CS 563 Advanced Topics in Computer Graphics
Russian Roulette - Sampling Reflectance Functions

by Alex White

Monte Carlo Ray Tracing

- Monte Carlo
  - In ray tracing, use randomness to evaluate higher dimensional integrals
  - But while correct on the average, a variance is induced
- Variance leads to a noisy image
- Convergence rate: quadruple samples to reduce variance by half
- Goal: Improve efficiency without increasing samples

What is efficiency?

- For an estimator *F* :

$$\in [F] = \frac{1}{V[F] \times T[F]}$$

- *V[F]* = its variance
- *T[F]* = running time to compute the value

Russian Roulette

- Improve efficiency by increasing likelihood that sample will have a significant contribution
- Spend less time on small contributions

Russian Roulette

- Direct lighting integral

$$L_o(p, w_o) = \int_{S^2} f_r(p, w_o, w_i) L_d(p, w_i) \left|\cos q_i\right| dw_i$$

- Estimator for N samples

$$\frac{1}{N}\sum_{i=1}^{N}\frac{f_r\left(p, w_o, w_i\right) L_d\left(p, w_i\right)\left|\cos q_i\right| dw_i}{p\left(w_i\right)}$$

- Most of the work comes from tracing a shadow ray
- How to optimize?

Russian Roulette

- If $f_r(p, w_o, w_i)$ is zero for the direction w_i we can skip the tracing work
- Why stop there?
- What about rays where this value is very small?
- Or when w_i is close to the horizon?

- We can’t ignore or we would underestimate the end result
- Answer: weighting!

Russian Roulette

- However, this never reduces variance
- Imagine a bunch of black pixels with a few very bright ones
- A technique known as “efficiency-optimized Russian roulette” attempts to rectify this
  - Keep track of average variance and average ray count
  - Use to compute threshold for each new sample

Battle of the Roulettes - Thiago Ize, University of Utah

Full scene – efficiency-optimized Russian roulette (10.5m)

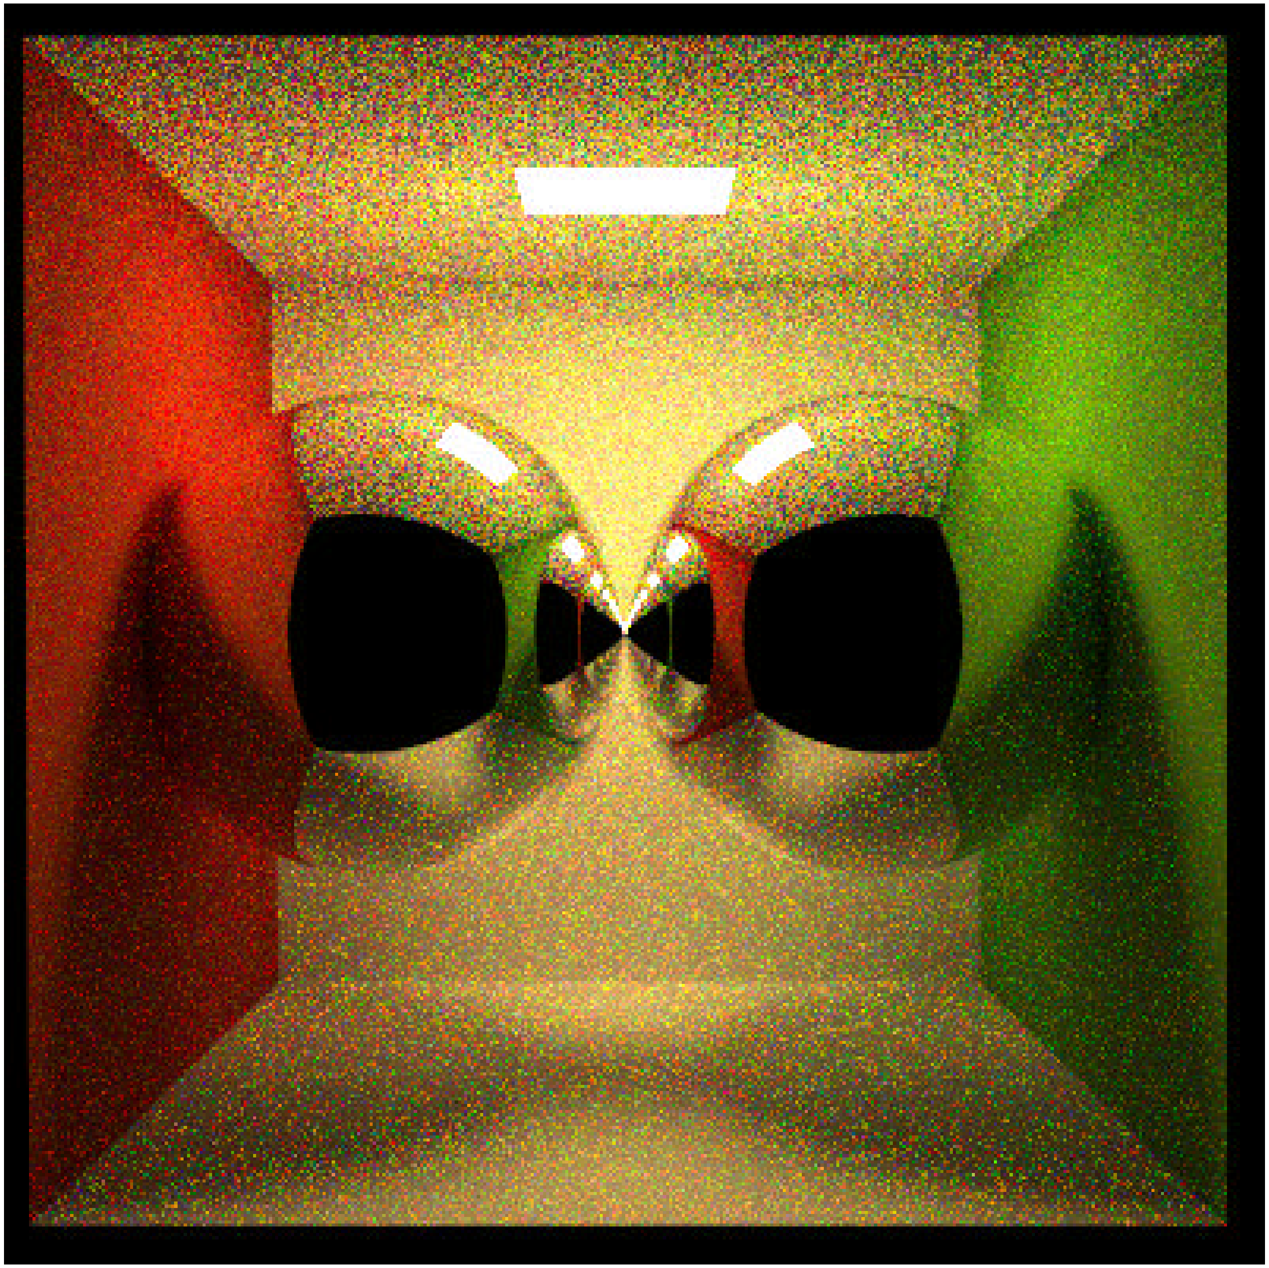

Battle of the Roulettes

Fixed Max Depth of 100 (53.2m)

