Light Fields and Camera Arrays

Image Manipulation and Computational Photography
CS294-69 Fall 2011
Maneesh Agrawala

[Some slides from James Hays , Derek Hoiem, Alexei Efros and Fredo Durand]

A1 Texture Synthesis

Due Mon Sep 26

Implement application of texture synthesis

Image analogies, Hole-filling, Patchmatch Structured hybrids ….

Adequate to implement, best solutions go beyond:

Every technique has some limitations (well written papers usually describe some of them). Develop techniques to address one or more limitations?

Sometimes different papers present different techniques for addressing the same problem Implement competing techniques and compare their strengths and weaknesses.

It may be possible to combine ideas from multiple papers to produce a new hybrid technique that addresses a new problem. Develop a new way to combine the texture synthesis techniques your have read about to solve a new problem.

1 person = 1 paper,
2 people = 1 paper + issue from list above or 2 papers,
3 people = 2 papers + issue from list above

A2 Lights Fields/HDR/Tone-Mapping Due Mon Oct 10

Implement Light Fields or HDR/Tone-Mapping

Adequate to implement, best solutions go beyond:

Every technique has some limitations (well written papers usually describe some of them). Develop techniques to address one or more limitations?

Sometimes different papers present different techniques for addressing the same problem Implement competing techniques and compare their strengths and weaknesses.

It may be possible to combine ideas from multiple papers to produce a new hybrid technique that addresses a new problem. Develop a new way to combine the texture synthesis techniques your have read about to solve a new problem.

1 person = 1 paper,

2 people = 1 paper + issue from list above or 2 papers,

3 people = 2 papers + issue from list above

What is Light?

Electromagnetic radiation (EMR) moving along rays in space

- $R(\lambda)$ is EMR, measured in units of power (watts)
  - λ is wavelength

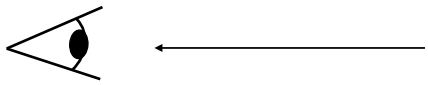

Useful things:

- Light travels in straight lines
- In vacuum, radiance emitted = radiance arriving
  - i.e. there is no transmission loss

What Do We See?

3D world

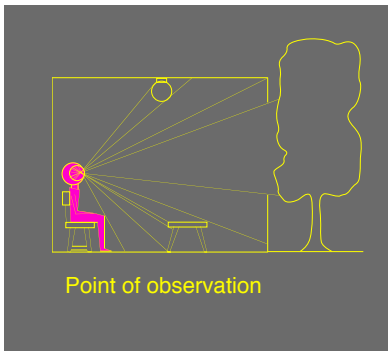

2D image

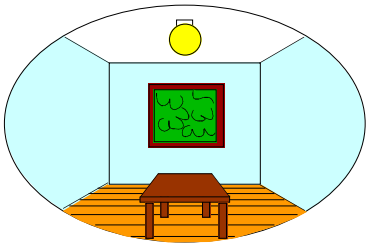

What Do We See?

3D world

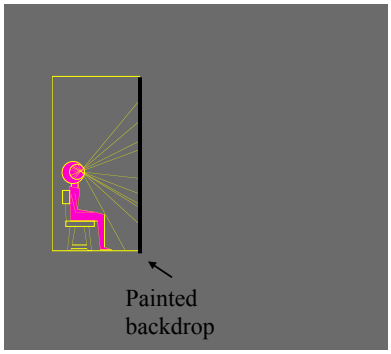

2D image

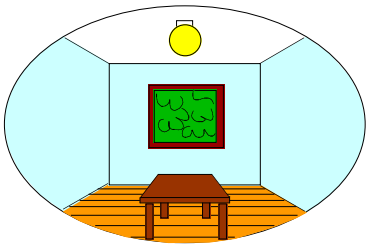

On Simulating the Visual Experience

Just feed the eyes the right data

- No one will know the difference!

Philosophy:

- Ancient question: "Does the world really exist?"

Science fiction:

- Many, many, many books on the subject, e.g. *slowglass* from "Light of Other Days"
- Latest take: *The Matrix*

Physics:

- *Slowglass* might be possible?

Computer Science:

- Virtual Reality

To simulate we need to know:

What does a person see?

The Plenoptic Function

Figure by Leonard McMillan

Q: What is the set of all things that we can ever see?

A: The Plenoptic Function (Adelson & Bergen)

Grayscale snapshot

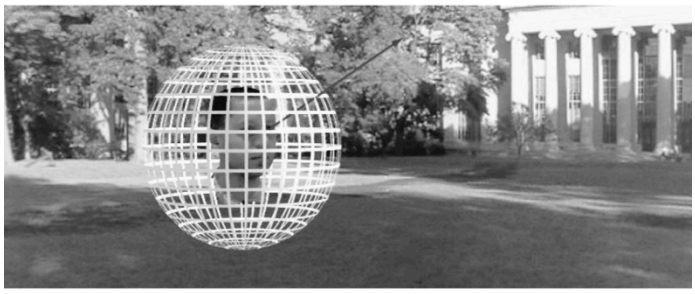

$$P(\theta,\phi)$$

intensity of light

- Seen from a single view point
- At a single time
- Averaged over the wavelengths of the visible spectrum

(can also do *P(x,y),* but spherical coordinate are nicer)

Color snapshot

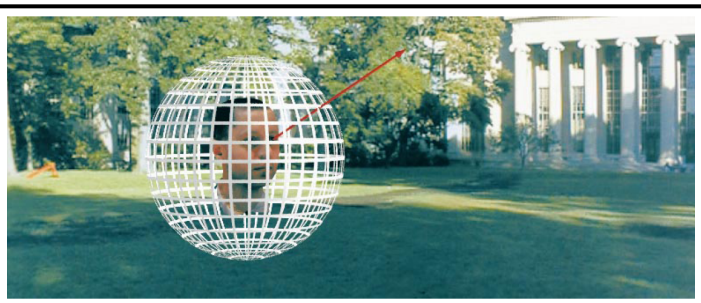

$$P(\theta,\phi,\lambda)$$

intensity of light

- Seen from a single view point
- At a single time
- As a function of wavelength

2D: Image

What is an image?

