Interactive Stable Ray Tracing

Alessandro Dal Corso[1,2], Marco Salvi[1], Craig Kolb[1],
Jeppe Revall Frisvad[2], Aaron Lefohn[1], David Luebke[1]

[1]NVIDIA, [2]Technical University of Denmark

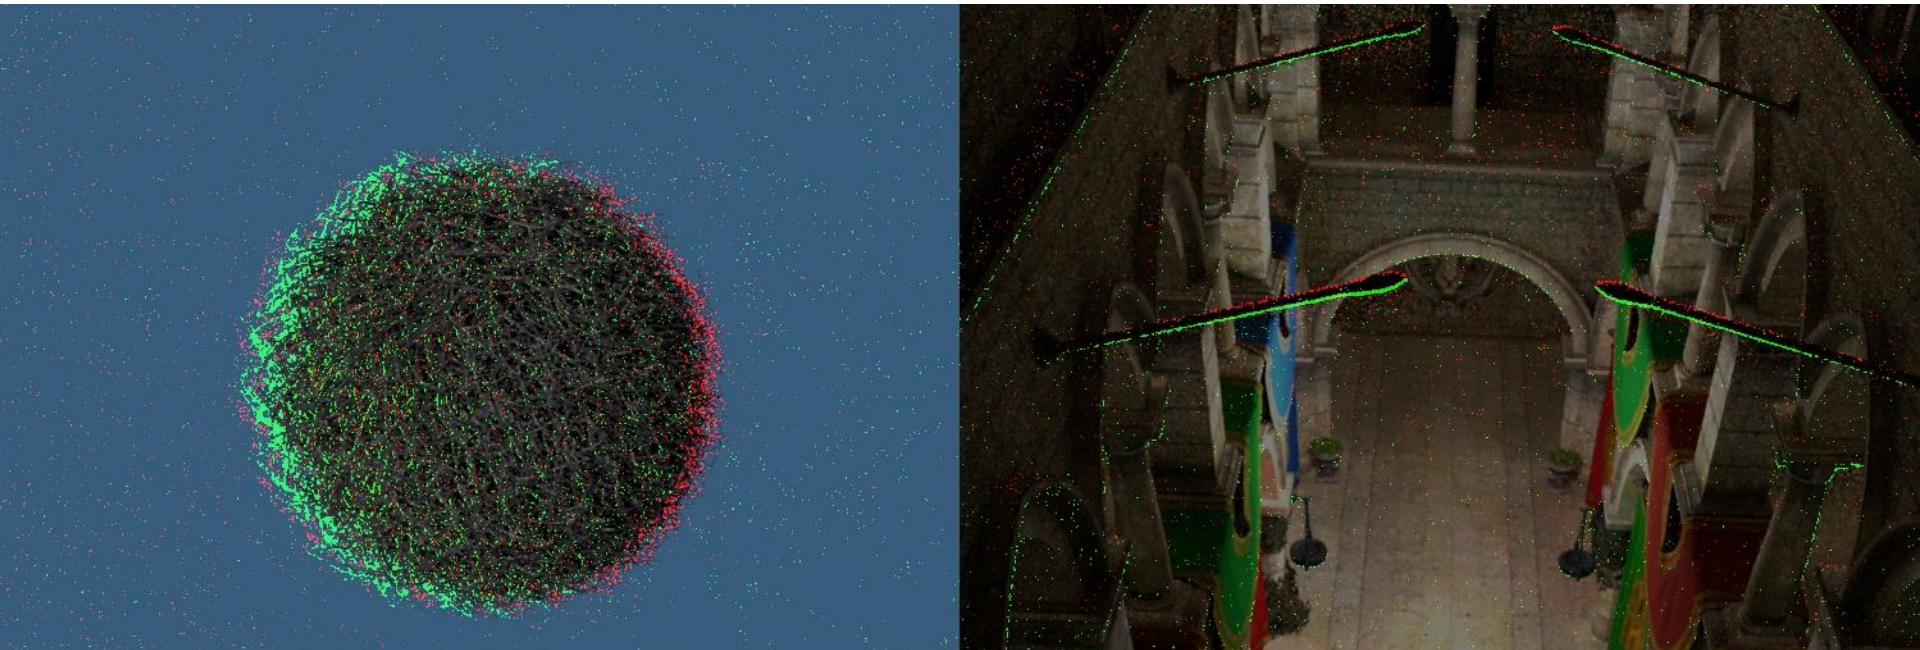

Comparison with previous work

Video

Problem:
How do we improve temporal stability?

Stable samples

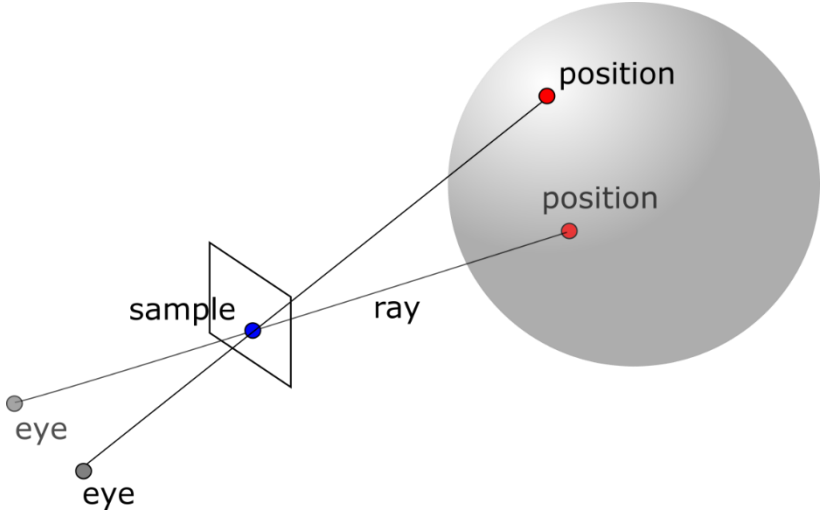

Supersampling

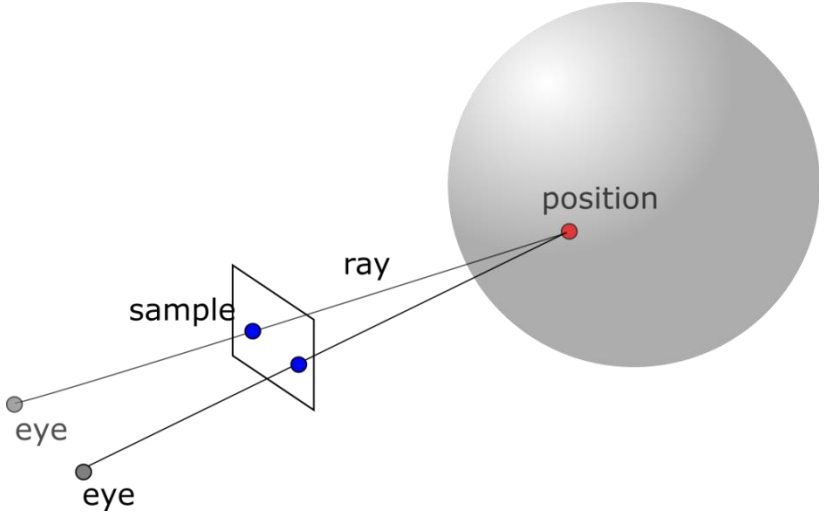

Stable ray tracing

Related work

- Sample reprojection [Badt 1988, Walter et al. 1999, 2002; Zhu et al. 2005]
  - Mostly for efficiency reasons
  - Stores shading result

