Light Field Rendering

Marc Levoy *Pat Hanrahan*

Computer Science Department
Stanford University

Light Field Rendering

Marc Levoy *Pat Hanrahan*

Computer Science Department
Stanford University

Apple's QuickTime VR

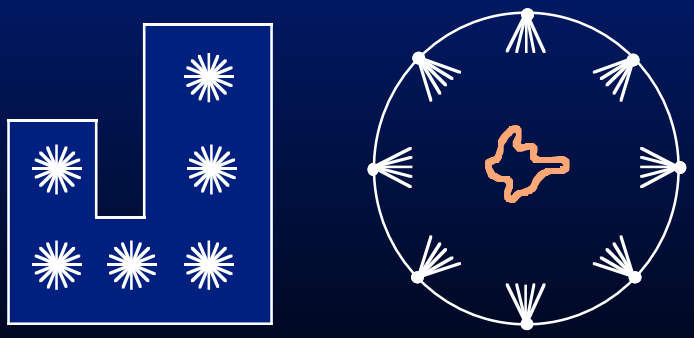

Outward Inward

Image-Based Rendering

Advantages

- Modest computation compared to classical C.G.
- Cost independent of scene complexity
- Imagery from real or virtual scenes

Limitations of precomputation

- Static scene geometry
- Fixed lighting
- Fixed look-from or look-at point

Generating New Views

Problem: fixed vantage point/center

One Solution: view interpolation

- Interpolating between range images (Chen and Willams, 1993)
- Correspondences and epipolar analysis (McMillan and Bishop, 1995)

-> Requires depths or correspondences:

must be extracted from acquired imagery

relatively expensive and error-prone morph

Related Work

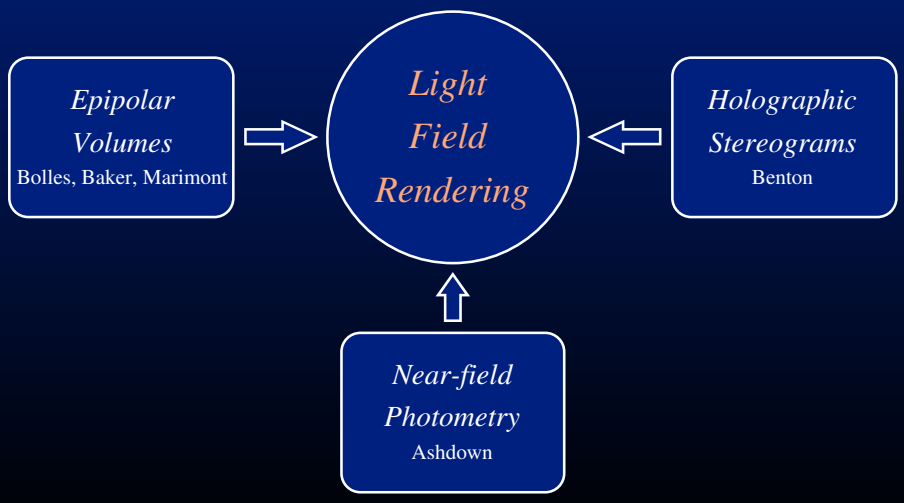

Light Fields

Gershun's and Moon's idea of a light field:

Radiance as a function of a ray or line: L (x, y, z, θ, ϕ)

- In "free space" (no occluders) 5D reduces to 4D
  - Exterior of the convex hull of an object
  - Interior of an environment
- Images are 2D slices
  - Insert acquired imagery
  - Extract image from a given viewpoint

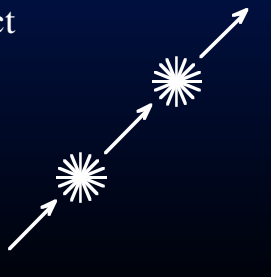

4D Light Field

Major Issues

- Representation
- Creation
- Display
- Acquisition
- Compression

Light Slab or Beam

Two-plane parameterization

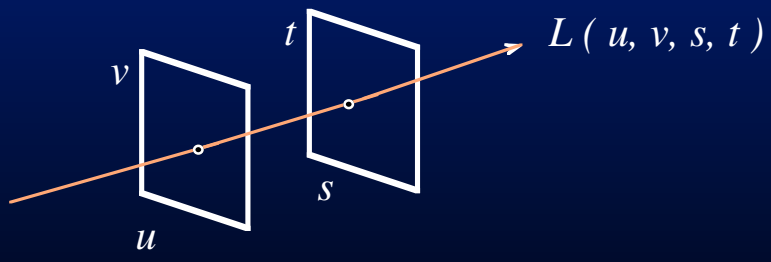

- Light flows from *uv* plane to *st* plane
- Planes in any position
- Parameters *uv* and *st* between 0 and 1