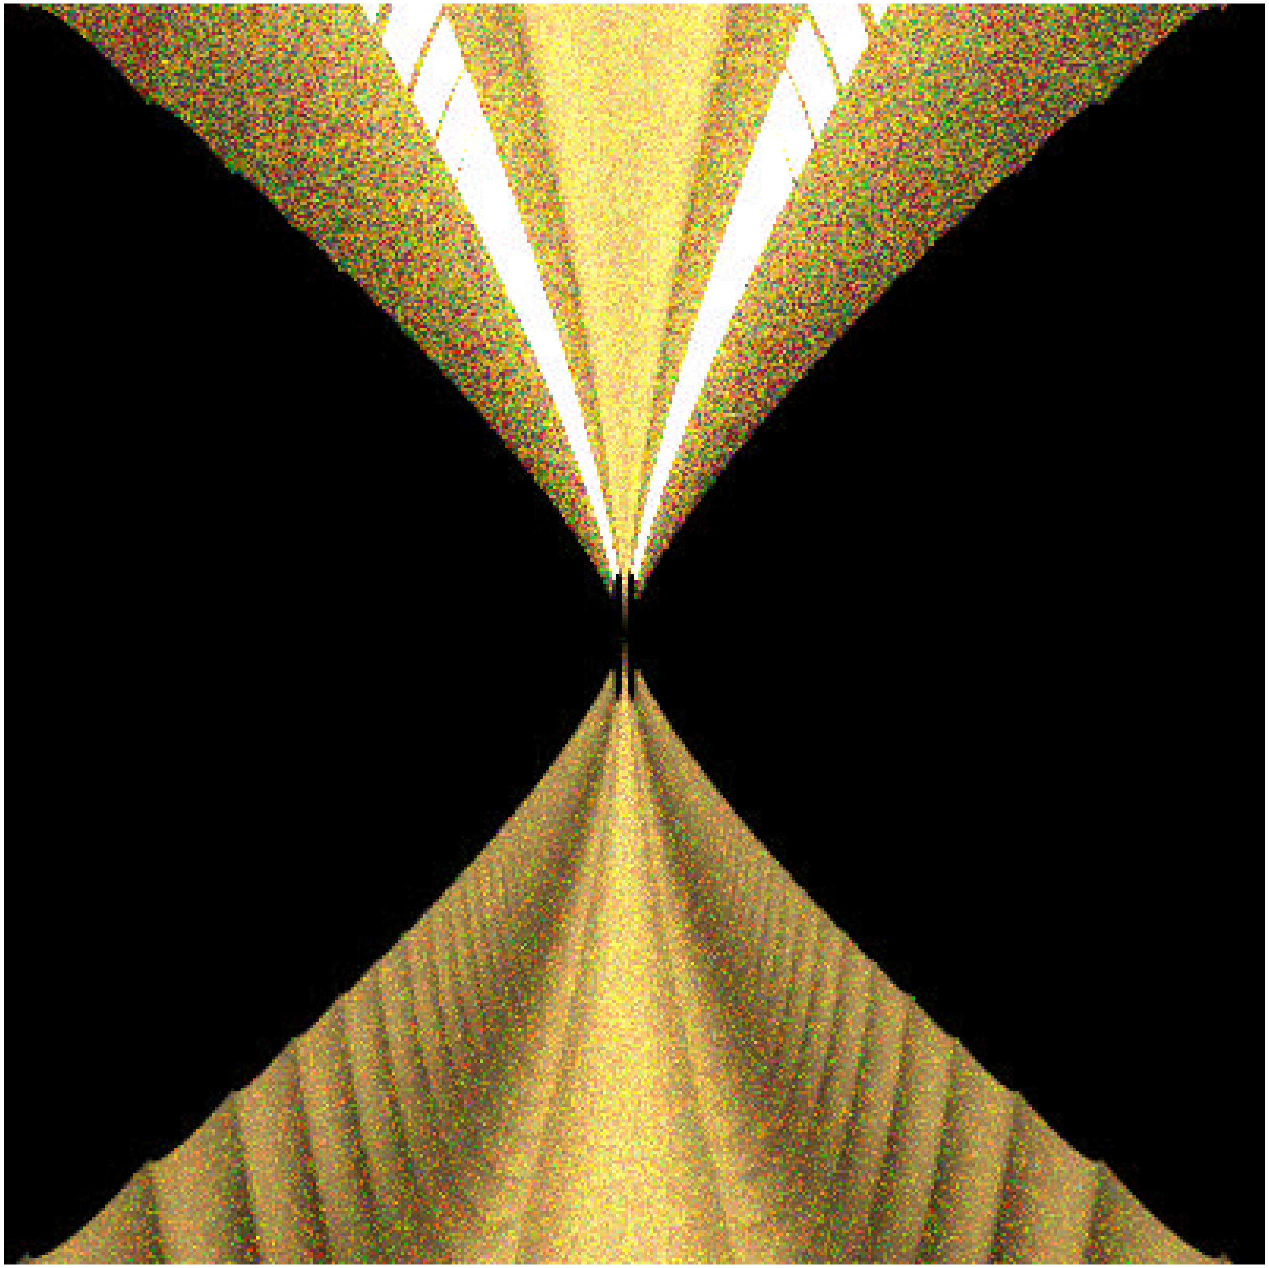

Battle of the Roulettes

Fixed Max Depth of 10 (10.6m)

