



Real-Time Stochastic Lightcuts

Daqi Lin & Cem Yuksel
University of Utah

I3D 2020





Real-Time Stochastic Lightcuts

Daqi Lin & Cem Yuksel
University of Utah


ISD 2020

Many-Lights Problem



Real Lights


[Image Credit: [Smale & Hansford](#)]



Virtual Point Lights


[Image Credit: Miloš Hašan,
[Boyadzhiev & Matzen](#)]

Traditional Real-Time Methods



Light Map

[Image Credit: [Nenad Nesovic](#)]



Tiled-based Lighting


[Image Credit: [Christina Coffin](#)]

Monte Carlo Real-Time Methods




Dynamic Many-Light Sampling (ATS)

[Moreau et al. 2019]



ReSTIR

[Bitterli et al. 2020]

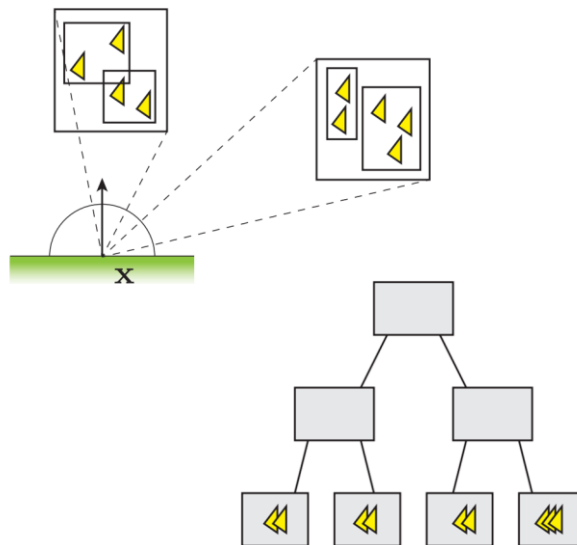



Monte Carlo Real-Time Methods



Dynamic Many-Light Sampling [Moreau et al. 2019]

Real-time version of Adaptive Tree Splitting (ATS)


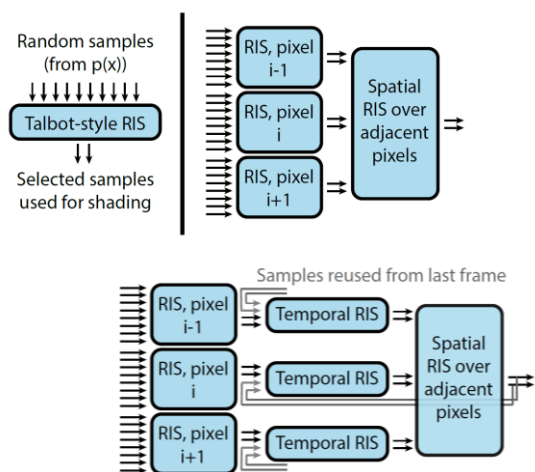


Monte Carlo Real-Time Methods



ReSTIR [Bitterli et al. 2020]


Reservoir-based Spatiotemporal Importance Resampling



Our method



ATS [Moreau 2019]



Real-Time Stochastic Lighthcuts

equal sampling time

Our method



- Extension of *stochastic lightcuts* [Yuksel 2019]
- Perfect binary tree
- New weight computation scheme
- Cut sharing



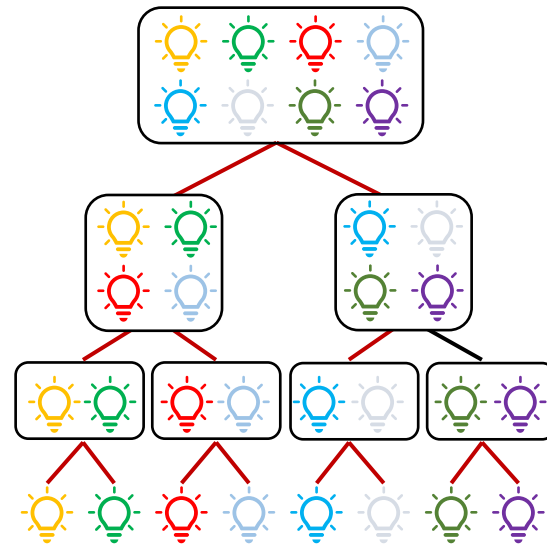

Related Work

Lightcuts [Walter et al. 2005]