Creating a Light Field

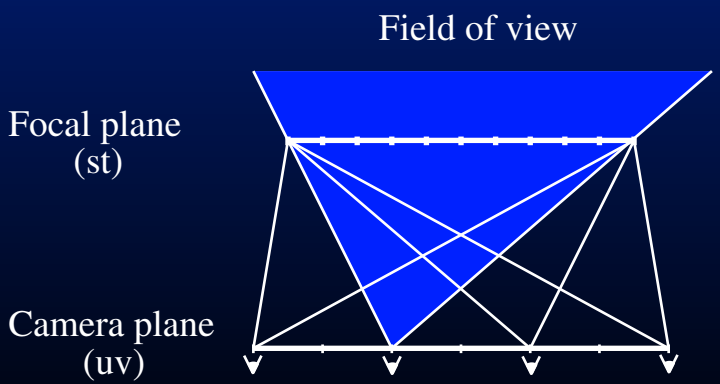

Array of Images

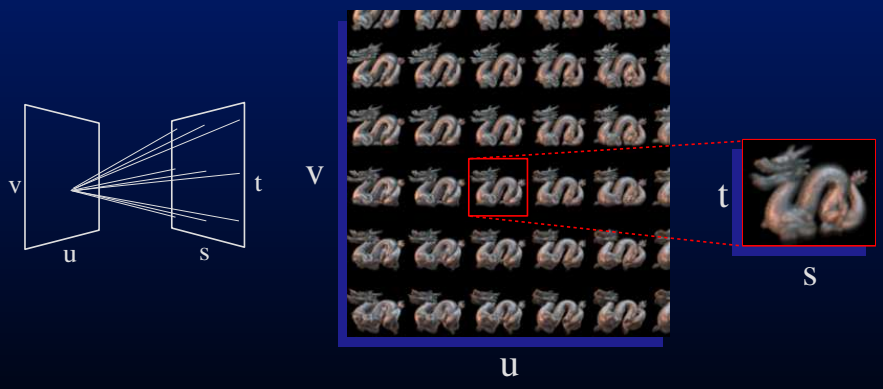

Fast Display

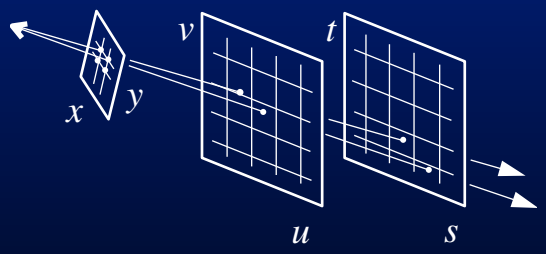


```
foreach x, y
    compute u,v,s,t
    I(x,y) = L(u,v,s,t)
```

