Monte Carlo Methods for physically based Volume rendering
SIGGRAPH 2018 Course
Advanced methods and acceleration data structures

Johannes Hanika
Karlsruhe Institute of Technology
Weta Digital

content

rainbows and heroes

- spectral tracking
- hero wavelengths

emissive volumes

- fires, explosions, forward next event estimation

acceleration data structures

- adaptive storage for heterogeneous media

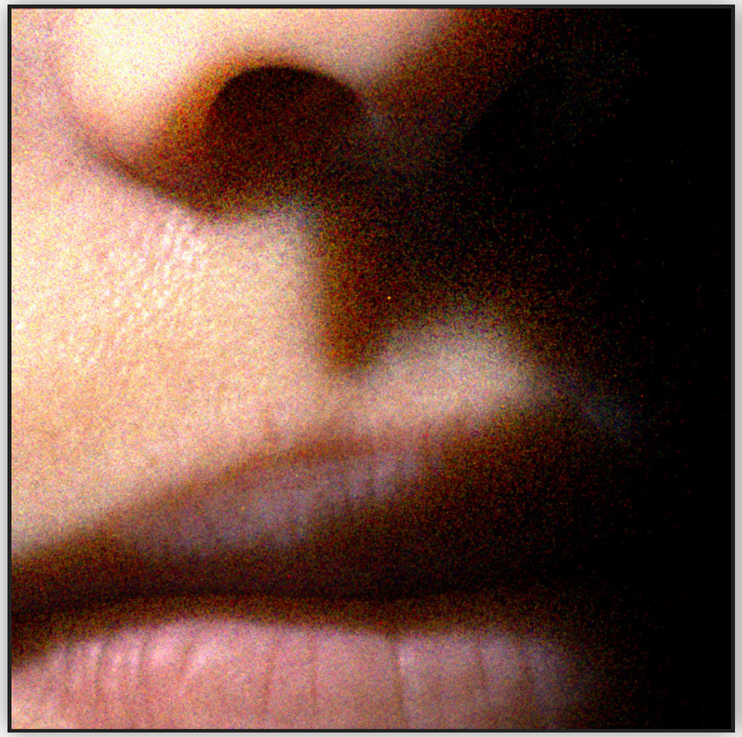
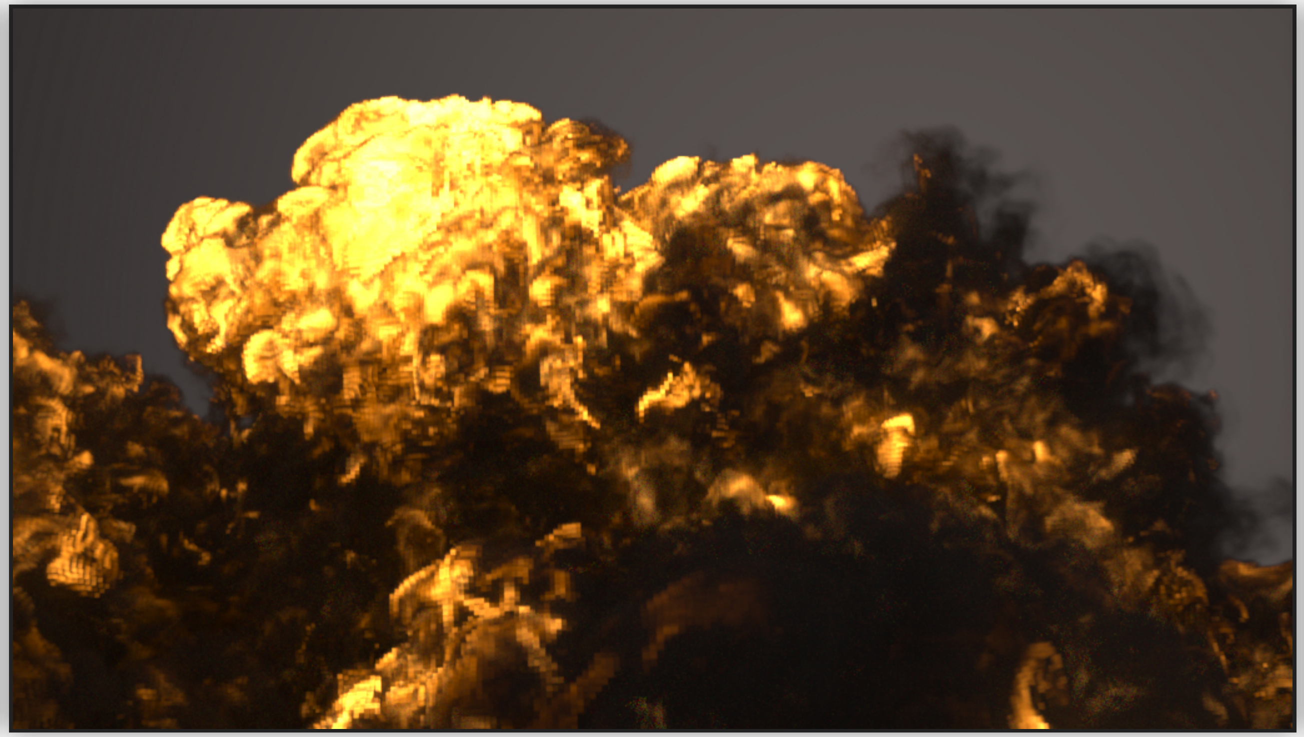
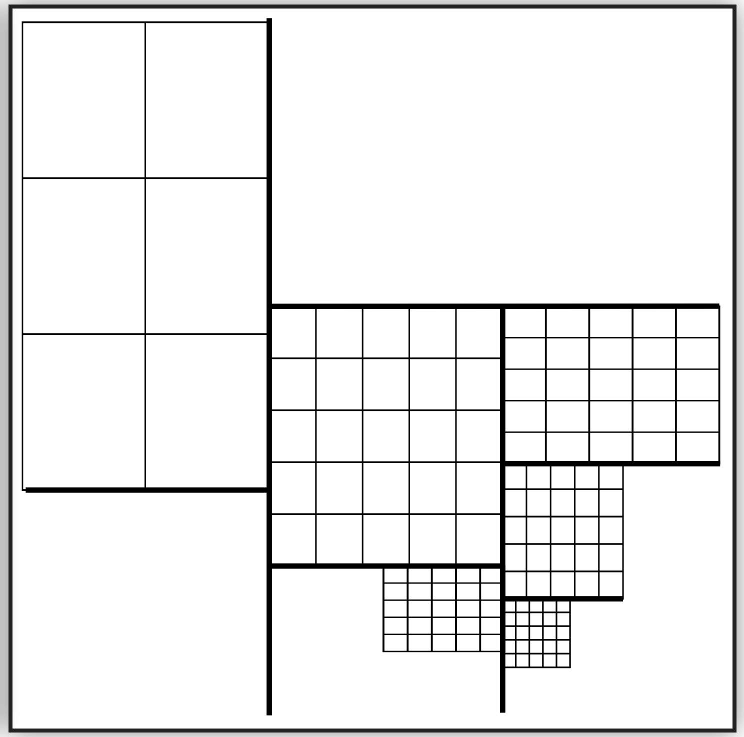

spectral tracking

- example material: human skin, insanely complex layered structure
  - expensive to model and trace through
- often used in graphics: approximate by homogeneous chromatic medium
  - captures main look features: red blur
  - low path vertex count for more efficient simulation

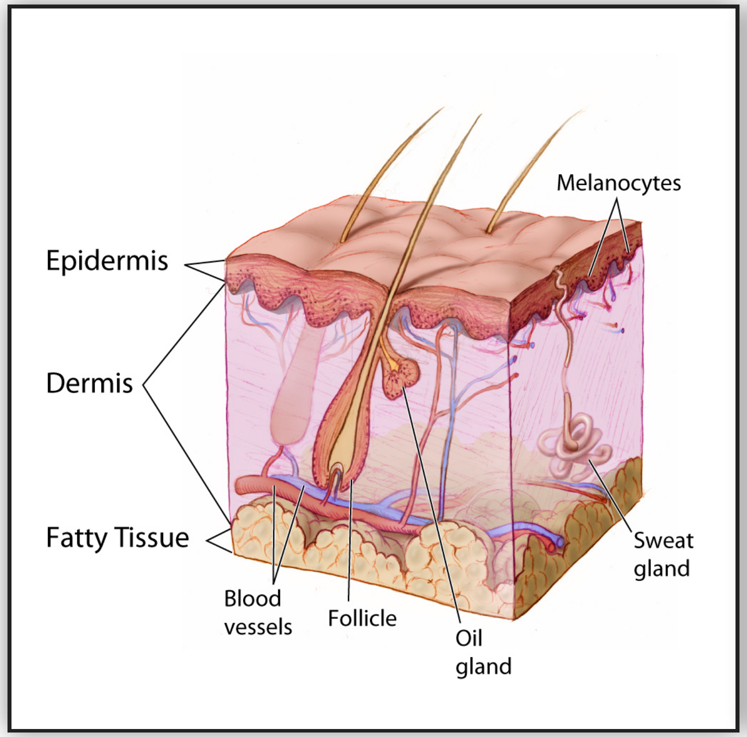

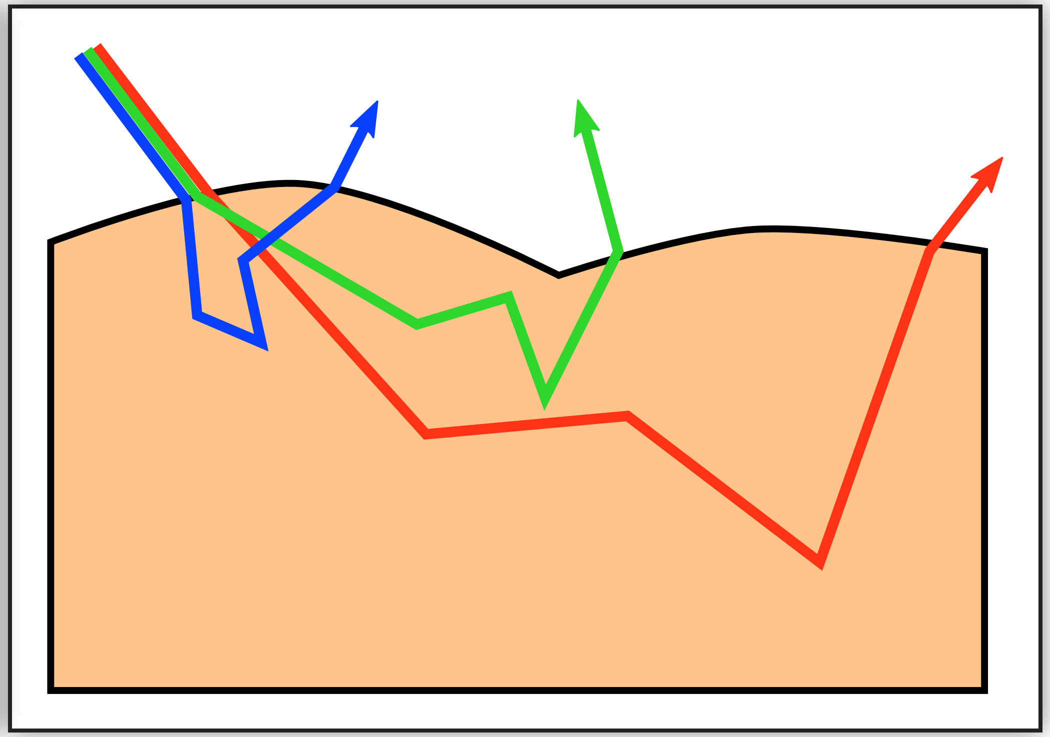

spectral tracking

- a problem often encountered in skin: chromatic media
  - collision coefficients μ depend on wavelength λ
  - for instance free flight distance much longer for long wavelengths:

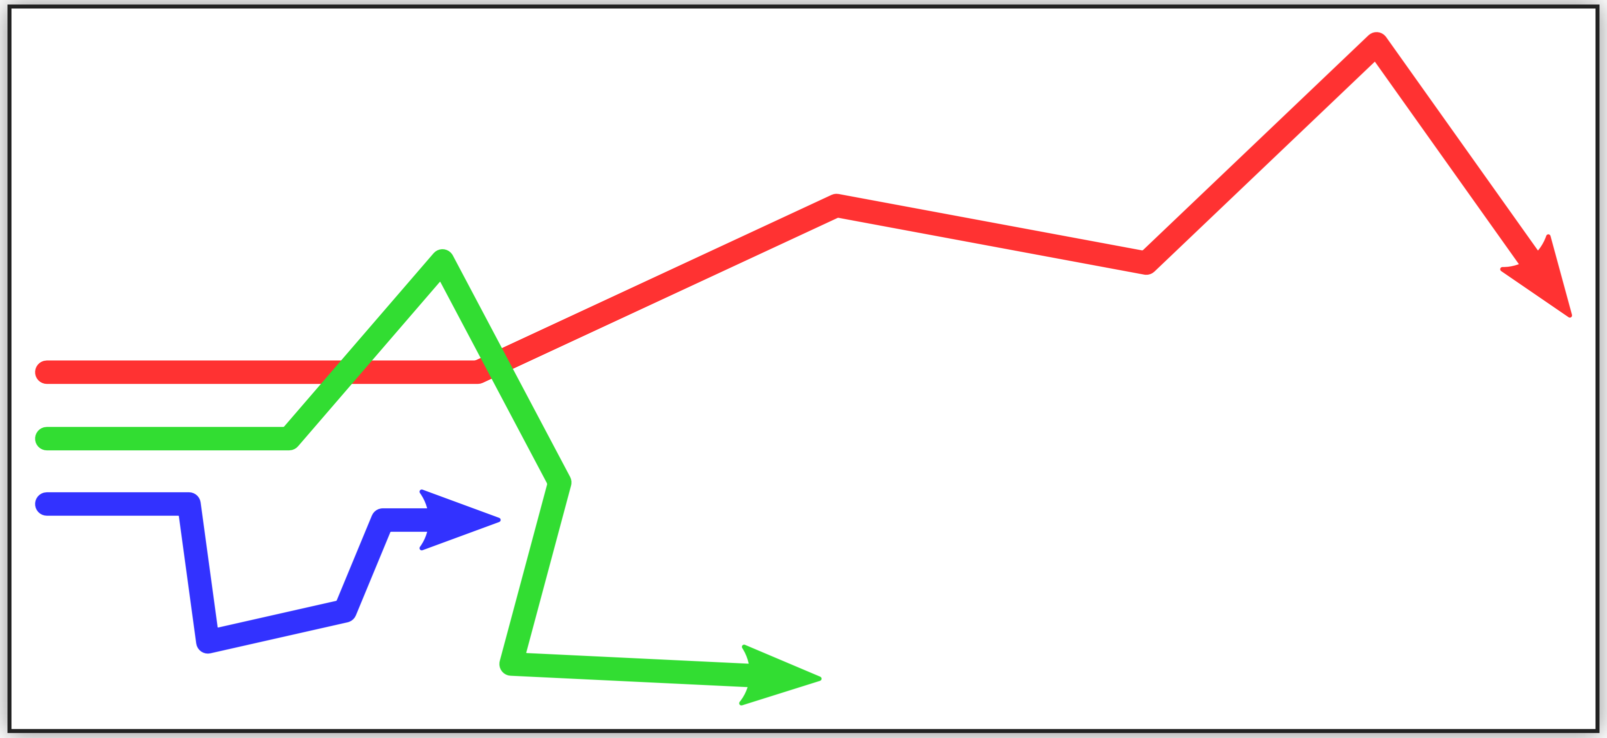

  - makes path invalid for different wavelength?
  - can we still exploit coherence?

spectral tracking

- a problem often encountered in skin: chromatic media
  - collision coefficients μ depend on wavelength λ
  - for instance free flight distance much longer for long wavelengths:

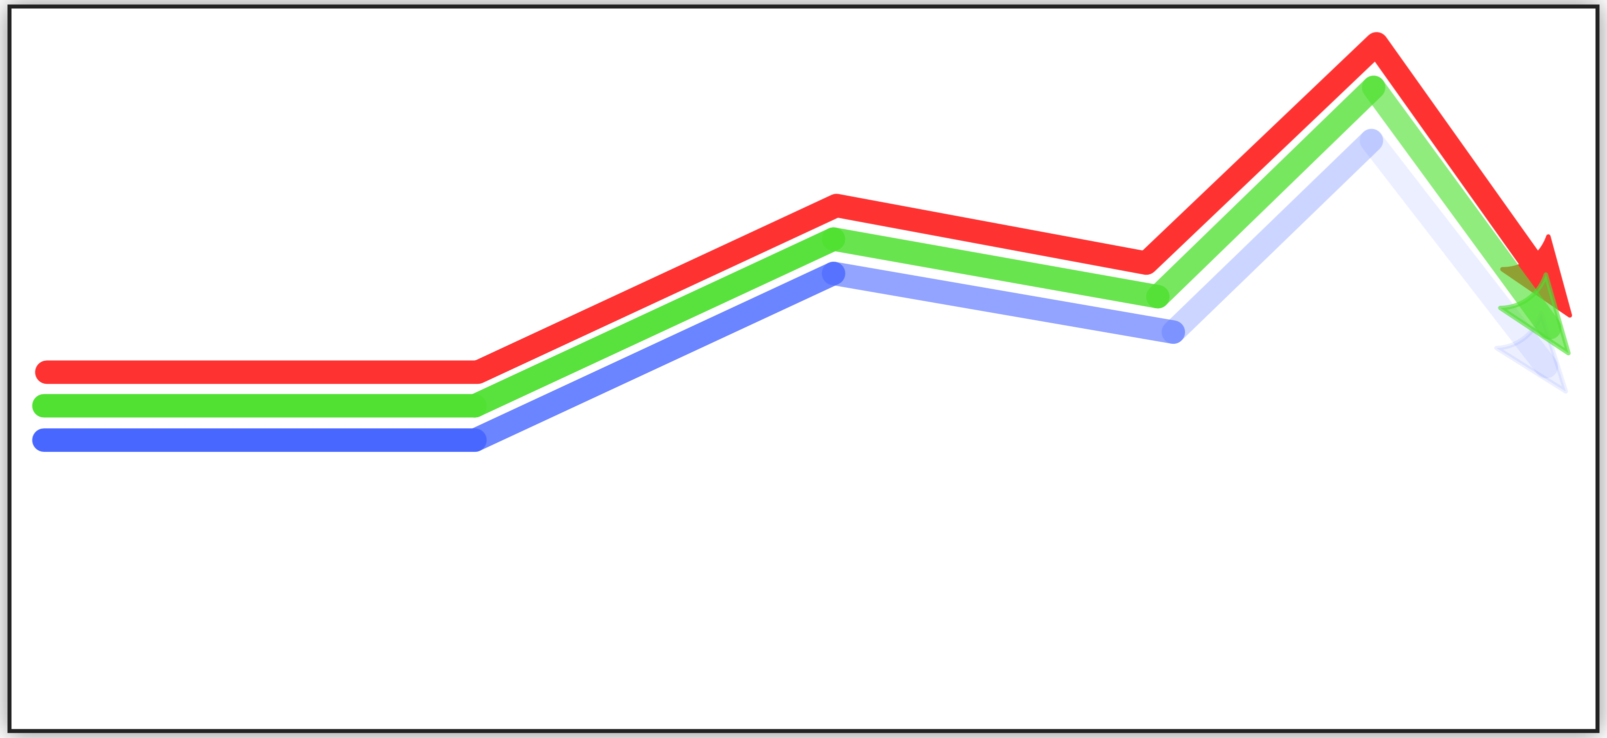

  - makes path invalid for different wavelength?
  - can we still exploit coherence?

spectral tracking via MIS

- **hero wavelength** sampling [WND∗14]

