Efficient Visibility Reuse for Real-time ReSTIR

Yusuke Tokuyoshi

Advanced Rendering Research Group

Spatiotemporal Reservoir Resampling (ReSTIR)

- Powerful importance resampling technique
- Reuse samples in spatial neighbors and past frames ➡ many candidate samples ☺
- Resample according a target distribution ≈ integrand (i.e., lighting with shadow rays)

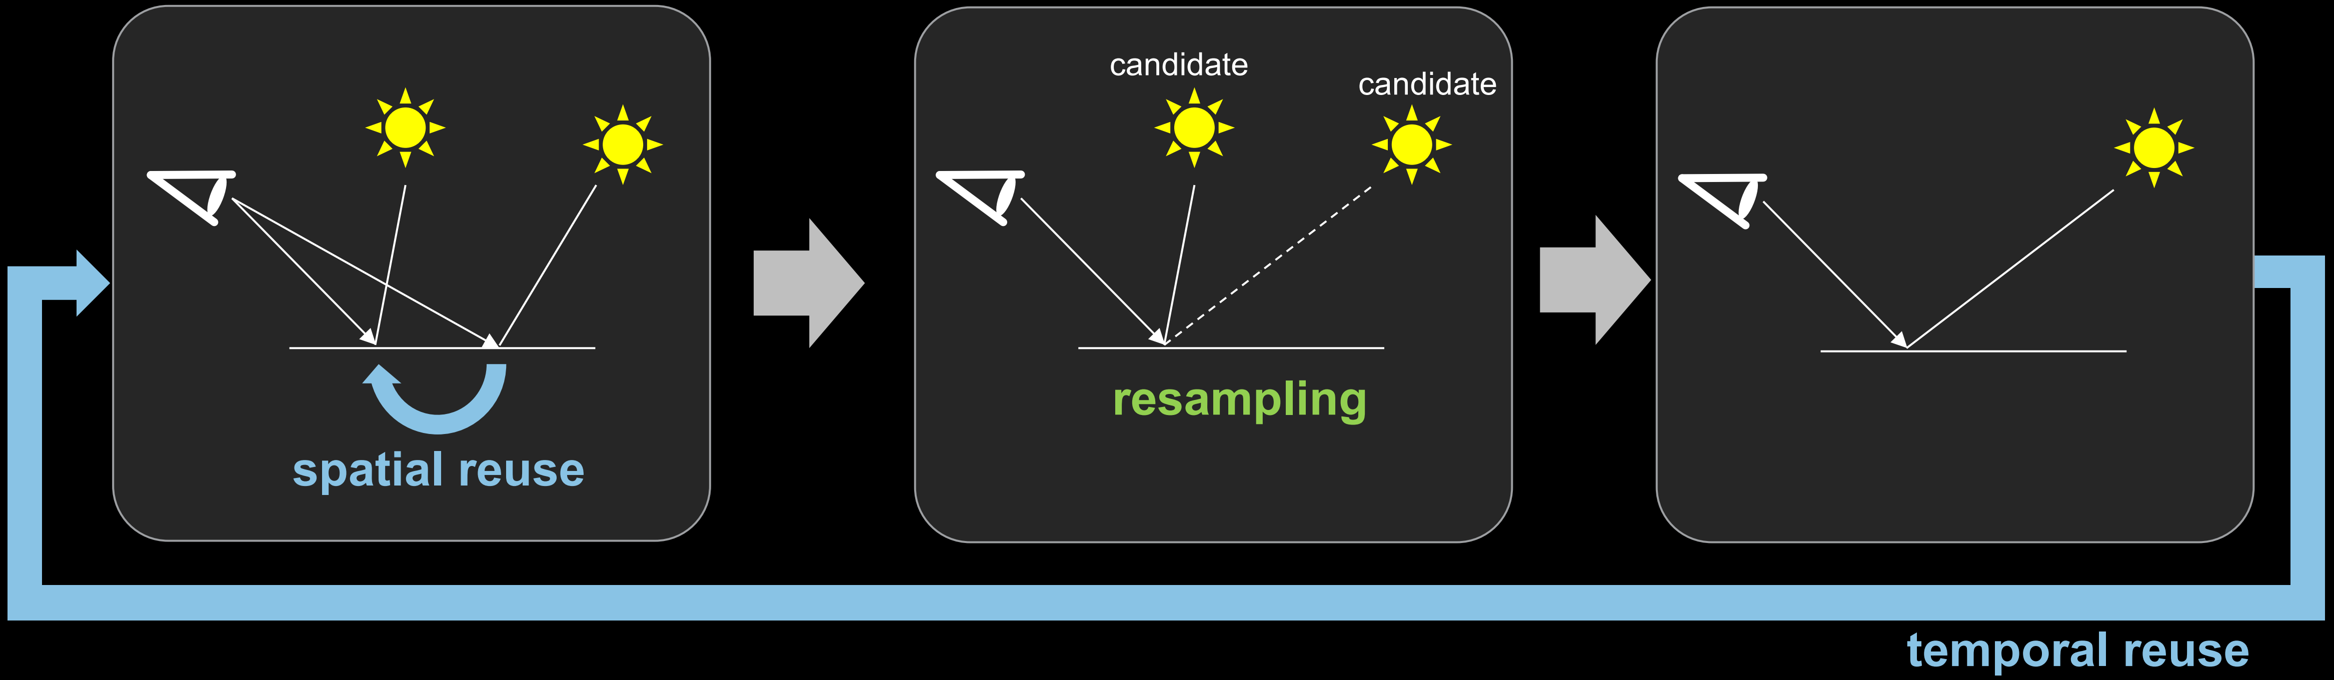

Shadow Noise in Real-time ReSTIR [Bitterli et al. 2020; Wyman and Panteleev 2021]

- Ignore shadow ray visibility for the target distribution
- Reduce the variance for the light contribution w/o shadow ray visibility ☺
- Inefficient to reduce shadow noise ☹

