

Spatiotemporal Blue Noise Masks

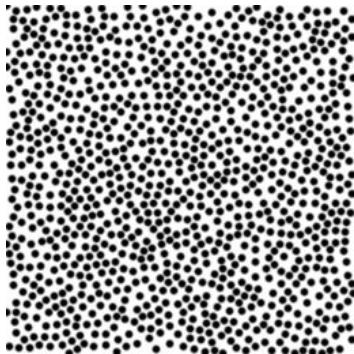
Alan Wolfe
Nathan Morrical
Tomas Akenine-Möller
Ravi Ramamoorthi

NVIDIA, EA SEED
NVIDIA, University of Utah
NVIDIA
NVIDIA, University of California, San Diego

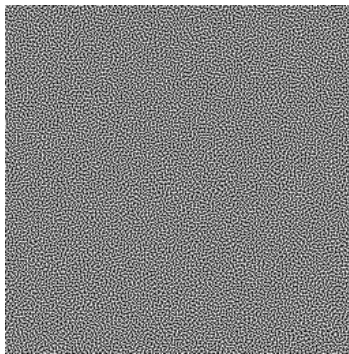
 @Atrix256
alan.wolfe@gmail.com
<https://blog.demofox.org/>

Blue Noise Samples vs. Masks

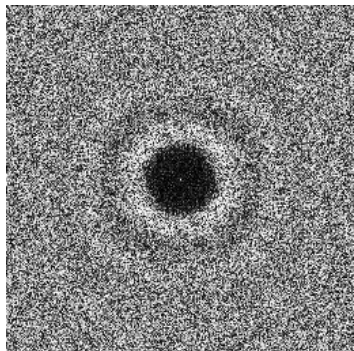
Samples



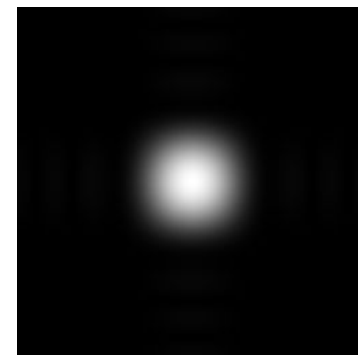
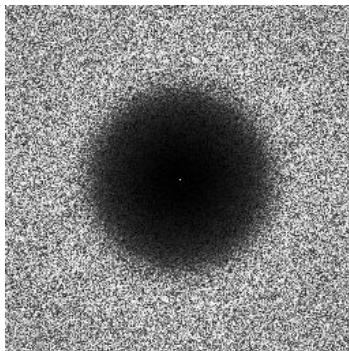
Mask



Samples
DFT



Mask
DFT



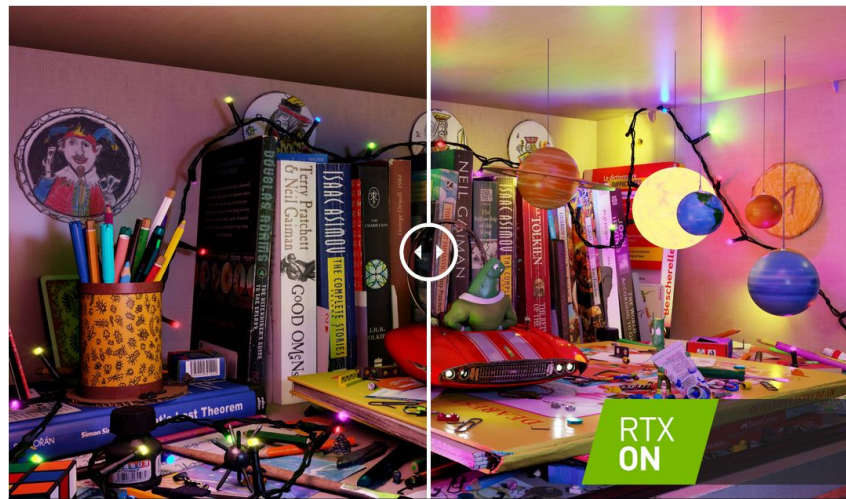
Gaussian Kernel DFT

Real Time: Always Wanting To Do More With Less



Real Time Global Illumination with RTXGI
<https://developer.nvidia.com/rtx/ray-tracing/rtxgi>

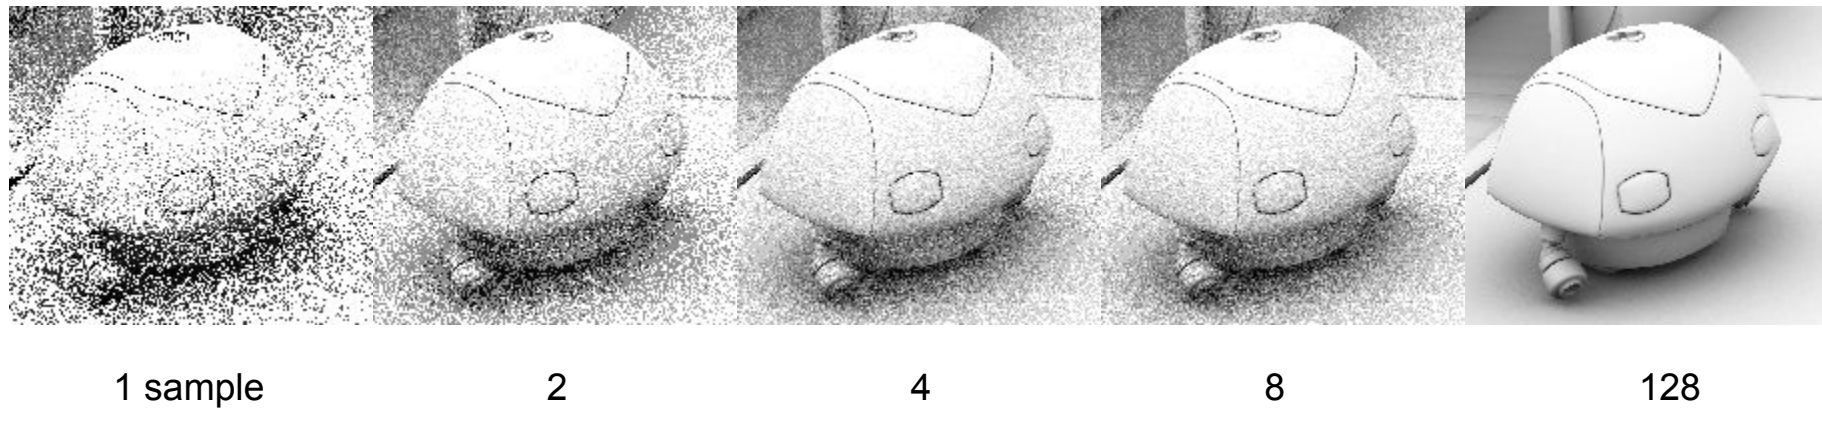
Fast probe grid GI + ray tracing probe updates



Real Time Direct Illumination with RTXDI
<https://developer.nvidia.com/rtx/ray-tracing/rtxdi>

Solves “many lights” direct illumination by making each pixel learn where to shoot rays for best results.