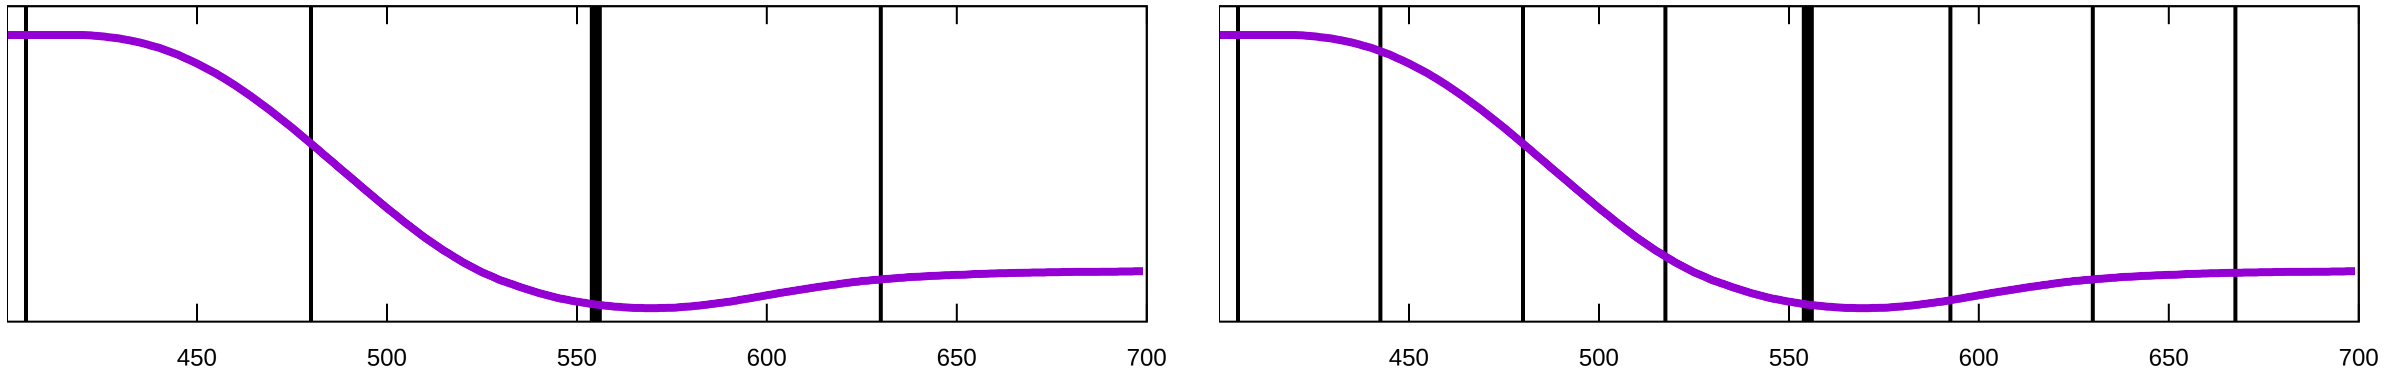

- sample perfectly for one single wavelength λ_0
- evaluate path for a stratified set of wavelengths λ_i at the same time
- optimally weighted combination via **MIS (balance heuristic)**
  - limited to **regular tracking** because it requires explicit evaluation of PDF

$$\frac{f(\bar{\mathbf{x}}, \lambda_i)}{\sum_j p(\bar{\mathbf{x}}, \lambda_j)}$$

[WND∗14] Wilkie A., Nawaz S., Droske M., Weidlich A., Hanika J.: Hero wavelength spectral sampling. CGF (Proc. EGSR) 33, 4 (June 2014), 123–131.

spectral tracking via MIS

image comparison 64spp

- skin material with 1 wavelength

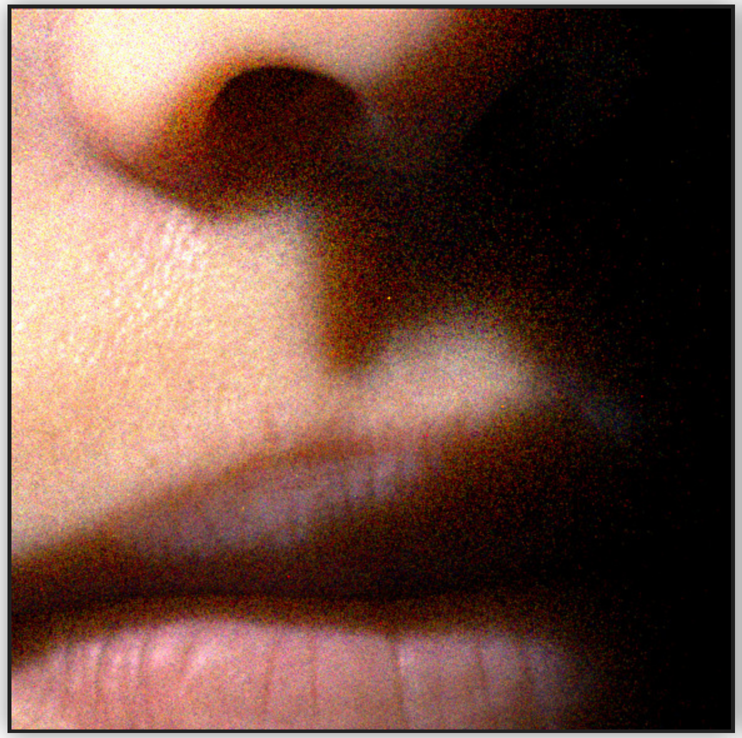

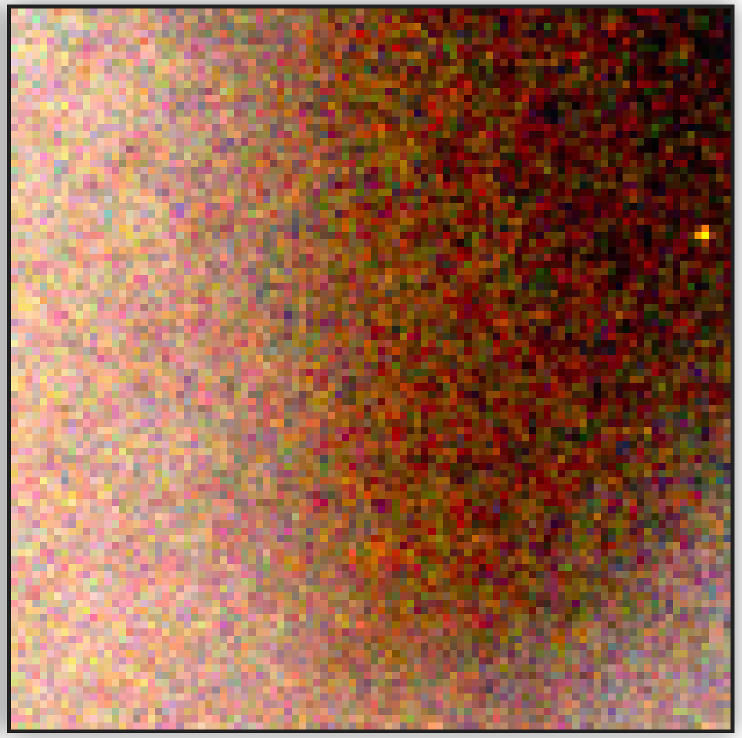

spectral tracking via MIS

image comparison 64spp

- skin material with 4 wavelengths (SSE)

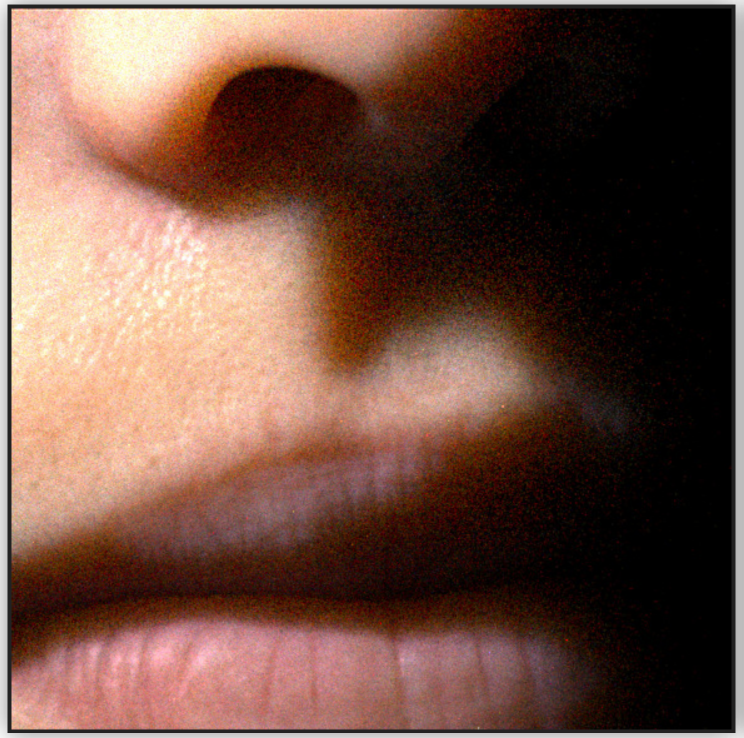

spectral tracking via MIS

image comparison 64spp

- skin material with 8 wavelengths (AVX)

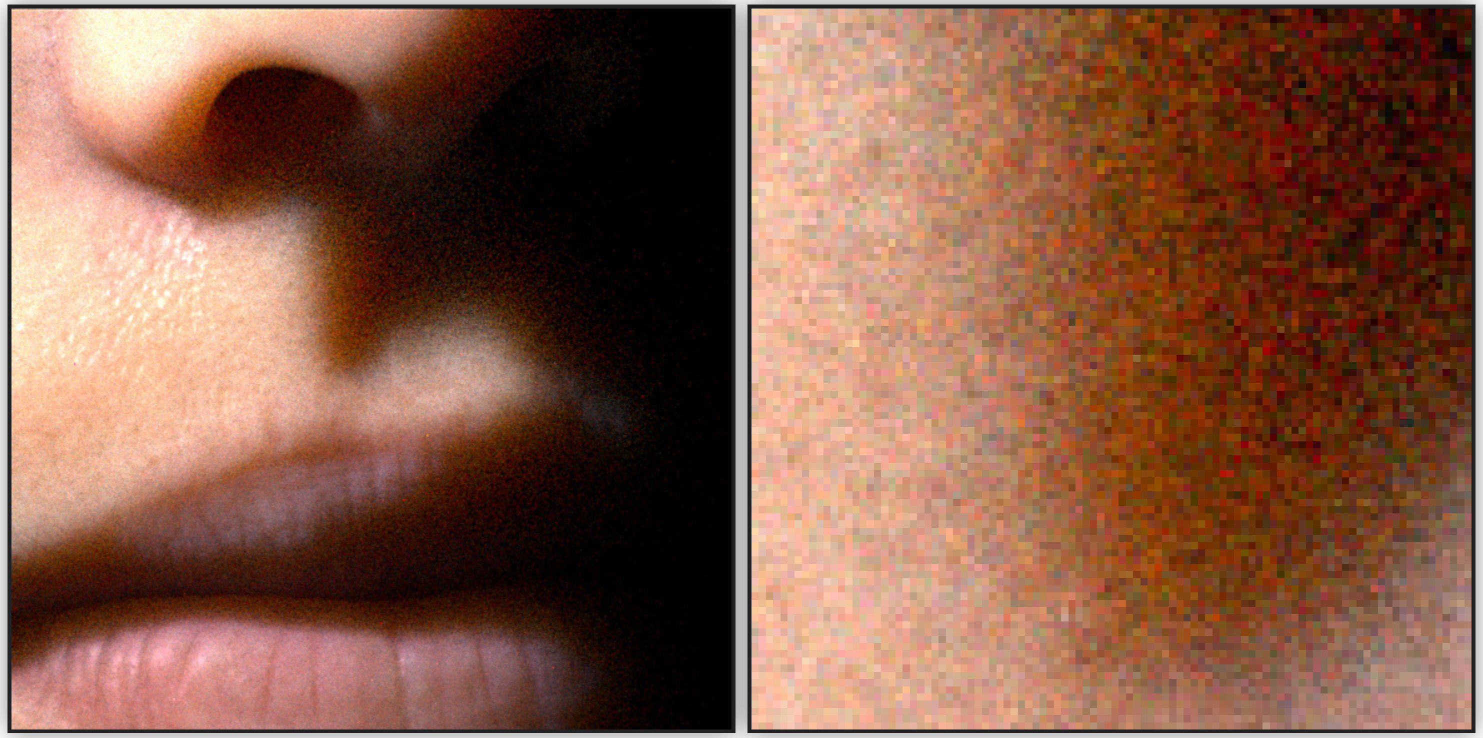

- note that all these images are using the exact same paths!

spectral tracking without PDF [KHLN17]

- sample fictitious event by common majorant $\bar{\mu}$
- how do decide for null collision, scattering, or absorption at next vertex $\mathbf{x}$?
- probability $P_\star$ according to

$$\mu_\star(\mathbf{x}, \lambda) \in \mu_n(\mathbf{x}, \lambda), \mu_s(\mathbf{x}, \lambda), \mu_a(\mathbf{x}, \lambda)$$

$$P_\star(\mathbf{x}) = \mathrm{reduce}(|\mu_\star(\mathbf{x}, \lambda)|)c$$

- c is a normalisation constant such that the $P_\star$ sum to one

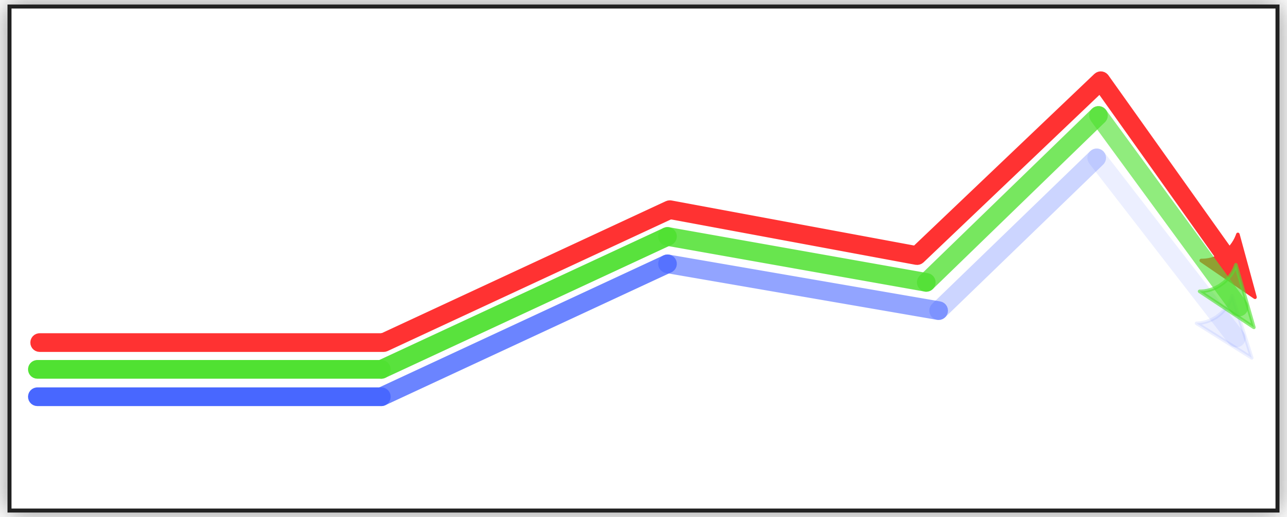

[KHLN17] Kutz P., Habel R., Li Y. K., Novák J.:
Spectral and decomposition tracking for rendering heterogeneous volumes.
ACM TOG (Proc. SIGGRAPH) 36, 4 (July 2017)

spectral tracking without PDF [KHLN17]

- probability $P_\star$ according to $\mu_\star(\mathbf{x},\lambda) \in \mu_n(\mathbf{x},\lambda), \mu_s(\mathbf{x},\lambda), \mu_a(\mathbf{x},\lambda)$
- pick by maximum over λ_i (always follow densest material)

$$P_\star(\mathbf{x}) = \max(|\mu_\star(\mathbf{x},\lambda)|)c$$

- pick by average weighted by spectral path throughput history $\bar{w}(\bar{\mathbf{x}},\lambda)$

$$P_\star(\mathbf{x}_j) = \text{reduce}(|\bar{w}(\bar{\mathbf{x}},\lambda)\mu_\star(\mathbf{x},\lambda)|)c$$

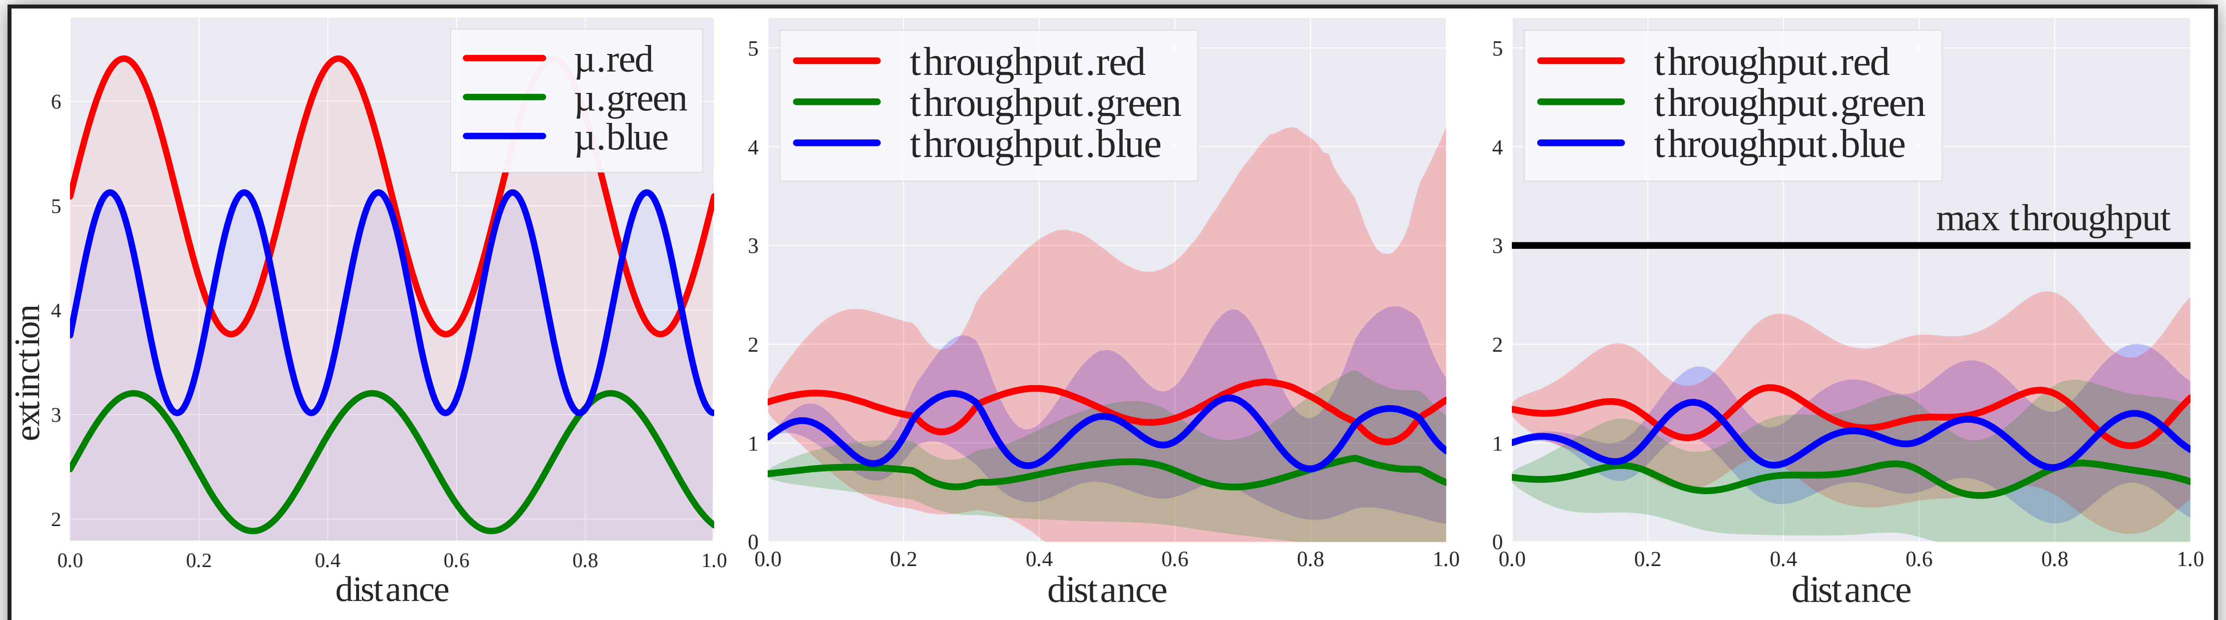

RGB medium | Maximum | History-aware avg.