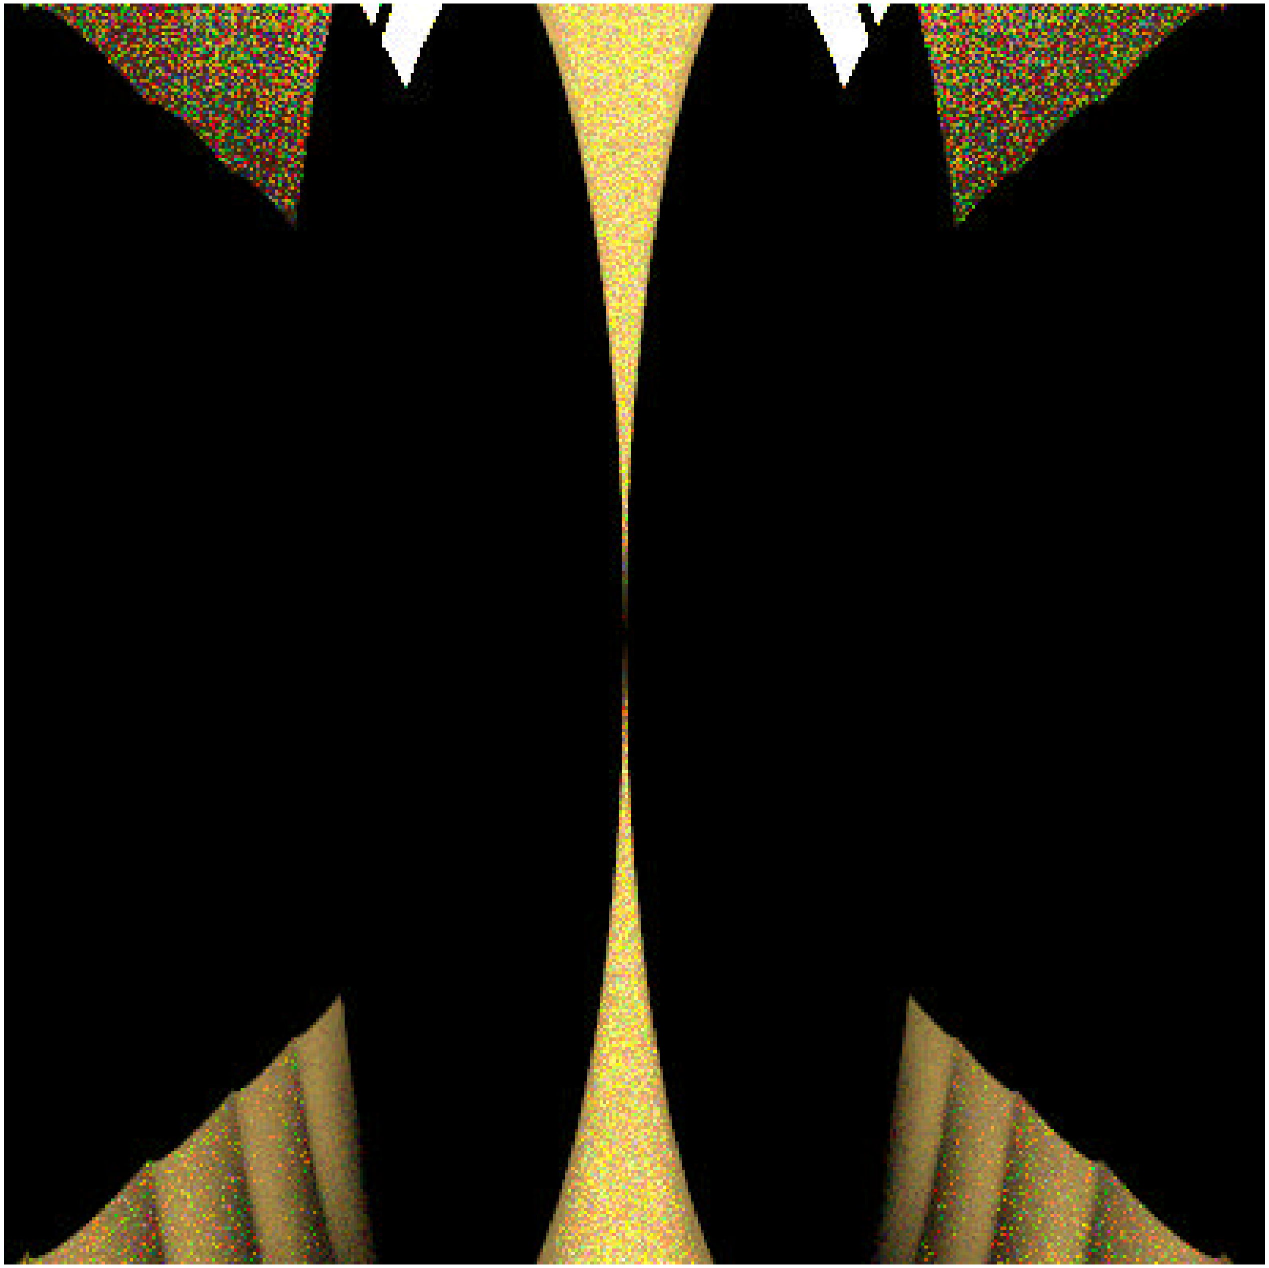

Battle of the Roulettes

Russian roulette (30% termination)

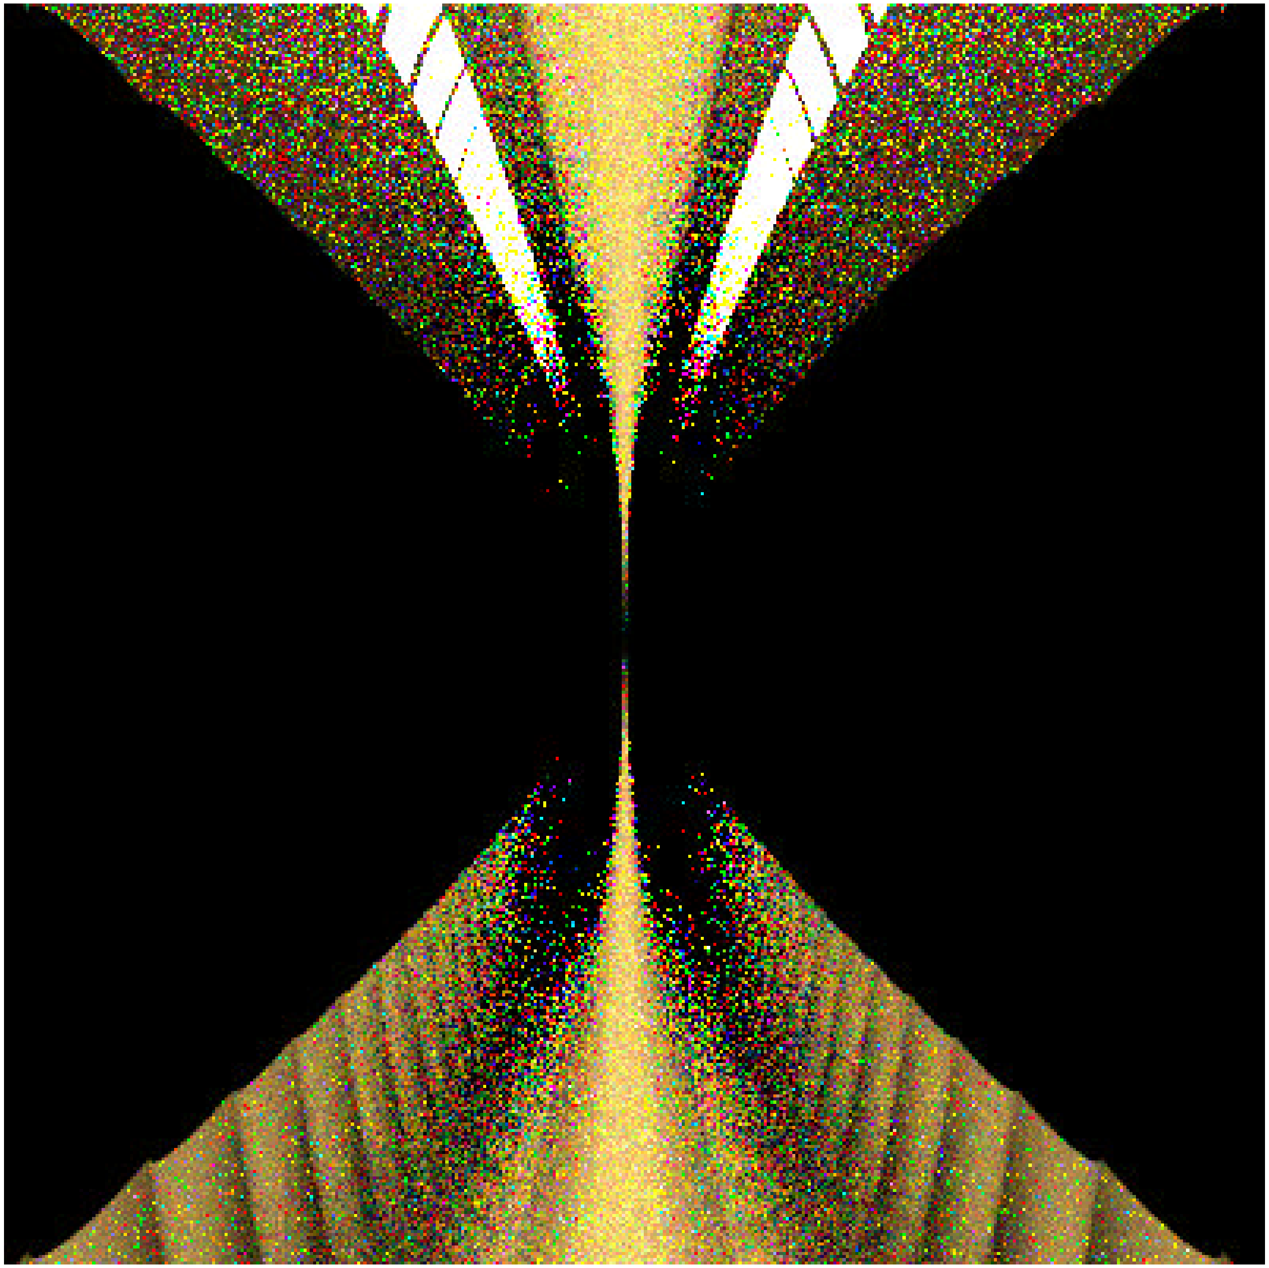

Battle of the Roulettes

Russian roulette (*q* proportional to “pathThroughput/0.5 (27.8m))