2-2D Projective Mappings

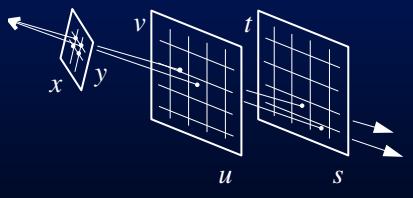

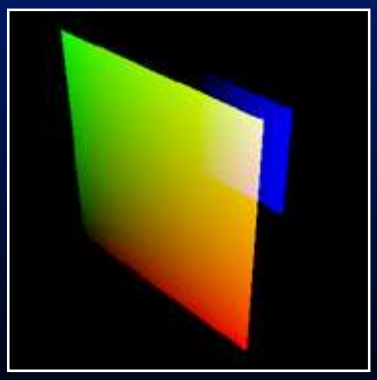

Point of View Rendering

4D Interpolation

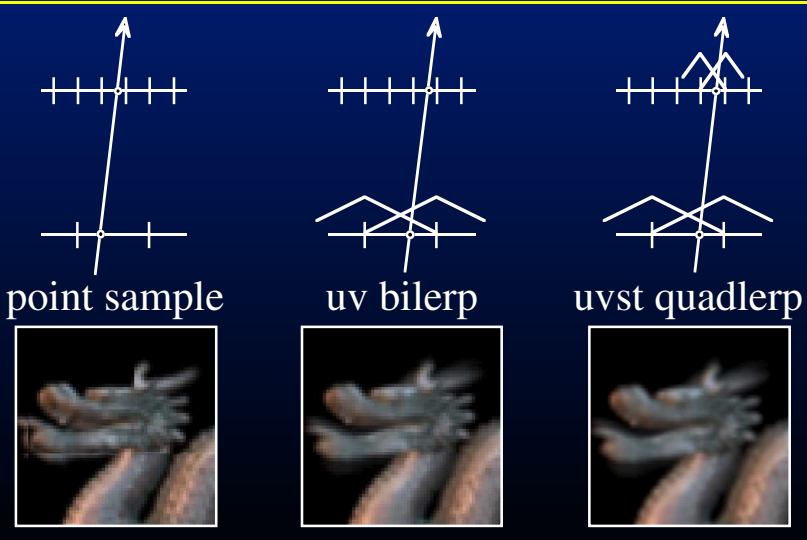

Point Space Interpretation

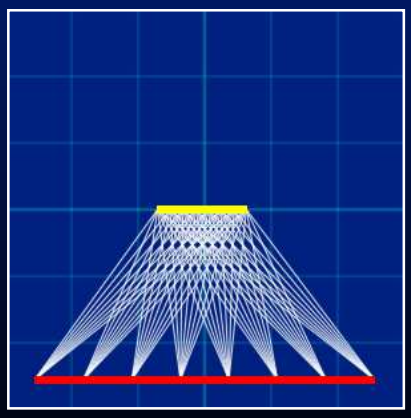

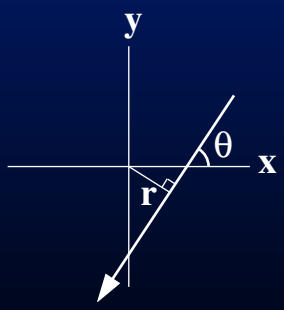

Parallel Planes

Line Space Interpretation

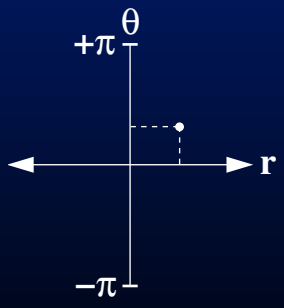

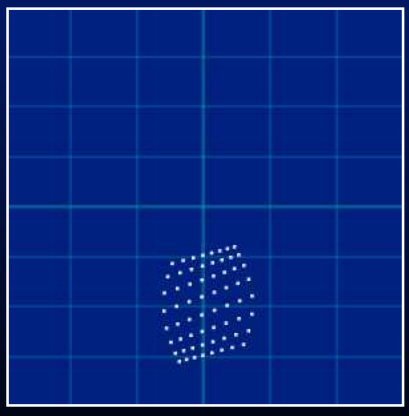

Parallel Planes

Line Space

- Coverage
- Uniform sampling
- Antialiasing: line prefilter = aperture+pixel filter

Pixel filter + Aperture filter = Ray filter

Aperture Filtering

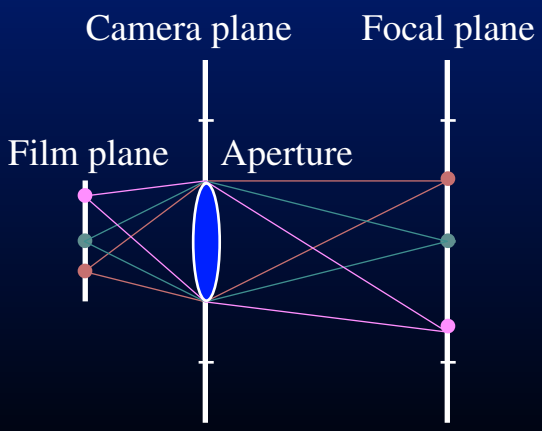

Digitization Issues

Design decisions

- Planar vs. spherical motion (Apple's ObjectMaker)
- Inward vs. outward looking
- Computer vs. human controlled