- Adaptive direct illumination sampling [Vévoda and Křivánek 2016]
- Bayesian online regression for adaptive direct illumination sampling [Vévoda et al. 2018]
- Adaptive Tree Splitting [Estevez and Kulla 2018]
- **Stochastic Lightcuts** [Yuksel 2019]



Stochastic Lightcuts [Yuksel 2019]





Stochastic Lightcuts [Yuksel 2019]



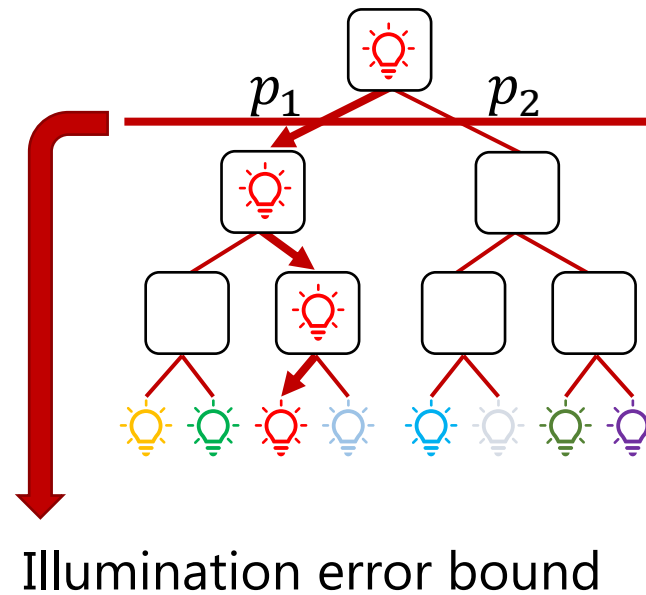

Stochastic Lightcuts [Yuksele 2019]





Stochastic Lightcuts [Yuksel 2019]



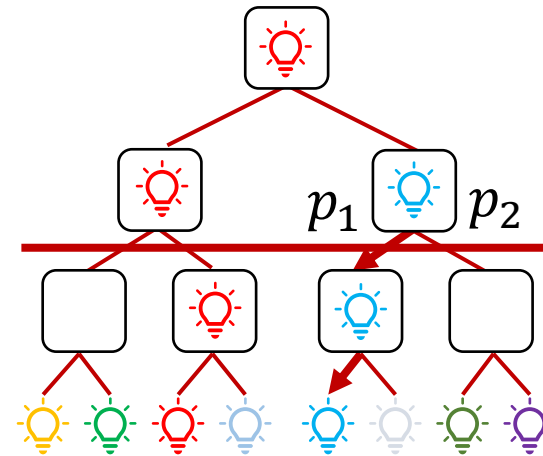

Stochastic Lightcuts [Yuksele 2019]





Stochastic Lightcuts [Yuksel 2019]





Stochastic Lightcuts [Yuksel 2019]





Stochastic Lightcuts [Yukse1 2019]



Stochastic Lightcuts [Yuksel 2019]



Stochastic Lightcuts [Yuksel 2019]



Stochastic Lightcuts [Yukse1 2019]



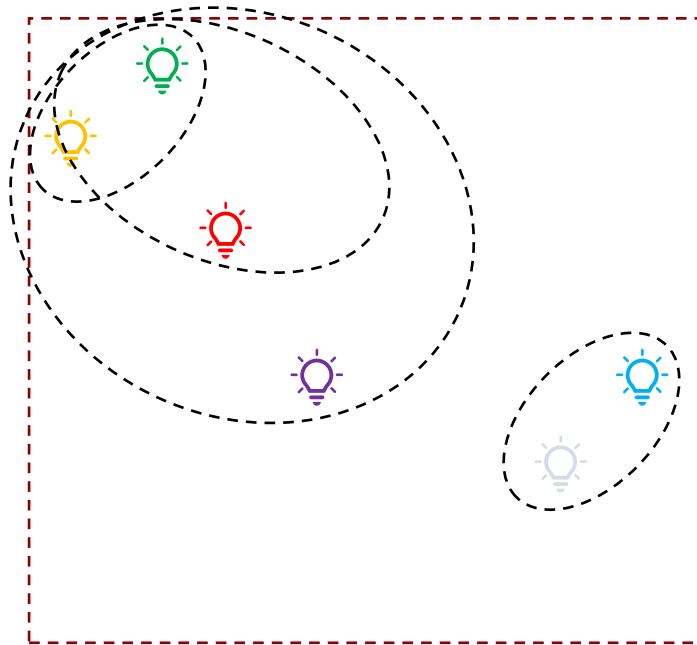
Benefits:

- Eliminates sampling correlation
- High sampling efficiency
- No restriction on light types


Stochastic Lightcuts [Yuksel 2019]



Problems:



Agglomerative clustering



Unbalanced tree

Stochastic Lightcuts [Yukse1 2019]



However, stochastic lightcuts allows using any tree





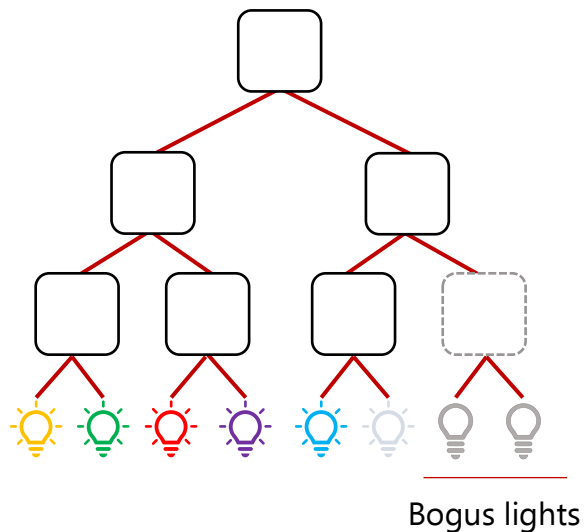

Real-Time Stochastic Lightcuts





Real-Time Stochastic Lightcuts



- Perfect Binary Tree





Real-Time Stochastic Lightcuts


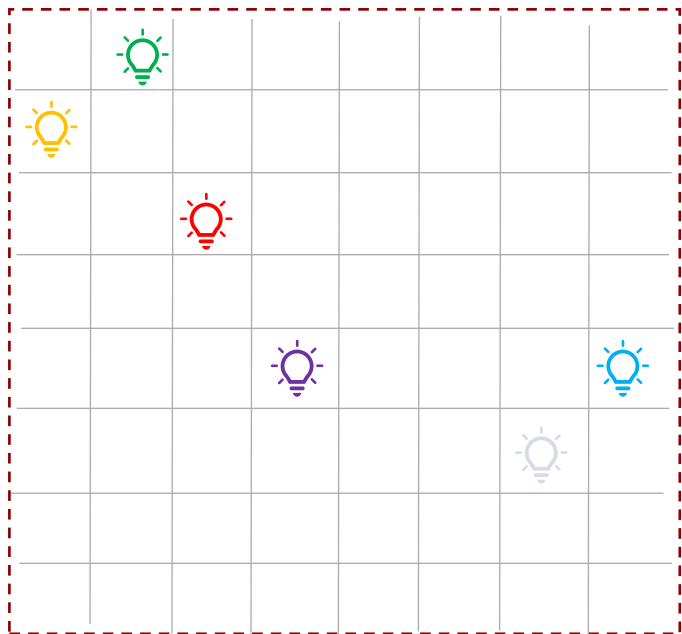
- Perfect Binary Tree





Real-Time Stochastic Lightcuts



- Perfect Binary Tree





Real-Time Stochastic Lightcuts



- Perfect Binary Tree





Real-Time Stochastic Lightcuts



- Perfect Binary Tree





Real-Time Stochastic Lightcuts


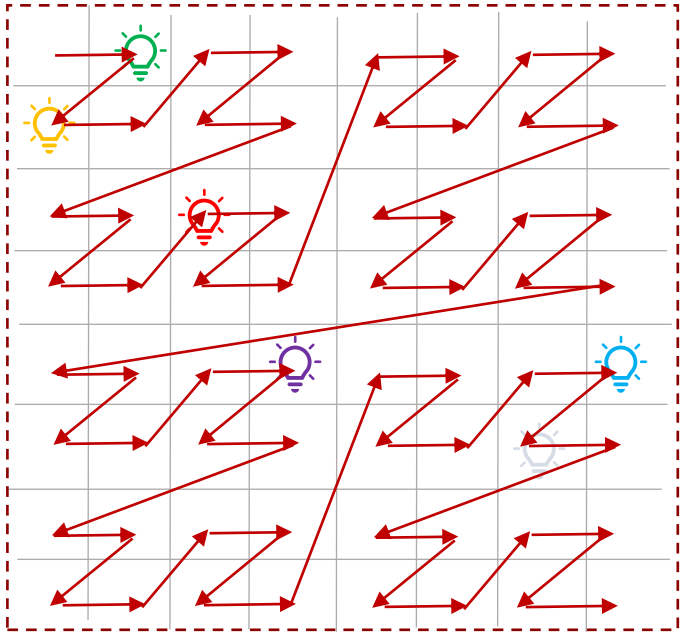
- Perfect Binary Tree





Real-Time Stochastic Lightcuts


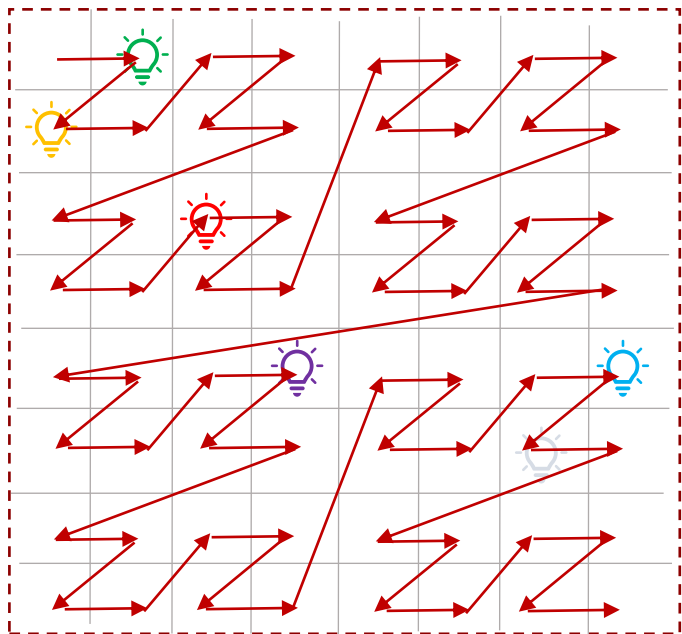
- Perfect Binary Tree





Real-Time Stochastic Lightcuts



