Sampling and Reconstruction of Visual Appearance: From Denoising to View Synthesis

CSE 274 [Fall 2021], Lecture 6

Ravi Ramamoorthi

http://www.cs.ucsd.edu/~ravir

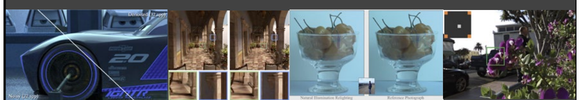

Basics of Denoising, Frequency Analysis

Monte Carlo Rendering (biggest application)

- Basic idea of denoising
- Frequency analysis one key concept
- Presentation of key papers at next class
- Relevant to other applications as well

Cook et al. [1984] results

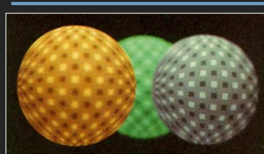

depth of field

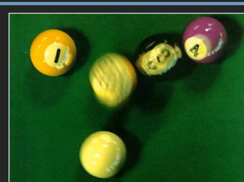

motion blur

soft shadows

glossy reflection

PRADEEP SEN SIGGRAPH2015

Motivation

- Distribution effects (depth of field, motion blur, global illumination, soft shadows) are slow. Many dimensions sample

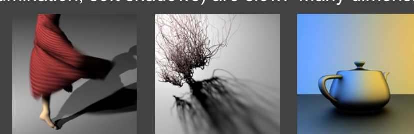

- Ray Tracing physically accurate but slow, not real-time
- Can we adaptively sample and filter for fast, real-time?

Sample result

Path-traced scene / Filtered scene

[Kalantari et al. 2015]

scene by Jo Ann Elliott

4 samples/pixel (48.8 sec)

using only post-process filter!

PRADEEP SEN SIGGRAPH2015

Adaptive sampling + reconstruction

- **These algorithms use 2 kinds of noise reduction strategies, sometimes combined:**
  1. **Adaptive sampling algorithms**
     - **Use information from renderer to position new samples better to reduce noise**
  2. **Reconstruction (filtering) algorithms**
     - **Use information from renderer to remove MC noise directly**
- **Both methods have been explored in the past, but new algorithms make remarkable advances**

PRADEEP SEN SIGGRAPH2015

Multi-Dimensional Adaptive Sampling

- Hachisuka, Jarosz, … Zwicker, Jensen [MDAS 2008]
- Scenes with motion blur, depth of field, soft shadows
- Involves high-dimensional integral, converges slowly
- Exploit high-dimensional info to sample adaptively
- Sampling in 2D image plane or other dims inadequate
  - Need to consider full joint high-dimensional space

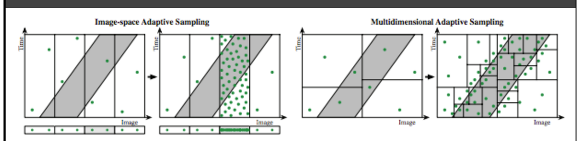

Multi-Dimensional Adaptive Sampling

Motion Blur and Depth of Field 32 samples per pixel

Resurgence (2008 -)

- Eurographics 2015 STAR report by Zwicker et al.
  - Papers below are key a-priori, frequency analysis methods
  - Many other approaches to be discussed in class
- [Durand et al. 2005] *Frequency analysis light transport*
  - Key theoretical ideas, but not initially very practical
- [Chai et al. 2000] Plenoptic Sampling (wedge spectrum)
- [Egan et al. 2009] First practical a-priori frequency method

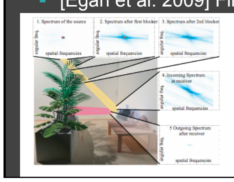

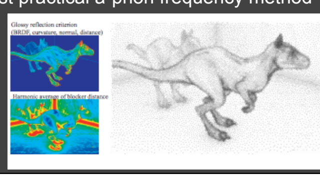

Background: Fourier Analysis

Analysis in the frequency (not spatial) domain

- Sum of sine waves, with possibly different offsets (phase)
- Each wave different frequency, amplitude

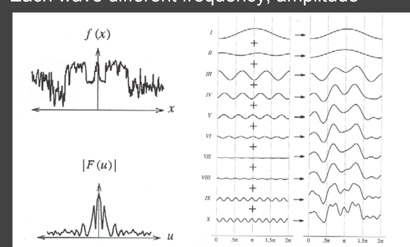

Figure 2.6 Wolberg

Fourier Transform

- Tool for converting from spatial to frequency domain

$$f(x) = \sum_{u=-\infty}^{+\infty} F(u) e^{2\pi i u x}$$

$$e^{2\pi i u x} = \cos(2\pi u x) + i \sin(2\pi u x)$$

$$i = \sqrt{-1}$$

- Or vice versa
- One of most important mathematical ideas
- Computational algorithm: Fast Fourier Transform
  - One of 10 great algorithms scientific computing
  - Makes Fourier processing possible (images etc.)
  - Not discussed here, but look up if interested

Fourier Transform

- Simple case, function sum of sines, cosines

$$f(x) = \sum_{u=-\infty}^{+\infty} F(u) e^{2\pi i u x}$$

$$F(u) = \int_0^1 f(x) e^{-2\pi i u x} dx$$

- Continuous infinite case

Forward Transform: $F(u) = \int_{-\infty}^{\infty} f(x) e^{-2\pi i u x} dx$

Inverse Transform: $f(x) = \int_{-\infty}^{+\infty} F(u) e^{2\pi i u x} du$

Fourier Transform

- Simple case, function sum of sines, cosines

$$f(x) = \sum_{u=-\infty}^{+\infty} F(u) e^{2\pi iux}$$

$$F(u) = \int_0^1 f(x) e^{-2\pi iux} dx$$

- Discrete case

$$F(u) = \sum_{x=0}^{x=N-1} f(x)\left[\cos(2\pi ux/N) - i\sin(2\pi ux/N)\right], \qquad 0 \le u \le N-1$$

$$f(x) = \frac{1}{N}\sum_{u=0}^{u=N-1} F(u)\left[\cos(2\pi ux/N) + i\sin(2\pi ux/N)\right], \qquad 0 \le x \le N-1$$

Fourier Transform: Examples 1

Single sine curve
(+constant DC term)