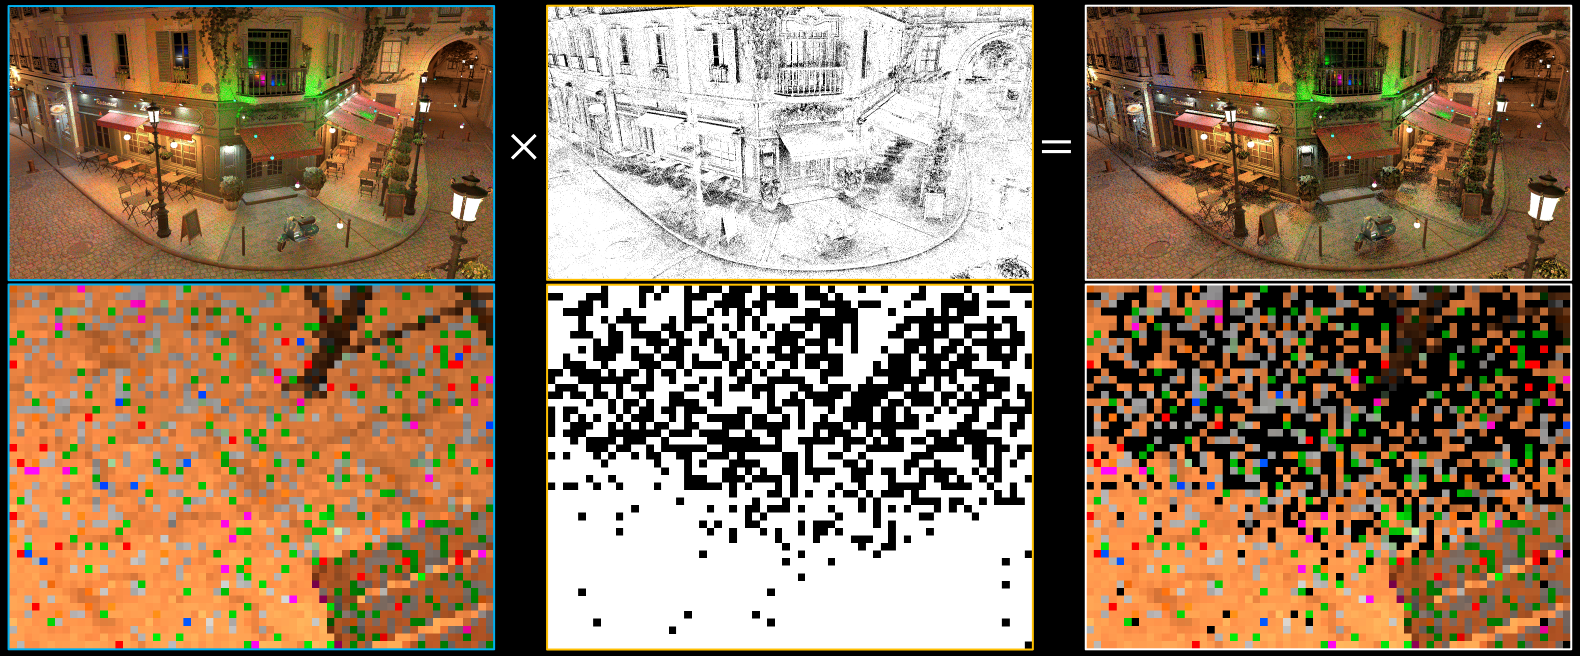

Light contribution w/o visibility $f(X)W_X$ × Visibility $V(X)$ = Result $f(X)V(X)\,W_X$

Reduce the shadow variance reusing visibilities

Previous Biased Visibility Reuse

- Reuse the initial sample visibility [Bitterli et al. 2020]
  - Trace a shadow ray for the initial sample and reuse it for resampling weights
  - Significantly reduce the variance inside shadows ☺
  - Bias = darkening + temporal delay
  - Ineffective around shadow edges ☹

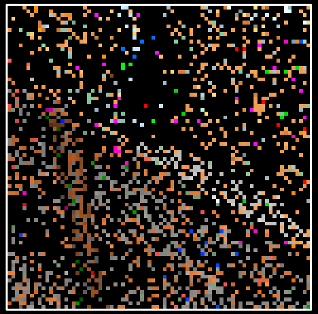
w/o vis. reuse

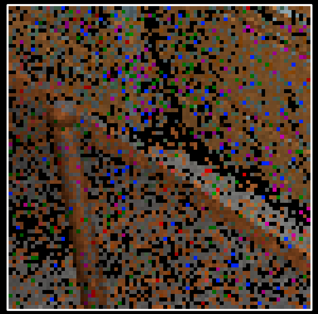
with vis. reuse

- Reuse the spatial sample visibility [Wyman and Panteleev 2021]
  - Trace a shadow ray for a (distant) spatial sample and reuse it for the integrand
  - Can control the shadow ray count (performance vs quality) ☺
  - Bias = temporal delay (+ shadow edge disappearance for small ray counts)
  - Does not reduce the shadow variance

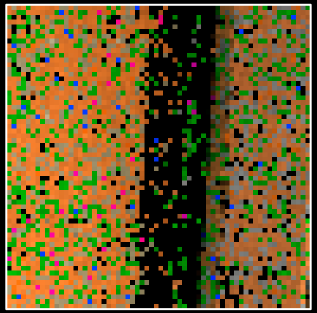
2 rays/pixel

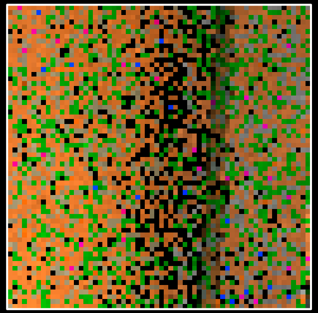
1.25 rays/pixel

Our Method

- Variance reduction method based on Wyman and Panteleev's visibility reuse [2021]
  - Easy to implement ☺
- Additional bias for typical real-time implementations
- Discuss conditions to obtain unbiased results