- Perfect Binary Tree





Real-Time Stochastic Lightcuts

- Perfect Binary Tree



Real-Time Stochastic Lightcuts



- Perfect Binary Tree
 - Extremely fast to build
 - High sampling performance
(compensates for loss of tree quality)

Real-Time Stochastic Lightcuts





- Perfect Binary Tree
 - Extremely fast to build
 - High sampling performance (compensates for loss of tree quality)
 - Better computation/memory behavior
 - No child pointer required



Real-Time Stochastic Lightcuts

- Two-Level Perfect Tree




Each leaf is a perfectly binary tree




Real-Time Stochastic Lightcuts

- Light Tree Traversal




$$p_1 = \frac{w_1}{w_1 + w_2} \quad p_2 = \frac{w_2}{w_1 + w_2}$$





Real-Time Stochastic Lightcuts

- Light Tree Traversal



$$p_1 = \frac{w_1}{w_1 + w_2} \quad p_2 = \frac{w_2}{w_1 + w_2}$$

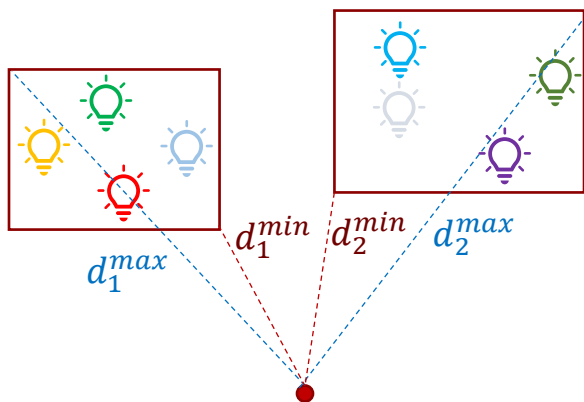
$$w_j = \frac{F_j(\mathbf{x}, \omega) \|\mathbf{I}_j\|}{(d_j(\mathbf{x}))^2}$$