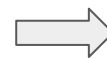
Stochastic Algorithms Allow Tuning Quality vs Speed



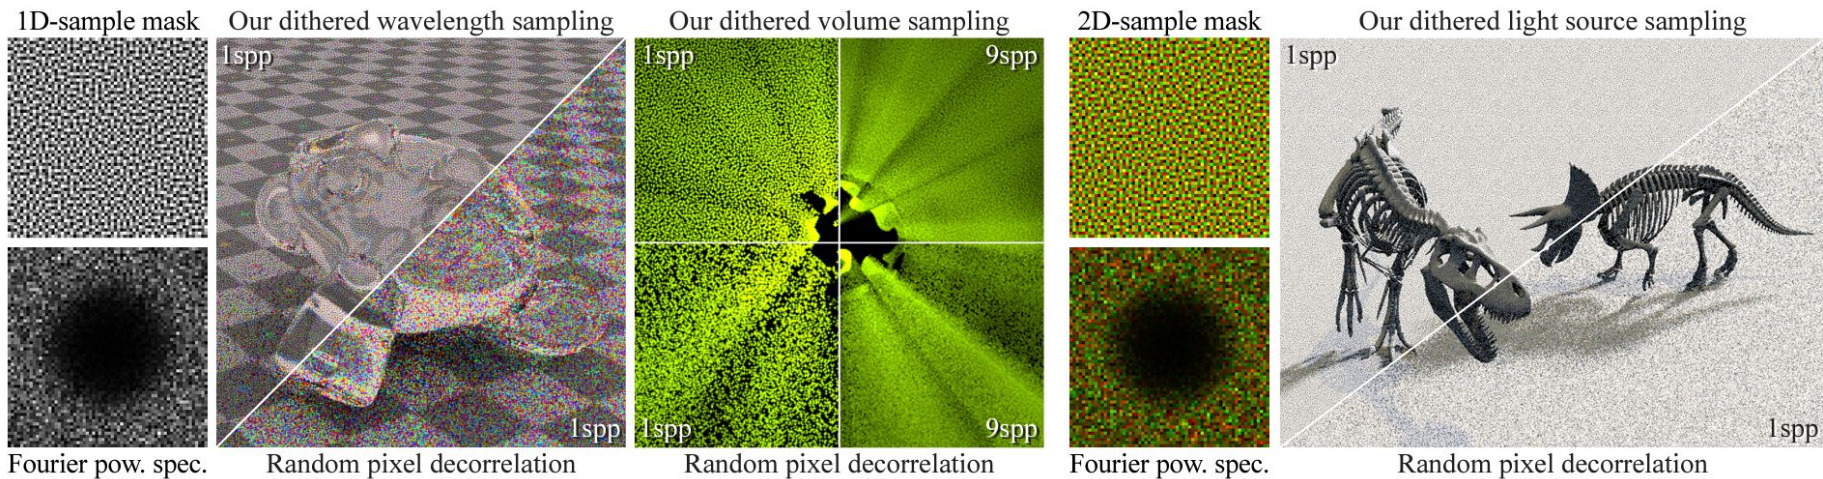
More Speed



Higher Quality



Anti Correlated Blue Noise >> Independent White Noise



Anti Correlated Blue Noise >> Independent White Noise



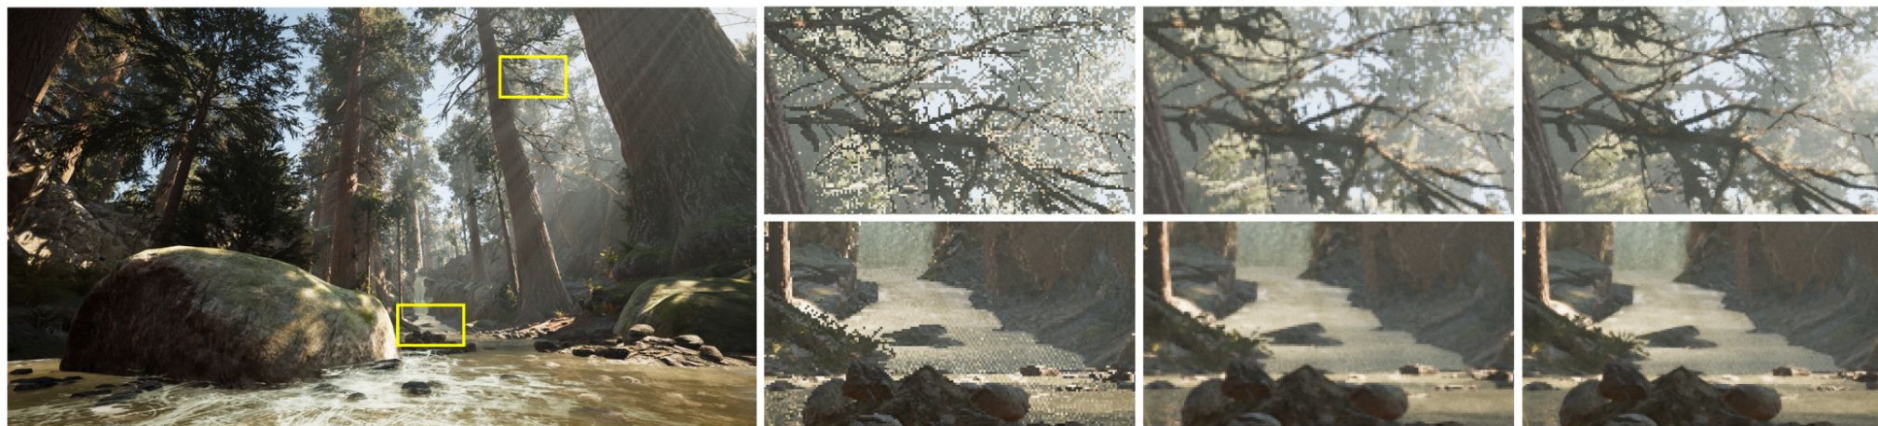
White Noise Dithered



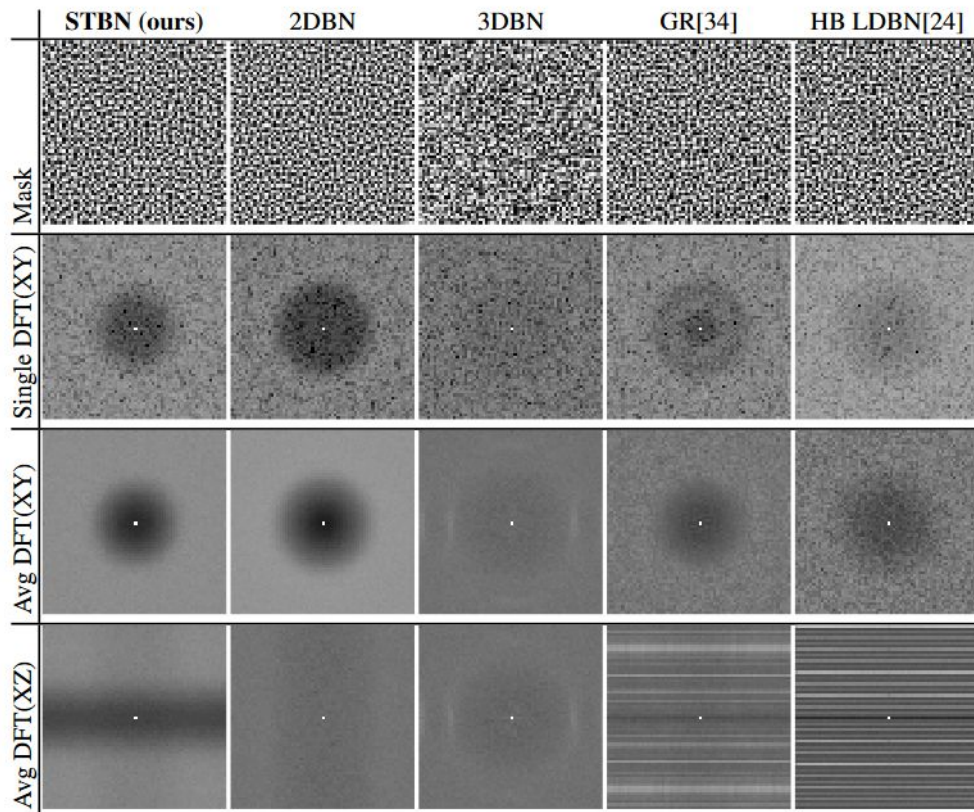
Blue Noise Dithered

1 bit per color channel, Gaussian blurred

Real Time Is Animated & Amortized Over Time



Animated Blue Noise - Frequencies



Void And Cluster: Scalar Noise Algorithm

Simplified algorithm (see paper for full version):

1. Start with zero energy field
2. Place a point at the lowest energy and update energy field
3. Repeat #2 until all points are filled in
4. Order of point insertion determines pixel value (remap to texture range).