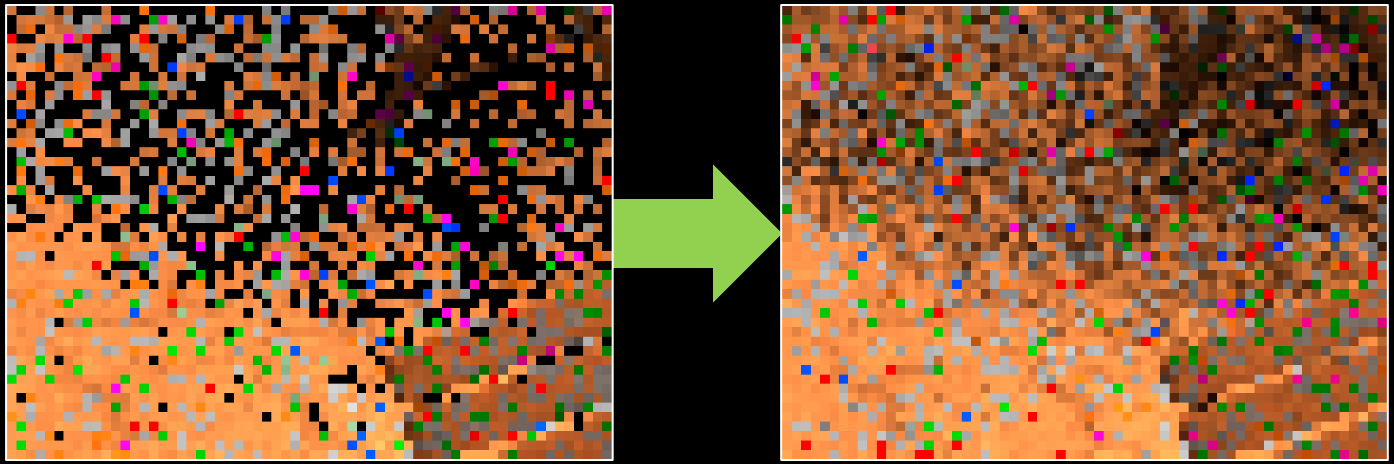

RMSPE: 80.3%

Wyman and Panteleev [2021]
(2 rays/pixel)

RMSPE: 59.6%

Our method
(2 rays/pixel)

Previous method


CPU 0.078 ms, GPU 8.869 ms, 60 Hz
Engine Tuning
+ Application/...
[X] Display Frame Rate
[-] Display Profiler
- Experiments/...
[X] Culling
Ours (2 rpp) ReSTIR
[-] Reference
+ Graphics/...
+ Renderer/...
+ Timing/...
Le Petit Coin
MENU
Our method

Our Method

Our Visibility Reuse

- Reuse spatiotemporal sample visibilities for the integrand as in Wyman and Panteleev [2021]

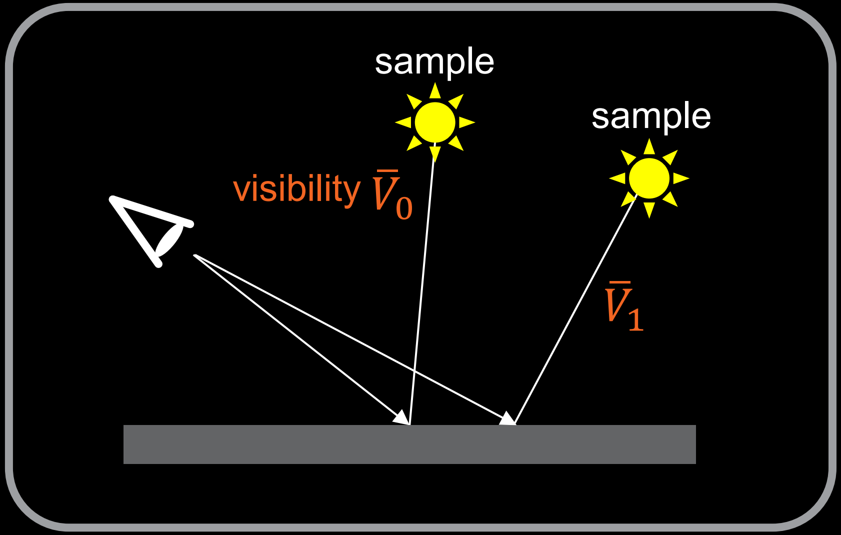

Our Visibility Reuse

- Reuse spatiotemporal sample visibilities for the integrand as in Wyman and Panteleev [2021]

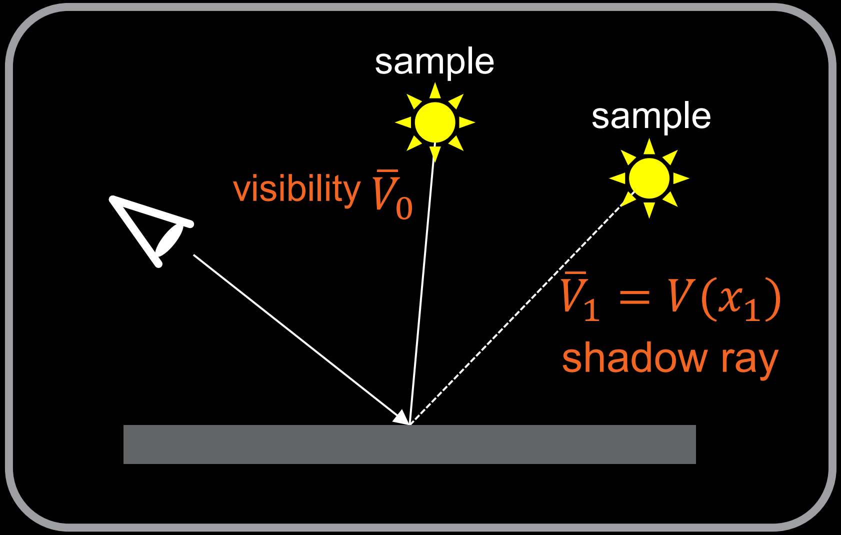

Our Visibility Reuse

- Reuse spatiotemporal sample visibilities for the integrand as in Wyman and Panteleev [2021]