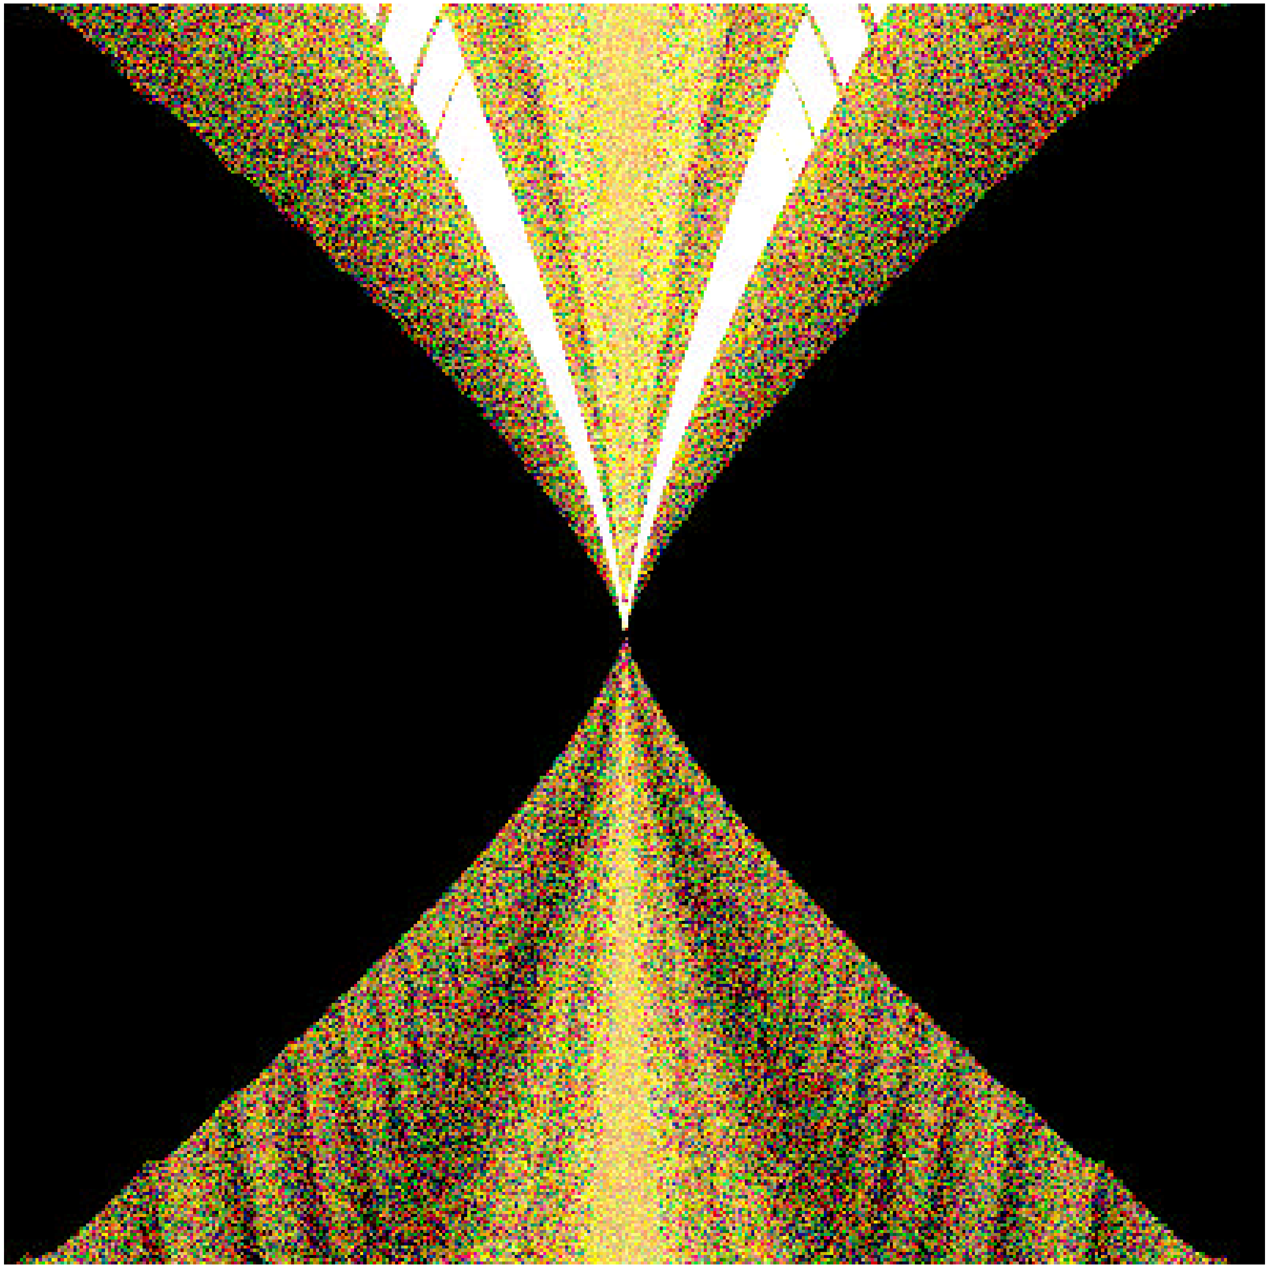

Battle of the Roulettes

Russian roulette (*q* proportional to “pathThroughput/0.1” (31.3m))