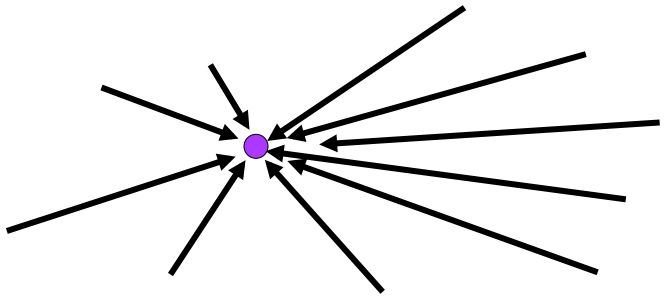

All rays through a point

Panorama?

Slide by Rick Szeliski and Michael Cohen

Image

Image plane

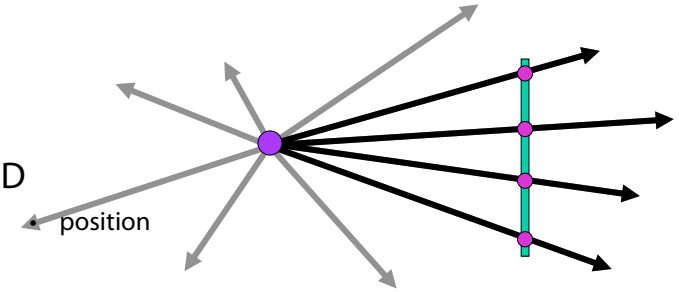

2D

- position

Spherical Panorama

See also: 2003 New Years Eve
http://www.panoramas.dk/fullscreen3/f1.html

All light rays through a point form a ponorama

Totally captured in a 2D array -- $P(\theta,\phi,\lambda)$

Where is the geometry???

Panorama Viewer

http://www.panoramas.dk/US/golden-gate.html

A movie

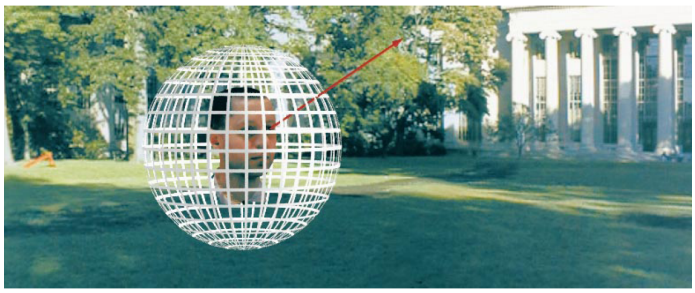

$$P(\theta,\phi,\lambda,t)$$

is intensity of light

- Seen from a single view point
- Over time
- As a function of wavelength

Holographic movie

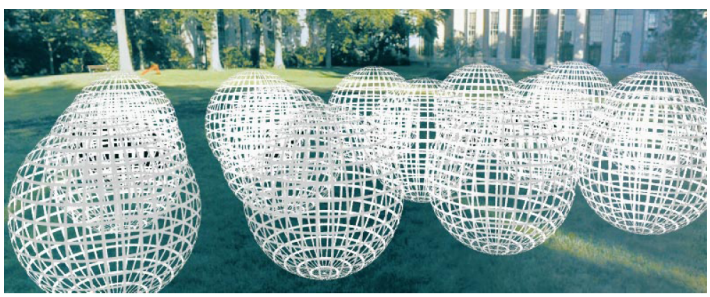

$$P(\theta,\phi,\lambda,t,V_X,V_Y,V_Z)$$

intensity of light

- Seen from ANY viewpoint
- Over time
- As a function of wavelength

The Plenoptic Function

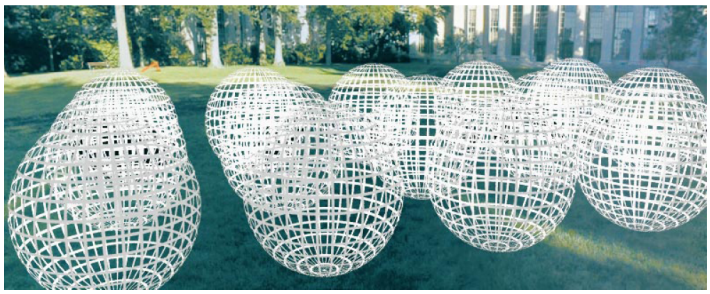

$$P(\theta,\phi,\lambda,t,V_X,V_Y,V_Z)$$

Can reconstruct every possible view, at every moment, from every position, at every wavelength

Contains every photograph, every movie, everything that anyone has ever seen! it completely captures our visual reality! Not bad for a function…

Sampling Plenoptic Function (top view)

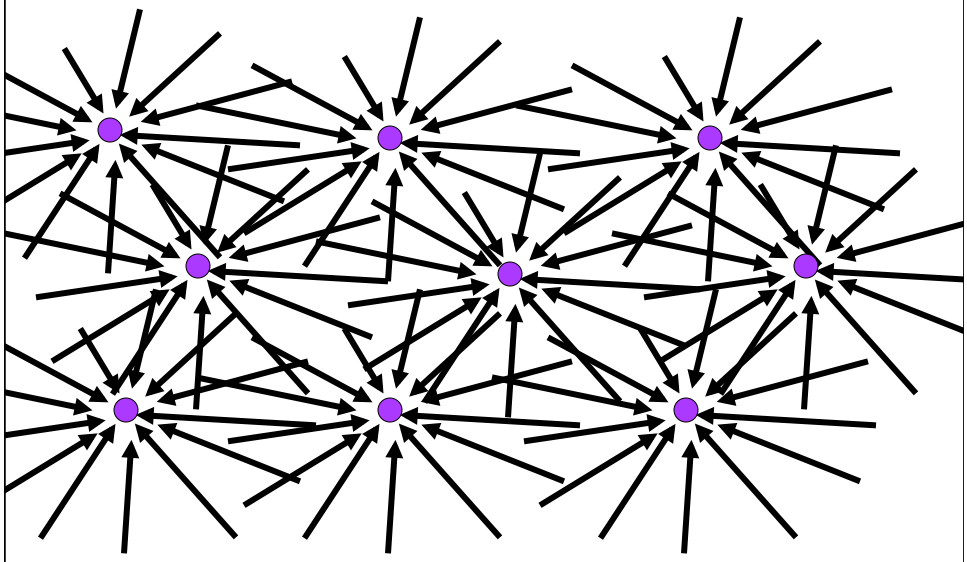

Just lookup -- Quicktime VR

Ray

Let's not worry about time and color:

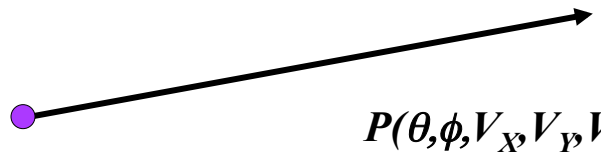

$P(\theta,\phi,V_X,V_Y,V_Z)$

5D

- 3D position
- 2D direction

Slide by Rick Szeliski and Michael Cohen

How Can We Use This?