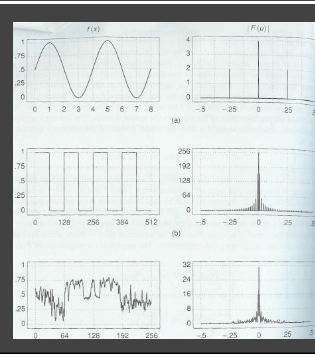

$$f(x) = \sum_{u=-\infty}^{+\infty} F(u) e^{2\pi iux}$$

$$F(u) = \int_0^1 f(x) e^{-2\pi iux} dx$$

Fourier Transform Examples 2

Forward Transform: $F(u) = \int_{-\infty}^{\infty} f(x) e^{-2\pi iux} dx$

Inverse Transform: $f(x) = \int_{-\infty}^{+\infty} F(u) e^{2\pi iux} du$

- Common examples

$f(x)$	$F(u)$
$\delta(x - x_0)$	$e^{-2\pi iux_0}$
1	$\delta(u)$
e^{-ax^2}	$\sqrt{\pi/a}\, e^{-\pi^2 u^2/a}$

Fourier Transform Properties

Forward Transform: $F(u) = \int_{-\infty}^{\infty} f(x) e^{-2\pi iux} dx$

Inverse Transform: $f(x) = \int_{-\infty}^{+\infty} F(u) e^{2\pi iux} du$

- Common properties
  - Linearity: $F(af(x) + bg(x)) = aF(f(x)) + bF(g(x))$
  - Derivatives: [integrate by parts] $F(f'(x)) = \int_{-\infty}^{\infty} f'(x) e^{-2\pi iux} dx = 2\pi iuF(u)$
  - 2D Fourier Transform
    Forward Transform: $F(u,v) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(x,y) e^{-2\pi iux} e^{-2\pi ivy} dxdy$
    Inverse Transform: $f(x,y) = \int_{-\infty}^{\infty}\int_{-\infty}^{+\infty} F(u,v) e^{2\pi iux} e^{2\pi ivy} dudv$
- Convolution (next)

Sampling Theorem, Bandlimiting

- A signal can be reconstructed from its samples, if the original signal has no frequencies above half the sampling frequency – Shannon
- The minimum sampling rate for a bandlimited function is called the Nyquist rate

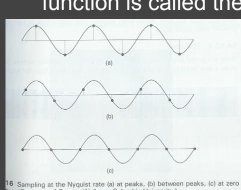

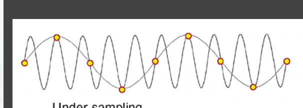

Sampling Theorem, Bandlimiting

- A signal can be reconstructed from its samples, if the original signal has no frequencies above half the sampling frequency – Shannon
- The minimum sampling rate for a bandlimited function is called the Nyquist rate
- A signal is bandlimited if the highest frequency is bounded. This frequency is called the bandwidth
- In general, when we transform, we want to filter to bandlimit before sampling, to avoid aliasing

Antialiasing

- Sample at higher rate
  - Not always possible
  - Real world: lines have infinitely high frequencies, can’t sample at high enough resolution
- Prefilter to bandlimit signal
  - Low-pass filtering (blurring)
  - Trade blurriness for aliasing

Ideal bandlimiting filter

- Formal derivation is exercise
  - Frequency domain

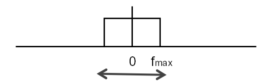

  - Spatial domain

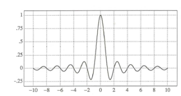

if full width $f_{max} = 1$

$$Sinc(x) = \frac{\sin \pi x}{\pi x}$$

Figure 4.5 Wolberg

Convolution 1

- Spatial domain: output pixel is weighted sum of pixels in neighborhood of input image
  - Pattern of weights is the “filter”

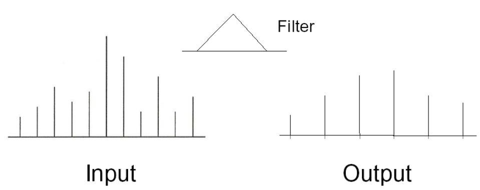

Convolution 2

- Example 1:

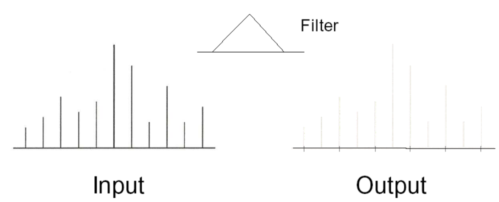

Convolution 3

- Example 1:

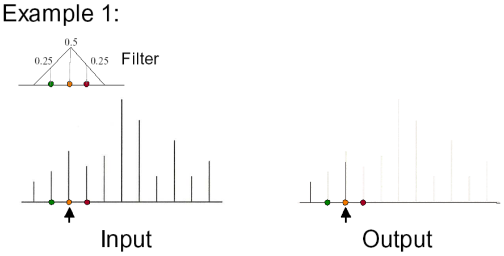

Convolution 4

- Example 1:

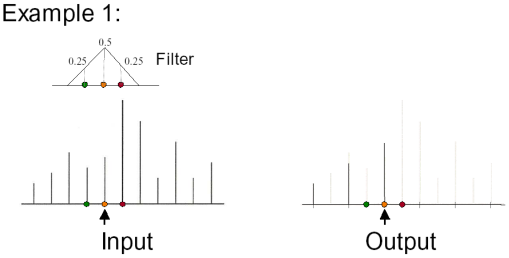

Convolution 5

- Example 1:

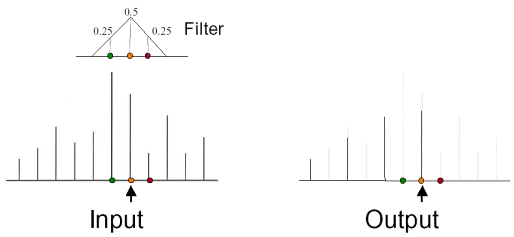

Convolution in Frequency Domain

Forward Transform: $F(u)=\int_{-\infty}^{\infty} f(x)e^{-2\pi iux}dx$

Inverse Transform: $f(x)=\int_{-\infty}^{+\infty} F(u)e^{2\pi iux}du$

- Convolution (f is signal ; g is filter [or vice versa])

$$h(y)=\int_{-\infty}^{+\infty} f(x)g(y-x)dx=\int_{-\infty}^{+\infty} g(x)f(y-x)dx$$

$$h=f*g \quad \text{or} \quad f\otimes g$$

- Fourier analysis (frequency domain multiplication) $H(u)=F(u)G(u)$

A Frequency Analysis of Light Transport

F. Durand, MIT CSAIL
N. Holzschuch, C. Soler, ARTIS/GRAVIR-IMAG INRIA
E. Chan, MIT CSAIL
F. Sillion, ARTIS/GRAVIR-IMAG INRIA