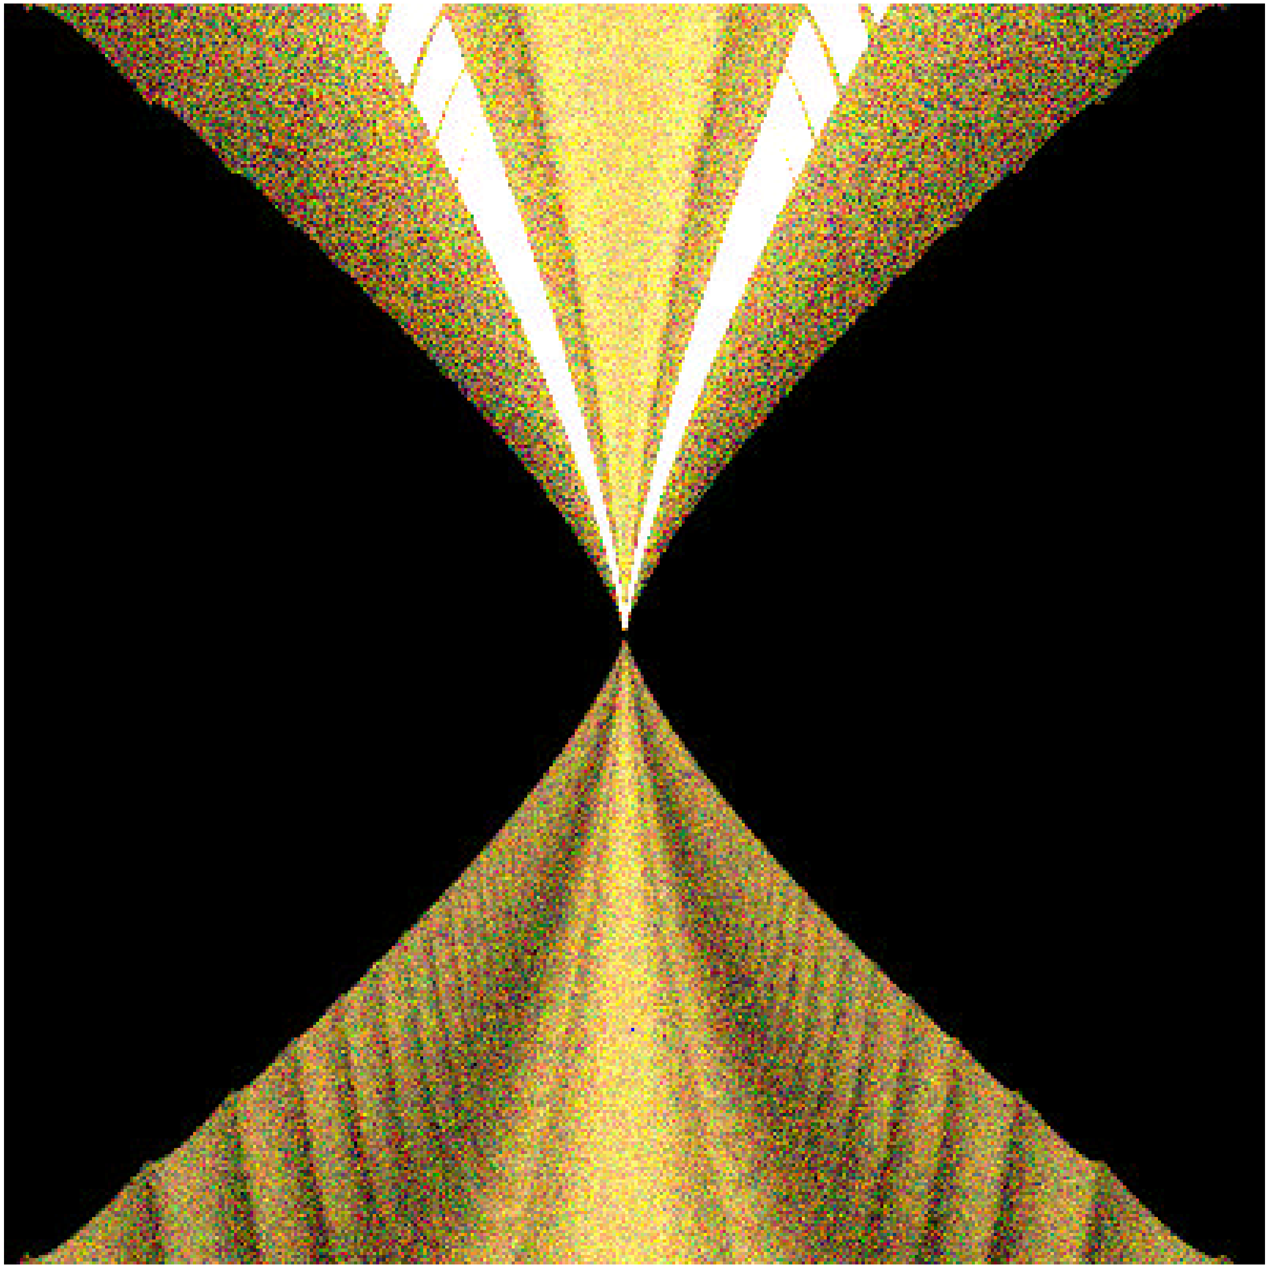

Battle of the Roulettes

Russian roulette (*q* proportional to "pathThroughput/0.01" (32.6m))

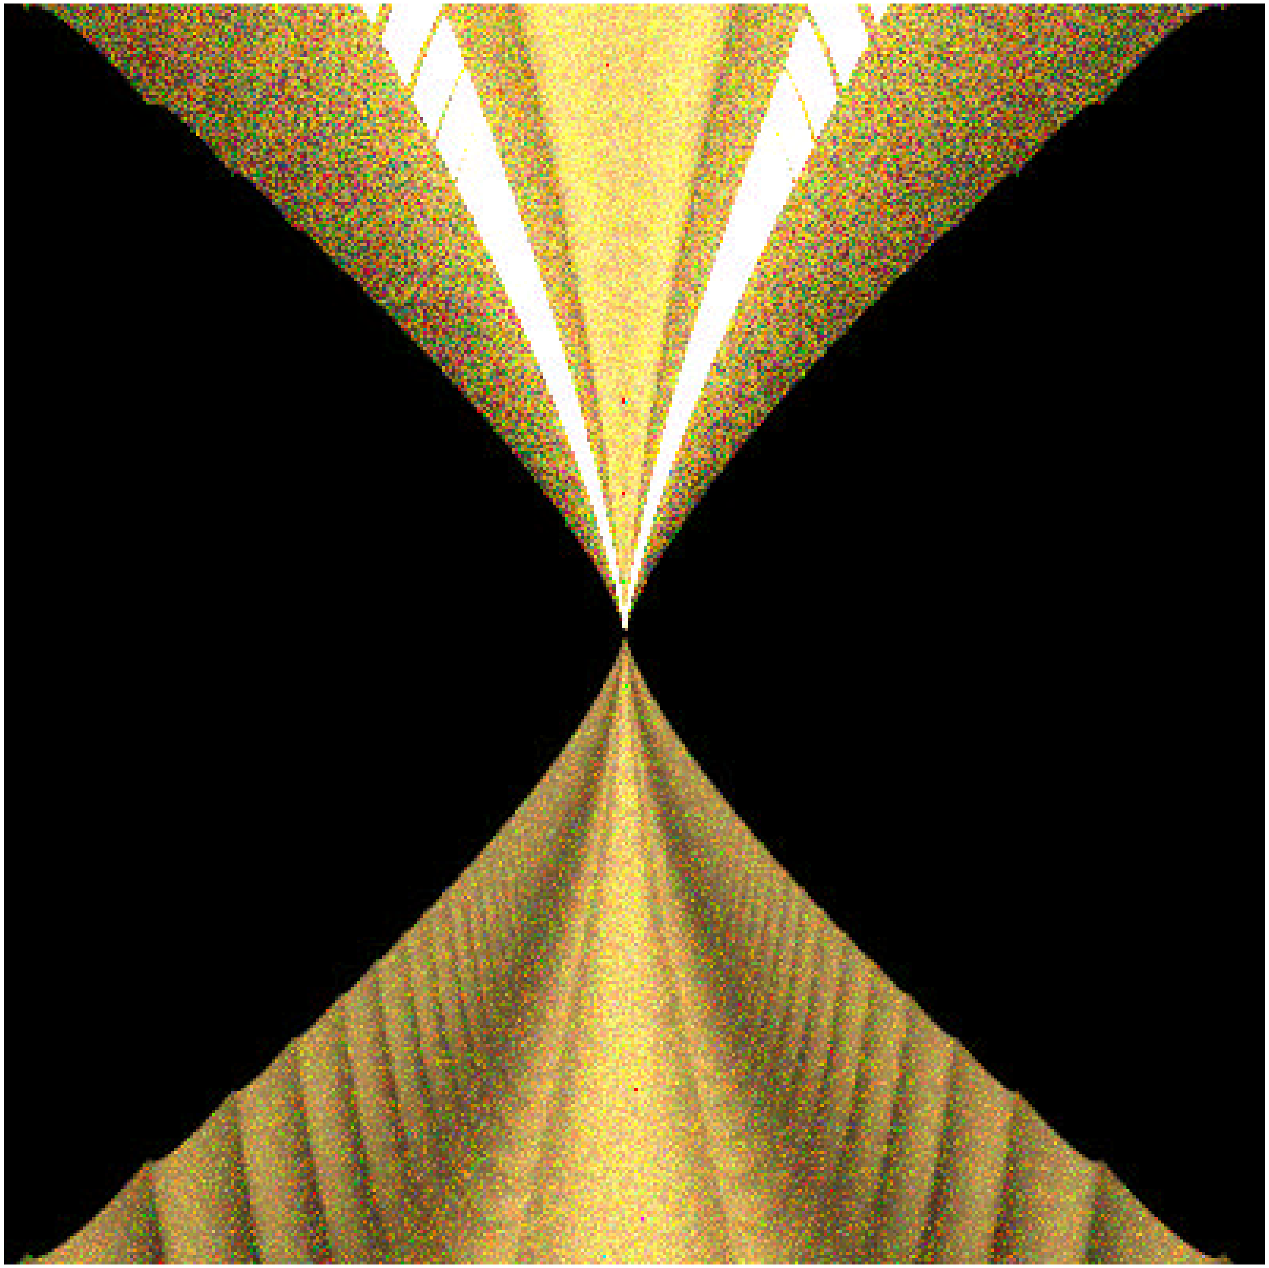

Battle of the Roulettes

Efficiency-optimized Russian roulette (28.8m)

