q2vkpt

Quake 2

Vulkan

Path Tracer

Christoph Schied

@c_schied

http://brechpunkt.de/q2vkpt

Path Tracing

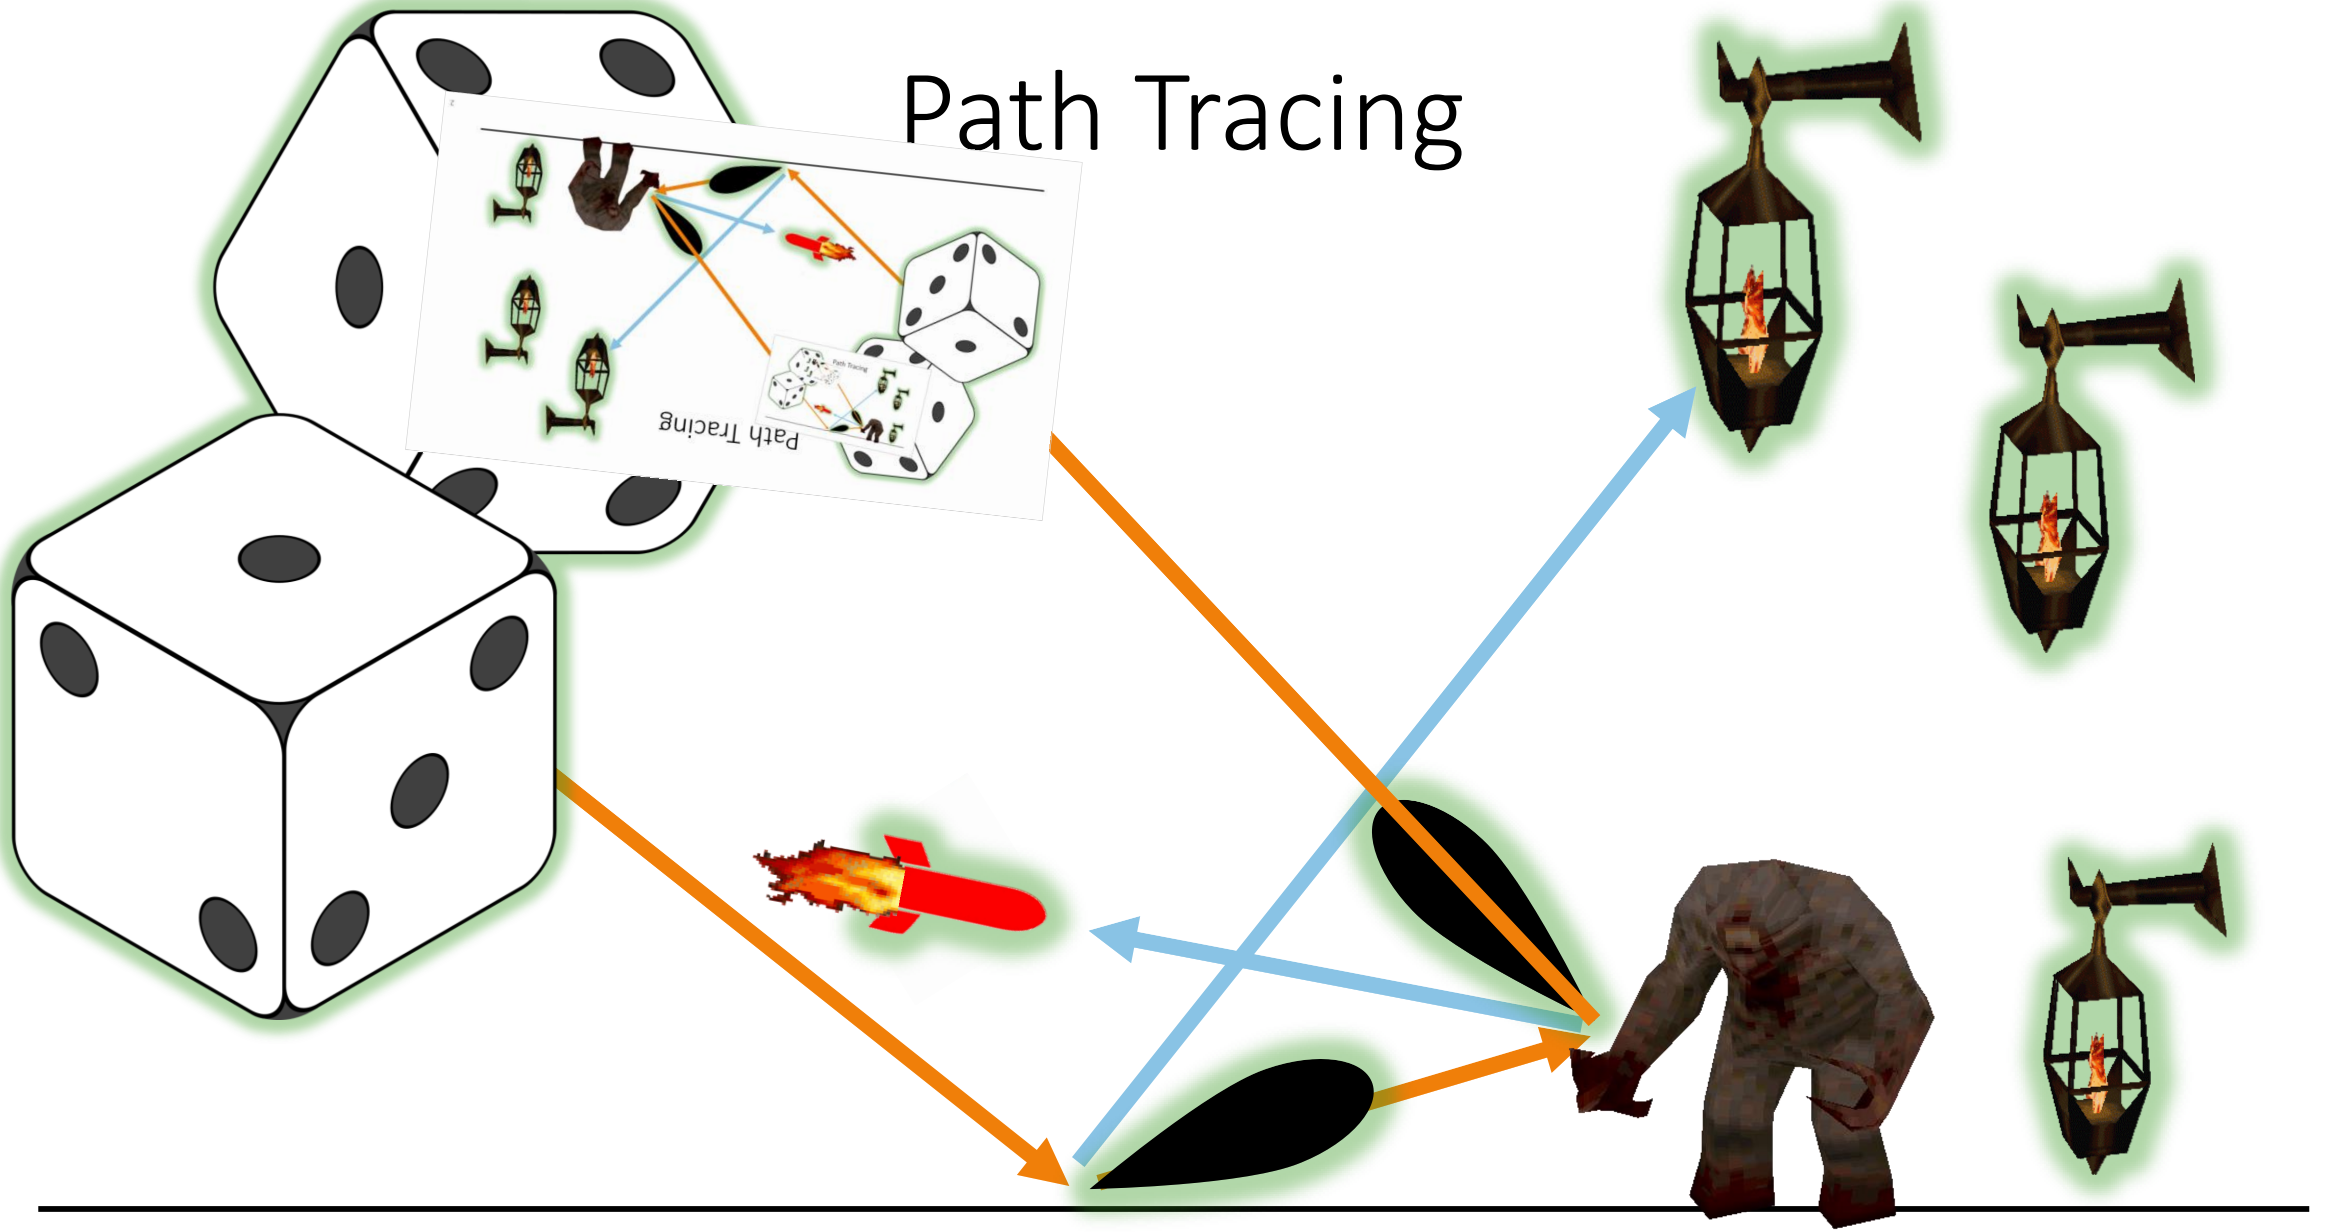

Path Tracing

Number of samples, n=1 in q2vkpt

$$L \approx \frac{1}{n}\sum_{i=1}^{n}$$

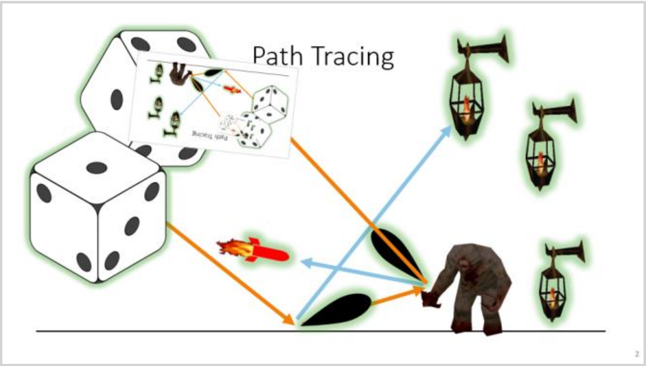

Main challenges

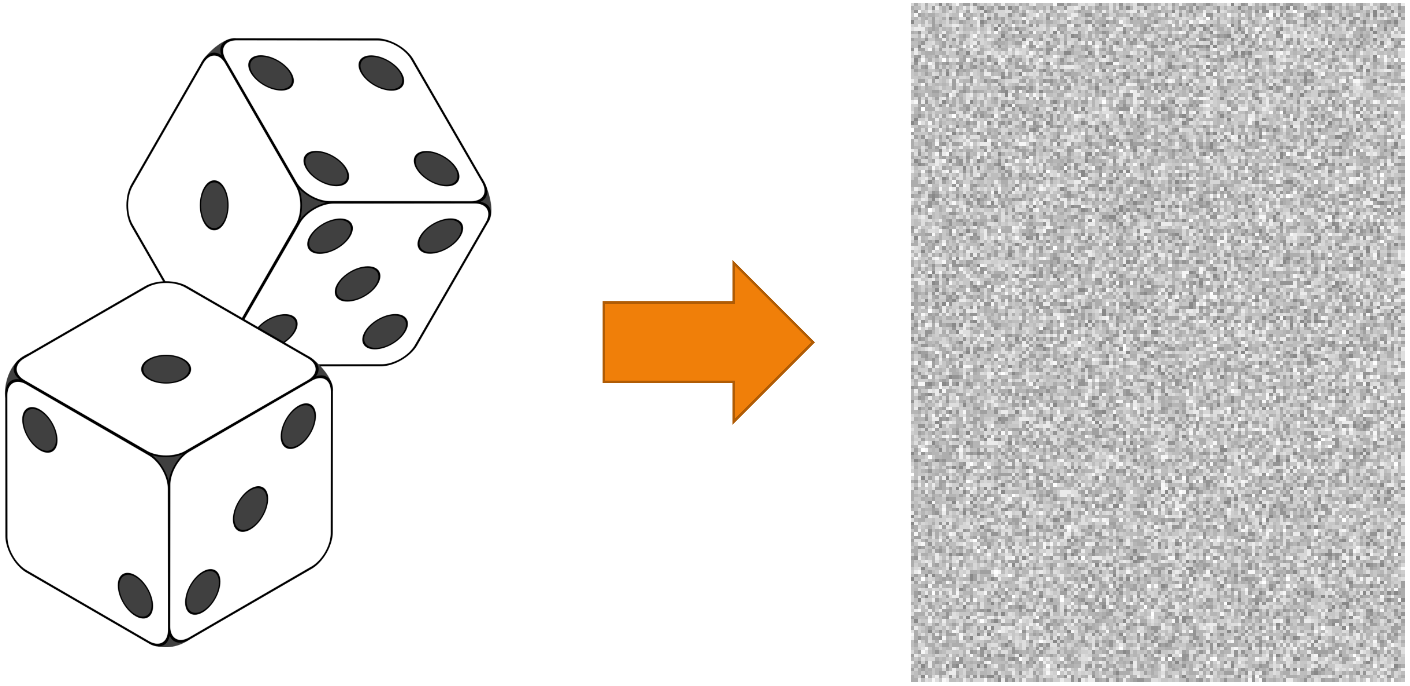

- Better Sampling → less noise
- Denoising

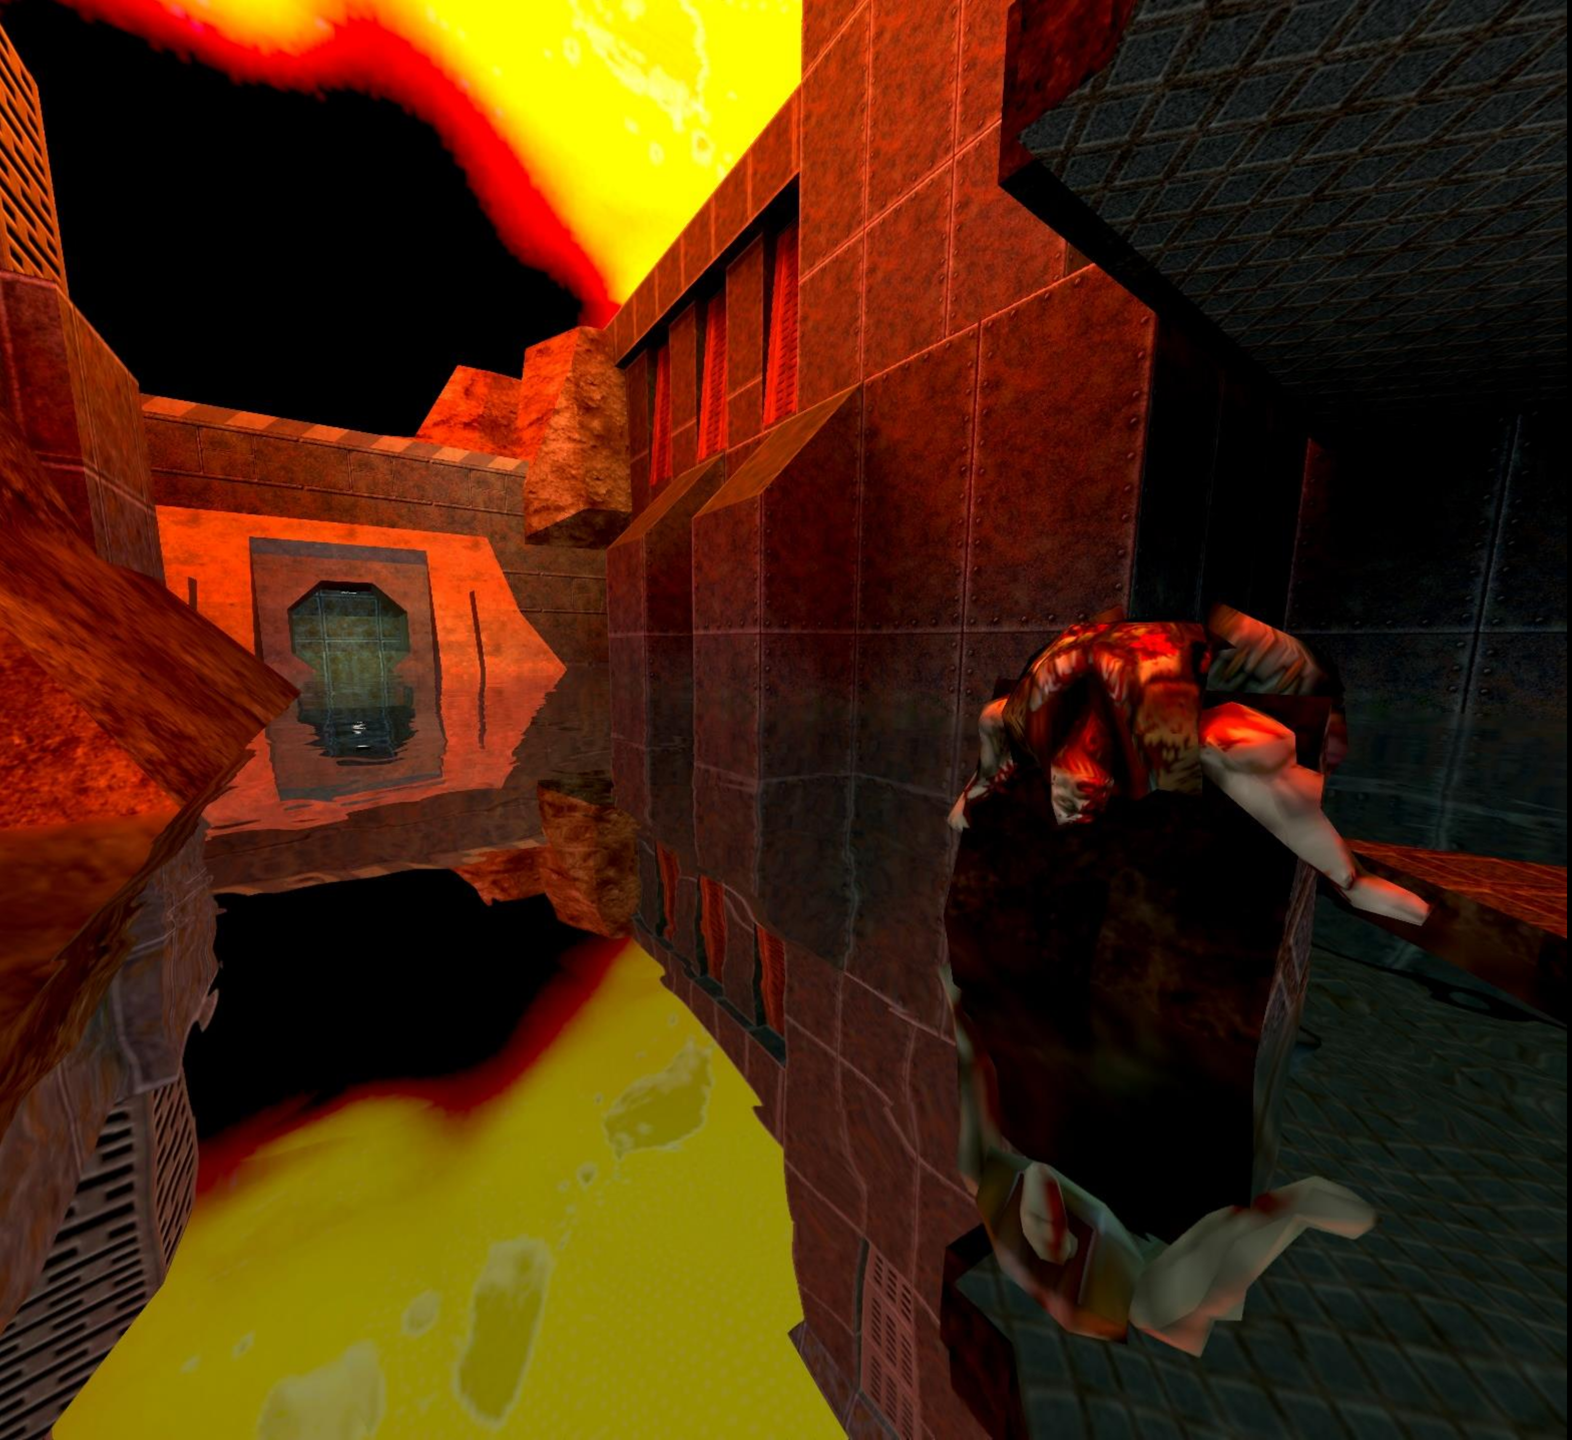

Q2VKPT

- Research prototype to evaluate current state of real-time path tracing
- Open source https://github.com/cschied/q2vkpt
- Entirely raytraced
- Real-time path tracing (one indirect bounce)
- C99, Vulkan, GLSL, RTX