distance to the light cluster



Real-Time Stochastic Lightcuts

- Our New Sampling Weight Formulation



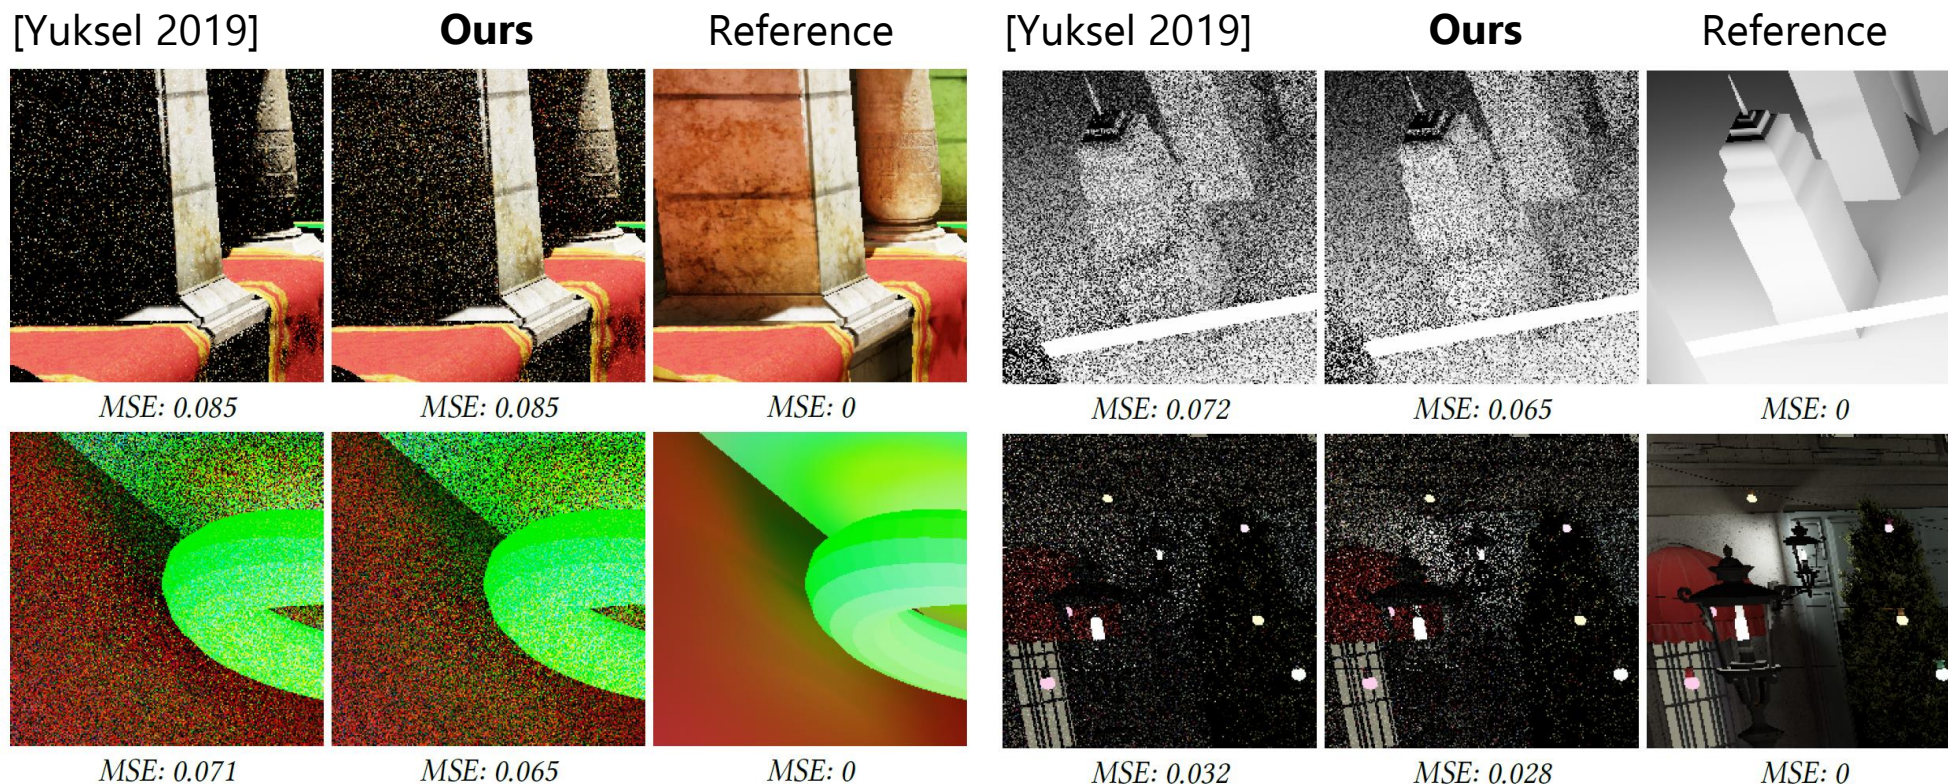
$$p_1 = \frac{w_1}{w_1 + w_2} \quad p_2 = \frac{w_2}{w_1 + w_2}$$

