# Lecture 11: The Light Field and Image-Based Rendering

Visual Computing Systems CMU 15-869, Fall 2014

## Demo (movie)



Royal Palace: Madrid, Spain

# Image-based rendering (IBR)

- So far in course: rendering = synthesizing an image from a 3D model of the scene
  - Model of scene includes: cameras, geometry, lights, materials

### **Today:** synthesizing novel views of a scene from <u>existing images</u>

- Where do the input images come from?
  - Previously synthesized by a renderer
  - Acquired from the real world (photographs)



# Why does this view look wrong?



# What's going on? consider image resampling



### $\odot$ = Required image sample for view 2

Scene: as represented by image rendered from view 1

(Keep in mind that although I've drawn little squares the image is just a collection of samples)

### Synthesizing novel scene views from an existing image

- If scene lies approximately on a plane, affine transform of the image from view 1 yields accurate image of scene from view 2
- This transform is just texture mapping
- Have I made assumptions about the scene in addition to the assumption of planar scene geometry?
  - Hint: think about surface reflectance

### Non-planar scene



 $\odot$  = Image sample for view 2

### Scene: (represented by image rendered from view 1)

### Non-planar scene



 $\odot$  = Image sample for view 2

# Synthesis of novel scene views with correct perspective requires non-affine transformation of original image

### Scene: as represented by image + depth rendered from view 1

## Artifact: undersampled source image



### Scene: represented by image + depth rendered from view 1

Undersampling: The image generated from view 1 contains a sparse sampling of surface regions oriented at a grazing angle to view 1.

 (the yellow surface region
View 2 projects to one pixel in view 1, but several pixels in view 2)

X

### Artifact: disocclusion



Disocclusion: surface region visible from view 2 was not visible at all from view 1



## **Disocclusion examples**



[Credit: Chaurasia et al. 2011]



[Credit: Chen and Williams 93]

## View interpolation



### Combine results from closest pre-existing views Question: How to combine?

## Sprites



Novel view of object synthesized from rendering sprites

**Prerendered Textures** 





## Microsoft Talisman

Proposed GPU designed to accelerate image-based rendering for interactive graphics (motivating idea: exploit frame-to-frame temporal locality by not rendering entire scene each frame)



(renders geometric models to images)

### [Torborg and Kajiya 96]

Implements image transform and image compositing operations

## **Microsoft Talisman**



### Each object is rendered separately into its own image layer (intent: high-quality renderings from 3D model, but not produced at real-time rates)

Image layer compositor runs at real-time rates As scene changes (camera/object motion, etc.), image layer compositor transforms each layer accordingly, then composites image layers to produce complete frame

System detects when image warp likely to be inaccurate, makes request to re-render layer

[Torborg and Kajiya 96]

### Image-based rendering in interactive graphics systems

- **Promise: render complex scenes efficiently by manipulating** preexisting images
- **Reality: systems have always suffered from artifacts that have** prevented them from being a viable replacement for rendering from a detailed, 3D scene description
  - Not feasible to prerender images for all possible scene configurations and views
  - Decades of research on how to minimize artifacts from missing information (intersection of graphics and vision: understanding what's in an image helps fill in missing information... and vision is unsolved)

# **Temporal reprojection anti-aliasing**

- **Common post-processing technique in many games** 
  - Goal: achieve higher quality shader results without increasing shading cost.
- Idea: combine rendered image pixels (or multi-sample buffer samples) from last frame with rendered output from current frame
- **Requires application to store previous frame's color buffer**
- To synthesize final frame at time *t*:
  - **Render scene at** *t* **as usual to produce color\_buffer**<sub>t</sub>
  - For each pixel p, compute velocity vector (based on camera and/or object movement
  - Use velocity vector to compute screen location of corresponding surface fragment's position *p*<sub>t-1</sub> in the previous frame ("reprojection" of surface to previous time)
  - Sample from color\_buffer<sub>t-1</sub> at  $p_{t-1}$ . Combine results with color\_buffert[*p*] to get final value for *p* in the current frame
- **Result:** better anti-aliasing of shading function: two shading samples per surface per pixel at a cost of only one shade per frame.
- Many heuristics to detect disocclusions, minimize errors in reproduction





### Good system design: efficiently meeting goals, subject to constraints

- **New graphics application goals:** 
  - Map the world
  - Navigate popular tourist destinations
  - Non-goal: virtual reality experience (artifact-free, real-time frame rate, viewer can navigate anywhere in the scene)
  - **Changing constraints:**
  - Cannot pre-render all scene configurations?
    - **Ubiquity of cameras**
    - **Cloud-based graphics applications: enormous storage capacity**
    - Bandwidth now available to access server-side capacity from clients

## **Google Street View**



Goal: orient/familiarize myself with 16th and Valencia, San Francisco, CA

Imagine complexity of modeling and rendering this scene (and then doing it for all of the Mission, for all of San Francisco, of California, of the world...)

### **Google Street View**



### Imagine if your GPU produced images that had geometric artifacts like this!

### Imagine if moving through a 3D rendered environment in a game had transitions like Google Maps

### **Photo-tourism** (now Microsoft Photosynth)





Input: collection of photos of the same scene

### Goal: get a sense of what it's like to be at Notre Dame Cathedral in Paris



### **Output:** sparse 3D representation of scene, 3D position of cameras for all photos

### **Alternative projections**



Pushbroom projection



[Image credit: Roman 2006]

### Each pixel column in image above is column of pixels from a different photograph

### **Result is orthographic projection in X**, perspective projection in Y

## **The Light Field**

[Levoy and Hanrahan 96] [Gortler et al., 96]



# Light-field parameterization

Light field is a 4D function (represents light in free space: no occlusion)



**Efficient two-plane parameterization** 

Line described by connecting point on (u,v) plane with point on (s,t) plane If one of the planes placed at infinity: point + direction representation

Levoy/Hanrahan refer to representation as a "light slab": beam of light entering one quadrilateral and exiting another



[Image credit: Levoy and Hanrahan 96]

## Sampling of the light field



(full light field is 4D function)

## Line-space representation

Each line in Cartesian space<sup>\*\*</sup> represented by a point in line space



**\*\*** Shown here in 2D, generalizes to 3D Cartesian lines

[Image credit: Levoy and Hanrahan 96]

## Sampling lines



# To be able to reproduce all possible views, light field should uniformly sample all possible lines

### Lines sampled by one slab

# Four slabs sample lines in all directions

## Acquiring a light field



### Measuring light field by taking multiple photographs (In this example: each photograph: constant UV)

[Image credit: Levoy and Hanrahan 96]

### Light field storage layouts



**(a)** 



**(b)** 



u

t





S

### [Image credit: Levoy and Hanrahan 96]

## Light field inside a camera



### **Ray space plot**

U

# Question: what does a pixel measure?

X

## Light field inside a camera





U

## Readings

### M. Levoy and P. Hanrahan. Light Field Rendering. SIGGRAPH 1996