[KHLN17] Kutz P., Habel R., Li Y. K., Novák J.:
Spectral and decomposition tracking for rendering heterogeneous volumes.
ACM TOG (Proc. SIGGRAPH) 36, 4 (July 2017)

spectral tracking without PDF [KHLN17]

- sample by common majorant $\bar{\mu}$
- how do decide for null collision, scattering, or absorption?
- probability according to $\mu_n(\lambda)$, $\mu_s(\lambda)$, $\mu_a(\lambda)$
  - pick by maximum over λ_i
  - pick by average weighted by spectral path throughput history
  - results in different noise patterns:

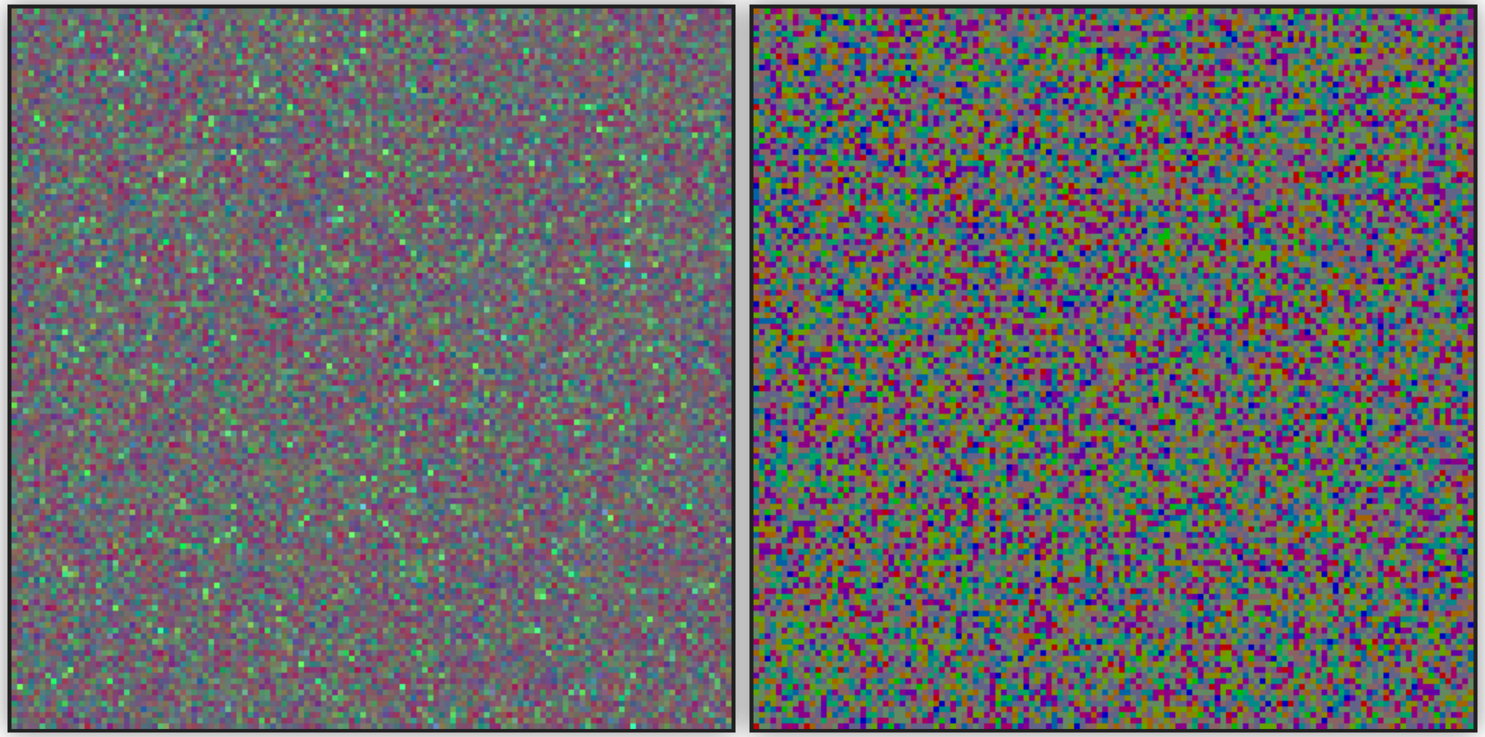

spectral tracking

a few differences:

- sampling optimal for hero wavelength
- sampling carfully balanced for all wavelengths

most important difference:

- require PDF
- cannot provide PDF
- means considering one or the other you need to balance requirements of your system
  - do you need to mix in other importance sampling strategies?
  - for instance equi-angular sampling?

emissive media

thin/dense media make a difference

- no event inside the medium means we cannot pick up emission:

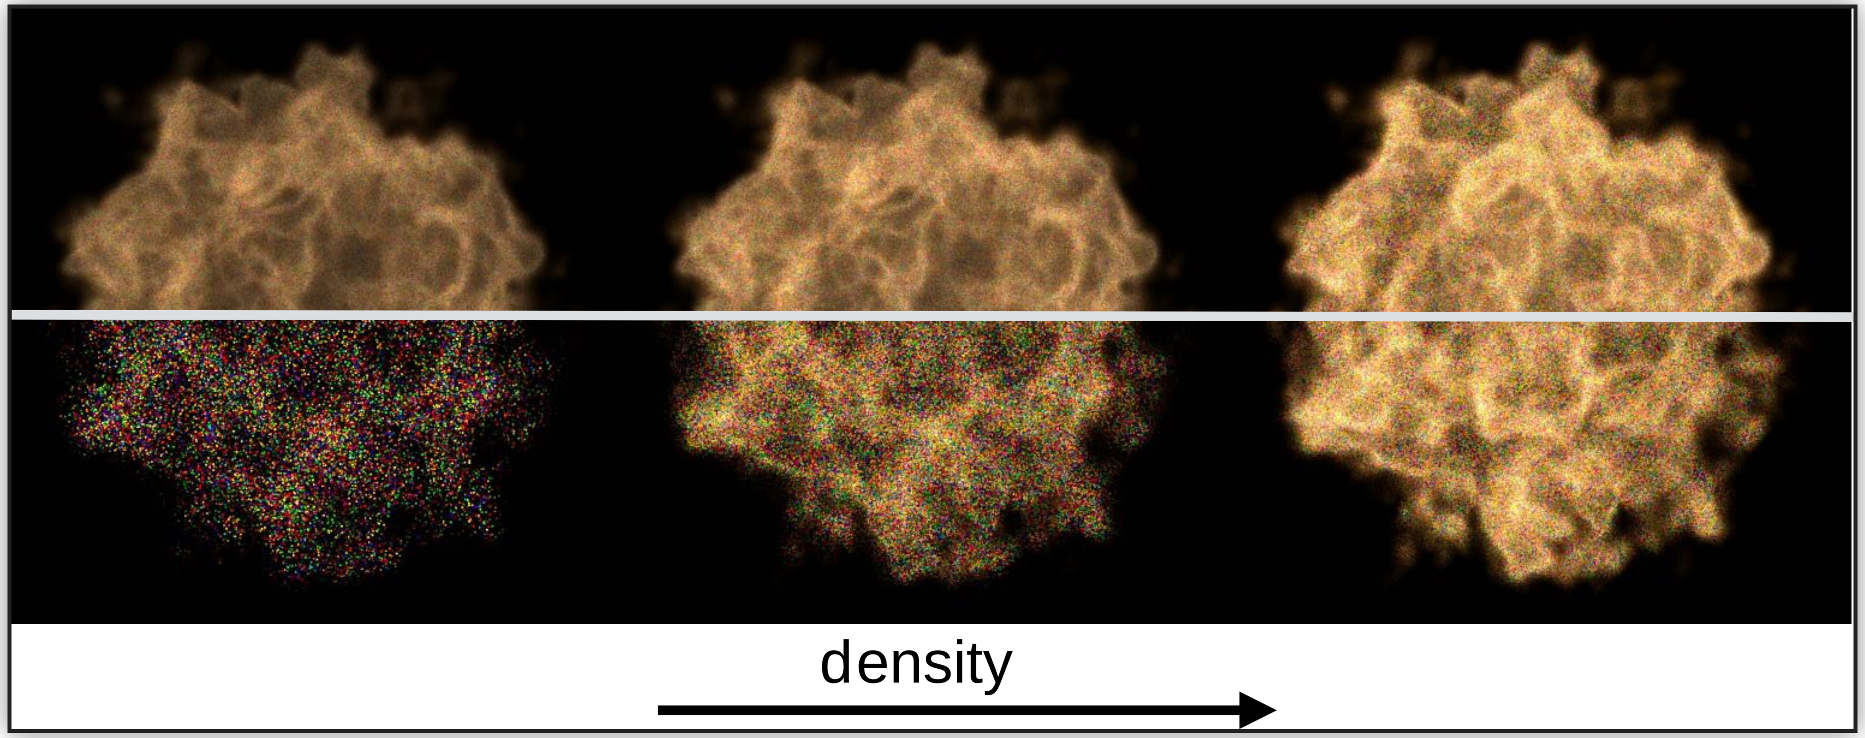

emissive media

thin/dense media make a difference

- follow the idea of beams, collect emission along a ray

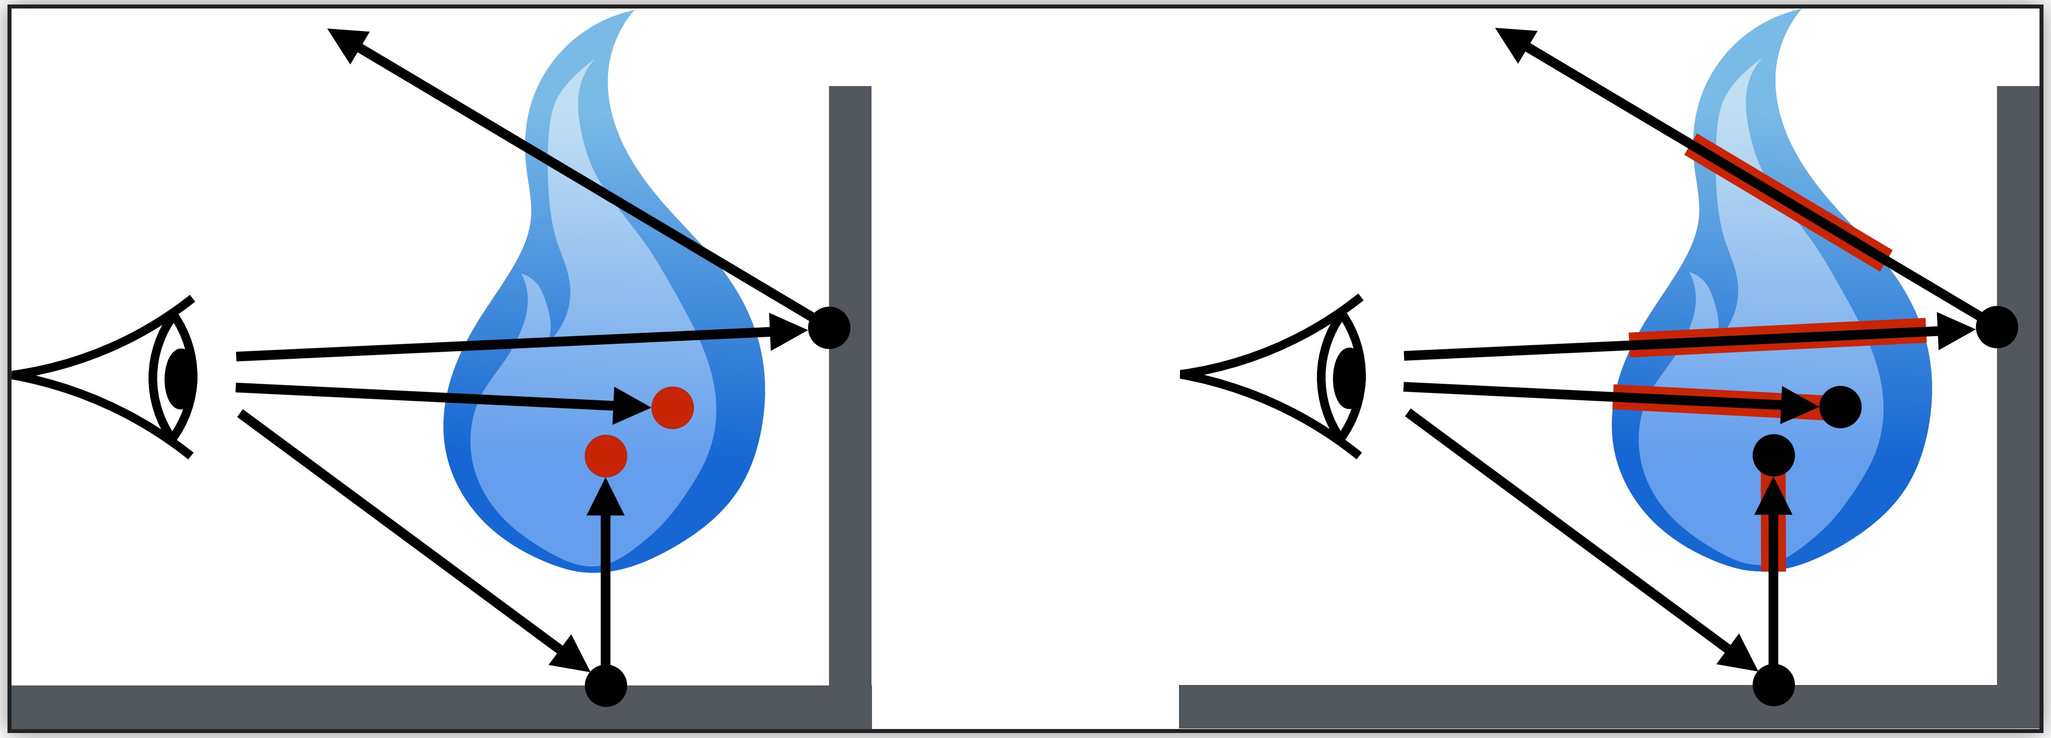

- particularly well suited for regular tracking, touching all voxels anyways

emissive media

thin/dense media make a difference

- direct MIS combination with NEE [VH13] **introduces noise**:

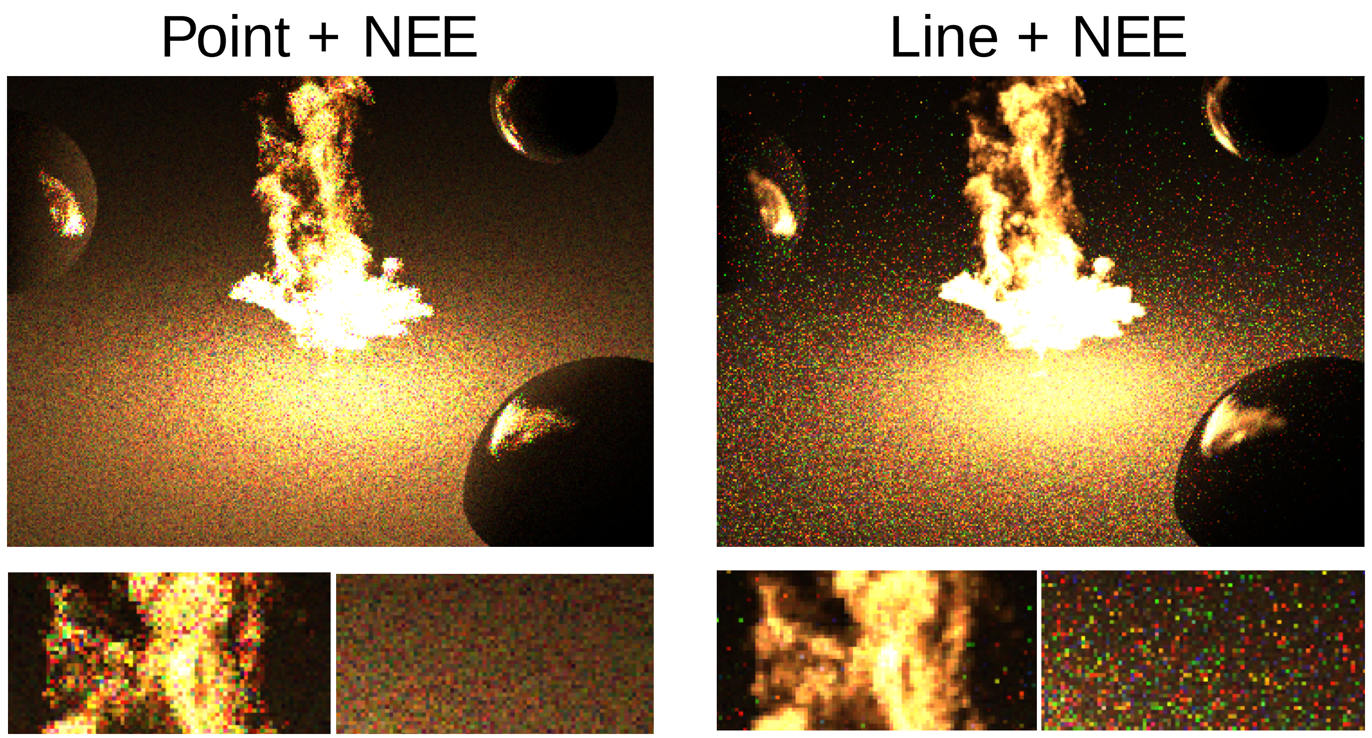

- reason: NEE cannot create paths with end point outside the medium
- forward scattering PDF is poor, however, and now it picks up line emission!

[VH13] Villemin R., Hery C.: Practical illumination from flames.
Journal of Computer Graphics Techniques 2, 2 (2013).

emissive media

where do the additional firefly samples come from?

- example of a problematic path
  - NEE cannot create these paths
  - forward scattering is same bad as usual, but now has a contribution from the segment!