Our modified energy function:

$$E(\mathbf{p}, \mathbf{q}) = \begin{cases} \exp\left(-\frac{\|\mathbf{p}-\mathbf{q}\|^2}{2\sigma^2}\right), & \text{if } \mathbf{p}_{xy} = \mathbf{q}_{xy} \text{ or } p_z = q_z \\ 0, & \text{otherwise,} \end{cases}$$

Simulated Annealing: Vector Noise Algorithm

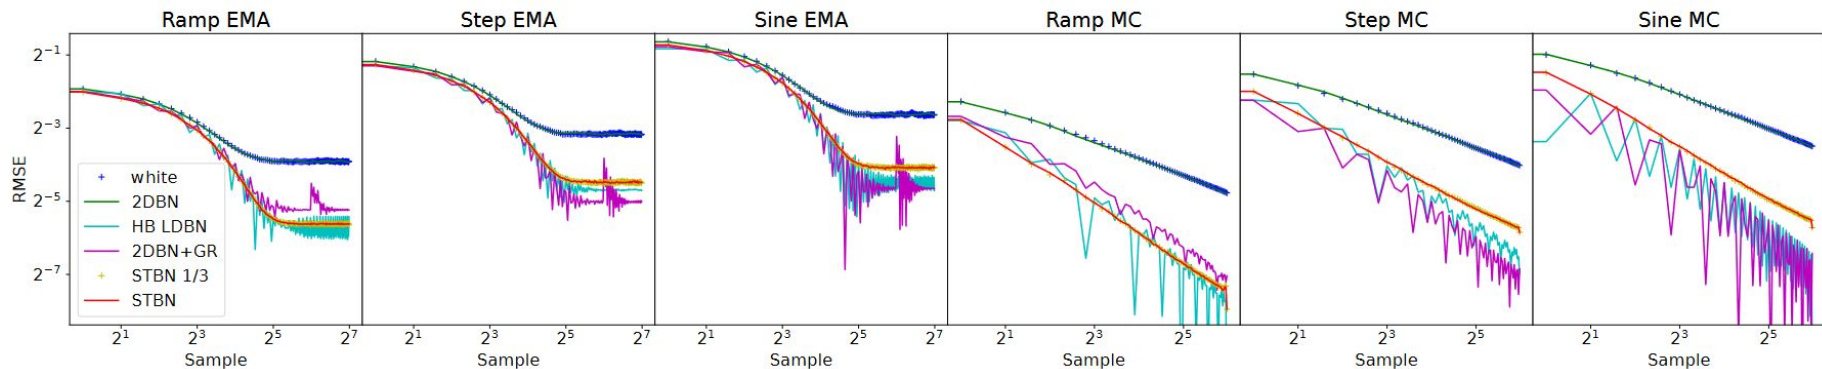
1. Initialize texture to uniform white noise vectors (*)
2. Pick 2 pixels at random, swap them if doing so improves overall energy
3. Repeat #2 until overall energy low enough, or enough swaps have occurred.

Our modified energy function:

$$E(\mathbf{p}, \mathbf{q}) = \begin{cases} \exp\left(-\frac{\|\mathbf{p}-\mathbf{q}\|^2}{\sigma_i^2} - \frac{\|\mathbf{V}_p-\mathbf{V}_q\|^{d/3}}{\sigma_s^2}\right), & \text{if } \mathbf{p}_{xy} = \mathbf{q}_{xy} \text{ or } p_z = q_z \\ 0, & \text{otherwise.} \end{cases}$$

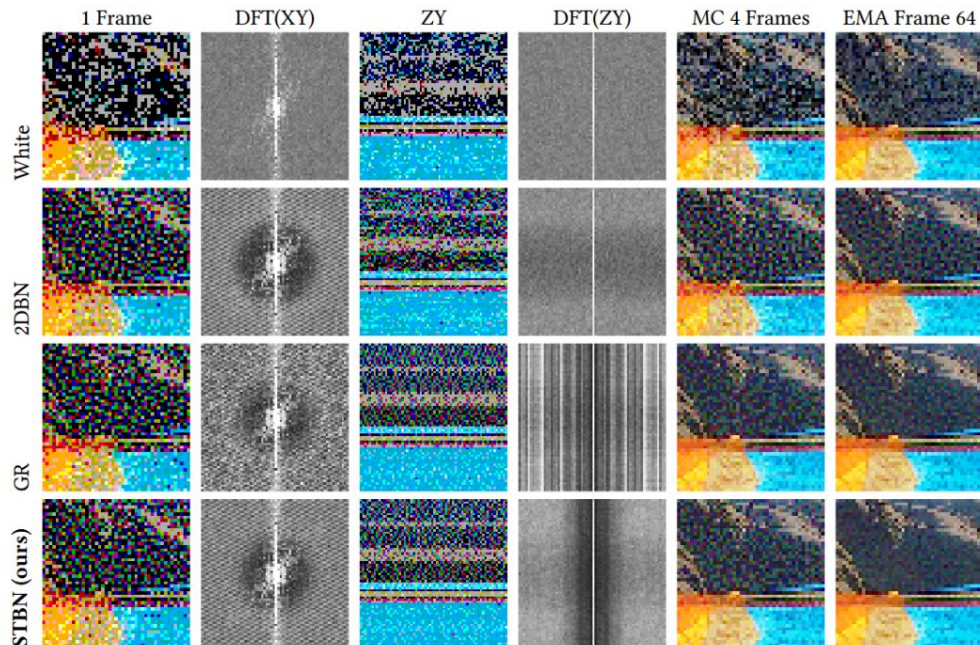
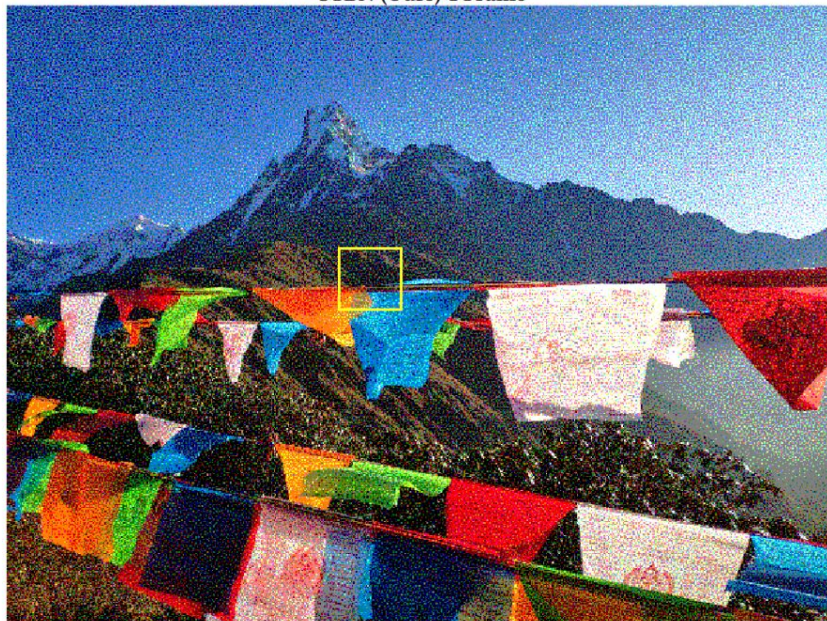
* Can use non uniform noise for importance sampled masks. More on that later!