https://www.youtube.com/watch?v=vrq1T93uLag

2560x1440, RTX2080 Ti

Denoising

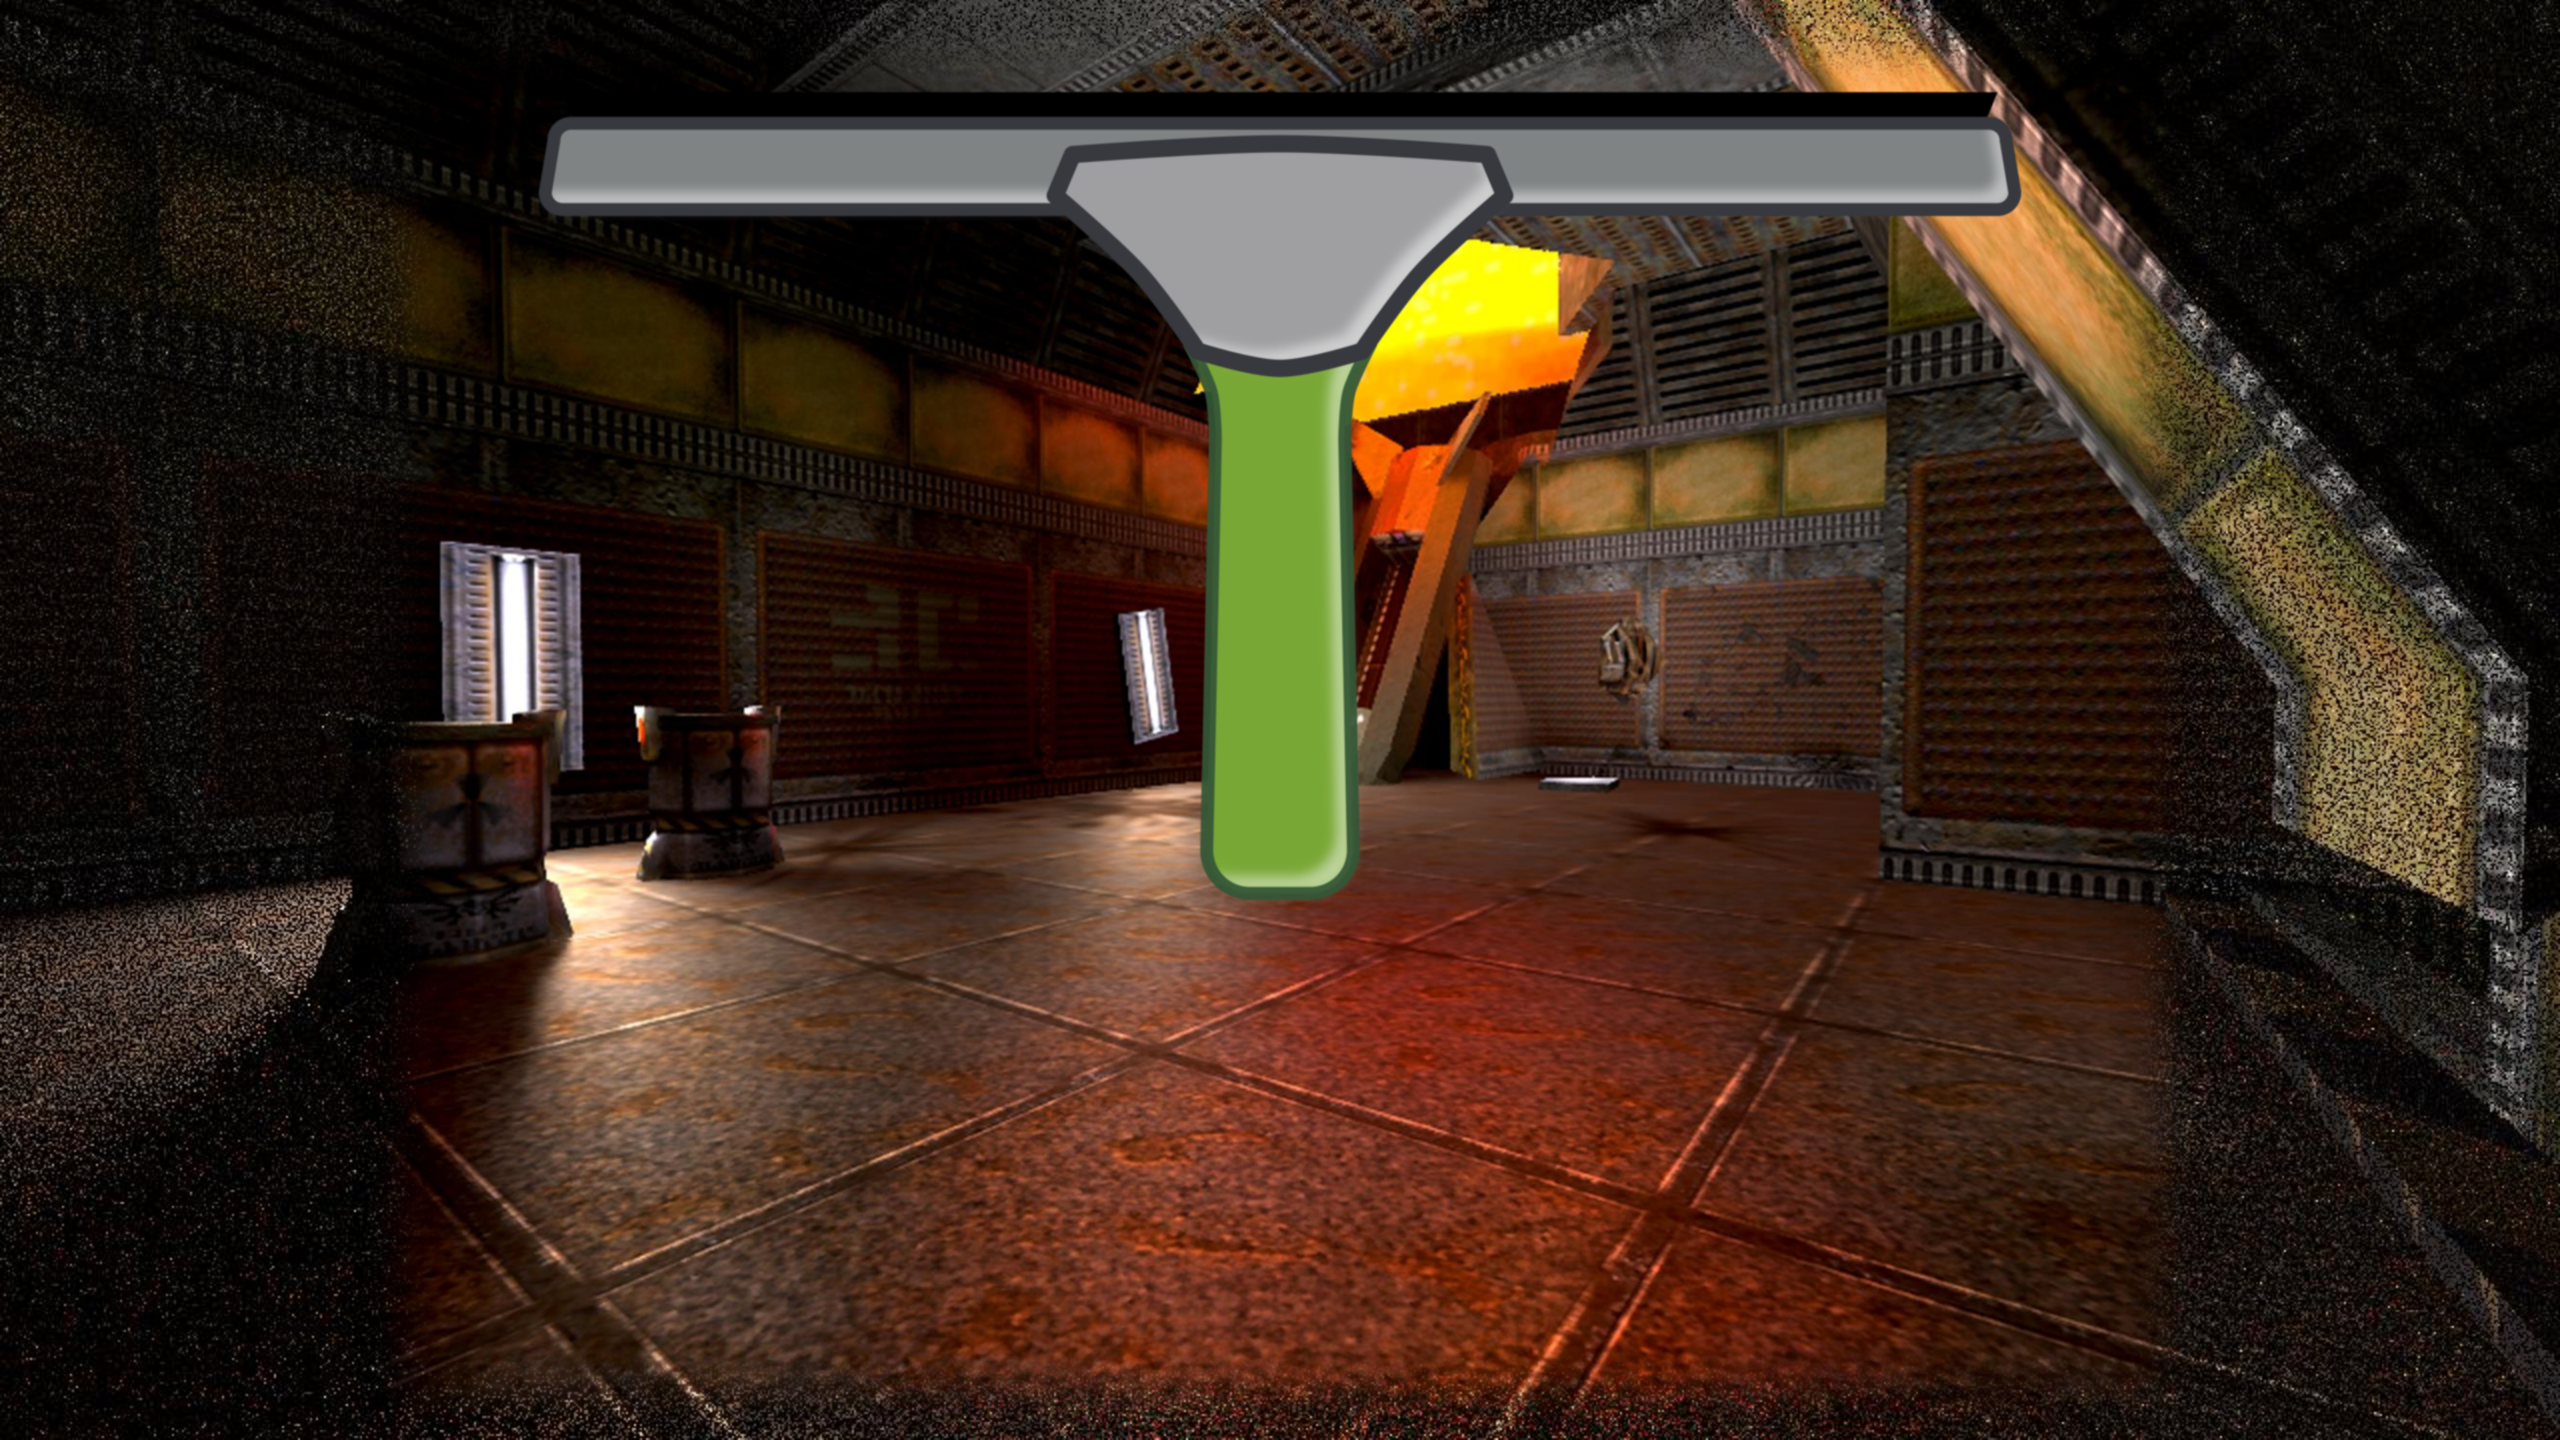

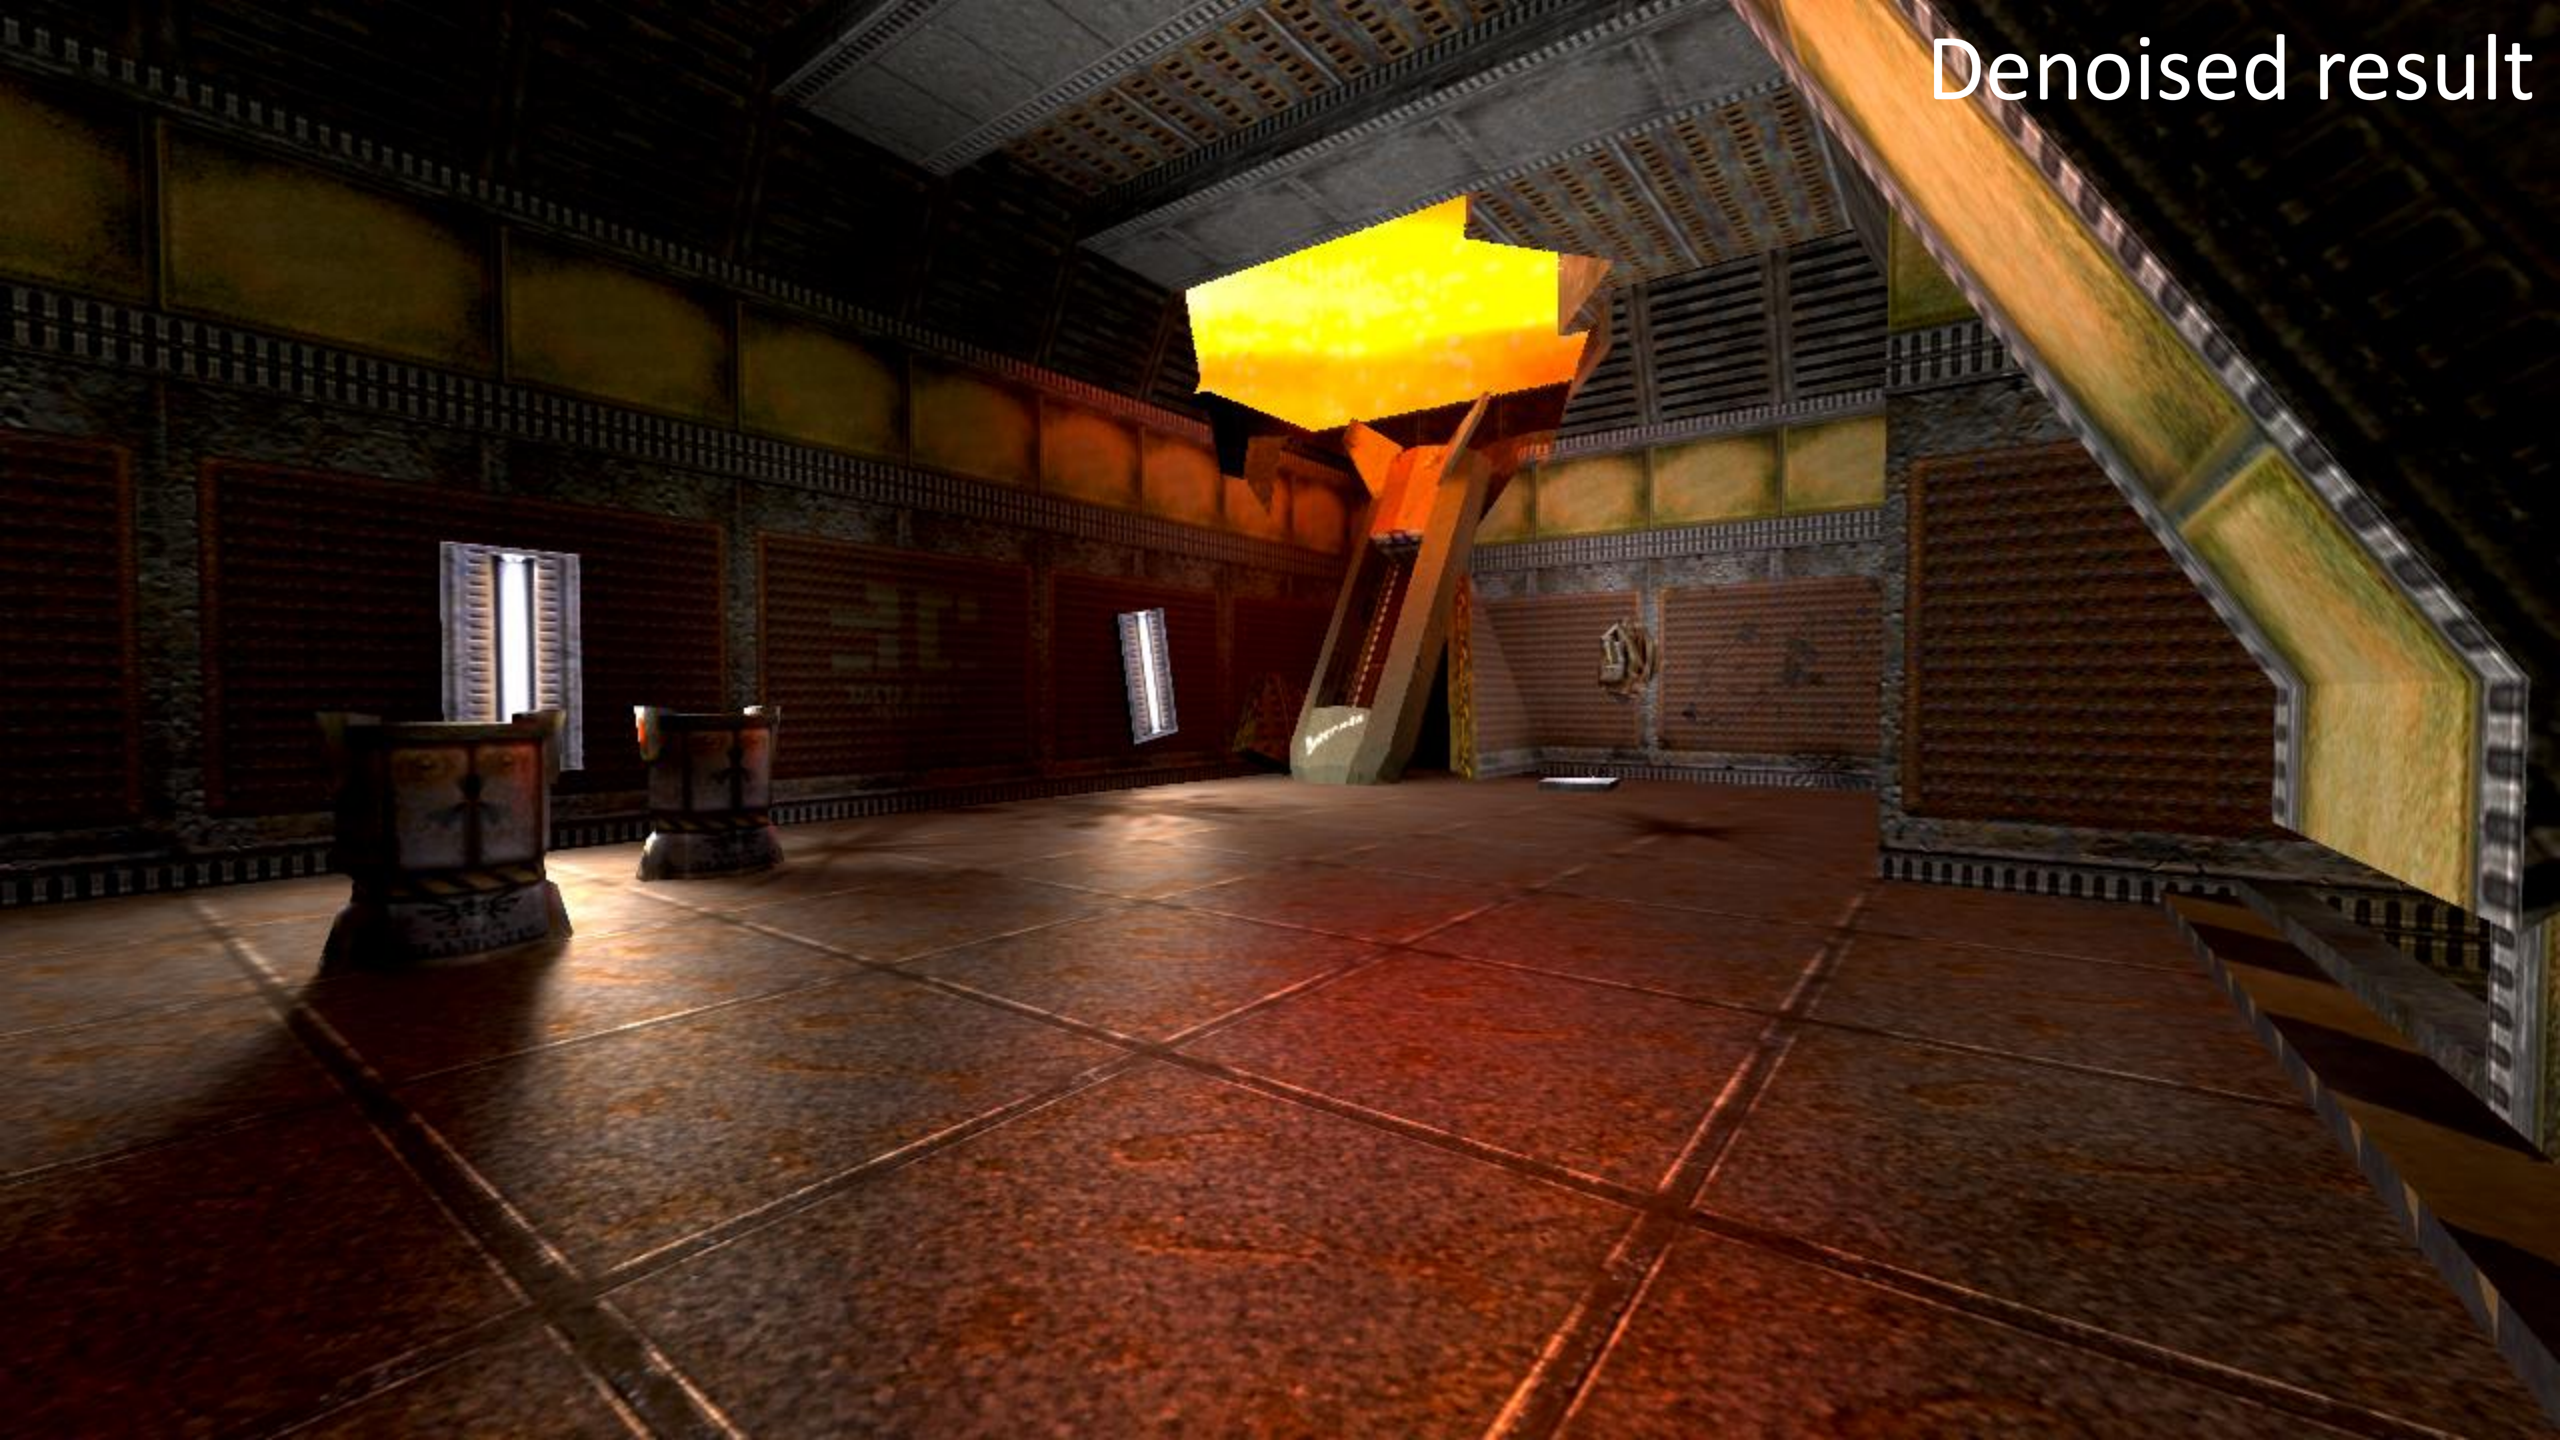

Denoised result


Path tracer output
(1spp)

Denoising (A-SVGF)

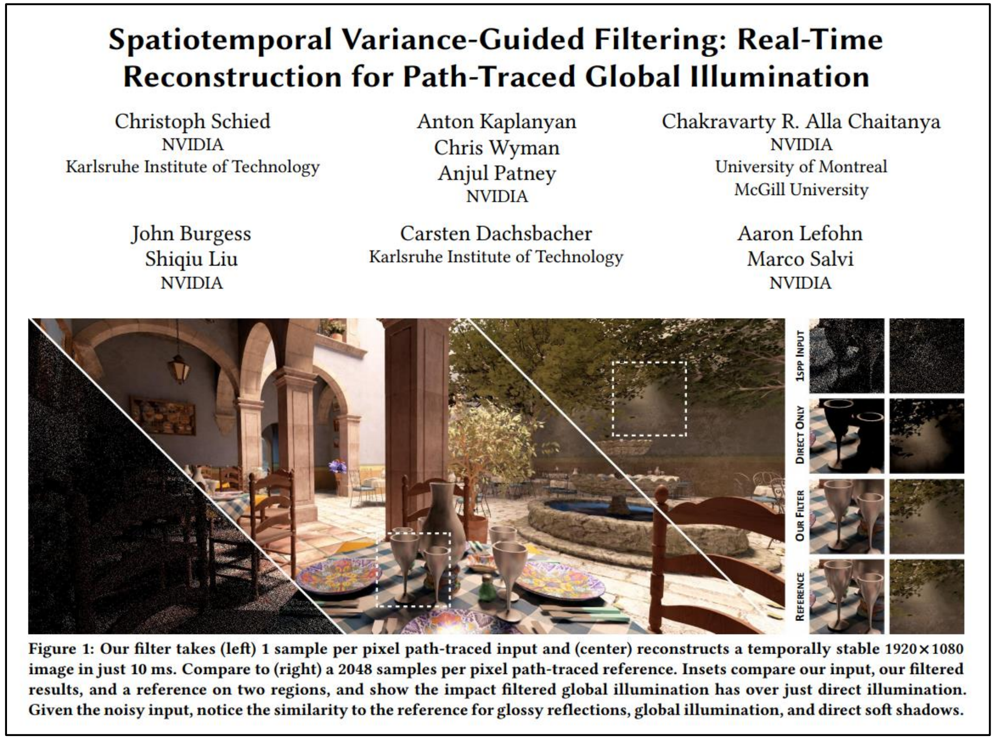

Spatiotemporal Variance-Guided Filtering: Real-Time Reconstruction for Path-Traced Global Illumination

Christoph Schied
NVIDIA
Karlsruhe Institute of Technology

Anton Kaplanyan
Chris Wyman
Anjul Patney
NVIDIA

Chakravarty R. Alla Chaitanya
NVIDIA
University of Montreal
McGill University

John Burgess
Shiqiu Liu
NVIDIA

Carsten Dachsbacher
Karlsruhe Institute of Technology

Aaron Lefohn
Marco Salvi
NVIDIA

1spp Input
Direct Only
Our Filter
Reference

Figure 1: Our filter takes (left) 1 sample per pixel path-traced input and (center) reconstructs a temporally stable 1920×1080 image in just 10 ms. Compare to (right) a 2048 samples per pixel path-traced reference. Insets compare our input, our filtered results, and a reference on two regions, and show the impact filtered global illumination has over just direct illumination. Given the noisy input, notice the similarity to the reference for glossy reflections, global illumination, and direct soft shadows.