Illumination effects

- Blurry reflections:

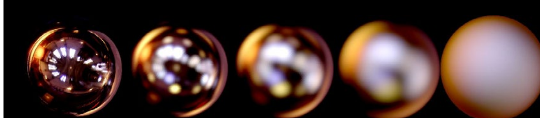

From [Ramamoorthi and Hanrahan 2001]

Illumination effects

- Shadow boundaries:

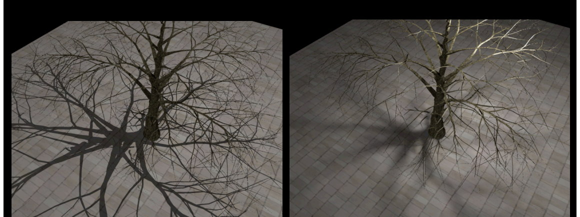

Point light source — Area light source

Problem statement

- How does light interaction in a scene explain the frequency content?
- Theoretical framework:
  - Understand the frequency spectrum of the radiance function
  - From the equations of light transport

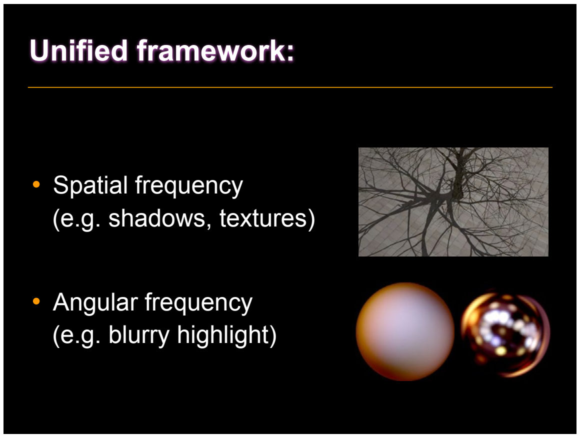
Unified framework:
• Spatial frequency
(e.g. shadows, textures)
• Angular frequency
(e.g. blurry highlight)

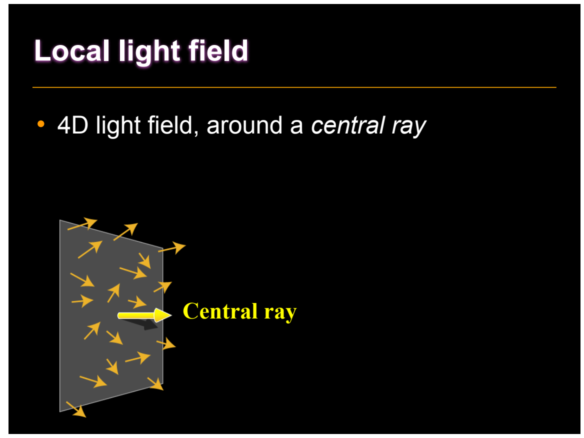
Local light field
• 4D light field, around a *central ray*
Central ray

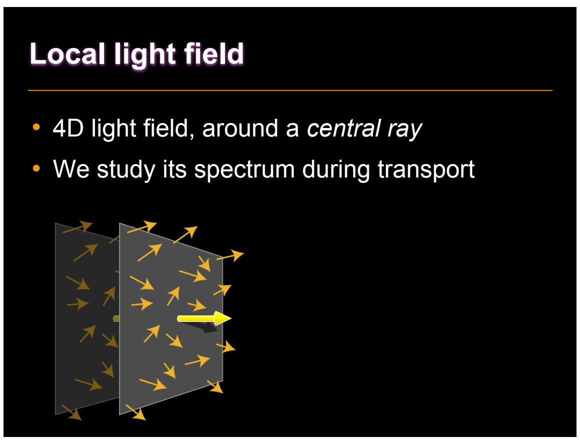
Local light field
• 4D light field, around a *central ray*
• We study its spectrum during transport

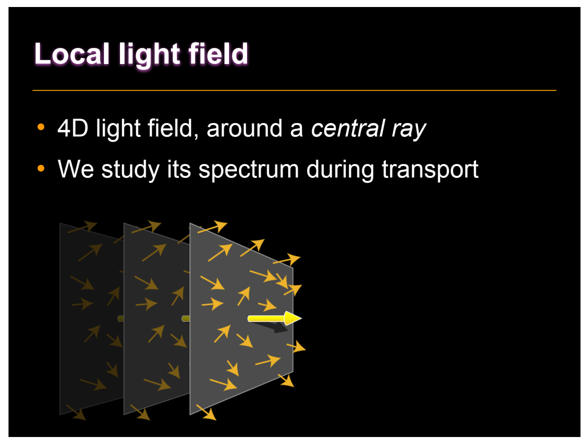
Local light field
• 4D light field, around a *central ray*
• We study its spectrum during transport

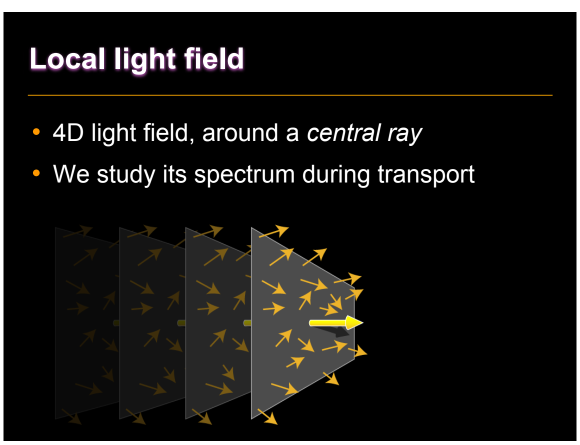
Local light field
• 4D light field, around a *central ray*
• We study its spectrum during transport

