humans
Carranza *et al*., SIGGRAPH ‘03

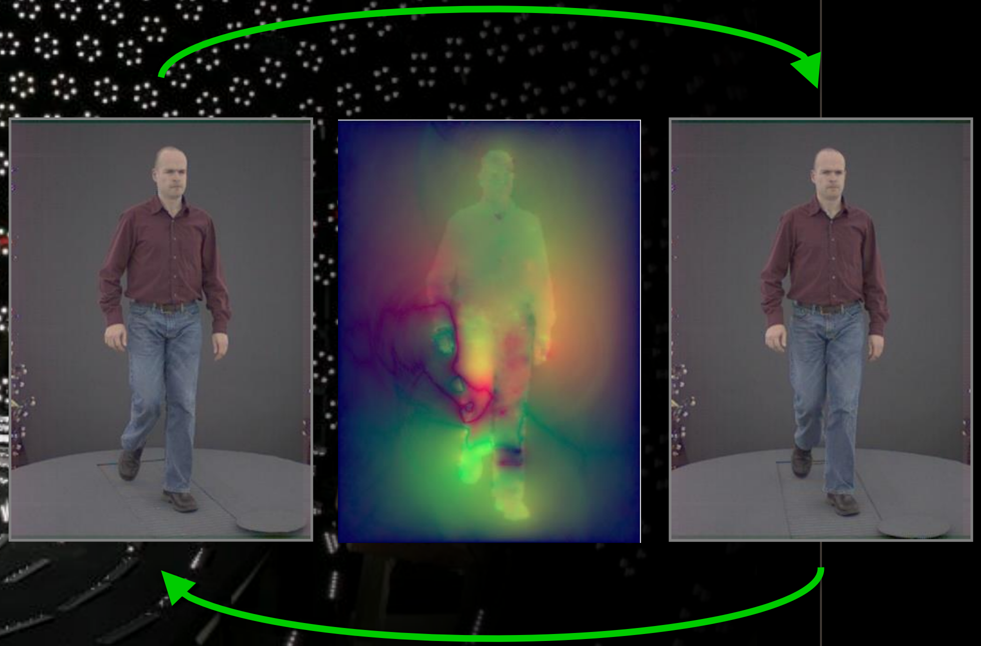

Einarsson et al. “Relighting Human Locomotion with Flowed Reflectance Fields”, EGSR 2006

M. Werlberger, T. Pock, and H. Bischof: *Motion Estimation with Non-Local Total Variation Regularization*, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), San Francisco, CA, USA, June 2010.