Animated Blue Noise - Simple Function Convergence



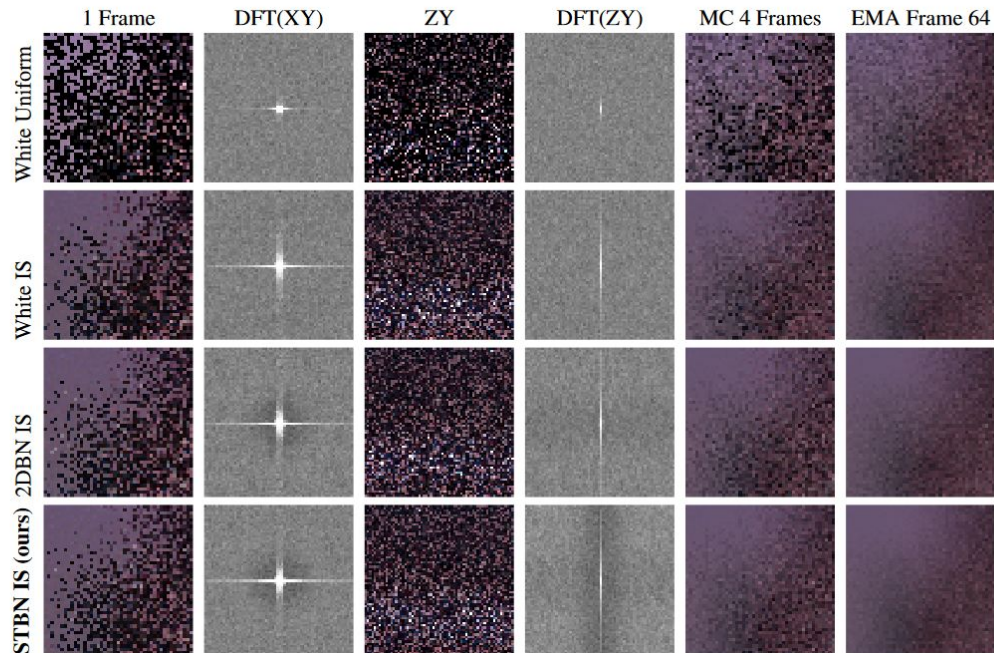
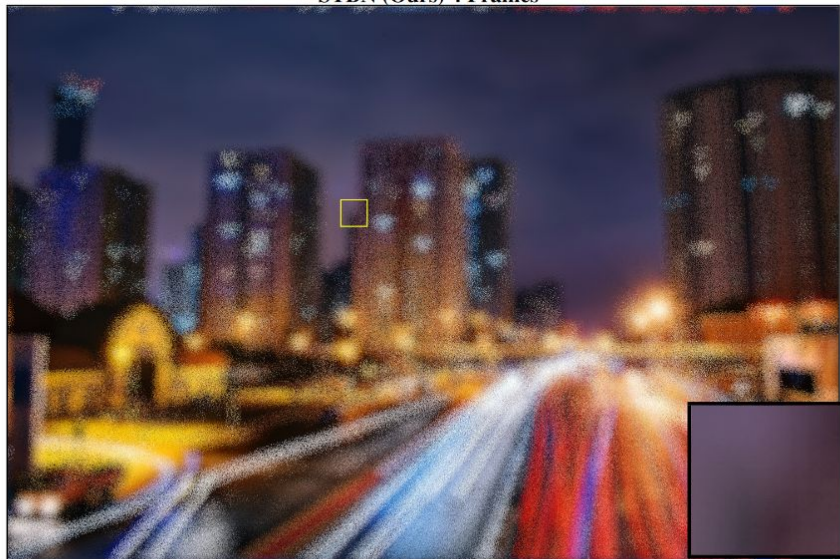
Rendered Results - Dithering

STBN (Ours) 1 Frame



Rendered Results - Stochastic Convolution

STBN (Ours) 4 Frames

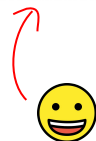


Rendered Results - Ray Traced Ambient Occlusion

STBN (Ours) 4 Frames



	1 Frame	DFT(XY)	ZY	DFT(ZY)	MC 4 Frames	TAA Frame 64
White						
White x Sobol						
2DBN						
STBN (ours)						



Rendered Results - Heitz and Belcour



(a) Heitz and Belcour w/ STBN



(b) 2D BN



(c) STBN (Ours)

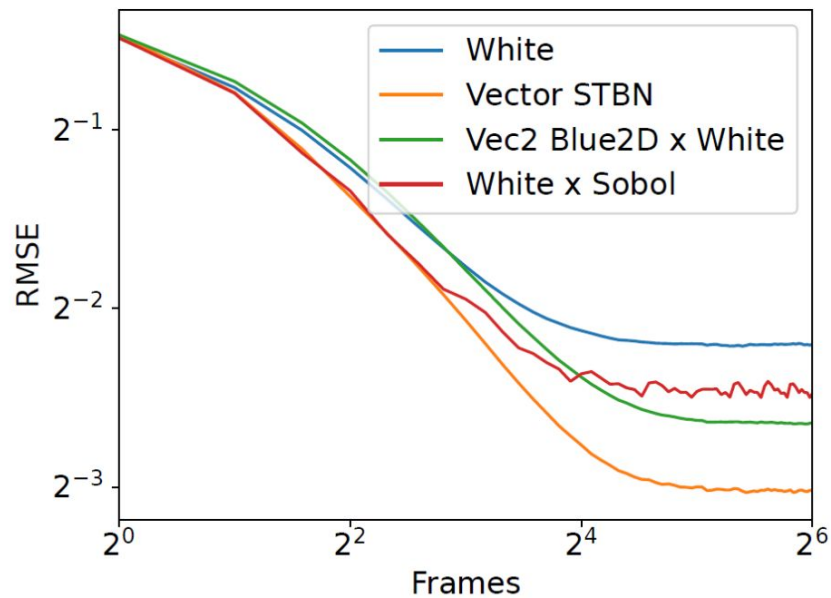


(d) Ground Truth

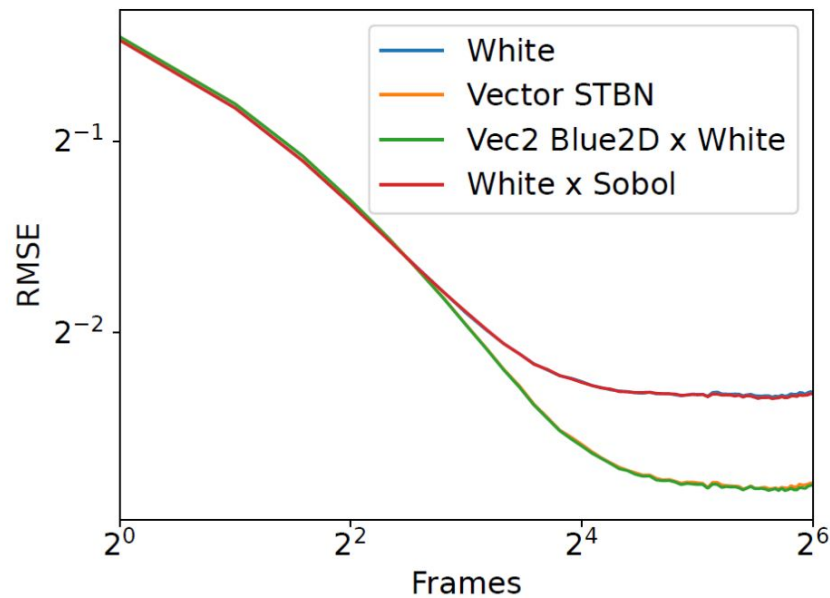
“Distributing Monte Carlo Errors as a Blue Noise in Screen Space by Permuting Pixel Seeds Between Frames”
Heitz and Belcour, 2019