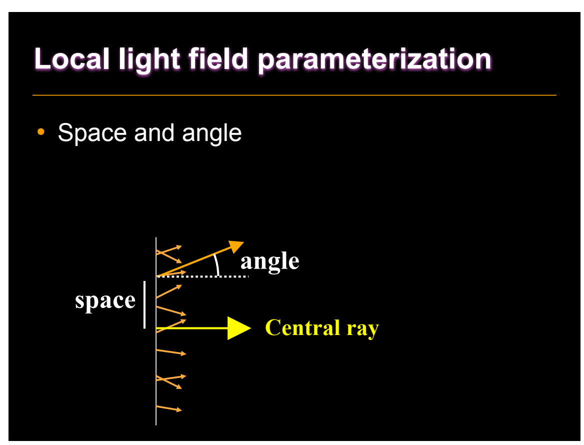
Local light field parameterization
• Space and angle
angle
space
Central ray

Local light field representation

- Density plot:

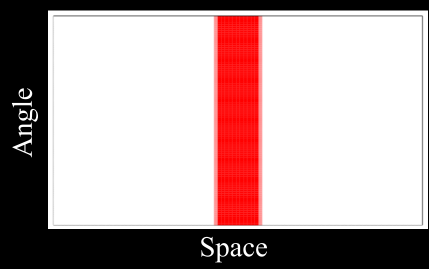

Local light field Fourier spectrum

- We are interested in the Fourier spectrum of the local light field
- Also represented as a density plot

Local light field Fourier spectrum

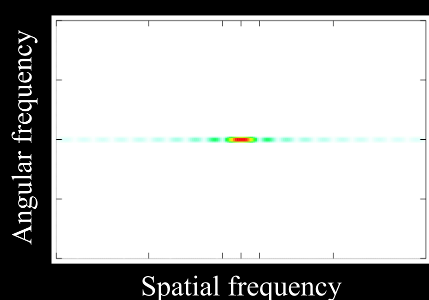

Fourier analysis 101

- Spectrum corresponds to blurriness:
  – Sharpest feature has size ~ $1/F_{max}$
- Convolution theorem:
  – Multiplication of functions: spectrum is convolved
  – Convolution of functions: spectrum is multiplied
- Classical spectra:
  – Box becomes sinc
  – Dirac becomes constant

Transport

- Shear: $x' = x - v\,d$

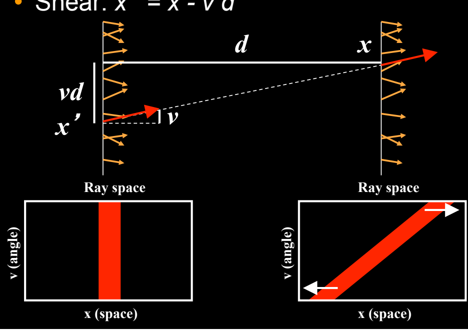

Transport in Fourier space

- Shear in primal: $x' = x - v\,d$
- Shear in Fourier, along the other dimension

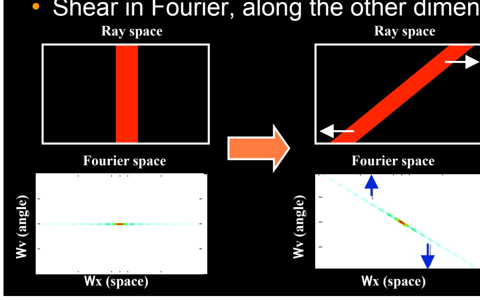

Transport becomes Shear

- This is consistent with light field spectra
  [Chai et al. 00, Isaksen et al. 00]

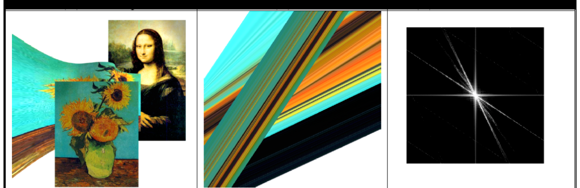

(d1)Scene image | (d2) EPI | (d3) Fourier transform of EPI

From [Chai et al. 2000]

BRDF integration

- Ray-space: ***convolution***
  - Outgoing light:
    convolution of incoming light and BRDF
  - For rotationally-invariant BRDFs
- Fourier domain: ***multiplication***
  - Outgoing spectrum: multiplication of incoming spectrum and BRDF spectrum

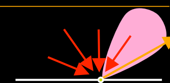

Adaptive shading sampling

- Per-pixel prediction of max. frequency (bandwidth)
  - Based on curvature, BRDF, distance to occluder, etc.
  - No spectrum computed, just estimate max frequency

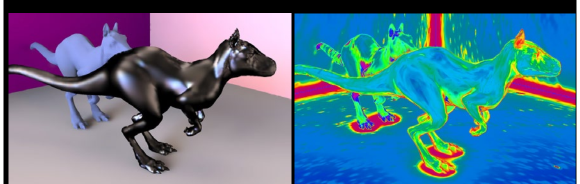

Per-pixel bandwidth criterion

Adaptive shading sampling

- Per-pixel prediction of max. frequency (bandwidth)
  - Based on curvature, BRDF, distance to occluder, etc.
  - No spectrum computed, just estimate max frequency

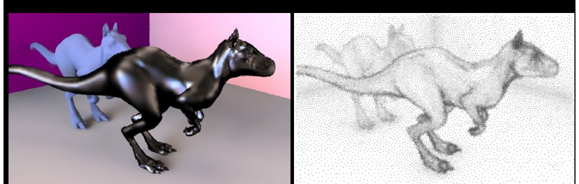

Shading samples

Uniform sampling

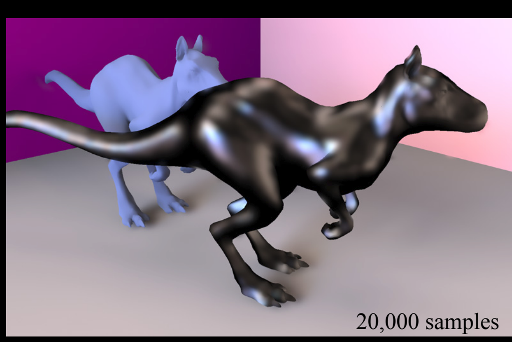

Adaptive sampling

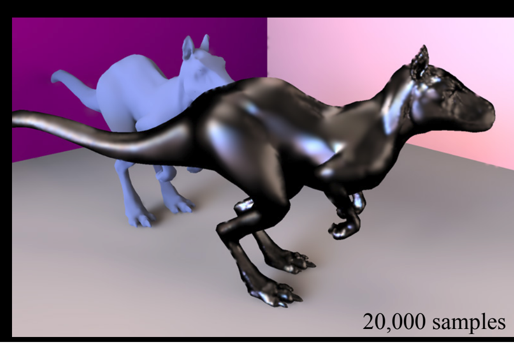

Plenoptic Sampling

- Plenoptic Sampling. *Chai, Tong, Chan, Shum 00*
- Signal-processing on light field
- Minimal sampling rate for antialiased rendering
- Relates to depth range, Fourier analysis
- Fourier spectra derived for 2D light fields for simplicity. Same ideas extend to 4D
- Key paper in many newer methods on sheared and axis-aligned filtering for adaptive sampling