Experimental Light Field Gantry

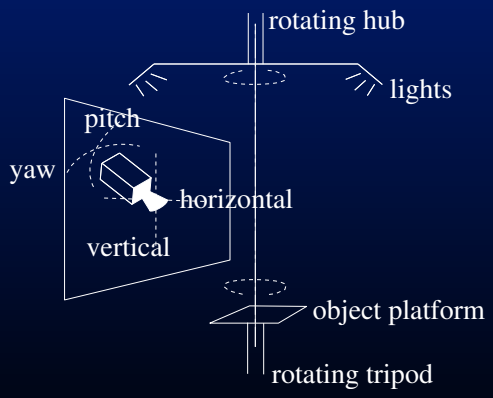

Compression Issues

- Need ≥ 100:1
- 4D redundancy
  - Nearby views are similar
  - Reflection/texture may be smooth
- Fast decoding
- Decompress 2D subset

Solution similar to: Beers, Agrawala, Chaddha, "Rendering from compressed textures"

Vector Quantization + Entropy Coding

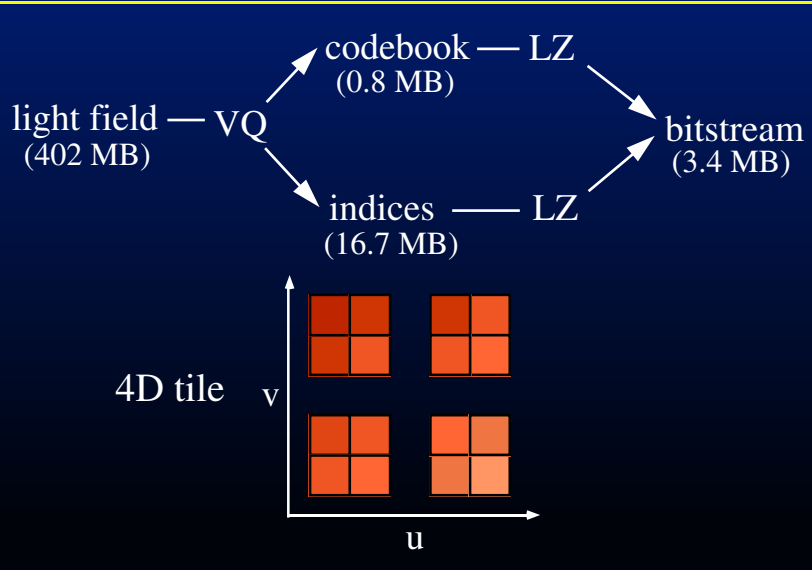

Compression Example

Original

Compressed 120:1

Performance summary

Digitization	
number of slabs	4
images per slab	32 x 16
pixels per image	256 x 256
raw size	402 MB
Compression	
VQ coding	24 : 1
entropy coding	5 : 1
total compression	120 : 1
compressed size	3.4 MB

Summary

- Represent 4D light field using two-planes
- Creation
  - Aperture equivalent to line space filter
  - Planar gantry for automatic acquisition
- Display
  - Simple projective map + linear resampling
  - Cost roughly 2x cost of texture mapping
- 100:1 compression makes system practical

Discussion

Limitations

- Sampling density must be high
- Large, densely occluded environments
- Fixed illumination, static scenes

Future Directions

- Light fields as first-class graphics primitives
- Rethink representations of light in graphics/vision
- Planar array of cameras
- Measuring and representing reflectance

Acknowledgements

Lucas Pereira, Li-Yi Wei, Ben Zhu	Light Field Software
David Addleman, George Dabrowski	Cyberware Retrofit
James Davis, Craig Kolb	Calibration
Navin Chaddha	VQ Compressor
Eric Chen, Michael Chen	Apple ObjectMaker
Brian Curless	Scanned Models
Julie Dorsey, Matt Pharr	Rendering
Tamara Munzner	Video
Celeste Fowler, Tom Funkhouser ,Seth Teller	Soda Hall Radiosity
Nina Amenta, Alain Fourner, Bob Lewis	Ideas
Michael Cohen, Gavin Miller	Motivation
Frank Crow at Interval	Money

Try It!

http://www-graphics.stanford.edu/