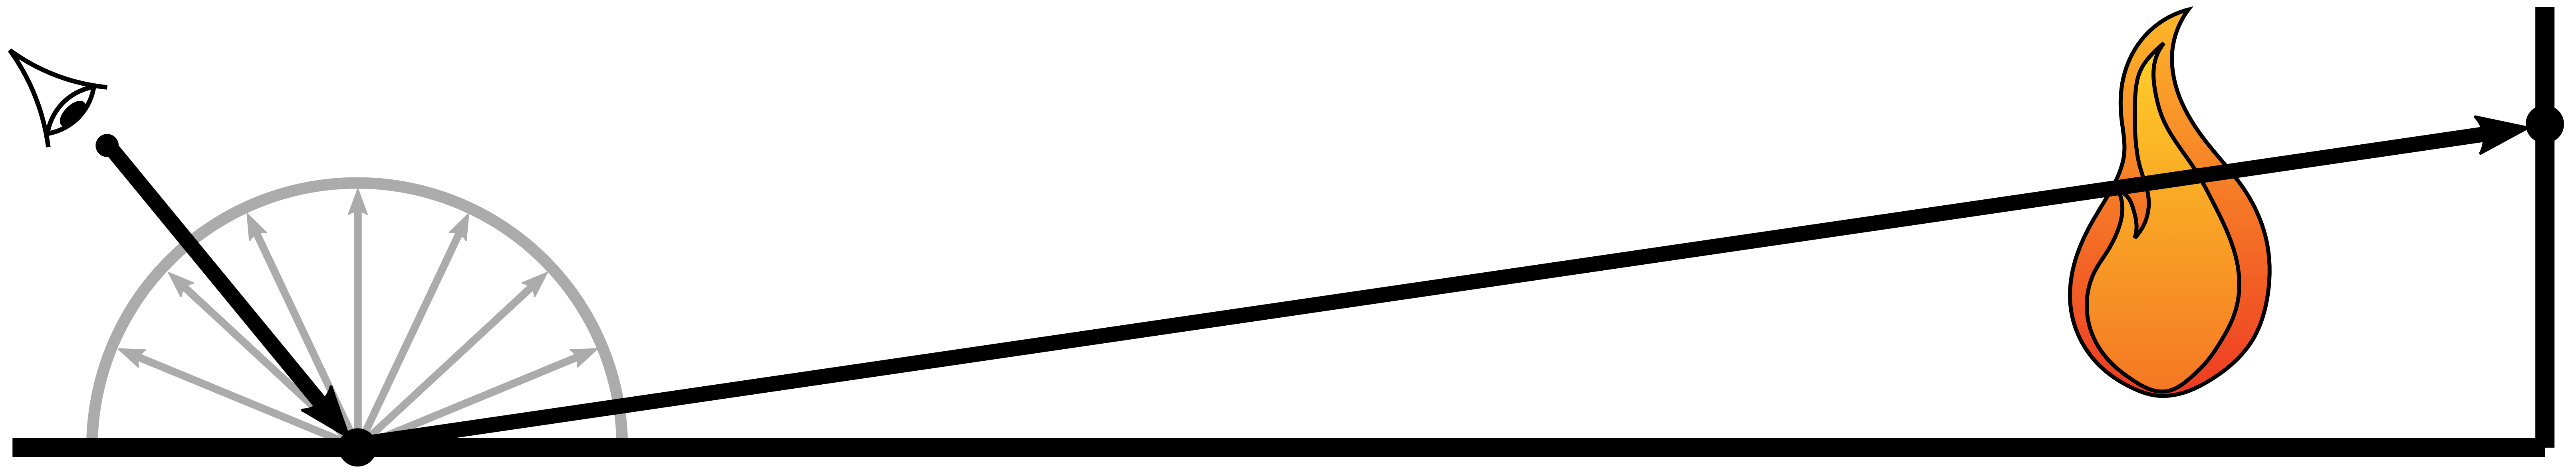

[SHZD17] Simon F., Hanika J., Zirr T., Dachsbacher C.:
Line integration for rendering heterogeneous emissive volumes.
CGF (Proc. EGSR) 36, 4 (June 2017).

emissive media

forward next event estimation (FNEE)

- next event estimation which also considers line emission [SHZD17]:

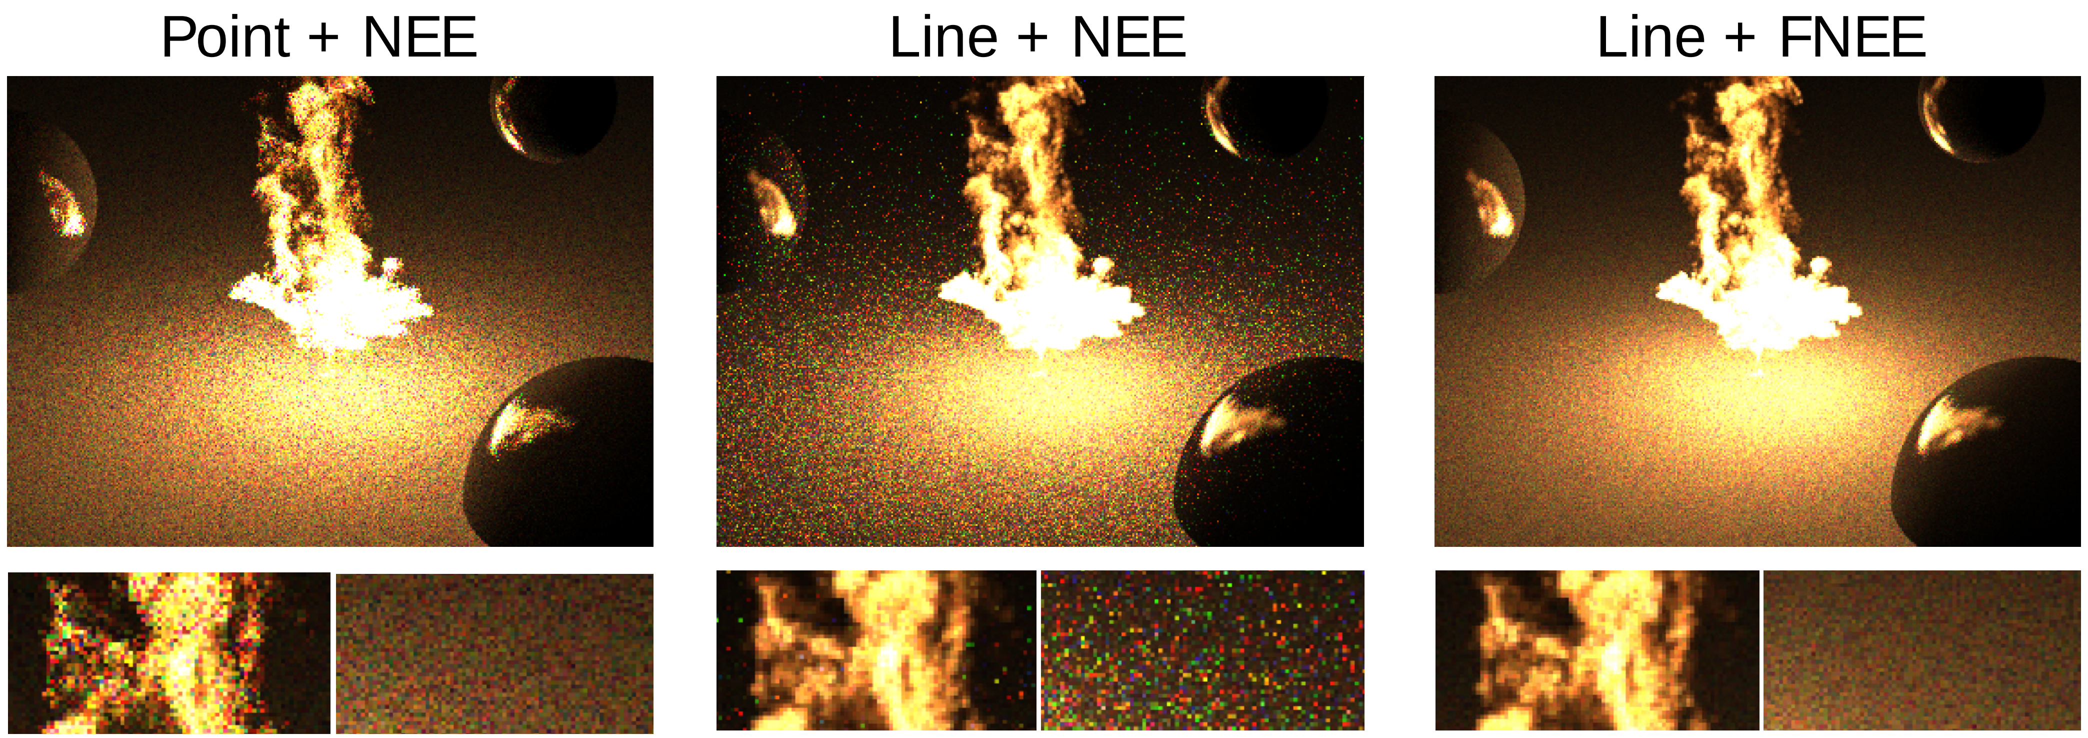

- NEE picks point in medium but only uses the direction to it
- line segment sampling completed by distance sampling

[SHZD17] Simon F., Hanika J., Zirr T., Dachsbacher C.:
Line integration for rendering heterogeneous emissive volumes.
CGF (Proc. EGSR) 36, 4 (June 2017).

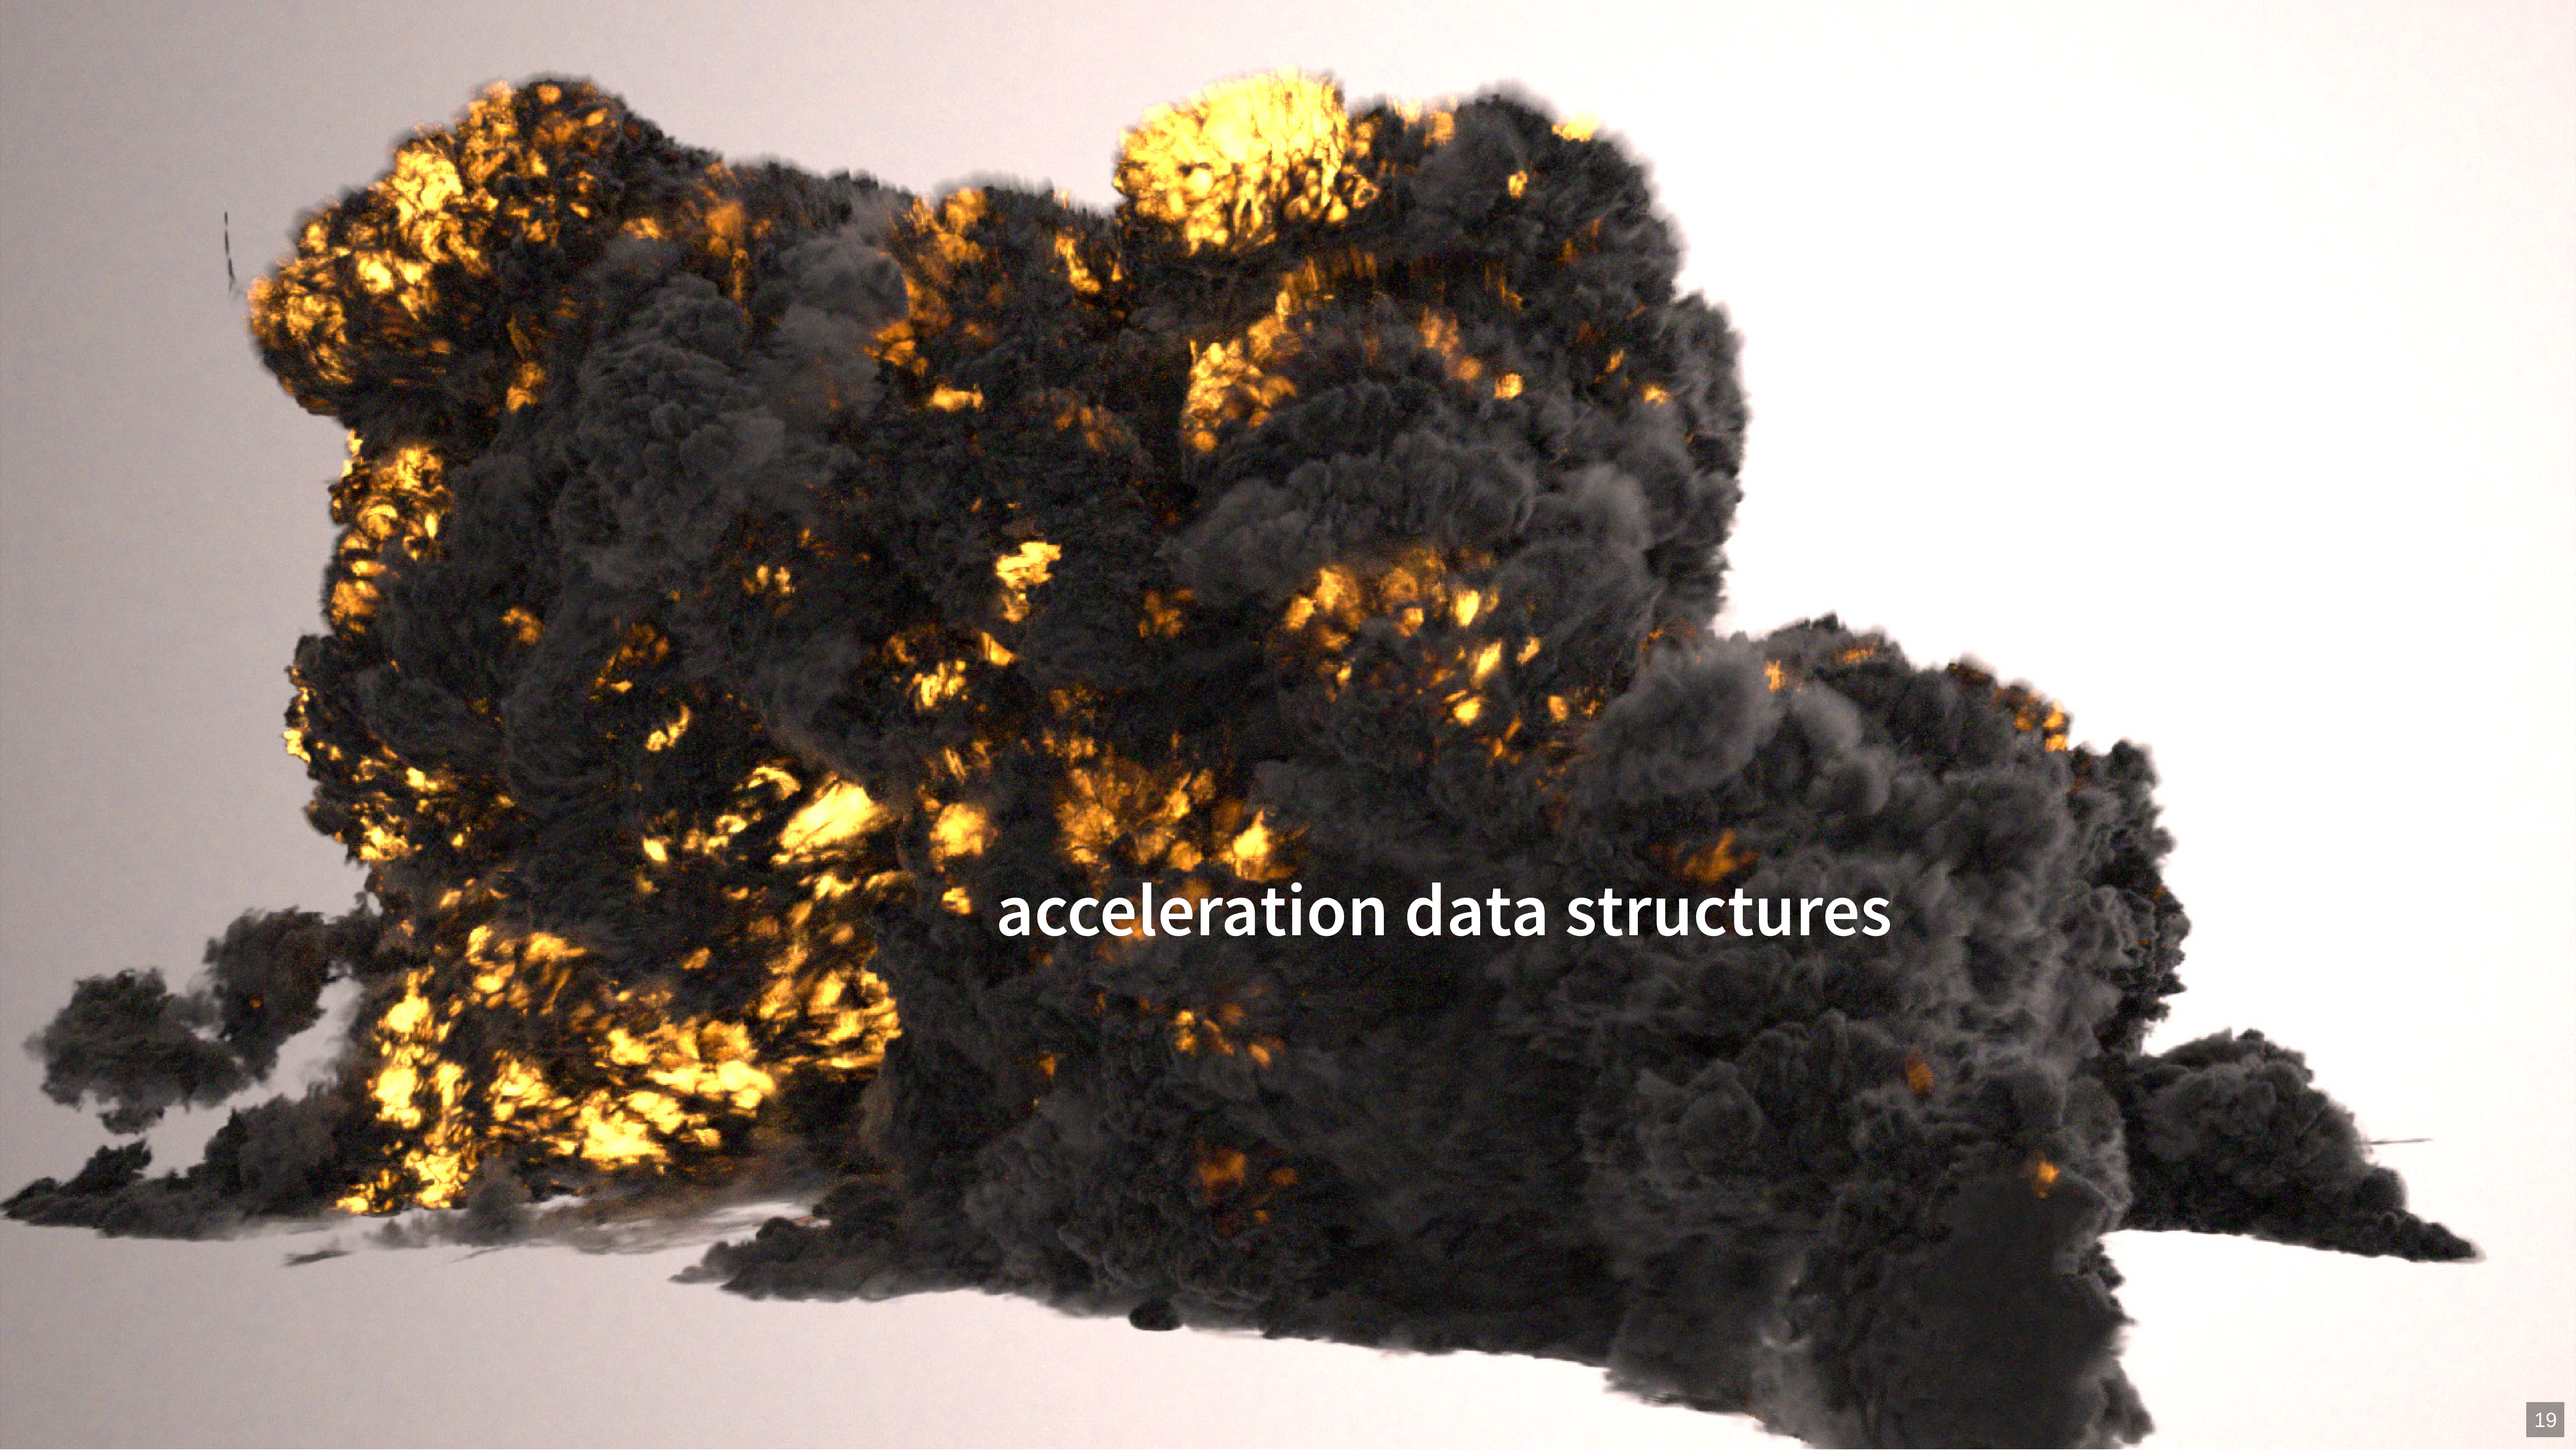
acceleration data structures

acceleration data structures

naive grid

- store coefficients $\mu_\star(\mathbf{x}, \lambda)$ per voxel
- only single (half) float density $\rho(\mathbf{x})$ and global cross sections $\sigma_\star$? If you can!
- motion blur?