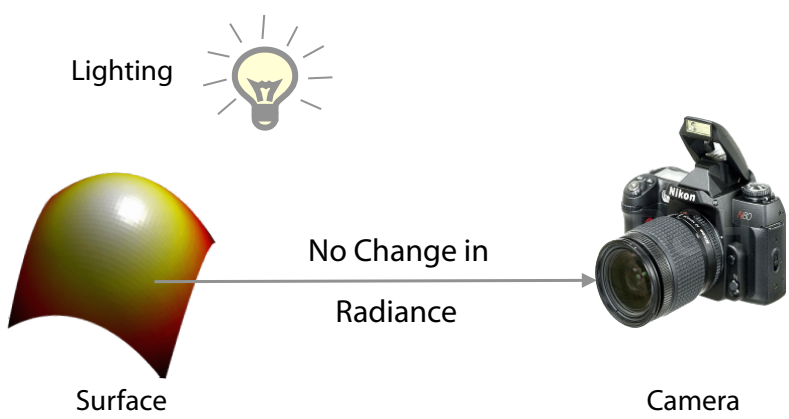

Ray Reuse

Infinite line

- Assume light is constant (vacuum)

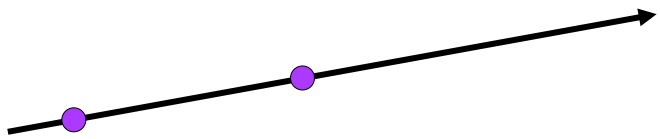

4D

- 2D direction
- 2D position
- non-dispersive medium

Slide by Rick Szeliski and Michael Cohen

Only Need Plenoptic Surface

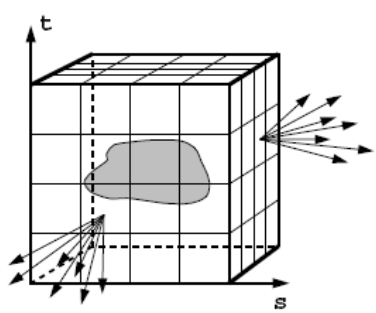

Figure 1: The surface of a cube holds all the radiance information due to the enclosed object.

