


# Real-Time High Quality Rendering

GAMES202, Lingqi Yan, UC Santa Barbara

## Lecture 12: Real-Time Ray Tracing 1



# Announcements

- GAMES101 resubmission
  - Starting next Monday!
- GAMES202 homework 3
  - Will be released soon hopefully
  - Your understanding is greatly appreciated

# Last Lectures

- Real-Time Physically-Based Materials
  - Microfacet BRDF, NDF, shadowing-masking
  - Kulla-Conty Approximation for multiple bounces
  - Disney principled BRDF
- Shading with microfacet BRDFs under polygonal lighting
  - Linearly Transformed Cosines (LTC)
- Non-photorealistic rendering (NPR)

# Some Arrangements

- Volumetric / scattering materials will not be covered in this course
  - Too many dependencies (RTE, BSSRDF, single/multiple scattering, etc.)
  - Will be fully covered in offline rendering, together with RTR techniques (delta tracking, dual scattering, layered materials, etc.)



[Final Fantasy VII Remake]



[Black Myth: Wukong]

# Some Arrangements

- Unreal Engine 5 early access is available now!
  - Again, both Nanite and Lumen are **TECHNICAL** breakthroughs
  - The underlying science is already understandable after learning this course
  - Will briefly analyze (or rather, guess) possible approaches in the last lecture



[UE5 Early Access Trailer  
(weakest boss ever)]

# Today

- New Topic: Real-Time Ray Tracing (RTRT)
  - Basic idea
  - Motion vector
  - Temporal accumulation / filtering
  - Failure cases
- Filtering techniques and implementation (next lecture)
  - Joint bilateral filtering
  - Spatiotemporal Variance-Guided Filtering (SVGF)

# RTRT is the Future


- In the real-time industry, people claim that

“Ray tracing is the future  
**and ever will be.**”

— The real-time industry

# RTRT is Happening

- In 2018, NVIDIA announced GeForce RTX series (Turing architecture)
  - Opening a \$250 billion market





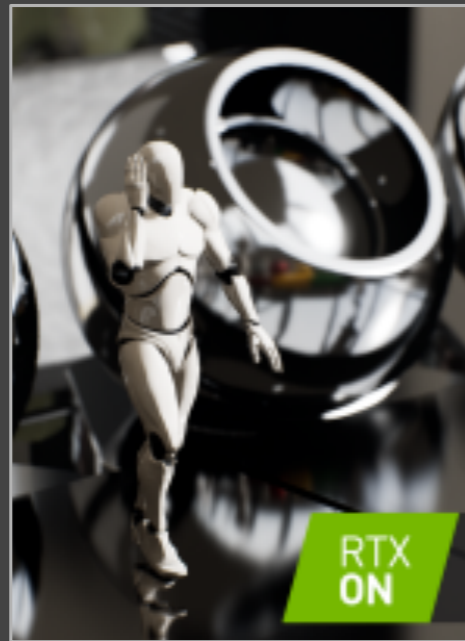
# RTRT is Happening

- What does RTX do?

Impressive demos of RTRT



Star Wars  
Reflections




RTX Demo



Porsche 70 Trailer



SOL




Rosewood Bangkok

# RTRT is Happening

- What does RTX actually do?

Advanced **ray traced** effects



Shadows



Reflections & Specular



Ambient Occlusion



Global Illumination


# RTRT is Happening

- What does RTX really do?

10 Giga rays per second == **1 sample per pixel**  
(for real time applications)

# RTRT is Happening

- What does RTX actually do?
- 1 SPP path tracing =
  - 1 rasterization (primary) +
  - 1 ray (primary visibility) +
  - 1 ray (secondary bounce) +
  - 1 ray (secondary vis.)



# RTRT is Happening

- 1 SPP = Extremely noisy results
- Key technology
  - **Denoising**



Fun image on Twitter

# State of the Art\* Denoising Solution



# Before we proceed...

- Goals (**with 1 SPP**)
  - Quality (no overblur, no artifacts, keep all details...)
  - Speed (< 2 ms to denoise one frame)
- **Mission impossible**
  - Sheared filtering series (SF, AAF, FSF, MAAF, ...)
  - Other offline filtering methods (IPP, BM3D, APR, ...)
  - Deep learning series (CNN, Autoencoder, ...)