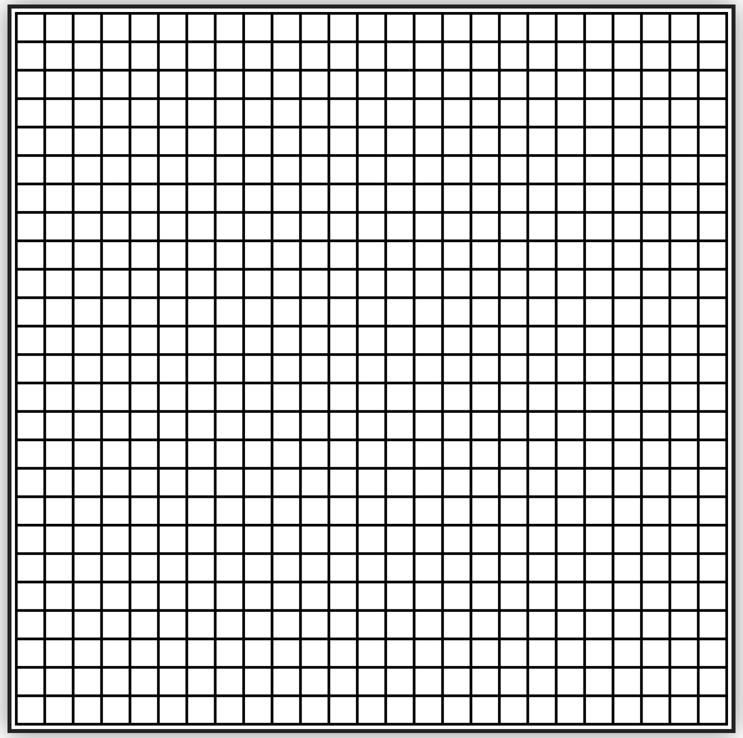

acceleration data structures

motion blur

- store time resolved volume in 4D [Wre16]
- rasterise temporal output of simulation directly

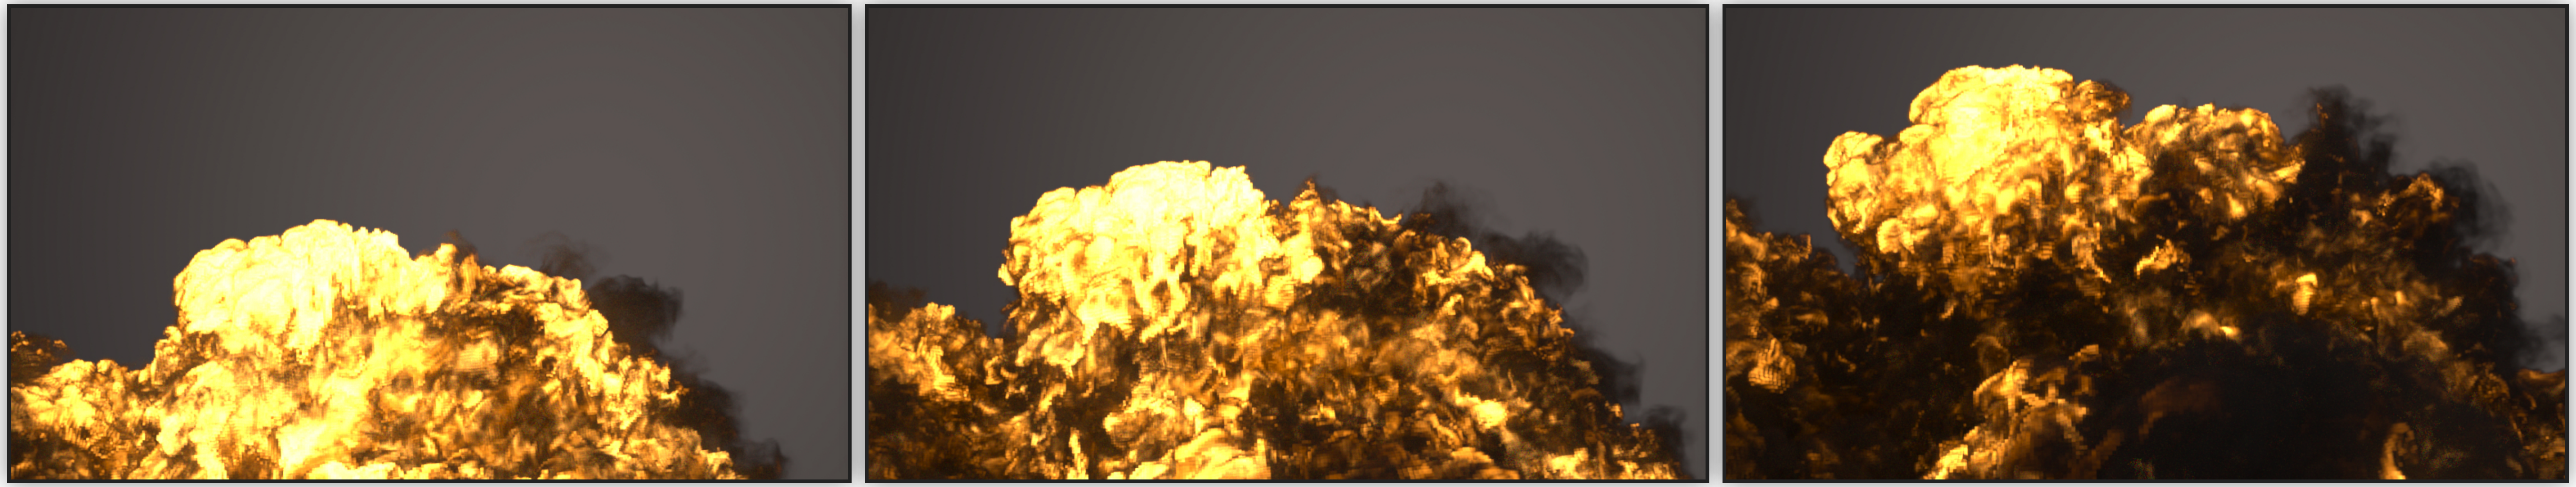

[Wre16] Wrenninge M.: Efficient rendering of volumetric motion blur using temporally unstructured volumes. Journal of Computer Graphics Techniques (JCGT) 5, 1 (January 2016), 1–34.

acceleration data structures

motion blur

- store time resolved volume in 4D [Wre16]
- Ramer Douglas Peucker line compression per voxel supports non-linear motion
  - for instance up: density $\rho(\mathbf{x}, t)$, right: time t:

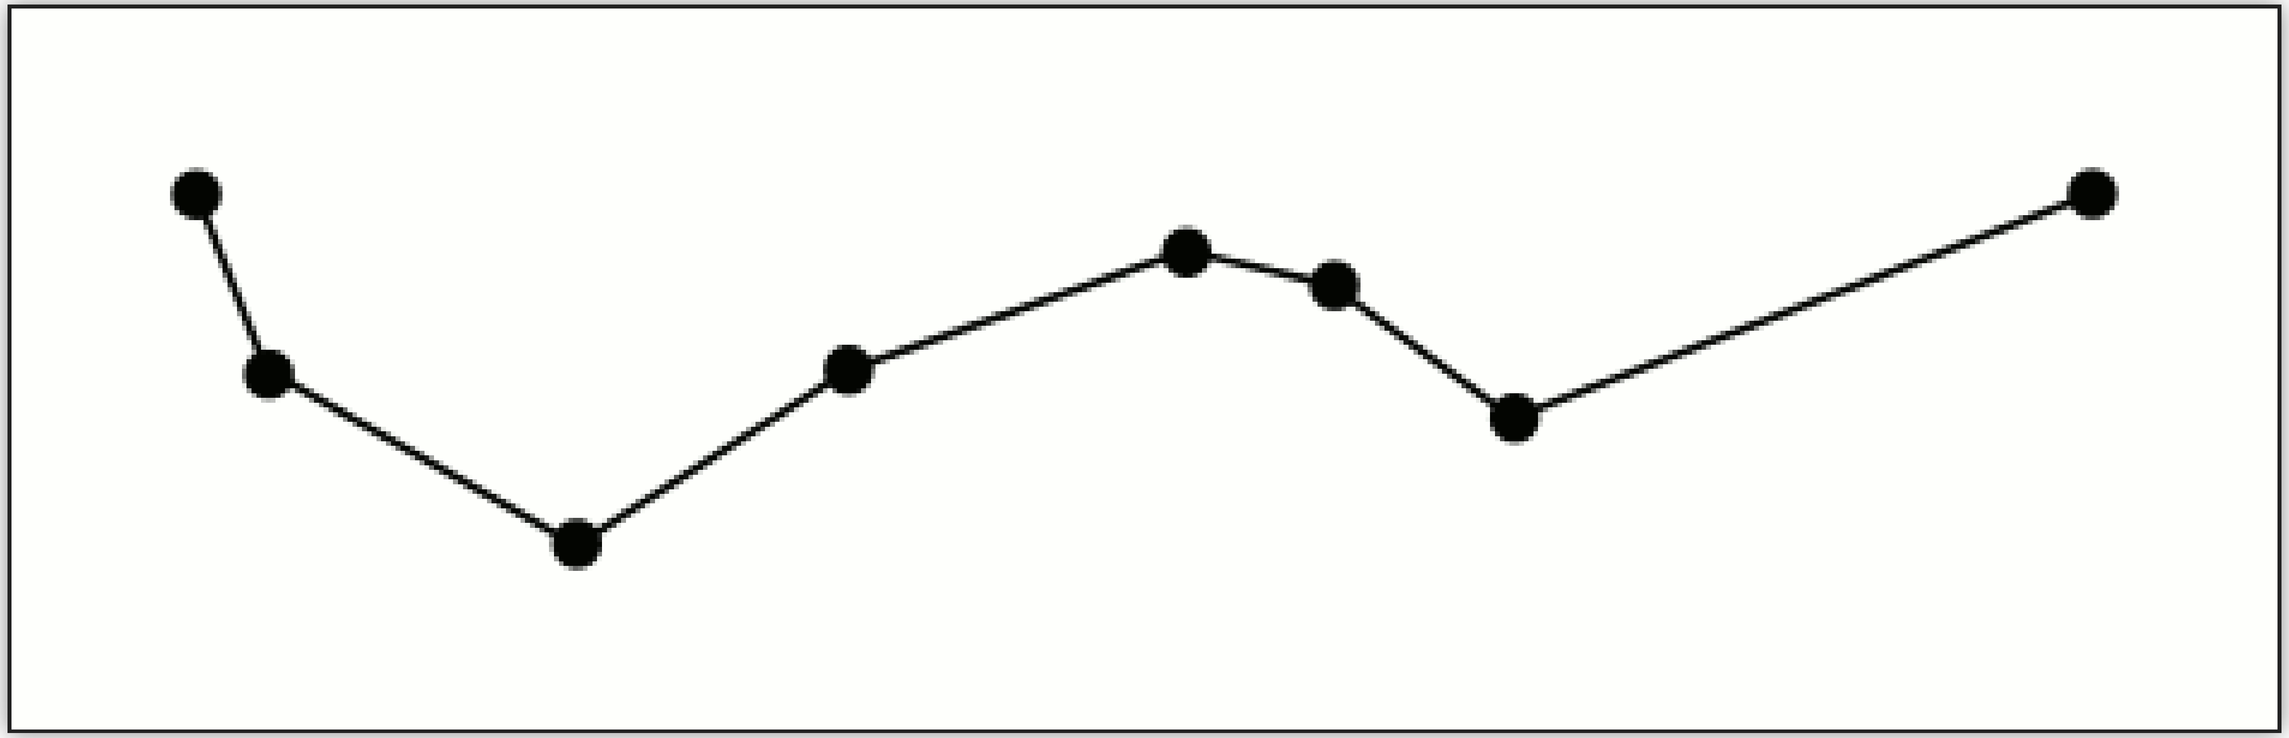

https://en.wikipedia.org/wiki/Ramer%E2%80%93Douglas%E2%80%93Peucker_algorithm

acceleration data structures

motion blur

- store time resolved volume in 4D [Wre16]
- supports non-linear motion better than Eulerian motion blur [KK07]


[Wre16] Wrenninge M.: Efficient rendering of volumetric motion blur
using temporally unstructured volumes. Journal of Computer Graphics
Techniques (JCGT) 5, 1 (January 2016), 1–34.

[KK07] Kim D., Ko H.-S.: Eulerian motion blur.
In Eurographics Workshop on Natural Phenomena (2007).

acceleration data structures

hierarchical grid (aka super voxels [SKTM11])

- hierarchical traversal, local memory access
- well suited for null collision based trackers:
  - store majorants on coarse grid, fine data on fine level

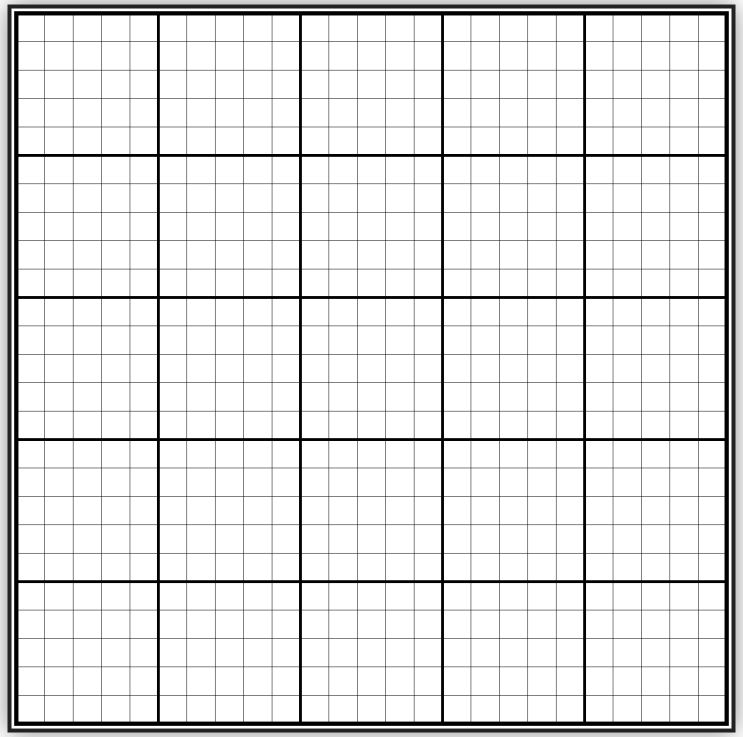

[SKTM11] Szirmay-Kalos L., Tóth B., Magdics M.:
Free path sampling in high resolution inhomogeneous participating media.
CGF 30, 1 (2011), 85–97.

acceleration data structures

regular tracking

- multi-level 3D DDA/Bresenham to enumerate all voxels pierced by ray
- example: free flight distance sampling

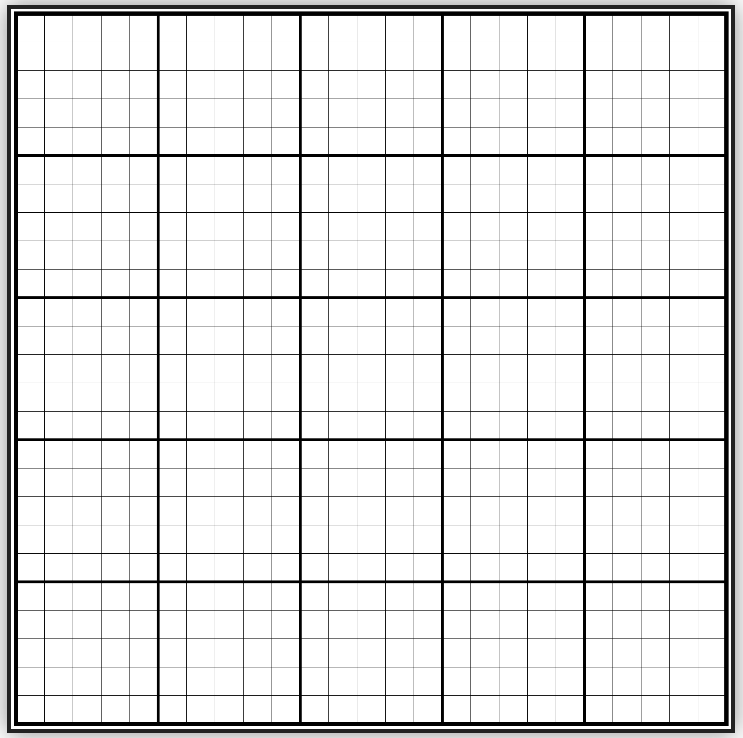

acceleration data structures

regular tracking

- multi-level 3D DDA/Bresenham to enumerate all voxels pierced by ray
- walk along ray direction

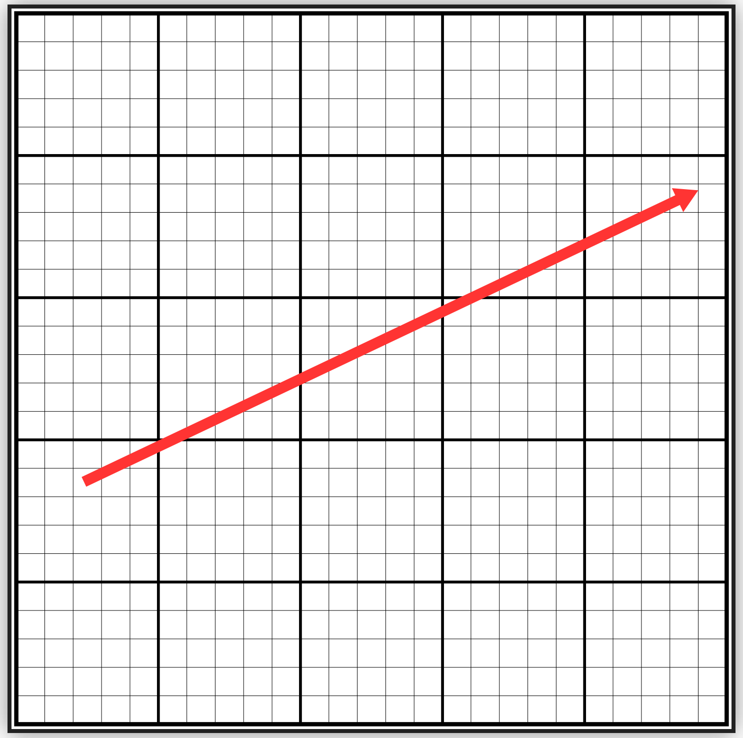

acceleration data structures

regular tracking

- multi-level 3D DDA/Bresenham to enumerate all voxels pierced by ray
- memory access to super voxel block

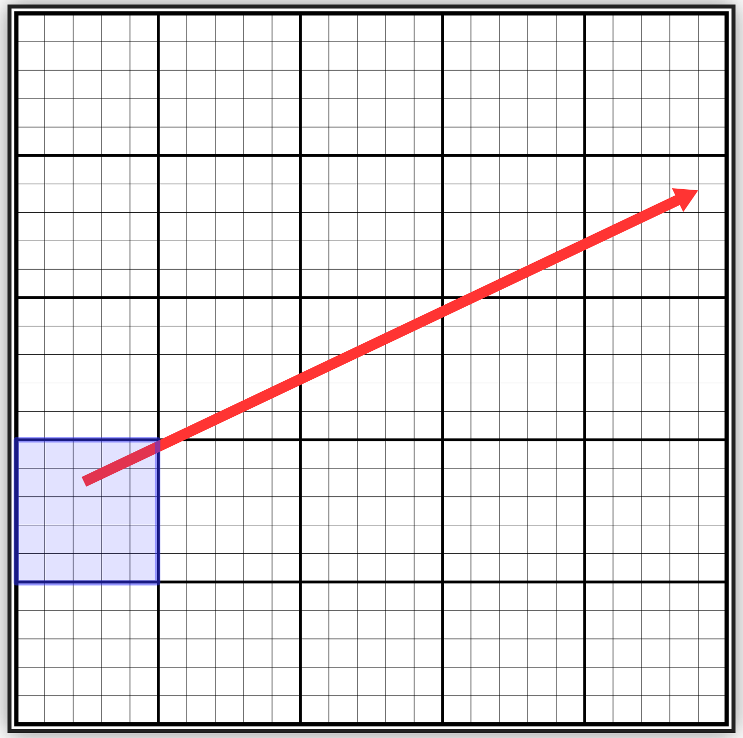

acceleration data structures

regular tracking

- memory access to voxels stored in brick
- all voxels are typically local in memory now

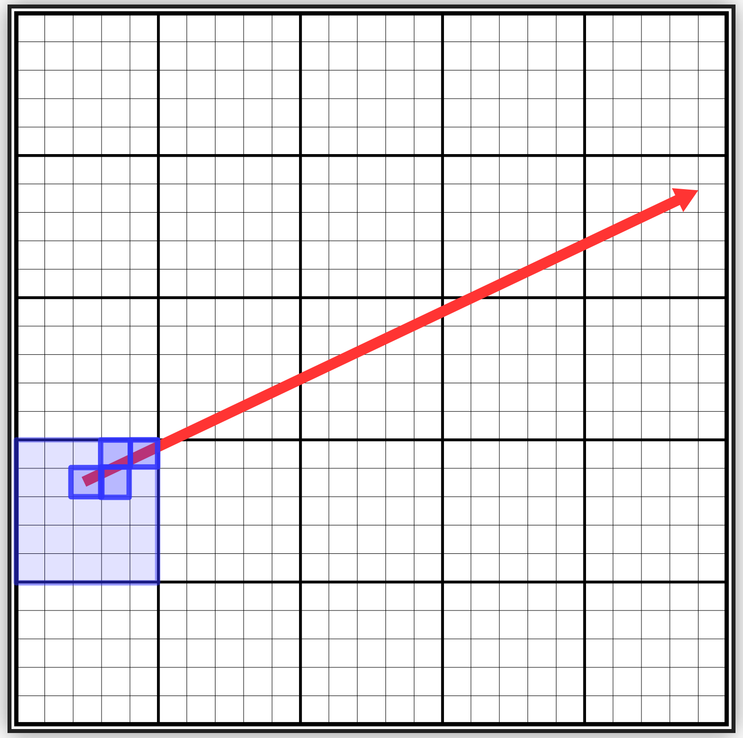

acceleration data structures

regular tracking

- if no intersection is found in this block
- continue to next block

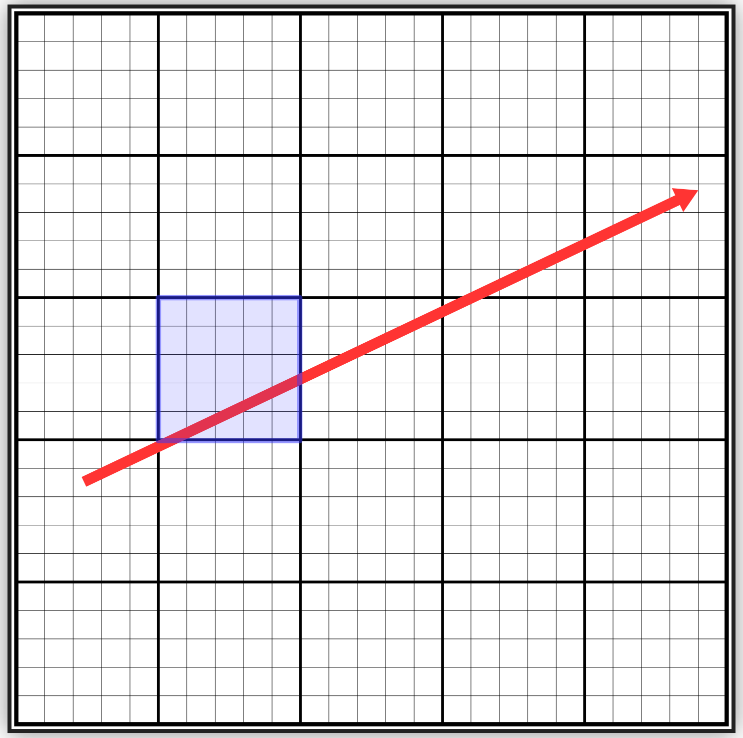

acceleration data structures

regular tracking

- accumulate optical thickness $\tau \leftarrow \tau + \rho(\mathbf{x}) \cdot \sigma_t$
- if $\tau >= -\log{(1-\xi)}$ found free flight distance, sub voxel accuracy by assuming homogeneous voxel