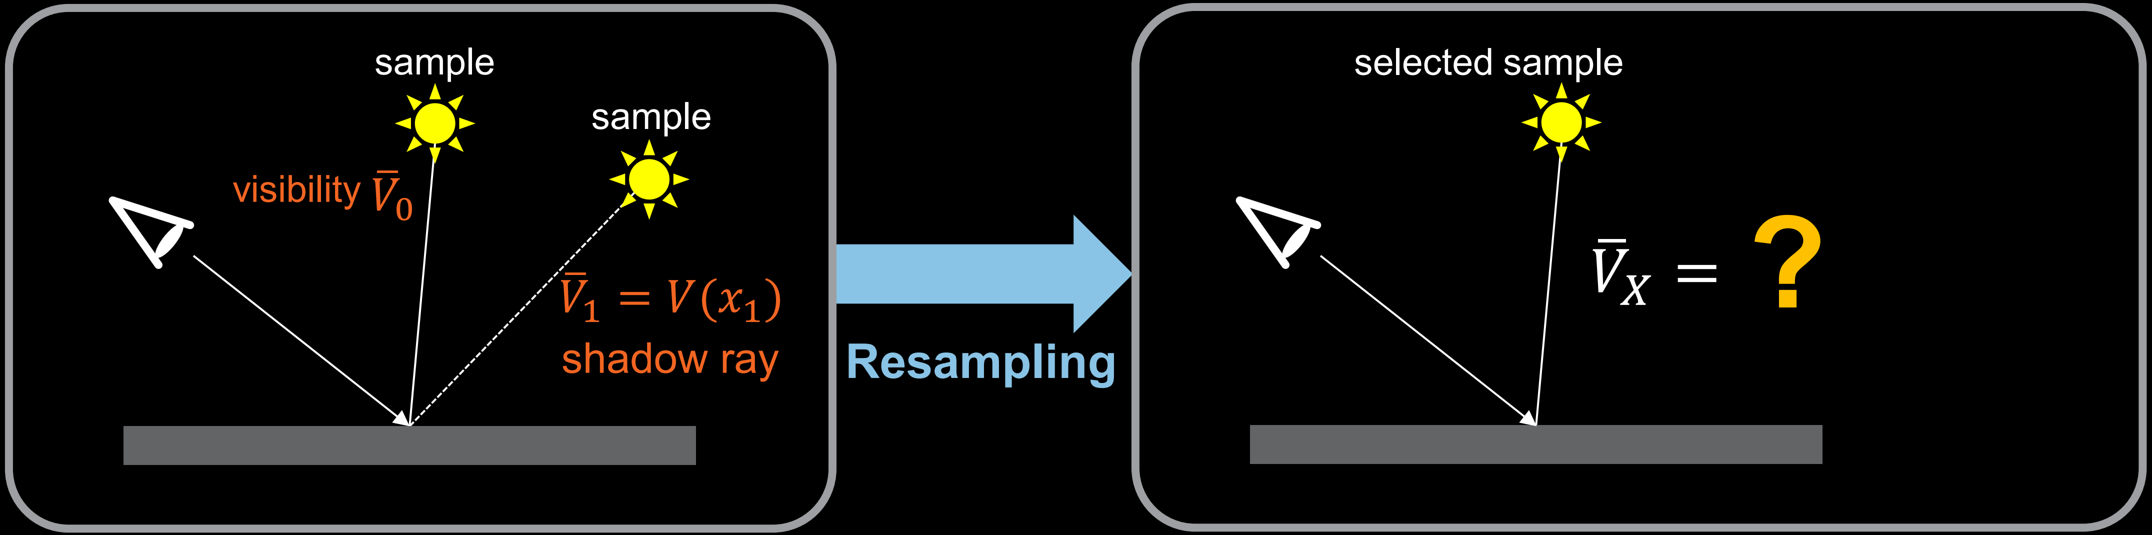

Our Visibility Reuse

- Reuse spatiotemporal sample visibilities for the integrand as in Wyman and Panteleev [2021]
- Estimate a new visibility using a weighted average of visibilities unlike the previous method