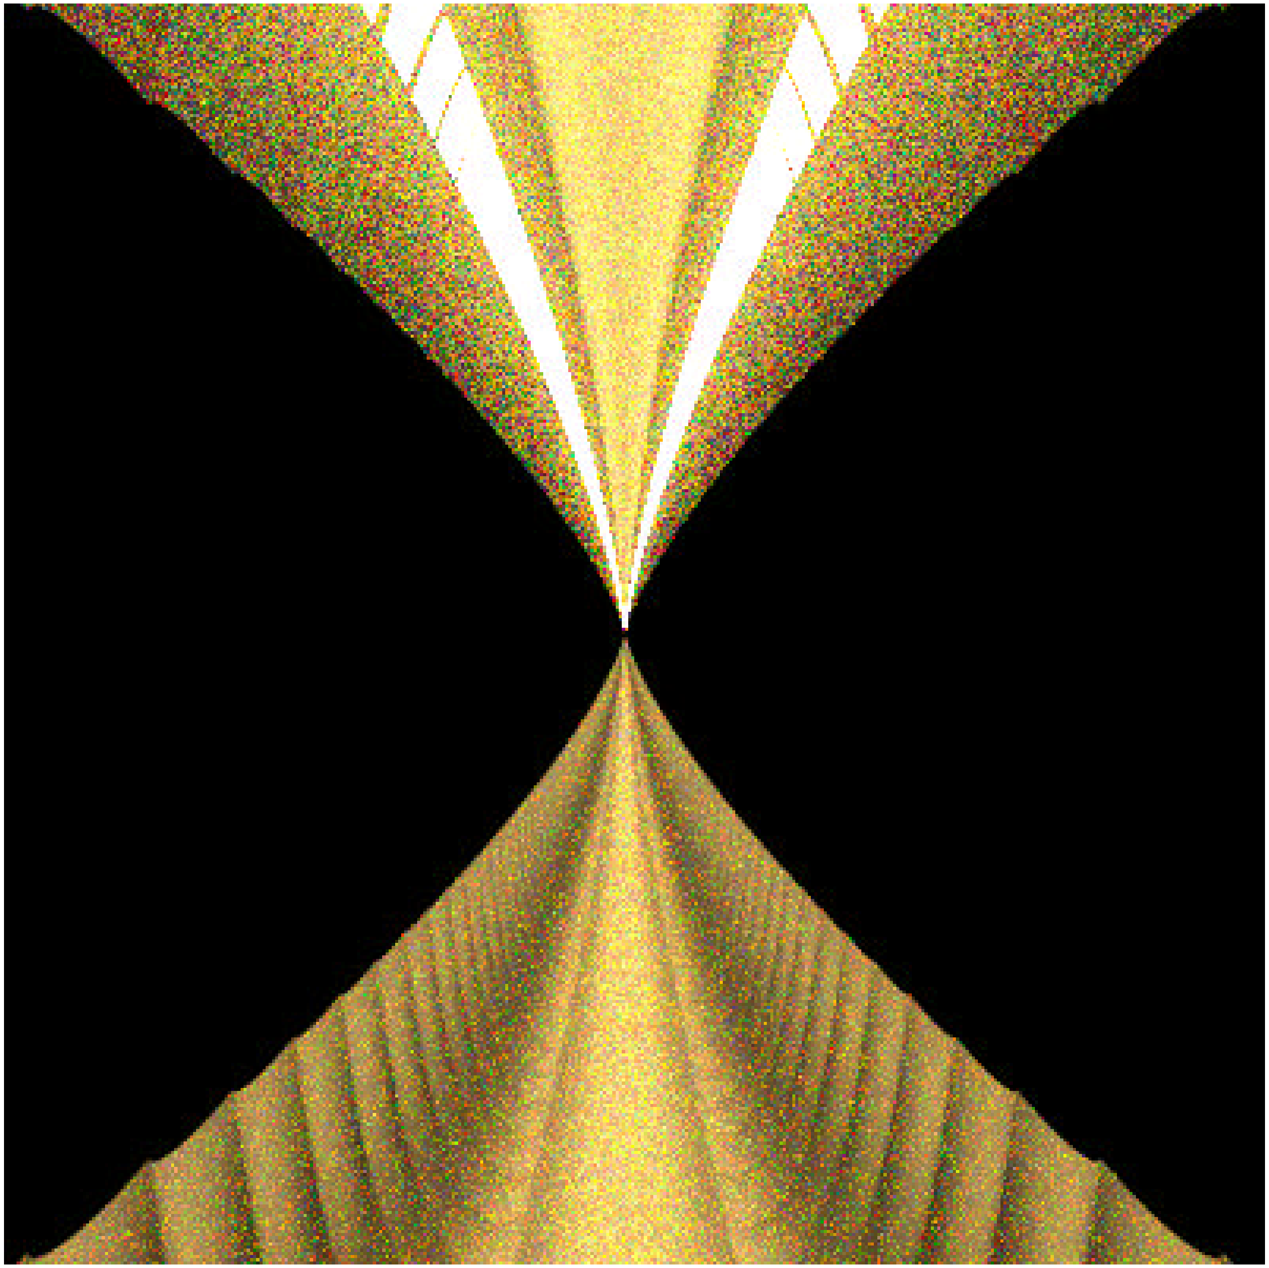

Battle of the Roulettes

Russian roulette (*q* proportional to “pathThroughput/0.5 (27.8m))