A Challenge With Moving Pixels & 1 SPP

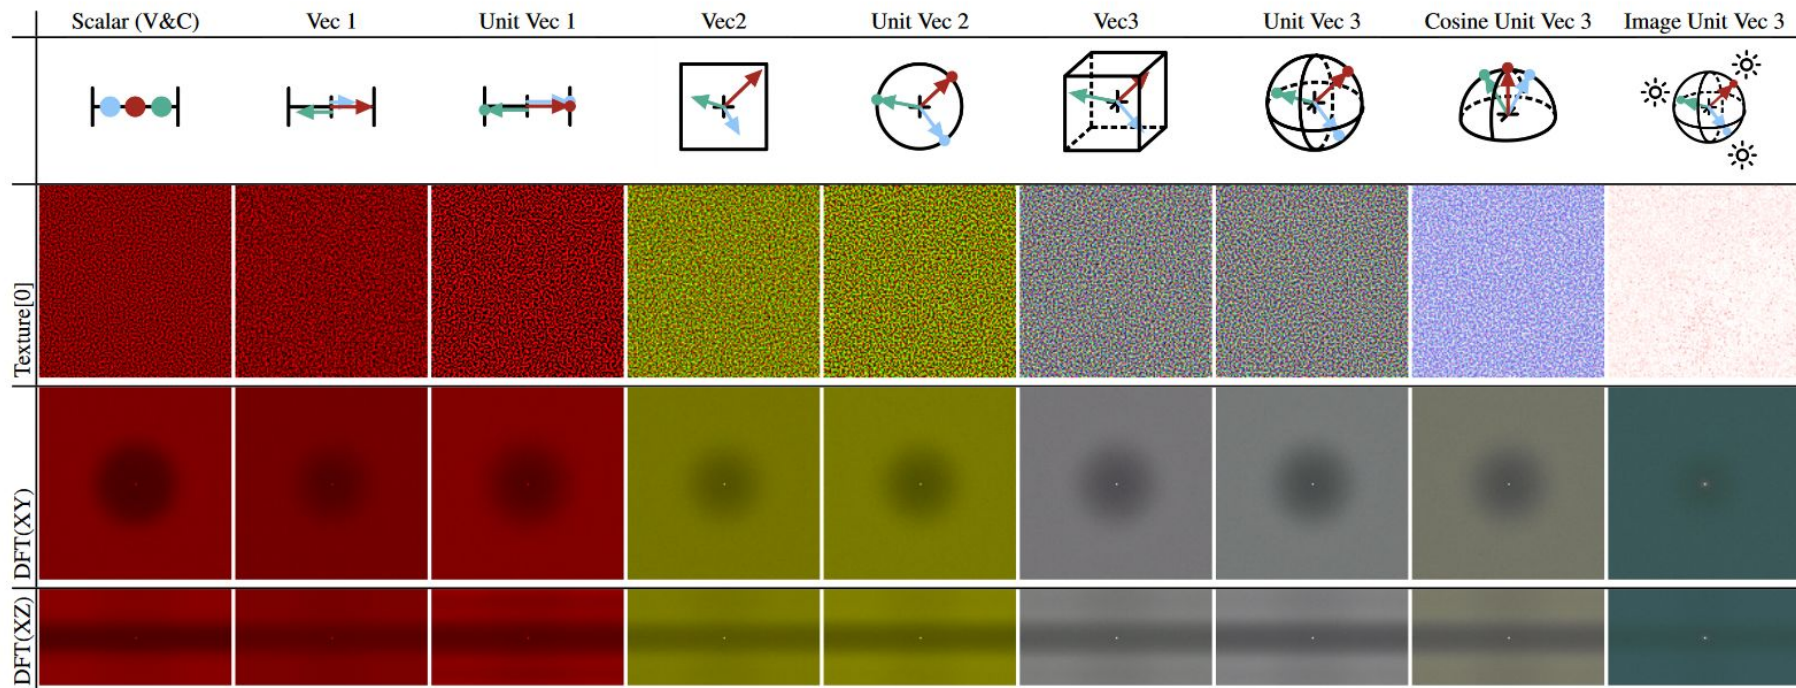
TAA (Still)



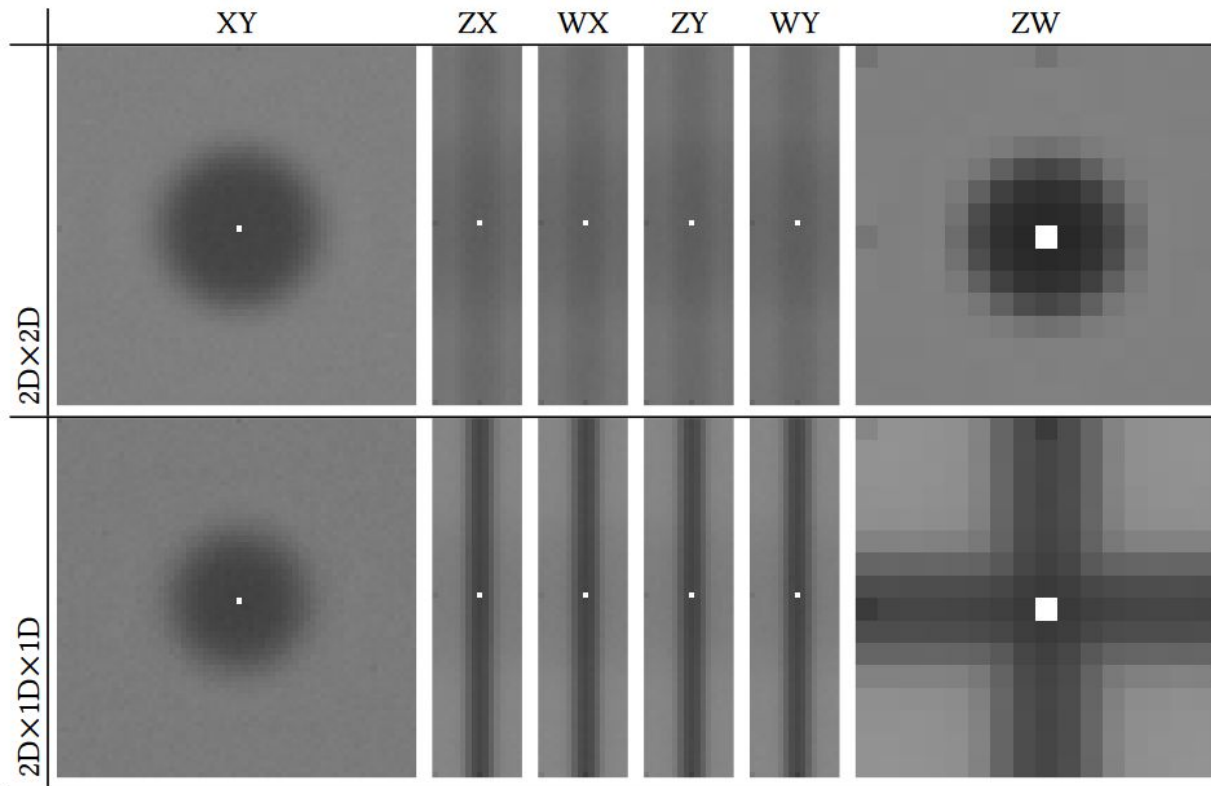
TAA (Moving)



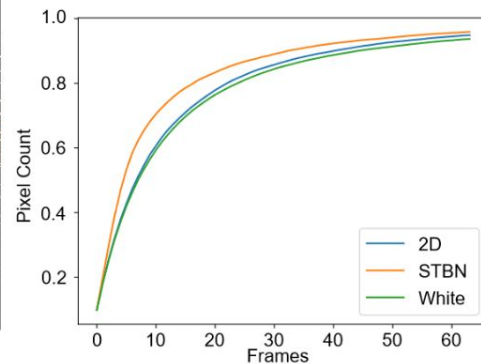
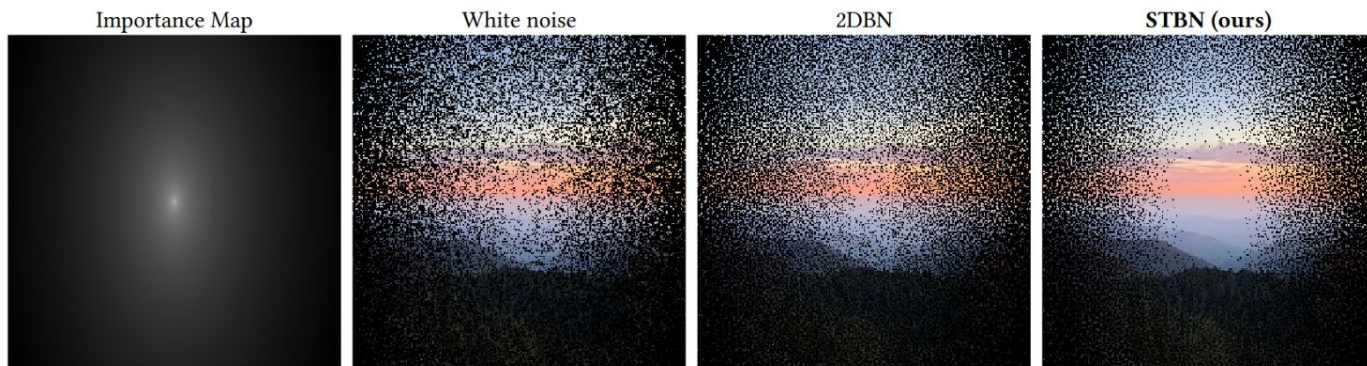
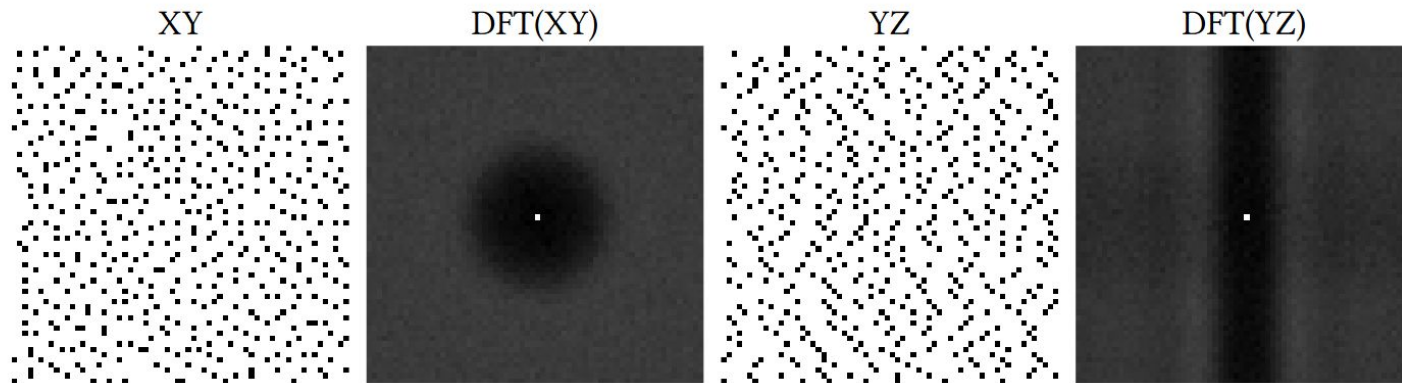
Other Results: Scalars, Vectors and Importance Sampling