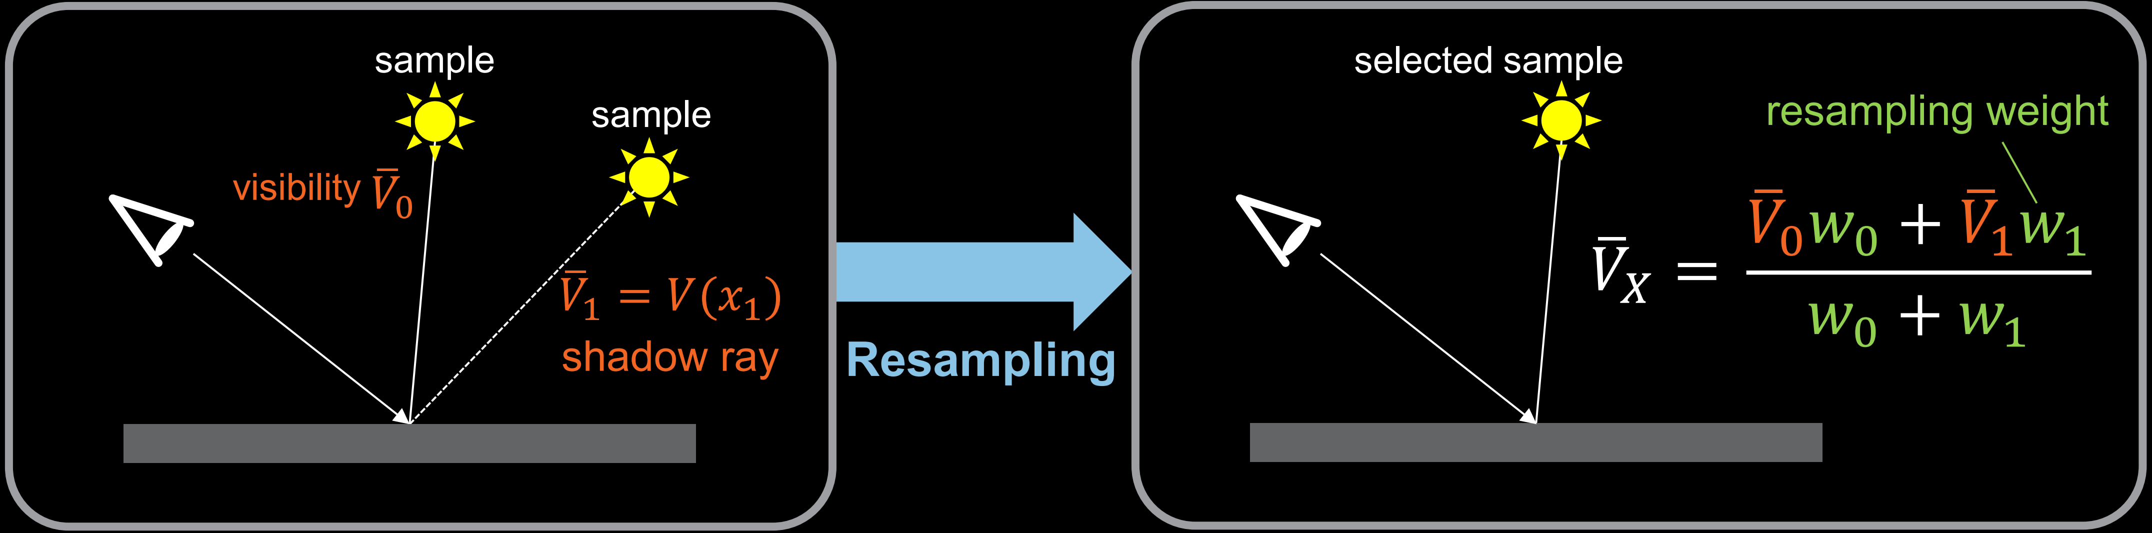

Simple ☺

Our Visibility Reuse

- Reuse spatiotemporal sample visibilities for the integrand as in Wyman and Panteleev [2021]
- Estimate a new visibility using a weighted average of visibilities unlike the previous method
- A variant of *weighted importance sampling (WIS)* [Bekaert et al. 2000] (a.k.a. ratio estimator [Heitz et al. 2018])
  - Canonical sample visibility is also estimated by WIS in a chained manner
  - Biased estimator

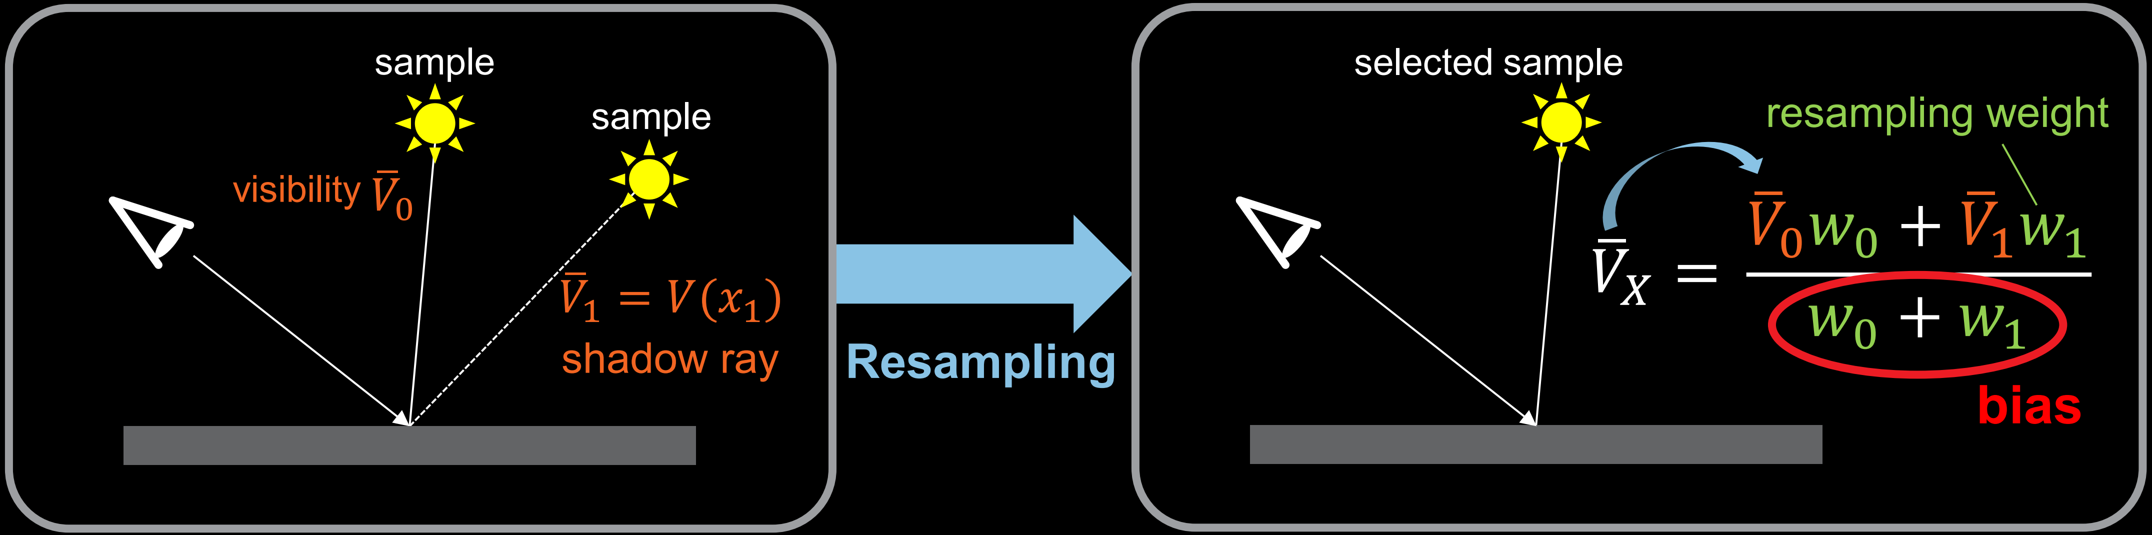

Simple ☺

Advanced Rendering Research Group

Bias Cancellation

- WIS is biased due to the normalization
- This normalization is recursively cancelled in ReSTIR ☺

$$\int_{\Omega} f(x)V(x)\mathrm{d}x \approx f(X)\bar{V}_X W_X = f(X)\frac{\sum_i \bar{V}_i w_i}{\cancel{\sum_i w_i}}\frac{\cancel{\sum_i w_i}}{\hat{p}_c(X)} = \frac{f(X)}{\hat{p}_c(X)}\sum_i \bar{V}_i w_i$$

$$\bar{V}_i w_i = m_i(T_i(x_i))\hat{p}_c(T_i(x_i))\frac{\sum_k \bar{V}_k w_k}{\cancel{\sum_k w_k}}\frac{\cancel{\sum_k w_k}}{\hat{p}_i(x_i)}\left|\frac{\partial T_i}{\partial x_i}\right|$$