- Sample reprojection for temporal stabilty [Adelson and Hodges 1995, Martin et al. 2002]
  - One sample per pixel only
  - Need post processing to remove residual temporal instability

- Temporal stability: Temporal supersampling [Karis 2014; Patney et al. 2016]
  - Post processing filter
  - Increases blurriness

Interactive stable ray tracing

- The final goal is temporal stability, retaining sharpness

- We allow a non integral number of samples per pixel
- Sample density estimation
- Making it efficient for modern GPU pipelines

- Application: caching of global illumination

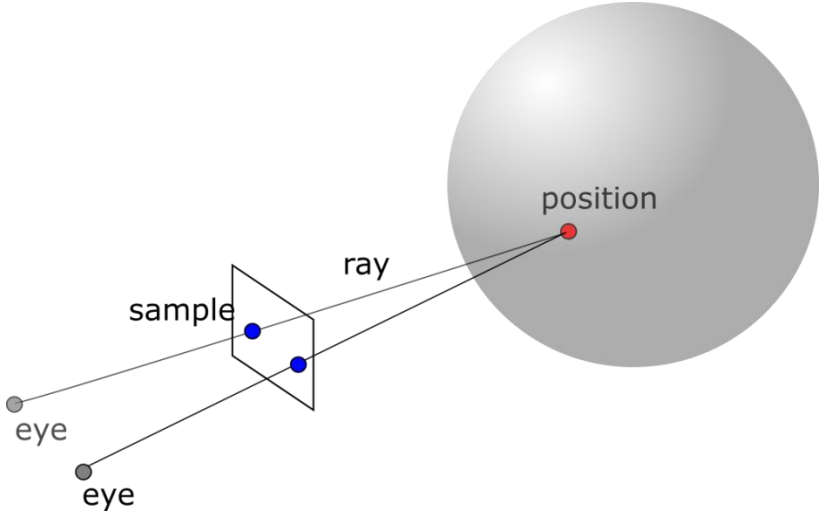

Interactive stable ray tracing

A generic stable ray tracing algorithm

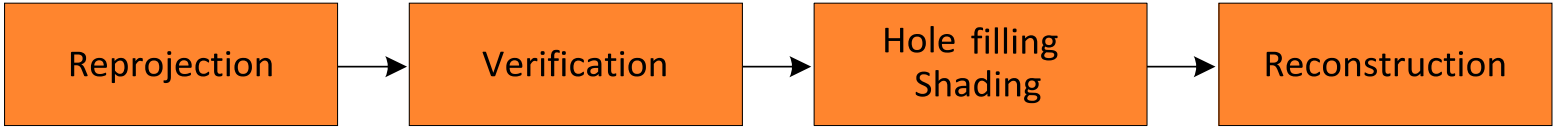

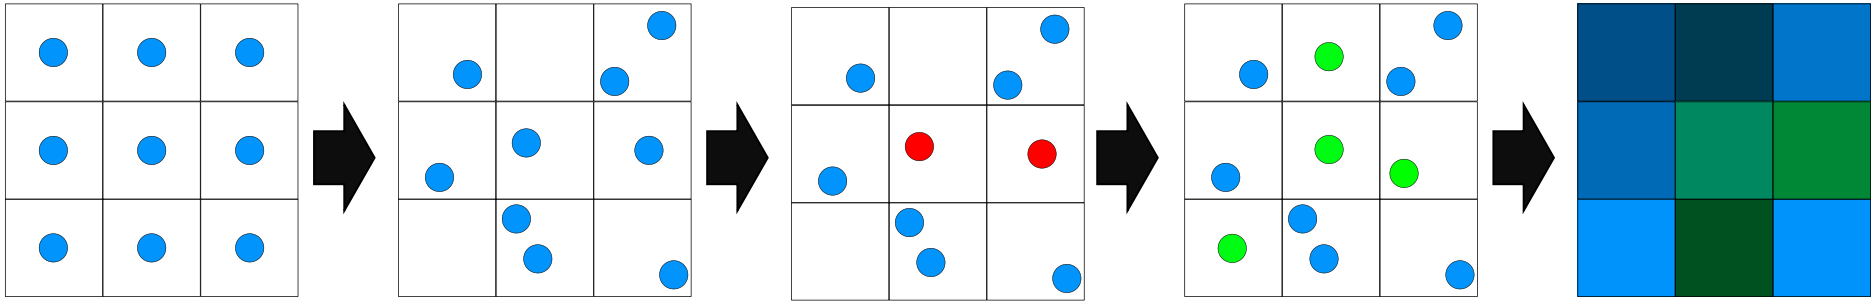

For an efficient implementation

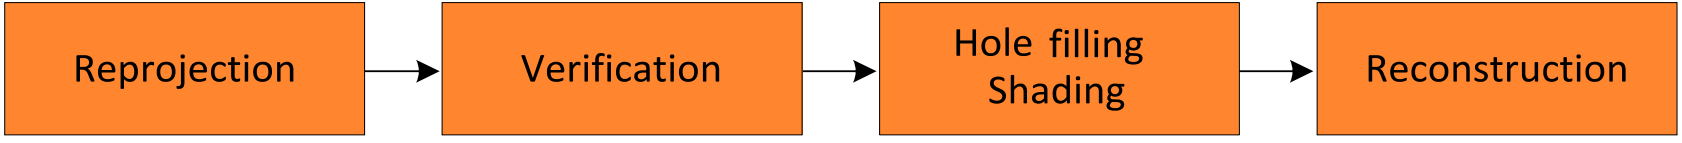

Our implementation:

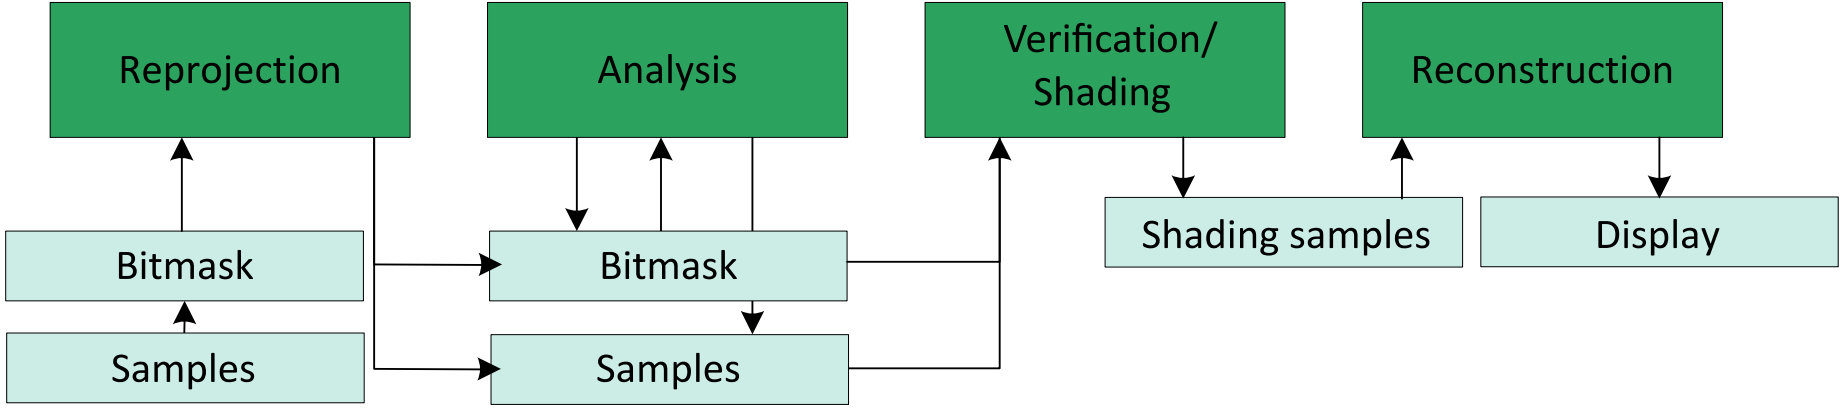

Our reprojection

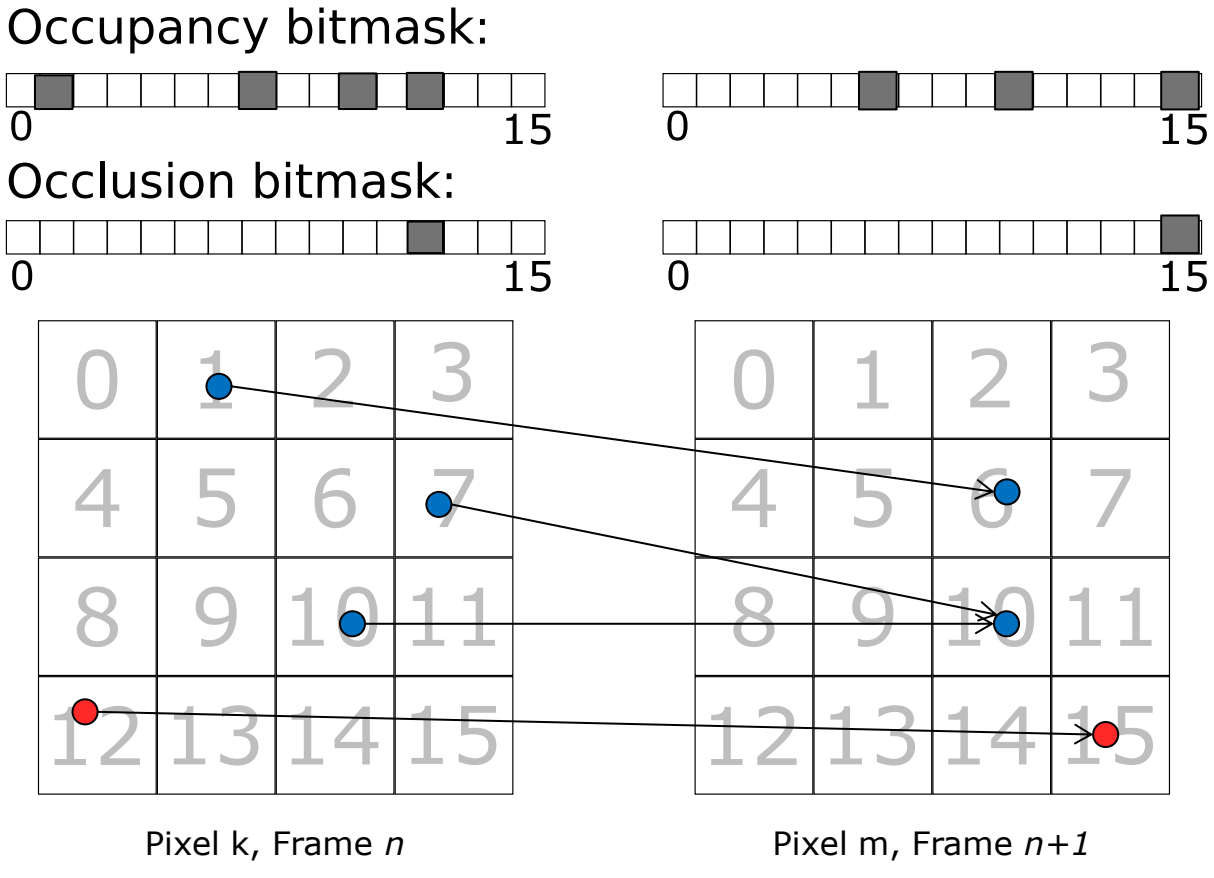

Analysis, first to estimate density

- User-defined target density (1 spp, 2 spp, etc.)