Other Results: Higher dimensionality



Other Results: Thresholded Spatiotemporal Point Sets



Summary

- STBN is a drop-in replacement for traditional Blue Noise Masks
- Optimized temporally to be “better than white noise”
- Can be scalar or vector, uniform or importance sampled

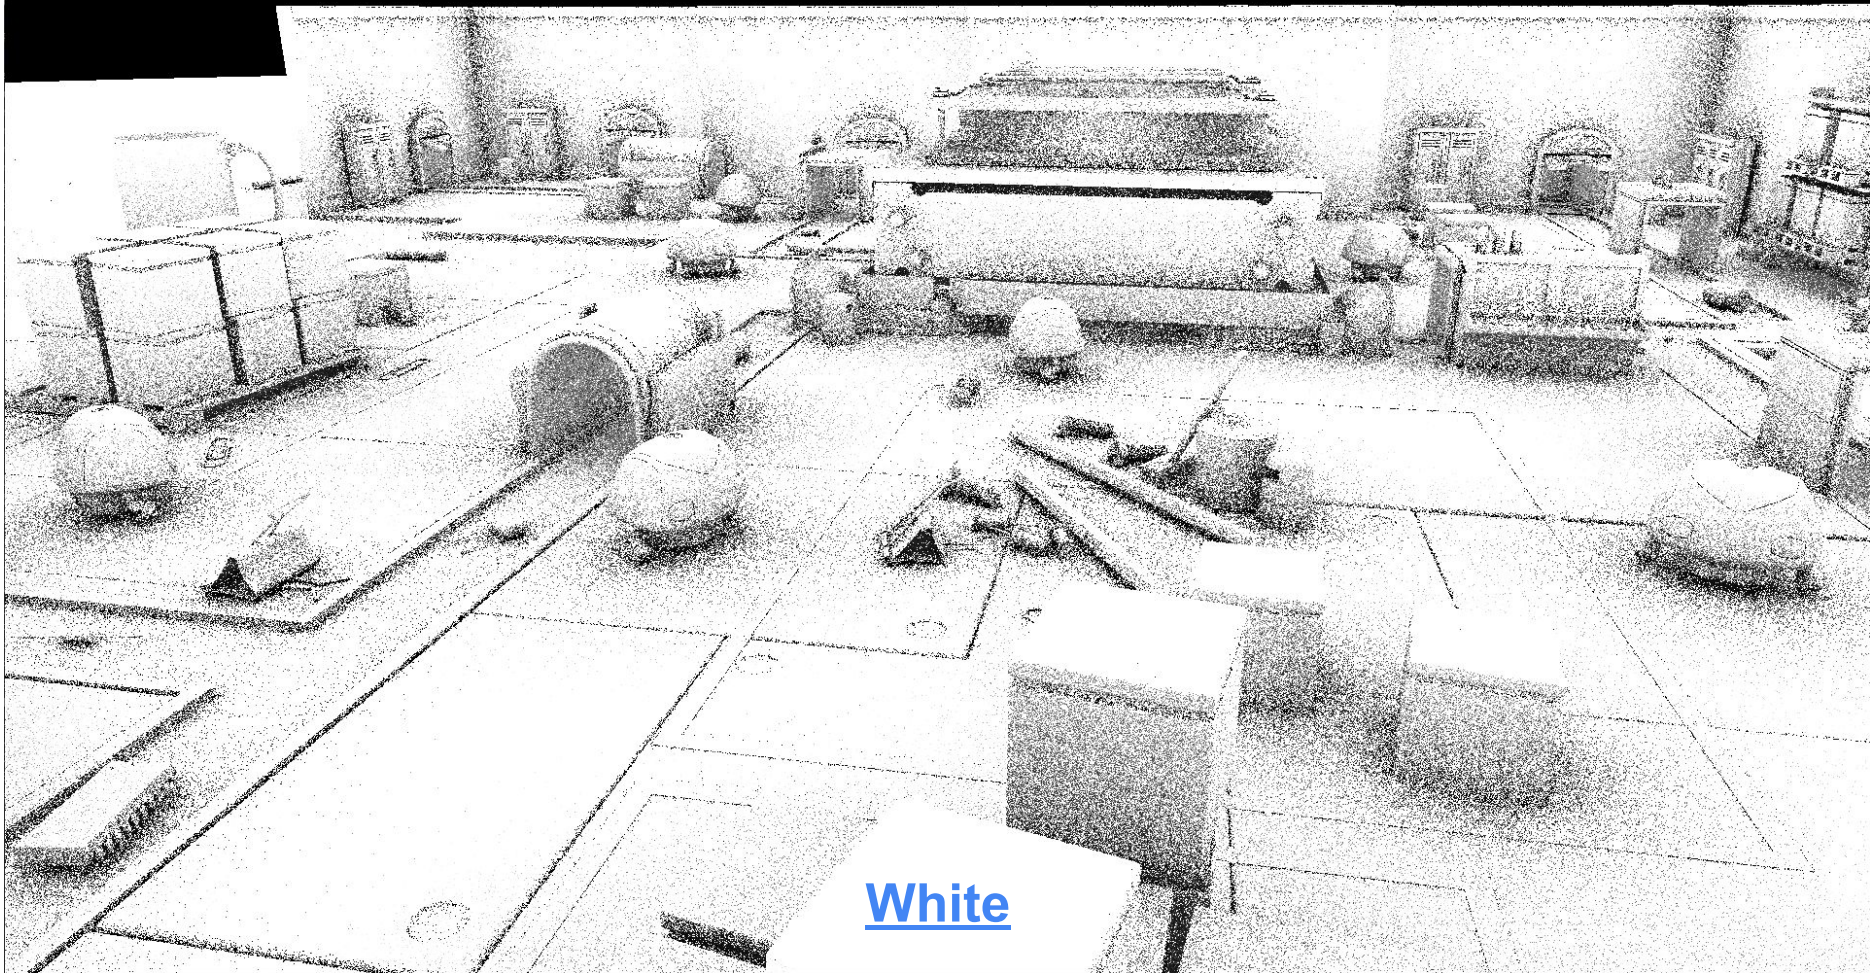