+

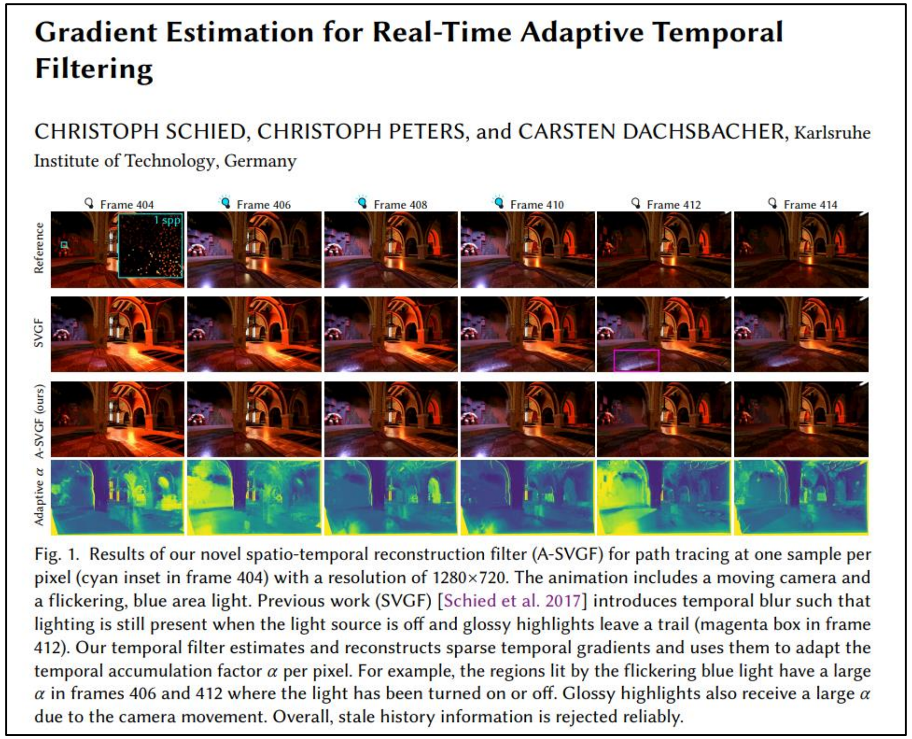

Gradient Estimation for Real-Time Adaptive Temporal Filtering

CHRISTOPH SCHIED, CHRISTOPH PETERS, and CARSTEN DACHSBACHER, Karlsruhe Institute of Technology, Germany

Frame 404 | Frame 406 | Frame 408 | Frame 410 | Frame 412 | Frame 414
1 spp
Reference
SVGF
A-SVGF (ours)
Adaptive α

Fig. 1. Results of our novel spatio-temporal reconstruction filter (A-SVGF) for path tracing at one sample per pixel (cyan inset in frame 404) with a resolution of 1280×720. The animation includes a moving camera and a flickering, blue area light. Previous work (SVGF) [Schied et al. 2017] introduces temporal blur such that lighting is still present when the light source is off and glossy highlights leave a trail (magenta box in frame 412). Our temporal filter estimates and reconstructs sparse temporal gradients and uses them to adapt the temporal accumulation factor α per pixel. For example, the regions lit by the flickering blue light have a large α in frames 406 and 412 where the light has been turned on or off. Glossy highlights also receive a large α due to the camera movement. Overall, stale history information is rejected reliably.

Denoising (A-SVGF)

Spatiotemporal Variance-Guided Filtering: Real-Time Reconstruction for Path-Traced Global Illumination

Christoph Schied
NVIDIA
Karlsruhe Institute of Technology

Anton Kaplanyan
Chris Wyman
Anjul Patney
NVIDIA

Chakravarty R. Alla Chaitanya
NVIDIA
University of Montreal
McGill University

John Burgess
Shiqiu Liu
NVIDIA

Carsten Dachsbacher
Karlsruhe Institute of Technology

Aaron Lefohn
Marco Salvi
NVIDIA

Figure 1: Our filter takes (left) 1 sample per pixel path-traced input and (center) reconstructs a temporally stable 1920×1080 image in just 10 ms. Compare to (right) a 2048 samples per pixel path-traced reference. Insets compare our input, our filtered results, and a reference on two regions, and show the impact filtered global illumination has over just direct illumination. Given the noisy input, notice the similarity to the reference for glossy reflections, global illumination, and direct soft shadows.

+

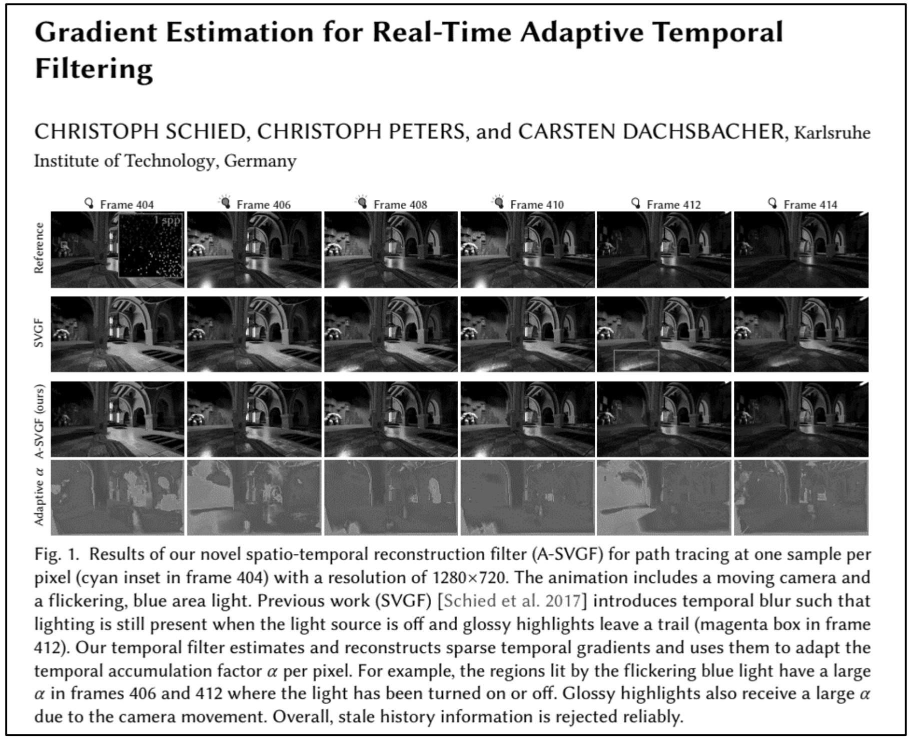

Gradient Estimation for Real-Time Adaptive Temporal Filtering

CHRISTOPH SCHIED, CHRISTOPH PETERS, and CARSTEN DACHSBACHER, Karlsruhe Institute of Technology, Germany

Fig. 1. Results of our novel spatio-temporal reconstruction filter (A-SVGF) for path tracing at one sample per pixel (cyan inset in frame 404) with a resolution of 1280×720. The animation includes a moving camera and a flickering, blue area light. Previous work (SVGF) [Schied et al. 2017] introduces temporal blur such that lighting is still present when the light source is off and glossy highlights leave a trail (magenta box in frame 412). Our temporal filter estimates and reconstructs sparse temporal gradients and uses them to adapt the temporal accumulation factor α per pixel. For example, the regions lit by the flickering blue light have a large α in frames 406 and 412 where the light has been turned on or off. Glossy highlights also receive a large α due to the camera movement. Overall, stale history information is rejected reliably.

Main concepts of SVGF

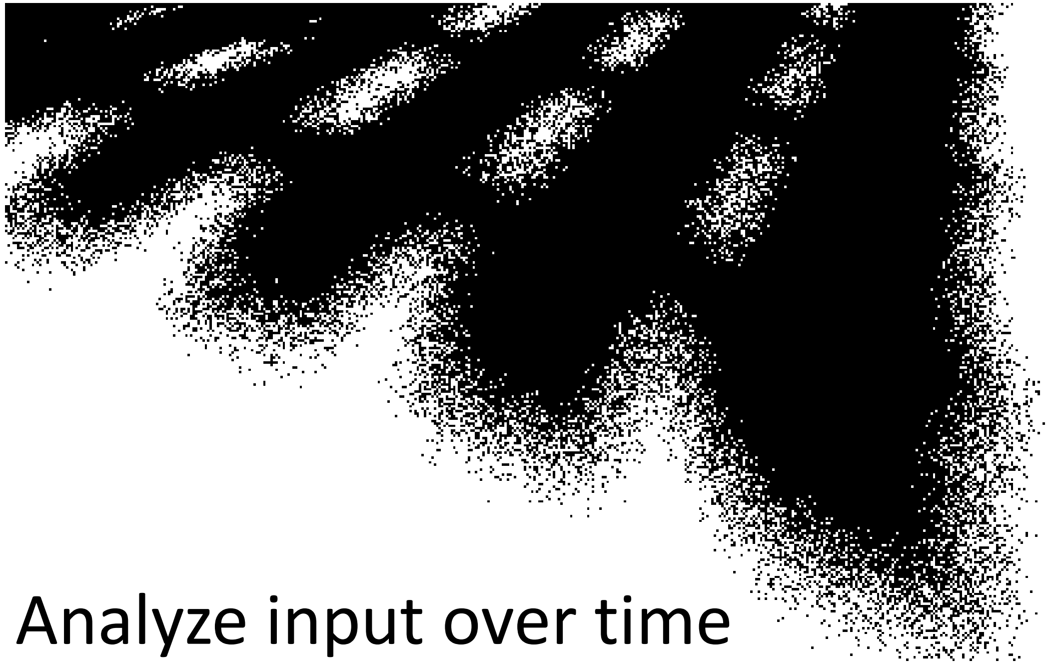

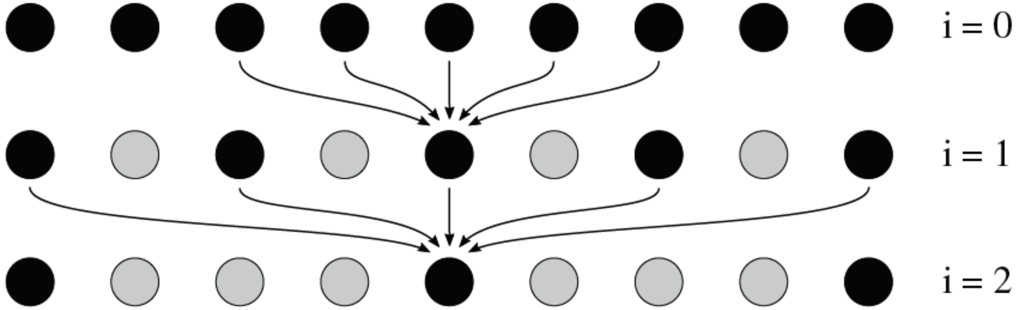

Filter hierarchically, starting small

- Estimate temporal stability after each filter iteration

→ Strong blur more likely in early iterations

Analyze input over time

- Temporally unstable → blur more
- Temporally stable → blur less

SVGF

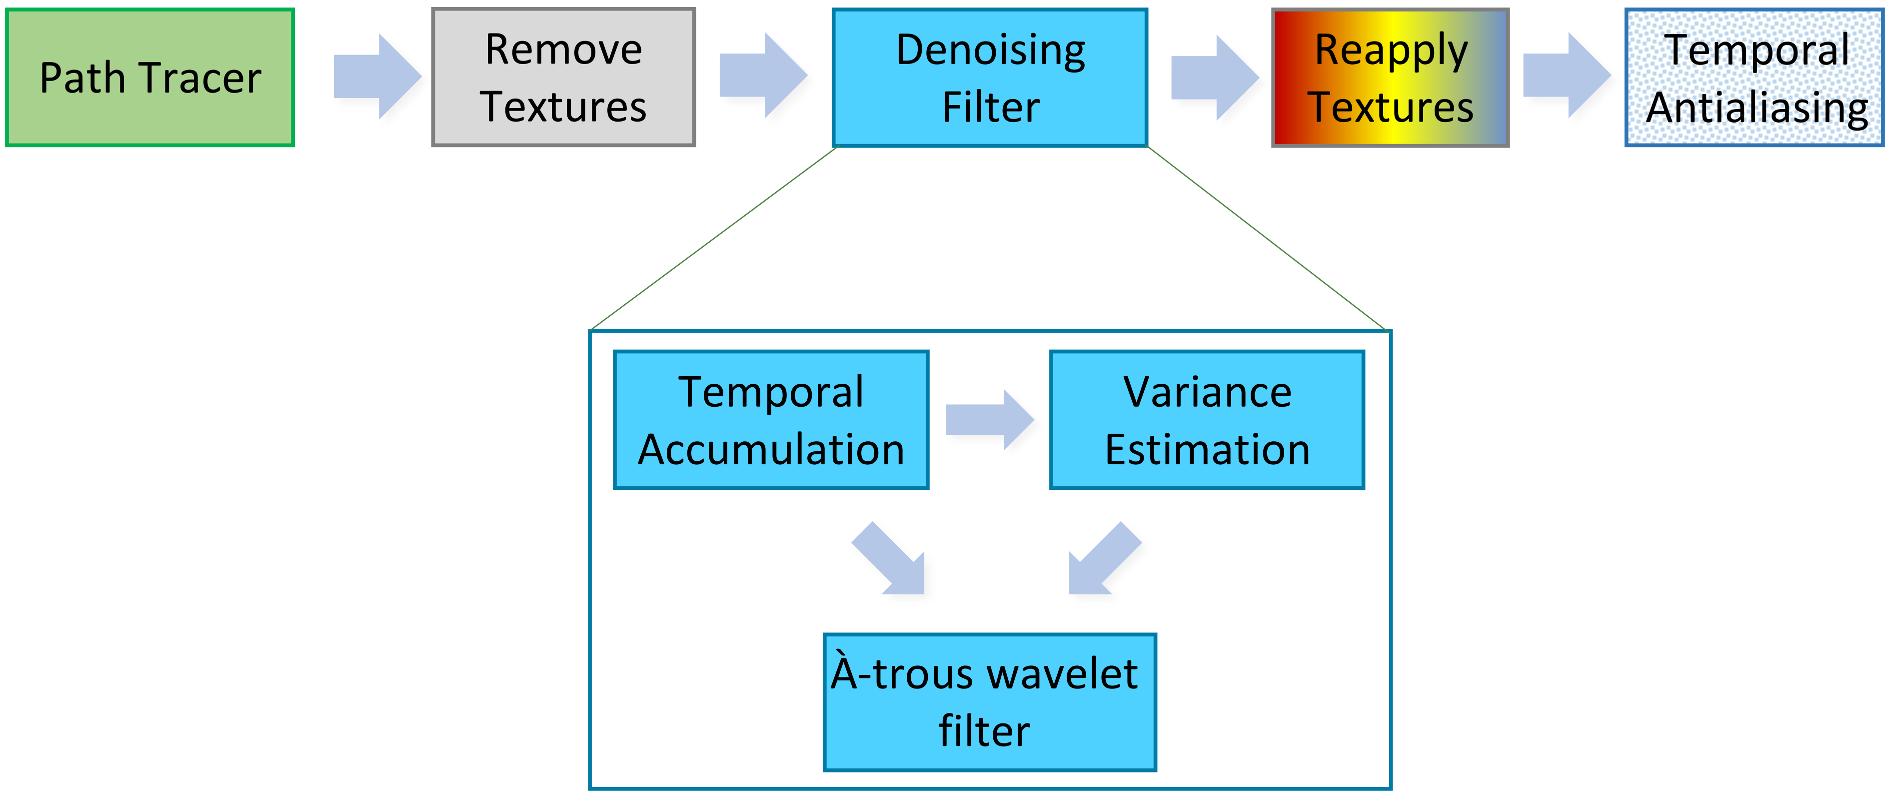
Path Tracer
Remove Textures
Denoising Filter
Reapply Textures
Temporal Antialiasing
Temporal Accumulation
Variance Estimation
À-trous wavelet filter

Edge-avoiding À-trous Wavelets

$\hat{c}_i(q)$

$h(q) \cdot w(p,q)$

$$\hat{c}_{i+1}(p) = \frac{\sum_{q \in \Omega} h(q) \cdot w(p,q) \cdot \hat{c}_i(q)}{\sum_{q \in \Omega} h(q) \cdot w(p,q)}$$

In q2vkpt:

- 3x3 box kernel
- 5 iterations

https://www.youtube.com/watch?v=Fv-jStEsCpE&t=48s

SVGF

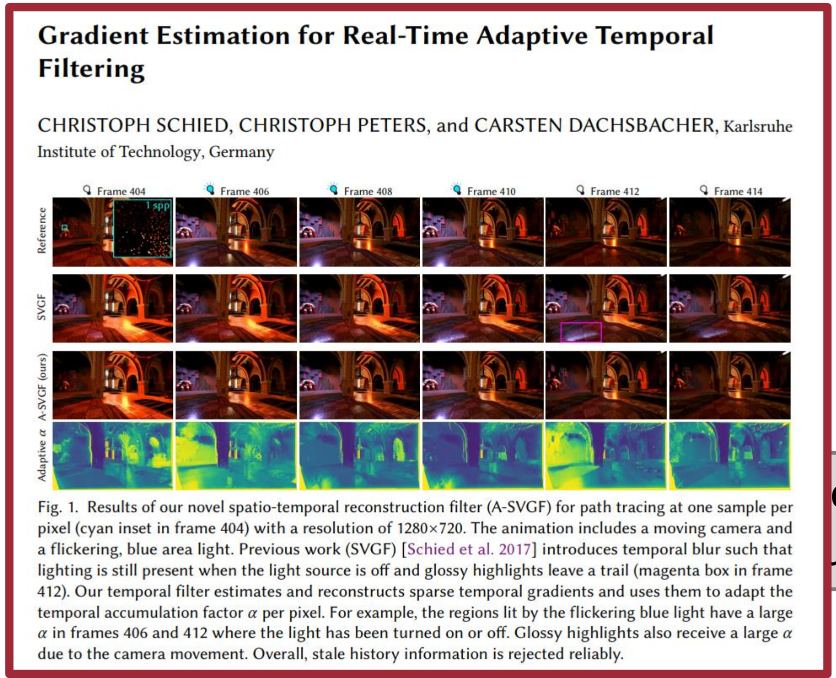

Gradient Estimation for Real-Time Adaptive Temporal Filtering

CHRISTOPH SCHIED, CHRISTOPH PETERS, and CARSTEN DACHSBACHER, Karlsruhe Institute of Technology, Germany

Fig. 1. Results of our novel spatio-temporal reconstruction filter (A-SVGF) for path tracing at one sample per pixel (cyan inset in frame 404) with a resolution of 1280×720. The animation includes a moving camera and a flickering, blue area light. Previous work (SVGF) [Schied et al. 2017] introduces temporal blur such that lighting is still present when the light source is off and glossy highlights leave a trail (magenta box in frame 412). Our temporal filter estimates and reconstructs sparse temporal gradients and uses them to adapt the temporal accumulation factor α per pixel. For example, the regions lit by the flickering blue light have a large α in frames 406 and 412 where the light has been turned on or off. Glossy highlights also receive a large α due to the camera movement. Overall, stale history information is rejected reliably.