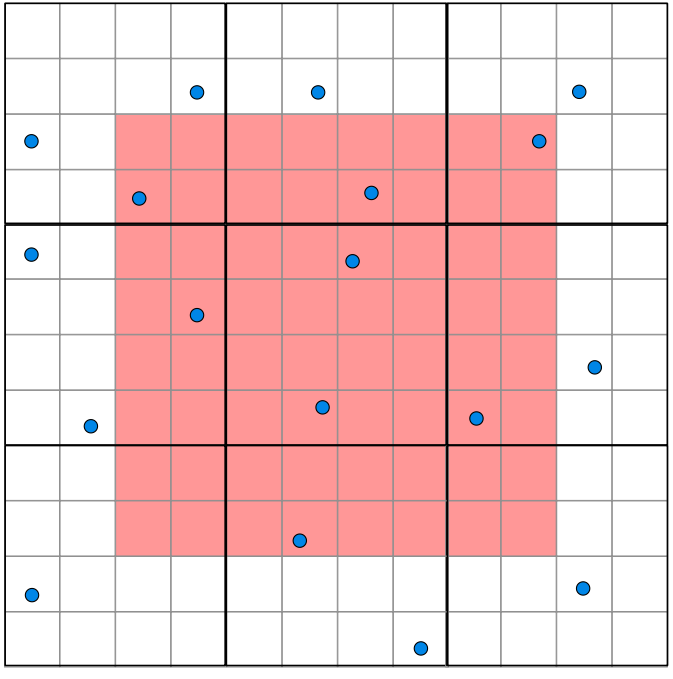

Density estimation radius

Analysis, to add and remove samples

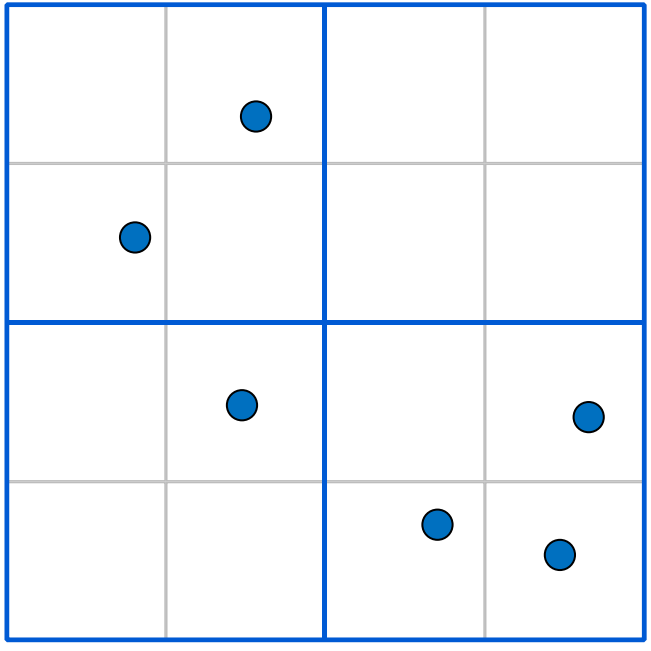

Sample removal

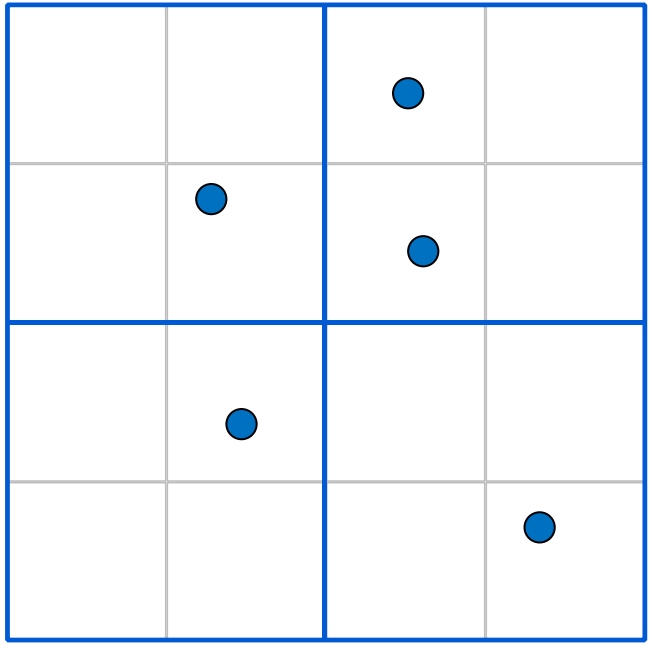

Sample addition

Verification, with a cached sample

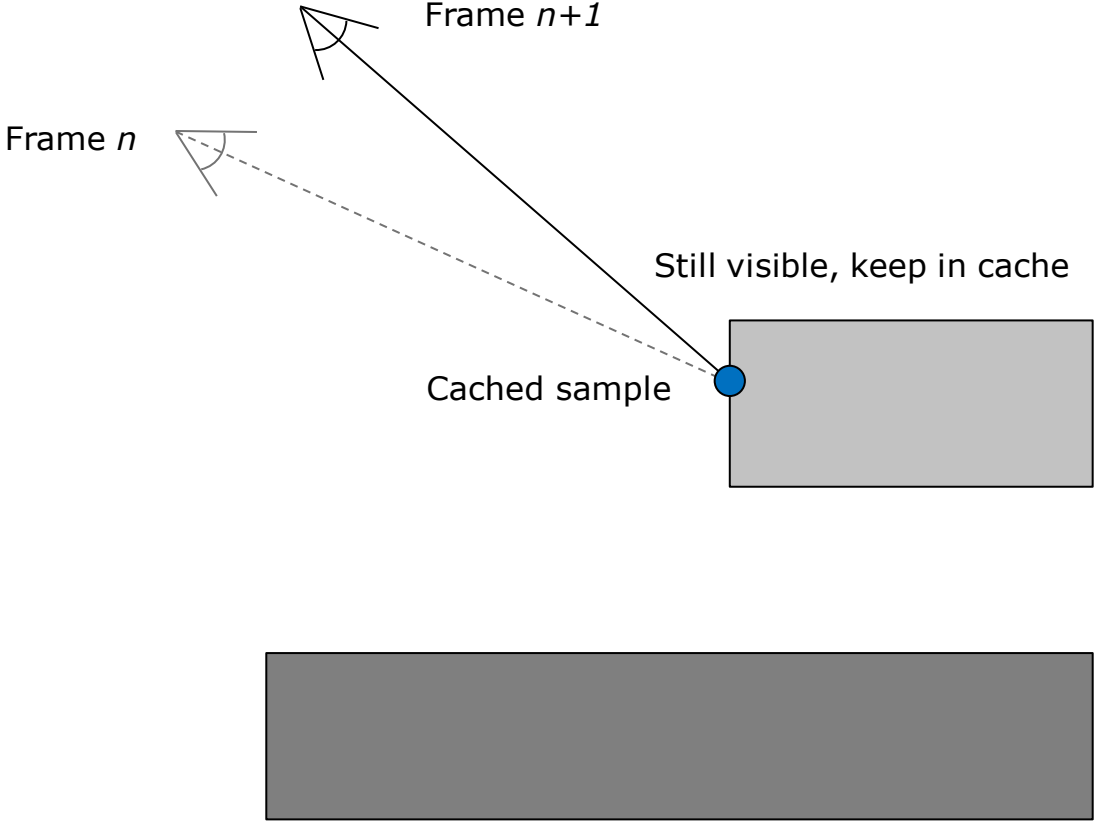

Verification, with an occluded sample

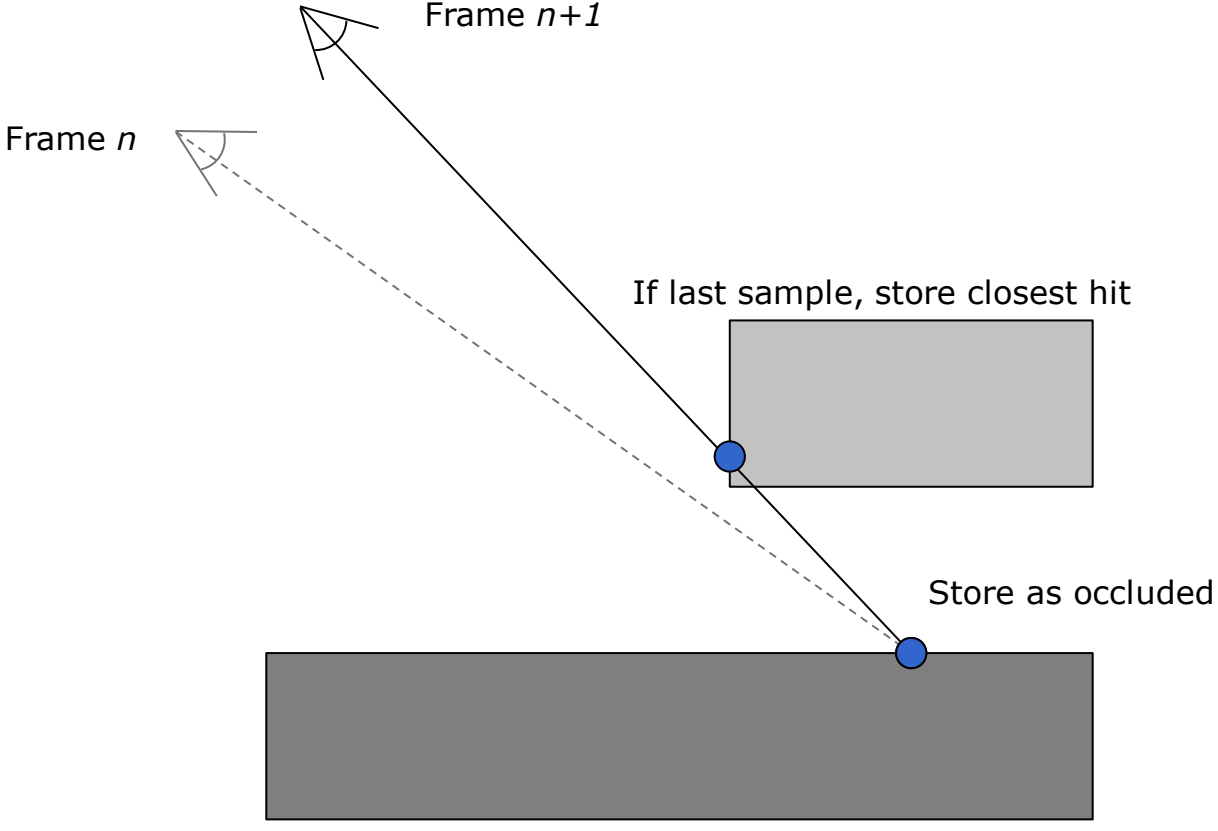

Reconstruction

- 3x3 gaussian filter

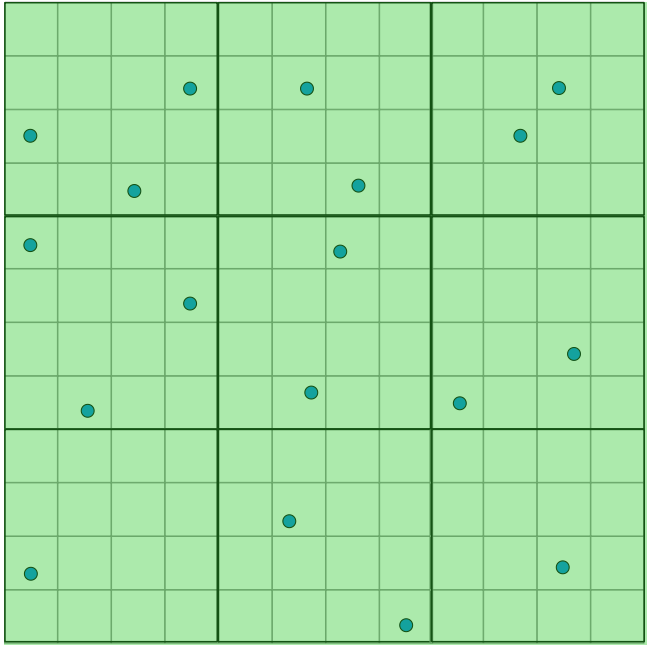

- It is possible to apply post processing such as temporal integration / antialiasing [Patney et al. 2016]