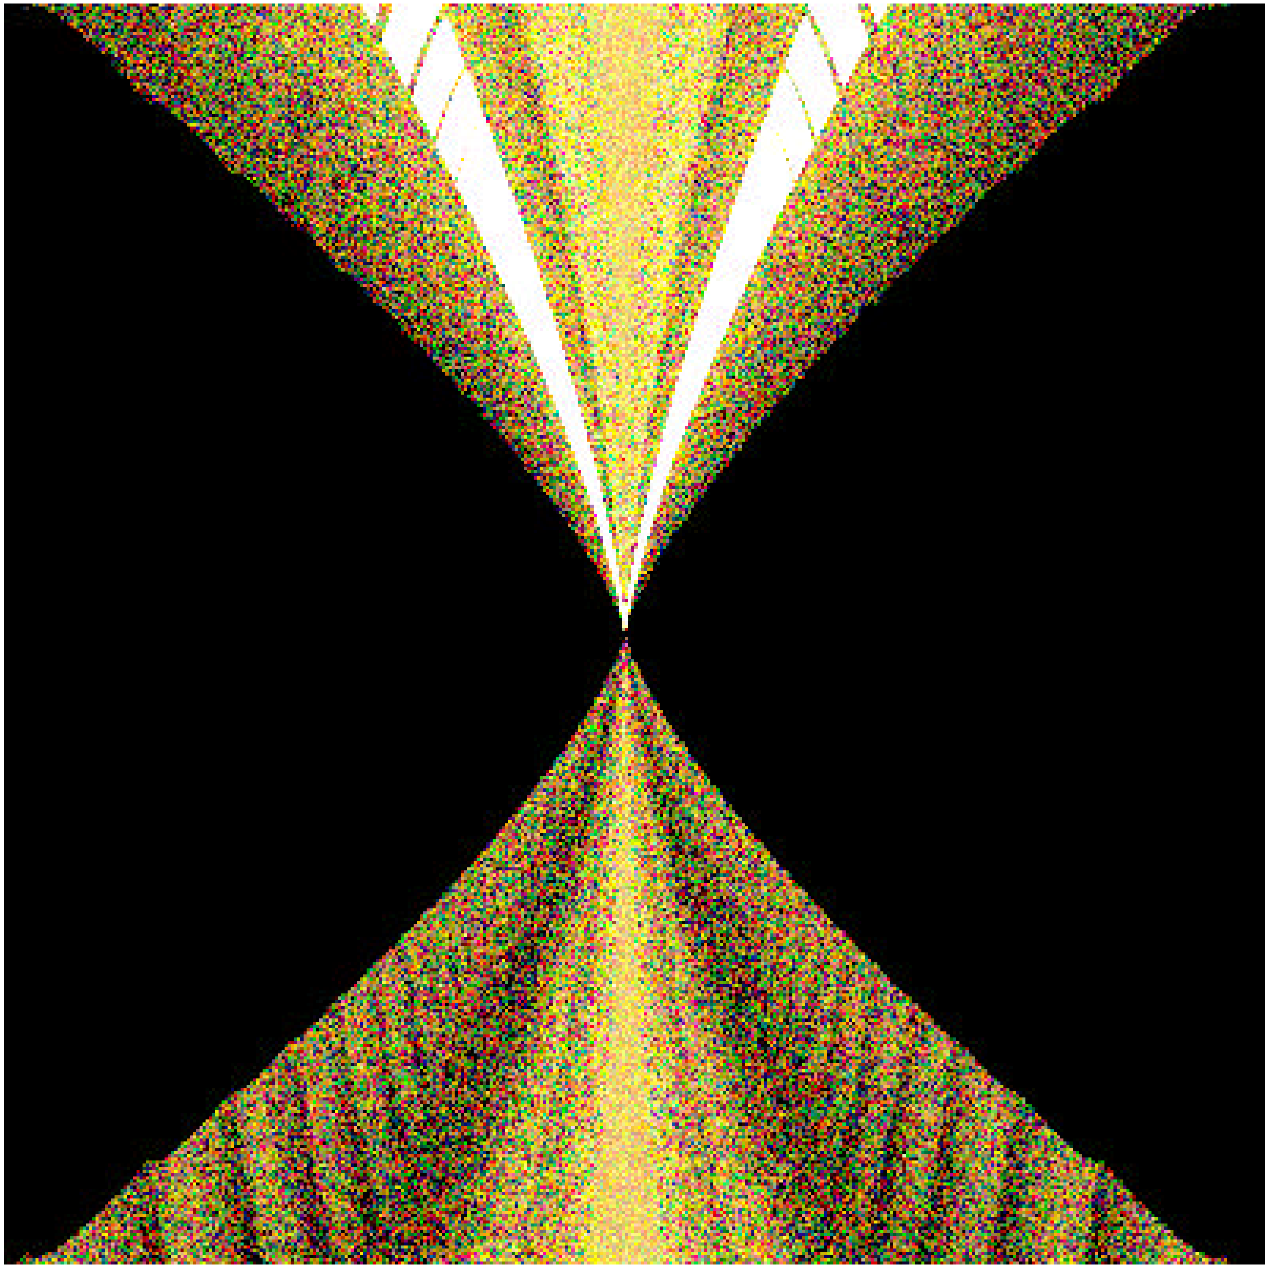

Splitting

- Russian roulette – reduces effort spent evaluating unimportant samples

- Splitting – increases important samples to improve efficiency
  - Each sample = 1 camera ray + 1 shadow ray
  - Important means shadow rays in many cases

Careful Sample Placement - Stratified Sampling Revisited

- Divide a domain into non-overlapping regions

- Stratified sampling can never increase variance
- Large strata contain more variation and will have individual means closer to the real mean

- Why not keep making strata smaller?
  - "Curse of dimensionality"
  - Possible to stratify some dimensions independently
  - Latin Hypercube sampling