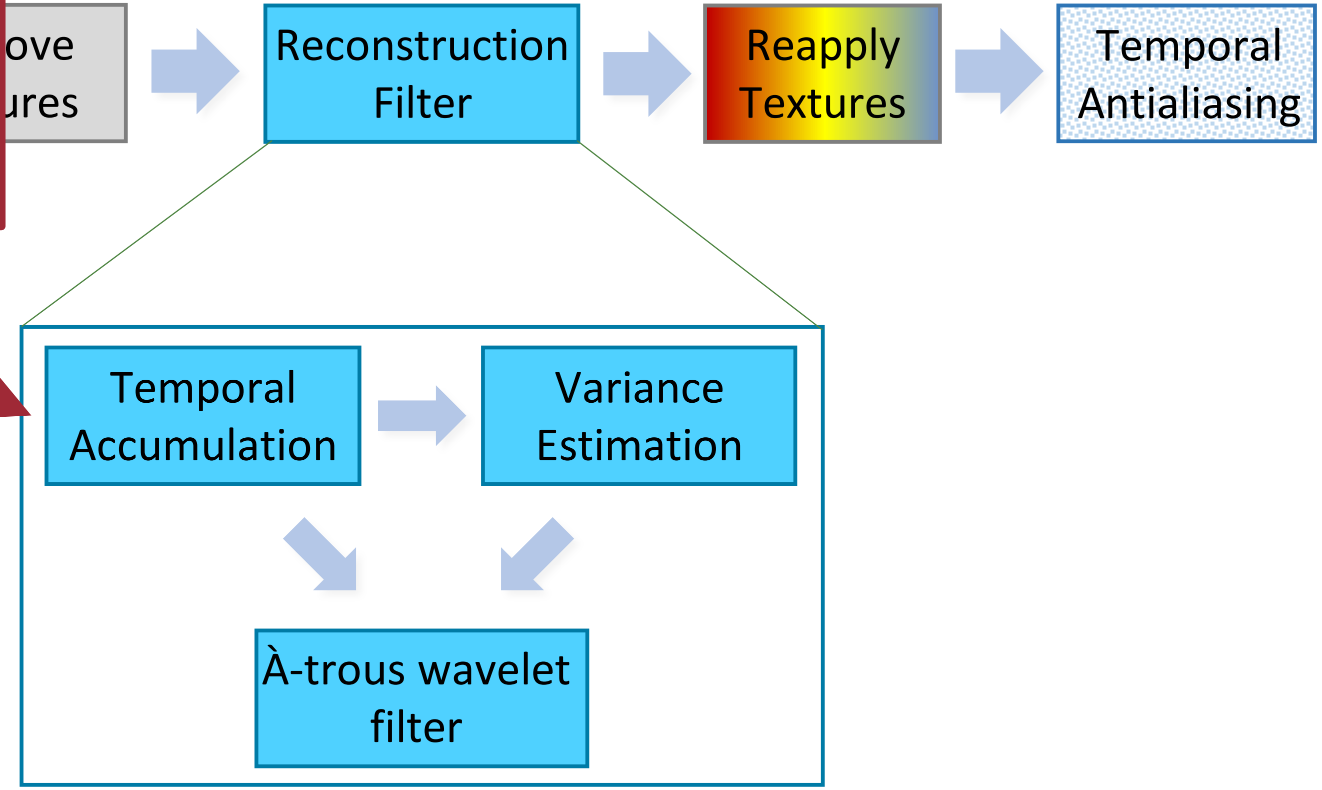

Denoising (A-SVGF)

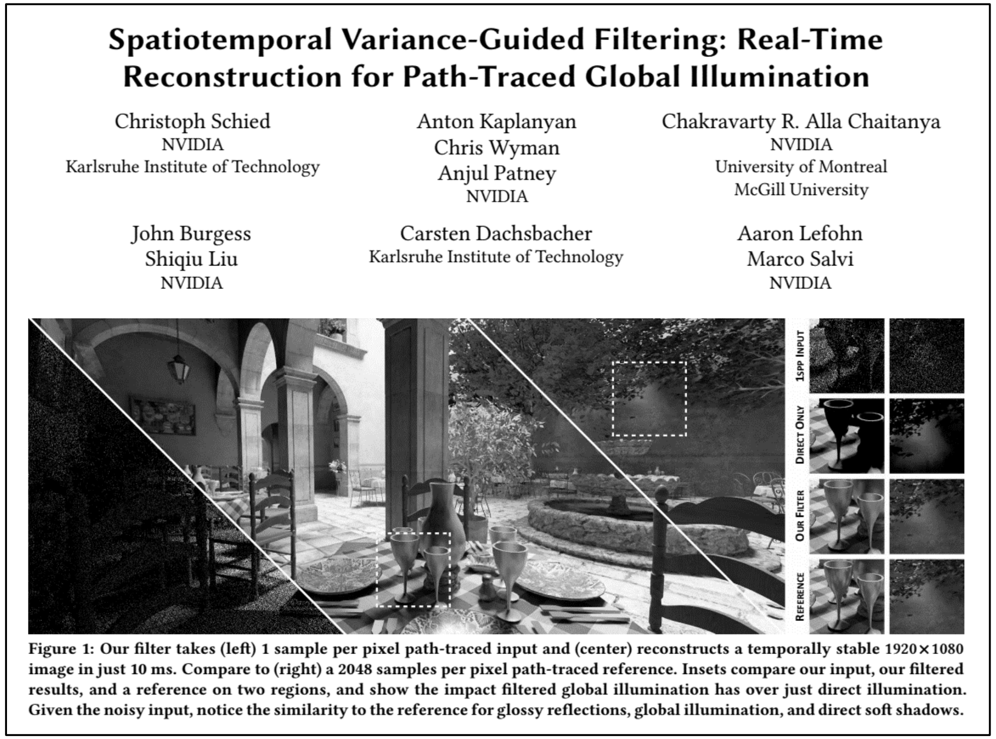

Spatiotemporal Variance-Guided Filtering: Real-Time Reconstruction for Path-Traced Global Illumination

Christoph Schied
NVIDIA
Karlsruhe Institute of Technology

Anton Kaplanyan
Chris Wyman
Anjul Patney
NVIDIA

Chakravarty R. Alla Chaitanya
NVIDIA
University of Montreal
McGill University

John Burgess
Shiqiu Liu
NVIDIA

Carsten Dachsbacher
Karlsruhe Institute of Technology

Aaron Lefohn
Marco Salvi
NVIDIA

Figure 1: Our filter takes (left) 1 sample per pixel path-traced input and (center) reconstructs a temporally stable 1920×1080 image in just 10 ms. Compare to (right) a 2048 samples per pixel path-traced reference. Insets compare our input, our filtered results, and a reference on two regions, and show the impact filtered global illumination has over just direct illumination. Given the noisy input, notice the similarity to the reference for glossy reflections, global illumination, and direct soft shadows.

+

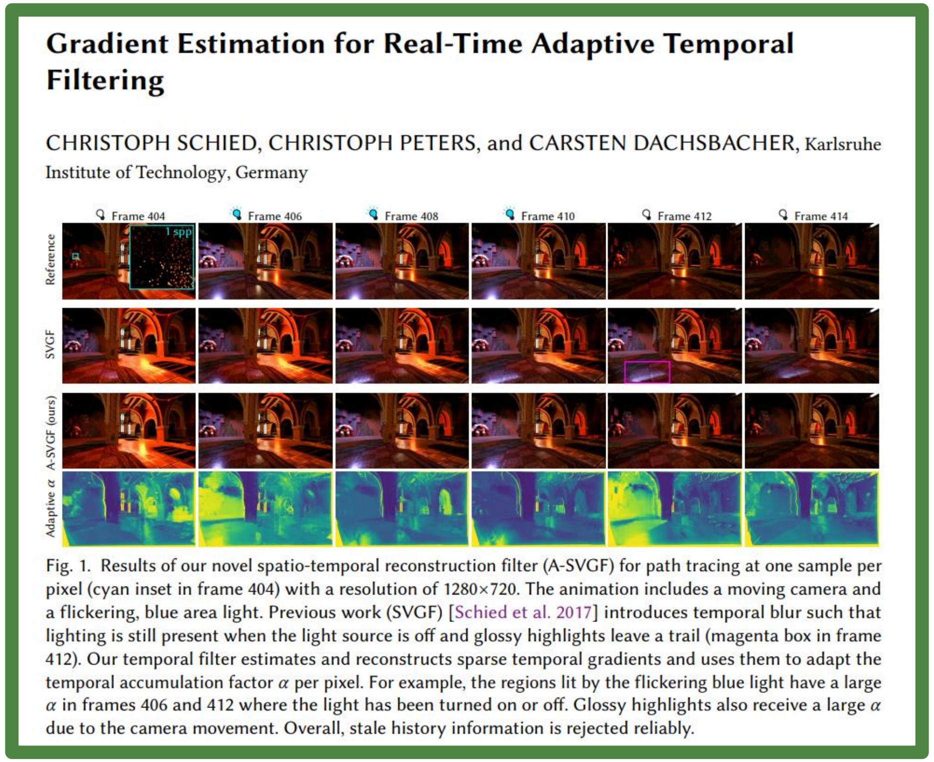

Gradient Estimation for Real-Time Adaptive Temporal Filtering

CHRISTOPH SCHIED, CHRISTOPH PETERS, and CARSTEN DACHSBACHER, Karlsruhe Institute of Technology, Germany

Fig. 1. Results of our novel spatio-temporal reconstruction filter (A-SVGF) for path tracing at one sample per pixel (cyan inset in frame 404) with a resolution of 1280×720. The animation includes a moving camera and a flickering, blue area light. Previous work (SVGF) [Schied et al. 2017] introduces temporal blur such that lighting is still present when the light source is off and glossy highlights leave a trail (magenta box in frame 412). Our temporal filter estimates and reconstructs sparse temporal gradients and uses them to adapt the temporal accumulation factor α per pixel. For example, the regions lit by the flickering blue light have a large α in frames 406 and 412 where the light has been turned on or off. Glossy highlights also receive a large α due to the camera movement. Overall, stale history information is rejected reliably.

Screen-space Reprojection

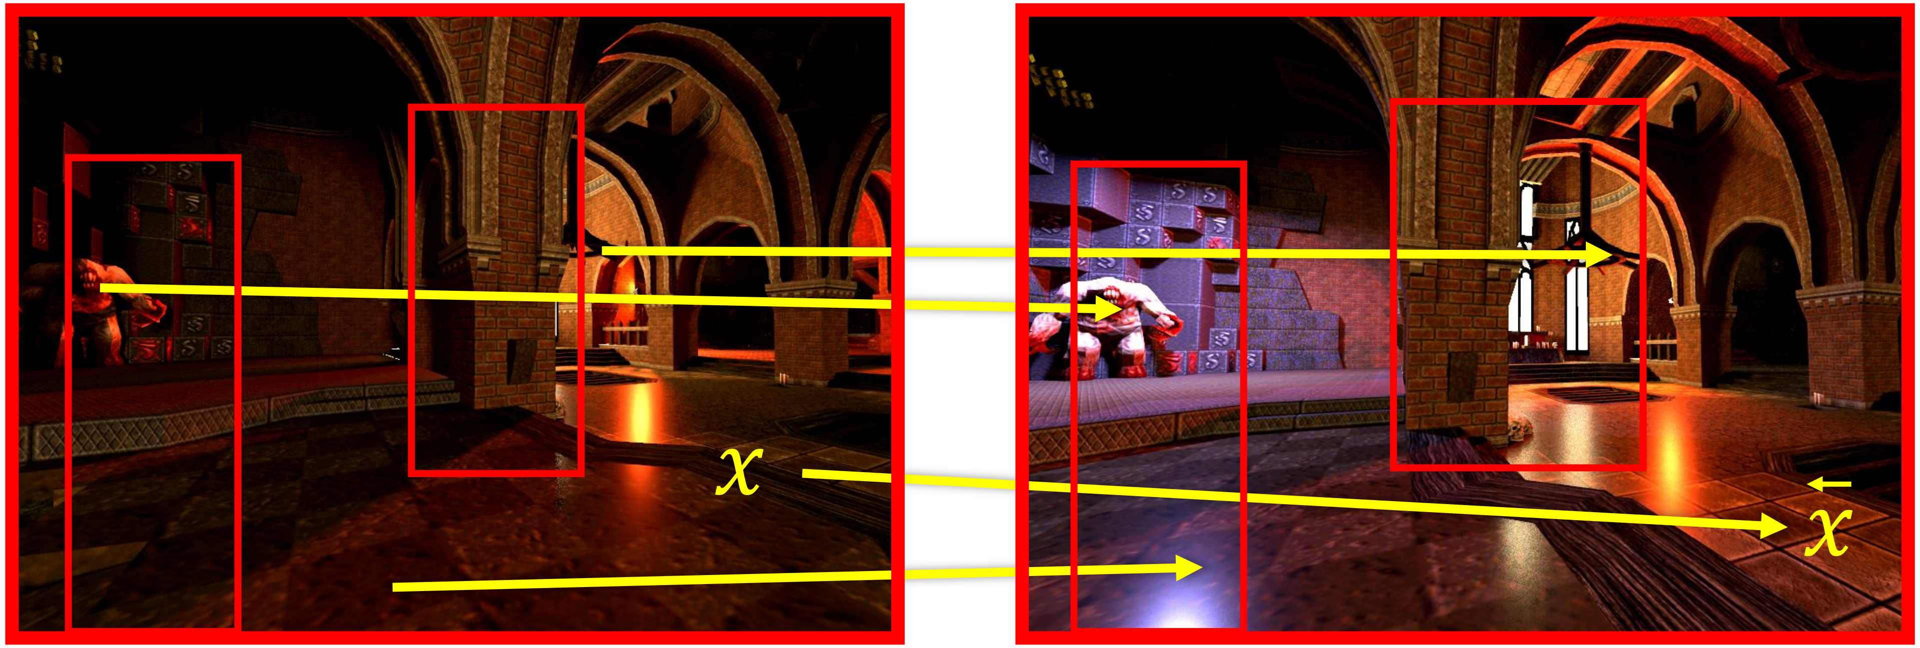

Current frame c_i Previous filtered frame $\hat{c}_{i-1}$

$$\hat{c}_i(x) = \alpha \cdot \boxed{c_i}\boxed{(x)} + (1-\alpha) \cdot \boxed{\hat{c}_{i-1}}\boxed{(\overleftarrow{x})}$$

Adaptive Temporal Filtering

$$\hat{c}_i(x) = \boxed{\alpha} \cdot c_i(x) + \boxed{(1-\alpha)} \cdot \hat{c}_{i-1}(\overleftarrow{x})$$

- Set α according to changes of the shading function
  - Moving shadows, glossy highlights, flickering light sources, …