Caching global illumination

- Our samples are suitable to store the result of indirect illumination
- In our implementation: unidirectional PT, Exponential moving average

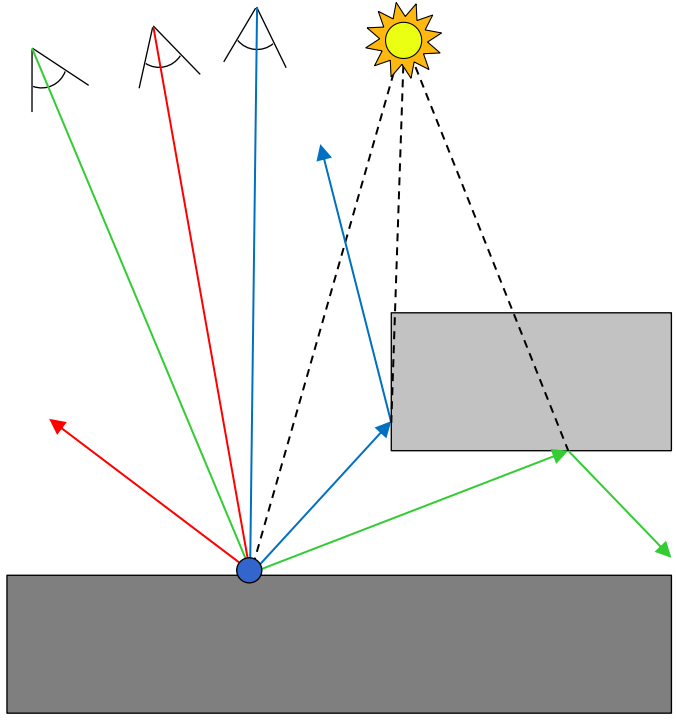

Results

Questions

- Reconstruction quality
  - What is the impact on reconstruction quality?
- Temporal stability
  - Do we improve temporal stability at iso performance?
- Sharpness improvement
  - Do we achieve better sharpness compared to temporal integration / antialiasing?
- Performance impact
  - What is the performance impact?