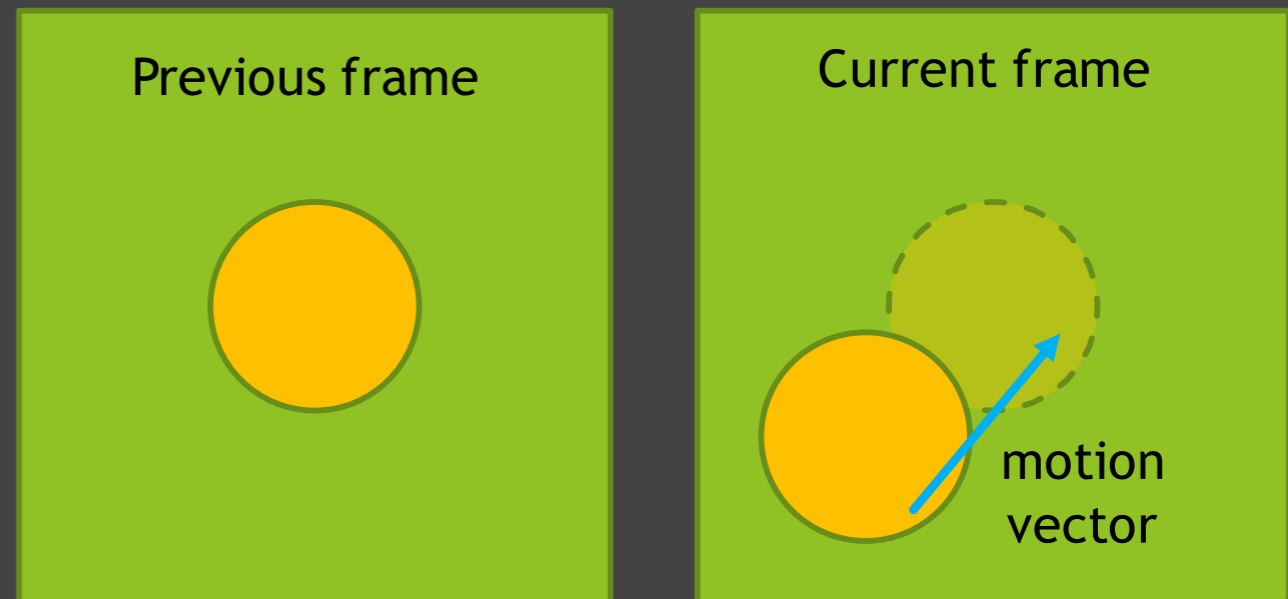
# Industrial Solution

- 3 most important ideas

- Temporal!
- **Temporal!!**
- **Temporal!!!**

- Key idea

- Suppose the previous frame is denoised and reuse it
- Use **motion vectors** to find previous locations
- Essentially increased SPP
- Spatial?






# The G-Buffers

- Geometry buffer
  - The auxiliary information acquired **FOR FREE**\* during rendering
  - Usually, per pixel depth, normal, world coordinate, etc.
  - Therefore, only **screen space** info



Direct illumination




Normal




Albedo

# Back Projection

- Pixel  $x$  in the current frame  $i$ 
  - ~~Where was **it** in the last frame  $i - 1$ ?~~
  - What pixel in frame  $i - 1$  contains **the same place/point that you see through pixel  $x$  in frame  $i$ ?**



frame  $i-1$




frame  $i$

# Back Projection

- Pixel  $x$  in the current frame  $i$ 
  - Where was it in the last frame  $i - 1$ ?
- Back projection
  - If world coord  $s$  is available as a G-buffer, just take it
  - Otherwise,  $s = M^{-1}V^{-1}P^{-1}E^{-1}x$  (still require z value)
  - Motion is known:  $s' \xrightarrow{T} s$ , thus  $s' = T^{-1}s$
  - Project world coord in frame  $i - 1$  to its screen:  
$$x' = P'V'M's'$$

# Temporal Accum./Denoising

- Let's denote:
  - $\sim$  : unfiltered
  - $-$  : filtered



- This frame (i-th frame)

$$\bar{C}^{(i)} = \text{Filter}[\tilde{C}^{(i)}]$$

$$\bar{C}^{(i)} = \alpha \bar{C}^{(i)} + (1 - \alpha) C^{(i-1)}$$

80%-90% contributions  
from last frame(s)!

$$\alpha = 0.1 - 0.2$$

17 FPS ( 59.03 ms ), 1spp

Video Capture

H.264 Codec

60 Video FPS

Codec Options

Capture UI (?)