- Make α per-pixel weight for local adaptivity

- **Need information about changes of shading (temporal gradient)**

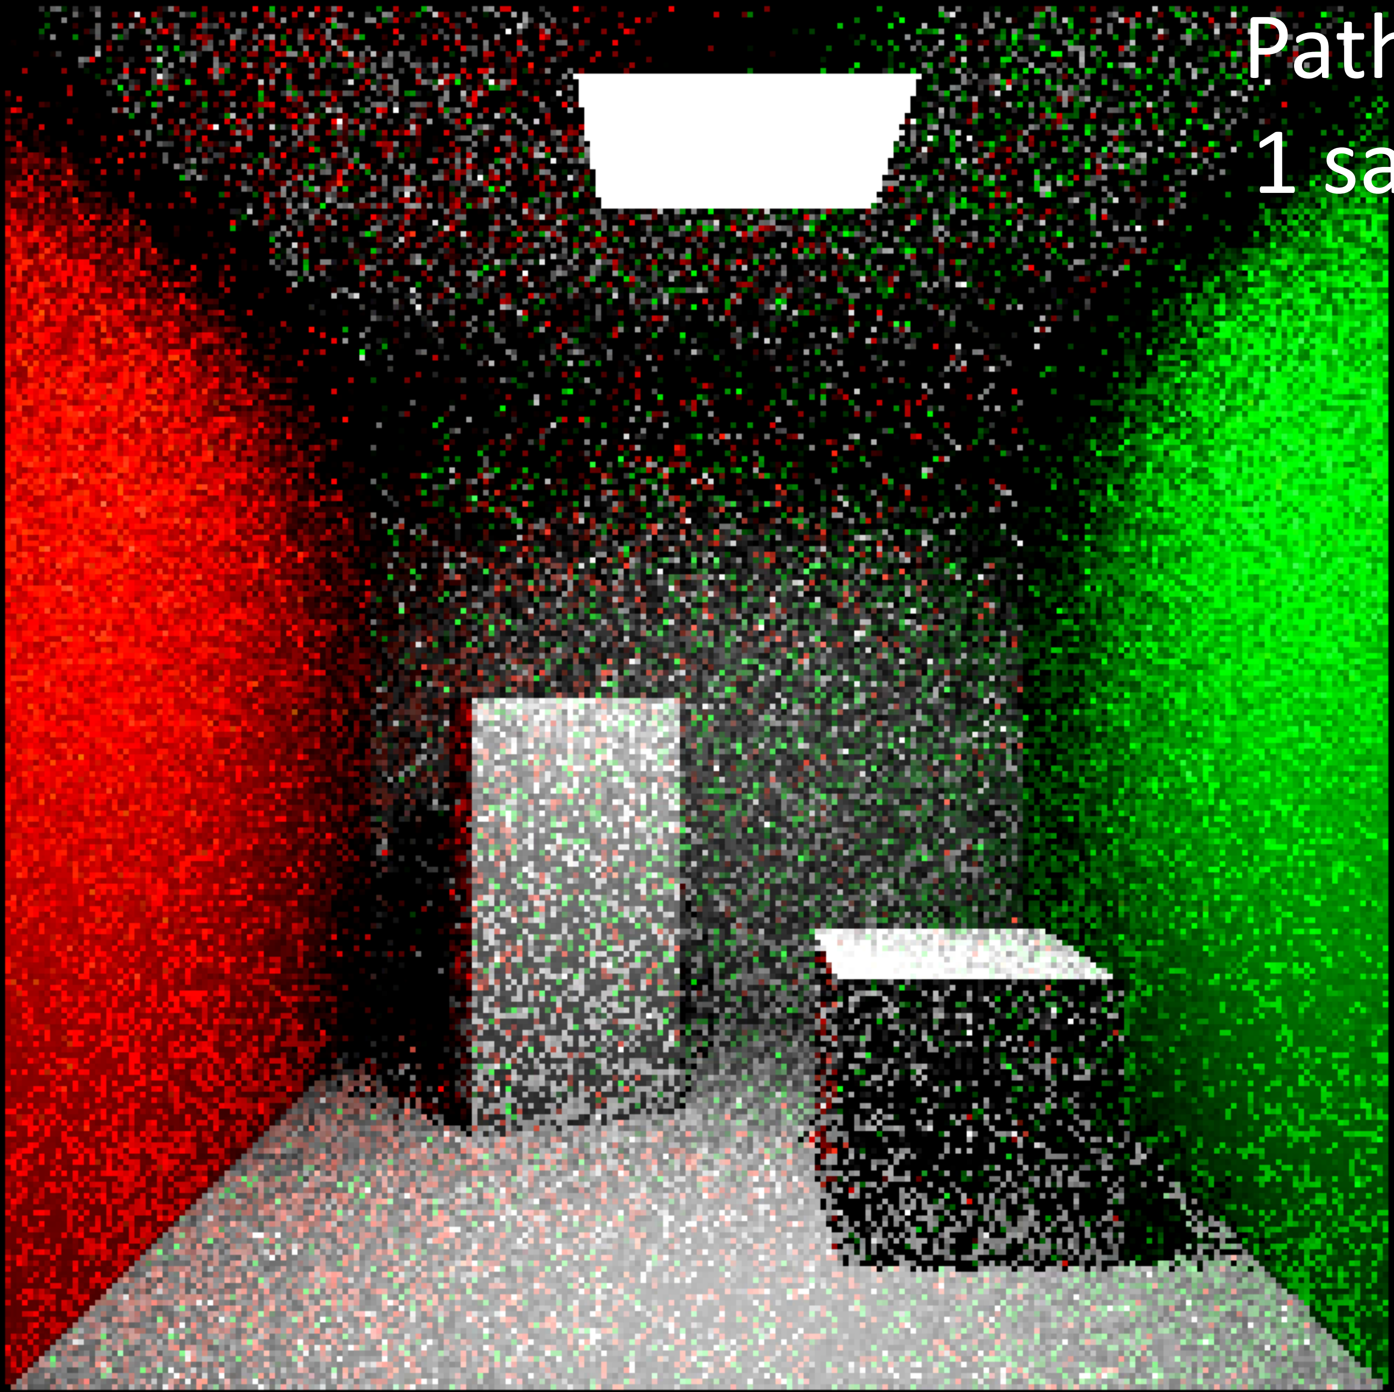

Path tracer output
1 sample per pixel

Difference of luminance
green positive
red negative

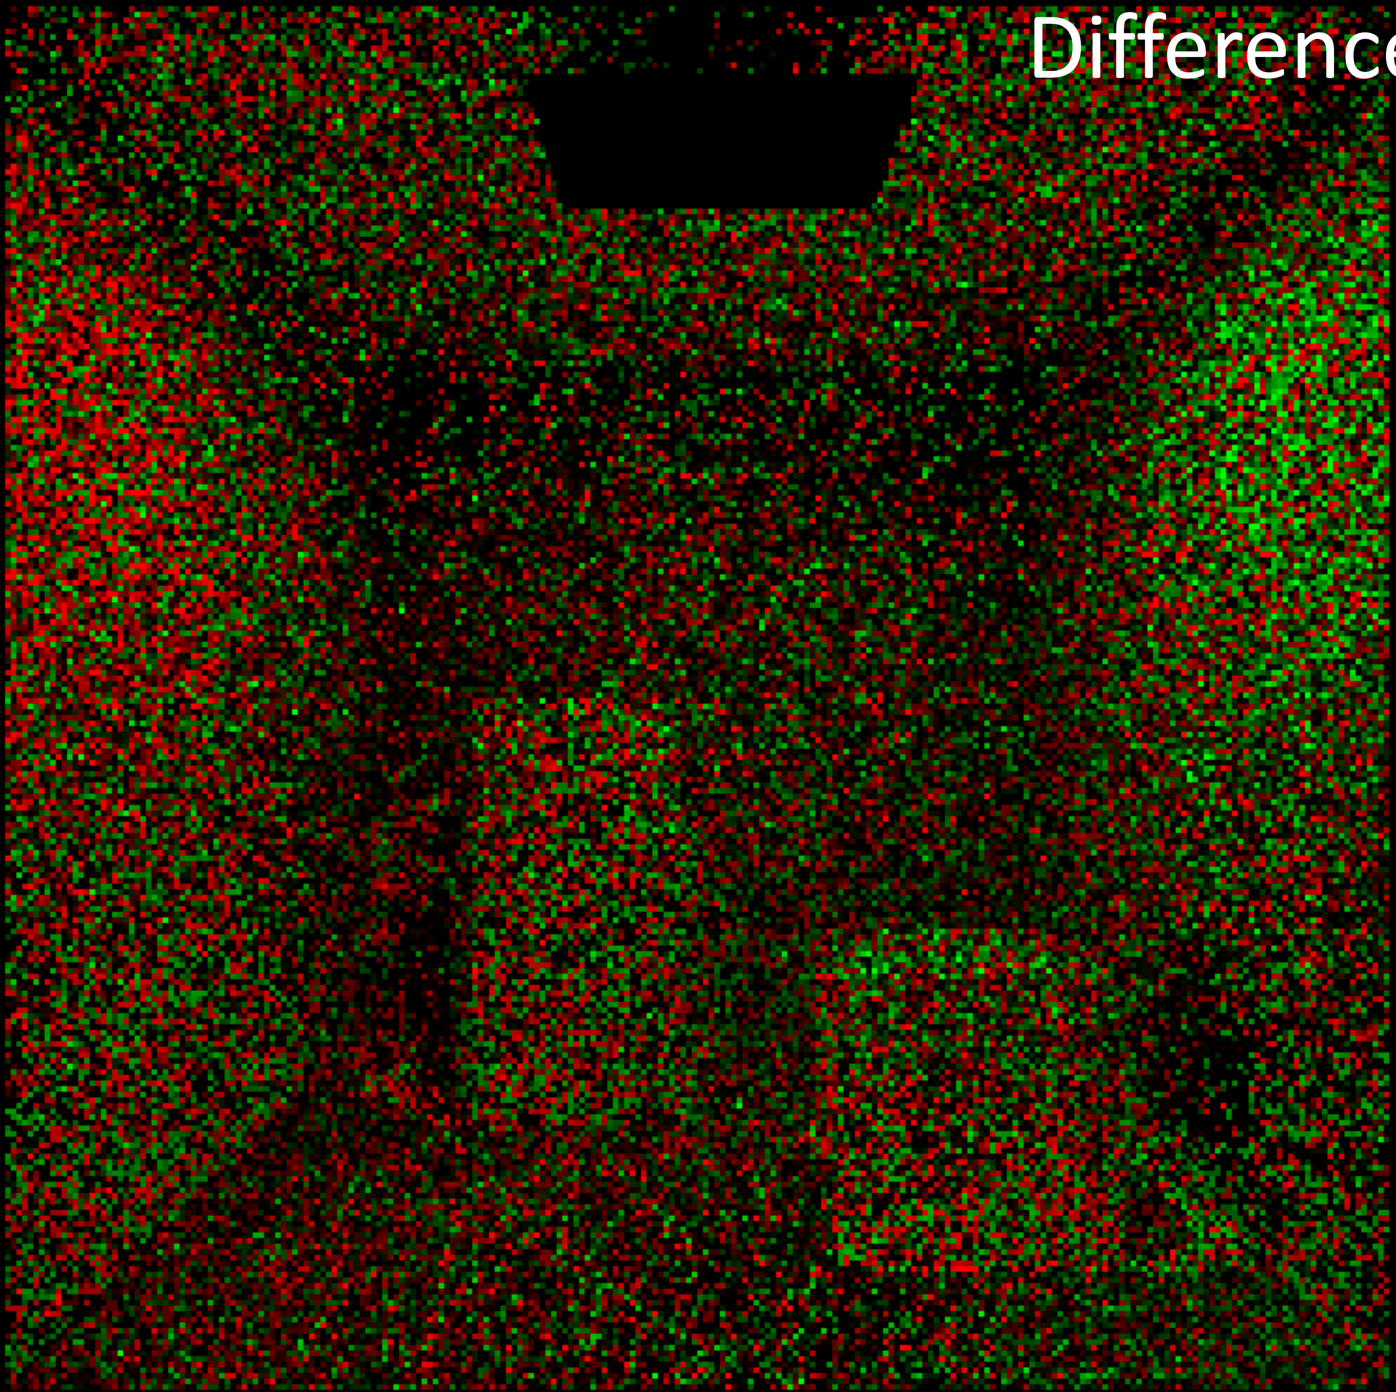

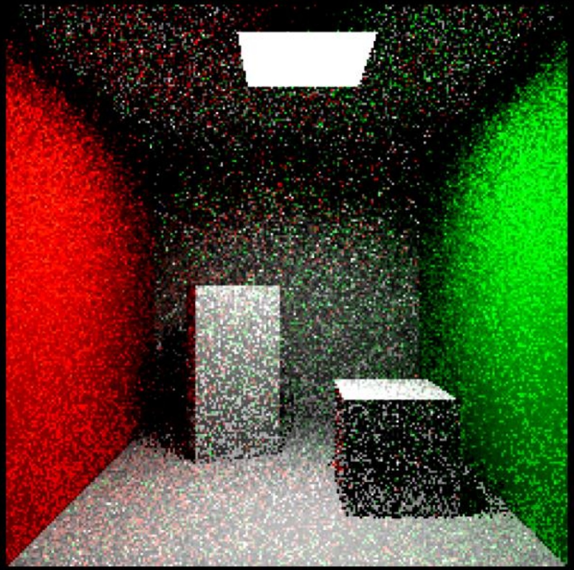

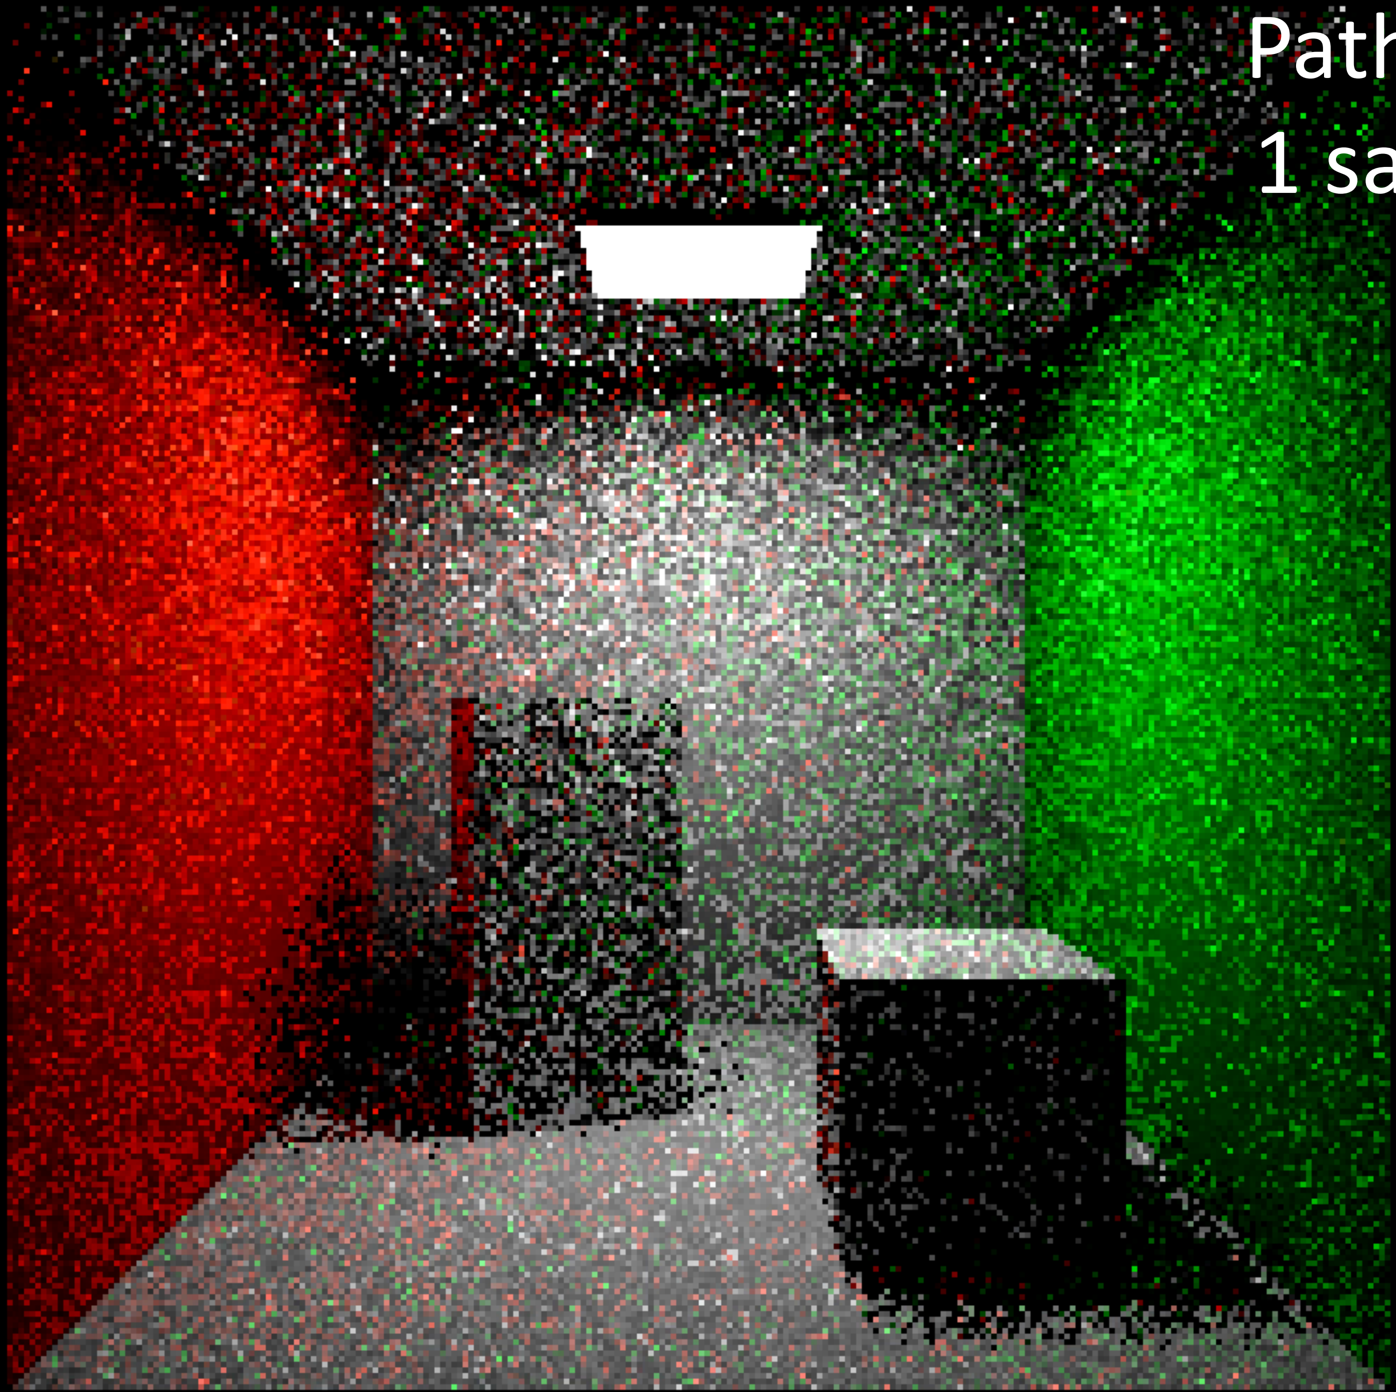
Path tracer output
1 sample per pixel

Difference of luminance
green positive
red negative

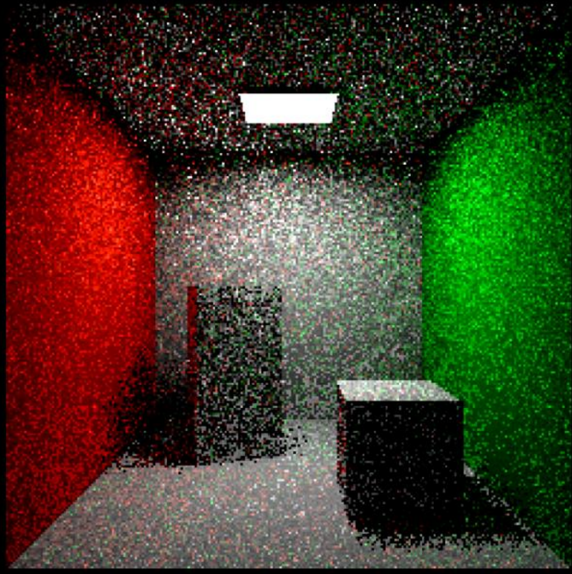

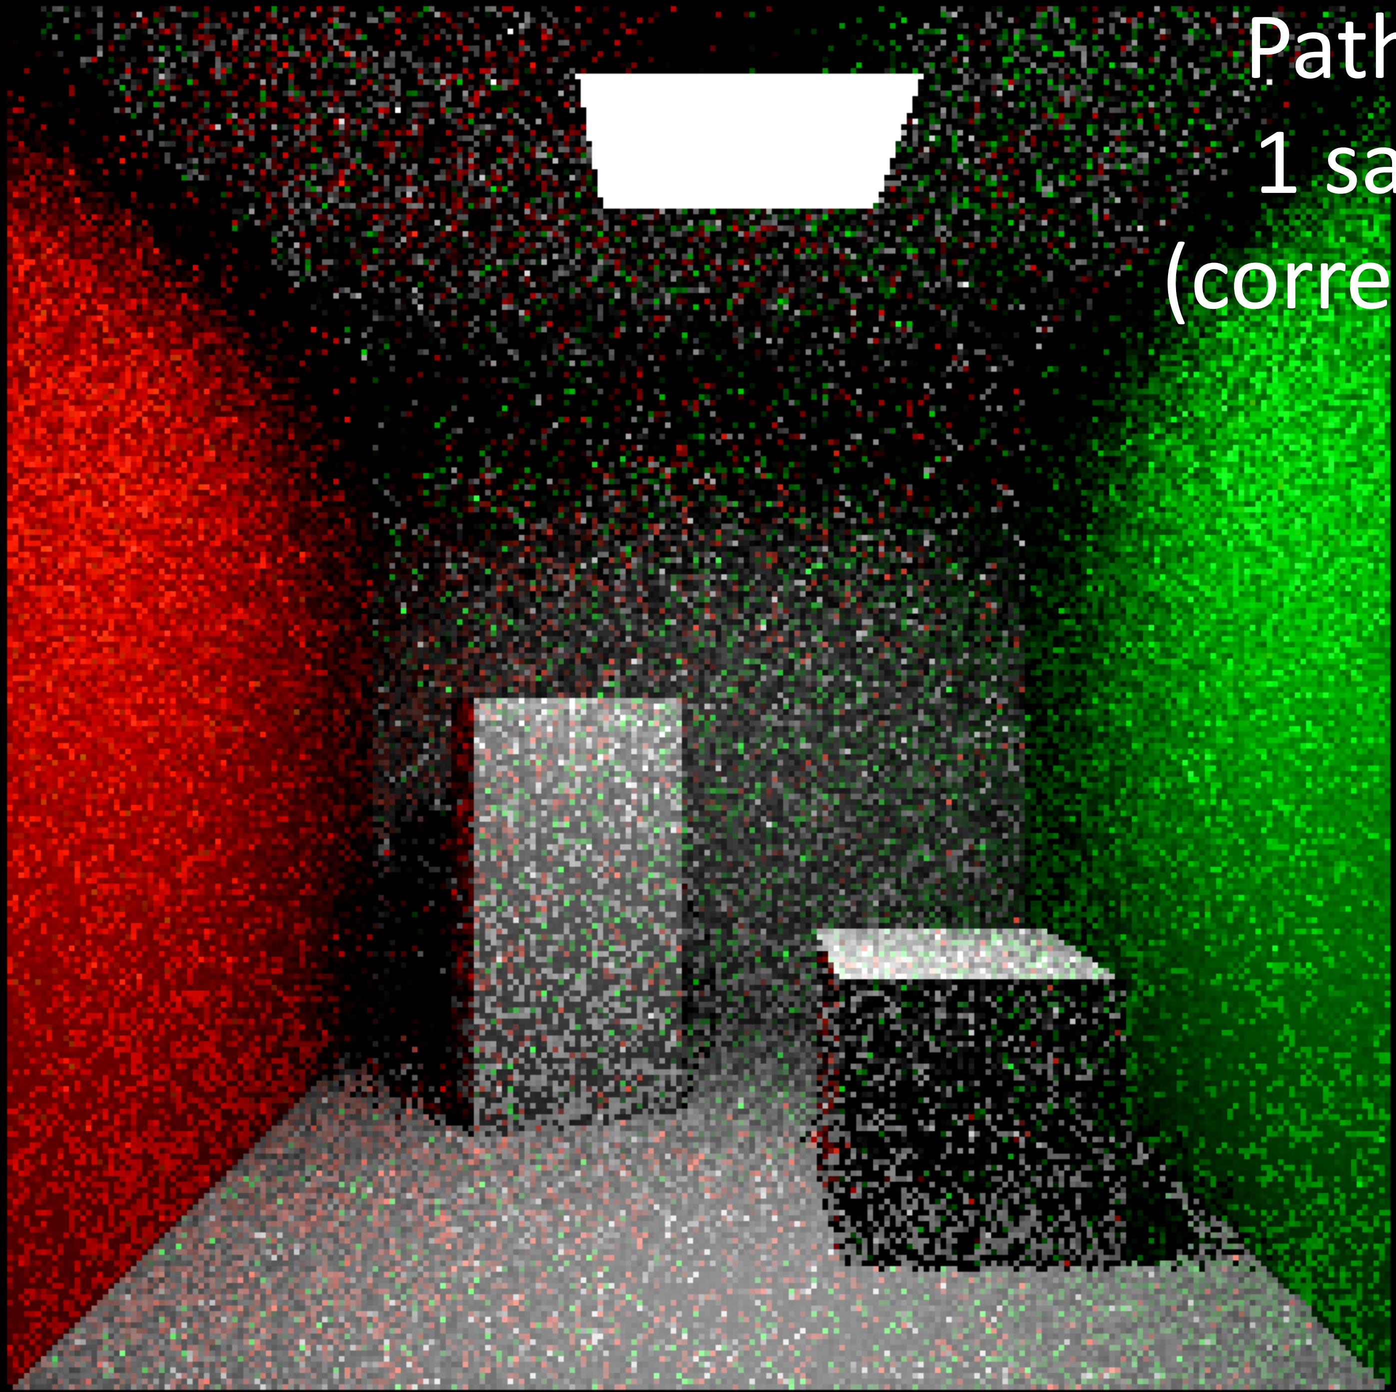

Path tracer output
1 sample per pixel
(correlated samples)