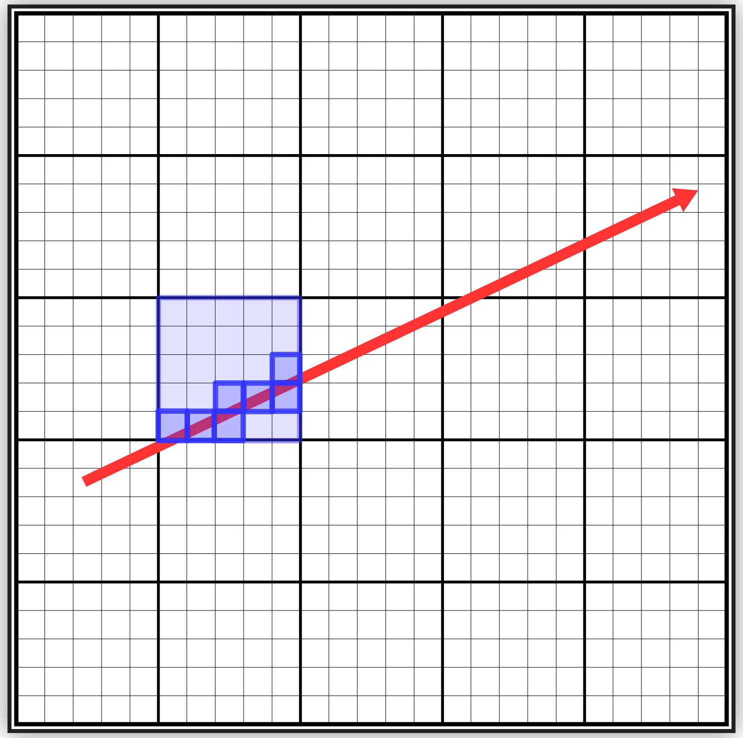

acceleration data structures

null collision tracking

- perform **regular tracking** on the coarse level [SKTM11]
- start of the algorithm very much the same as regular tracking

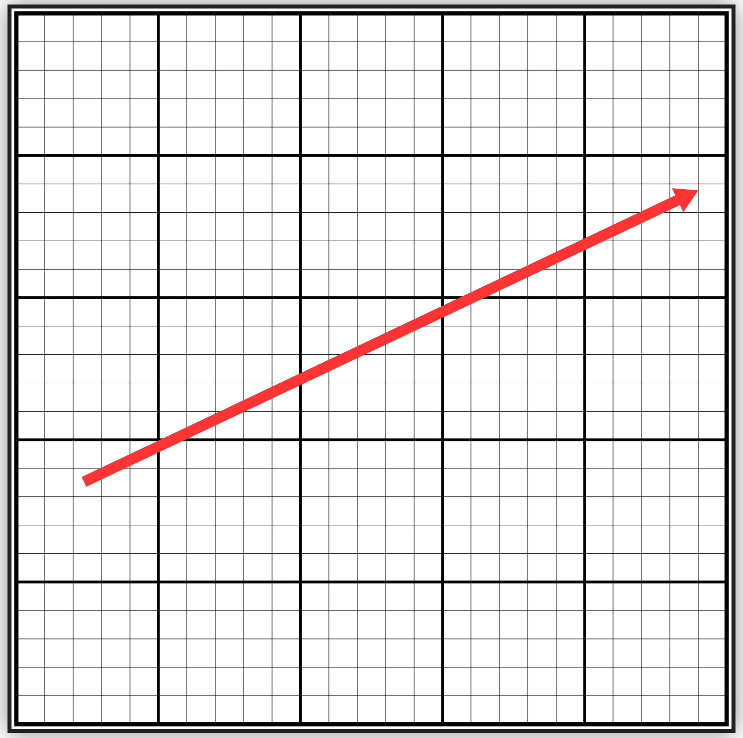

acceleration data structures

null collision tracking

- perform **regular tracking** on the coarse level [SKTM11]
- accumulate fictitious optical thickness $\tau \leftarrow \tau + \bar{\rho}(\mathbf{x}) \cdot \sigma_t$

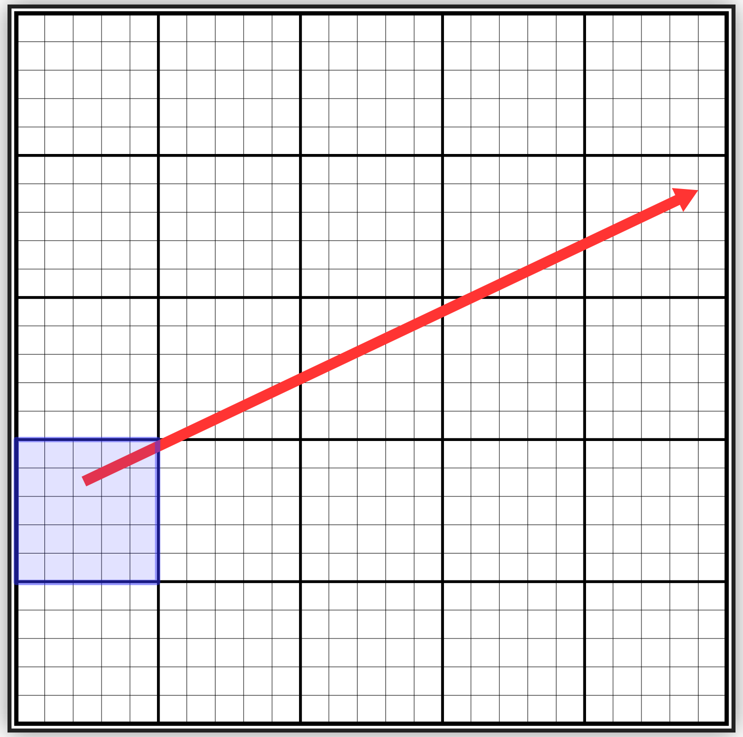

acceleration data structures

null collision tracking

- perform **regular tracking** on the coarse level [SKTM11]
- if $\tau >= -\log(1-\xi)$ found free flight distance inside this block

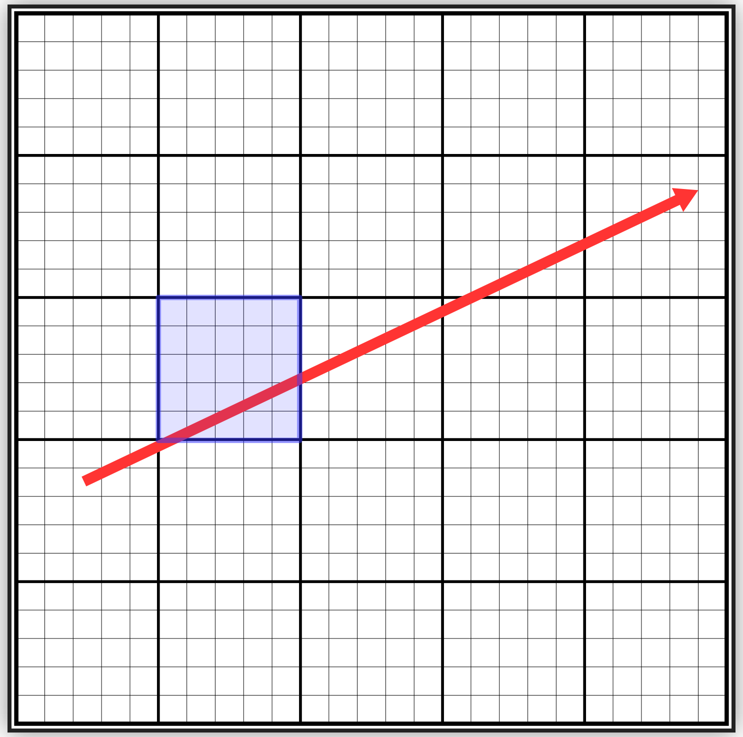

acceleration data structures

null collision tracking

- if $\tau >= -\log(1-\xi)$ found free flight distance inside this block
- find voxel position $\mathbf{x}$ assuming homogeneous fictitious matter in this block

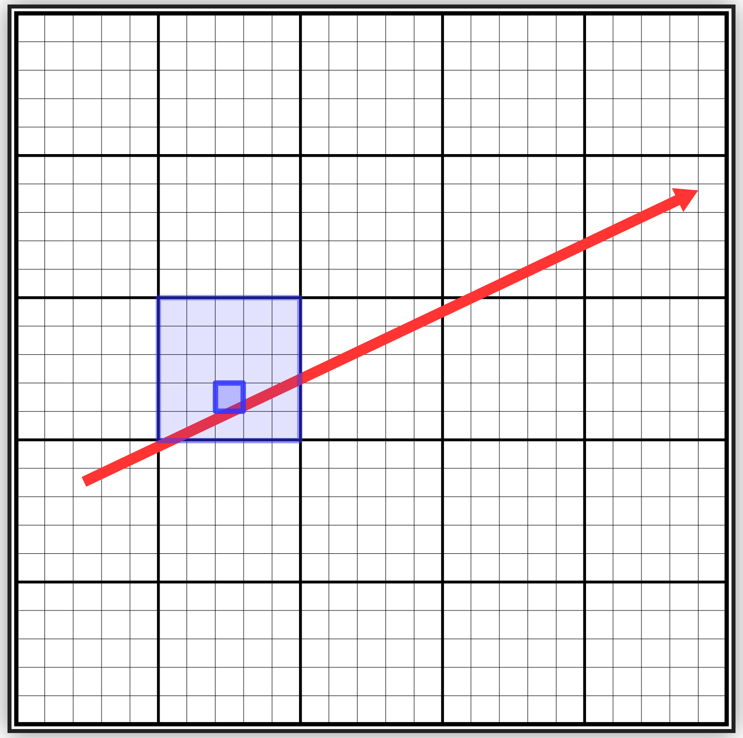

acceleration data structures

null collision tracking

- find voxel position $\mathbf{x}$ assuming homogeneous fictitious matter in this block
- terminate if actual collision, i.e. $\zeta < \mu_t(\mathbf{x})/\bar{\mu}_t(\mathbf{x})$ (else restart tracking)

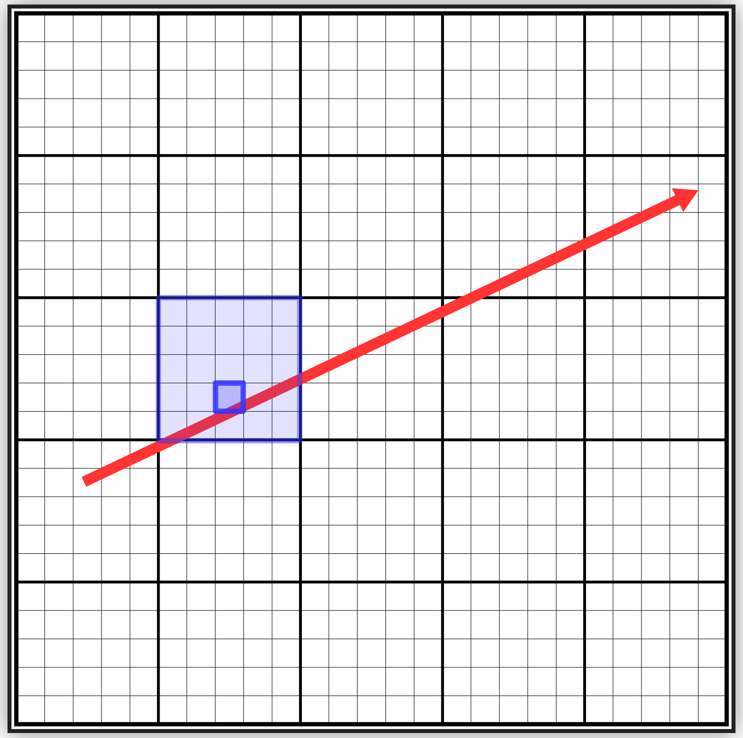

- this last point is critical for performance, we'll get back to this!

acceleration data structures

is a uniform grid the best we can do?

- regular tracking:
  - minimise visible integration error:
  - larger voxels in uniform areas, out of frustum, and in depth
- null collisions:
  - minimise amount of fictitious matter
  - means minimise steps taken until a real event is found
- end result:
  - fewer, larger voxels in uniform areas, empty space culled away
  - null collisions only care about coarse level

acceleration data structures

kd-trees

- space partitioning scheme by finding *largest empty rectangle* [YIC*11]

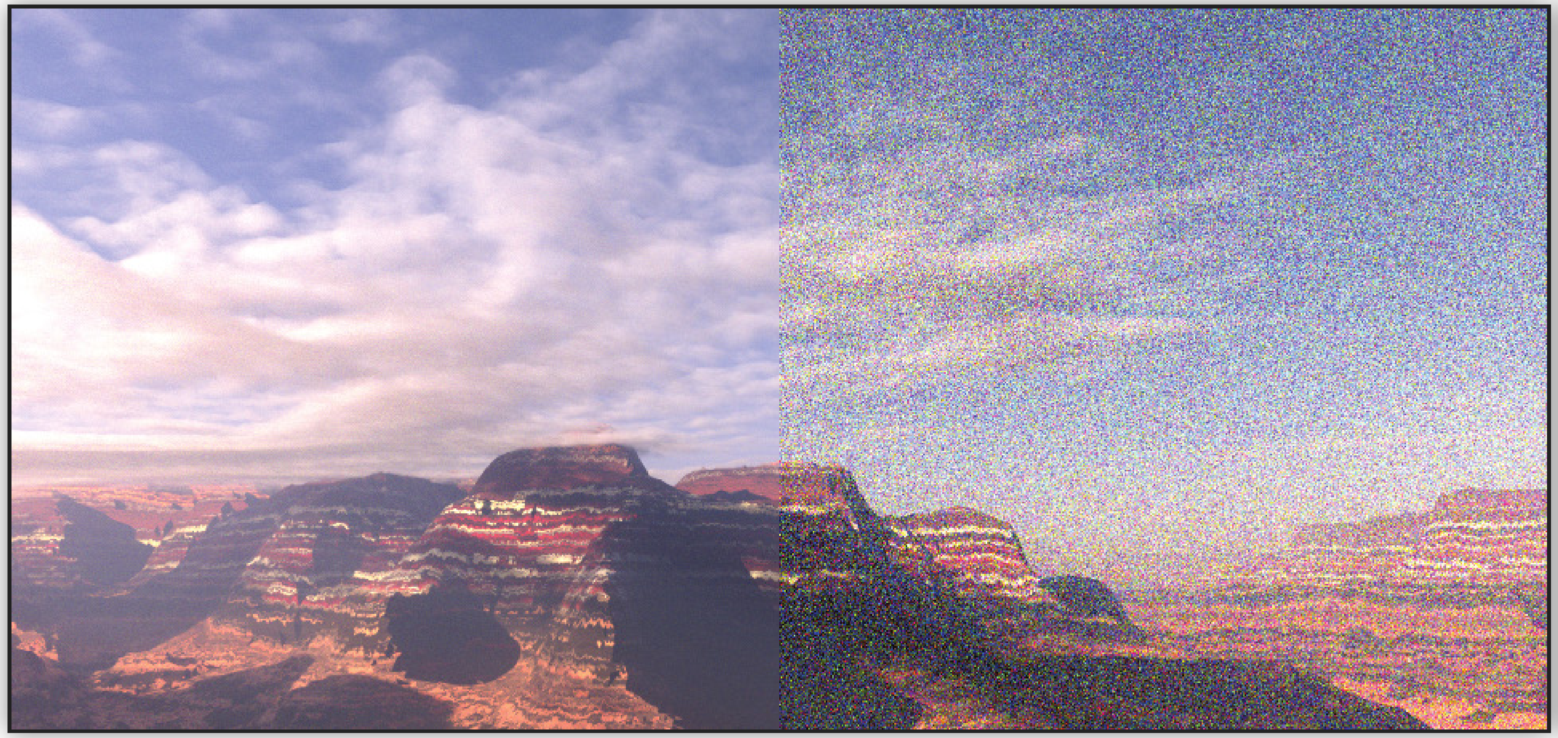
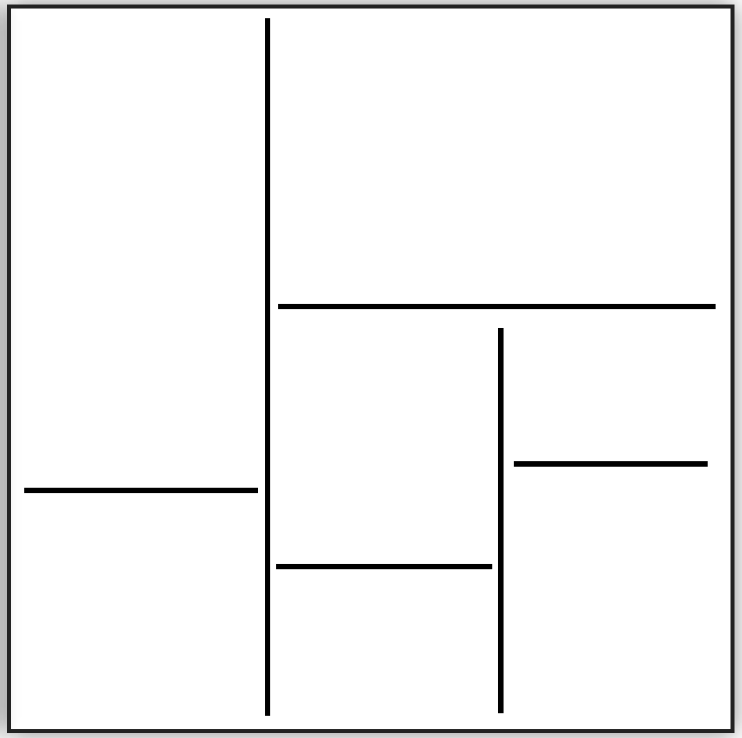

[YIC∗11] Yue Y., Iwasaki K., Chen B., Dobashi Y., Nishita T.:
Toward optimal space partitioning for unbiased, adaptive free path sampling of inhomogeneous participating media.
CGF (Proc. Pacific Graphics) 30, 7 (2011), 1911–1919.