Use Time-Range

Time Range

0.000 Start Time

23.000 End Time

Start Recording Cancel

Global Controls

0.000 Time

1.000 Time Scale

Reset Play Stop

Screen Capture Video Capture

Load scene Save scene Load model

Load filter stack Save filter stack

Renderer

Path Tracer Renderer

1 Subsample

Increase Indirect Roughness

0.050 MIS thresho

2 Bounces

Rasterize primary rays

1 Spp

Pixel subsampling

Direct lighting

16 Sample loop

Visualize buffers

Accumulate samples

2.120 Exposure

Pixel peep

FilterStack

Enable Reset

[0] Remove Outliers

Remove Outliers Filter

Cut Debug Output

Enable

Clamping Method


Replace outlier with Median

7 Filter Si

FilterStack/Filter\_0Average

FilterStack/Filter\_0Clamping

[1] DAF



# 1spp Ray Traced Global Illumination

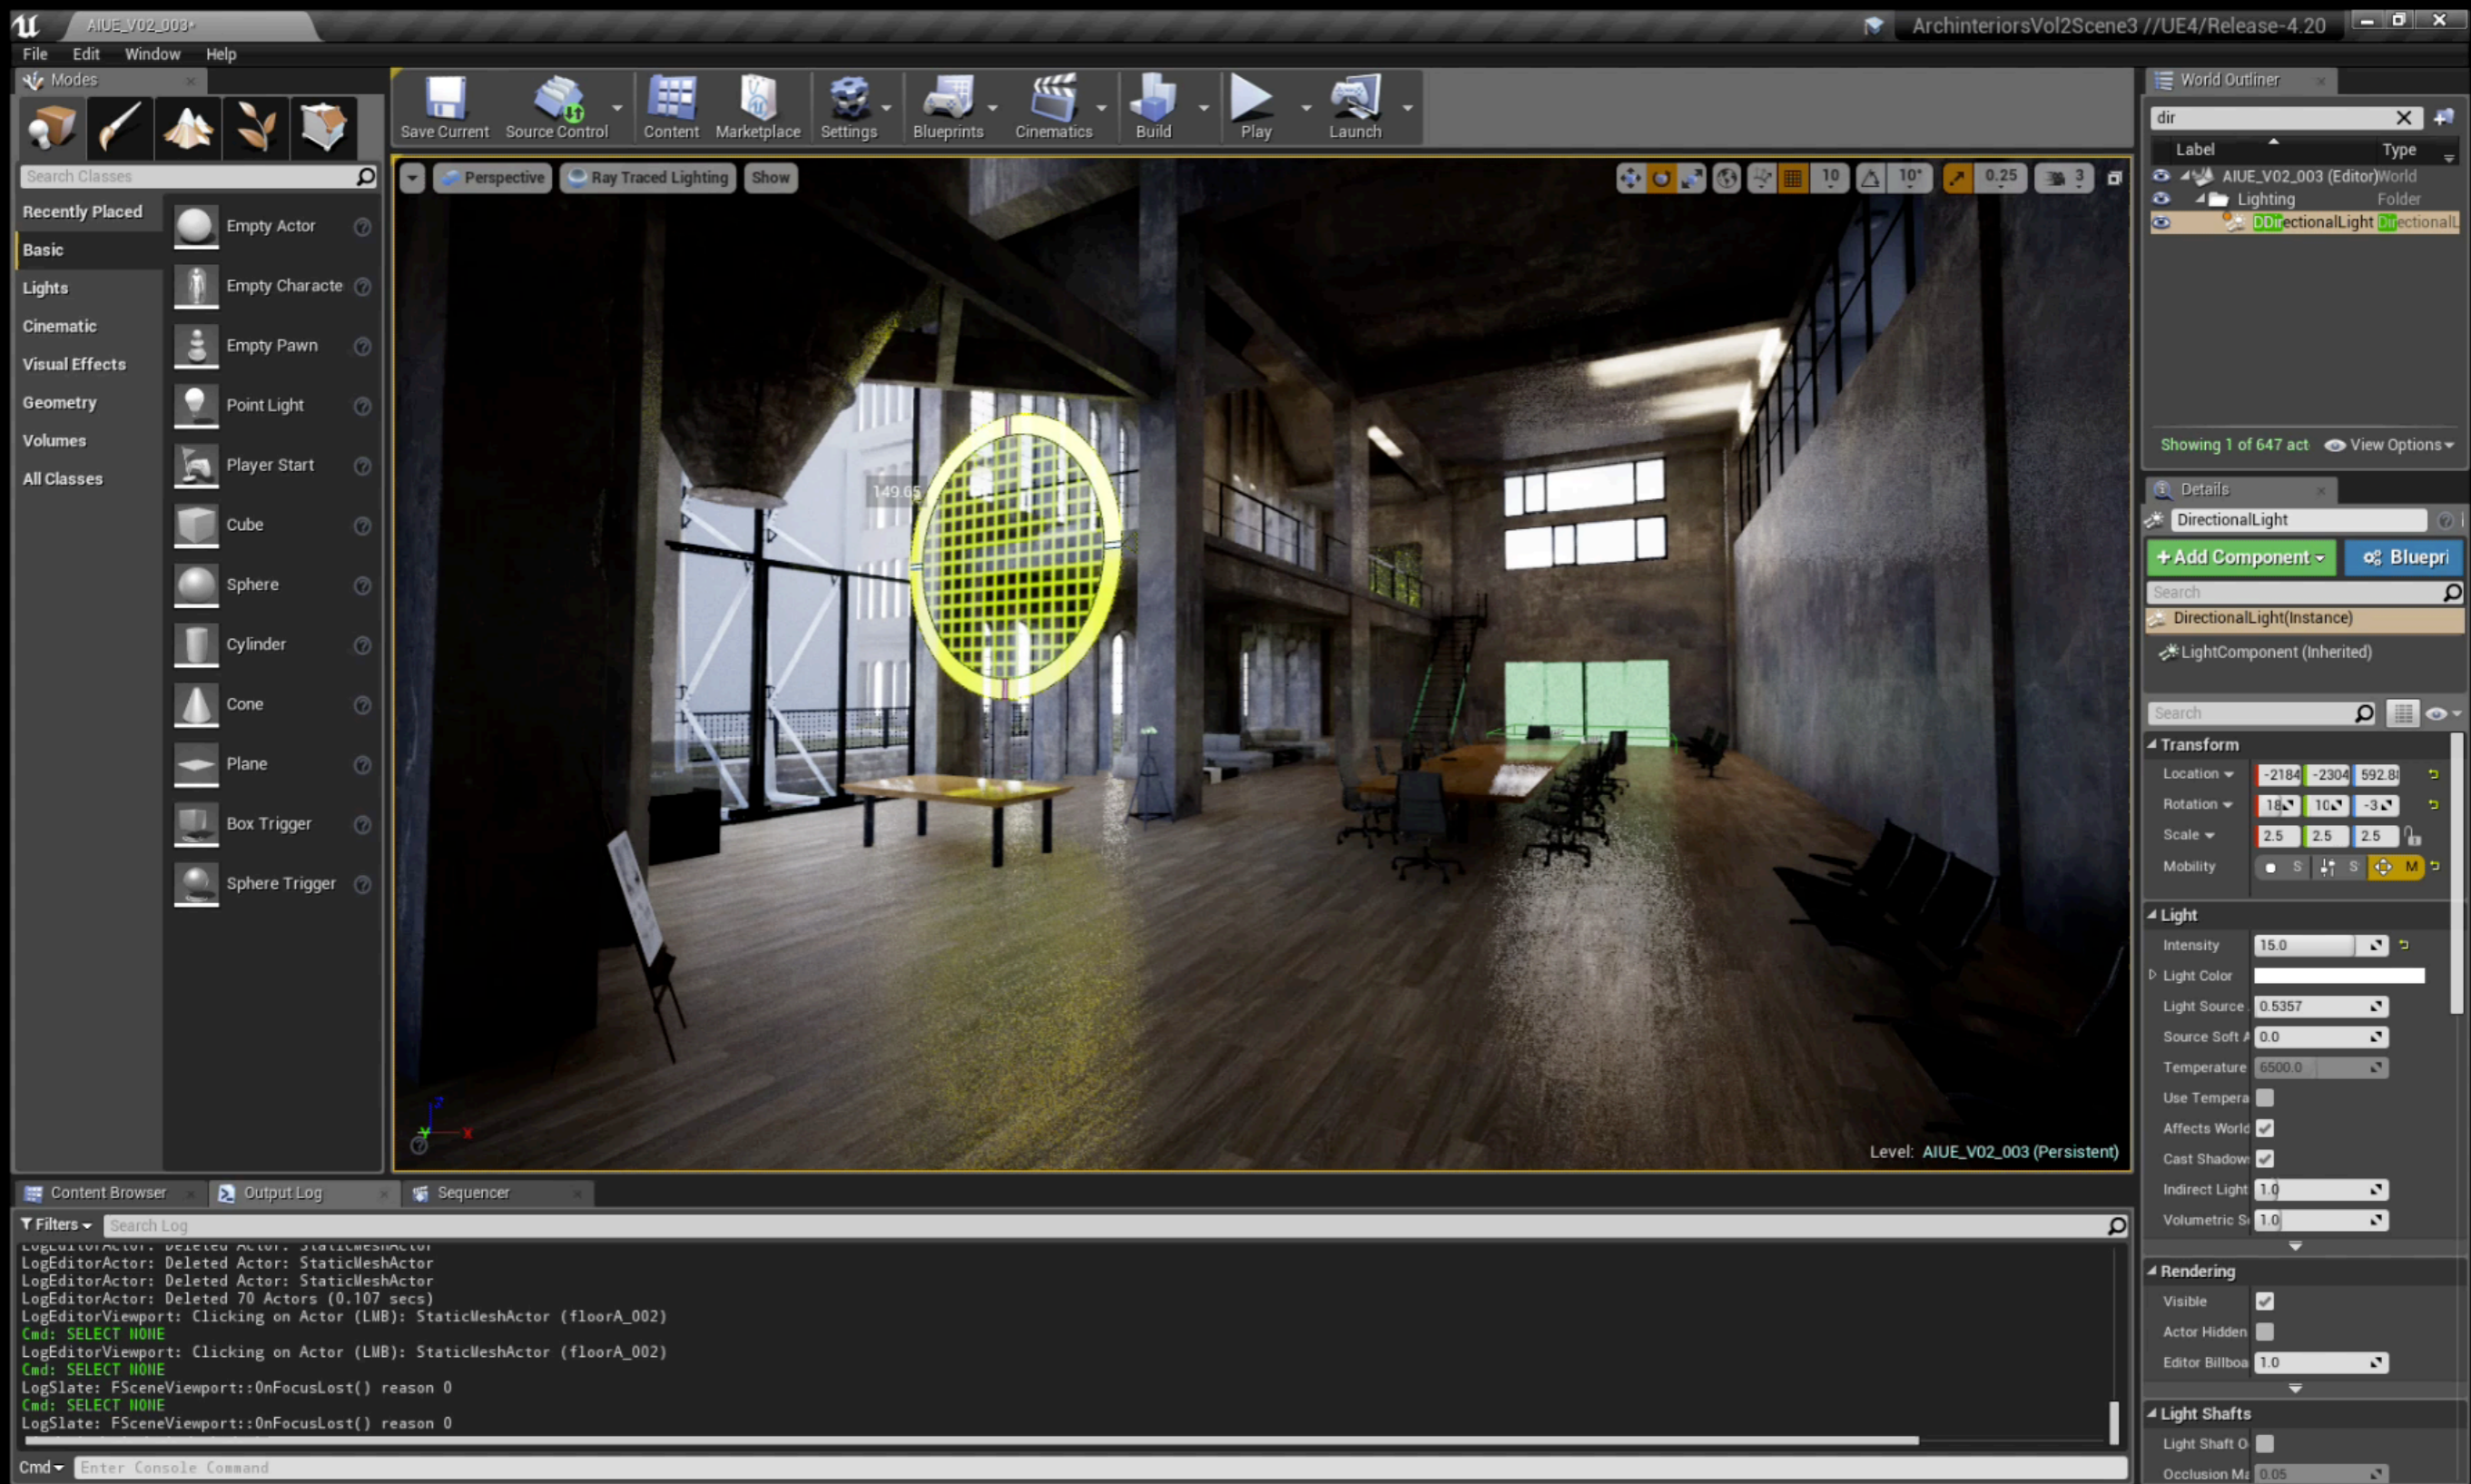


# 1spp Ray Traced Global Illumination + Denoising



# Ground Truth





# Temporal Failure

- Temporal info is not always available
  - Failure case 1: switching scenes  