Future Work

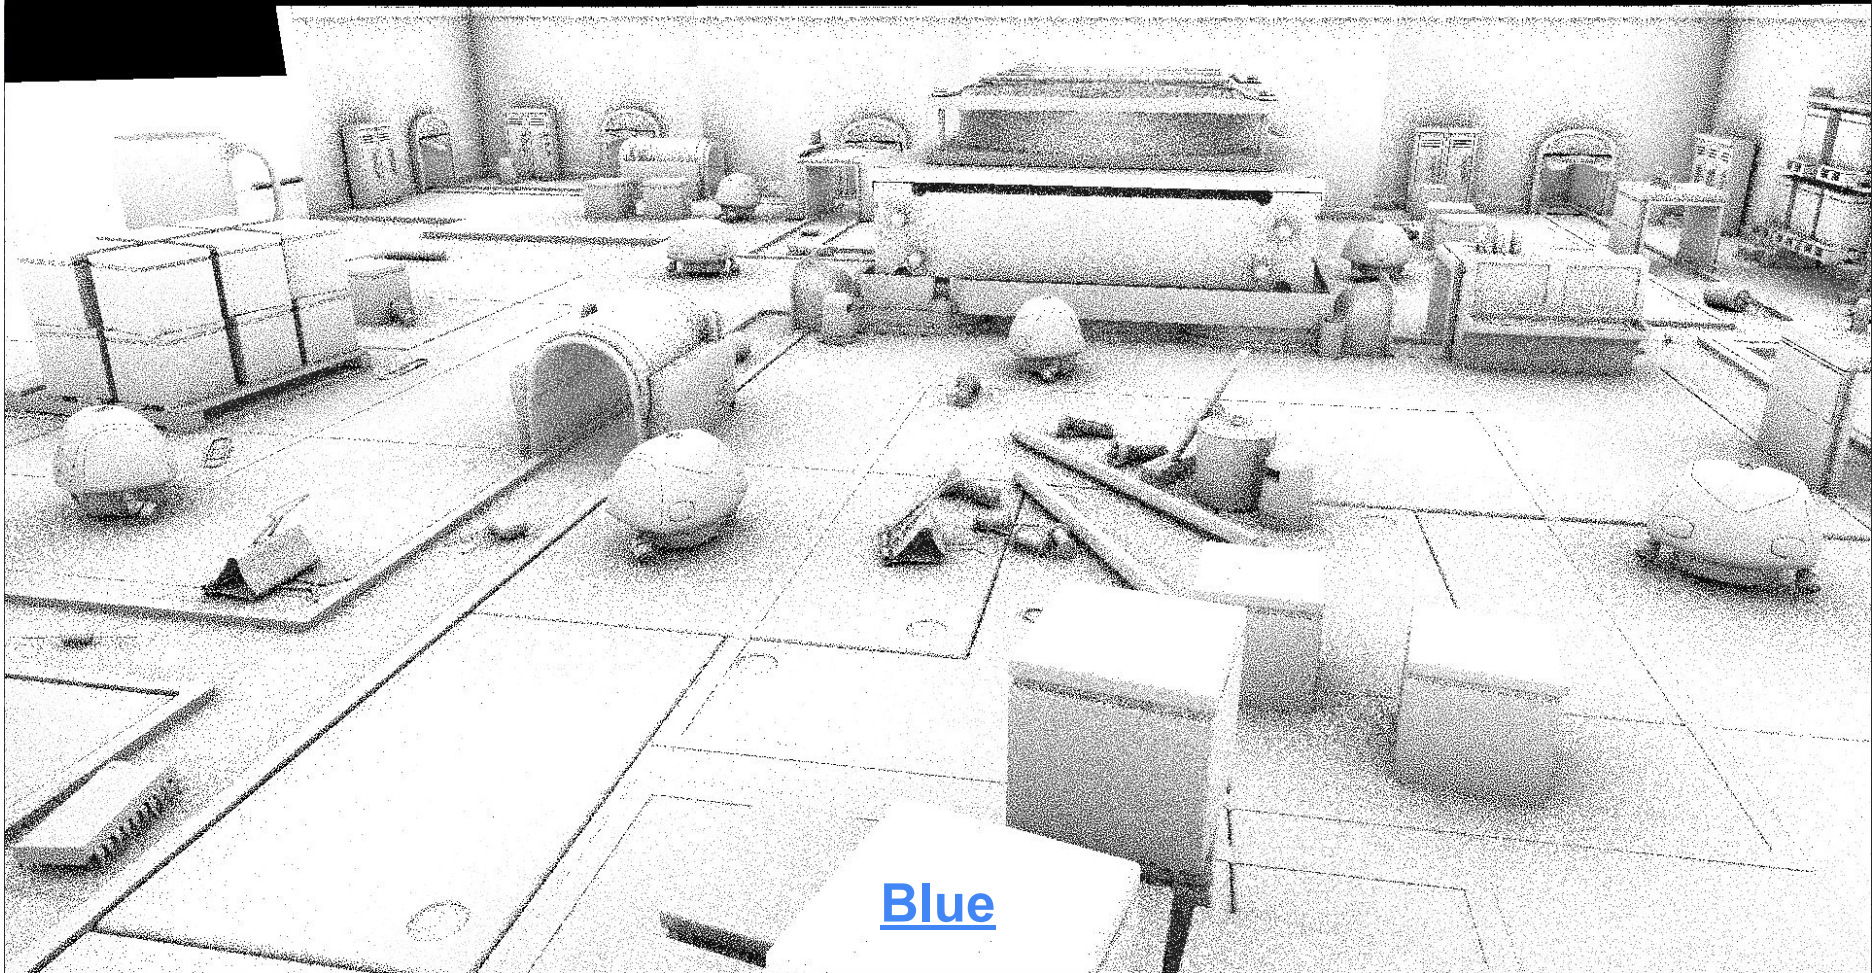
- Optimize for other filters (box blur? Unsharp mask?), rendering techniques, content
- Temporal optimization targeting TAA / EMA or others better
- Address convergence of moving pixels

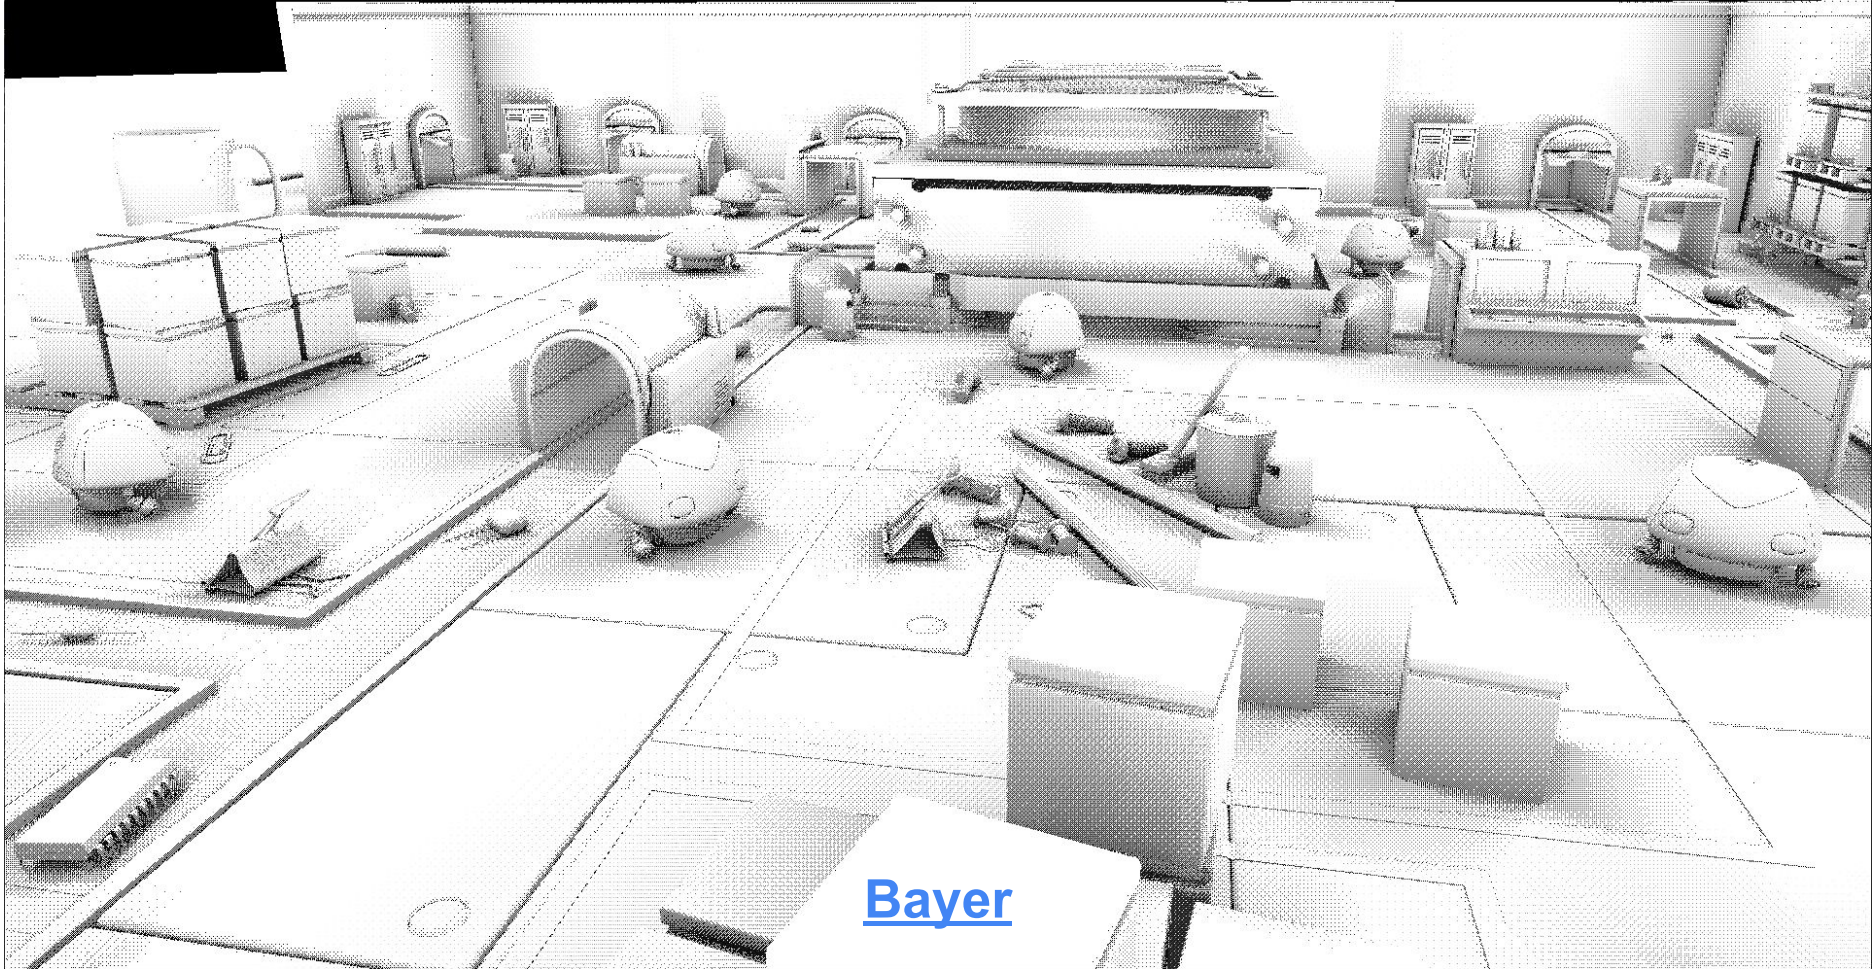
How/Why Do Blue Noise Masks (Textures) Work?