Quality impact mostly on low spp

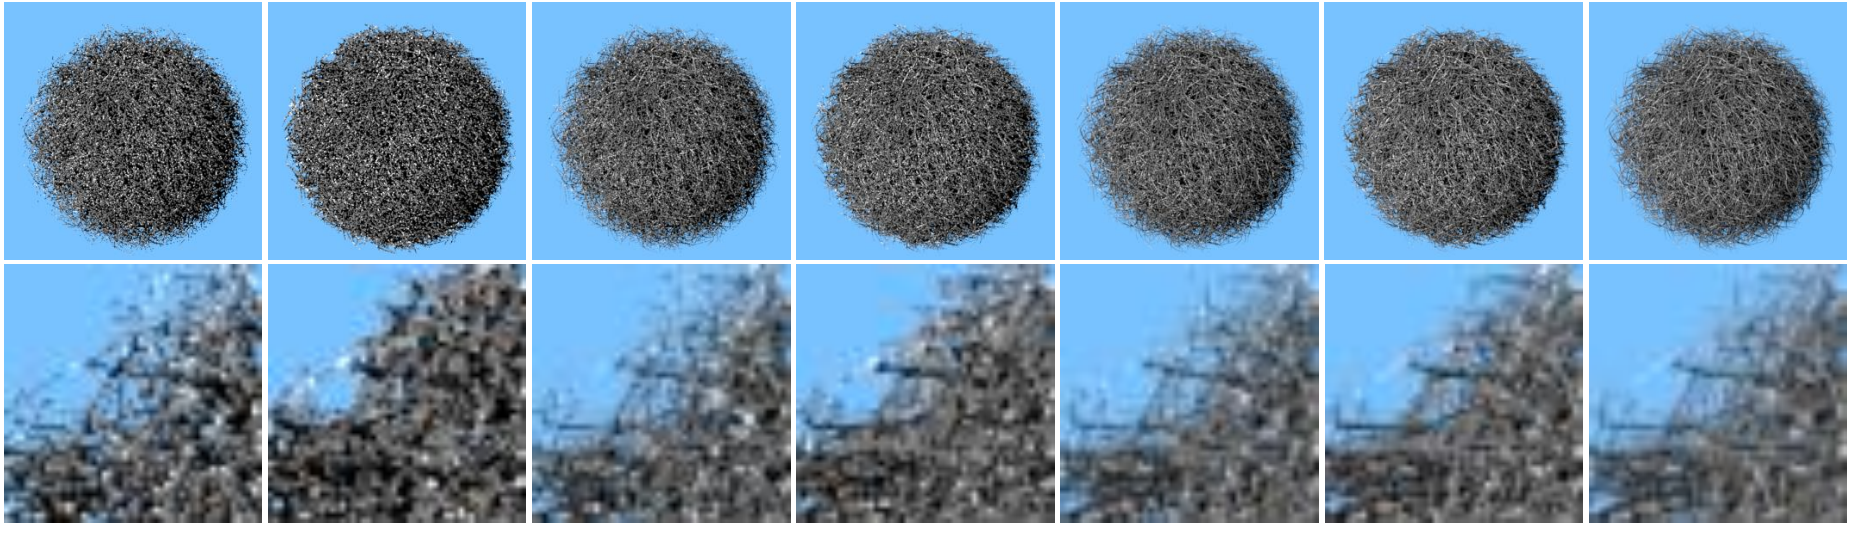

SS, 1 spp | SRT, 1 spp | SS, 2 spp | SRT, 2 spp | SS, 4 spp | SRT, 4 spp | SS, 32 spp

SS = Supersampling, SRT = Stable ray tracing

Temporal stability, hairball

Video

Impact on sharpness

- Sharpness: CPBD-based metric (higher is sharper)

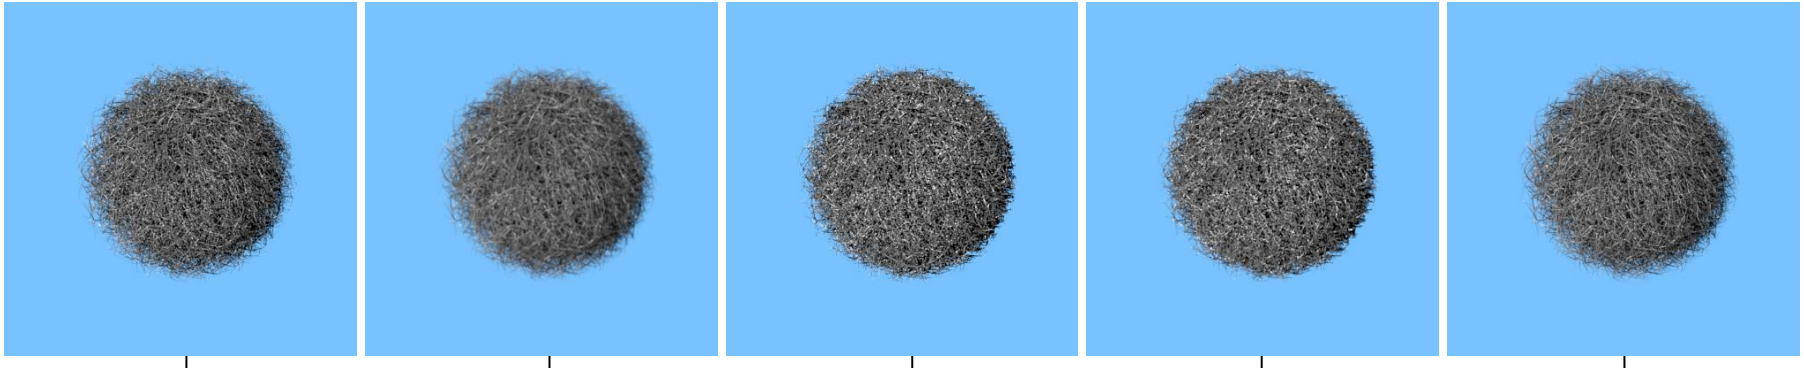

Supersampling, 4 spp
Sharpness: 0.8142

Supersampling, 4 spp
+ temporal antialiasing
Sharpness: 0.6610

Stable ray tracing, 2 spp
Sharpness: 0.8056

Stable ray tracing, 2 spp
+ temporal integration
Sharpness: 0.7783

Supersampling, 32 spp
Sharpness: 0.8054

Comparing with previous work

Video

A test with global illumination

Video

Global illumination, sharpness

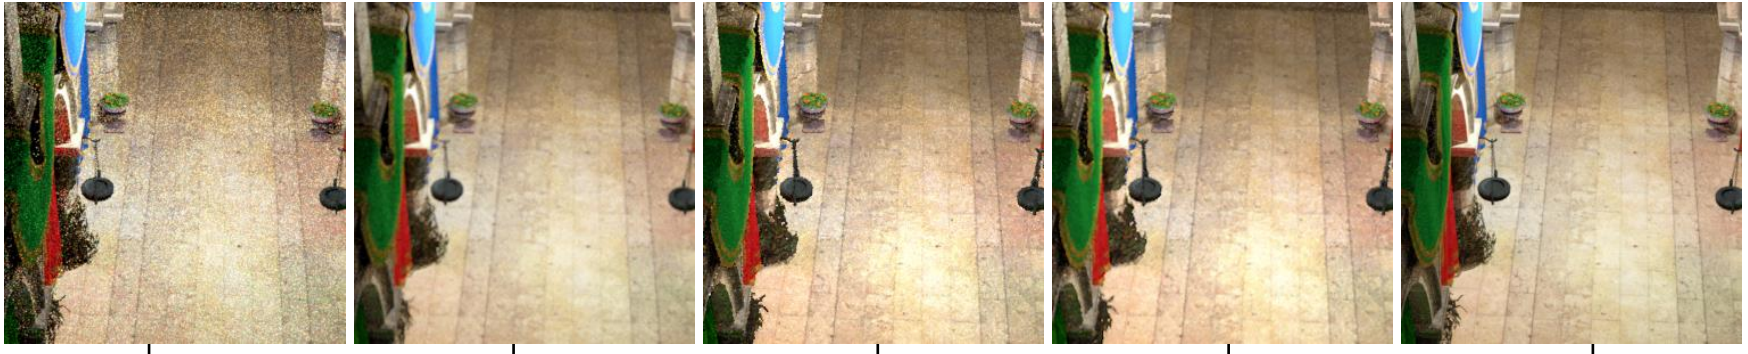

Supersampling, 2 spp

Sharpness: 0.7924

Supersampling, 2 spp

+ temporal antialiasing

Sharpness: 0.5348

Stable ray tracing, 1 spp

Sharpness: 0.7085

Stable ray tracing, 1 spp

+ temporal integration

Sharpness: 0.6060

Supersampling, 32 spp

Sharpness: 0.6771

Performance results

Technique	Reprojection	Analysis	Verification/ Shading	Reconstruction	Total
Stable ray tracing, d_{target}= 1 spp	1.05 ms	0.28 ms	18.91 ms	0.72 ms	**20.94 ms**
Supersampling, 1 spp	-	-	13.35 ms	0.21 ms	**13.56 ms**
Supersampling, 2 spp	-	-	20.94 ms	0.38 ms	**21.32 ms**