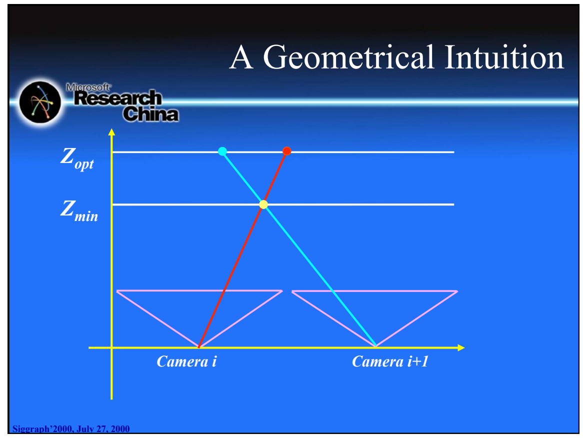

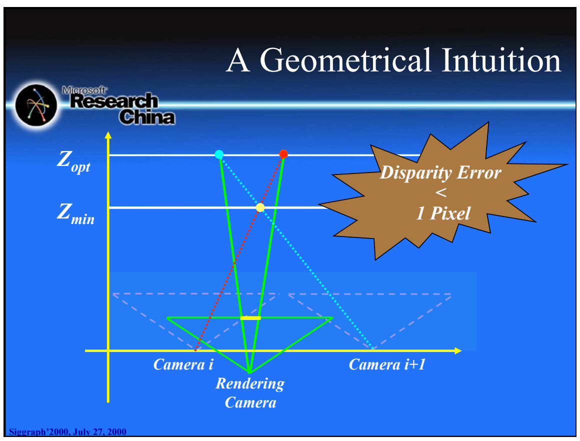

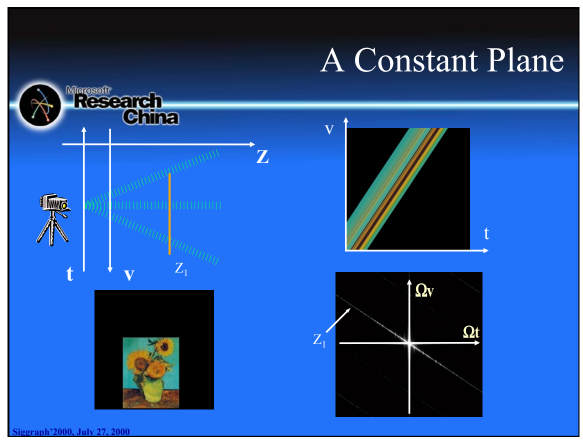

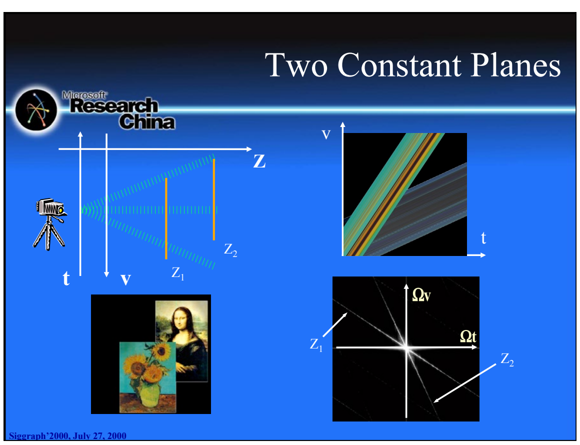

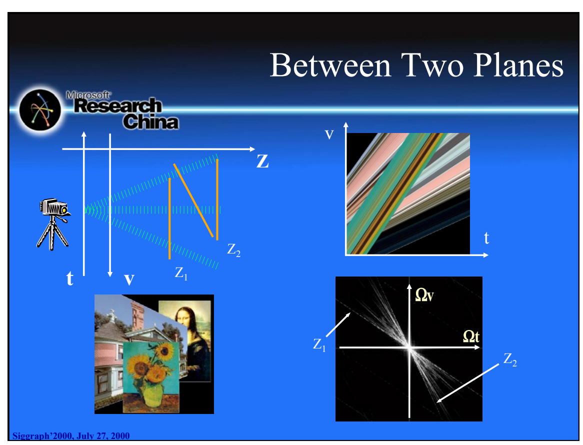

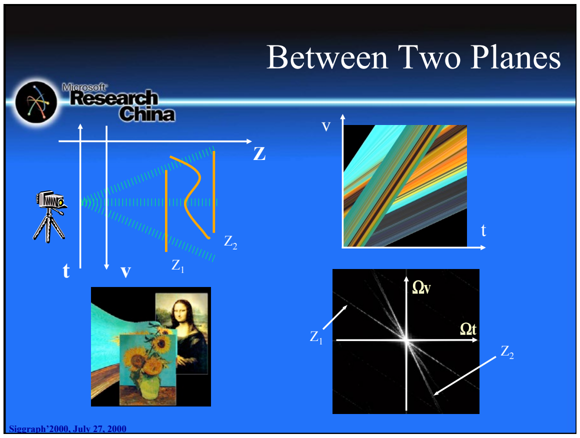
Between Two Planes
Microsoft Research China
Z
v
t
t
v
Z_1
Z_2
Ωv
Ωt
Z_1
Z_2
Siggraph'2000, July 27, 2000