AMD

Remaining Two Biases

- Temporal shadow delay inherited from the previous method
  - In other words, $\Sigma_i \bar{V}_i w_i$ is unbiased for static shadows
- Ratio of the light contribution $f(X)$ and target distribution $\hat{p}_c(X)$
  - If this ratio is constant, we can avoid the bias

$$\int_\Omega f(x)V(x)\mathrm{d}x \approx f(X)\bar{V}_X W_X = \frac{f(X)}{\hat{p}_c(X)} \Sigma_i \bar{V}_i w_i$$

$$\Sigma_i \bar{V}_i w_i \approx \int_\Omega \hat{p}_c(x)V(x)\mathrm{d}x$$

unbiased for static shadows

Biased correction term

$$\int_\Omega \hat{p}_c(x)V(x)\mathrm{d}x \Rightarrow \int_\Omega f(x)V(x)\mathrm{d}x$$

Using $\hat{p}_c(X) \propto f(X)$, our method is unbiased for static shadows

Avoid the Bias for Colored Light Sources?

- The light contribution $f(X)$ can have different distributions between RGB channels

- Generate samples for each color channel?
- Randomly sample a color channel?

Expensive for real-time ☹

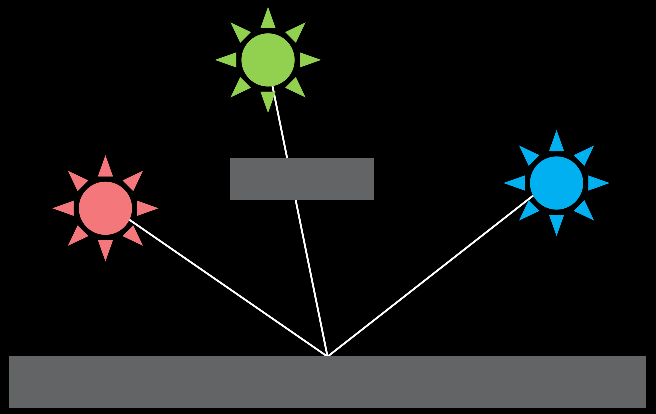

Biased Visibility Reuse for Real-time Rendering

- The light contribution $f(X)$ can have different distributions between RGB channels
- Approximate $\hat{p}_c(X)$ by the luminance of $f(X)$
- Color of invisible lights can leak into shadows, but it is still unbiased in luminance

Invisible light color

$$\int_{\Omega} f(x)V(x)\mathrm{d}x \approx \frac{f(X)}{\hat{p}_c(X)} \Sigma_i \bar{V}_i w_i$$

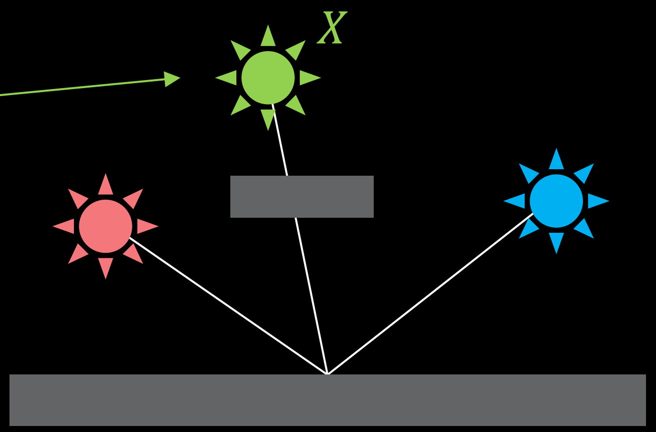

Invisible lights can be used as the correction term

Biased Visibility Reuse for Real-time Rendering

- The light contribution $f(X)$ can have different distributions between RGB channels
- Approximate $\hat{p}_c(X)$ by the luminance of $f(X)$
- Color of invisible lights can leak into shadows, but it is still unbiased in luminance
- Reduce this color leak bias by combining Bitterli et al.'s visibility reuse [2020]
  - Reduce the sampling of invisible lights

Invisible light color

$$\int_{\Omega} f(x)V(x)\mathrm{d}x \approx \frac{f(X)}{\hat{p}_c(X)} \Sigma_i \bar{V}_i w_i$$

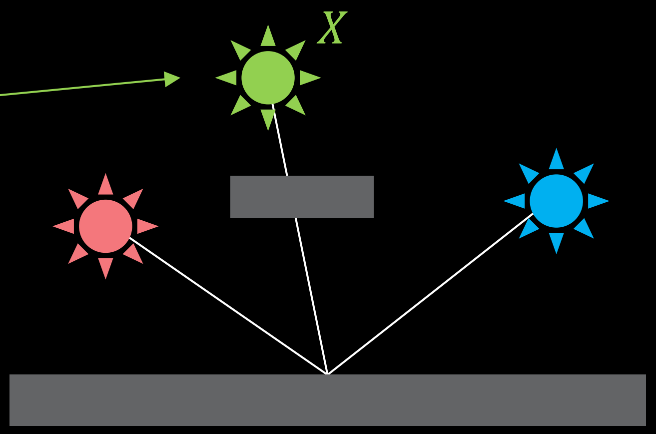

Biased Visibility Reuse for Real-time Rendering

- The light contribution $f(X)$ can have different distributions between RGB channels
- Approximate $\hat{p}_c(X)$ by the luminance of $f(X)$
- Color of invisible lights can leak into shadows, but it is still unbiased in luminance
- Reduce this color leak bias by combining Bitterli et al.'s visibility reuse [2020]
  - Reduce the sampling of invisible lights

visible light color

$$\int_{\Omega} f(x)V(x)\mathrm{d}x \approx \frac{f(X)}{\hat{p}_c(X)} \Sigma_i \bar{V}_i w_i$$