acceleration data structures

- combine super voxels [SKTM11] and kd-trees [YIC∗11] with adaptive blocks

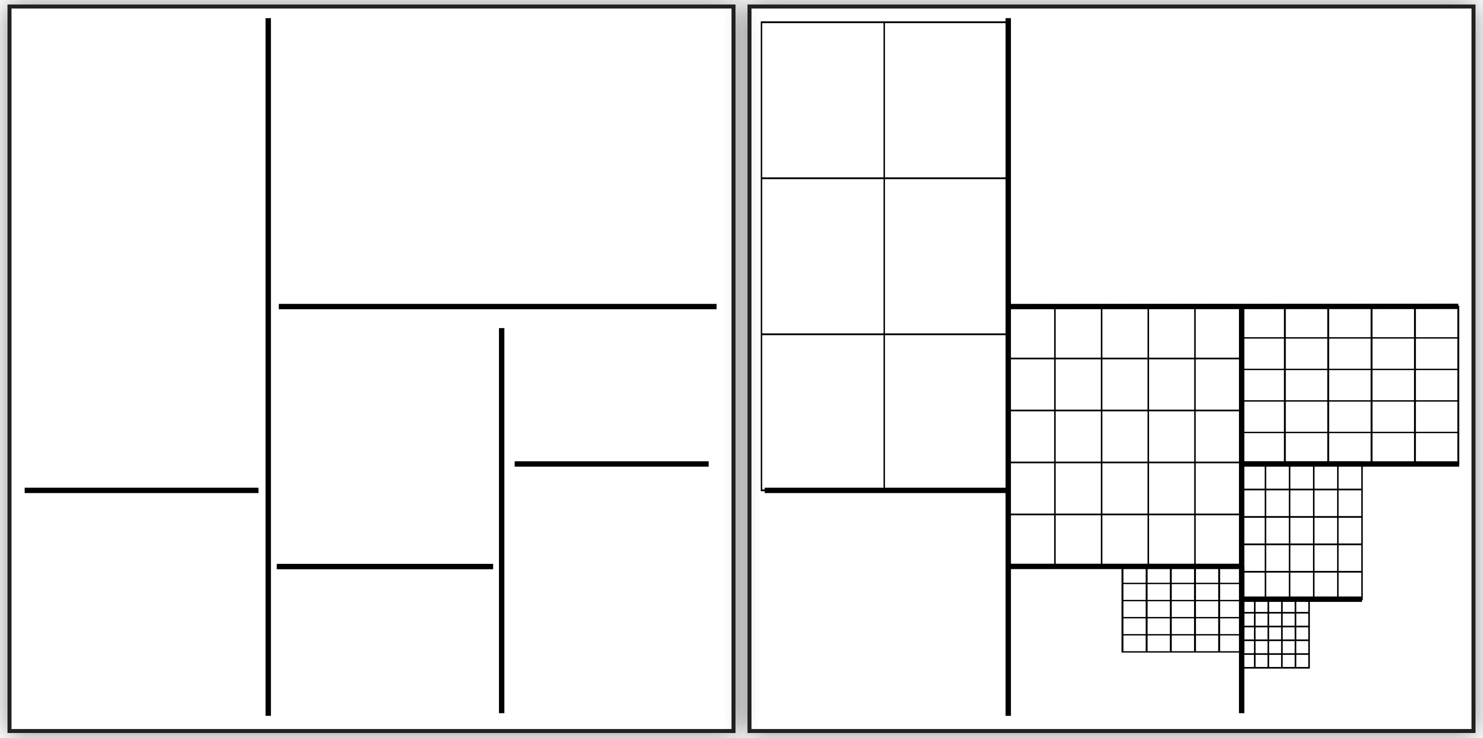

- adaptivity driven by
  - pixel footprint / camera tessellation
  - heterogeneity / variation
- just as super voxels: kd nodes store majorants $\bar{\mu}$ in coarse blocks

acceleration data structures

- combine super voxels [SKTM11] and kd-trees [YIC∗11] with adaptive blocks

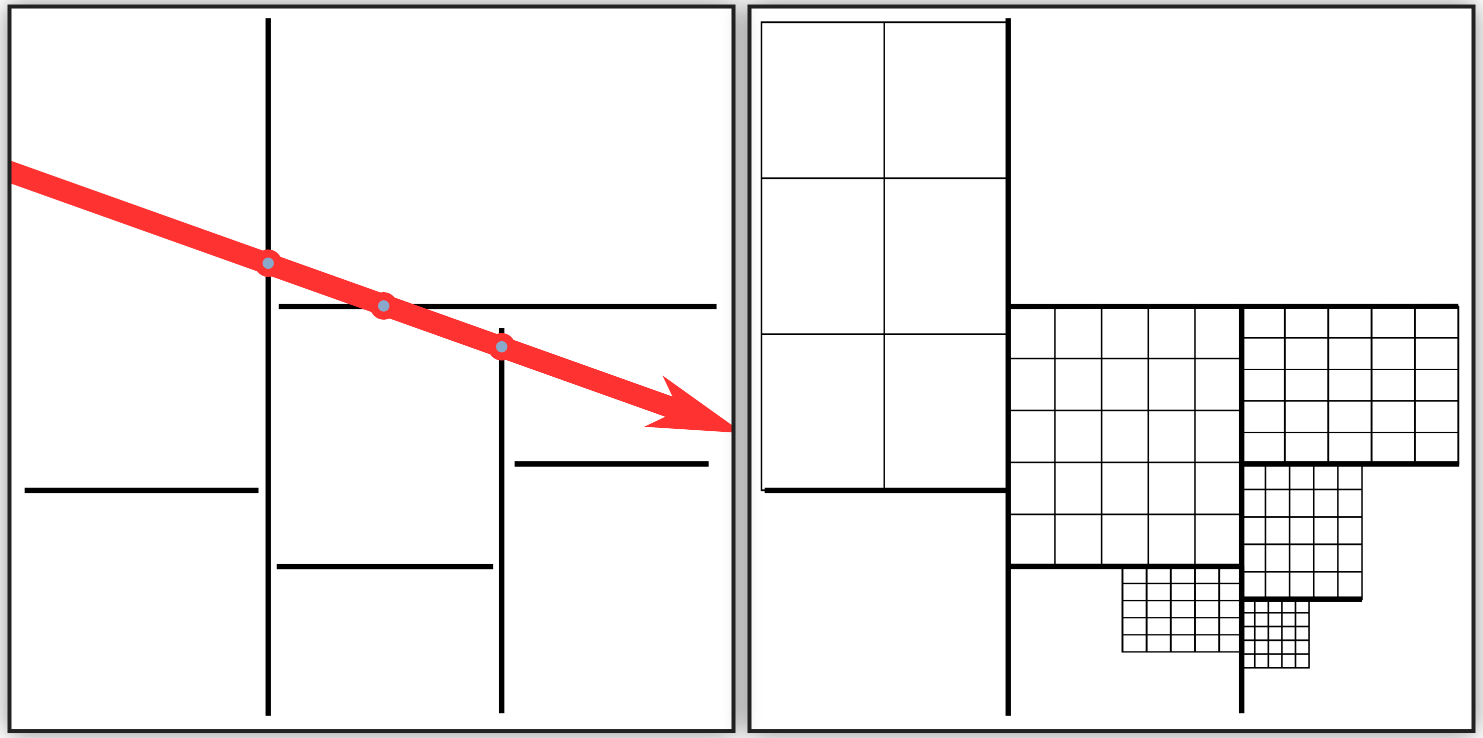

- adaptivity driven by
  - pixel footprint / camera tessellation
  - heterogeneity / variation
- just as super voxels: kd nodes store majorants $\bar{\mu}$ in coarse blocks
  - perform **regular tracking** on coarse blocks [SKTM11]

acceleration data structures

- combine super voxels [SKTM11] and kd-trees [YIC∗11] with adaptive blocks

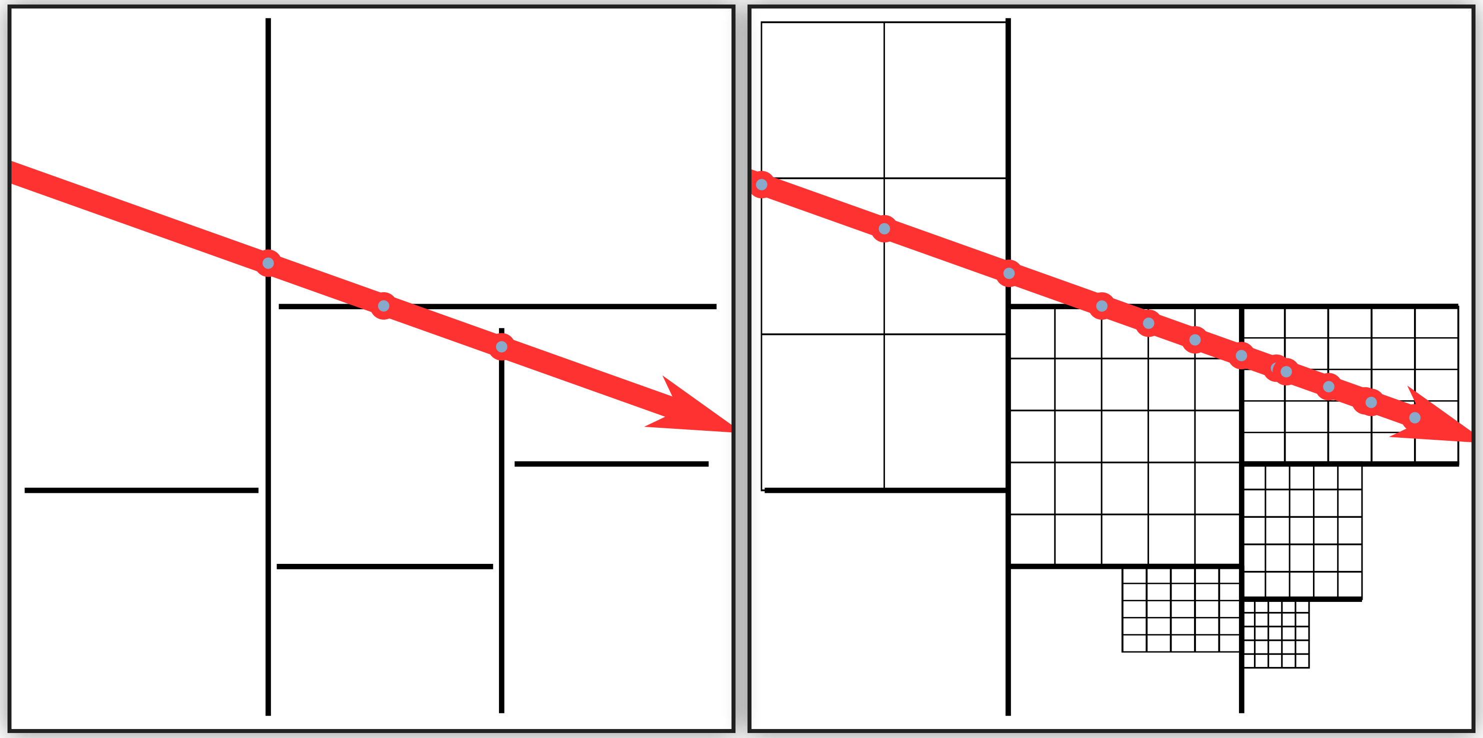

- adaptivity driven by
  - pixel footprint / camera tessellation
  - heterogeneity / variation
- just as super voxels: kd nodes store majorants $\bar{\mu}$ in coarse blocks
  - perform **regular tracking** on coarse blocks [SKTM11]
  - access $\mu_s(\lambda), \mu_a(\lambda)$ on fine levels to sample collision type

acceleration data structures

regular tracking

- needs to step through *every* voxel, bad for fine tessellations
- well chosen tessellation is a big advantage!

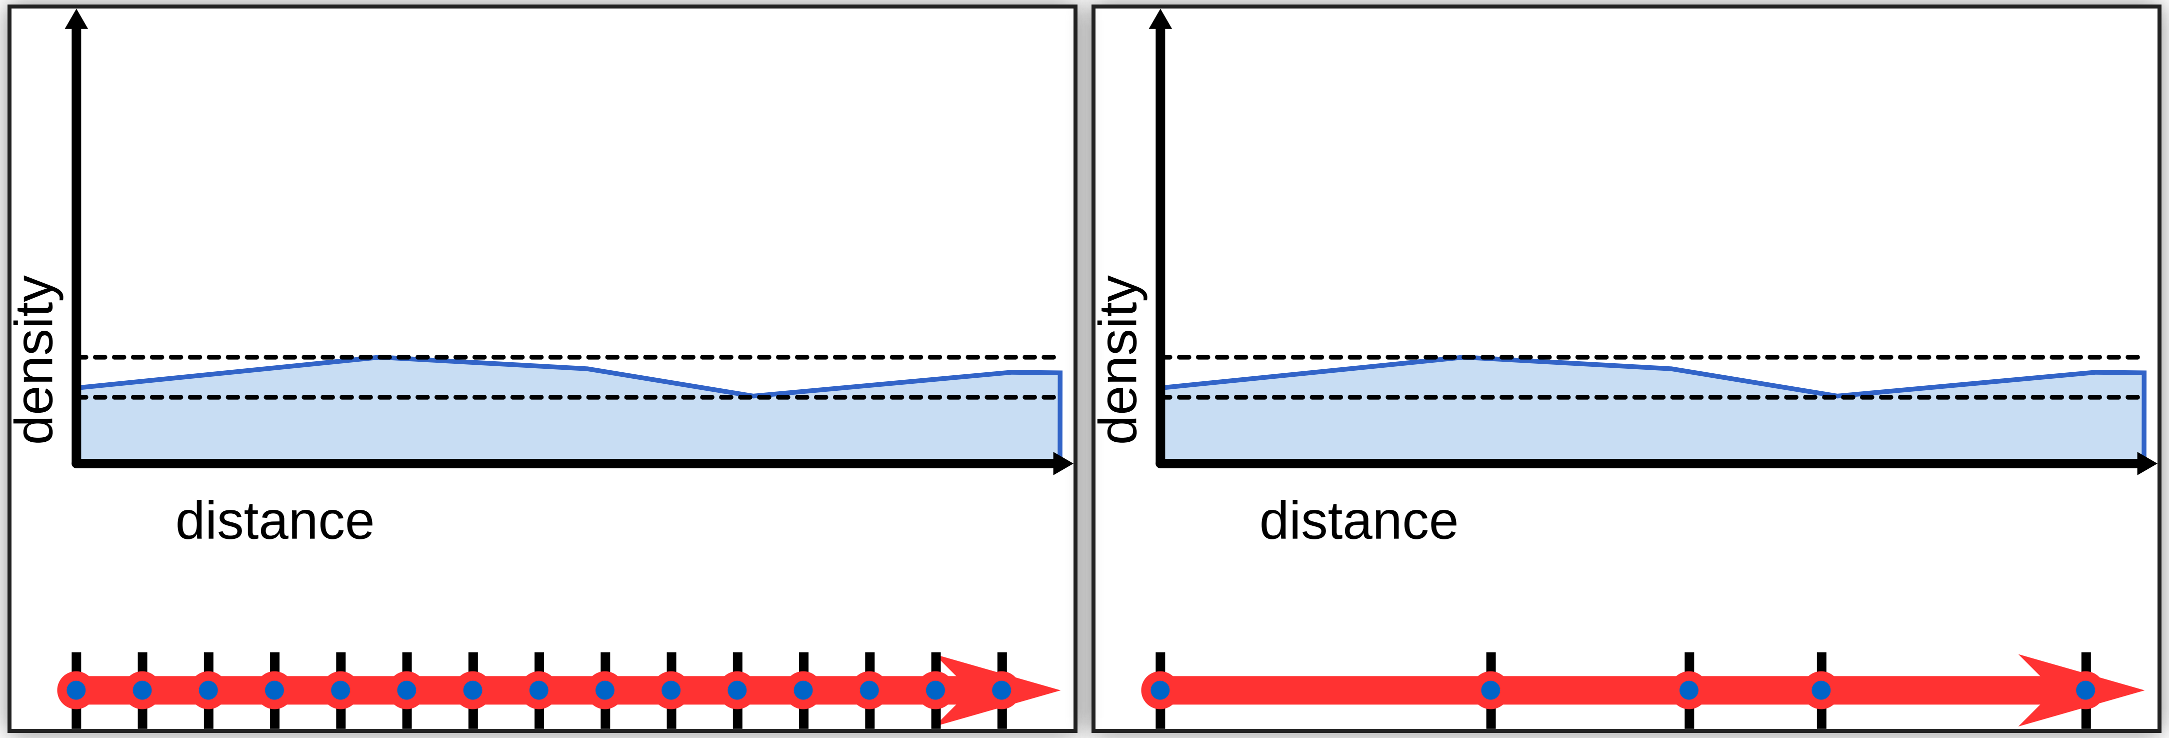

acceleration data structures

null collision-based tracking

- independent of tessellation
- efficient in tightly bounded, thin media (mean free path longer than voxel width)

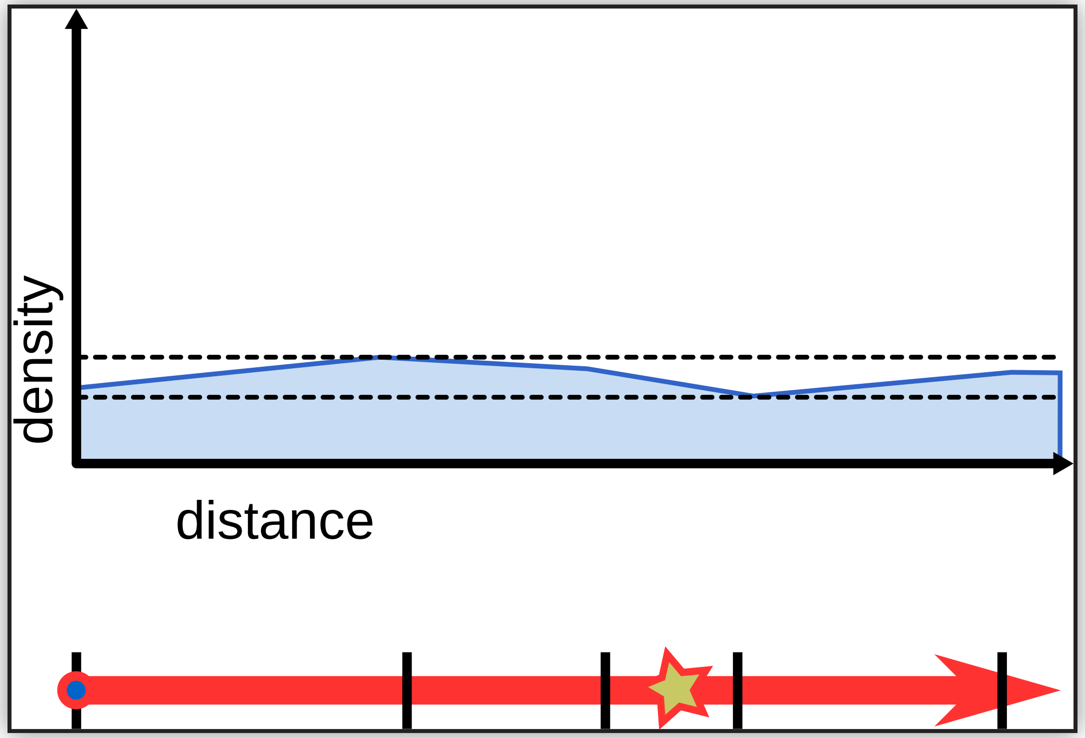

acceleration data structures

null collision-based tracking

- independent of tessellation
- inefficient for loose bounds, avoid during hierarchy construction!