$$p_j = (p_j^{min} + p_j^{max})/2$$



Real-Time Stochastic Lightcuts



- Comparison of different weight formulations





Real-Time Stochastic Lightcuts

- Cut Sharing







Real-Time Stochastic Lightcuts

- Cut Sharing

Random representative pixel

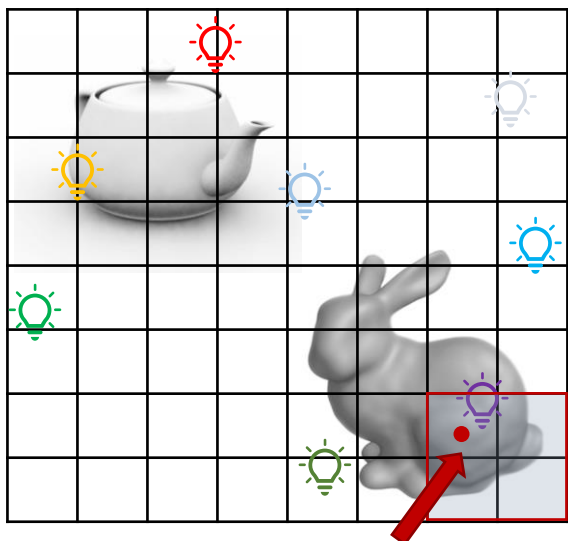


Cut shared by the 2x2 pixel block




Real-Time Stochastic Lightcuts

- Cut Sharing



Random representative pixel




Cut shared by the 2x2 pixel block




Real-Time Stochastic Lightcuts

- Cut share size vs performance




Real-Time Stochastic Lightcuts



Cut sharing block size $k = 8$

Interleaved Sampling





[Segovia et al. 2006b]



Real-Time Stochastic Lightcuts


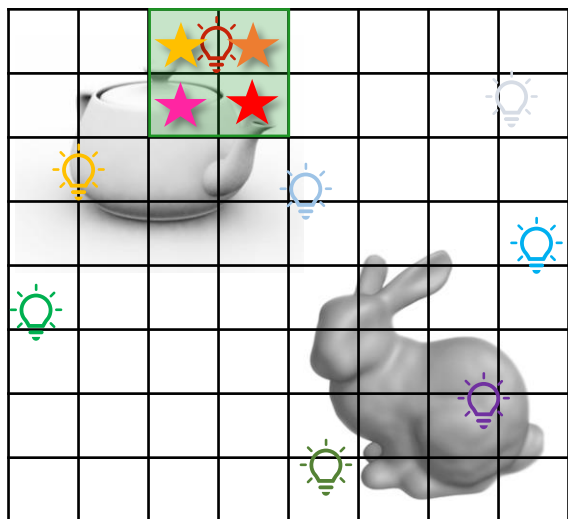
- No interleaved sampling (each pixel samples from all subtree roots)