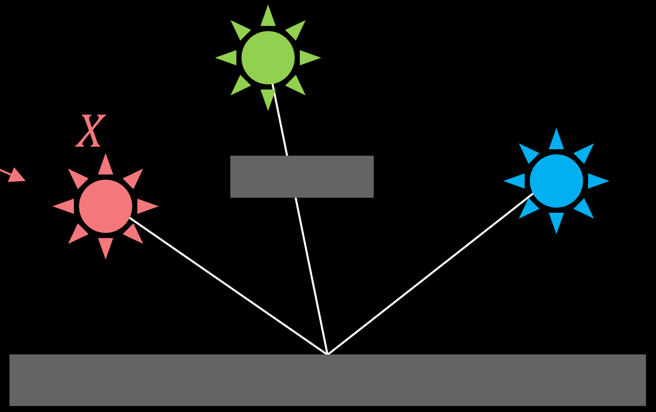

Results and Conclusion

Biased Visibility-Reuse ReSTIR, 2 rays/pixel

Previous method (7.85 ms, RMSPE: 80.3%)

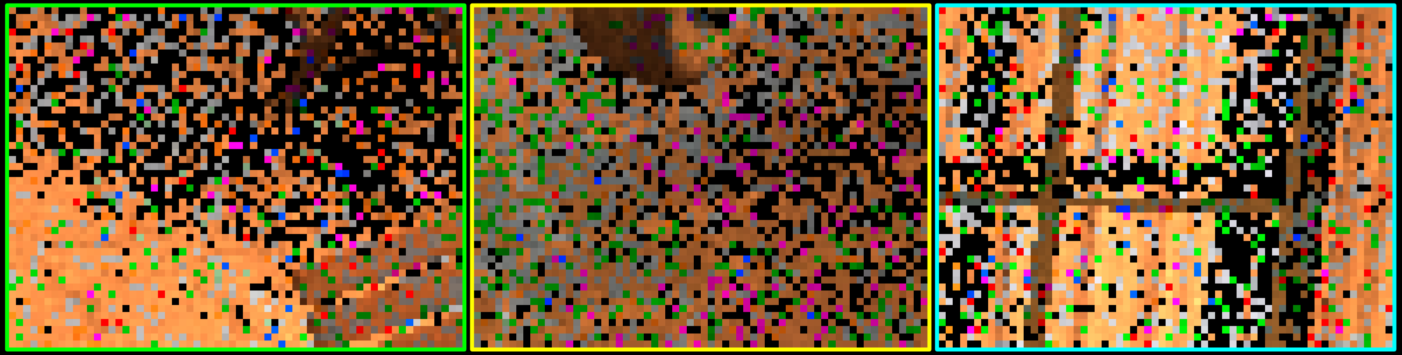

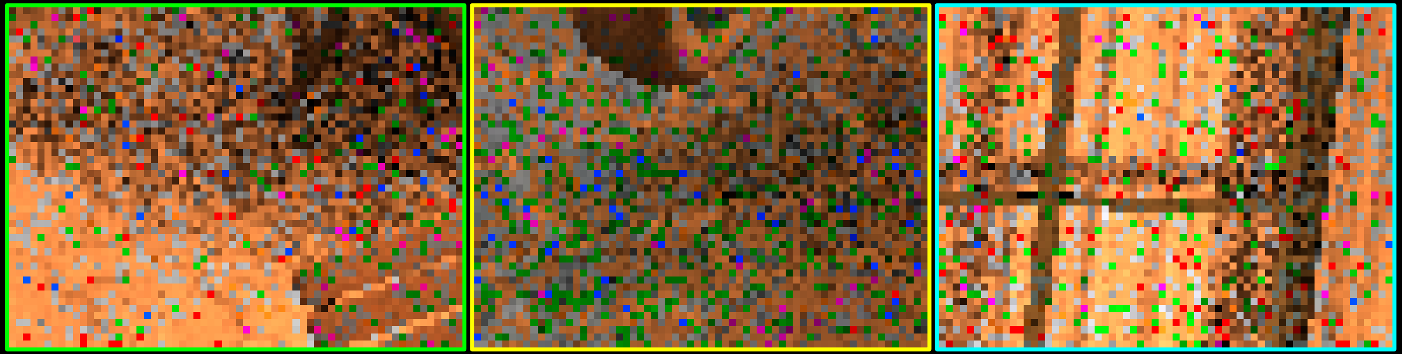

Our method (7.86 ms, RMSPE: 59.6%)