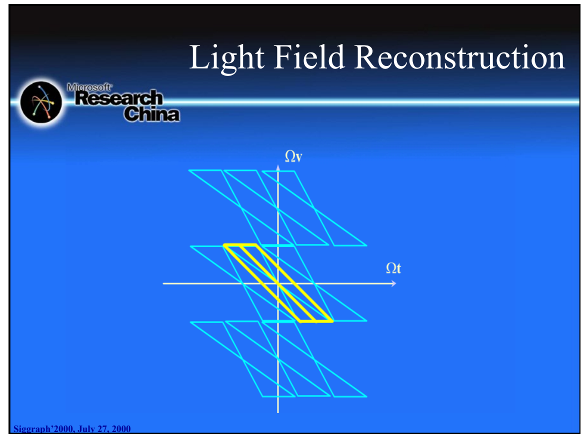
Light Field Reconstruction
Microsoft Research China
Ωv
Ωt
Siggraph'2000, July 27, 2000

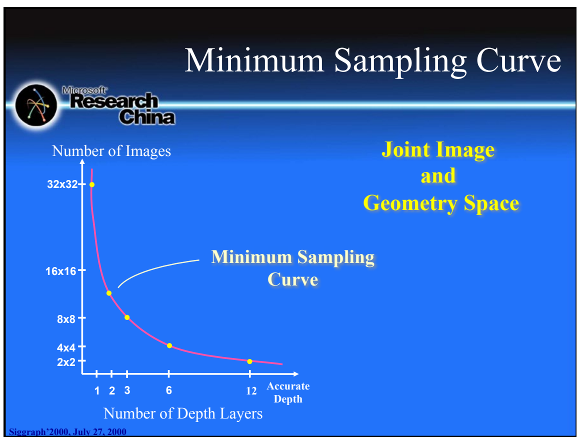
Minimum Sampling Curve
Microsoft Research China
Number of Images
32x32
16x16
8x8
4x4
2x2
1 2 3 6 12
Accurate Depth
Number of Depth Layers
Joint Image and Geometry Space
Minimum Sampling Curve
Siggraph'2000, July 27, 2000

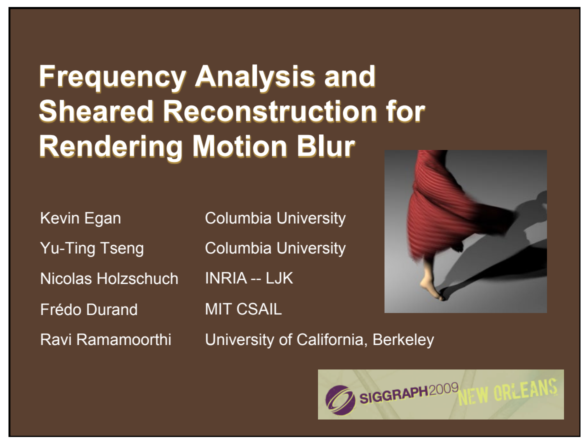
Frequency Analysis and Sheared Reconstruction for Rendering Motion Blur
Kevin Egan Columbia University
Yu-Ting Tseng Columbia University
Nicolas Holzschuch INRIA -- LJK
Frédo Durand MIT CSAIL
Ravi Ramamoorthi University of California, Berkeley
SIGGRAPH2009 NEW ORLEANS

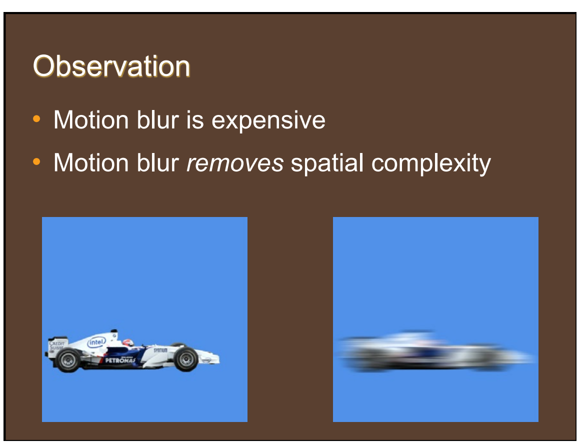
Observation
• Motion blur is expensive
• Motion blur *removes* spatial complexity

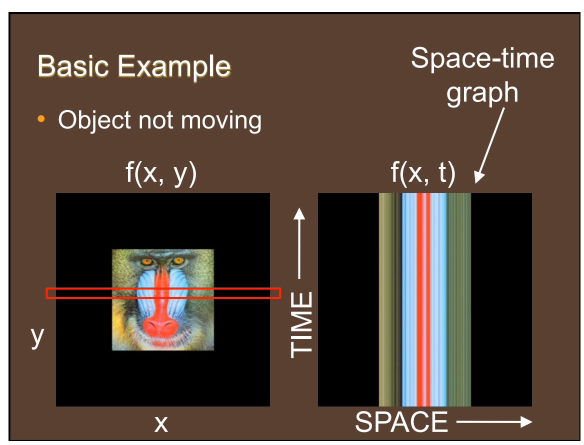
Basic Example
• Object not moving
Space-time graph
f(x, y)
f(x, t)
y
x
TIME
SPACE

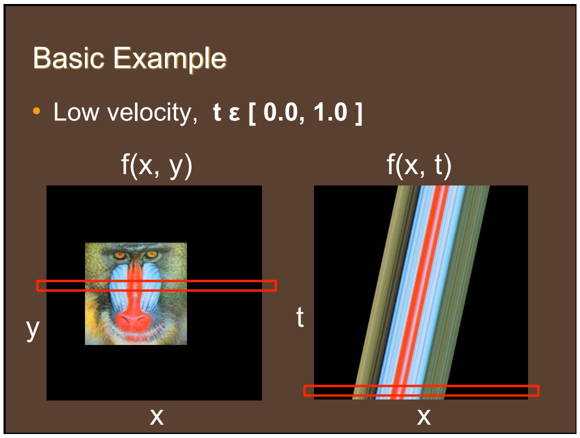
Basic Example
• Low velocity, t ε [0.0, 1.0]
f(x, y)
f(x, t)
y
x
t
x