Real-Time Stochastic Lightcuts

- Interleaved Sampling







8 light samples
18.7 ms



8 light samples (2x2 interleaved)
21.0 ms



8 light samples
18.7 ms





8 light samples (2x2 interleaved)
21.0 ms



Real-Time Stochastic Lightcuts


- Interleaved Sampling






Real-Time Stochastic Lightcuts

- Interleaved Sampling




Frame 1






Real-Time Stochastic Lightcuts

- Interleaved Sampling



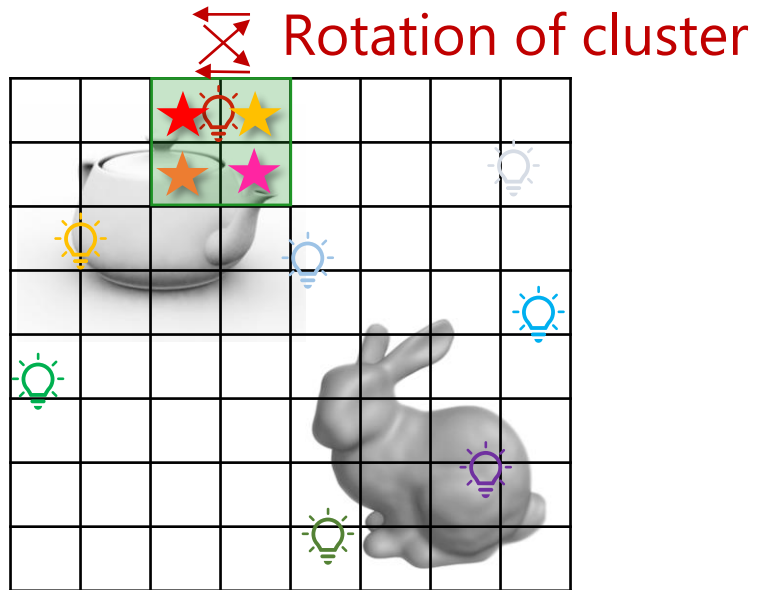
Frame 2



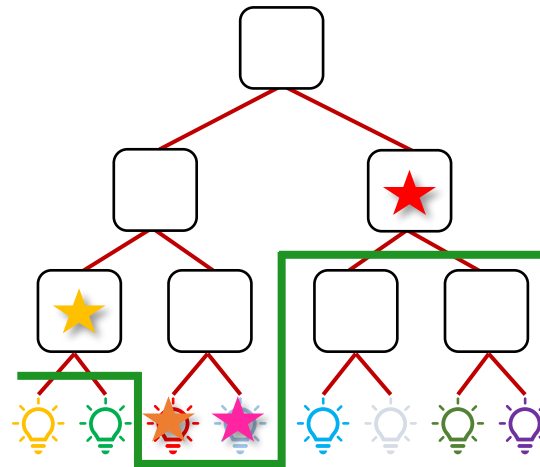


Real-Time Stochastic Lightcuts

- Interleaved Sampling



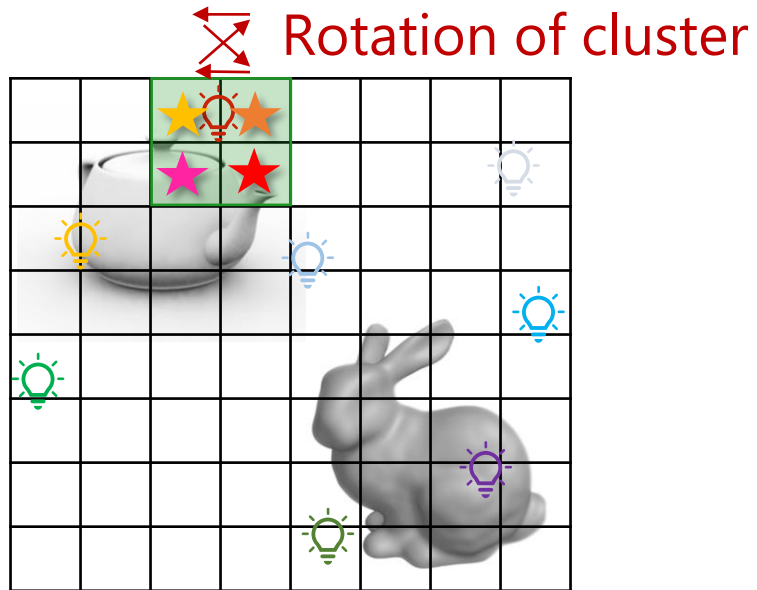
Frame 3






Real-Time Stochastic Lightcuts

- Interleaved Sampling



Frame 4



Real-Time Stochastic Lightcuts



- Perfect binary tree
- New weight computation scheme
- Cut sharing
 - Allows interleaved sampling