Figure 15.02a

Random Sampling

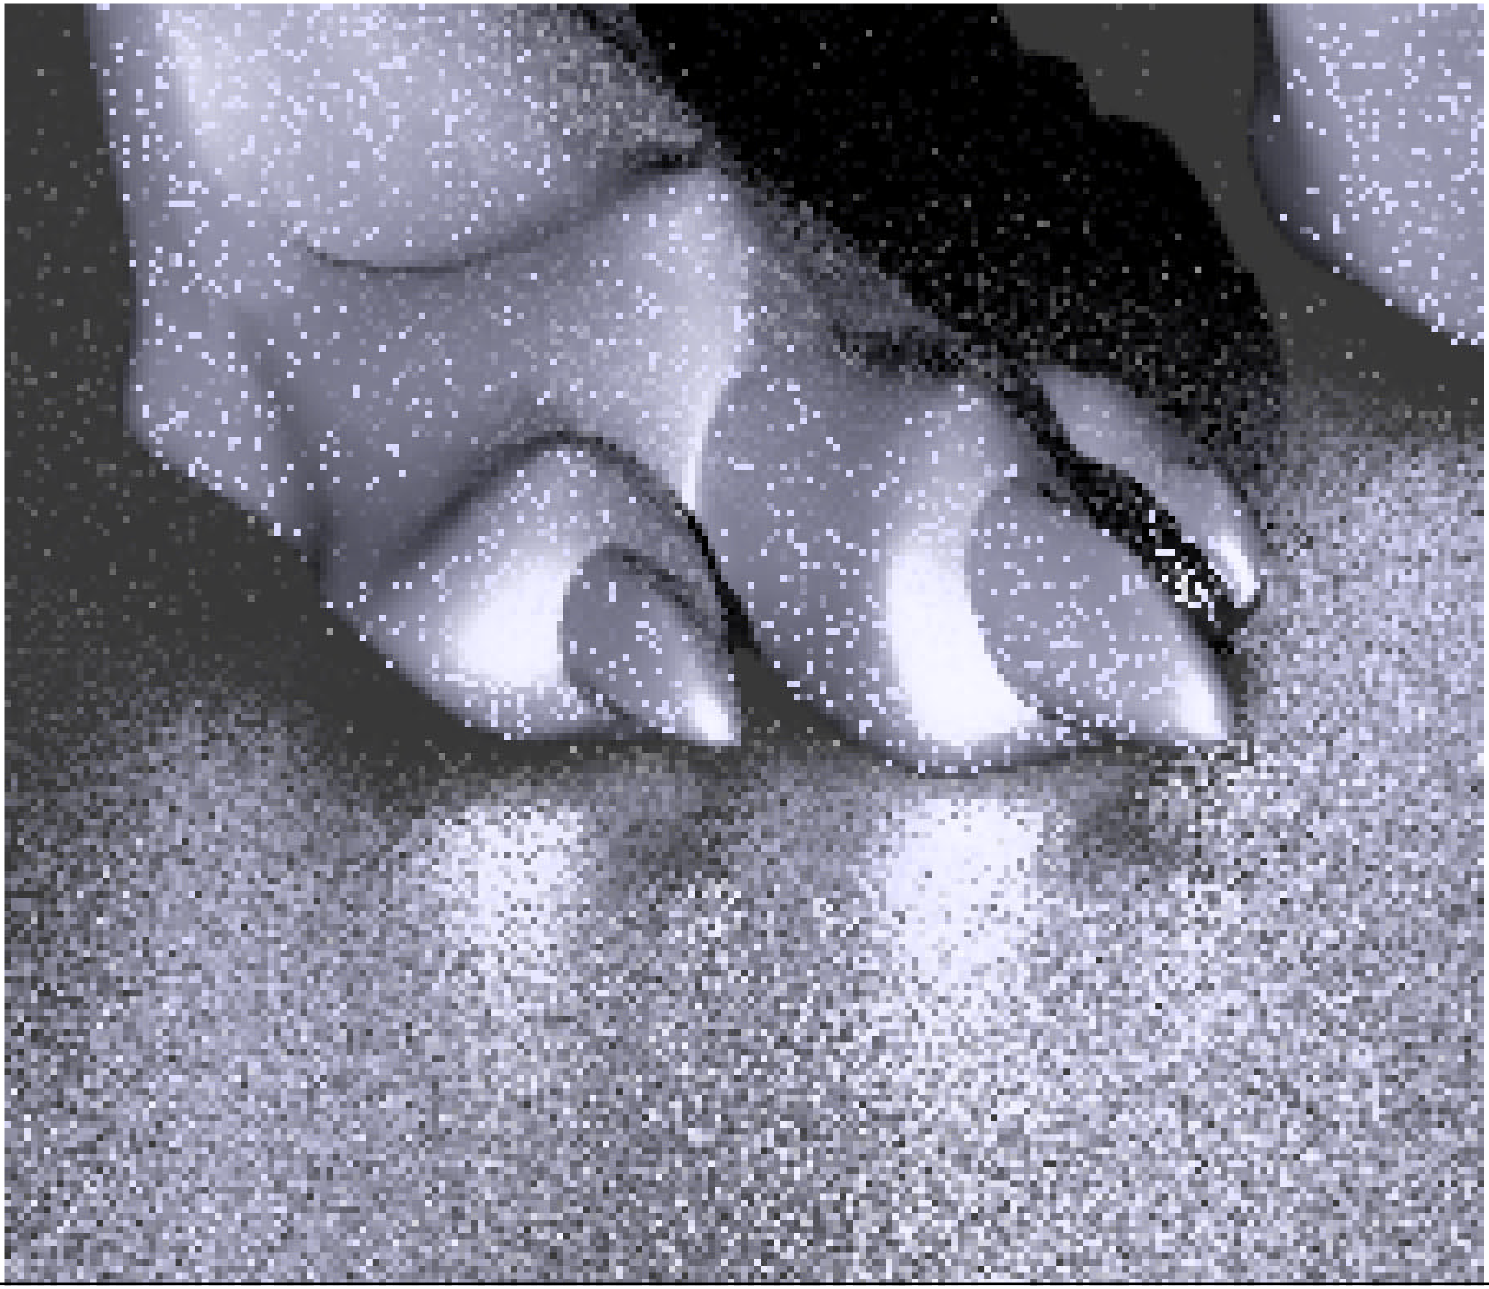

Figure 15.02b

Stratified Sampling

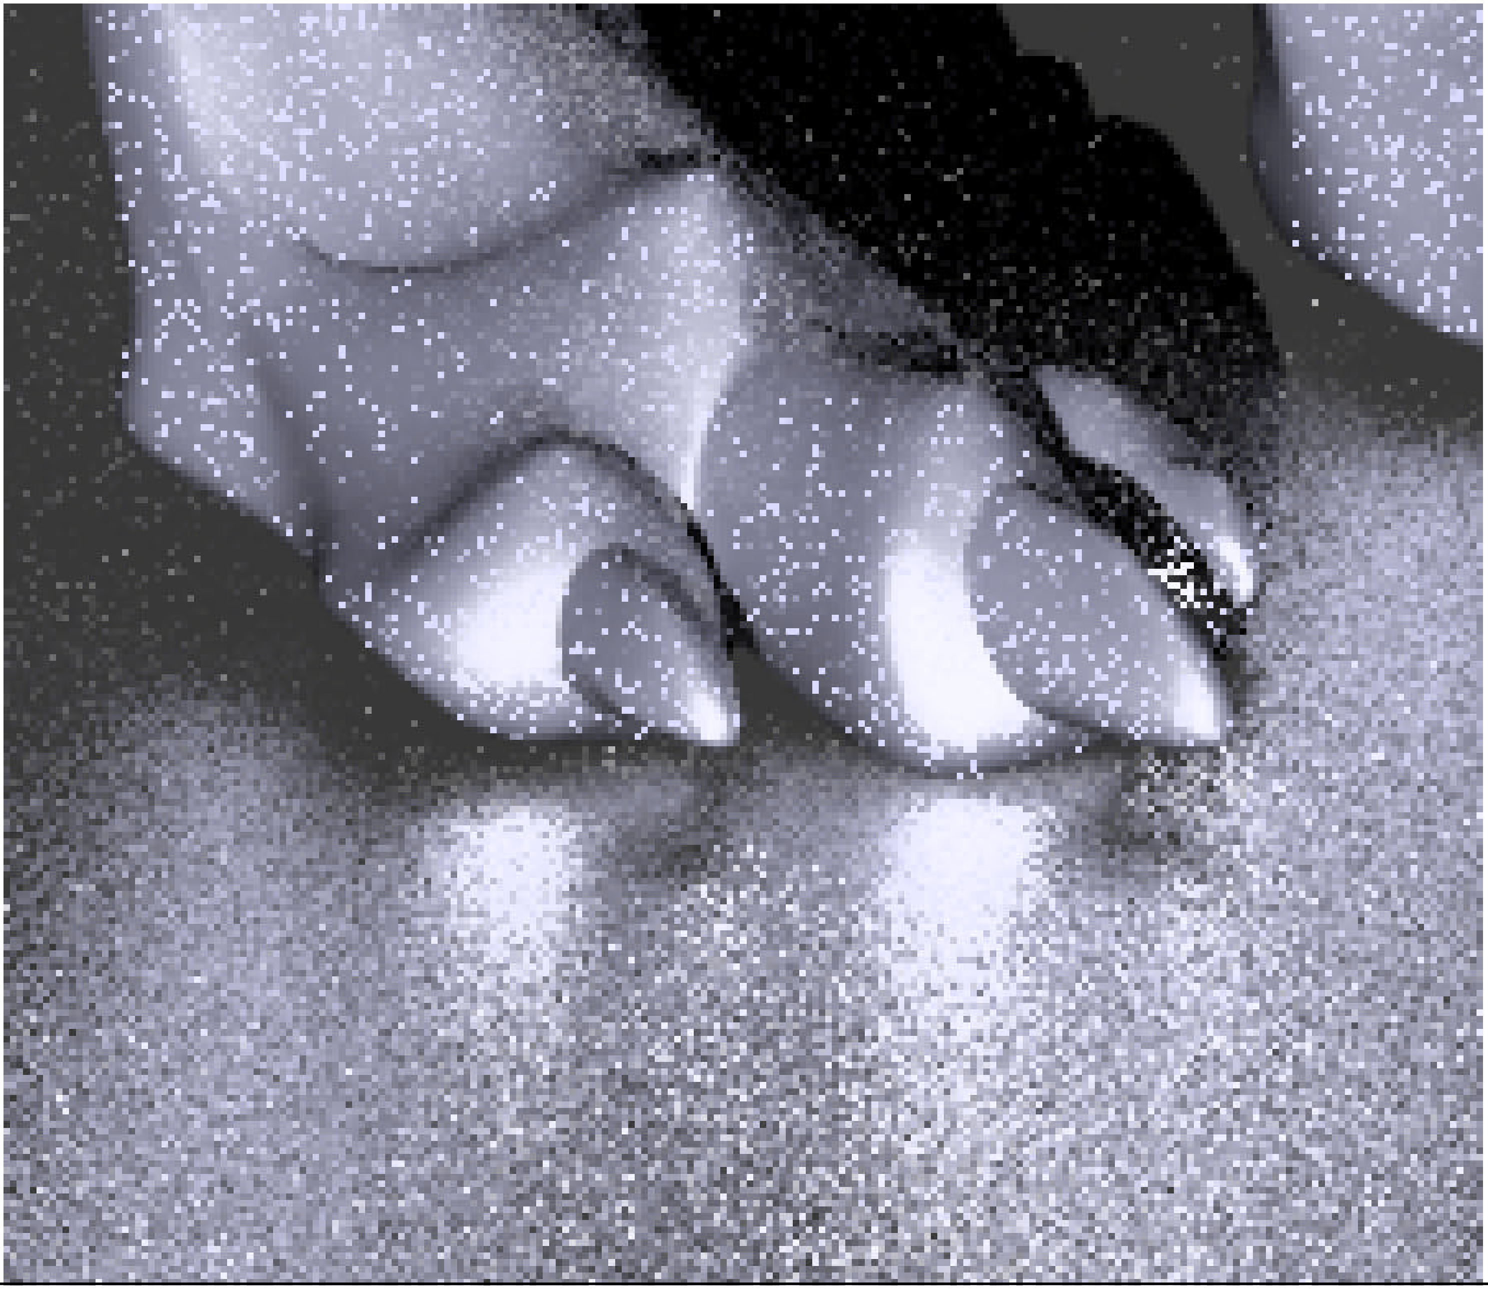

Quasi Monte Carlo

- Replace random numbers with low-discrepancy point sets generated by carefully designed deterministic algorithms
- Advantage: faster rates of convergence

- Works better with a smooth integrand not characteristic of graphics
  - performs slightly better in practice
- Works better with smaller dimensions
- A hybrid technique “randomized quasi-Monte Carlo” extends the benefits to larger dimensions

Warping Samples

- Sample points lie within $[0,1]^2$
- PBRT uses algorithms to transform to map to light sources
- Mapping must preserve benefits of stratification

- (0,2)-sequence still well distributed
- Random stratified is not

Bias

- Sometimes, picking an estimator whose expected value does NOT equal the actual, can lead to lower variance
- Bias (b):

$$b = E[F] - F$$

- How can this be good?

Importance Sampling

- Importance sampling
  - Choose a sampling distribution that is similar in shape to the integrand
  - Samples tend to be taken from “important” parts where the value is higher

  - One of the most frequently use variation reduction techniques
  - It’s not difficult to fine an appropriate distribution in graphics
    - Often integrand a product of more than one function
    - Finding one similar to one multiplicand is easier
    - However a bad choice can be worse than uniform

Multiple Importance Sampling

- But often we can find one similar to multiple terms
- Which one do we use?

References

- *Physically Based Rendering*. Pharr and Humphreys.
- “Efficiency-Optimized Russian Roulette.” Thiago Ize, University of Utah. http://www.cs.utah.edu/~thiago/cs7650/hw12/index.html.
- “Quasi-Monte Carlo method.” *Wikipedia*. http://en.wikipedia.org/wiki/Quasi-Monte_Carlo_method