View Interpolation

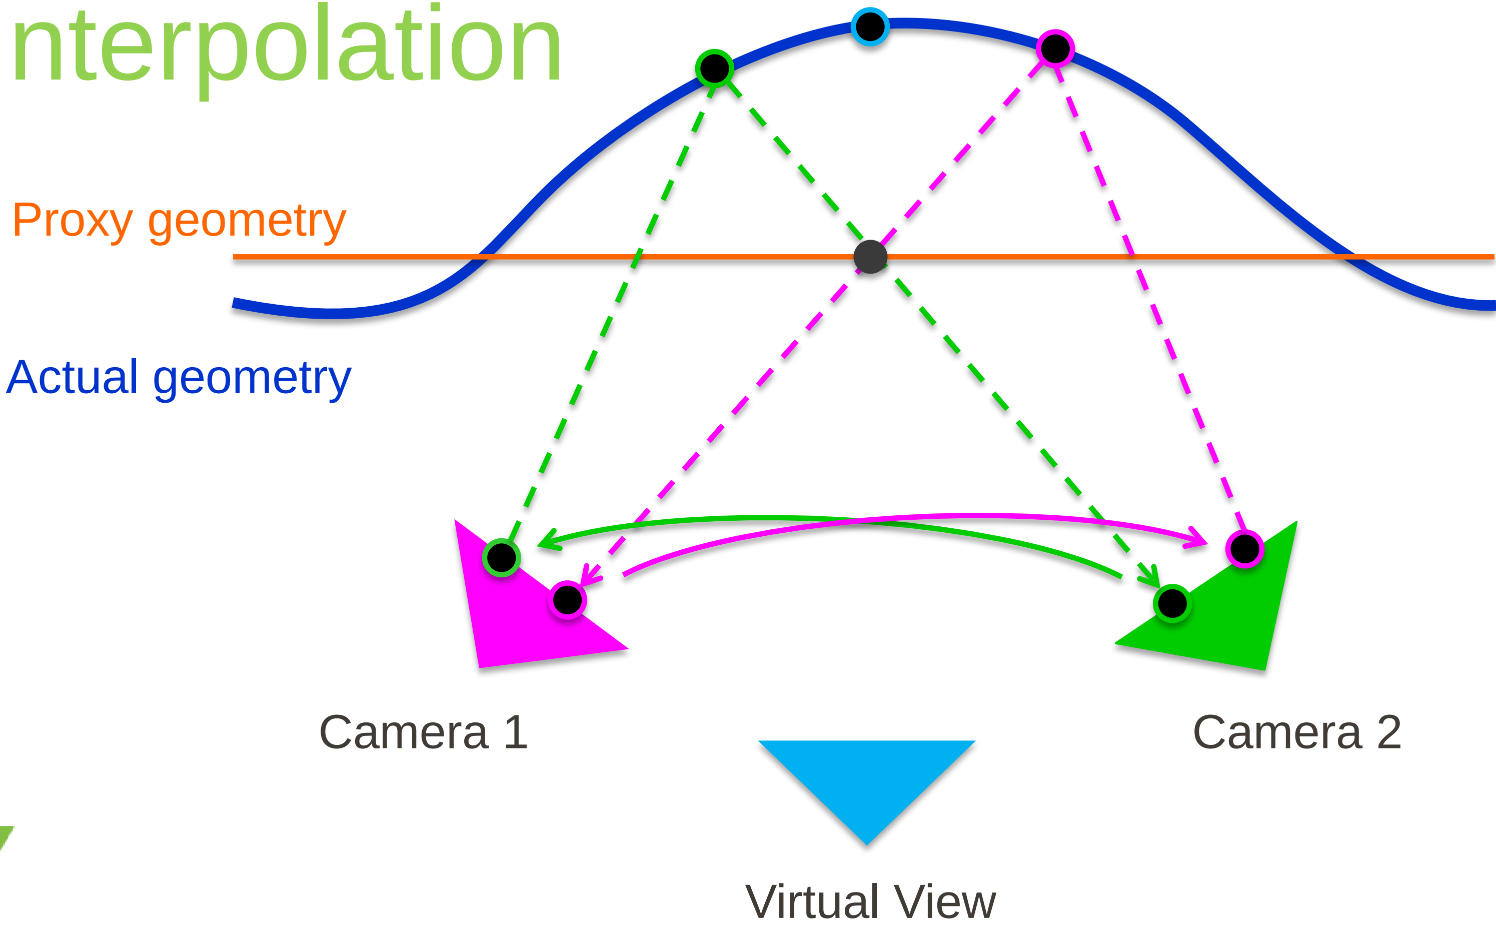

View Interpolation

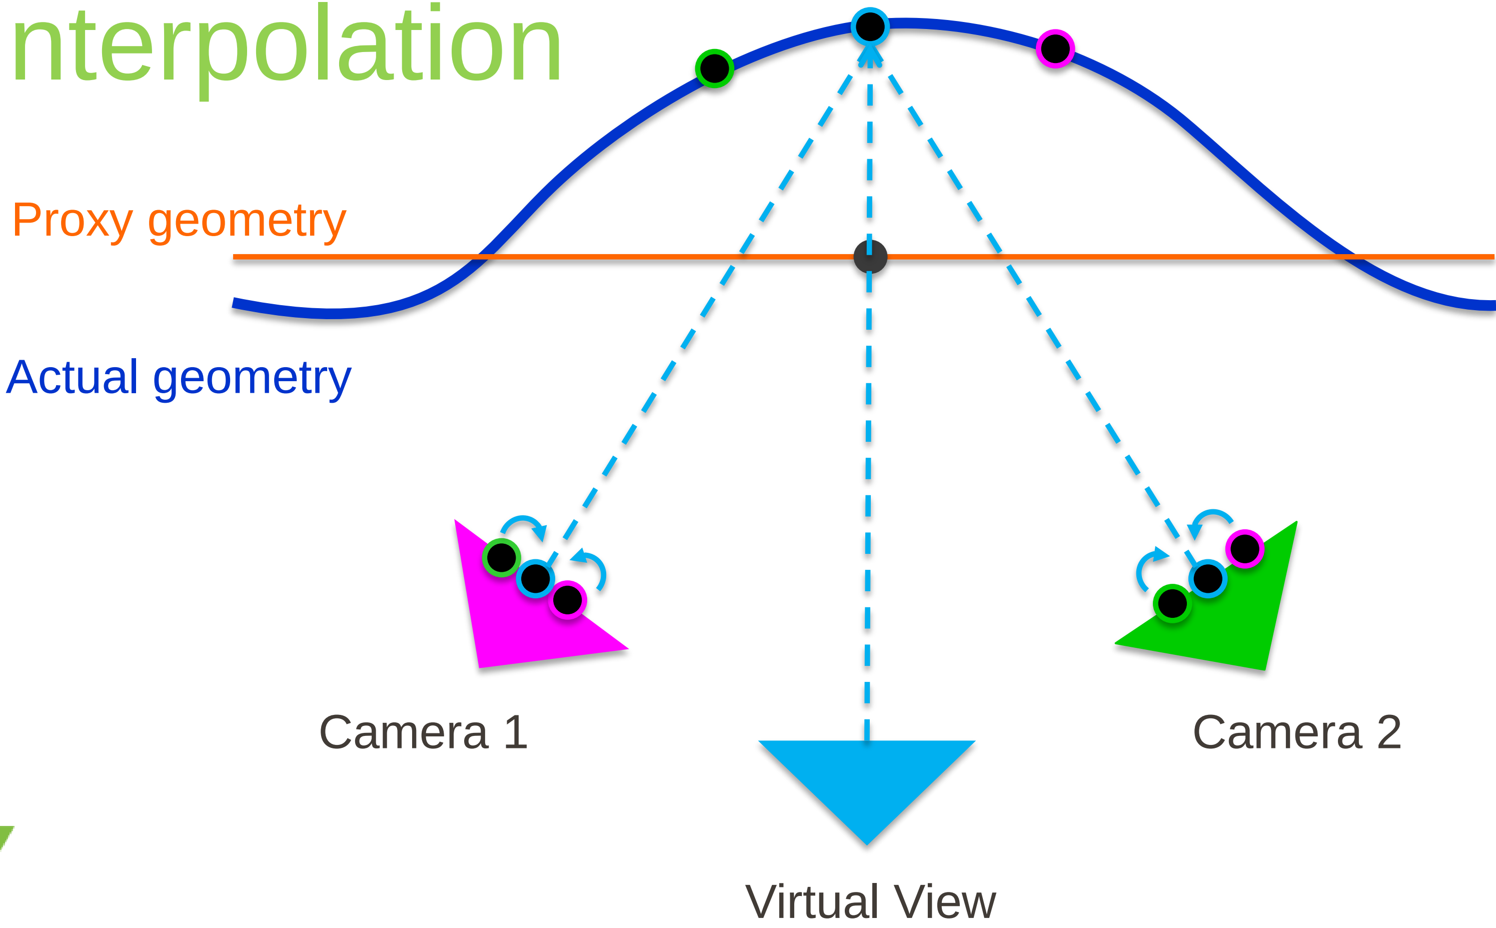

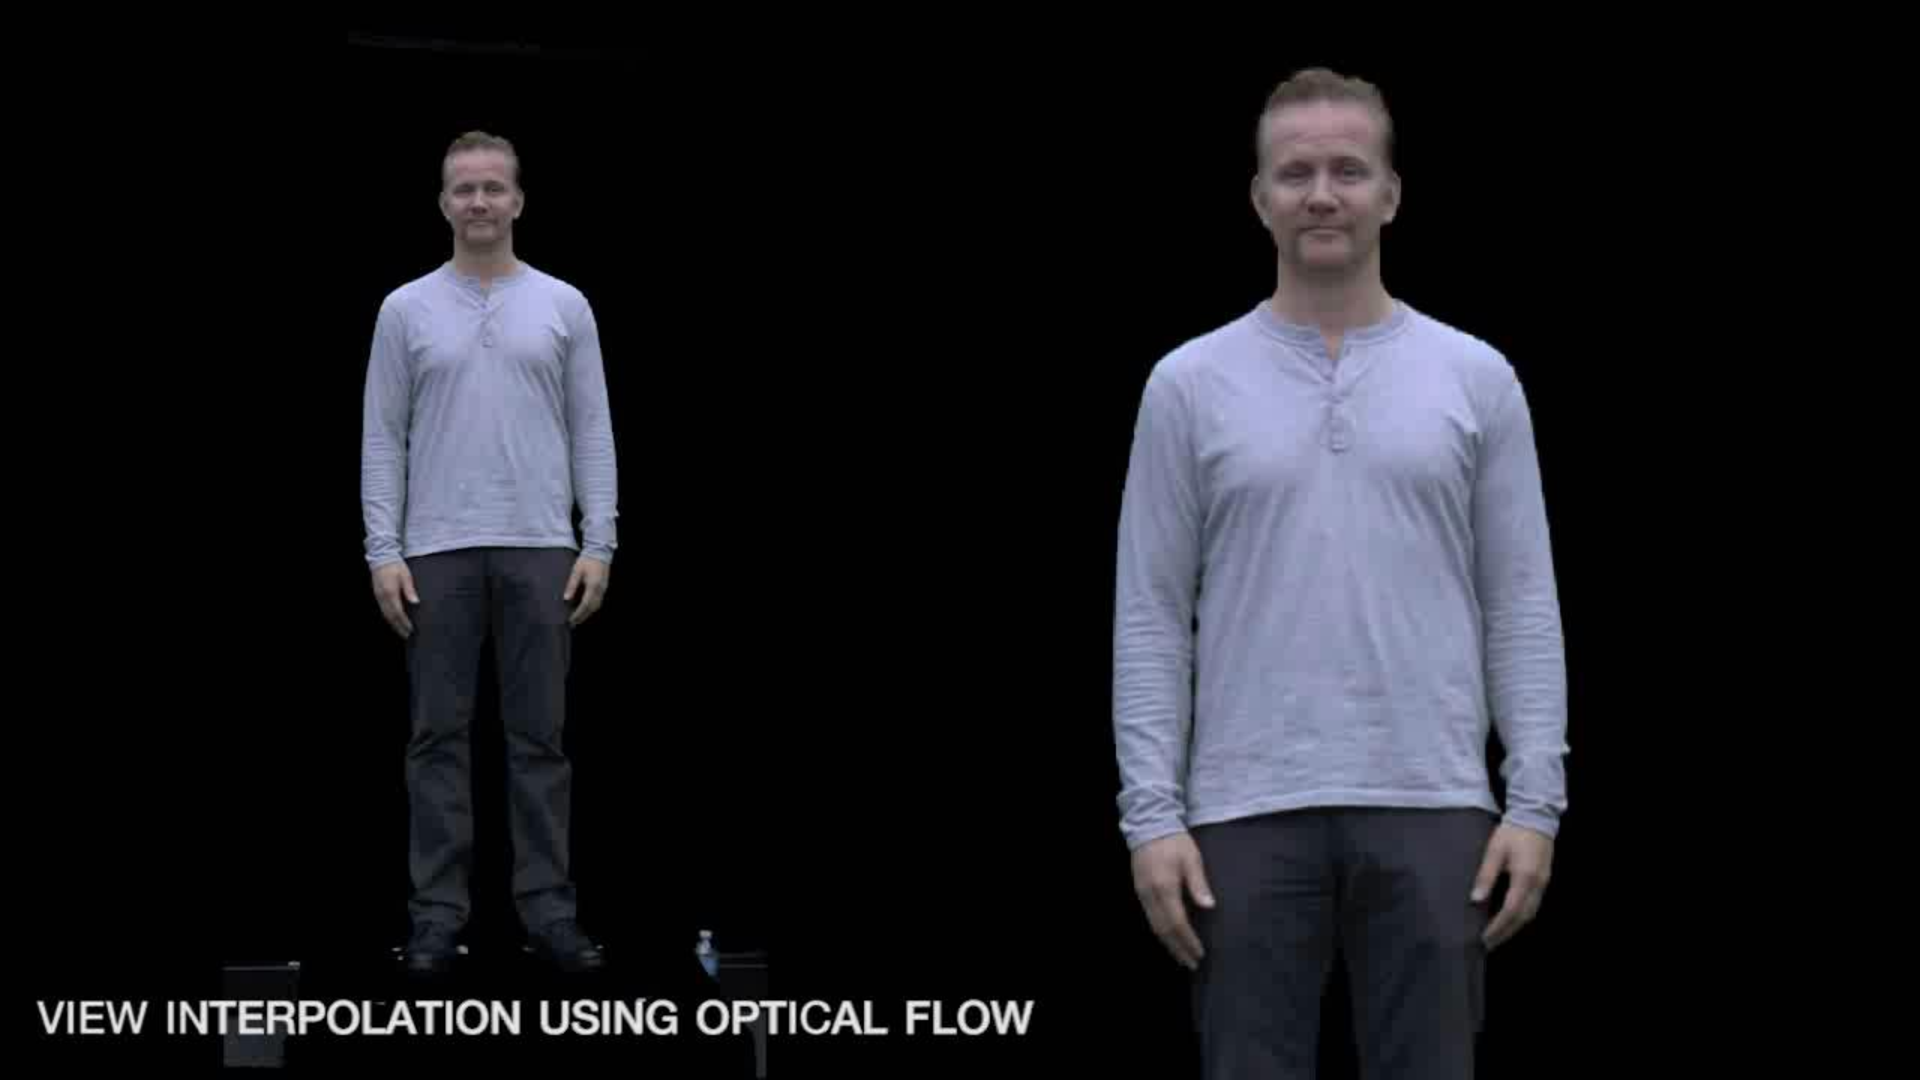
VIEW INTERPOLATION USING OPTICAL FLOW