Synthesizing Novel Views

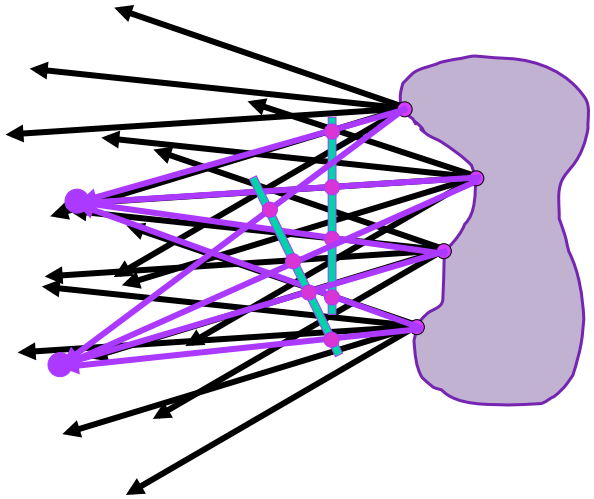

Slide by Rick Szeliski and Michael Cohen

Lumigraph / Lightfield

Outside convex space

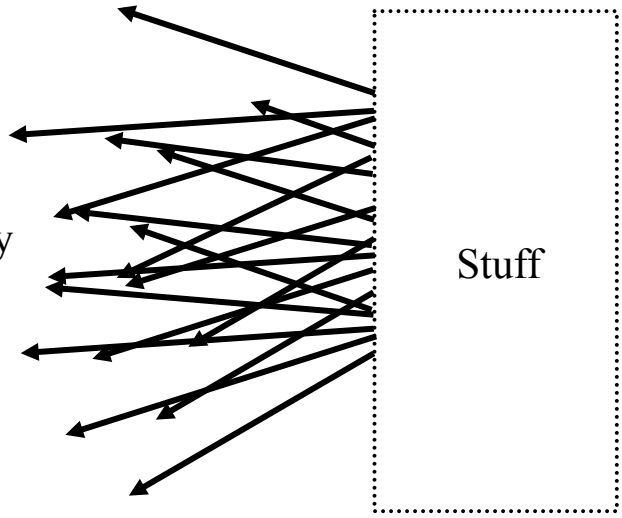

4D

Slide by Rick Szeliski and Michael Cohen

Lumigraph - Organization

2D position

2D direction

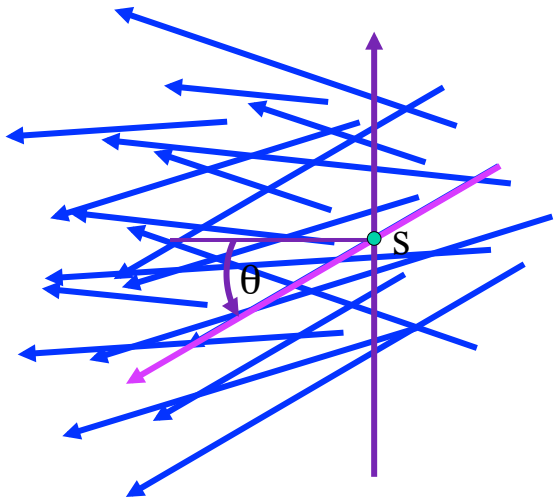

Slide by Rick Szeliski and Michael Cohen

Lumigraph - Organization

2D position

2D position

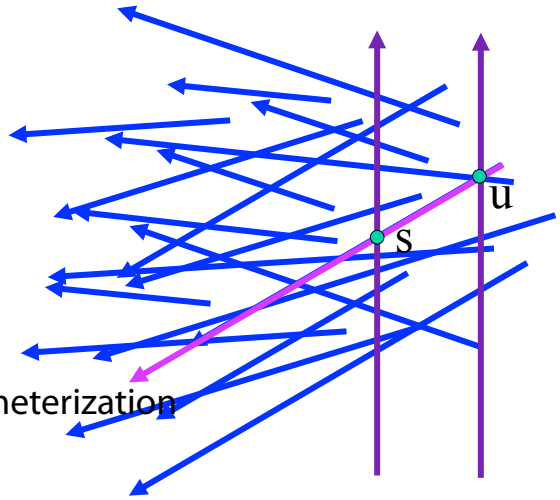

2 plane parameterization

Slide by Rick Szeliski and Michael Cohen

Lumigraph - Organization

2D position

2D position

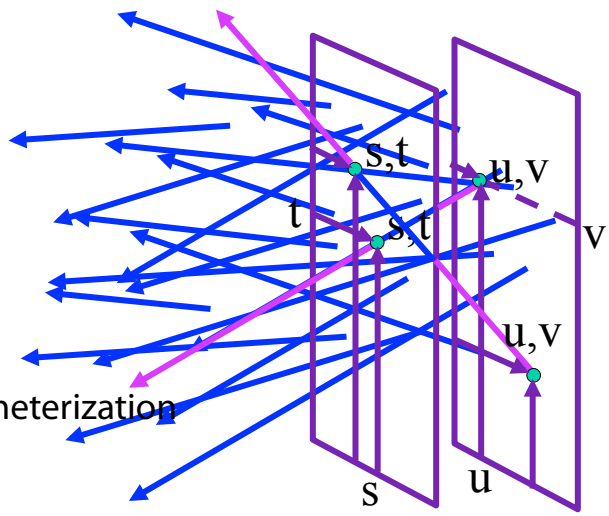

2 plane parameterization

Slide by Rick Szeliski and Michael Cohen

Lumigraph - Organization

Hold s,t constant

Let u,v vary

An image

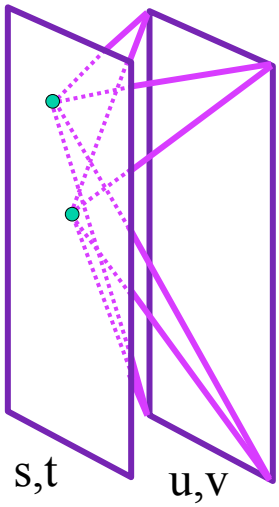

Slide by Rick Szeliski and Michael Cohen

Lumigraph / Lightfield

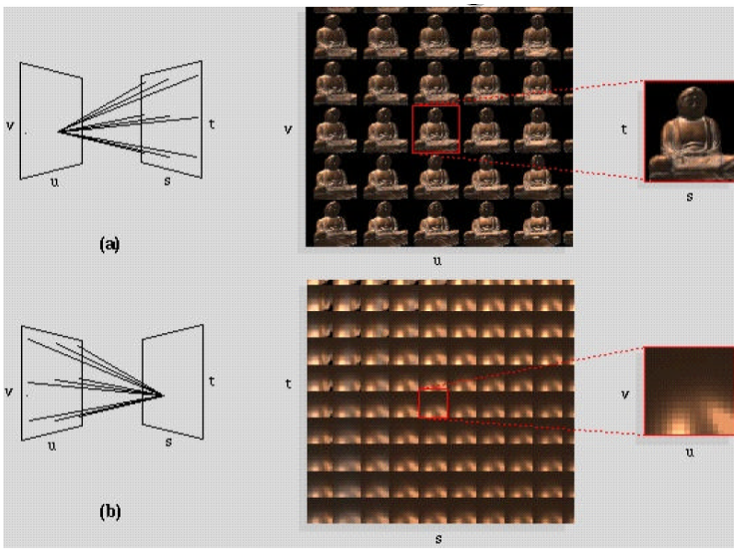

Lumigraph - Capture

Idea 1

- Move camera carefully over s,t plane
- Gantry (see Lightfield paper)

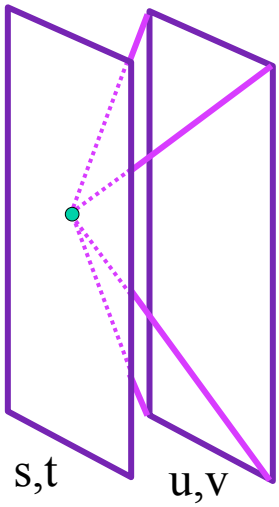

Slide by Rick Szeliski and Michael Cohen

Lumigraph - Capture

Idea 2

- Move camera anywhere
- Rebinning (see Lumigraph paper)