Difference of luminance
(correlated samples)
green positive
red negative

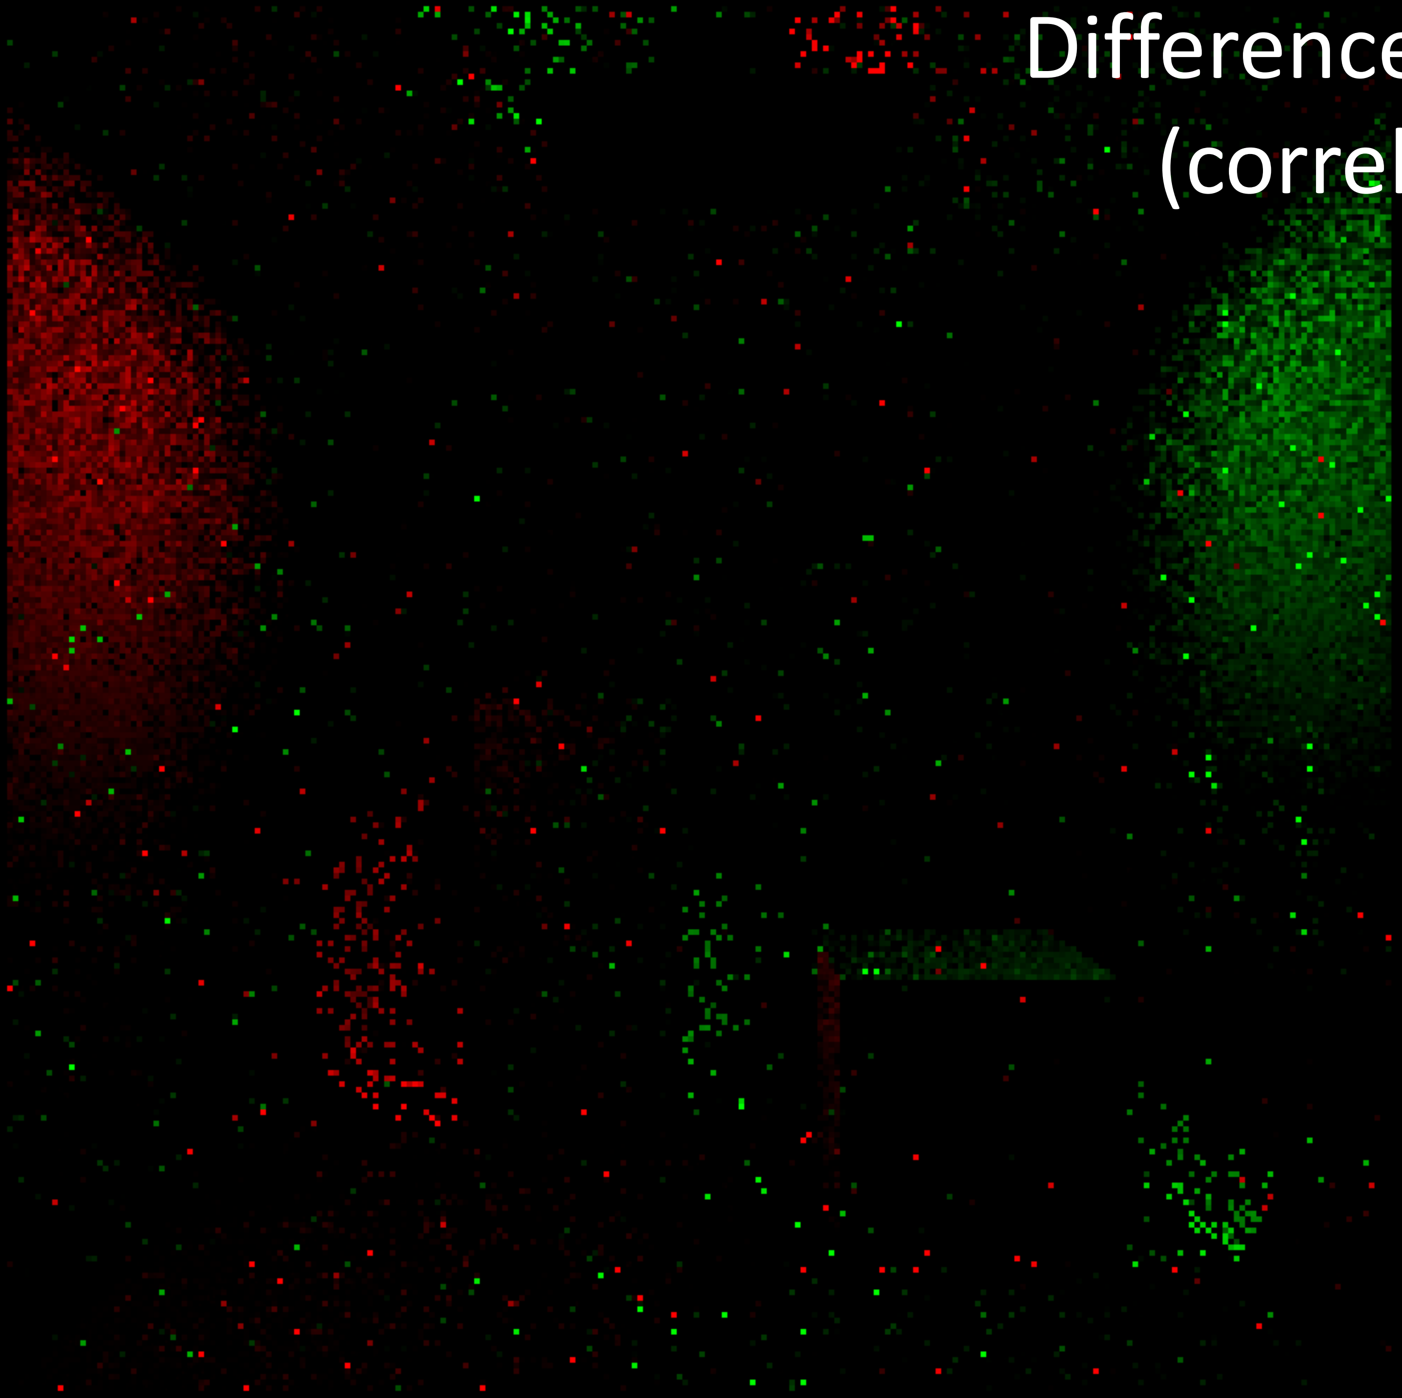

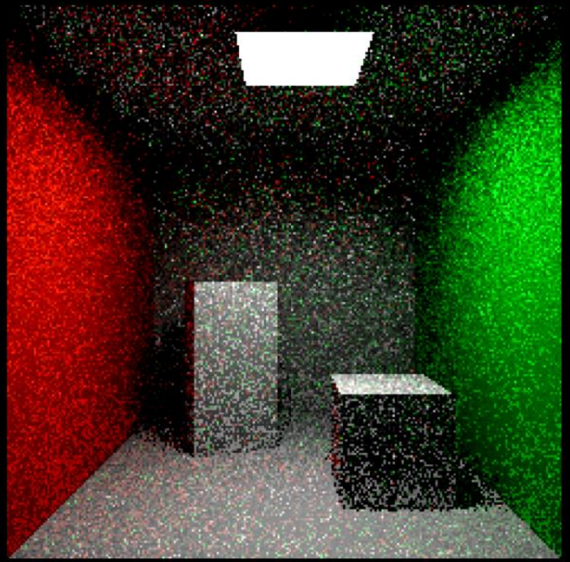

Adaptive temporal filter weight

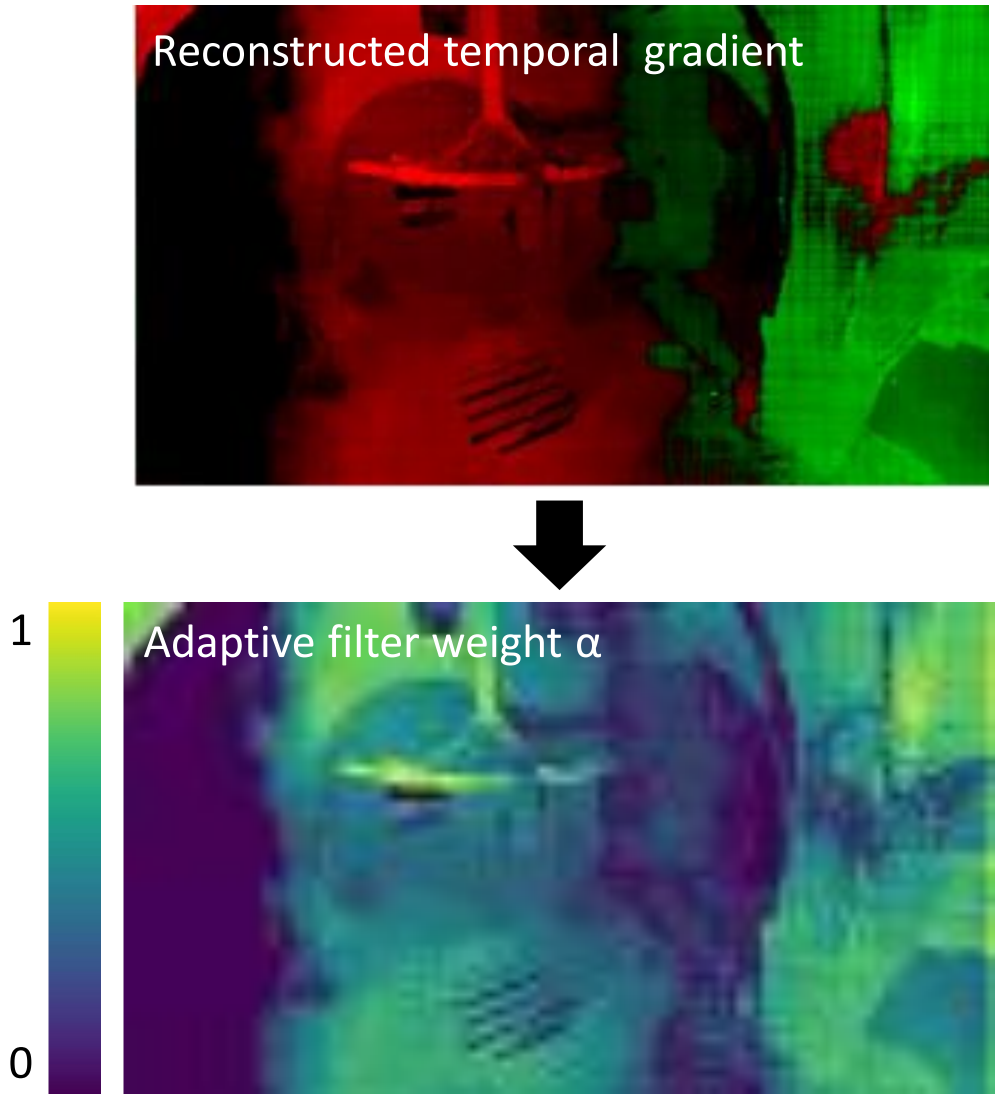

- Sample and reconstruct temporal gradient

- Change α according to relative rate of change https://www.youtube.com/watch?v=HP4Xwp0ETfc

Path Tracer

Treat each triangle of light meshes as individual area light

Light selection / sampling

Tried:

- Light hierarchy

Issues:

- Speed
- Inconsistent quality under animation

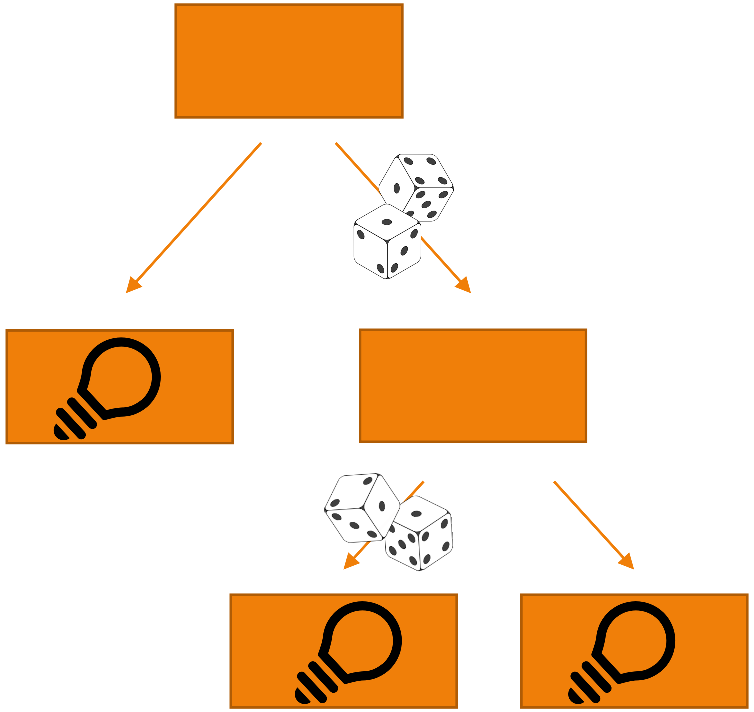

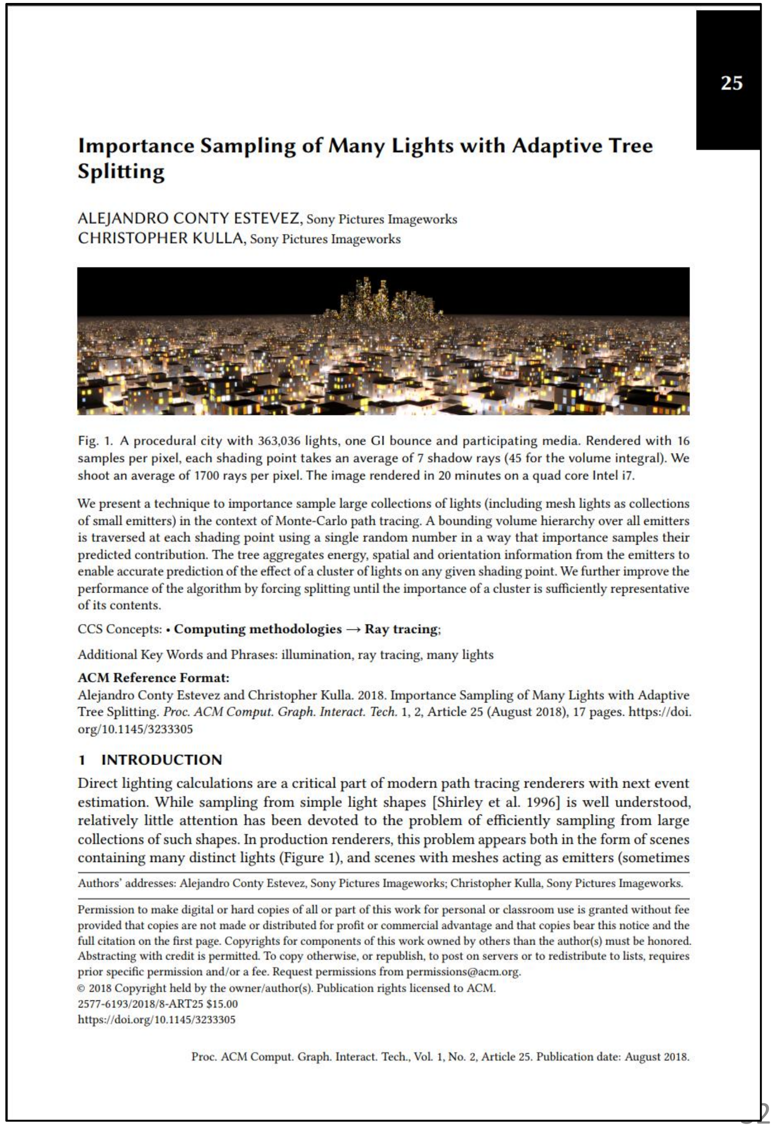

25

Importance Sampling of Many Lights with Adaptive Tree Splitting

ALEJANDRO CONTY ESTEVEZ, Sony Pictures Imageworks
CHRISTOPHER KULLA, Sony Pictures Imageworks

Fig. 1. A procedural city with 363,036 lights, one GI bounce and participating media. Rendered with 16 samples per pixel, each shading point takes an average of 7 shadow rays (45 for the volume integral). We shoot an average of 1700 rays per pixel. The image rendered in 20 minutes on a quad core Intel i7.

We present a technique to importance sample large collections of lights (including mesh lights as collections of small emitters) in the context of Monte-Carlo path tracing. A bounding volume hierarchy over all emitters is traversed at each shading point using a single random number in a way that importance samples their predicted contribution. The tree aggregates energy, spatial and orientation information from the emitters to enable accurate prediction of the effect of a cluster of lights on any given shading point. We further improve the performance of the algorithm by forcing splitting until the importance of a cluster is sufficiently representative of its contents.

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

Additional Key Words and Phrases: illumination, ray tracing, many lights

ACM Reference Format:
Alejandro Conty Estevez and Christopher Kulla. 2018. Importance Sampling of Many Lights with Adaptive Tree Splitting. *Proc. ACM Comput. Graph. Interact. Tech.* 1, 2, Article 25 (August 2018), 17 pages. https://doi.org/10.1145/3233305

1 INTRODUCTION

Direct lighting calculations are a critical part of modern path tracing renderers with next event estimation. While sampling from simple light shapes [Shirley et al. 1996] is well understood, relatively little attention has been devoted to the problem of efficiently sampling from large collections of such shapes. In production renderers, this problem appears both in the form of scenes containing many distinct lights (Figure 1), and scenes with meshes acting as emitters (sometimes

Authors' addresses: Alejandro Conty Estevez, Sony Pictures Imageworks; Christopher Kulla, Sony Pictures Imageworks.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.
2577-6193/2018/8-ART25 $15.00
https://doi.org/10.1145/3233305

Proc. ACM Comput. Graph. Interact. Tech., Vol. 1, No. 2, Article 25. Publication date: August 2018.