1920×1080 pixels
AMD Radeon™ RX 7900 XTX GPU

Visualization of Visibility Term

1920×1080 pixels
AMD Radeon^{TM} RX 7900 XTX GPU

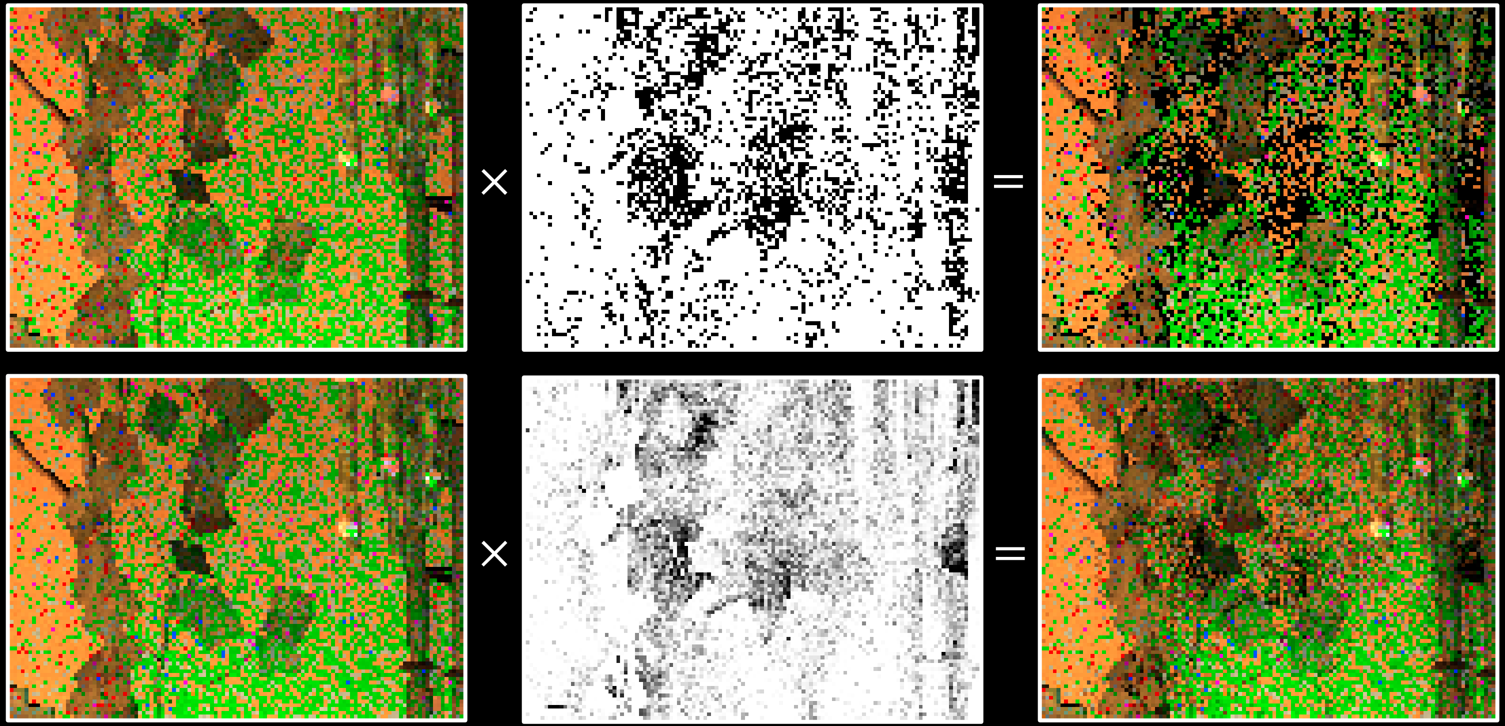

Indirect Illumination Using Virtual Point Lights [Keller 1997]

Limitation: Color Leak Bias

Reference

Previous method

RMSPE: 15.3%
(darkening bias)

Our method

RMSPE: 15.7%
(darkening + color leak)

- Color leak bias for practical real-time rendering
  - Biased in color, but unbiased in luminance
- This additional color leak bias (+0.4%) is significantly smaller than Bitterli et al.'s darkening bias (15.3%)
- Can avoid the bias by generating samples for each RGB channel or sampling a color channel

Conclusion

- Variance reduction technique based on WIS for visibility-reuse ReSTIR
  - Bias due to WIS is recursively cancelled in ReSTIR
- Color leak bias can remain for colored lights
  - Can reduce this bias by combining Bitterli et al.'s visibility reuse
  - Also discussed conditions to obtain unbiased results
- Simple and efficient
  - Suitable for real-time rendering

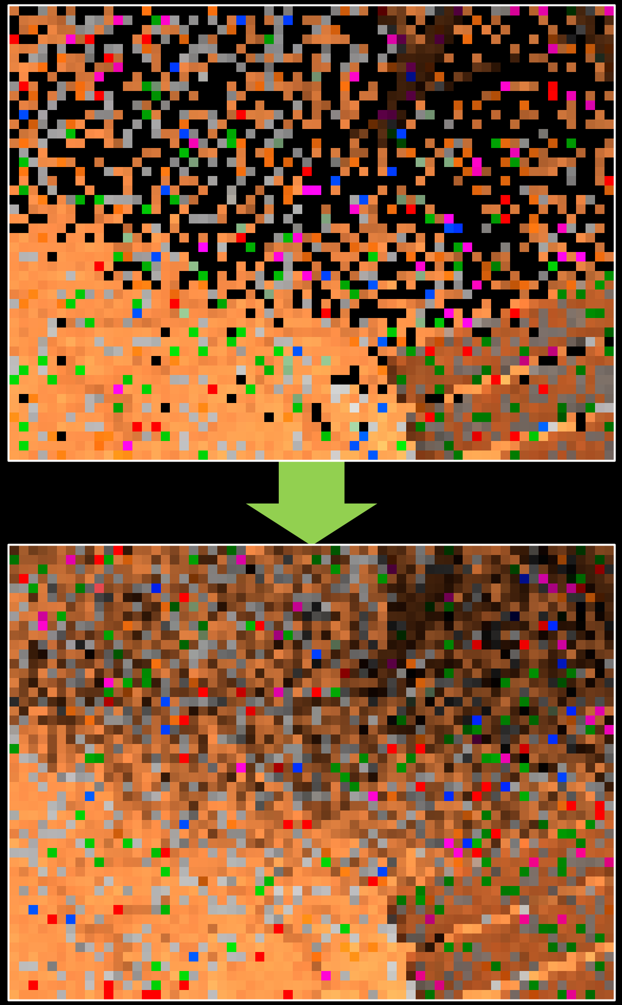

References

- Philippe Bekaert, Mateu Sbert, and Yves D. Willems. 2000. Weighted Importance Sampling Techniques for Monte Carlo Radiosity. In EGWR '00. 35–46.
- Benedikt Bitterli, ChrisWyman, Matt Pharr, Peter Shirley, Aaron Lefohn, and Wojciech Jarosz. 2020. Spatiotemporal Reservoir Resampling for Real-Time Ray Tracing with Dynamic Direct Lighting. ACM Trans. Graph. 39, 4, Article 148.
- Eric Heitz, Stephen Hill, and Morgan McGuire. 2018. Combining Analytic Direct Illumination and Stochastic Shadows. In I3D '18. Article 2.
- Alexander Keller. 1997. Instant radiosity. In SIGGRAPH ‘97. 49–56.
- Amazon Lumberyard. 2017. Amazon Lumberyard Bistro, Open Research Content Archive (ORCA).
- Chris Wyman and Alexey Panteleev. 2021. Rearchitecting Spatiotemporal Resampling for Production. In HPG ’21.

ARR

Advanced Rendering Research Group