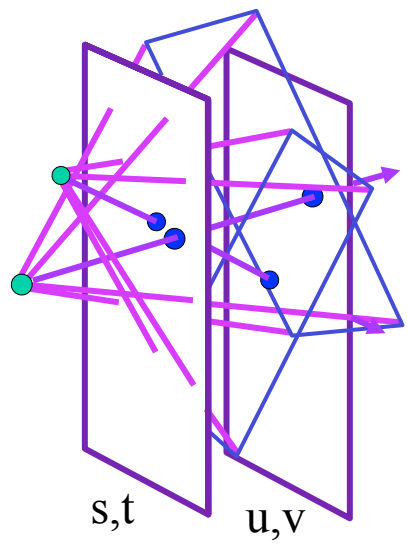

Slide by Rick Szeliski and Michael Cohen

Lumigraph - Rendering

For each output pixel

- determine s,t,u,v

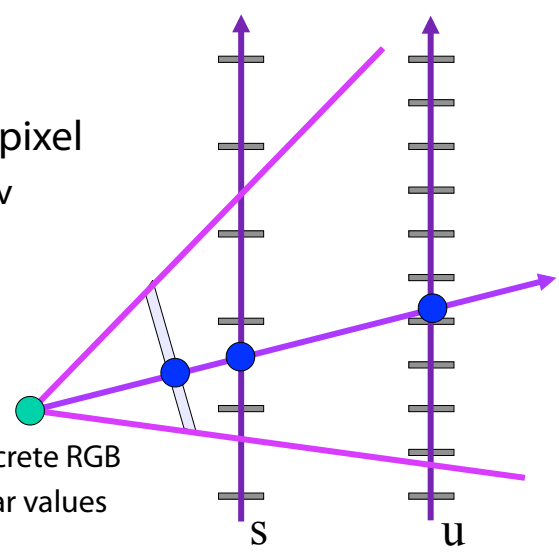

- either
  - use closest discrete RGB
  - interpolate near values

Slide by Rick Szeliski and Michael Cohen

Lumigraph - Rendering

Nearest

- closest s
- closest u
- draw it

Blend 16 nearest

- quadrilinear interpolation

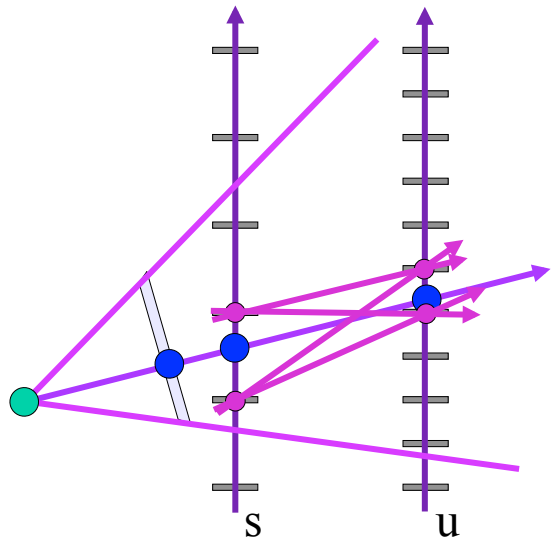

Slide by Rick Szeliski and Michael Cohen

Light field photography using a handheld plenoptic camera

Ren Ng, Marc Levoy, Mathieu Brédif,
Gene Duval, Mark Horowitz and Pat Hanrahan

Conventional versus light field camera

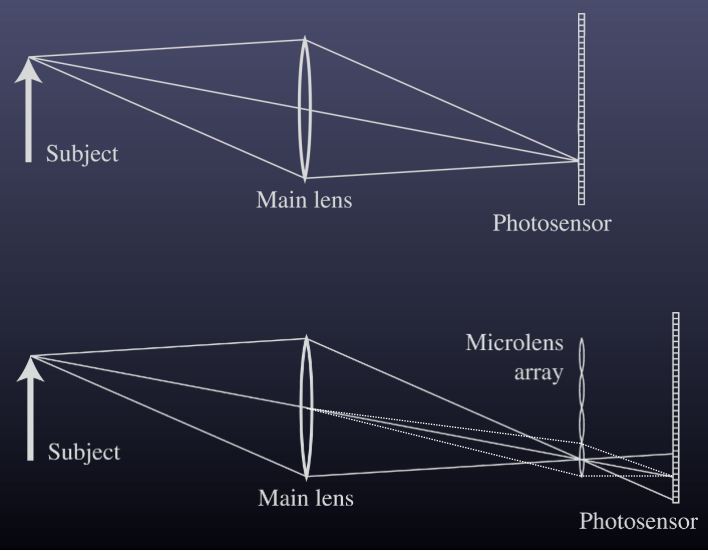

Conventional versus light field camera

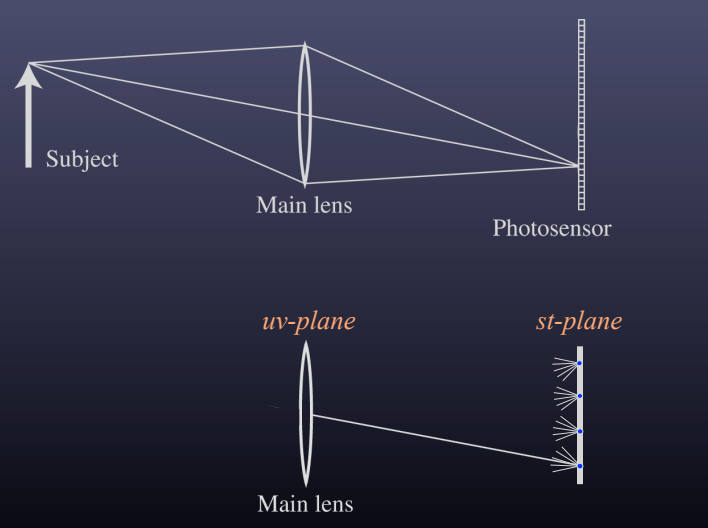

Prototype camera

Contax medium format camera

Kodak 16-megapixel sensor

Adaptive Optics microlens array

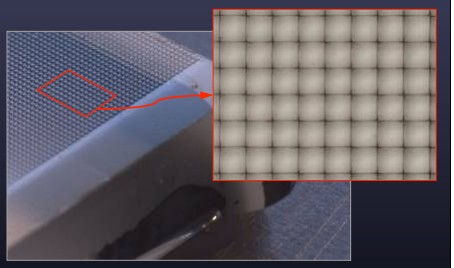

125μ square-sided microlenses