Light selection / sampling for static lights

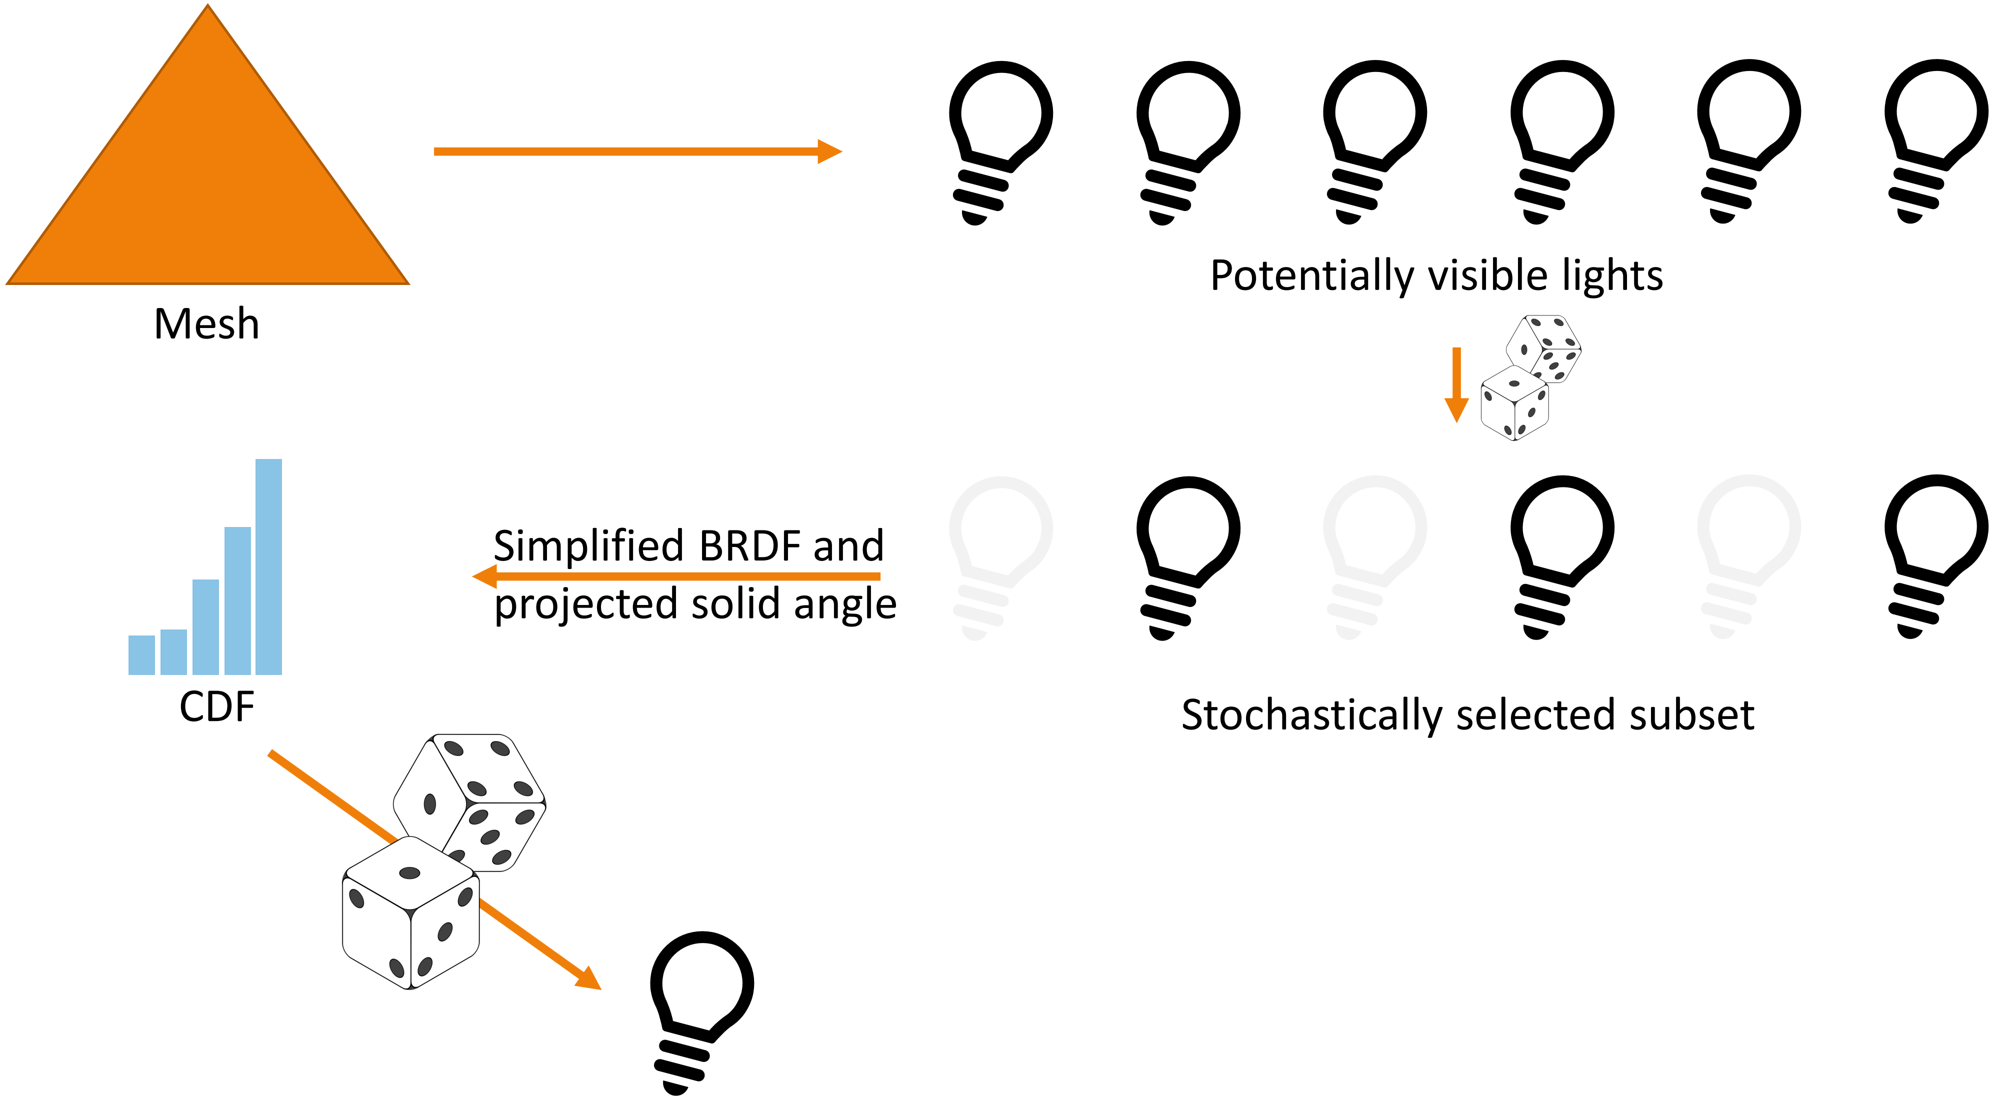

Light selection / sampling for dynamic lights

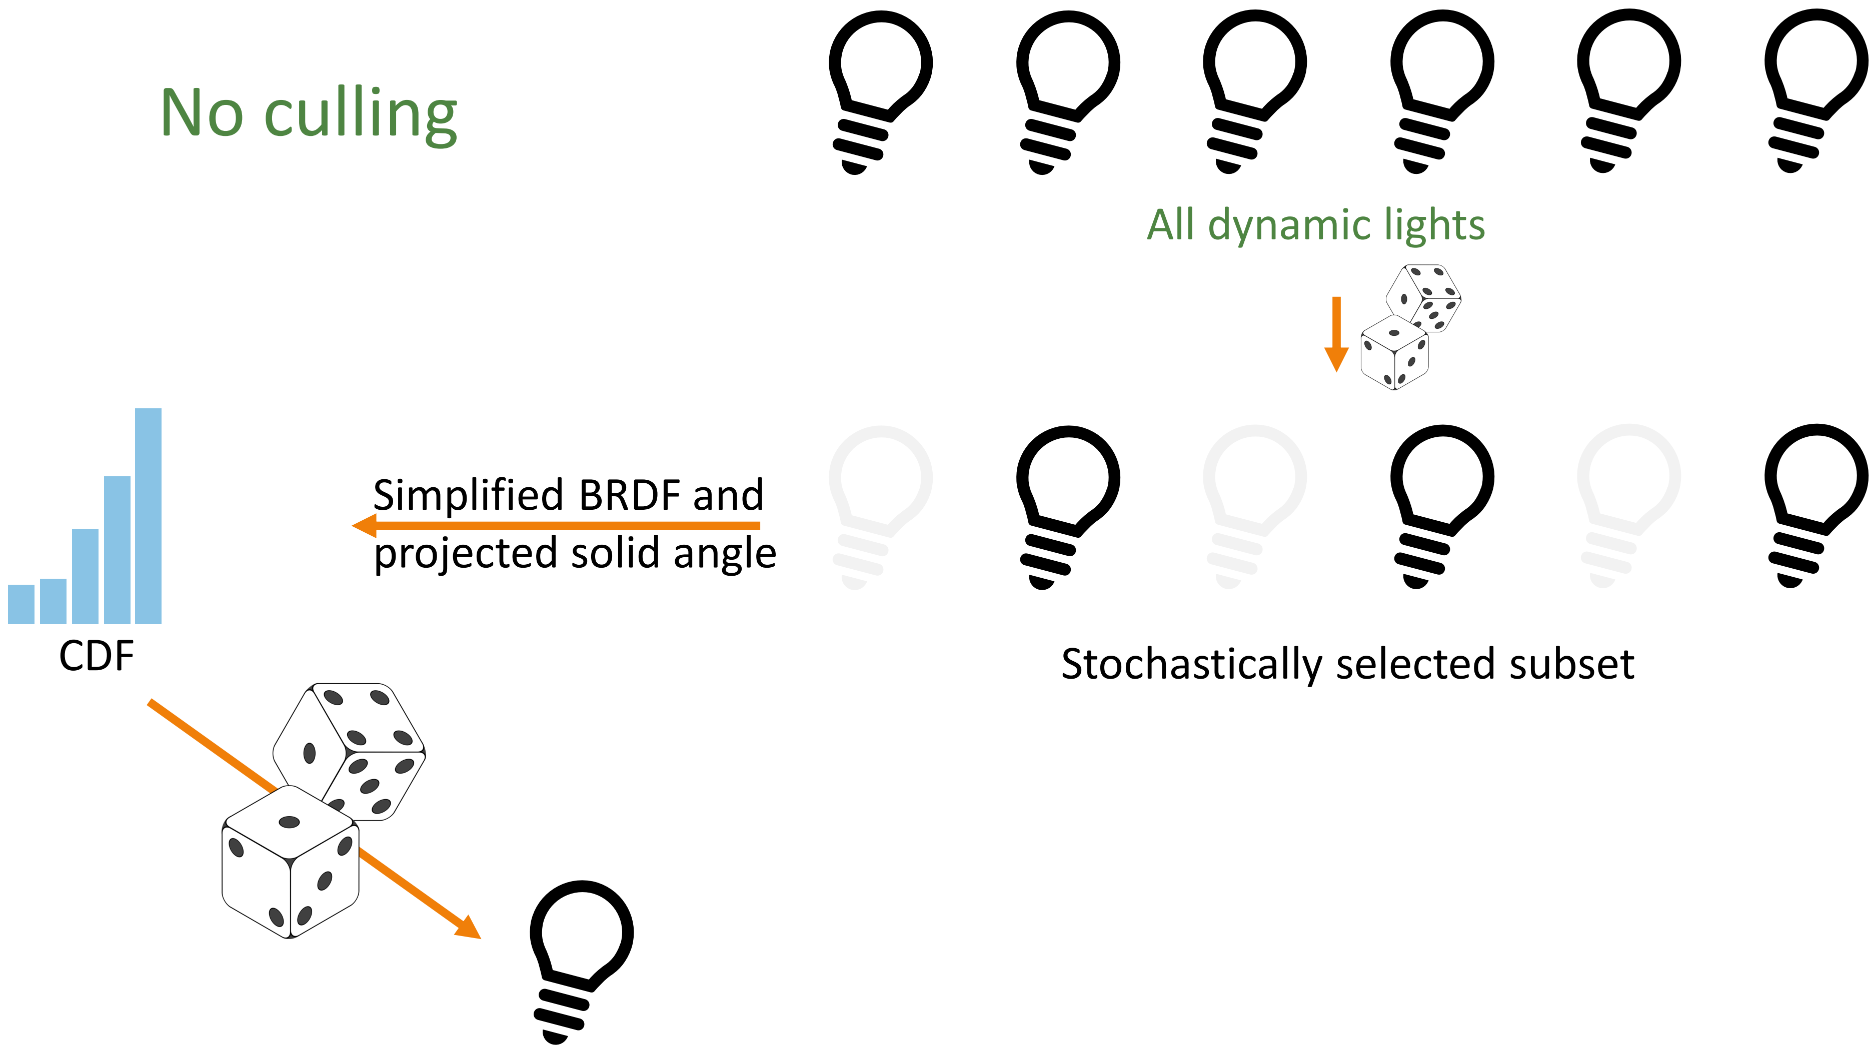

Light selection

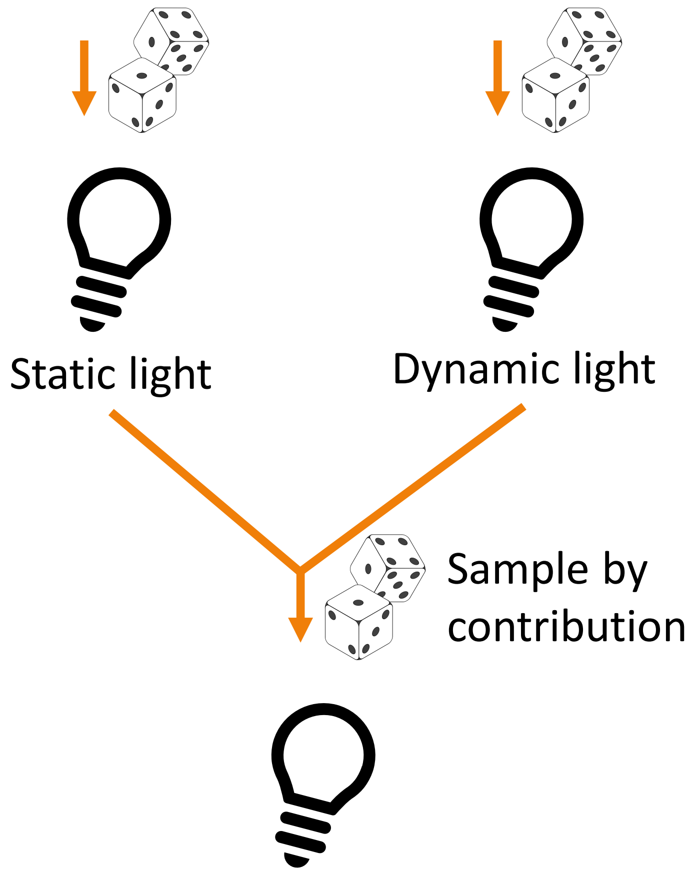

Path tracer

- One path per pixel
- One indirect bounce
- Two shadow rays

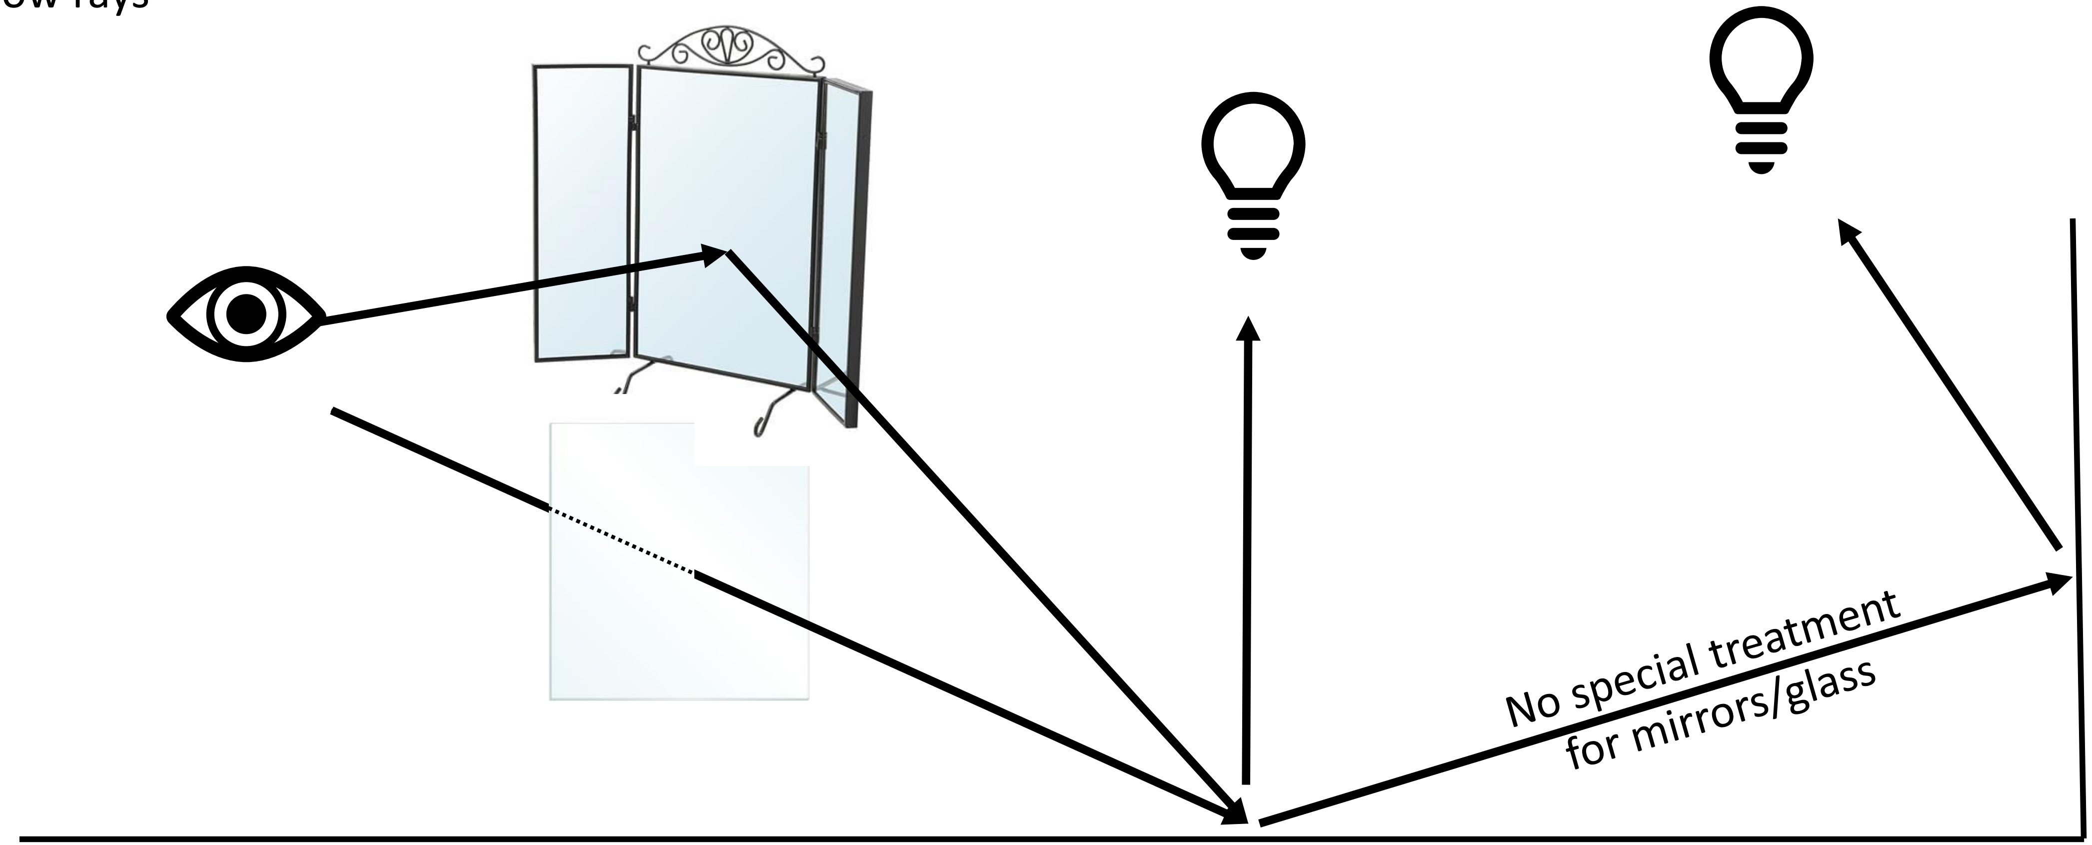

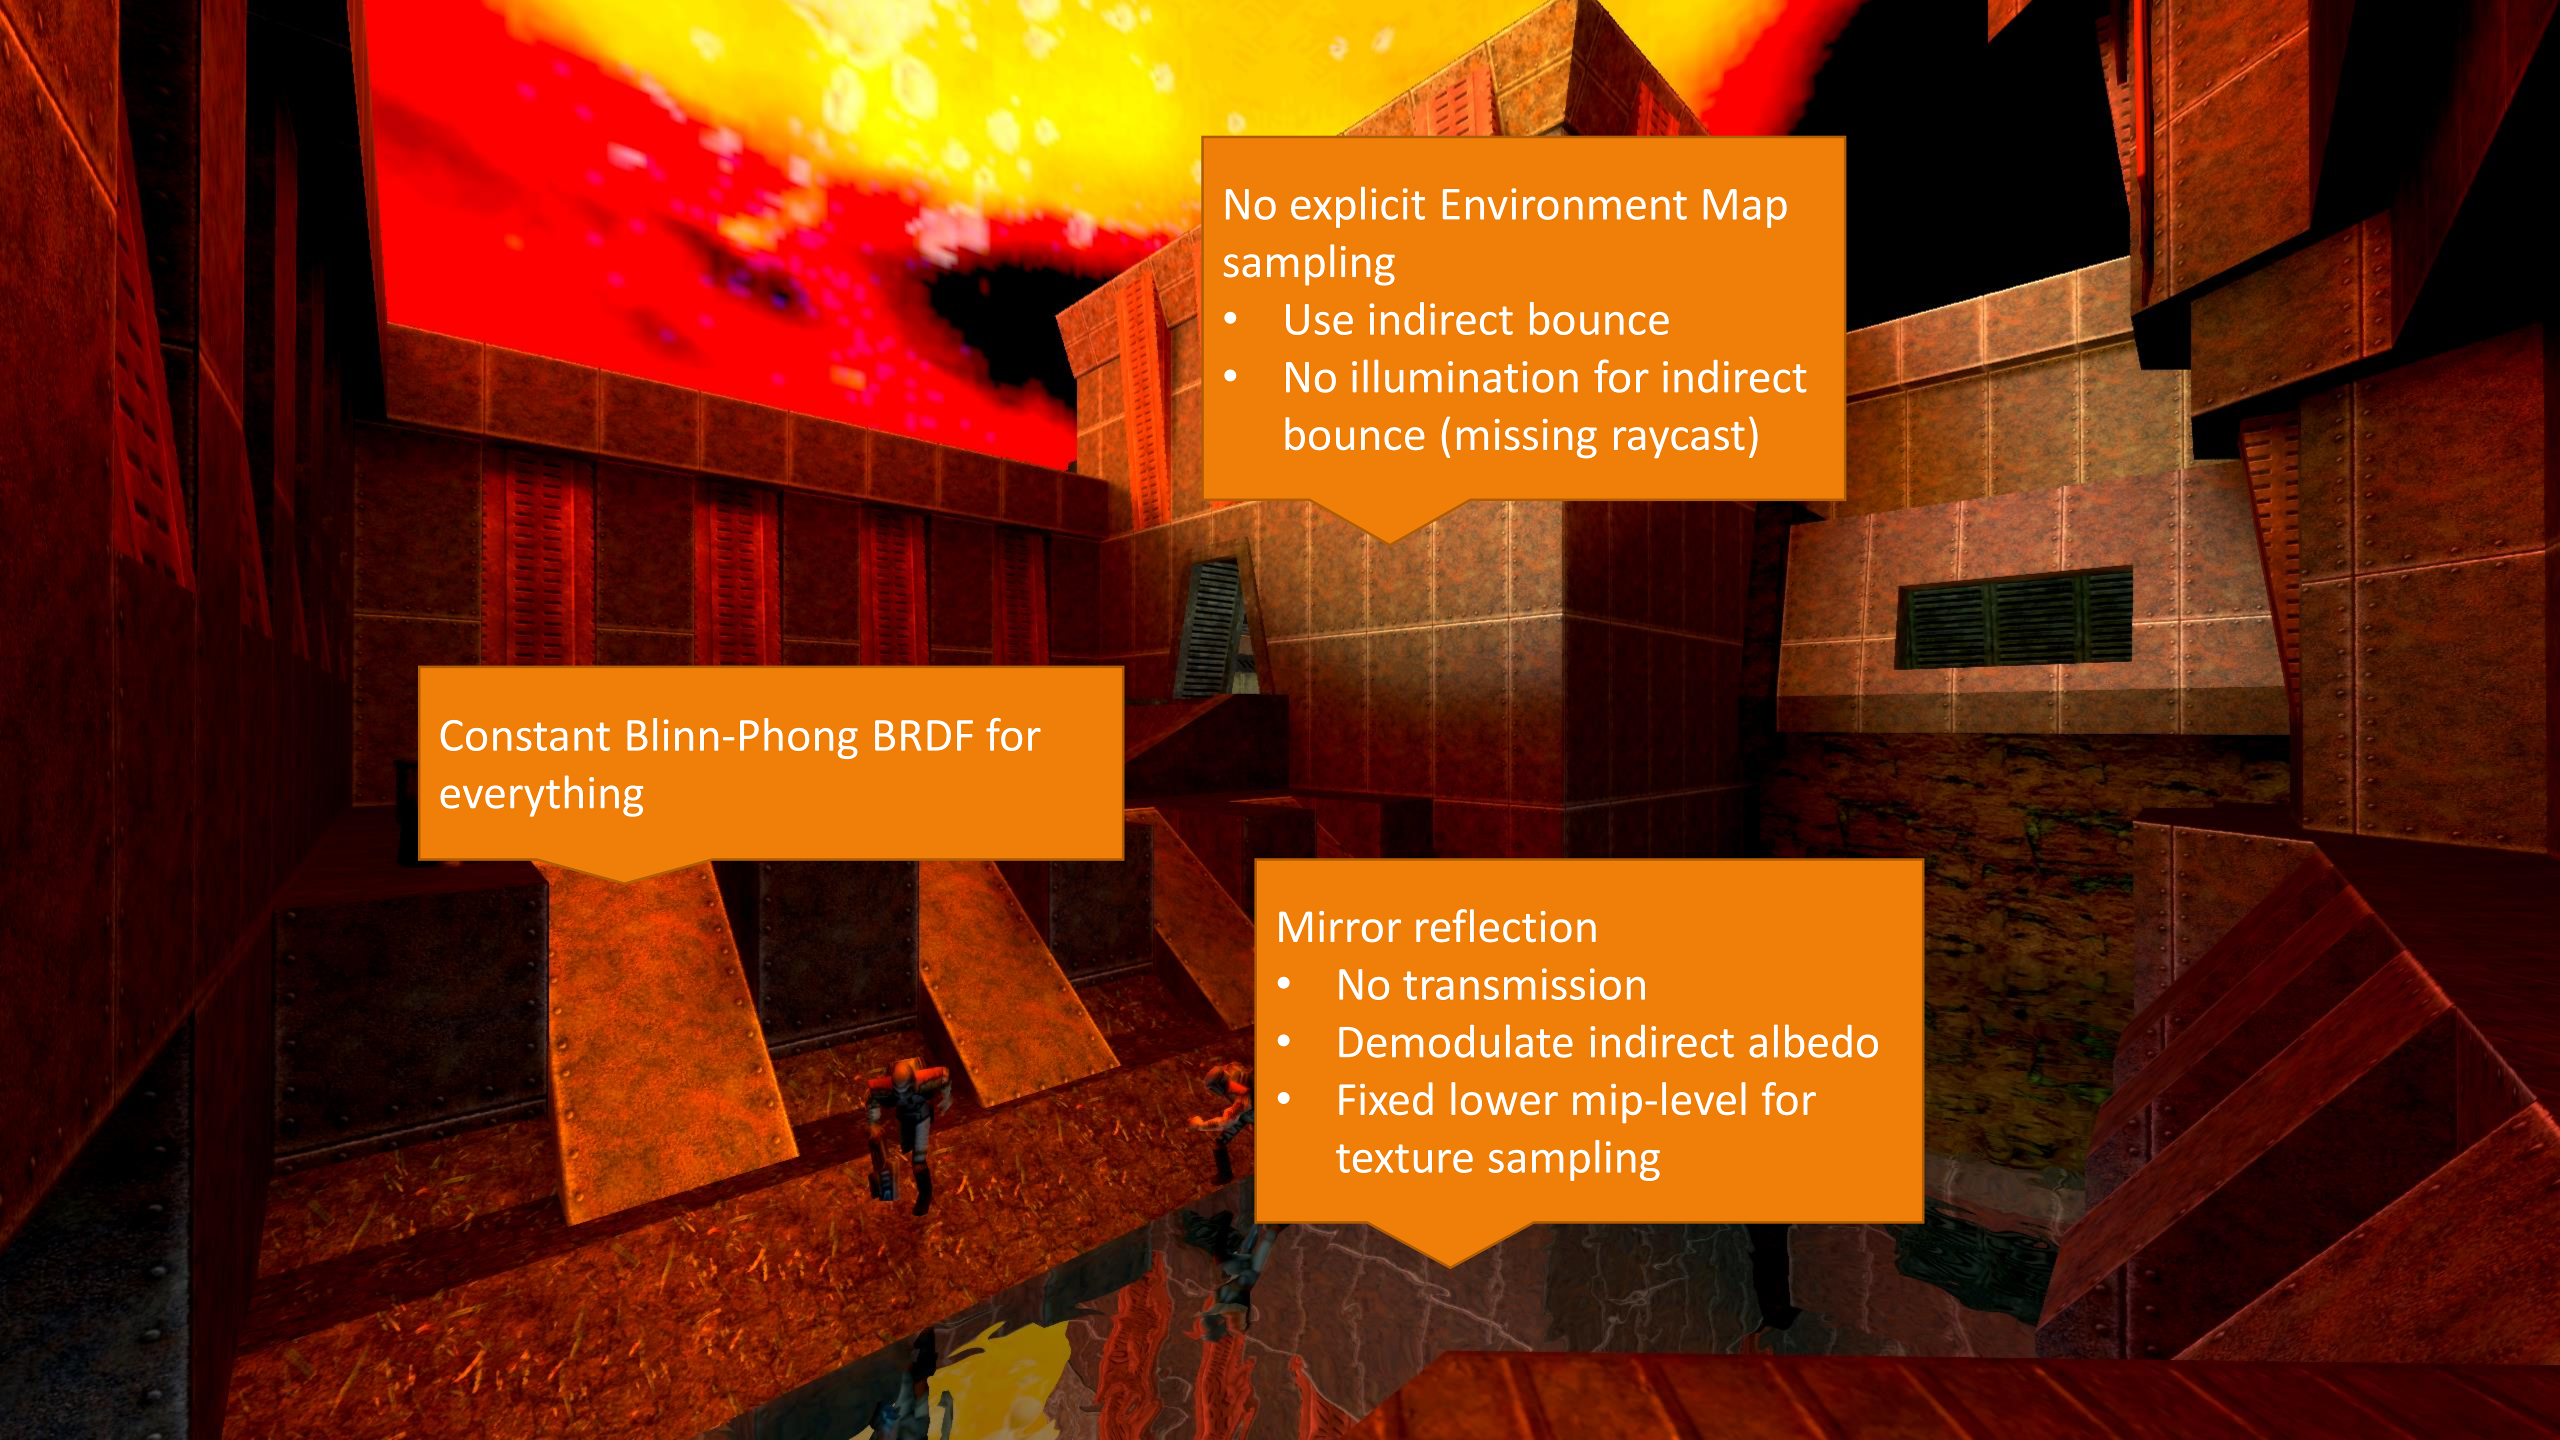
No explicit Environment Map sampling
• Use indirect bounce
• No illumination for indirect bounce (missing raycast)
Constant Blinn-Phong BRDF for everything
Mirror reflection
• No transmission
• Demodulate indirect albedo
• Fixed lower mip-level for texture sampling