**(burn-in period)**



[Monster Hunter Rise]

# Temporal Failure

- Temporal info is not always available
  - Failure case 1: switching scenes
  - Failure case 2: walking backwards in a hallway  
**(screen space issue)**

[Resident Evil Movie]



# Temporal Failure

- Temporal info is not always available
  - Failure case 1: switching scenes
  - Failure case 2: walking backwards in a hallway
  - Failure case 3: suddenly appearing background **(disocclusion)**



# Ignoring Temporal Failure?

- We can still blindly use temporal information
  - Of course, this is incorrect
  - But what kind of artifact will it bring?

- **Lagging!**



10x



# Traditional motion vectors (no clamp)

**Shadow**  
**Fence**  
Moving Objects

**Shadow**  
**Pink Room**  
Moving Light  
Changing Light Sizes

**Shadow** **Shadow**  
**Apples** **Fence**  
Curved Surfaces  
Multiple Lights

**Glossy**  
**Sun Temple**  
Moving Camera

**Glossy**  
**Restaurant**  
Moving Camera


**Glossy**  
**Restaurant**  
Curved Surfaces

**Occlusion**  
**PICA**  
Moving Objects

**Occlusion**  
**PICA 2**  
Moving Objects

# Adjustments to Temp. Failure

- Clamping  $\bar{C}^{(i)} = \alpha \bar{C}^{(i)} + (1 - \alpha) C^{(i-1)}$ 
  - Clamp previous toward current
- Detection
  - Use e.g. object ID to detect temporal failure
  - Tune  $\alpha$ , binary or continuously
  - Possibly strengthen / enlarge spatial filtering
- Problem: **re-introducing noise!**



10x

# Traditional motion vectors (clamp)

**Shadow**  
**Fence**  
Moving Objects

**Shadow**  
**Pink Room**  
Moving Light  
Changing Light Sizes

**Shadow**  
**Apples**  
Curved Surfaces

**Shadow**  
**Fence**  
Multiple Lights

**Glossy**  
**Sun Temple**  
Moving Camera

**Glossy**  
**Restaurant**  
Moving Camera

**Glossy**  
**Restaurant**  
Curved Surfaces


**Occlusion**  
**PICA**  
Moving Objects


**Occlusion**  
**PICA 2**  
Moving Objects



# More Temporal Failure

- Temporal failure can also happen in shading
  - Consider the “fence” scene with a moving light behind
  - What’s the motion vector of the **shadows**?





# Ground Truth

**Shadow**  
Fence  
Moving Objects

**Shadow**  
Pink Room  
Moving Light  
Changing Light Sizes

**Shadow** | **Shadow**  
Apples | Fence  
Curved | Multiple  
Surfaces | Lights

**Glossy**  
Sun Temple  
Moving Camera

**Glossy**  
Restaurant  
Moving Camera

**Glossy**  
Restaurant  
Curved  
Surfaces


**Occlusion**  
PICA  
Moving Objects

**Occlusion**  
PICA 2  
Moving Objects

# More Temporal Failure

- Temporal failure can also happen in shading
  - Consider the moving chairs
  - What's the motion vector of the **glossy reflected images**?





# Some Side Notes

- The temporal accumulation is inspired by Temporal Anti-Aliasing (TAA)

- They are very similar
- Temporal reuse essentially increases the sampling rate

- Is there any research on further alleviating temporal failure?