- *4000 × 4000 pixels ÷ 292 × 292 lenses = 14 × 14 pixels per lens*

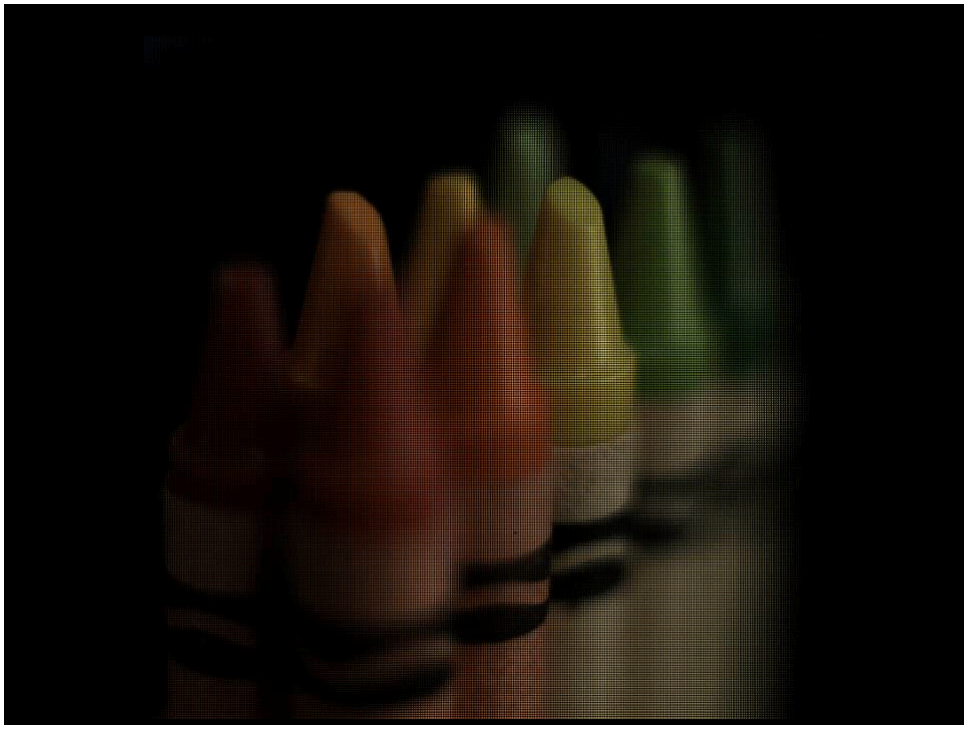

Digitally stopping-down

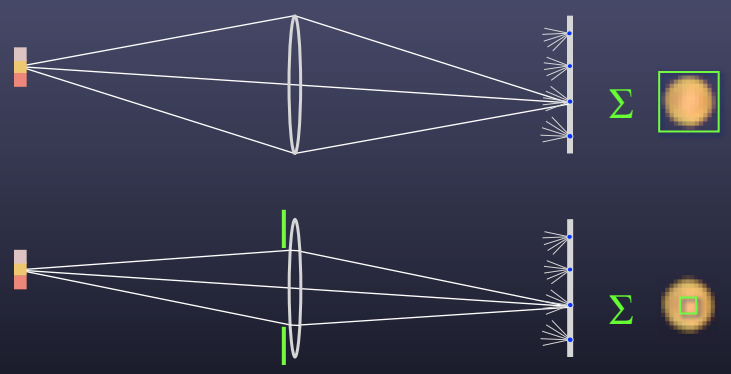

- stopping down = summing only the central portion of each microlens

Digital refocusing

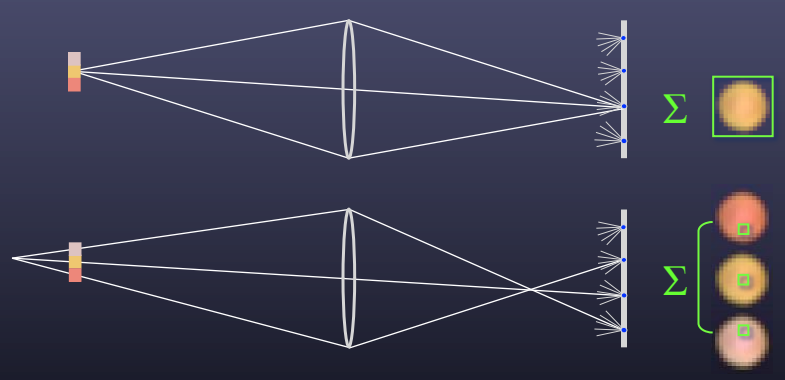

- refocusing = summing windows extracted from several microlenses

Example of digital refocusing

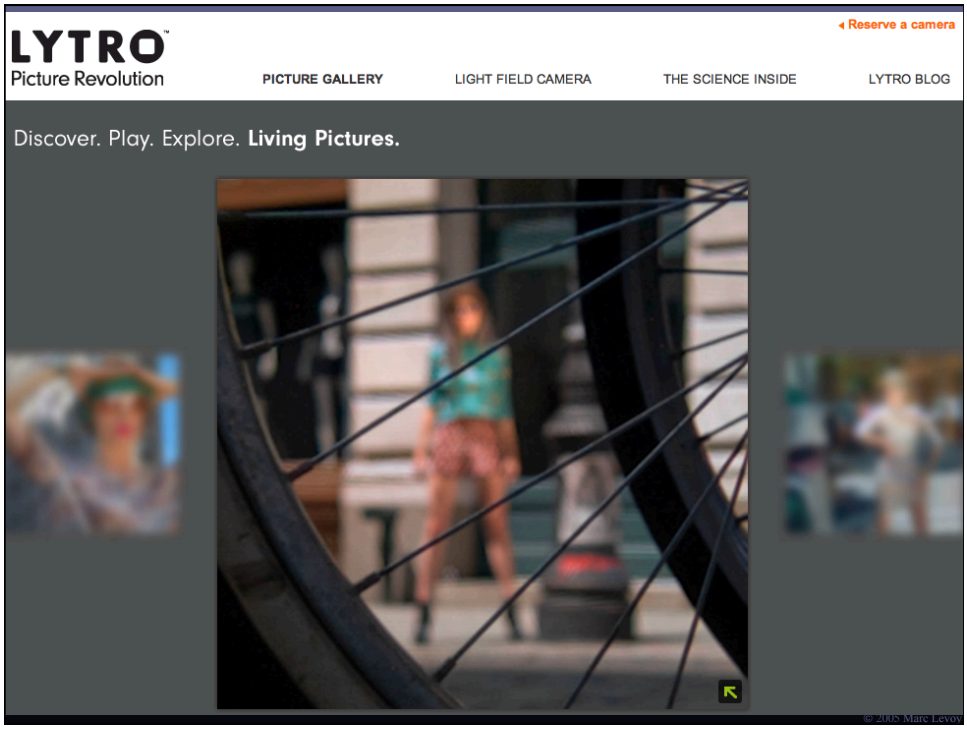

Digitally moving the observer

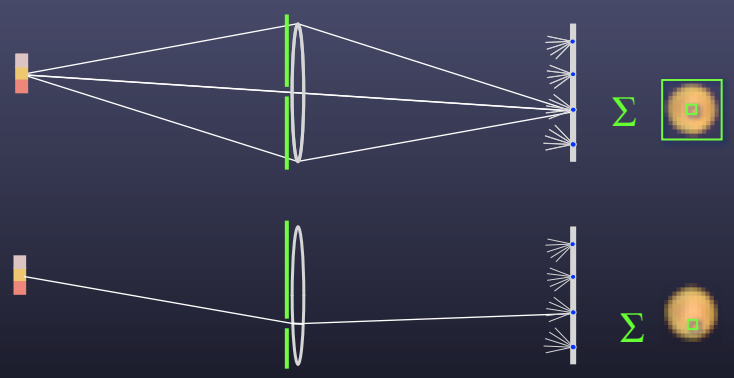

- moving the observer = moving the window we extract from the microlenses

Example of moving the observer

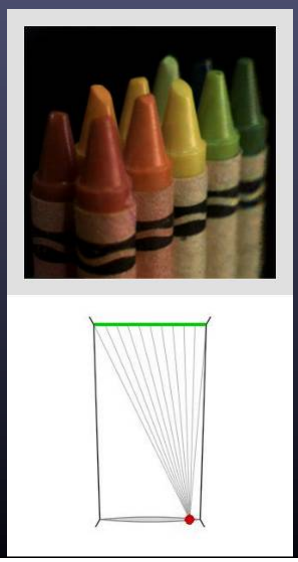

Moving backward and forward

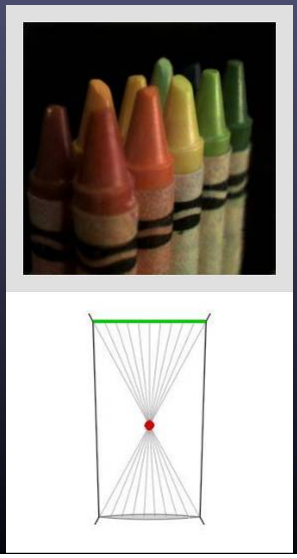

3D Lumigraph

One row of s,t plane

i.e., hold t constant

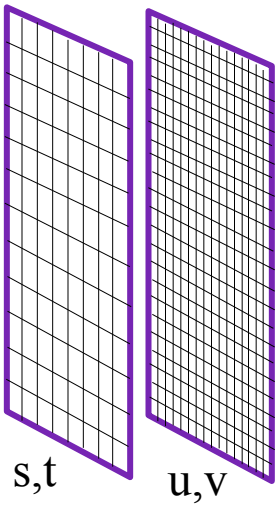

3D Lumigraph

One row of s,t plane

- i.e., hold t constant
- thus s,u,v
- a “row of images”