Results



Perfect Binary Tree Construction Time



~100,000 VPLs



12,146 emissive triangles

	<i>Crytek Sponza</i>		<i>Cornell Box</i>	
Geometry update	N/A	(0%)	0.01 ms	(6%)
Compute bounds	0.03 ms	(7%)	0.02 ms	(10%)
Morton code generation	0.01 ms	(3%)	0.01 ms	(4%)
Sorting	0.28 ms	(63%)	0.07 ms	(45%)
Building the leaf level	0.03 ms	(8%)	0.01 ms	(5%)
Building internal levels	0.08 ms	(19%)	0.05 ms	(29%)
Total Time	0.43 ms	(100%)	0.15 ms	(100%)

Compared to agglomerative clustering:

584x faster

147x faster

Perfect Binary Tree Construction Time



~100,000 VPLs



12,146 emissive triangles

	<i>Crytek Sponza</i>		<i>Cornell Box</i>	
Geometry update	N/A	(0%)	0.01 ms	(6%)
Compute bounds	0.03 ms	(7%)	0.02 ms	(10%)
Morton code generation	0.01 ms	(3%)	0.01 ms	(4%)
Sorting	0.28 ms	(63%)	0.07 ms	(45%)
Building the leaf level	0.03 ms	(8%)	0.01 ms	(5%)
Building internal levels	0.08 ms	(19%)	0.05 ms	(29%)
Total Time	0.43 ms	(100%)	0.15 ms	(100%)

Compared to agglomerative clustering:

584x faster

147x faster

ATS [Moreau et al. 2019]

Sampling Time: 21.9 ms


Sample Count: 5



Ours

Sampling Time: 22.1 ms


Sample Count: 11



ATS [Moreau et al. 2019]

Sampling Time: 48.1 ms


Sample Count: 11



Ours

Sampling Time: 22.1 ms

Sample Count: 11




ATS [Moreau et al. 2019]

Sampling Time: 48.1 ms

Sample Count: 11

Build Time: (CPU) 176.6 ms




MSE: 0.058

Ours

Sampling Time: 22.1 ms

Sample Count: 11

Build Time: (GPU) 0.4 ms



MSE: 0.064




ATS [Moreau et al. 2019]

Sampling Time: 18.4 ms

Sample Count: 4

Build Time: (CPU) 0.6 ms




Ours

Sampling Time: 18.1 ms

Sample Count: 13

Build Time: (GPU) 0.2 ms



ATS [Moreau et al. 2019]

Sampling Time: 18.4 ms

Sample Count: 4

Build Time: (CPU) 0.6 ms



ATS (unfiltered)

Ours

Sampling Time: 18.1 ms

Sample Count: 13

Build Time: (GPU) 0.2 ms



Ours (unfiltered)



ATS (filtered)



Ours (filtered)



Equal-Time Comparisons



Crytek Sponza (VPL)



Cornell Box




Arnold Buildings



Amazon Lumberyard Bistro





Random Sampling
12 light samples
MSE: 0.108


Unbalanced Tree + No Cuts
6 light samples
MSE: 0.098


Perfect Tree + No Cuts
9 light samples
MSE: 0.097


Unbalanced Tree + Lightcuts
7 light samples
MSE: 0.078

Perfect Tree + Lightcuts
17 light samples
MSE: 0.050

Reference







Random
Sampling
35 light samples
MSE: 0.046


Unbalanced Tree
+ No Cuts
12 light samples
MSE: 0.021

Perfect Tree
+ No Cuts
19 light samples
MSE: 0.033

Unbalanced Tree
+ Lightcuts
17 light samples
MSE: 0.011

Perfect Tree
+ Lightcuts
29 light samples
MSE: 0.019

Reference



Random
Sampling
13 light samples
MSE: 0.033

Unbalanced Tree
+ No Cuts
6 light samples
MSE: 0.025

Two-level Perfect
Tree + No Cuts
8 light samples
MSE: 0.028

Unbalanced Tree
+ Lightcuts
9 light samples
MSE: 0.023

**Two-level Perfect
Tree + Lightcuts**
16 light samples
MSE: 0.022

Reference

Conclusion



- Real-Time Stochastic Lightcuts
 - A perfect binary tree which is extremely fast to build
 - New weight computation scheme with lower error
 - Cut sharing technique to amortize cut computation cost

➔ More light samples within the same render time to achieve higher sampling quality

Future Work

- Importance sampling of light clusters
- Reduce sample evaluation cost





Paper, code and video available at:

<https://dqlin.xyz/pubs/2020-i3d-SLC/>



Real-Time Stochastic Lightcuts

Daqi Lin & Cem Yuksel

University of Utah

I3D 2020