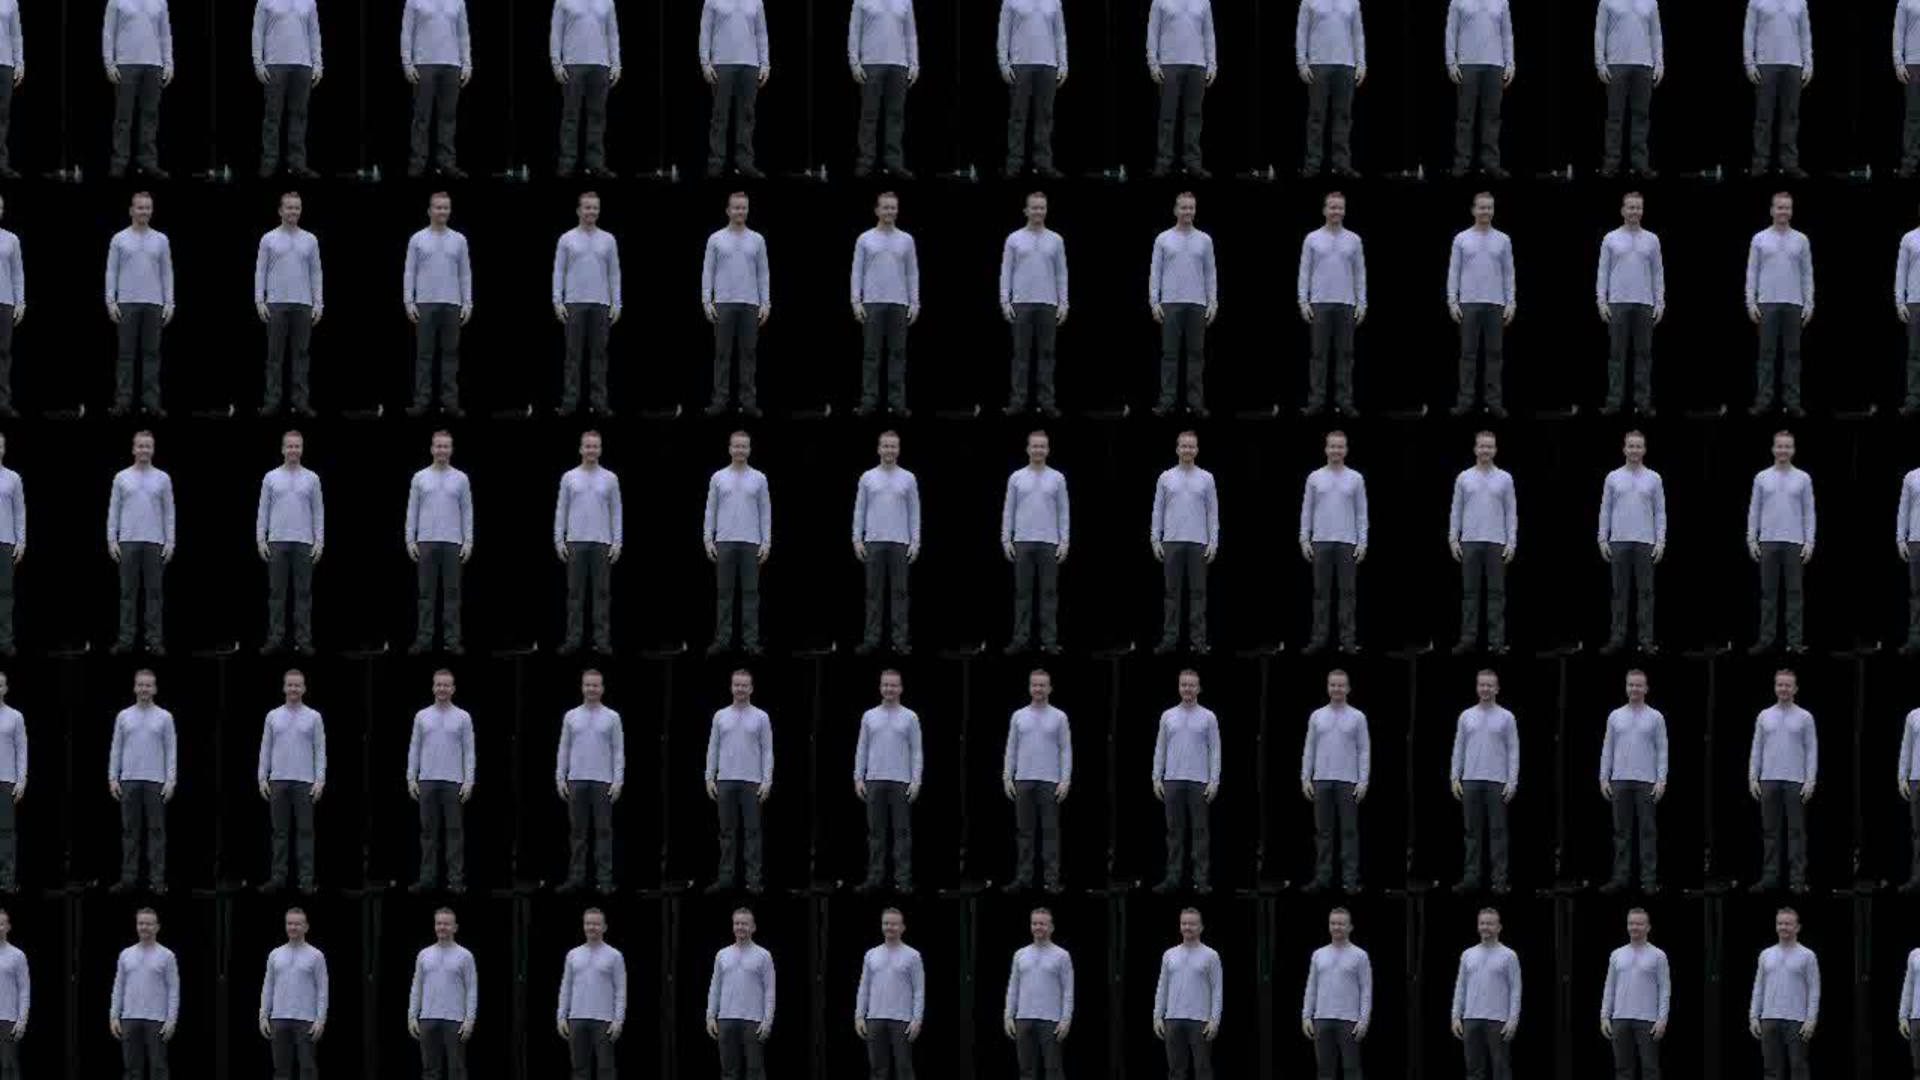

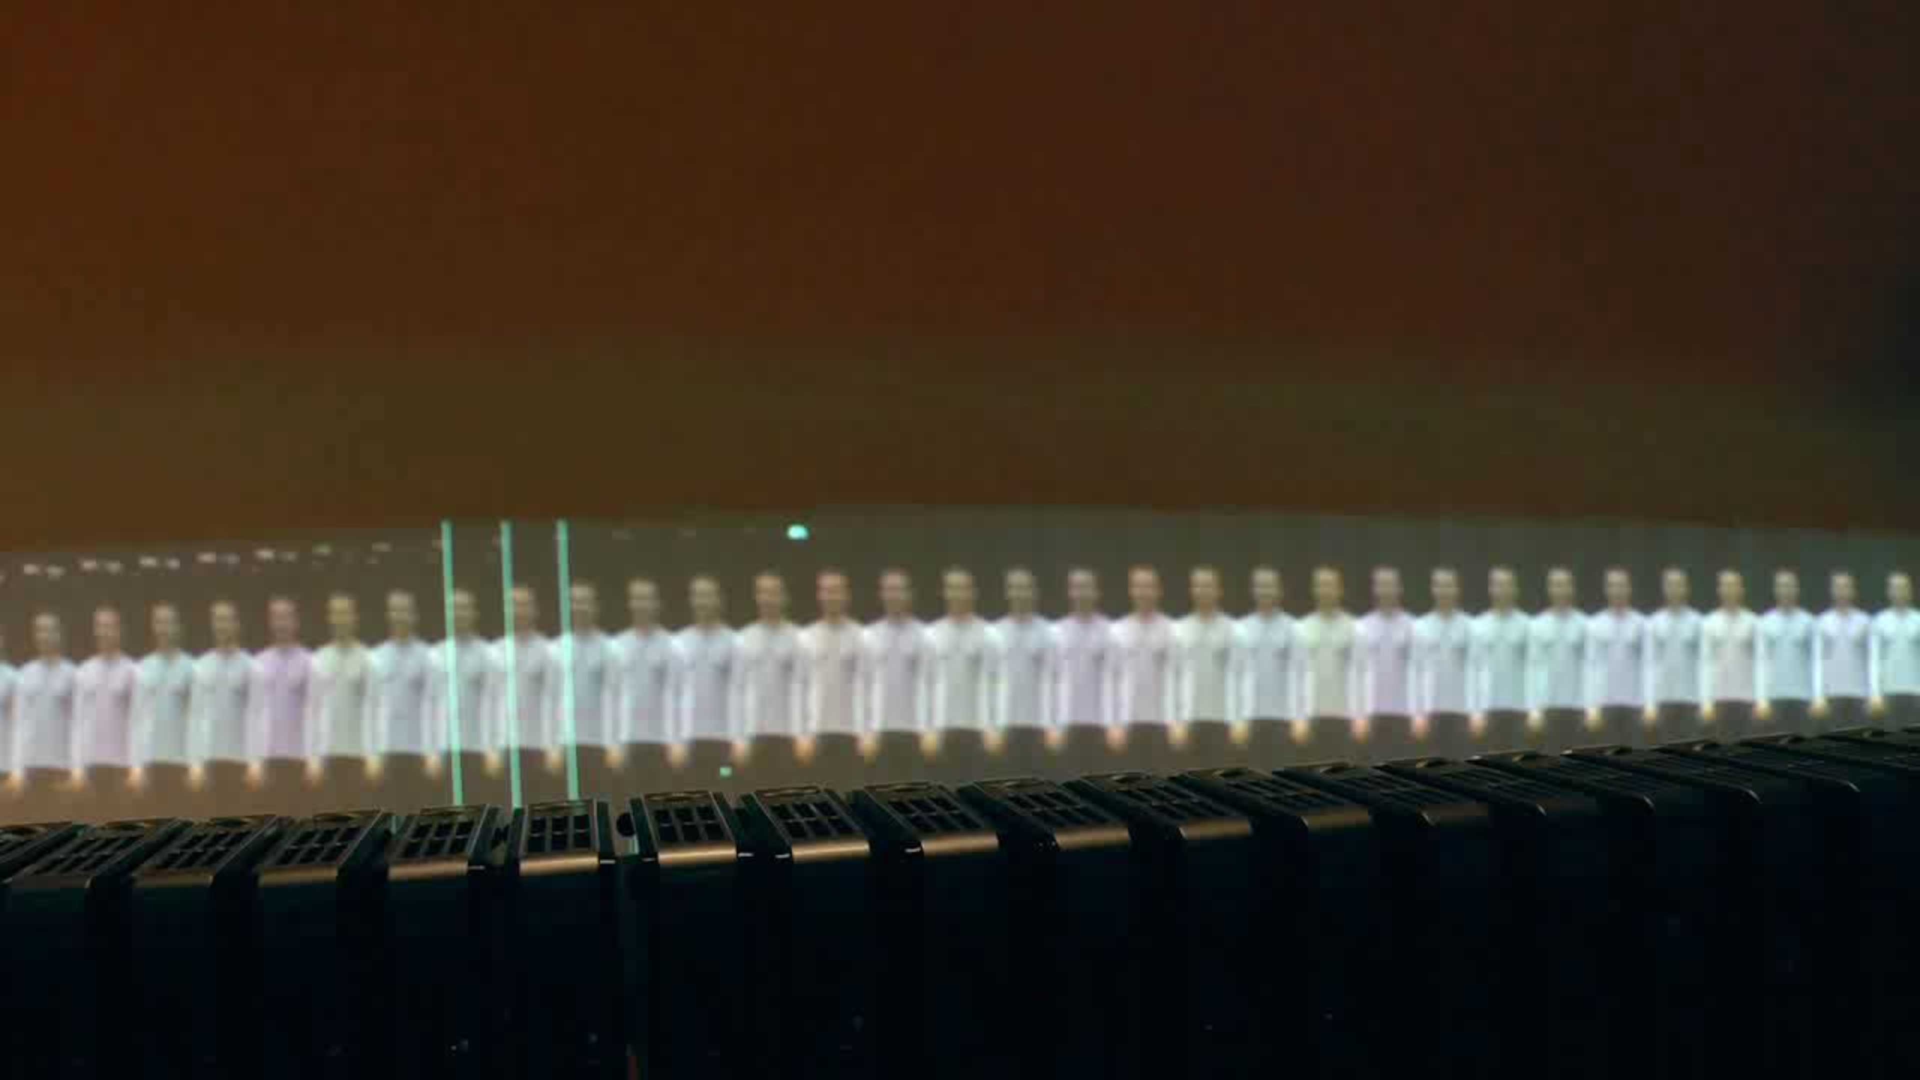

VIEW INTERPOLATION ON DISPLAY

Video Decoding

- 11 source videos, 20 optical flow videos per GPU

- CPU decoding FFMPEG (multi-core)
- GPU MPEG video decoding (NVCUVID)

Distributed rendering

Windows 7 Default:

commands sent to most single GPU and blitted across

Current solution: New instance of application per GPU

Next step: OS/Vendor specfic extensions to assign resources to GPUs (ie WGL_NV_gpu_affinity)

Shalini Venkataraman, "*Programming Multi-GPUs for Scalable Rendering*" GTC 2012

Ongoing Work

- Incorporate natural language processing / artificial intelligence
- Extend up to 30+ hours of interview

Arstein et al. "Time-Offset Interaction with a Holocaust Survivor", *Proceedings of International Conference On Intelligent User Interfaces (IUI)*, 2014

GPU TECHNOLOGY CONFERENCE

Conclusions

- Simple techniques for rendering geometry and light fields for automultiscopic displays
- Limited by GPU bandwidth
- Need new tools to exploit redundancy, and distribute resources across views

Questions

Thanks to CNN, Morgan Spurlock, Inside Man Productions, Shoah Foundation, Pinchas Gutter, Julia Campbell, Bill Swartout, Randall Hill, Randolph Hall, U.S. Air Force DURIP, and U.S. Army RDECOM

http://gl.ict.usc.edu/

GPU TECHNOLOGY CONFERENCE