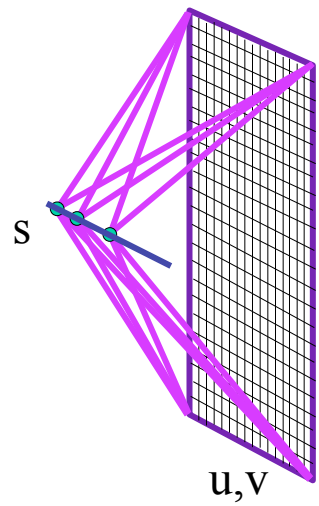

P(x,t) by David Dewey

Other ways to sample Plenoptic Function

Moving in time:

- Spatio-temporal volume: $P(\theta,\phi,t)$
- Useful to study temporal changes
- Long an interest of artists:

Claude Monet, Haystacks studies

Space-Time Images

Other ways to slice the plenoptic function…

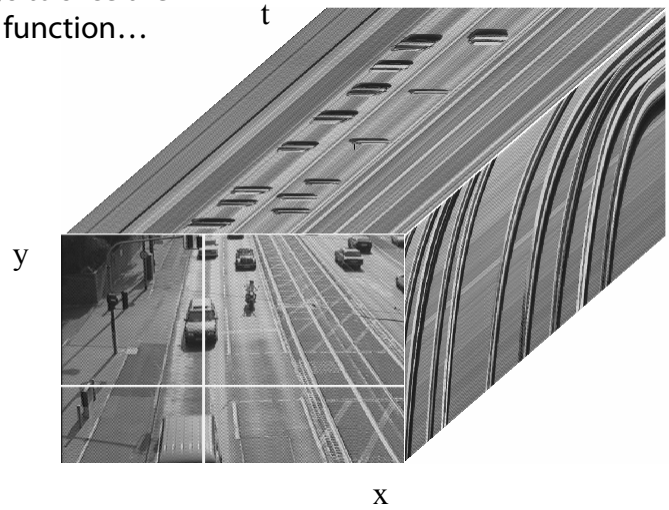

Stanford Multi-Camera Array

- 640 × 480 pixels × 30 fps × 128 cameras
- synchronized timing
- continuous streaming
- flexible arrangement

Stanford Multi-Camera Array

Demo #1: High Speed Video

- N cameras, each running at 30Hz
- Stagger the frames of cameras by $1/N^{th}$ of a frame
- Align images to single perspective

Video: 52 cameras, 1560 Hz

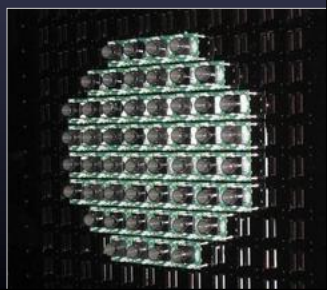

Video: 52 cameras, 1560 Hz

Demo #2: High Resolution Video

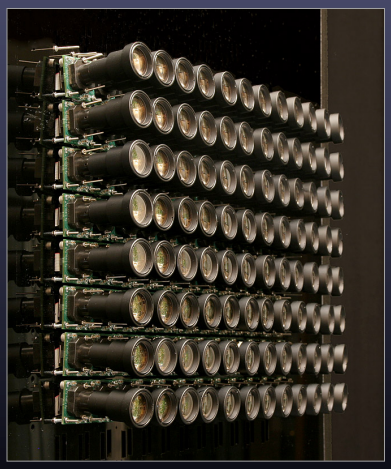

- 12 × 8 array of VGA cameras
- total field of view = 29° wide
- seamless stitching
- cameras individually metered