1. If x is a pixel's random seed and $f(x)$ is the shaded output of the pixel...
2. The assumption is that if x correlated with x' , then $f(x)$ will also be correlated with $f(x')$.
3. Blue noise mask values are anti correlated among neighbors, so the rendering results have that property too.

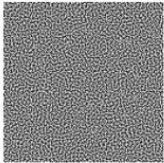
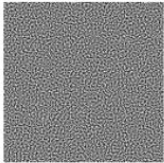
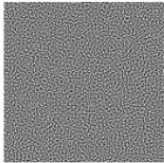
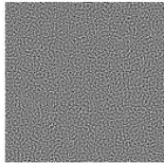
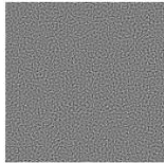
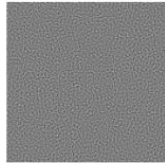
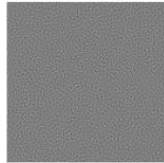
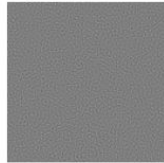
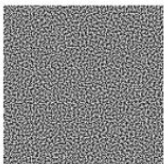
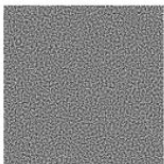
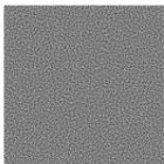
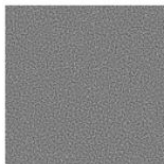
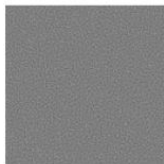



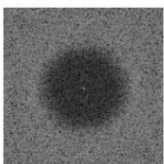
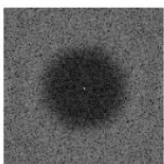
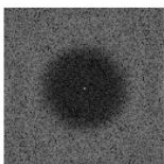
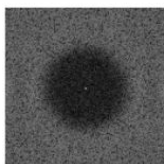
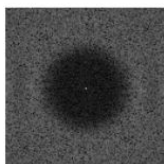
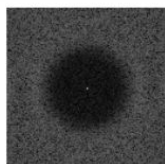
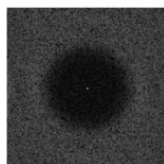
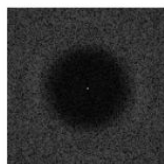
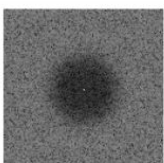
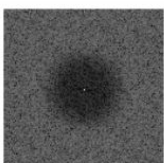
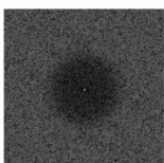
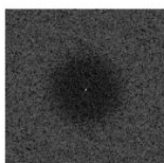
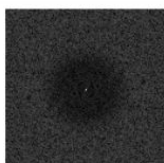
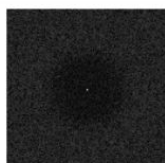

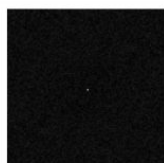
See next images...



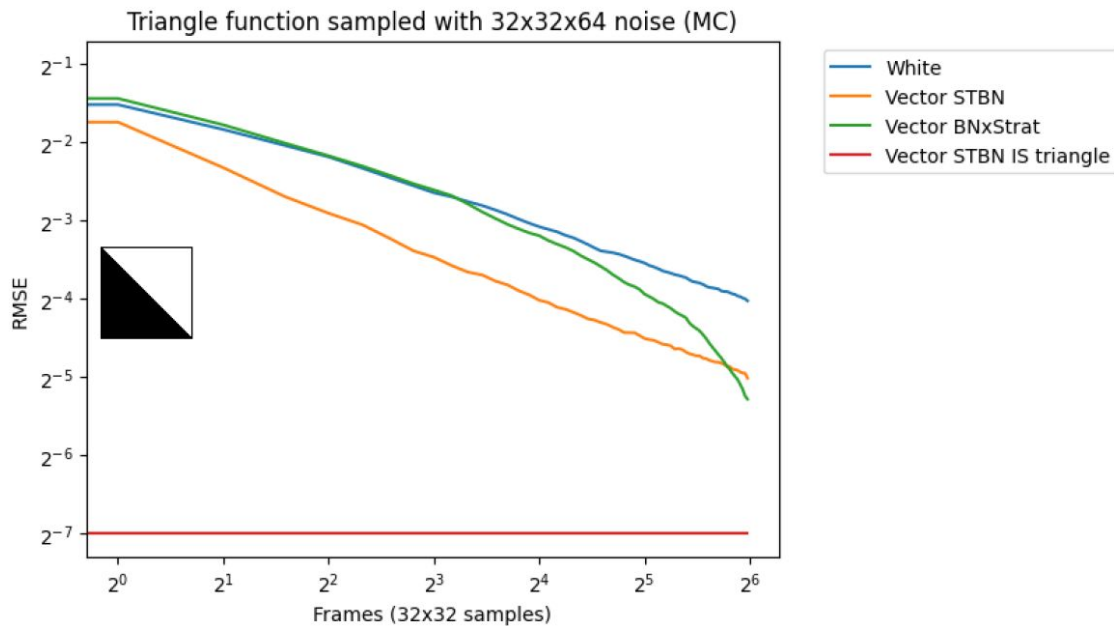




Does STBN Stay Blue As It Converges?

	1 sample	2	3	4	8	16	32	64
Blue2D								
STBN (Ours)								
Blue2D								
STBN (Ours)								

Blue Over Space, Stratified Over Time



Vector Blue Noise Through Curve Inversion