Technique	Reprojection	Analysis	Verification/ Shading	Reconstruction	Total
Stable ray tracing, d_{target}= 2 spp	1.23 ms	0.38 ms	28.88 ms	0.82 ms	**31.31 ms**
Supersampling, 3 spp	-	-	28.36 ms	0.54 ms	**28.90 ms**
Supersampling, 4 spp	-	-	35.86 ms	0.71 ms	**36.57 ms**

Limitation, edge thickening

- Caused by change in distributions of samples

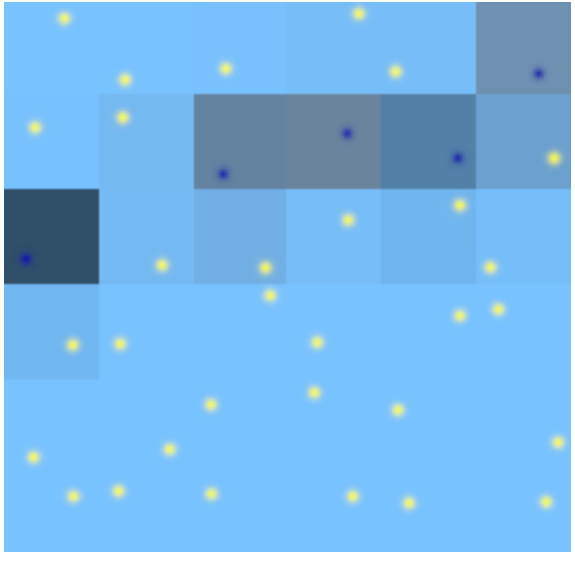

Frame 0, No thickening

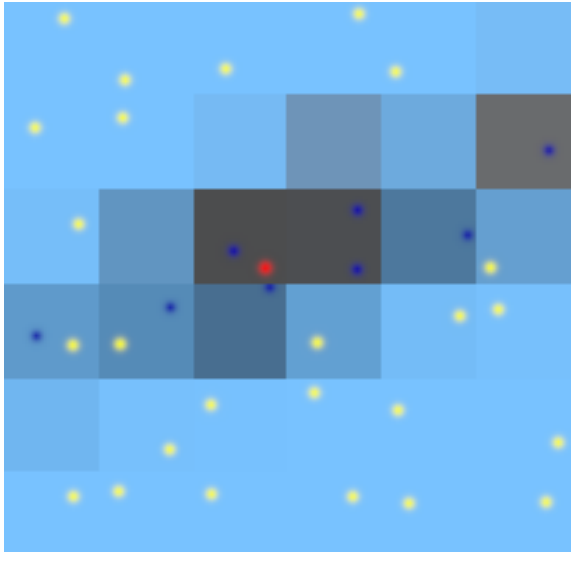

Frame 1, Thickened edge

- Legend: **Background sample**, **Hair sample**, **Occluded sample**

A new practical technique for stable shading

- Based on sample reprojection, with inexpensive analysis phase
- Balances temporal stability and image sharpness
- Sharp, fairly temporally stable result
- Allows us to cache global illumination effects

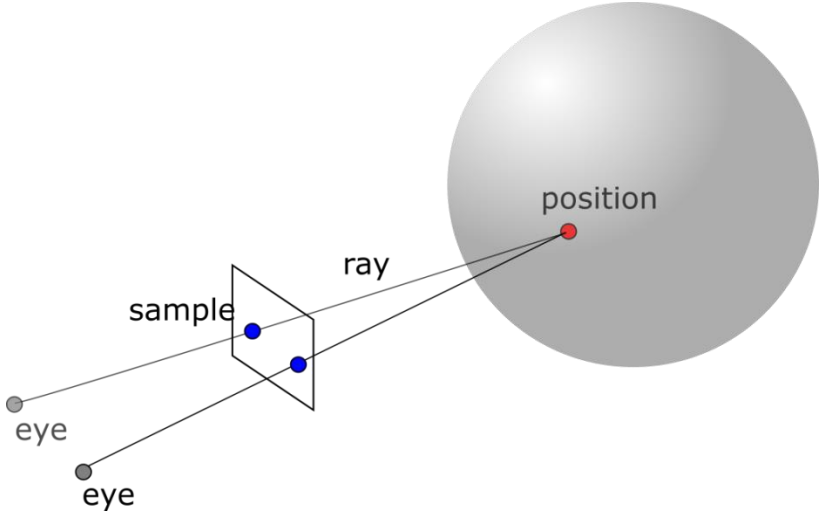

References

[Adelson and Hodges 1995] Stephen J. Adelson and Larry F. Hodges. 1995. Generating exact ray-traced animation frames by reprojection. IEEE Computer Graphics and Applications 15, 3 (1995), 43–52.

[Badt 1988] Sig Badt, Jr. 1988. Two algorithms for taking advantage of temporal coherence in ray tracing.The Visual Computer 4, 3 (1988), 123–132.

[Bala et al. 2003] Kavita Bala, BruceWalter, and Donald P. Greenberg. 2003. Combining edges and points for interactive high-quality rendering. ACM Transactions on Graphics (Proc. of SIGGRAPH 03) 22, 3 (July 2003), 631–640.

[Herzog et al. 2010] Robert Herzog, Elmar Eisemann, Karol Myszkowski, and Hans-Peter Seidel. 2010. Spatio-temporal upsampling on the GPU. In Proceedings of Interactive 3D Graphics and Games (I3D '10). ACM, 91–98.

[Karis 2014] Brian Karis. 2014. High-quality temporal supersampling. In Advances in Real-Time Rendering in Games, Part I. Number 10 in ACM SIGGRAPH 2014 Courses. http://advances.realtimerendering.com/s2014/

[Martin et al. 2002] William Martin, Peter Shirley, Steven Parker, William Thompson, and Erik Reinhard. 2002. Temporally coherent interactive ray tracing. Journal of Graphics Tools 7, 2 (2002), 41–48.

[Nehab et al. 2007] Diego Nehab, Pedro V. Sander, Jason Lawrence, Natalya Tatarchuk, and John R. Isidoro. 2007. Accelerating real-time shading with reverse reprojection caching. In Proceedings of Graphics Hardware (GH 2007).

[Patney et al. 2016] Anjul Patney, Marco Salvi, Joohwan Kim, Anton Kaplanyan, Chris Wyman, Nir Benty, David Luebke, and Aaron Lefohn. 2016. Towards foveated rendering for gaze tracked virtual reality. ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2016) 35, 6 (November 2016), 179:1–179:12.

[Scherzer et al. 2007] Daniel Scherzer, Stefan Jeschke, and Michael Wimmer. 2007. Pixel-correct shadowmaps with temporal reprojection and shadow test confidence. In Rendering Techniques 2007 (Proceedings of EGSR 2007). Eurographics Association, 45–50.