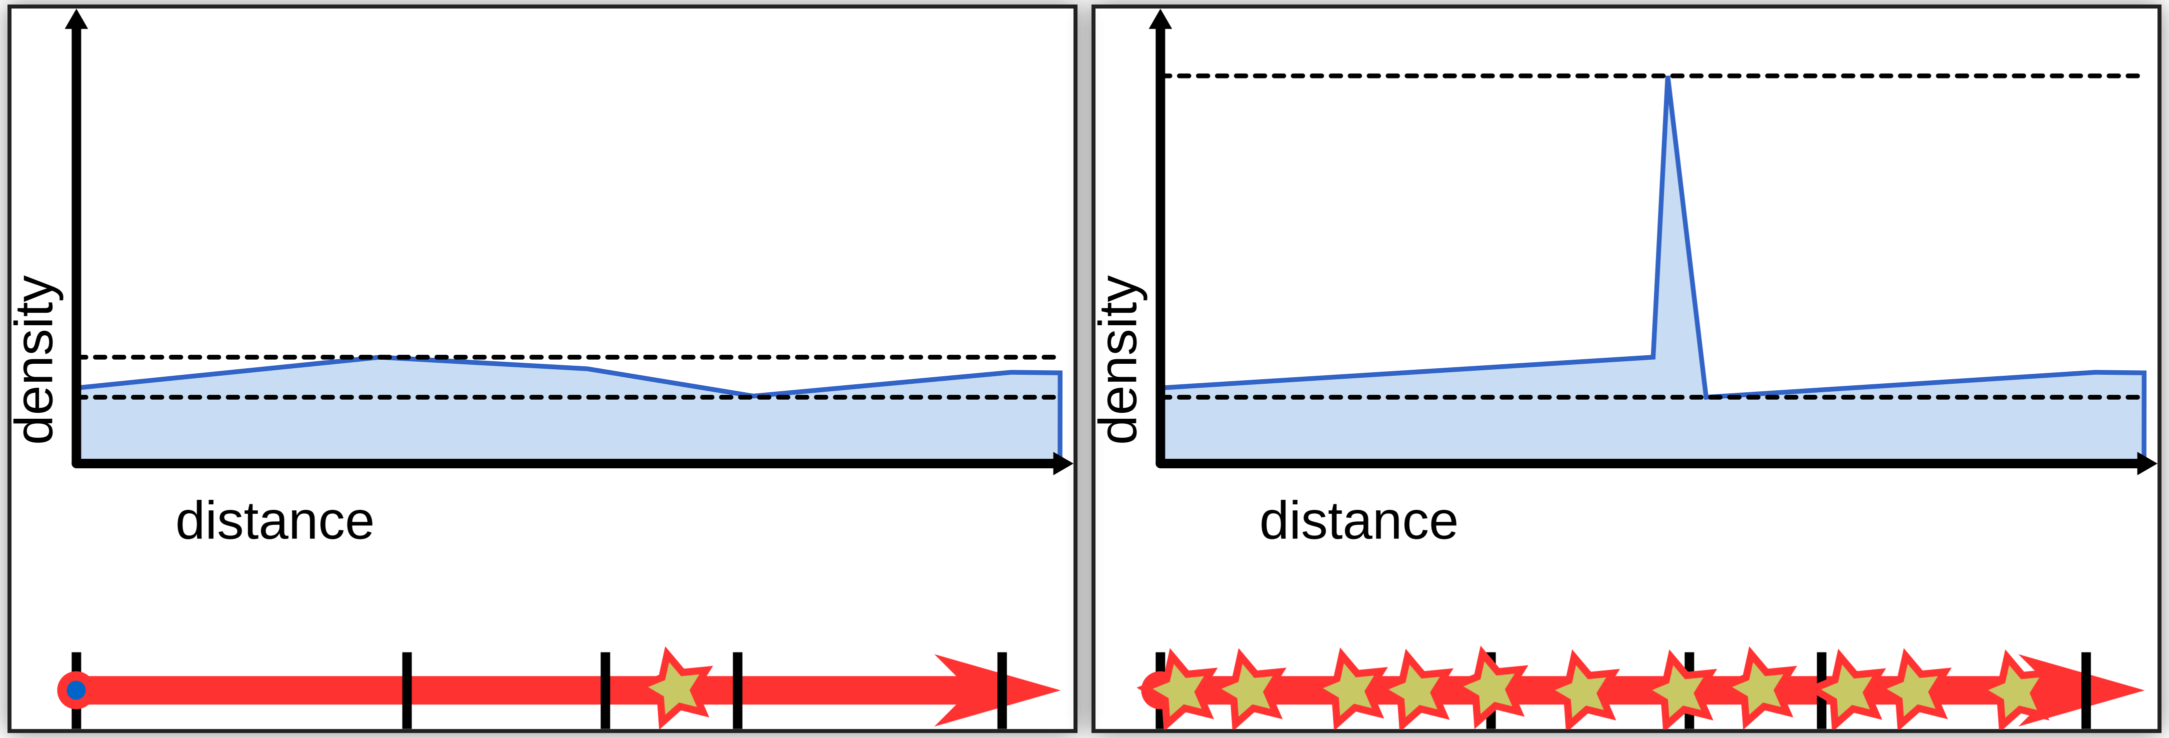

acceleration data structures

null collision-based tracking

- independent of tessellation
- high number of events in dense media, regardless of tessellation!

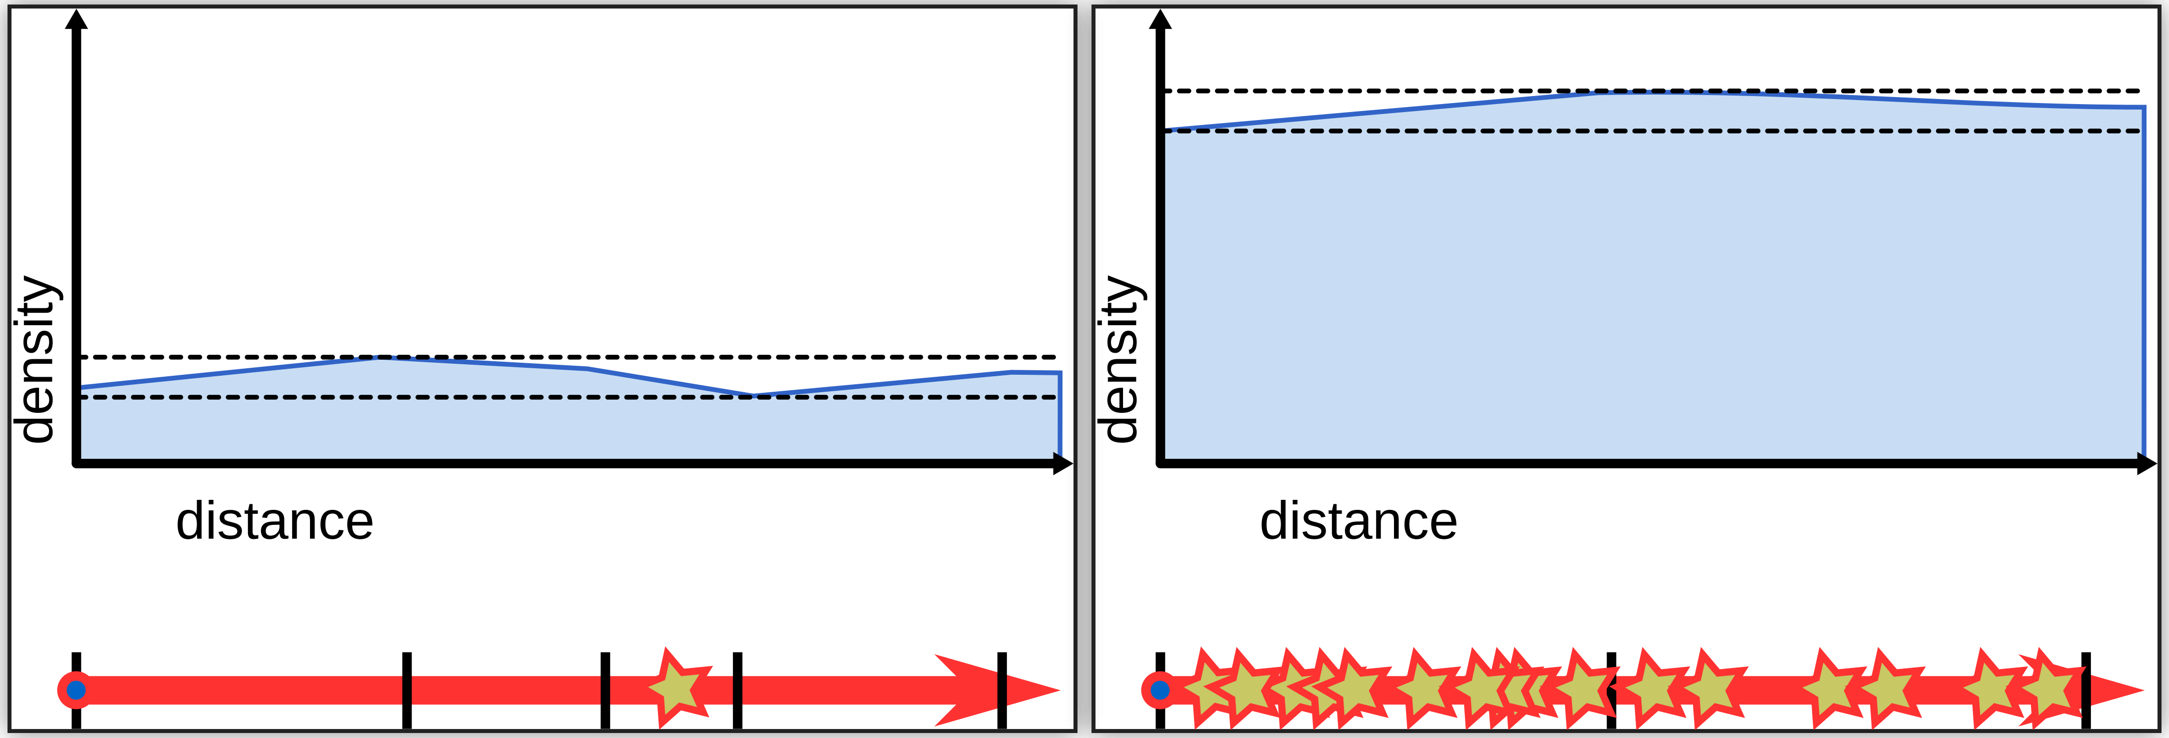

- accessing the memory within the same voxel is still expensive
- alleviated by decomposition tracking [KHLN17]

acceleration data structures

decomposition tracking [KHLN17]

- separate coefficients μ into sum of components:
  - coarse voxels for homogeneous parts
  - sparse fine details added on top

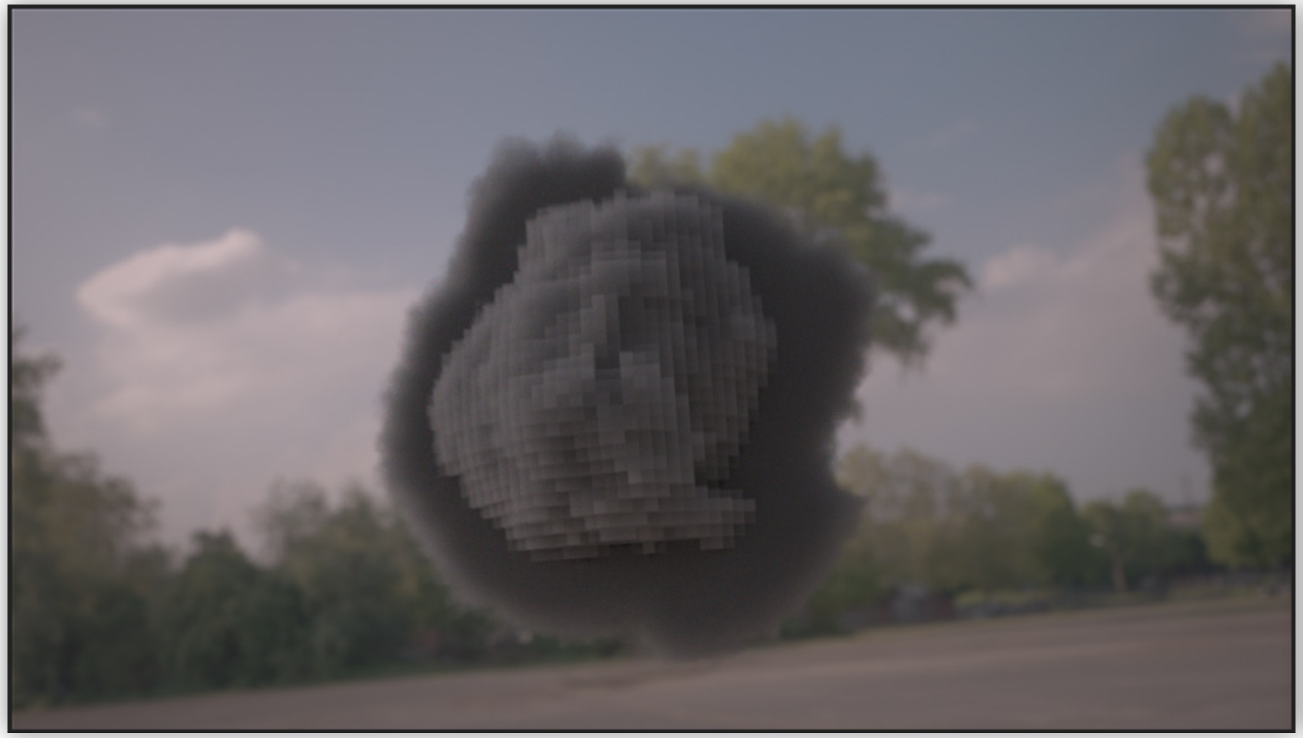


- track through coarse part first and use for early out
  - also profits regular tracking

end of my part

up next:

Jaroslav to present future work and outlook

summary

free flight distance sampling

- woodcock/delta tracking

transmittance estimation

- track-length
- residual ratio
- free flight versions

path sampling

- path space formulation
- summary of advanced methods

acceleration structures

- for regular tracking
- for null collisions (bottom-level)

open research problems

null collision algorithms and MIS

- missing link to integrate into powerful framework
  - for instance combine with equi-angular sampling
- can we estimate the PDF?
  - expectation and division do not commute!

$$X = \frac{f(\bar{\mathbf{x}})}{p(\bar{\mathbf{x}})}$$

open research problems

leverage recent advances in machine learning

- special purpose denoising
  - including a volume prior?
- path guiding for volumes?
  - joint importance sampling for multiple vertices?

open research problems

joint handling of surfaces and geometry

- still often surface transport is handled separately
  - makes inclusion of all interreflections hard
  - custom-cut algorithms increase maintenance cost
- represent surfaces as volumes, too?
  - and then ideally jointly downscale for LOD!

open research problems

generalisation to correlated scatterers

- core assumption of exponential path length: uncorrelated particles!
  - particle repulsion such as in cell growth is very correlated
  - really, no collision can be found inside the current particle (min distance)
  - some existing work

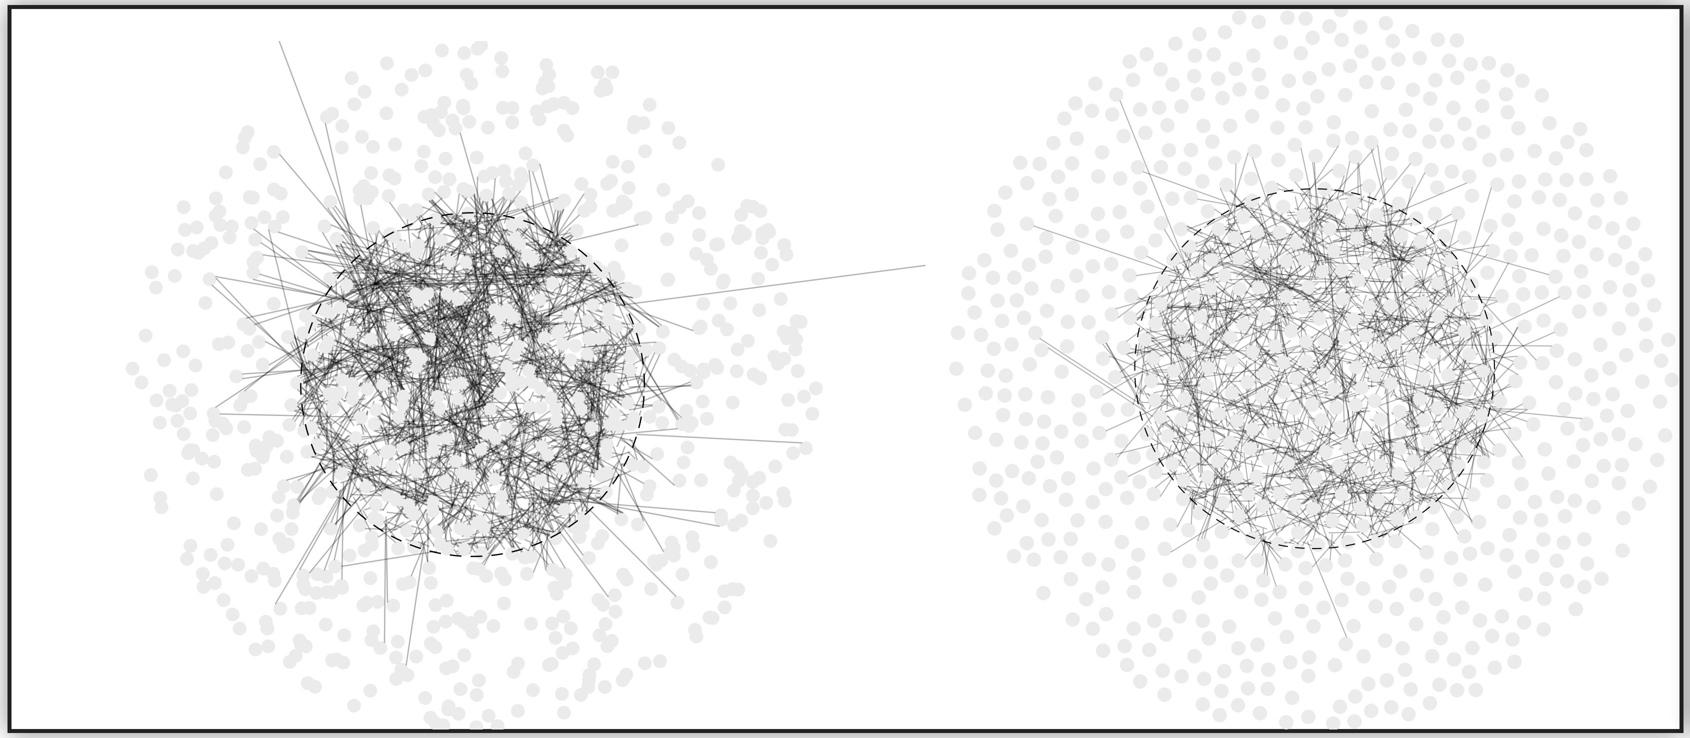

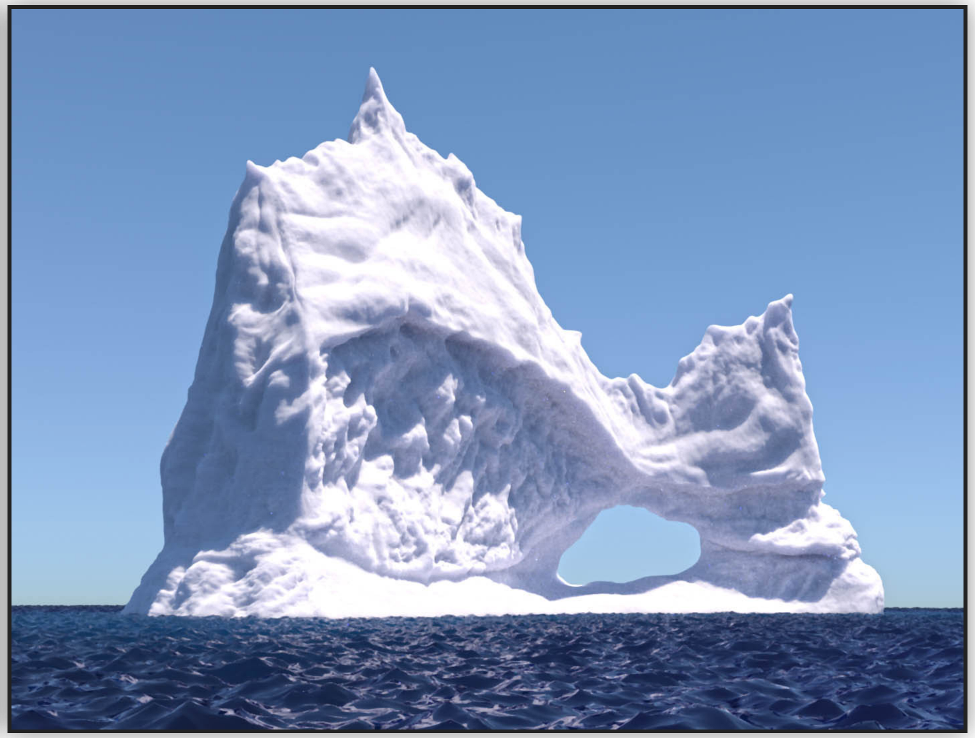

[d'Eon 2018, Jarabo et al. 2018, Bitterli et al. 2018]

thank you!

any questions?

acknowledgements:

- Peter Kutz
  - for tracing down many of early delta tracking papers
- reviewers
  - for feedback on the paper draft
- Maurizio Nitti
  - for help w/ illustrations