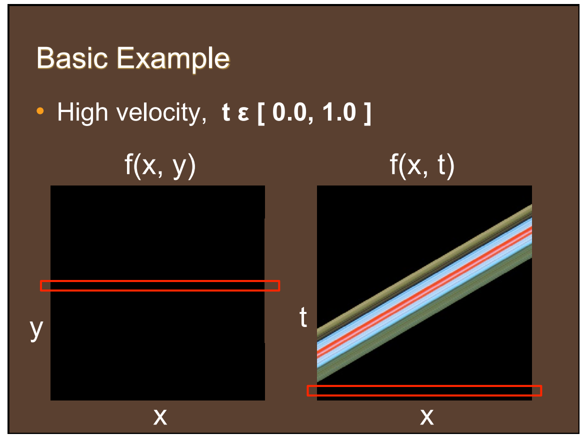
Basic Example
• High velocity, t ε [0.0, 1.0]
f(x, y)
f(x, t)
y
x
t
x

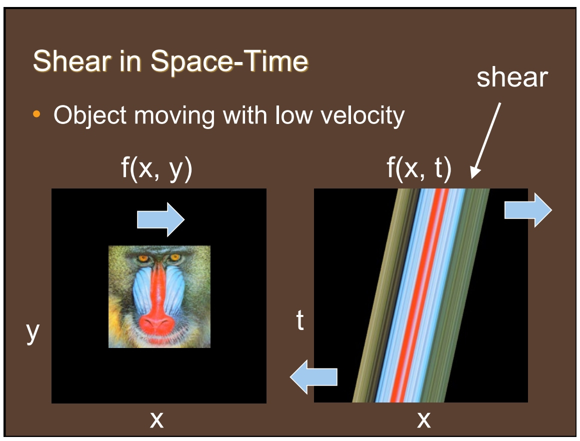
Shear in Space-Time
• Object moving with low velocity
shear
f(x, y)
f(x, t)
y
x
t
x

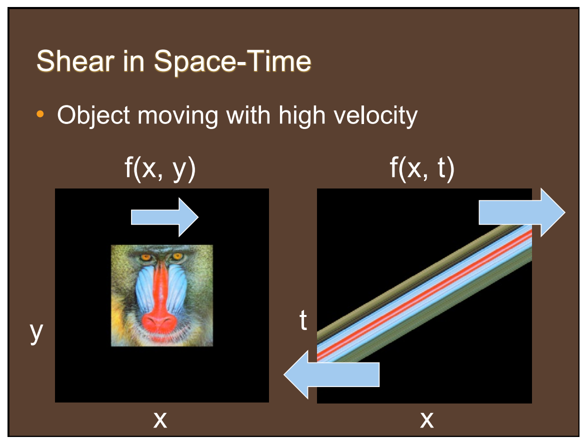
Shear in Space-Time
• Object moving with high velocity
f(x, y)
f(x, t)
y
x
t
x

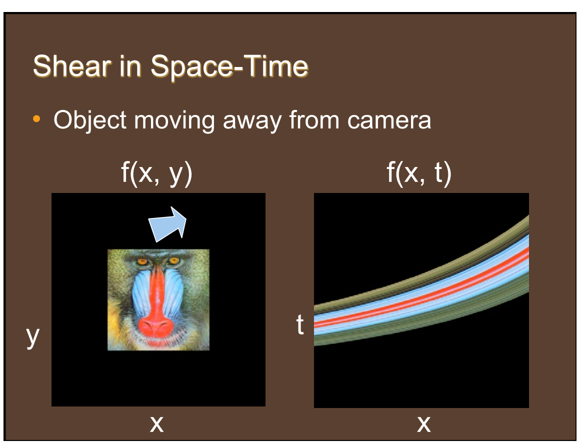
Shear in Space-Time
• Object moving away from camera
f(x, y)
f(x, t)
y
x
t
x

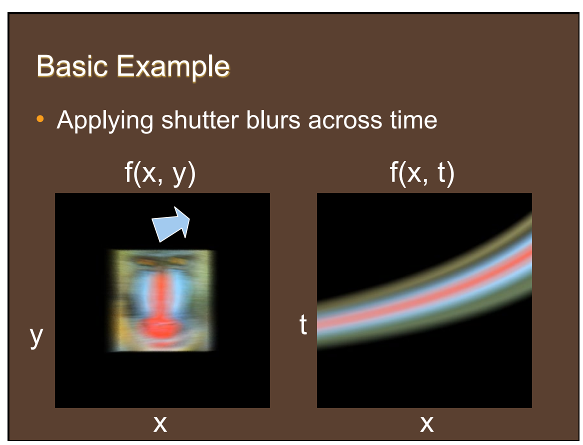
Basic Example
• Applying shutter blurs across time
f(x, y)
f(x, t)
y
x
t
x

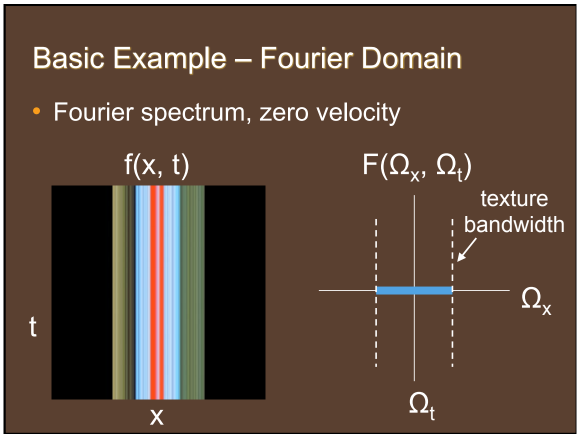
Basic Example – Fourier Domain
Fourier spectrum, zero velocity
f(x, t)
F(Ωx, Ωt)
texture bandwidth
t
x
Ωx
Ωt

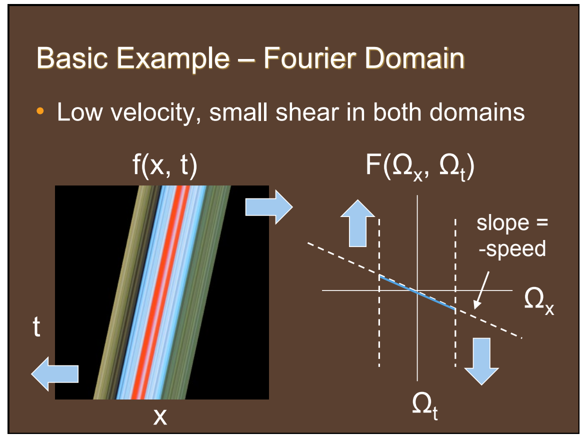
Basic Example – Fourier Domain
Low velocity, small shear in both domains
f(x, t)
F(Ωx, Ωt)
slope = -speed
t
x
Ωx
Ωt