[Velazquez-Armendariz et al. 2006] Edgar Velazquez-Armendariz, Eugene Lee, Kavita Bala, and Bruce Walter. 2006. Implementing the render cache and the edge-and-point image on graphics hardware. In Proceedings of Graphics Interface 2006 (GI '06). Canadian Information Processing Society, 211–217.

[Walter et al. 1999] Bruce Walter, George Dre.akis, and Steven Parker. 1999. Interactive rendering using the render cache. In Rendering techniques '99 (Proceedings of EGWR 1999). Springer, 19–30.

[Yang et al. 2009] Lei Yang, Diego Nehab, Pedro V. Sander, Pitchaya Si.hi-amorn, Jason Lawrence, and Hugues Hoppe. 2009. Amortized supersampling. ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2009) 28, 5 (December 2009), 135:1–135:12.

[Zhu et al. 2005] Tenghui Zhu, Rui Wang, and David Luebke. 2005. A GPU accelerated render cache. In Proceedings of Pacific Graphics 2005 (short paper).

Quality impact mostly on low spp

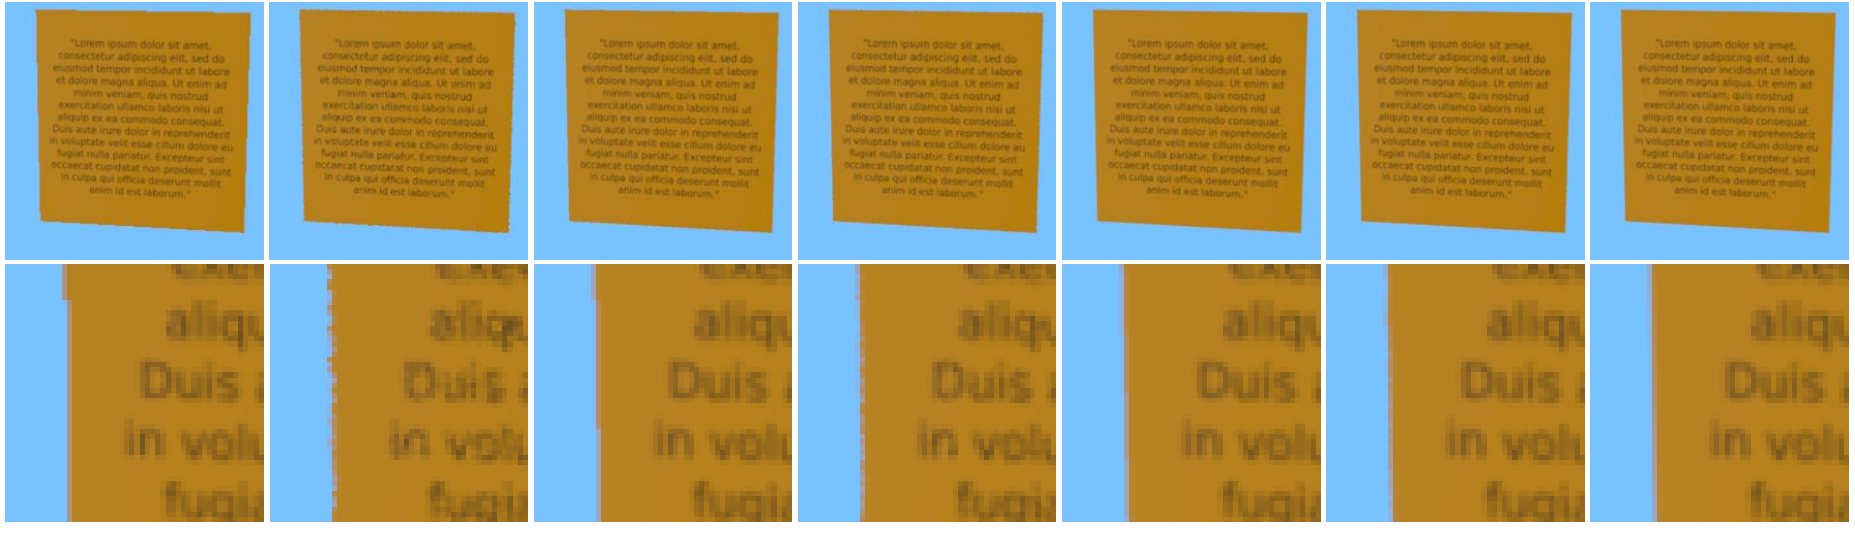

SS, 1 spp | SRT, 1 spp | SS, 2 spp | SRT, 2 spp | SS, 4 spp | SRT, 4 spp | SS, 32 spp

No averaging improves sharpness

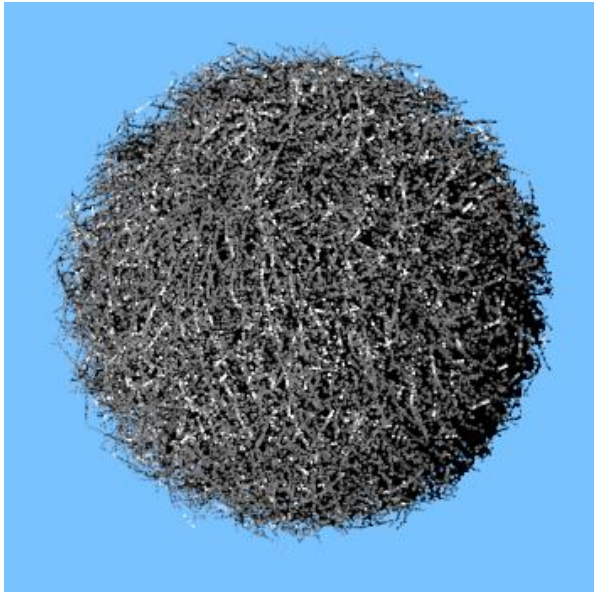

Stable ray tracing, 1 spp
sharpness: 0.8182

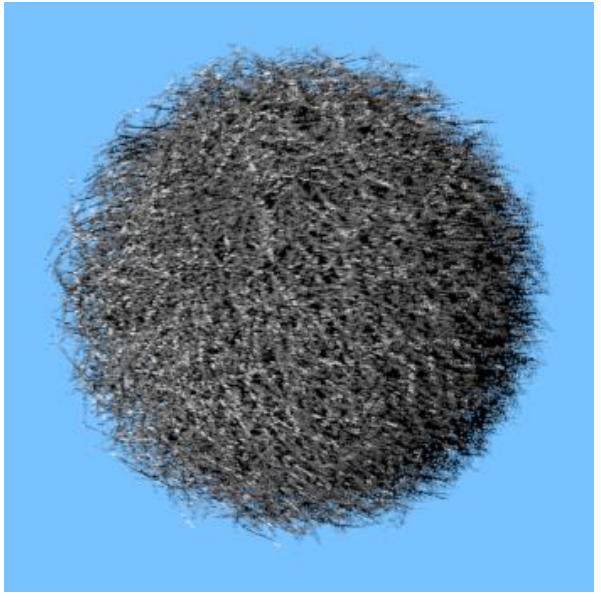

[Martin et al. 2002]
sharpness: 0.6957