- Yes! Our Eurographics (EG) paper “Temporally Reliable Motion Vectors for Real-time Ray Tracing”

EUROGRAPHICS 2021 / N. Mitra and I. Viola (Guest Editors) Volume 40 (2021), Number 2

## Temporally Reliable Motion Vectors for Real-time Ray Tracing

Zheng Zeng<sup>1</sup>, Shiqin Liu<sup>2</sup>, Jinglei Yang<sup>3</sup>, Lu Wang<sup>1†</sup> and Ling-Qi Yan<sup>1,3</sup>

<sup>1</sup>School of Software, Shandong University, China  
<sup>2</sup>NVIDIA, USA  
<sup>3</sup>University of California, Santa Barbara, USA

Scene	Method	RMSE	SSIM
Shadows (Pink Room)	Ours	0.1708	0.8385
	No denoise	0.2146	0.8081
	No clamp	0.2208	0.8350
	SVGF	0.4887	0.4350
Glossy Reflections (Restaurant)	Ours	0.2212	0.8109
	No clamp	0.2310	0.8032
	Temporal Grad	0.2322	0.8101
	Soc. w/ rms	0.3321	0.5821
Occlusions (PICA)	Ours	0.1410	0.8972
	No clamp	0.1410	0.8974
	Temporal Grad	0.1410	0.8974
	SVGF	0.1819	0.8119

**Figure 1:** We propose temporally reliable motion vectors for shadows (left, pink room scene), glossy reflections (middle, contemporary restaurant scene) and occlusions (right, PICA scene) that better explore and utilize temporal correspondences between adjacent frames. Our method significantly reduces noise, overblur and ghosting artifacts compared to the state of the art temporal reuse methods with traditional motion vectors. The quantitative evaluation metrics (RMSE and SSIM) are shown below the images.

**Abstract**  
Real-time ray tracing (RTRT) is being pervasively applied. The key to RTRT is a reliable denoising scheme that reconstructs clean images from significantly undersampled noisy inputs, usually at 1 sample per pixel as limited by current hardware's computing power. The state of the art reconstruction methods all rely on temporal filtering to find correspondences of current pixels in the previous frame, described using per-pixel screen-space motion vectors. While these approaches are demonstrated powerful, they suffer from a common issue that the temporal information cannot be used when the motion vectors are not valid, i.e. when temporal correspondences are not obviously available or do not exist in theory. We introduce temporally reliable motion vectors that aim at deeper exploration of temporal coherence, especially for the generally-believed difficult applications on shadows, glossy reflections and occlusions, with the key idea to detect and track the cause of each effect. We show that our temporally reliable motion vectors produce significantly better temporal results on a variety of dynamic scenes when compared to the state of the art methods, but with negligible performance overhead.

**CCS Concepts**  
• Computing methodologies → Rendering; Ray tracing;

**1. Introduction**  
It is generally believed that ray tracing has become the gold standard in rendering to generate realistic images in a physically correct way. But ray tracing is historically slow, usually taking hours to days to produce a single noise-free image. For this reason, real-time ray tracing (RTRT) that performs at > 30 frames per second was considered impractical for a long time, until the recent breakthrough of graphics hardware brought it into life. Nowadays, real-time ray tracing has not only attracted enormous interest in the


† Corresponding authors: luwang\_hciv@sdau.edu.cn, lingqi@cs.ucsb.edu

© 2021 The Author(s)  
Computer Graphics Forum © 2021 The Eurographics Association and John Wiley & Sons Ltd. Published by John Wiley & Sons Ltd.

# Spatial Denoising (Next Lec.)

- This frame (i-th frame)

$$\bar{C}^{(i)} = Filter[\tilde{C}^{(i)}]$$



- How to filter the current frame?
  - Bilateral filter? ([https://en.wikipedia.org/wiki/Bilateral\\_filter](https://en.wikipedia.org/wiki/Bilateral_filter))
  - Cross / joint bilateral filter (and their variants)
    - Taking more info into account
    - G-buffers: normal / depth / object ID, etc.

# Next Lecture

- Real-Time Ray Tracing 2  
(filtering techniques and implementation)



[Spatiotemporal Variance-Guided Filtering]

**Thank you!**