White noise

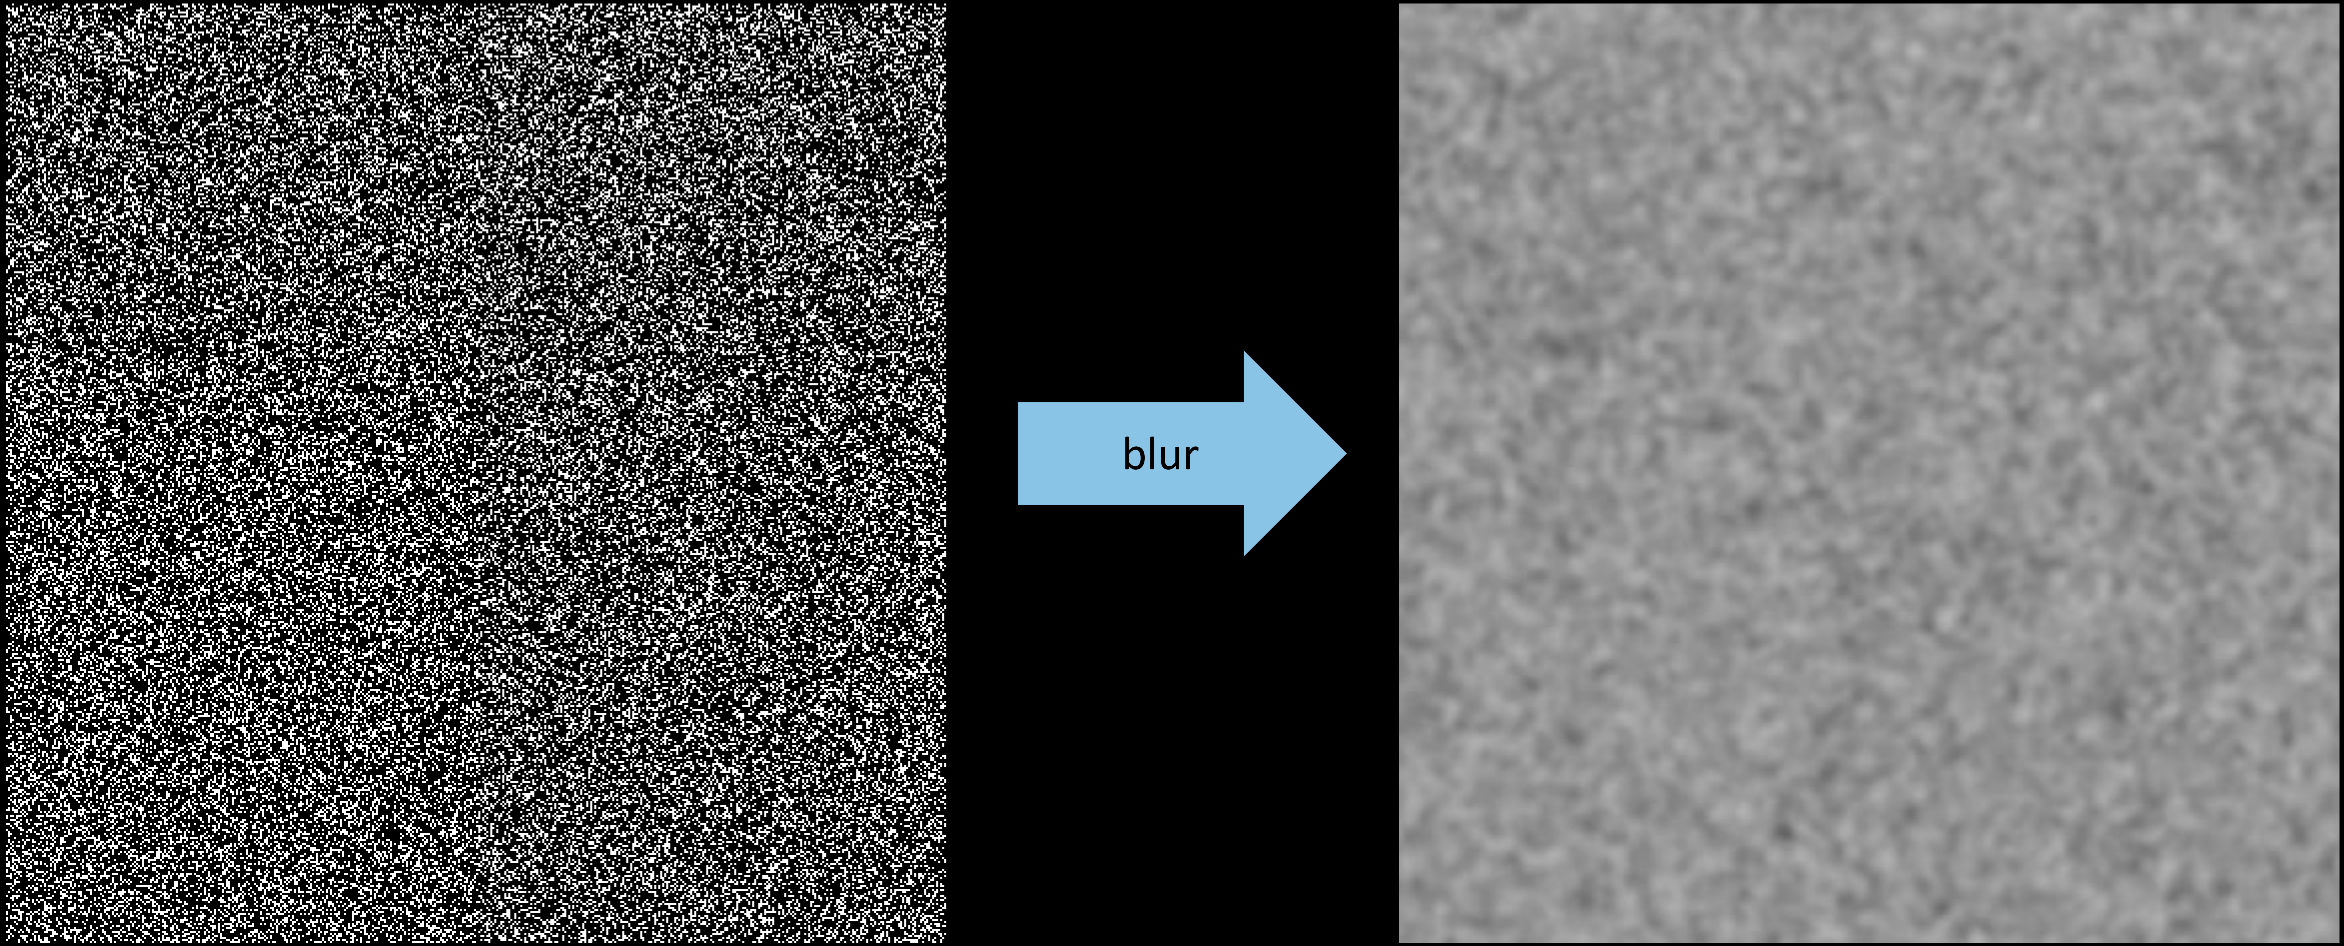

Source: http://momentsingraphics.de/?p=127

Blue noise

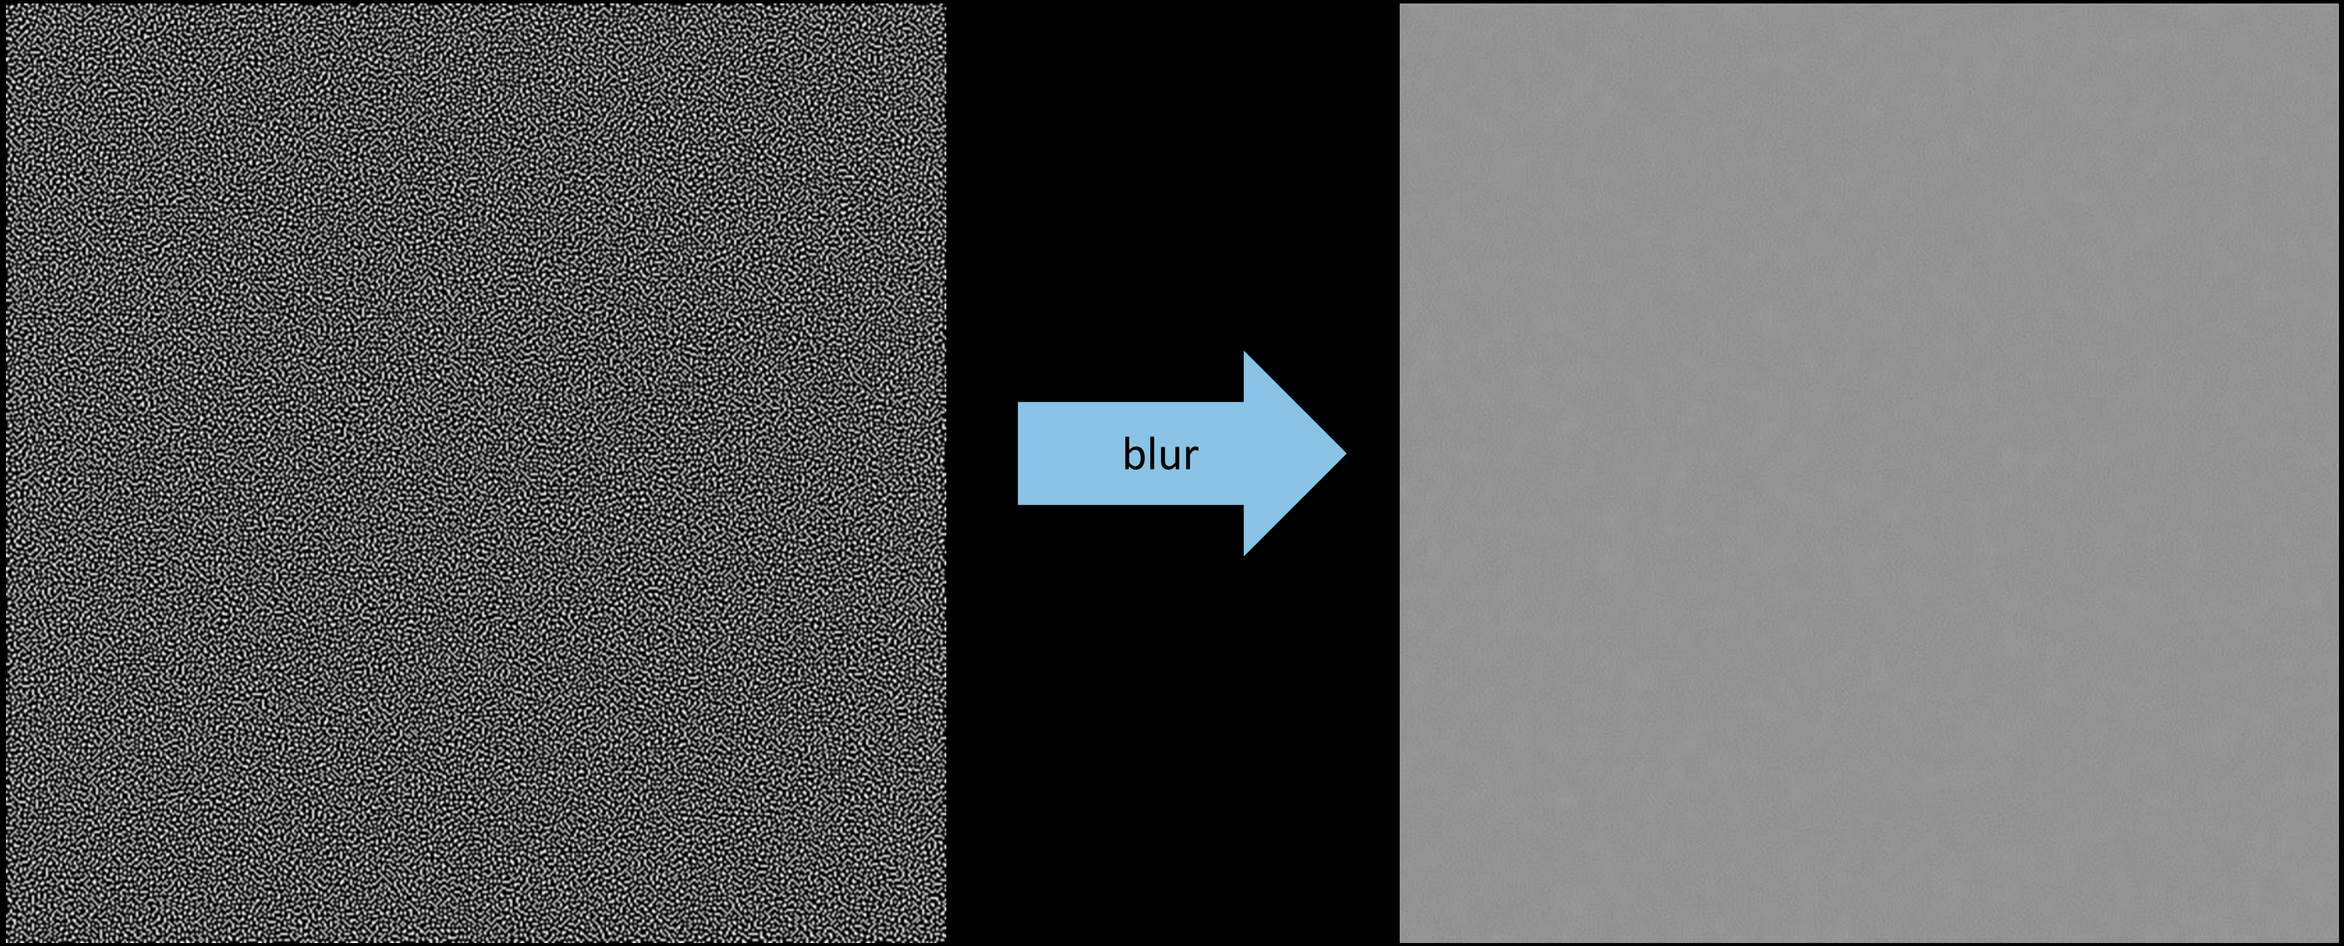

Source: http://momentsingraphics.de/?p=127

Magnitude of Fourier Transform

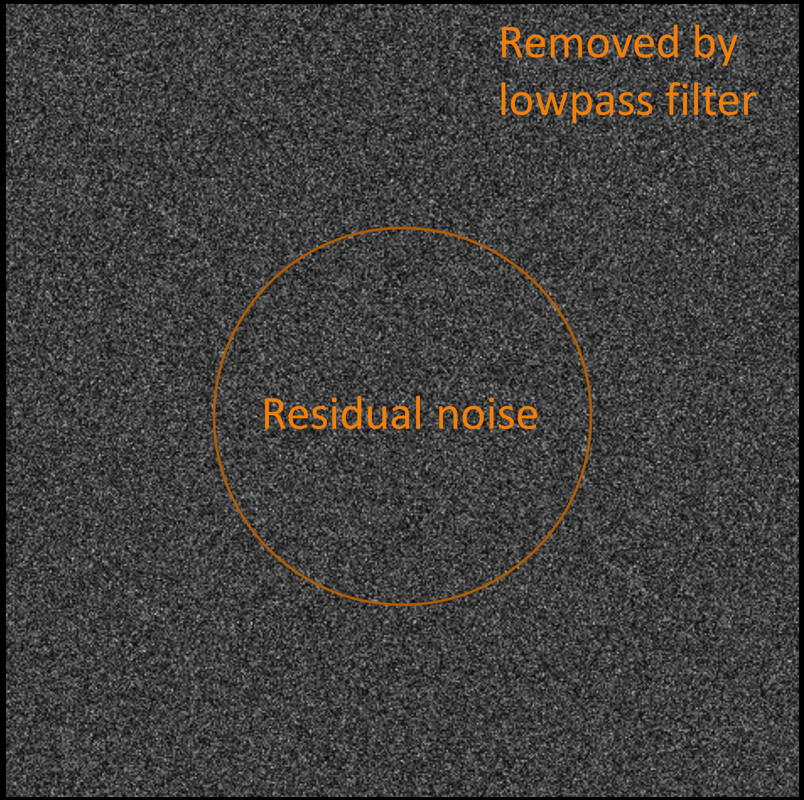

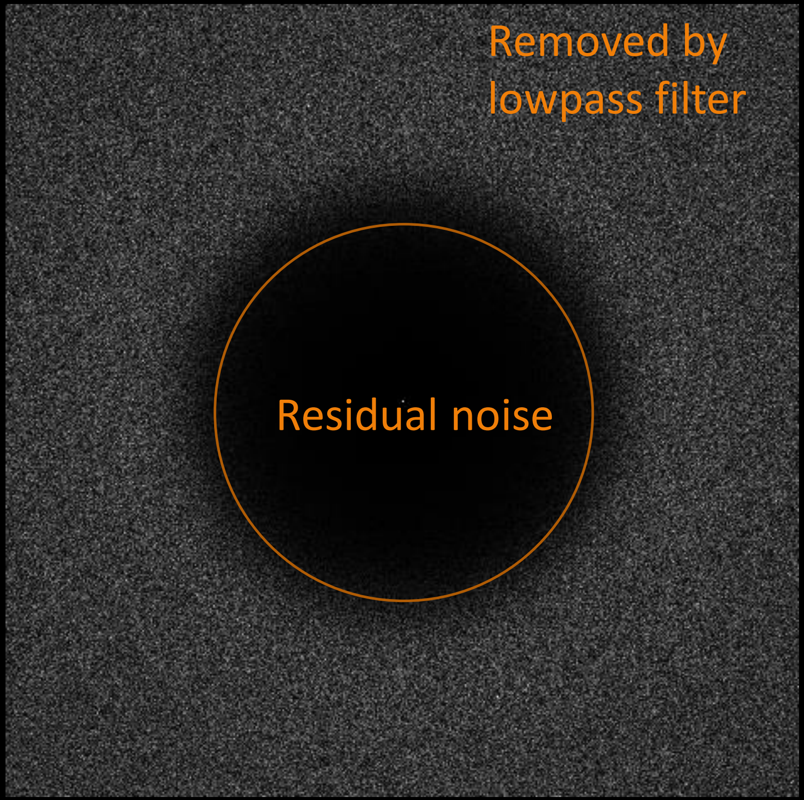

Source: http://momentsingraphics.de/?p=127

Implementation details

Acceleration structures

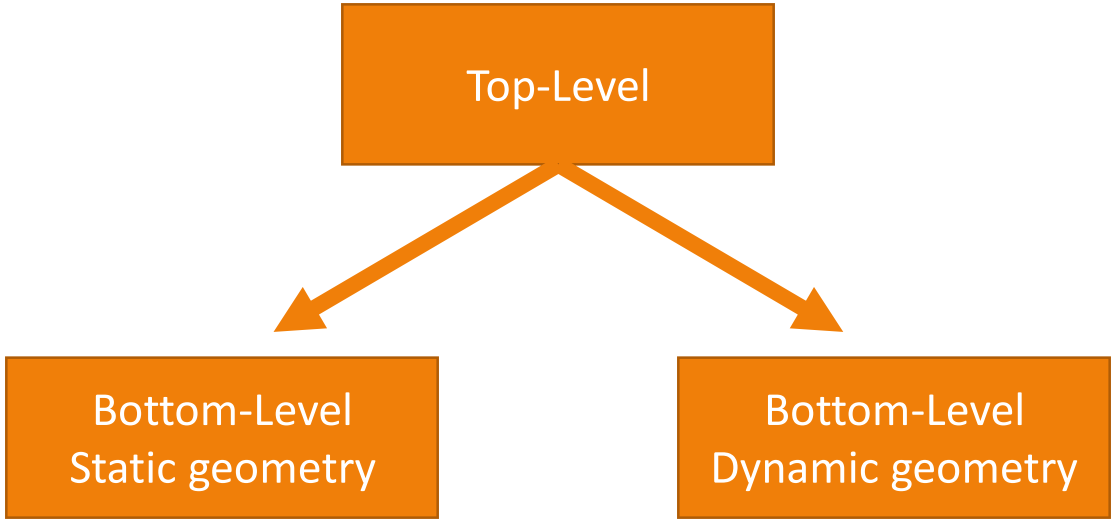

Built once on map-load

Rebuilt from scratch per frame

Forward / Backward projection

- Required for Adaptive Temporal Filtering
- Visibility buffer for forward projection
- Map instances between frames

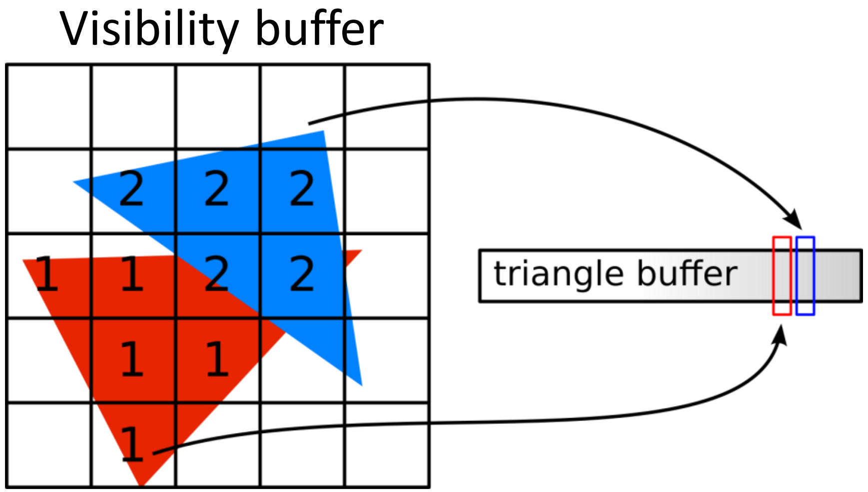

Conclusion

- Real-time path tracing is possible (in the near future)

- Transition difficult
  - Random access to everything
  - Tweaking of assets

- Need more research specifically tailored towards real-time rendering
  - Fast and robust importance sampling
  - Denoising

Thanks!

Addis Dittebrandt
Alisa Jung
Anton Kaplanyan
Christoph Peters
Florian Reibold
Johannes Hanika
Stephan Bergmann
Tobias Zirr

NVIDIA
id Software

http://brechpunkt.de/q2vkpt
https://github.com/cschied/q2vkpt
schied@brechpunkt.de
@c_schied