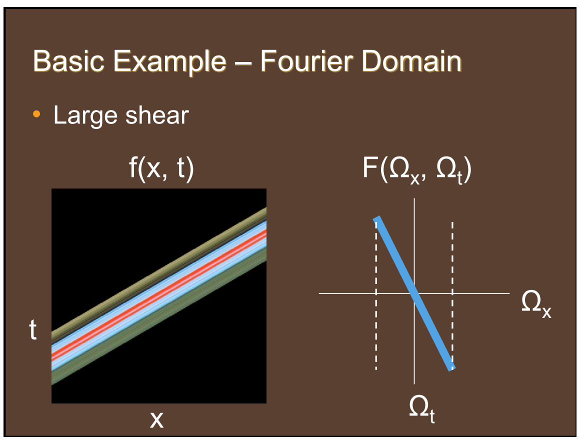
Basic Example – Fourier Domain
Large shear
f(x, t)
F(Ωx, Ωt)
t
x
Ωx
Ωt

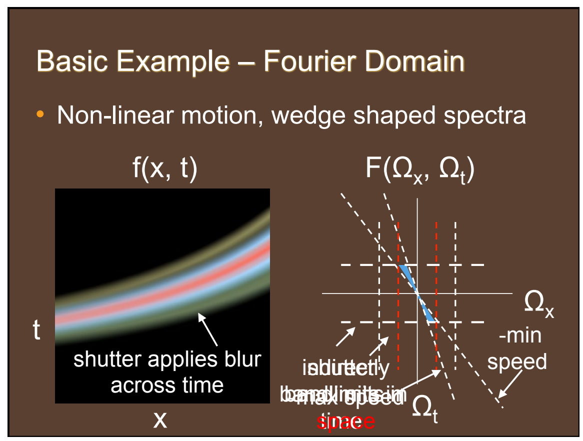
Basic Example – Fourier Domain
Non-linear motion, wedge shaped spectra
f(x, t)
F(Ωx, Ωt)
shutter applies blur across time
-min speed
t
x
Ωx
Ωt

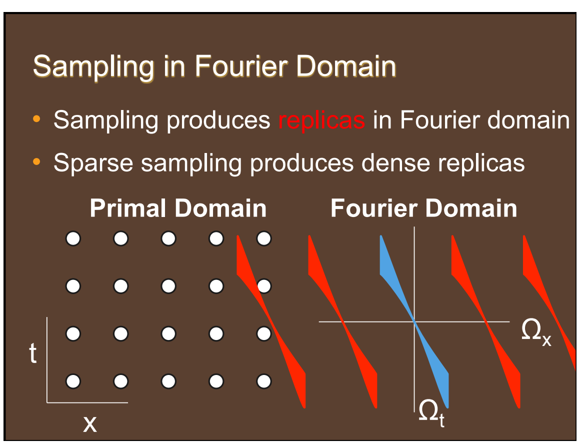
Sampling in Fourier Domain
Sampling produces replicas in Fourier domain
Sparse sampling produces dense replicas
Primal Domain
Fourier Domain
t
x
Ωx
Ωt

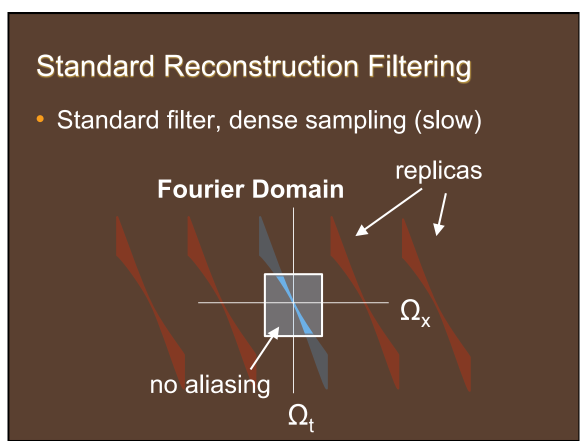
Standard Reconstruction Filtering
Standard filter, dense sampling (slow)
Fourier Domain
replicas
no aliasing
Ωx
Ωt

Standard Reconstruction Filter

- Standard filter, sparse sampling (fast)

Fourier Domain

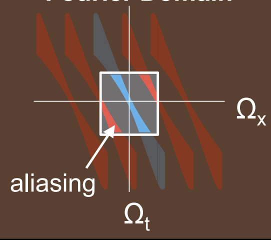

Sheared Reconstruction Filter

- Our sheared filter, sparse sampling (fast)

Fourier Domain

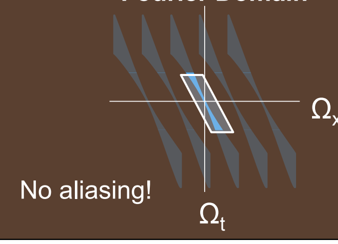

Sheared Reconstruction Filter

- Compact shape in Fourier = wide in primal

Primal Domain

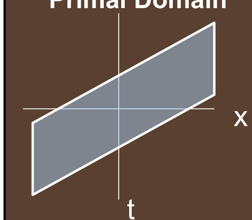

Fourier Domain

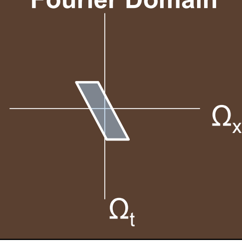

Car Scene

Our Method,
4 samples per pixel

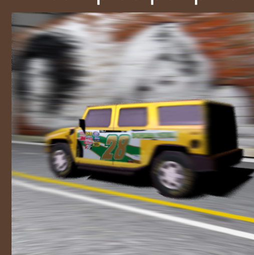

Stratified Sampling
4 samples per pixel

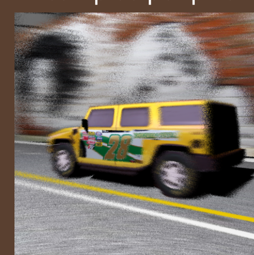

Teapot Scene

Our Method
8 samples / pix

motion blurred reflection

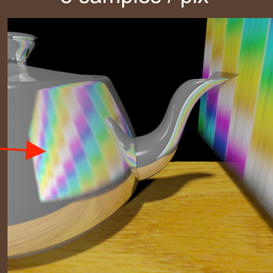

Ballerina Video