Tiled Video: 7 Megapixels

Tiled Video: 7 Megapixels

Camera Array: Portable Version

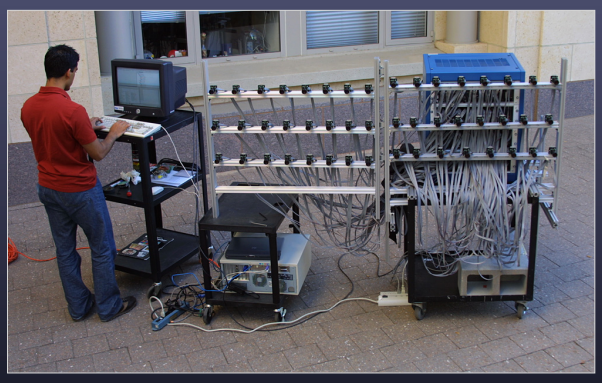

- 48 cameras in 16 x 3 layout
- 2m wide baseline

Synthetic Aperture Focusing: Scene

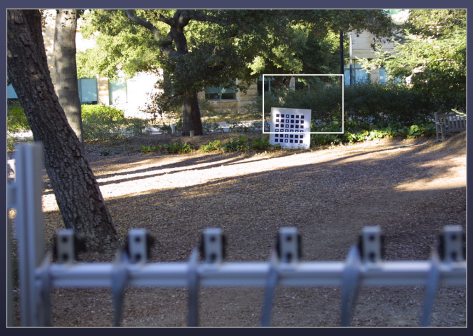

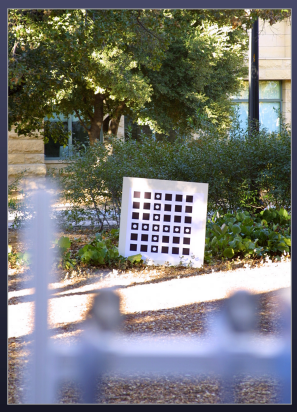

- distance to occluder 33m
- distance to targets 40m
- field of view at target 3m

Synthetic Aperture Focusing: Results

Synthetic aperture sequence

Synthetic Aperture Focusing: Results

Synthetic aperture sequence

Synthetic Aperture Focusing

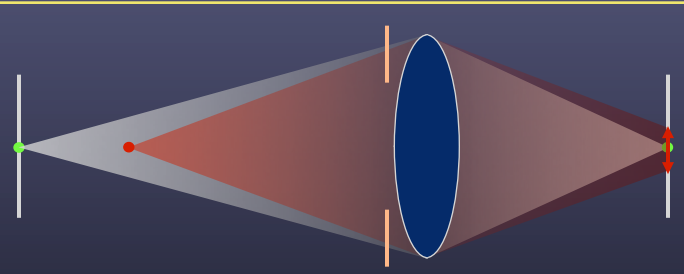

Synthetic Aperture Focusing

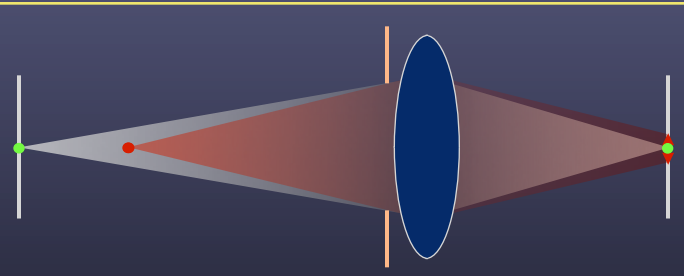

Digital refocusing

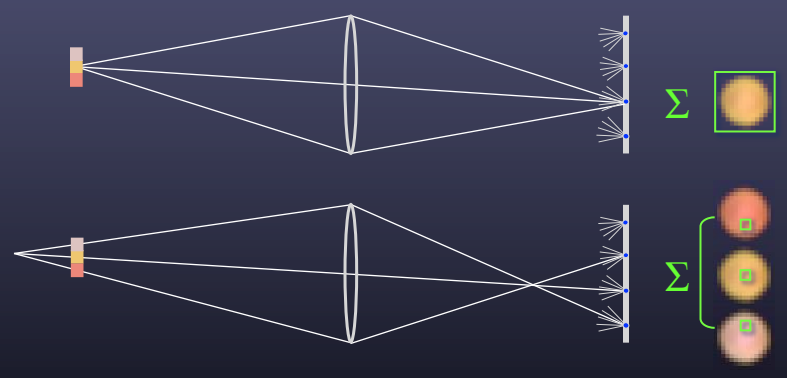

- refocusing = summing windows extracted from several microlenses

Suppressing Contributions of Occluders

(a) (b) (c)

Figure 11: Matted synthetic aperture photography. (a) A sample image from one of 90 cameras used for this experiment. (b) The synthetic aperture image focused on the plane of the people, computed by aligning and averaging images from all 90 cameras as described in the text. (c) Suppressing contributions from static pixels in each camera yields a more vivid view of the scene behind the occluder. The person and stuffed toy are more clearly visible.

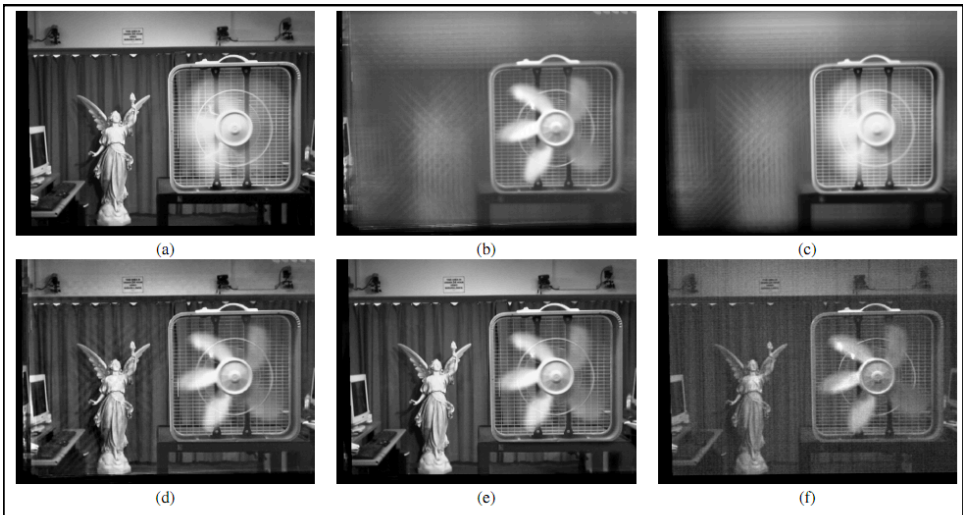

Figure 12: Hybrid synthetic aperture photography for combining high depth of field and low motion blur. (a-c) Images captured of a scene simultaneously through three different apertures: a single camera with a long exposure time (a), a large synthetic aperture with short exposure time (b), and a large synthetic aperture with a long exposure time. Computing (a+b-c) yields image (d), which has aliasing artifacts because the synthetic apertures are sampled sparsely from slightly different locations. Masking pixels not in focus in the synthetic aperture images before computing the difference (a + b - c) removes the aliasing (e). For comparison, image (f) shows the image taken with an aperture that is narrow in both space and time. The entire scene is in focus and the fan motion is frozen, but the image is much noisier.

Hybrid Synthetic Aperture – High DOF/Low Motion Blur